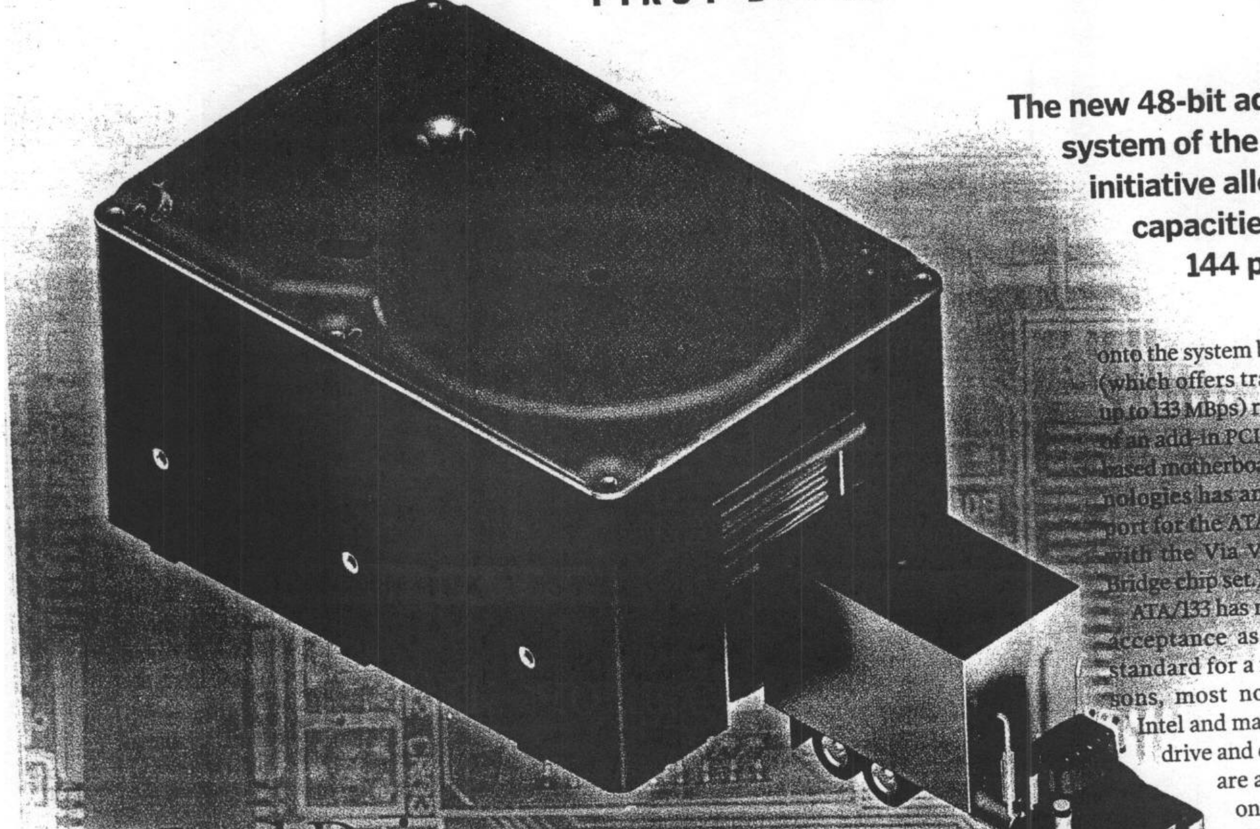

Hard Drives Reach the Next Plateau

The new 48-bit addressing system of the Big Drive initiative allows drive capacities of up to 144 petabytes.

BY JOHN R. DELANEY

Over time, most PC users realize that they can never have enough storage space. This is equally true for IT administrators who manage huge stores of data across networks, as well as for desktop users who work with digital imaging, large databases, and content-creation applications. As more users integrate their PCs into home entertainment centers, storing MP3 and streaming video files will eat up valuable hard drive space at an alarming rate.

Up to now, integrated drive electronics (IDE) drives topped out at 137GB of storage because of limitations inherent in the 28-bit addressing scheme developed for the current Advanced Technology Attachment (ATA) specification. But Maxtor's Big Drive technology tears down this 137GB barrier by incorporating a newly developed 48-bit addressing scheme, resulting in the industry's first 160GB IDE/ATA drive and opening the door for higher capacities in the not-too-distant future.

Led by Maxtor and other industry leaders such as Compaq, IBM, Microsoft, and Via Technologies, the Big Drive initiative solves the 28-bit hard drive addressing dilemma, which, in mathematical terms, allows for a maximum 268,435,456 sectors that hold 512 bytes of data. Simply put, 28-bit technology recognizes only 137GB of storage space. The new 48-bit addressing system changes the equation drastically, allowing drive capacities of up to 144PB (petabytes, or 144 million GB).

While petabyte drives are still years away, we can expect to see IDE hard drive capacities double in size within the next year. Having Microsoft on-board is significant, since its operating systems will have to support 48-bit addressing to work correctly.

Storage capacity is only half the story, though. As hard drives become bigger, they need to be able to move data more quickly through the interface to avoid becoming a bottleneck to system performance. Even though the newer "fast" IDE hard drives have speeds of 7,200 rpm (as opposed to 5,400 rpm), they are limited to peak transfer rates of 100 MBps through the ATA/100 interface.

At this point, the ATA/100 interface is still ahead of the pack as far as meeting the required Big/Fast drive data rates, but eventually, the interface will become a bottleneck. As a result, new interface technologies are emerging that will increase the interface speeds by as much as 50 percent in the coming year and by as much as 500 percent in the next four years.

One such technology, known as ATA/133, was recently introduced by Maxtor and is currently implemented in the company's two newest drives, the Maxtor Diamond Max Plus D740X and the Maxtor Diamond Max 540X.

Unlike ATA/100, which is supported by most motherboard chip sets and integrated onto the system board, ATA/133 (which offers transfer rates of up to 133 MBps) requires the use of an add-in PCI card for Intel-based motherboards. (Via Technologies has announced support for the ATA/133 interface with the Via VT8233A South Bridge chip set.)

ATA/133 has not gained wide acceptance as the next new standard for a number of reasons, most notably because Intel and many leading hard drive and chip set vendors are already focusing on the successor to ATA/100—the recently announced Serial ATA (SATA) interface, which will initially support 150-Mbps transfer rates. Original plans called for SATA devices to be available in the first half of 2002. We'll probably have to wait, however, until the first half of 2003 before SATA becomes the standard. Eventually, SATA will replace the ATA/100 interface, which uses bulky ribbon cables to transfer data in parallel mode. SATA will use thinner, longer cables, thus improving airflow though the PC chassis and enabling systems to run cooler. More important, future SATA revisions will enable transfer rates of up to 600 MBps in the coming years. Even so, Maxtor believes the ATA/133 interface will provide an interim solution should drive speeds breach the 100-MBps barrier before SATA is ready.

Maxtor has implemented Big Drive technology for the first time in its **Maxtor DiamondMax D540X** ($399.95 direct, *www.maxtor.com*) a mammoth 160GB IDE, 5,400-rpm drive that targets the network-attached storage

---

and e
The D
break
signe
and fe
It is al
incor
face.
Ultra
Promi
Set
trolle
GHz F
8200
forwa
133 ca
drive,
the in
utility
wizar
and fo
instal
2000.
unit is
BIOS r
out a
machi
date o
port t

IBM squa meg are

Bot
boast
data
ATA/1
bench
ed no
claim,
terfac
tested
conne
trolle
8,120
WinM
Busin
We th
the pri
on ou
and re
Disk
Busin
that th
has th
transf
hard
busin
for no

and enterprise server market. The D540X, the first IDE drive to break the 137GB barrier, is designed using four 40GB platters and features a 2MB cache buffer. It is also one of the first drives to incorporate the ATA/133 interface. (We tested it with the Ultra133TX2 add-in card from Promise Technology.)

Setting up the drive/controller combination on our 2.0-GHz P4–based Dell Dimension 8200 test system was straightforward. We installed the ATA/133 card, connected the D540X drive, booted the system with the included Maxtor MaxBlast II utility disc (which contains a wizard for easily partitioning and formatting the drive), and installed Microsoft Windows 2000. Since our Dell test-bed unit is a fairly new system, the BIOS recognized the drive without a problem, although older machines may need a BIOS update or may not be able to support the drive at all.

IBM hopes to hit 100 gigabits per square inch by 2003, which means that megacapacity drives of 400GB or more are a real possibility.

Both Promise and Maxtor boast a 33 percent increase in data rates because of the ATA/133 interface, but our benchmark test results provided no proof to support this claim, although the ATA/133 interface was slightly faster. We tested the D540X while it was connected to the Promise controller and received scores of 8,120 on our Business Disk WinMark 99 and 54.1 on our Business Winstone 2001 tests. We then tested the drive using the primary IDE port (ATA/100) on our system's motherboard and received scores of 7,810 on Disk WinMark and 53.5 on Business Winstone, indicating that the ATA/100 interface still has the overhead to handle the transfer rate of high-capacity hard drives running typical business applications, at least for now.

The **Maxtor DiamondMax Plus D740X** ($299 direct) features 80GB of storage capacity and also uses the ATA/133 interface, but it is a 7,200-rpm drive designed with two 40GB platters and 2MB of SDRAM for cache buffering. Geared towards the business and enthusiast market, the D740X is a performance drive.

The D740X features a ball-bearing motor or Maxtor's Quiet Drive technology to reduce the noise associated with ball-bearing drives. Quiet Drive uses a fluid dynamic bearing (FDB) motor to achieve a low-noise, high-speed storage solution for desktops, as well as gaming console systems.

Like the D540X, the D740X was easy enough to install using Maxtor's MaxBlast Plus II utility. Once again, we noticed only a slight performance increase when testing the drive with the ATA/133 card compared with the integrated IDE controller, with Disk WinMark scores of 9,400 and 8,900, respectively.

Western Digital is one of more than 70 companies that are throwing their weight behind the SATA standard and forgoing the implementation of ATA/133 technology in their product line. The company's family of 7,200-rpm Caviar drives, which targets the high-performance desktop market, gets a high-capacity boost with the **Western Digital Caviar 120 GB ATA/100** drive, the first 120GB 7,200-rpm drive to hit the market ($299 direct, *www.westerndigital.com*). Based on a three-platter (40GB each) architecture, the Caviar 120 includes a 2MB cache, but Western Digital plans to release an 8MB buffer version in the near future. Though the increased buffer size may be advantageous when used in server configurations, whether desktop users running normal business applications will see a notable performance boost compared with the 2MB version is not clear.

The Caviar 120 is bundled with Western Digital's Data Lifeguard utility, which includes tools for installation and diagnostics, as well as a set of reliability and data-protection tools for monitoring and repairing drive errors. We connected the drive to our system's integrated IDE port, ran the EZ installation program, and installed Microsoft Windows 2000 without a hitch. The Caviar 120 produced a score of 54.4 on our Business Winstone and 8,740 on our Business Disk WinMark tests. When we connected the drive to the Promise ATA/133 card, we received an identical Winstone result and a WinMark score of 8,900, which is a bit higher than that with the ATA/100.

The **IBM Deskstar 120GXP** ($369 direct, *www.ibm.com*) may not have been the first 120GB 7,200-rpm drive available, but our benchmark test results show it to be the best performer we've seen so far, if only by a small margin. For now, IBM is sticking with the ATA/100 interface while waiting for SATA to make its presence known. The Deskstar 120GXP, which uses three 40GB platters and includes a 2MB buffer, is IBM's highest-capacity IDE drive. It is well suited to desktop systems that require the speed and capacity needed for digital content creation and storage, as well as for storing large video and audio files.

The drive also implements IBM's AFC (antiferromagnetically coupled) media coating, better known as pixie dust, which the company claims improves data stability and increases areal density (the amount of data, in bits, that can fit onto a platter). Greater areal density enables higher-capacity drives. To put it in perspective, the areal density of the Deskstar 120GXP is just under 30 gigabits per square inch. But IBM hopes to hit 100 gigabits per square inch by 2003, which means that megacapacity drives of 400GB or more are a real possibility.

Using the included Disk Manager utility, we installed the Deskstar 120GXP on our test-bed system and loaded Windows 2000. Connected to the system's IDE controller, the 120GXP produced a Disk WinMark score of 8,840 and a Business Winstone score of 54.8. After we installed the Promise ATA/133 card and connected the drive, we received scores of 9,410 and 56.

All things considered, any of these drives would be a fine complement to a high-end system. Their capacities might seem like overkill today, but if multimedia is on your roadmap, you might just save yourself an upgrade headache farther down the line.

HARD DRIVES

High scores are best.
Bold type denotes first place.

	Interface	Business Winstone 2001	Business Disk WinMark 99 (Thousands of bytes per second)	High-End Disk WinMark 99 (Thousands of bytes per second)
IBM Deskstar 120GXP	ATA/133 ATA/100	**56.0** 54.8	**9,410** 8,840	**25,000** 23,600
Maxtor DiamondMax D540X	ATA/133 ATA/100	54.1 53.5	8,120 7,810	20,100 19,600
Maxtor DiamondMax Plus D740X	ATA/133 ATA/100	54.7 53.9	9,400 8,900	23,400 22,100
Western Digital Caviar 120	ATA/133 ATA/100	54.4 54.4	8,900 8,740	23,600 23,300

We ran all tests on a 2-GHz P4–based Dell Dimension 8200 with 512MB of RAM, running Windows 2000 (Service Pack 2). All drives were formatted for NTFS.

Does SCSI Have a Future?

The old-school interface is keeping pace with current standards and offers some compelling uses. **By Alfred Poor**

Computers now come with interfaces that were unheard of just a few years ago. The convenience of Universal Serial Bus (USB) and the high speed of IEEE 1394 (FireWire) make attaching peripheral devices and moving data at high rates easy. By contrast, SCSI interfaces may seem to be relics of the PC's past. SCSI devices are known for being complex to install, tricky to configure, and expensive. But SCSI is now faster than ever, and it may prove worthwhile for some common uses.

SCSI BASICS

SCSI stands for *Small Computer System Interface* and is a parallel interface. USB and FireWire are serial interfaces, which means that they send bits of data one at a time along a single signal line. By contrast, SCSI uses 8 or 16 separate data lines, and it sends data simultaneously on all the lines.

The SCSI interface is a tiny network within your computer. Each device has its own ID number, and the devices communicate with the controller and with one another. SCSI equipment must arbitrate conflicts when two or more devices attempt to use the SCSI bus at the same time. This requires more on-board intelligence and is one reason SCSI hardware often costs more than devices designed for other interfaces.

There are essentially three parameters that define a SCSI connection and distinguish the various versions: the bus width, the bus speed, and the type of electrical signaling (more on signaling below). SCSI-1, the original version of SCSI, uses an 8-bit-wide data bus running at 5 MHz. The transfer rate of up to 5 MBps compares favorably with that of USB, at 1.5 MBps.

SCSI-2 introduced two major changes. The first—called Fast SCSI—doubled the bus speed to 10 MHz, which doubled the maximum transfer rate of the 8-bit bus to 10 MBps. SCSI-2 also introduced a 16-bit bus—known as Wide SCSI—which again doubled throughput. Combine the two new features and you get Fast/Wide SCSI, capable of up to 20-MBps throughput—near the low end of FireWire's 12.5- to 50-MBps range.

SCSI-3—also known as UltraSCSI—raised the bar again by doubling the bus speed to 20 MHz. This gave 8-bit UltraSCSI a 20-MBps transfer rate and Wide UltraSCSI a 40-MBps rate.

The SCSI-3 specification became the basis for a series of SCSI Parallel Interface (SPI) standards. SPI-2 is also known as Ultra2 SCSI and is the last of the standards to include an 8-bit-wide bus. The bus speed doubled once more to 40 MHz, so Ultra2 SCSI can transfer up to 40 MBps, and Wide Ultra2 SCSI can handle 80 MBps—significantly faster than FireWire's maximum throughput.

ONGOING DEVELOPMENT

SPI is currently in its third version, known as Ultra3. To qualify for Ultra3, a controller must support at least one of five different features:

- *Fast-80 data transfer.* As its name implies, this feature speeds up the data-bus transfer rate. But rather than raising the rate by raising the bus clock from the 40 MHz used in Ultra2 SCSI, Fast-80 data transfer uses a technique called *double-transition clocking.* Up through Ultra2 SCSI, data is transferred only once per bus clock cycle. The double-transition-clocking approach transfers data on both the rise and fall of the bus clock signal, providing twice as many transfers. Ultra3 SCSI comes in a wide (16-bit-data-bus) version only, so the maximum transfer rate is 160 MBps.
- *CRC error checking.* This process uses a complex, 32-bit algorithm to compute a check value and will catch all but one in billions of errors. This is a far more effective technique than parity-checking.
- *Domain validation.* This ensures that data transfers occur at the maximum rate. The attached components' rates are determined at power-up, then the controller sends a message to each device, which each device returns. If the controller gets the same data, transmission can proceed at the previously set rate. If not, the controller steps the rate down until the exchange is successful.

Too Many Hang-Ups

My modem disconnects frequently. My ISP blames line noise. My phone company says there is no problem. How can I find out who is right and who's passing the buck? And if the phone company is right and there's no line noise, what might the problem be?

KEVIN GINSBERG

If the problem is noise, you should hear it during voice calls. If you have a dedicated data line, plug in a phone and use it for a few days. If you don't hear any noise, then noise is not the problem.

A Call Waiting tone will also cause a modem to hang up. If you have Call Waiting, find out the code in your area to turn it off (your telephone book should have the information). Include the code in your Dial-Up Networking setup. Also try dialing in with another modem, preferably an external version, to rule out a problem with your current modem. If you're using a 56K modem, try setting the modem to connect at slower speeds and see whether that eases the problem.—*M. David Stone*

Out in the Cold

My computer is in an office in my garage, and the temperature in there drops to about 55 degrees at night in the winter. In the morning, the computer doesn't recognize my Zip Drive until the room's temperature reaches about 60 degrees. Why would the temperature affect my computer system this way?

JOE ROMANO

This sounds like a classic case of something not quite connecting until some metal component expands slightly as it heats up. A connector may not be seated well, or there may be a hairline crack in the cable to the drive or in the connector itself.

First, try removing the power and data connectors from the drive and reseating them. If that doesn't solve the problem, try a different power connector. You may need to replace the data cable if it's badly damaged.—*MDS*

KVM Switches Revisited

I read with interest your article on KVM (keyboard, video, and mouse) switches in your issue of May 9, 2000 ("Make Room on Your Desk," page 137). The information was great, but you missed one tip. Although most laptops have only a single PS/2 connector, a Y-cable lets you plug a keyboard and mouse into the same port. Y-cables are hard to find, but they work with most laptops.

DAVID CHARLES BAKER

Y-cables can indeed be useful, but they don't always work with KVM switches. In my tests, using a Y-cable with a KVM switch confused my Toshiba laptop, making mouse movement seem random. The best compromise is to plug the keyboard connector from the KVM switch into the Y-cable but plug a second mouse directly into the Y-cable, instead of using the mouse cable from the switch.—*MDS*

Unwanted Scan Lines

After not using my HP PhotoSmart scanner for several months, I tried scanning some 35-mm negatives. Every scan had a reddish-purplish streak across the image. The negatives aren't scratched, so this seems to be a problem with the scanner. What can I do to fix this problem?

GEORGE WILSON

A scanner's sensor array reads only one line across at a time. If a single sensor is blocked by, say, a speck of dust, it will create a streak across the scan. Reddish-purplish is one way to describe magenta; that suggests the scanner is reading only the red and blue components—but not the green component—for a particular pixel.

Try opening the top panel of the PhotoSmart scanner and blowing compressed air—available at most photo supply stores—across the scan area. This should blow away the dust and solve the problem—assuming that's the reason for the streak.—*MDS*

The Many Faces of SCSI

Interface	Maximum data-transfer rate (MBps)	Bus width (bits)	Maximum cable length by signaling method (meters) Single-ended	Differential	LVD	Maximum devices supported
SCSI-1	5	8	6	25	N/A	8
Fast SCSI	10	8	3	25	N/A	8
Fast/Wide SCSI	20	16	3	25	N/A	16
UltraSCSI	20	8	1.5	N/A	N/A	8
Wide UltraSCSI	40	16	1.5	25	N/A	8 or 16*
Ultra2 SCSI	40	8	N/A	25	12	8
Wide Ultra2 SCSI	80	16	N/A	25	12	16
Ultra160 SCSI	160	16	N/A	N/A	12	16
Ultra320 SCSI	320	16	N/A	N/A	12	16

N/A—Not applicable: The technology is typically not manufactured in this configuration. * 8 devices for single-ended signaling, 16 for differential.

- *Quick arbitration and selection.* QAS speeds up the procedure that allocates use of the bus. One device can relinquish control directly to another without having to repeat an arbitration process.
- *Packetization.* This transmits data in Internet-style packets. Multiple commands, status messages, and chunks of data can be combined in a single stream, eliminating the overhead required to send this data in separate transmissions.

Efforts have been made to define interfaces that support more than one of these five features. Ultra160 (or Ultra160/m) must support Fast-80, CRC, and domain validation, and Ultra160+ requires support for all five features.

Work on the SPI-4 specification—called Ultra-320 SCSI—is under way. Data-transfer rates will hit 320 MBps. Ultra320 SCSI relies on double-transition clocking, but the base clock rate is raised from 40 Hz to 80 Hz. At full speed, this bus can transfer the equivalent of a full CD-ROM in just 2 seconds. Devices supporting Ultra320 SCSI are just beginning to appear on the market.

BUS WIDTH AND SIGNALING

Bus width determines how many devices you can support. The 8-bit bus is limited to 8 SCSI IDs, but the 16-bit bus supports 16 IDs. With a wide controller, you can have up to 15 SCSI devices. (A SCSI controller uses one ID.)

SCSI capability is also affected by signaling. Signal quality declines as cables get longer and as clock speed increases. The evolving SCSI specification required devices to be closer to their hosts. To keep pace with other enhancements, signaling had to improve.

The original SCSI specification called for *single-ended* signaling. Each bit is sent as a signal over an individual wire. The voltage on this wire is compared with that of a paired ground wire. If the voltages are the same—0—the bit is read as 0. If the voltage on the signal wire is not 0, the bit is a 1. With this system, SCSI-1 cables could stretch up to 6 meters (nearly 20 feet). SCSI-2 was limited to just 3 meters (about 10 feet). UltraSCSI and Wide UltraSCSI could handle 3-meter runs, but not with more than four devices. Eight devices limited the cable to 1.5 meters (about 5 feet).

The solution was *differential signaling,* which uses two wires for each bit. To send a 0, the system applies no voltage to either wire. For a 1, it sends a positive voltage on one wire and a negative voltage on the other, creating a larger differential voltage than in the single-ended method. The extra circuitry used for differential signaling is more expensive and requires more power. The big advantage is that differential SCSI buses can be as long as 25 meters (about 82 feet).

Ultra2 SCSI introduced a new signaling standard called *low-voltage differential* (LVD). This standard calls for a smaller voltage difference across paired wires. The result is a system that is limited to 12-meter cable runs (nearly 40 feet) but much less costly. The circuitry required for signaling can be integrated on the SCSI interface chip, so no extra components are required.

WHITHER SCSI?

So where does this leave the SCSI interface? Is it a practical tool for computers, or is it a relic?

The extreme demands on modern servers point to one of the main uses for SCSI. Servers must store and furnish huge amounts of data. RAID (Redundant Array of Inexpensive Disks) technology lets you load a server with multiple hard drives and access them simultaneously. But today's high-speed drives can have spindle speeds in excess of 10,000 rpm and sustained transfer rates above 40 MBps. You'd need Ultra160 SCSI or better to handle five of these monsters.

SCSI may also prove helpful if you want to add several storage devices to your home system. Wide UltraSCSI is about equal to FireWire in throughput, and it can handle far more data than USB. Newer SCSI versions are better yet.

Chances are good that personal computing devices will rely more on USB and FireWire, because they're cheaper and easier to configure. Eventually, SCSI may be relegated to use in servers alone. For now, SCSI is a reasonable alternative when you're trying to stuff a lot of devices into a single system.

Alfred Poor is a contributing editor of PC *Magazine.*

MAGNETO-OPTICAL DATA STORAGE

Combining optical techniques and magnetic storage media, MO technology will likely be used in magnetic drives to improve data density. The result: the continued fall of the cost of storing bits.

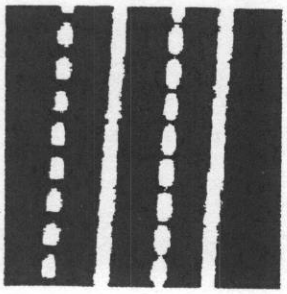

ATOMIC-FORCE MICROSCOPE SCAN OF A RECORDED BIT ON A DISK'S PHASE CHANGE MATERIAL COATING, SIROS TECHNOLOGIES, SAN JOSE, CA

OPTICALLY ASSISTED WINCHESTER TECHNOLOGY, SEAGATE TECHNOLOGY, SCOTTS VALLEY, CA

COMPUTER USERS TEND TO TAKE FOR GRANTED THE EXISTENCE OF MAGNETIC AND OPTICAL MEMORY DEVICES THAT OFFER BREATHTAKING DATA-STORAGE CAPACITY WITH EXTREMELY HIGH RELIABILITY AT VERY LOW COST. THE TECHNOLOGICAL WONDERS OF THESE RUGGED DEVICES ROUTINELY ESCAPE USERS' NOTICE, BECAUSE THE INTERNAL FUNCTIONING OF DISK DRIVES IS HIDDEN FROM VIEW. NEVERTHELESS, MAGNETIC AND OPTICAL drives are arguably among the world's most sophisticated electromechanical devices. Equally astonishing is the rapid advance of this technology, far outpacing the highly visible technology sectors of semiconductor electronics and computer software.

Magneto-optical (MO) data storage represents a combination of optical data-storage techniques and magnetic storage media. This article aims to show that this hybridization is an ongoing process, with the most recent developments being among the most exciting in the 40-year history of the technology. Two of the most striking physical characteristics of magnetic and optical drive technology are the exponential growth of the data-storage density in commercial products and the mechanical performance of recording heads following a track of recorded information on the medium.

The storage density metric is called "areal" density, or bit count per unit area on the medium. It has recently been doubling every 8 to 12 months in magnetic disk storage and about every 24 months in MO storage, according to DISK/TREND, Inc., a market research firm in Mountain View, Calif. This advance translates into an exponential drop in the cost of a stored bit of information—value to the customer unheard of in any other industry.

The mechanical operation of disk drives is stunning when one realizes that a "flying" magnet head in the drive is analogous (through scaling) to a 747 aircraft flying over rolling terrain a few inches off

TERRY MCDANIEL

Figure 1(a). Writing a track of information in optical recording, (b) definitions of wavelength 2, (c) numerical aperture (NA), and (d) spot full-width at half maximum intensity (FWHM).

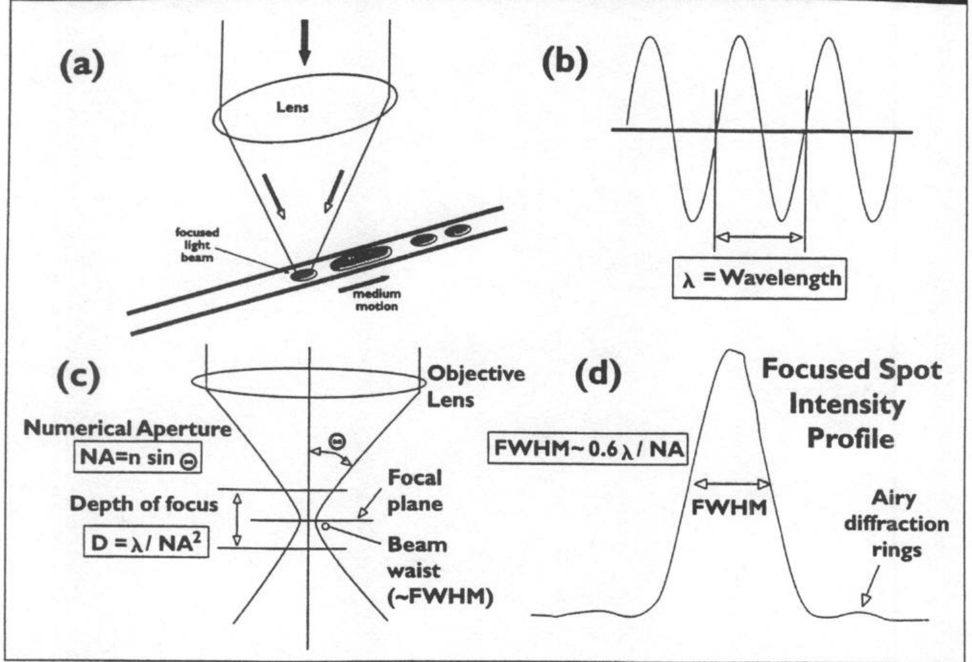

the ground. The interface between the aerodynamically supported "slider" and the disk surface containing the user's precious data must be maintained reliably (without crashes) over the entire average five-year life of the drive. Optical drives are perhaps equally amazing in tracking data "remotely," with about 50,000 times greater separation between the head and the disk. A mechanical servo-control system provides for a 1µm-diameter focused light beam to follow a comparable-size feature on the disk surface to within 0.1µm accuracy while the surface undergoes vibratory motions in two dimensions with amplitudes in the range of 100µm.

In the 25 years since its earliest commercial development, optical data storage has become an industry that generates about 15% of the total annual data storage market revenues of roughly ~$60 billion worldwide, and it probably accounts for a similar fraction of the total installed worldwide digital data storage capacity (~10^{18}B, according to DISK/TREND and the National Storage Industry Consortium of San Diego, Calif.). The most common optical storage devices today are the ubiquitous compact disc (CD) and the emerging digital versatile disc (DVD) used in audio, video, and computer applications for distributing content and information.

As optical disc storage has developed into a robust technology in the removable media sector over the past 15 years, rewritable optical discs have also progressed. The leading rewritable optical media types today are MO and phase-change, which share somewhat complementary attributes. MO technology has led the way in opening markets for rewritable, removable optical storage, with several generations of International Standards Organization (ISO) drives (commonly found in optical jukeboxes) and the Sony MiniDisc (commonly found in consumer applications primarily in the Far East and Europe). This article offers an overview of MO data storage technology [5] and looks ahead to emerging directions that will make aspects of MO technology more important in magnetic disk drives in the future.

The usual advantage of optical disc storage is media removability combined with large-capacity random-access data storage. MO is the rewritable optical-recording method best suited for high-performance, extended-lifetime, high-density applications. MO storage material is a thin magnetic film (similar to magnetic disks). The recording process is thermally assisted magnetic recording (atomic magnet reorientation), known to data storage engineers as a rapid and dependable process with practically infinite cyclability. Phase-change media, while providing a popular and viable low-cost alternative write-once or rewritable option, are inferior to MO in both raw performance (data recording rate and storage density) and cyclability. In phase-change recording, the recording/erasure processes involve crystalline-amorphous atomic structural changes (atoms move around); therefore these processes are slower and more prone to wear-out phenomena.

The similarity of MO and magnetic recording is the basis for an expectation of an exciting synergy in the future. The merger and hybridization of these heretofore independently successful technologies is a possibility, as reflected in a string of research and development efforts publicized over the past five years [1–4, 6, 7].

Role of Optical Data Storage

Optical recording emerged as a viable technology in the 1960s and 1970s when it was made technically feasible by the development of low-cost, compact light sources of adequate power, namely solid-state diode lasers. In classical optical recording, the system designer arranges for an intense, focused light beam to interact with a storage medium (see Figure 1a). Focusing light implies that the location of the light focal point in the medium is relatively remote (a few millimeters) from the optical components that generate and guide the light beam. These so-called "far-field" optics represent a key distinction from magnetic recording, whereby the head is placed very close to the recording medium (today less than 0.1μm).

The optical readback process is performed with the same light beam, though with optical power reduced from the writing level by a factor of from 3 to 6. In reading, the light is used to sense the induced physical change from the writing process, usually through some combination of optical reflection, transmission, and change in the state of the light beam's polarization, or orientation of the internal electric field in the electromagnetic wave.

Some variants of this standard form of optical recording are known to optical storage engineers. One approach is "near-field recording," whereby an optical advantage can be gained by utilizing the light in very close proximity to its zone of emergence from the light generation and guidance device. Another approach is to use "volumetric" optical storage. In one example, the storage medium has multiple recording layers. A focused, far-field light beam is easily refocused at variable depths in the medium, making feasible multilayer storage media, provided the individual layers are sufficiently transparent. Volumetric storage can multiply the storage capacity of a disk or tape significantly, since the third dimension of the storage medium is used, instead of just spreading the data over a plane.

Optical Storage Context

To place MO recording in a useful context, the characteristics of today's common optical storage methodologies need to be defined. Figure 1(a) outlines the process of a focused light beam writing a track of information on a moving optical storage medium. This writing is carried out by modulating the optical power sent to the focusing lens, resulting in local physical change in the recording material. The substrate material carrying the recording medium may be a disk, tape, or card.

The recording configuration in Figure 1(a) is a serial one, because the recorded marks are created sequentially on a track as the medium moves under the focused beam. In a serial digital recording application, this scheme might be called "bit-by-bit" recording. This approach is contrasted with a possibly more parallel method in which multiple beams record multiple data streams concurrently or in which blocks or frames of data might be recorded optically at one instant at various 3D positions in the medium by an extended beam, as in holographic recording (see Orlov's "Volume Holographic Data Storage" in this section). This article is limited to serial recording systems.

Types of optical recording can be differentiated further according to the functionality provided to the user and by the reversibility of the writing process. Audio and computer CDs are the most common types of read-only optical media (often called ROM for "read-only memory"). In this case, the media are pre-written at the production factory, with specified information (such as recorded music or a set of computer program files) replicated thousands of times. The original information content was recorded onto a master disc by a process very much like that in Figure 1(a), then accurately copied by a mass replication process onto low-cost media the end user can read but cannot write on. The replication process for ROM discs is highly parallel; a single disc is made from the master in a few seconds (~10^9–10^{10}B replicated simultaneously). The physical embodiment is a sequence of small, light-scattering pits along the disc track.

ROM optical media are contrasted with two forms of writable media: write-once, read-many (WORM) and rewritable, erasable, or write/read (W/R). WORM media is exemplified by CD-recordable media (CD-R) allowing the user to write information on the disc one time, then read it back an unlimited number of times. WORM media is preferred when a vendor wants to create a certified body of information that cannot be altered without extraordinary means. Consequently, professional-quality WORM optical media has become a legal standard as an information repository when an audit trail must be established and preserved, since ROM and WORM optical media involve essentially irreversible writing processes.

In order to allow multiple recordings, rewritable optical media must be readable at every stage. Therefore, the recording process must be highly physically reversible. However, few physical processes are perfectly reversible, so a rewritable medium also needs a measure of cyclability, that is, an expected count of the number of reliable rewrites the medium can sustain. Because different physical processes are available to implement optical rewritability, engineers find there are different regimes of cyclability for different embodiments of W/R media. MO media excels by meeting an ISO specification to sustain at least 10^6

Figure 2. (a) Writing; (b) readout; and (c) erasure processes in MO recording.

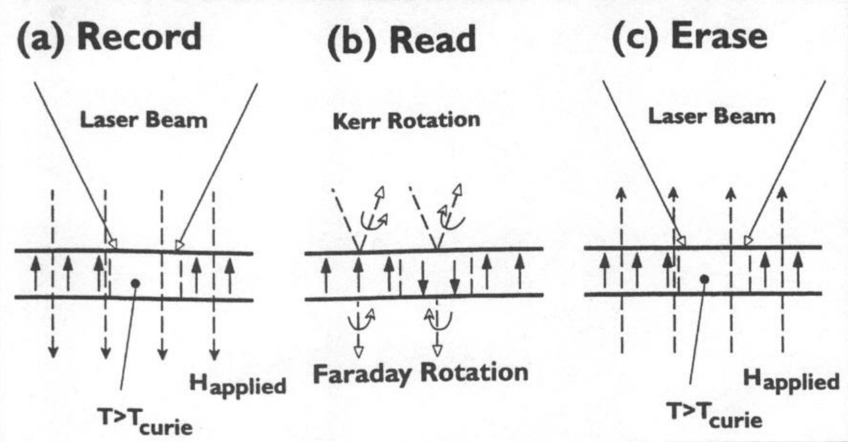

writing cycles and at least 10^7 reading cycles without unacceptable recording degradation. By contrast, CD-RW exemplifies a consumer-rewritable optical-disc product based on a phase-change medium that may have a factor of 10^2–10^4 times lower cyclability than typical MO media.

Note that W/R media (magnetic or optical) are used in ROM or WORM mode when adequate software or hardware protections are in place to prevent inadvertent overwriting or erasure of data.

MO Recording

MO recording is a form of magnetic recording in which light is used as a source of medium heating in the writing and erasure processes (thermomagnetic recording) and as a probe of the magnetic state of the medium in reading (see Figure 2). In each of the figure's three panels, a thin film of magnetic material deposited on a smooth substrate (such as a disc and tape carrier) is shown in cross-section. The magnetic polarization of the MO material is oriented perpendicularly to the film plane, a necessity for MO readout to work. This magnetic orientation requires careful material selection and processing, since it is usually energetically favorable for the magnetic polarization to lie in the plane of the film (longitudinal orientation), as in magnetic disks.

Heating MO media is necessary for writing and erasing, as in Figure 2(a) and (c). In general, magnetic recording is achieved when an applied magnetic field overcomes the medium's resistance to switching, called its "coercivity." All magnetic materials steadily lose their magnetic properties, however, including coercivity, as their temperature is elevated. To record information at high density on a surface implies that the region of controlled switching is very small. In MO recording, this condition is achieved by combining a relatively uniform magnetic field from a coil device with strong localized heating from a focused light beam (see Figure 3). (The applied magnetic field is about 600 times stronger than the Earth's compass-influencing field.) When the medium cools to room temperature, the freshly reversed magnetic polarization is said to be "frozen in."

Two distinct means of thermomagnetic recording of a magnetization pattern along a track on the moving medium are found in MO drives, one using laser intensity (power) modulation (LIM), the other using magnetic field modulation (MFM) (see Figure 4). In LIM writing, the magnetic field is held constant; in MFM recording, the laser power can be kept on continuously or pulsed at exactly the data clock rate.

A binary data bit sequence can be encoded in magnetic domains on the medium in two ways: In pulse position modulation (PPM), the drive records a binary 1 or 0 corresponding to the existence or absence, respectively, of a small circular magnetic domain. Alternatively, the drive may use pulse width modulation (PWM), whereby a binary 1 or 0 corresponds to the existence or absence, respectively, of a magnetization transition from + to – or from – to +. This method is used in magnetic recording (see Figure 4 for a comparison of these methods). PWM encoding has advantages for achieving greater linear-bit density, since even the smallest circular domain encodes binary bits at both its leading and trailing edges. Consequently, information can be packed more densely with an equivalent number of written features.

Magneto-Optical and Magnetic Recording Compared

Because the commercial application of MO recording is mostly in disk drive devices, this article compares MO only with magnetic disk drives. For many years, personal computers have included both floppy disk drives and hard disk drives (HDDs). A floppy disk drive is a relatively low-cost, low-performance device that supports removable magnetic media of relatively low capacity (typically 1.44MB per cartridge). Removable media are convenient for nonelectronic transport of user data files between computers.

An HDD, by contrast, is a device within which the magnetic disk(s) is normally sealed; it provides much higher performance (faster data access and through-

Figure 3. Layout of an MO write/read head-disk system.

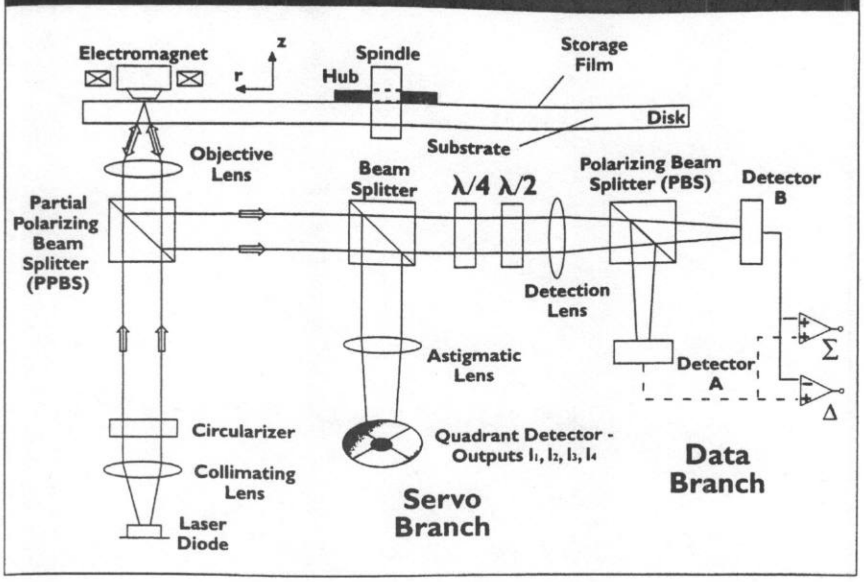

put) and higher capacity than a floppy drive. Magnetic recording involves close-range interaction between the head and medium. A clean environment is critical for maintaining this mechanical interface. Media removability is an obvious convenience to users but introduces a significant reliability risk for the storage device itself. There have been some widely used products with removable hard disks, and as expected, they offer significantly improved performance and capacity compared to floppy drives but are less dependable than conventional HDDs.

Compare this situation to MO drives. All optical drives excel in media removability with superb reliability—a direct consequence of the far-field head being well removed from the disk. When light is focused through the substrate onto the disk's second surface, as in Figure 3, the outer surface need not be pristine. Some amount of surface dust and dirt is tolerated by the optical system, because the light beam is roughly a million times more diffuse when passing through the exposed entry surface of the disk than it is at the focal point. The use of second-surface focusing means it is tolerable to not include a protective cartridge around a CD (though gross contamination causes light blockage, absorption, or scattering that is eventually detrimental to performance).

Besides media removability and device reliability, three other characteristics are of interest to users of disk storage devices: capacity, speed, and cost. When comparing the storage capacity of magnetic HDDs and MO drives, areal density has to be considered. Figure 5 compares average areal density over time for HDDs and MO drives. Until recently, MO recording had an advantage, because optical drives had a much higher track density (number of tracks per unit radial distance on the disk). Since 1990, however, magnetic

Figure 4. Two means—LIM and MFM—of thermomagnetic recording of a magnetization pattern along a track on the moving medium. In PPM coding, binary 1s are represented by the center position of the small nearby circular marks. In PWM coding, binary 1s are represented by the edge marks, which can be of N clock lengths, whereby the range of N depends on the particular code being used.

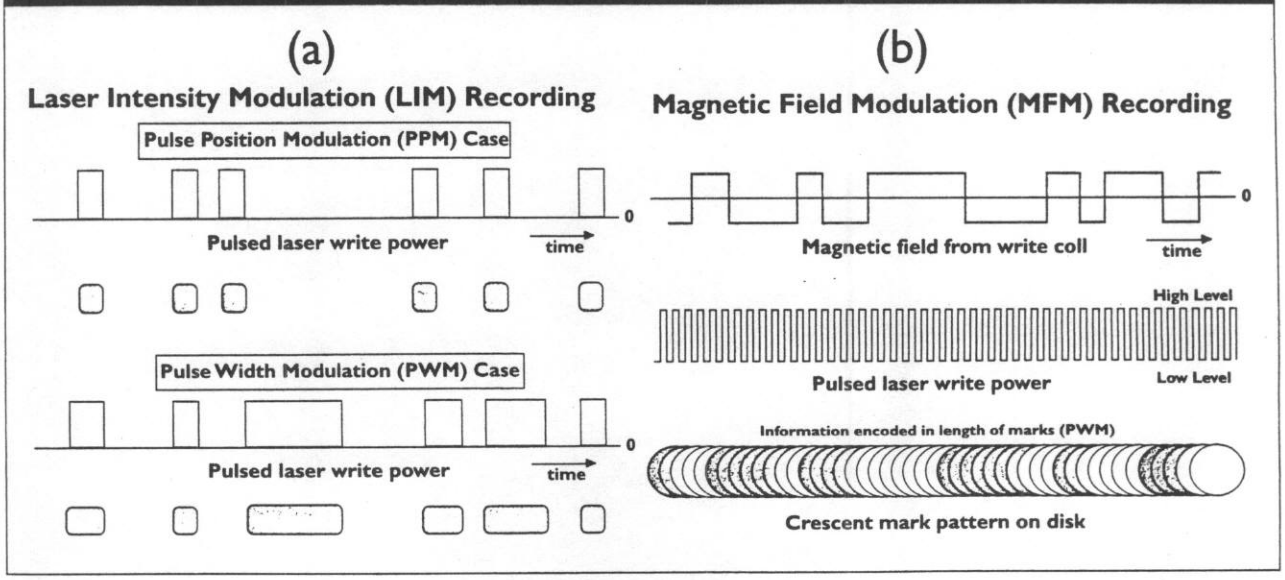

Figure 5. Areal density growth trends in MO (research and products) and HDD recording (products), 1985–2010.

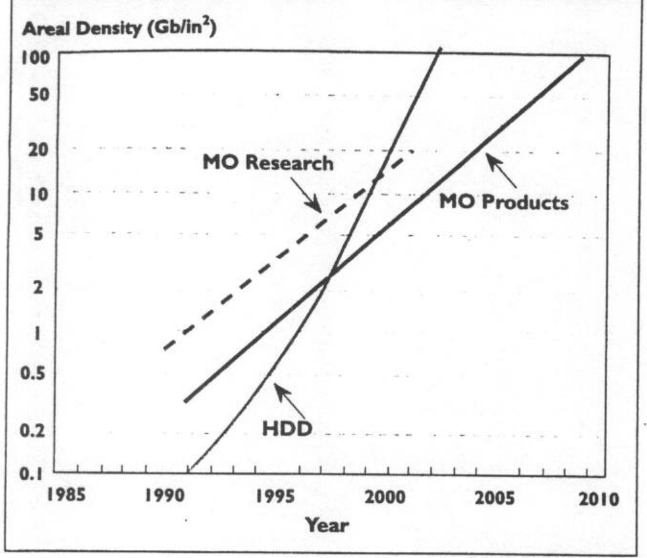

Source: M. Kume, Sanyo

recording has shown a significant increase in the annual growth rate of areal density, initially doubling from the historical average rate of about 30%–60% with the widespread introduction of magnetoresistive head technology. More recently, that rate has increased to 100%–200% annually, depending on the product. The MO recording industry has been managed more conservatively, reflecting an annual areal density growth rate of about 40% from 1992 to 2000; the same rate is now projected out to 2008.

When comparing MO drives and HDDs, it is important to remember that a huge storage capacity advantage for removable media systems results from the fact that capacity per drive is limited in part by the number of disk cartridges associated with a particular drive—a number without theoretical limit. Drive cost comparisons are somewhat complicated, since an HDD is bundled with dedicated media capacity, while an MO drive is not. MO drives are generally considered to be the most expensive of the optical disk drives, because an optical head for MO readout involves specialized optics for sensing the polarization state of reflected light. Moreover, MO drive manufacturing has probably not yet capitalized on economies of scale, generally due to limited market penetration. (Cost is both the cause and the result of relatively modest product volumes.)

Optimal Readout Process

Readout in MO recording is a matter of detecting the pattern of magnetization in the storage medium utilizing rotation of the plane of polarization upon reflection of polarized light from a magnetic mirror—called the magneto-optic Kerr effect, as in Figures 2(b) and 3. The drive's reading power is increased enough to raise the signal-to-noise ratio (SNR) as high as possible without heating the track to a level that degrades the written magnetic information. A system's designer tries to maximize the MO signal while keeping the reflectivity below 30%. The remaining fraction of light power is absorbed during writing and must be high enough to heat the disk efficiently.

A viable data channel in a recording device requires adequate SNR; therefore, the designer must consider noise minimization. The principal noise components (called "media noise") in an MO recording system are associated with the laser, the read-channel electronics, the light detectors, and the disk. Differential detection, as in Figure 3, cancels some but not all of these noise components. SNR is perhaps the single most important parameter governing data-channel performance, and thus has the greatest influence over the storage system performance metrics of interest to the user—capacity and data-throughput rate. In general, great care must go into a disk's optical, thermal, and magnetic designs in order to achieve sufficient and balanced system performance.

New Applications

Since the mid-1990s, several new research thrusts have begun to promise enhanced MO technology applicability and value from this form of rewritable optical storage. There is an important distinction between the schemes to extend the diffraction limit in readout discussed earlier and the more revolutionary approaches to MO recording in terms of system architecture. The following are some of the most notable developments.

High-density MO drive. In late 1995, TeraStor of San Jose, Calif., began developing a novel MO drive using focusing optics with a solid immersion lens—a moderately high refractive index lens placed in close proximity (<<λ) to the disk recording surface (*not* a near-field technique) [4]. This design effectively increases the numerical aperture of the light-incident medium, reducing the spot size, as in Figures 1(c) and (d), and increasing the recording areal-density potential. However, TeraStor disbanded in 2000 without shipping a product.

Flying MO head technology. In 1996, Quinta, a data storage startup company in San Jose, Calif., began developing flying MO head technology. When Seagate acquired Quinta in 1997, this technology was named Optically Assisted Winchester. In it, micro-optics and a microcoil are attached to a flying slider (the carrier for a HDD head), along with an optical-fiber light-delivery system to achieve a low-mass, low-cost means of realizing the first-surface configuration for MO recording. This design introduced MO recording technology into an HDD architecture while preserving HDD performance. Seagate has exhibited prototype drives but has not yet marketed a product based on the original concept.

Hybrid schemes. Hybrid schemes combining magnetic and optical recording elements show promise in the laboratory of outperforming traditional magnetic-only component combinations. Hitachi, Philips, and Sharp have publicized their work in this area [3, 6, 7], while other companies have expressed interest in such approaches. A common element in these methods is the incorporation of laser light into the recording and readout processes. Momentary heating of the recording medium is an effective way to mitigate a looming problem of writability in ordinary magnetic recording by using the thermomagnetic recording concept from MO technology. (Media coercivity in HDDs is being increased steadily to preserve the magnetic stability of ever-shrinking magnetic bit cells; writing heads have a definite physical limitation in their output magnetic switching fields.)

Near-field optical schemes. Research in the early 1990s demonstrated MO recording in domains smaller than 0.1μm by using a near-field optical source. Such an approach overcomes the diffraction limit of light, allowing electromagnetic field dimensions to be determined by the physical extent of the source, such as an aperture in a waveguide and the width of a laser cavity. While this technique allows much smaller optical "spots," such spots exist only in the proximity of the source, with the allowable separation typically on the same order as the source dimension. Moreover, such near-field schemes normally have the drawback that the transmission efficiency for optical power is extremely low. Therefore, much more efficient near-field optical sources need to be developed to produce energy flux sufficient for thermomagnetic recording.

Outlook

Some of these recent thrusts could extend optical recording beyond its historical regimes. MO storage has already proven itself in applications requiring very rugged, highly reliable, removable, rewritable optical media. The devices commercialized so far have won acceptance in the professional and consumer markets that demand high storage capacity and moderate random data-access performance with media removability. These applications are fundamentally different from those addressed by HDDs.

Implementation of new approaches in conventional MO recording shows promise in boosting areal density 10-fold. For example, a number of techniques—magnetic super-resolution (MSR), magnetic amplifying MO system (MAMMOS), and domain well dynamic displacement (DWDD)—improve MO readout resolution without needing a smaller focused light spot, which would be physically limited by the available light wavelength and objective lens. Meanwhile, the convergence of optical storage approaches with magnetic storage is an exciting new development. MO recording offers the key technologies of thermomagnetic recording and patterned media, which may be instrumental in alleviating a slowdown in HDD advances [2, 6]. The HDD industry is beginning to appreciate the potential of these approaches to help sustain its historical progress in cost and performance. Such hybridization would be a remarkable development for the HDD industry, which for the past 40 years has relentlessly pursued incremental progress based on scaling, punctuated by the timely introduction of stepwise improvements in component technology. MO technology is well-positioned to continue the flow of engineering solutions to users with an insatiable appetite for storing information.

REFERENCES

1. Alex, M., et al. Optically assisted magnetic recording (paper WeF-03). In *Proceedings of the Joint Magneto-Optical Recording International Symposium/Asia-Pacific Data Storage Conference* (Nagoya, Japan, Oct. 30–Nov. 2, 2000).
2. Davis, J. Beyond the superparamagnetic limit II: Far-field recording. *Data Stor. 4,* 2 (Feb. 1998), 33.
3. Katayama, H., et al. New magnetic recording method using laser-assisted read/write technologies. In *Proceedings of Magneto-Optical Recording International Symposium on Optical Memory and Optical Data Storage 1999* (Monterey, Calif., Jan. 10–13 1999).
4. Kino, G. Optical background for near-field recording. In *Proceedings of Magneto-Optical Recording International Symposium 1999* (Monterey, Calif., Jan 10–13, 1999).
5. McDaniel, T. and Victora, R., Eds. *Handbook of MO Data Recording,* Noyes Publishers, Westwood, N.J., 1997.
6. Ruigrok, J., et al. Disk recording beyond 100 Gb/in^2: Hybrid recording? *J. Appl. Phys. 87,* 9 (May 1, 2000), 5398–5403.
7. Saga, H., et al. A new recording method combining thermomagnetic writing and flux detection (paper Pd-08). In *Proceedings of the International Symposium on Optical Memory* (Tsukuba, Japan, Oct. 20–22, 1998), 188–189.

Terry McDaniel (Terry.W.McDaniel@seagate.com) is a research staff member in the Research Division of Seagate Technology in San Jose, Calif.

TAPE TECHNOLOGY ROUNDUP

New Cartridges for the Tape Backup Batt

By Kristen Newton

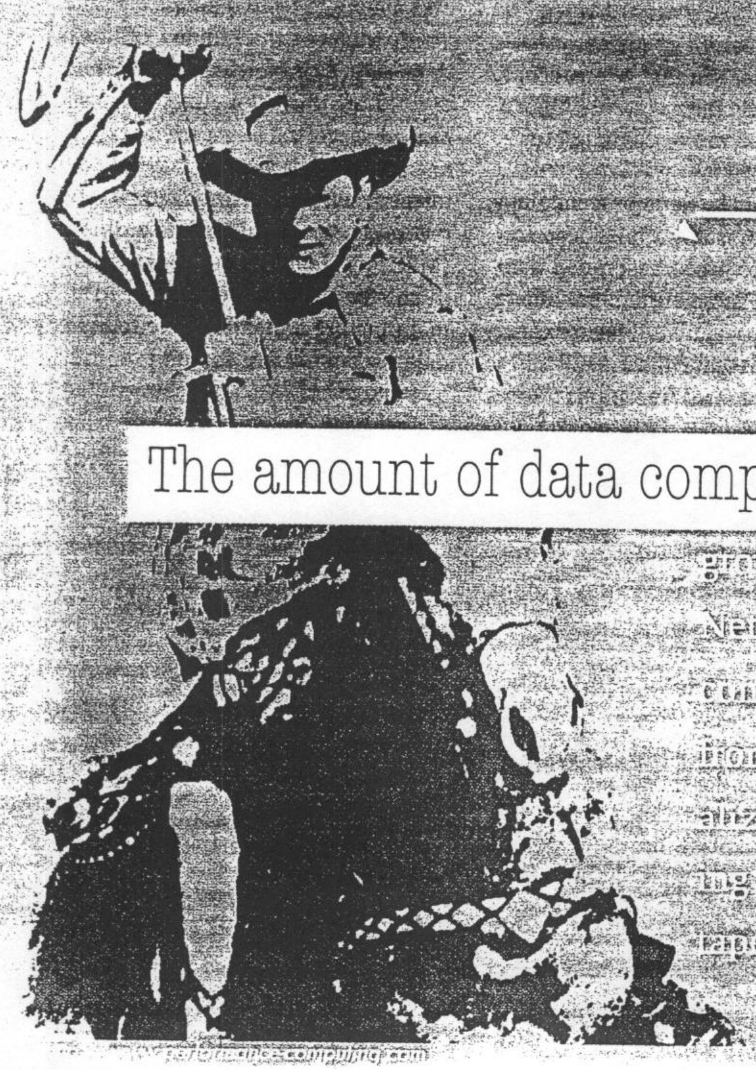

Illustration By Dwight Been

The amount of data companies must maintain is growing at incredibly rapid rates. Network storage is centralizing and consolidating. Storage is separating from the server and becoming virtualized. These changes are generating a great deal of activity in the tape market. As a result, many new

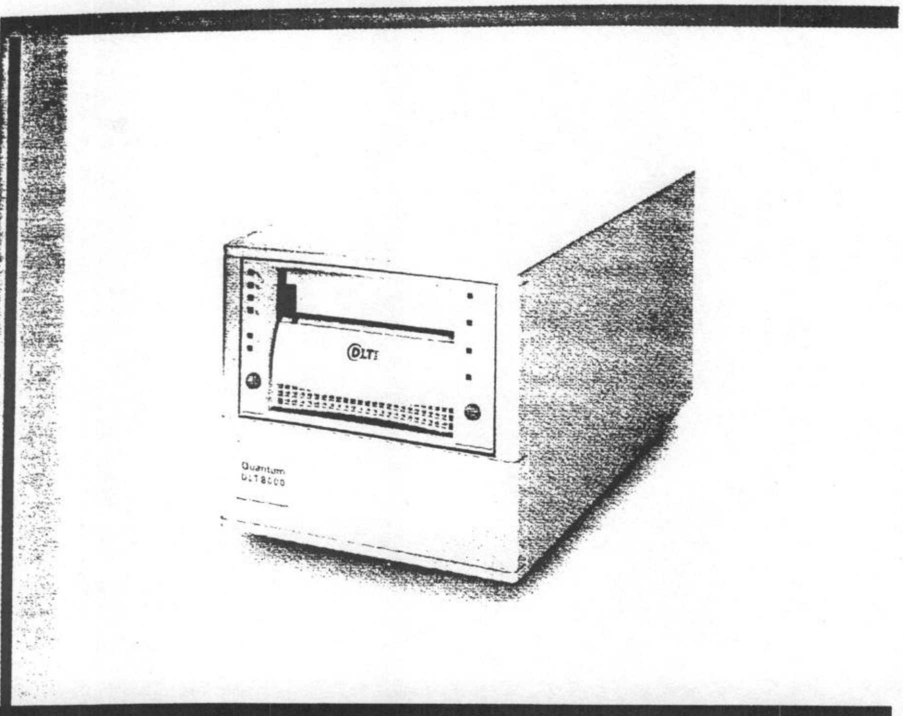

Quantum's DLT8000

tape technologies have been developed. Some are next-generation versions of existing products and others are completely new technologies. Tape technology, with its low cost per megabyte, is the standard solution for data-storage requirements, such as backup and archiving. Tape systems have been pressed to keep up with the increased need for inexpensive storage, as well as address shrinking backup windows and demands for reduced restore times. The upcoming tape technologies all offer increased capacity, performance, and reliability, so how can one sort through all the new product announcements? Here we will present an overview of some of the newest tape systems.

DIGITAL LINEAR TAPE

Quantum is currently shipping the third generation of DLTtape, the market-dominating midrange tape technology. Quantum developed DLT from technology gained from its purchase of Digital Equipment Corp. in September 1994. DLTtape's success is directly tied to its reliability, and high capacity and transfer rates. DLTtape systems use half-inch metal particulate tape in a single reel cartridge. The drive mechanism extracts the tape out of the cartridge using a leader pin and winds it onto the drive reel. Data is written in a linear fashion along the length of the tape.

Quantum's latest tape technology, the DLT8000, is not a huge leap forward from its predecessor, the DLT7000. In fact, it offers only incremental improvements—native capacity has increased 5GB to 40GB (up to 80GB compressed) and transfer rates are improved from 5MB/sec to 6MB/sec. Reliability has also been improved. The DLT8000 offers a recording head life of 50,000 hours and a mean time between failures (MTBF) of 250,000 hours at a nonstop duty cycle. Quantum expects that the improved reliability will decrease the total cost of ownership. The DLT8000 drive will be read/write compatible with all versions of DLTtape. DLT8000 will be competitively priced, but Quantum is not expected to turn the DLT7000 into a value-line product.

Currently, Quantum is positioning the DLT8000 as its flagship product, but this appears to be a defensive move to protect the DLTtape market share against incursions by Mammoth-2, AIT-2, and LTO's Ultrium, until the release of SuperDLT. Expect the DLT8000 to be the last of the "standard" DLTtape drives. To achieve the increases in performance and capacity of SuperDLT, several new and unique technologies have been deployed.

SuperDLT builds on DLTtape technology, but will have huge capacity increases. The first generation is expected to have a capacity of 100GB native, advancing to 1TB in future releases. SuperDLT is based on Quantum's Laser Guided Magnetic Recording (LGMR) technology, which uses a combination of optical and magnetic technologies to significantly increase the number of tracks on a tape. The increases in track density are derived by implementing optical servo tracks on the back, unused side of the media, leaving all of the front side for magnetic data recording. Magneto Resistive Cluster (MRC) heads use small tape heads in clusters, which are cheaper to manufacture and contribute to faster data-transfer rates. To build on the advantages of DLTtape's huge installed base, SuperDLT will be back-

Benchmark's DLT1 Drive

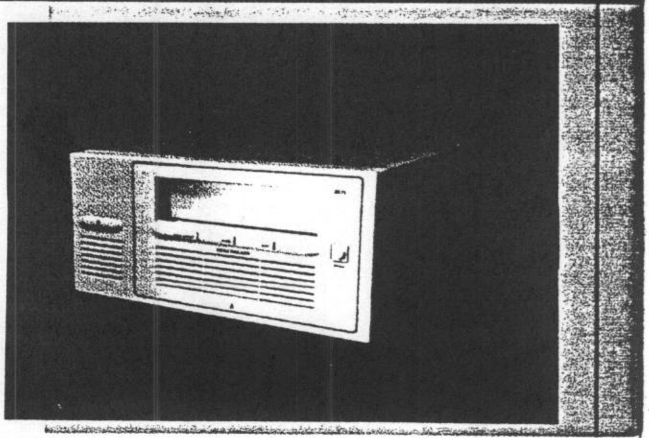

wards-read-compatible with DLT IV cartridges. Test drives are expected to be shipped to OEMs by the end of 1999, with end-user shipments in late 2000. Super DLT will use Advanced Metal Powder media to store high-density data.

Benchmark Tape Systems, through a technology transfer from Quantum, produces a DLTtape-based technology aimed at the NT network server and workstation markets. Benchmark recently began shipping its DLT1 tape drive. Similar to Quantum's DLT8000, the DLT1 has a native capacity of 40GB (up to 80GB compressed). Its transfer rate is half that of the DLT8000 at 3MB/sec, which is more than adequate for the low-end server market. The DLT's reliability figures match those of the DLT7000: 30,000 hours of head life and 200,000 hours MTBF on a continuous duty cycle. Benchmark's DLT1 is backwards-read-compatible with the DLT4000. The DLT1 will not compete with the DLT8000 on product features, however. Its true value is in its price—less than one-fourth that of the DLT8000, an ideal price point for the low-end server market. Benchmark also produces a lower-cost version of DLTtape IV media.

9840

StorageTek is dominant in the half-inch tape and automation markets. Earlier this year it released the 9840 tape format. The 9840 is the first in a series of next-generation tape drives from StorageTek aimed at the server and mainframe tape storage markets. It is the first tape technology intended to bridge the mainframe and open-systems environments. It is a high-performance, relatively low-capacity drive targeted at storage-intensive applications. Its native capacity is 20GB, with up to 80GB compressed. The native transfer rate is 10MB/sec. One highlight is its midpoint load, which reduces time-to-data during restore operations. Average access time is 11 seconds. To accommodate both mainframe and open-systems environments, the 9840 comes with either an ESCON or a SCSI connection. The 9840 fits into StorageTek's automated library architecture. These drives can be added to new or existing libraries, helping StorageTek maintain its hold on the data-center tape market.

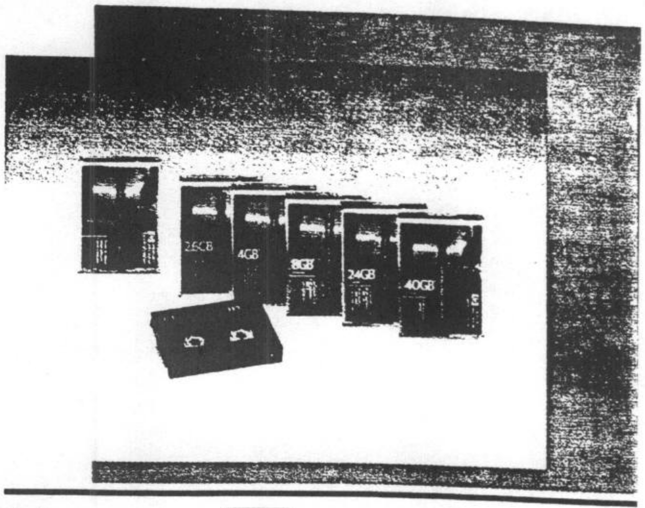

HP's DDS-4 family

LINEAR TAPE-OPEN

LTO began as a collaboration between HP, IBM, and Seagate Technology. These companies saw the need for a truly open tape standard for high-performance and high-capacity tape. The LTO recording format uses multichannel linear recording, where data is written in tracks running the length of the tape. Tracks are grouped into recording bands, bordered by a servo track on the top and bottom of each band. Servo tracks are written onto the media in the manufacturing process. Data is written to the innermost bands first, the most physically stable area on the tape. To allow for read after write, each write head has a corresponding read head. Compression techniques have also been improved upon in the LTO specifications. The compression algorithm used in LTO detects already compressed files and does not attempt further compression. LTO will be available in two formats. Accelis is the high-performance, rapid-data-access version; Ultrium is designed for high capacity and high reliability.

Four generations of each product have been defined. Native capacity and transfer rates are projected to double in each subsequent generation. The recording method for both will be RLL 1, 7 (Run Length Limited, a method of encoding common to magnetic disks) in the first generation and PRML (Partial Response Maximum Likelihood, a way of differentiating valid bits from those generated by noise) in the later generations. Metal Particle media is currently used. The fourth generation will employ thin-film media. Both drives come in four- and eight-head versions. (We will focus on the LTO eight-head flagship drive.)

The technical information required to build compatible LTO products will be made available to licensees. So far, 16 companies (including HP, IBM, and Seagate) have been approved for LTO licenses. Licensees will be periodically tested to verify format compliance to maximize data interchangeability on cartridges between multiple vendors' tape mechanisms. Compliance with the LTO standards ensures a competitive open market that will let consumers make product selection based on price, quality, and vendor relationships.

How LTO will affect the future product plans of HP and IBM remains to be seen. Both companies seem to be aggressively expanding their other tape product lines. LTO appears to be another offering in their tape technology strategies, not their sole solution.

Ultrium

Ultrium, another LTO technology that is based on DLT, retains the better features and overcomes the limiting ones of DLT. Like DLT, Ultrium uses half-inch tape in a single reel design. The tape-loading mechanism is similar to that of DLT: the tape is drawn out of the cartridge by the tape drive through a leader pin and wound onto the drive reel. In the recording format for Ultrium, 384 tracks are divided into four bands of 96 tracks each. The first-generation Ultrium products have a native capacity of 100GB, with data-transfer rates of 10MB–20MB/sec, reaching 800GB and 160MB–180MB/sec by the fourth generation. Ultrium is available in four different cartridge capacities: 10GB, 30GB, 50GB, and 100GB, in its initial release. Ultrium is aimed at applications that require high storage and performance, such as server backup, and will be a direct competitor to DLT. HP, IBM, and Seagate are hoping that the strength of their company names and the open LTO standard will entice end users to shift from DLT.

Accelis

Accelis is similar to IBM's Magstar MP technology. It uses 8mm tape in a two-reel tape cartridge. The tape never leaves the cartridge, which eliminates tape threading (as in Ultrium) and decreases load times. Time-to-data is further reduced through midpoint loading. Equal amounts of tape are wound onto each reel, allowing for the cartridge to be loaded in the middle of the tape, halving the average seek time compared to beginning-of-tape loads. Accelis's recording format has 256 tracks, divided into two bands of 128 tracks each. The first generation Accelis products will have a native capacity of 25GB, a data-transfer rate of 10–20MB/sec, and an average data-access time of less than ten seconds. Later generations will increase native capacity to 200GB and the data-transfer rate to 160–180MB/sec, while decreasing average data-access time to less than seven seconds. Accelis is positioned for use in libraries, offering high-speed access to relatively small files. It is well-suited for near-line applications such as image retrieval, data inquiry, search, and data mining, as well as backup and restore, especially when restore time is important.

HELICAL-SCAN (8MM AND 4MM)

Helical-scan tape technology originated in the audio and video markets. In the late 1980s, this technology migrated to the data-storage market, providing a ten-fold increase in storage capacity over existing products. Helical-scan technology records data onto relatively short tracks that lie at an angle to the edge of the tape. Because it is easier to write tracks in a straight line over short distances, helical scan technology allows for high track density. The tape moves slowly along a high-speed rotating head. Both Exabyte and Sony have been advancing the 8mm product line.

Mammoth

Exabyte's MammothTape technology was designed specifically for backup and data storage, an improvement over previous versions of 8mm technology, which remain closely linked to audio/video applications. Mammoth uses Advanced Metal Evaporative (AME) media in a two-reel cartridge. The first generation Mammoth drive offers a native capacity of 20GB, with a 3MB/sec data-transfer rate. Mammoth-2 will have a native capacity of 60GB, and a data-transfer rate of 12MB/sec. Mammoth-2 is aimed at the UNIX and NT network server markets.

Comparing Tape Formats

Format	Manufacturer	Native Capacity	Compressed Capacity	Native Transfer Rate	Media Type	Drive MSRP	Availability
DLT 8000	Quantum	40GB	80GB	6MB/sec	Metal Particle	$6,000	July 2000
DLT1	Benchmark Tape Systems	40GB	80GB	3MB/sec	Metal Particle	$1,299	Currently Available
Super DLT	Quantum	100GB	200GB	10MB/sec	Advanced Metal Powder	TBA	Late 2000
Ultrium	TBA (LTO)	100GB	200GB	10–20MB/sec	Metal Particle	TBD	Late 1999
Accelis	TBA (LTO)	25GB	50GB	10–20MB/sec	Metal Particle	TBD	Late 1999
9840	StorageTek	20GB	80GB	10MB/sec	Advanced Metal Particle	$27,400	Currently Available
Mammoth-2	Exabyte	60GB	120GB	12MB/sec	AME	$4,995	Fall 1999
AIT-2	Sony	50GB	100GB	6MB/sec	AME	$2,499	N/A
DDS-4	Sony, Hewlett-Packard	20GB	40GB	1–3MB/sec	Digital Audio Tape	$1,149	Currently Available
VXA	Ecrix	33GB	60GB	3MB/sec	AME	$1,295	Currently Available

Exabyte has some credibility issues to overcome, due to its slow time-to-market with Mammoth. However, it has laid out a credible migration path for the Mammoth family. The next generation is expected to have a native capacity of 120GB and a data transfer rate of 18MB/sec.

AIT

Sony developed its Advanced Intelligent Tape (AIT) format to handle the high-capacity requirements of large, data-intensive applications. AIT's differentiating feature is its Memory-in-Cassette (MIC), a 16Kbit memory chip built into the tape cartridge. The MIC chip provides improved data access by letting the tape be loaded or unloaded at any of its 256 partitions. AIT-2 utilizes AME media and has a native capacity of 50GB and a 6MB/sec data-transfer rate. Sony intends to release subsequent generations of AIT every two years, with capacity and transfer rates doubling in each generation. AIT-3 is projected to have a native capacity of 100GB and a 12MB/sec data-transfer rate.

DDS

The 4mm Digital Data Storage (DDS) tape format was co-developed by HP and Sony in the late 1980s. DDS is the market leader in the low-end server market, with over 5.5 million drives installed. Both HP and Sony have announced DDS-4 drives and appear to be advancing the product line. The DDS-4 drive will have a native capacity of 20GB and a transfer rate of 1MB–3MB/sec, depending on the drive manufacturer. To achieve the increased capacity, the DDS-4 tape has been lengthened and the tracks reduced in width. Benefiting the large installed base, the DDS-4 will be backwards-read-compatible with DDS-DC and backwards-read- and write-compatible with DDS-2 and DDS-3 tapes.

VXA

Ecrix's VXA products use a new technology that overcomes the fixed-tape-speed limitations of traditional streaming tape technologies (both helical-scan and linear tape formats). Typically, servers are unable to deliver data at a constant speed or at the rate expected by the drive. Streaming drives are designed to take data from the server and fill the drive buffer, which then feeds data to the drive. In practice, the drive frequently drains the buffer faster than the server can fill it. This causes interruptions to the flow of data to the drive, causing the tape to stop, wait for the buffer to fill, back up, re-accelerate to the proper speed, and then continue recording data. This process (called backhitching) reduces the effective data-throughput rate, increases backup and restore times, and decreases media life. The performance and reliability of streaming technology depends on the geometry of the data tracks. Changes to the straightness of the tracks and the alignment between a track and the edge of the tape caused by normal wear and tear can result in an unreadable tape. Additionally, the alignment of the drive heads can drift, decreasing the ability of like drives to exchange data.

VXA does not use streaming technology. Rather, it reads and writes data in packet form. VXA uses an algorithm to break long strings of data into data packets. The packets comprising a single data string may arrive in the buffer at different times, but the drive is able to reassemble them in the proper order. Each data packet includes 64 bytes of data, a synchronization marker, unique address information, and a Cyclical Redundancy Check and Error Correction Code. Each track contains 387 data packets.

VXA also utilizes Variable Speed Operation, which alters the tape speed to match the host's data-transfer rate, thereby eliminating backhitch problems. Because the drive matches the host-transfer rates, a complete stop of the drive is unlikely, but may occur. In this situation, reads and writes are allowed from "Ready Mode," which operates the drive below normal operating speed. Ready Mode increases performance because the drive does not need to fully accelerate before it can begin reading or writing. Over-Scan Operation (OSO) is used to eliminate issues with track geometry. In OSO, the drive heads cover more than just the recorded area on the tape, ensuring that all data packets are read.

VXA has a native capacity of 33GB and a data-transfer rate of 3MB/sec. Similar to Mammoth-2 and AIT-2, VXA uses AME media.

MAKING CHOICES

Obviously, DLTtape and Ultrium will be direct competitors for market share in the high- to midrange-server-backup market. Quantum is expecting DLTtape's reputation and large installed base to make it the winner. However, Ultrium does pose a challenge with its open standard and re-engineering of the DLTtape design. Both Exabyte and Sony have credible offerings and are making a run at the high-capacity drive market, but their impact is yet to be seen. VXA tape competes with the DLTtape, Ultrium, Mammoth, and AIT options. It is priced to be competitive with current-generation DDS products, which makes it attractive to budget-conscious businesses. DDS-4 competes by extending the life of the DDS-installed base, reducing the chances of those consumers migrating to AIT, DLTtape, Mammoth, or Ultrium. Benchmark's DLT1 could prove to be a wildcard in this arena if it demonstrates proven reliability and a plausible migration path.

Late delivery of Mammoth by Exabyte left a void in the high capacity tape market. DLTtape took advantage of this void and became the market leader. Many of the products discussed here are slated for release in similar time frames. How each fares may be determined by which one reaches the market first. The one you select should be determined by your application requirements and budget.

Ralph Barker

TALE OF THREE TAPES

In last month's special report on tape technologies ("Tape Technology Roundup," p. 48), we provided an overview of advances in that segment of the market. However, we also wanted to test some of the new drives available within tape families that have a history in high-end-computing markets, or are likely to have appeal for enterprise-level backup. Although Exabyte Corp.'s Mammoth 2 drives were not shipping at press time, we were able to bring three other new drives into our lab for testing—Ecrix Corp.'s new VXA-1 drive, Quantum Corp.'s DLT 8000 drive, and Sony Corp.'s DDS-4 drive. All of the tested units were standalone external drives.

As noted in last month's special report, the DLT 8000 and DDS-4 extend their respective technologies to new capacities and data rates. The DLT 8000's capacity is 40GB (uncompressed), with a data-rate specification of 6MB/sec. The DDS-4's capacity is 20GB (uncompressed), and records at 1–3MB/sec according to specs. Ecrix's VXA-1 is the new kid on the block, using an entirely new packet-recording technology, and has a capacity of 33GB (uncompressed). While the DLT 8000 and DDS-4 drives are traditional streaming technologies that depend on the host's ability to supply data at a rate sufficient to keep the drive active, the VXA-1's data rate varies based on the host system's rate of data flow—if that flow is slow, the VXA-1 decreases the recording speed to match. The stated maximum data rate for the VXA-1 is 3MB/sec.

While data rates quoted by tape-drive manufacturers are good for comparison purposes, those specifications rightly ignore the variables associated with actual use of the drive in a production environment. To get a sense of the levels of performance you might see in a non-testbed situation, we elected to test the drives by performing normal backups on a conventional UNIX workstation. Additionally, we were interested in providing a baseline for performance expectations, and selected a system that is decidedly less than state-of-the-art—a Sun SPARCstation 5, equipped with Fujitsu's 170MHz TurboSPARC CPU. For our tests, we used both the system's onboard Narrow SCSI-2 controller, and a SunSwift SBus adapter (a Fast Wide SCSI-2 controller with a built-in 10/100Base-T Ethernet port).

SETUP AND OPERATION

Although most tape drives can utilize the standard st driver, it is best to modify the */kernel/drv/st.conf* file with the geometry parameters for the drive you are using. Doing so enables the creation of the appropriate device nodes during a `boot -r` operation. Both Ecrix and Sony include the necessary information in the documentation that comes with the drives, and Quantum has the data available on its Web site (*http://www.quantum.com/app_notes/app_note_018.htm*). Note that the syntax of the *st.conf* file lets you define a maximum of four densities for a particular drive designation. For example, the DLT 8000 uses a value of 0x88 for its 40GB uncompressed mode, and 0x89 for 80GB compressed. Additional values can be selected based on the types of cartridges and densities you

The VXA-1, DLT 8000, and DDS-4 drives

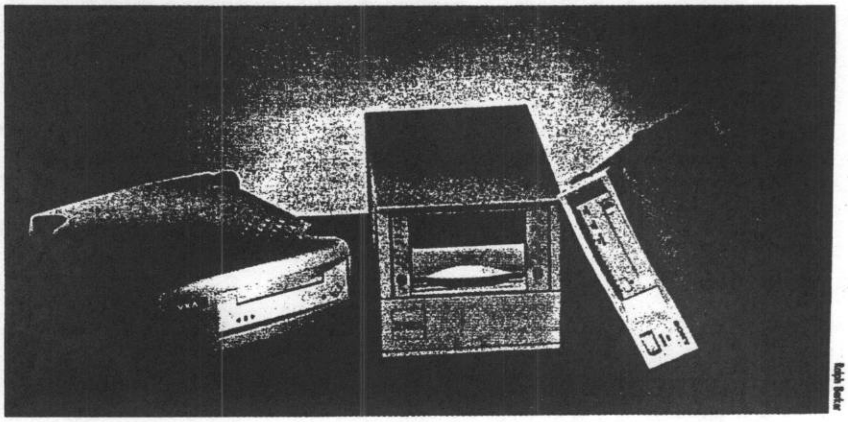

Ralph Barker

need to support within the DLT family.

None of the drive installations presented any problems for us during our testing. Remember, however, that the *st.conf* file's syntax, like any kernel configuration file, is highly sensitive to typographical errors. If you introduce a syntax error when modifying the file, you will get a series of error messages during the subsequent reboot. To correct such problems, fix the syntax error and reboot again.

One of the less scientific elements of systems administration is getting to know the clicks and whirs of the tape drives you use. After a few backups with a particular drive, your subconscious takes over, and you can tell whether your backup operation is progressing well based on the sounds coming from the drive. These new drives are more challenging, as they are generally quieter than their predecessors. The VXA-1 from Ecrix is particularly quiet, so you will likely find yourself relying more on the front-panel indicators than on the sound of the drive. On the other hand, if your work environment is already noisy enough to suit your tastes, you may welcome the quieter characteristics of the new drives.

PERFORMANCE

As mentioned earlier, the performance specifications from tape drive manufacturers do not always reflect real-life figures. These specs describe the maximum capability of the drive under lab conditions. Actual performance will be affected by a combination of factors, including system bus speed, the type of SCSI interface used, and the level of other activity on the system. Additionally, greater variance among these factors will likely be encountered for single-drive configurations, as individual drives will often be used for standalone, older-generation systems. In contrast, multiple drives in library configurations will frequently benefit from upgrades performed on the servers to which the tape libraries are attached.

To understand these variations, we used both the native SCSI interface on our test host and a newly installed SCSI adapter of more recent vintage. The upgraded SPARC Station 5's native adapter is a Narrow SCSI-2, and our SCSI upgrade was a SunSwift adapter—Fast Wide SCSI-2. In backing up the same 1.3GB file system, we saw throughputs of 2.2MB/sec for the Ecrix VXA-1, 2.1MB/sec for the Quantum DLT 8000, and 2.2MB/sec for the Sony DDS-4, all averaged across the results for ten runs on the native adapter. Moving the drives to the SunSwift adapter produced average throughputs that were about ten percent faster—2.4MB/sec for the Ecrix VXA-1, 2.3MB/sec for the Quantum DLT 8000, and 2.4MB/sec for the Sony DDS-4.

CONCLUSIONS

While our tests showed average data rates substantially lower than manufacturer specifications, it is important to view those results in the proper perspective. We intentionally chose an older host as a means of providing something close to baseline performance expectations. Newer systems with faster CPUs and faster system buses will likely produce substantially better results—typically somewhere between our results and the manufacturer's specs. Your choice among these drives, however, will seldom be based on speed alone.

Excluding tape drives chosen for newly formed companies with newly acquired systems, all tape systems need to fit into an existing backup architecture, and long-term plans for the evolution of that architecture. As such, your decision to upgrade tape drives will be strongly influenced by your existing tape-drive selections, so there is no clear-cut "winner" among the drives we tested. The Quantum DLT 8000 and the Sony DDS-4 represent logical evolutions of those two technologies, mostly in terms of increased capacity.

It is unlikely, however, that you would be making a choice between these two technologies, as these drives are aimed at different segments of the market. The 20GB uncompressed capacity of the DDS-4, and backward-read and -write compatibility, make the Sony DDS-4 drive a good choice for workstations and small servers, particularly in organizations that have existing 4mm DDS drives. In contrast, Quantum's DLT technology is intended for larger installations, and is used most often in library configurations, where the duty cycles on tape drives are considerably heavier.

The improvements in the DLT 8000 over its immediate predecessor, the DLT 7000, are less dramatic than those of the DDS-4 compared to the DDS-3. The DLT 8000's capacity of 40GB (uncompressed) at 98,250bpi is an incremental change to the DLT 7000's capacity of 35GB at 85,937bpi—not a striking improvement. While the DLT 8000 is a good choice for new libraries that need to be installed in the short term, upgrading existing libraries is a different matter. Many sites committed to the DLT family may prefer to wait for SuperDLT before making upgrades to libraries currently equipped with DLT 7000 drives.

The Ecrix VXA-1 is the wild card here. The VXA-1 represents a new approach to recording technology—that of packet recording. Ecrix's packet technology substantially reduces the risk of failed restores due to spot damage to the media. As proof of concept, Ecrix offers the 100%-restore results of tests where VXA-1 tapes have been boiled, frozen, and subjected to coffee spills (see *http://www.vxatape.com/extreme/* for that story). The VXA-1 technology was sufficiently innovative to win one of our Outstanding Product Awards in 1999 (December 1999, p.14). With suggested list prices of $899 for internal and $989 for external drives, the VXA-1 presents a compelling alternative to other technologies for workstation and small server backups. So far, however, none of the major library vendors have released large systems based on the VXA-1 drives. We expect that to change, since the VXA-1 is an 8mm technology, and it should be a simple matter to adapt existing 8mm libraries to the VXA-1 drives. But only time and market pressures will decide when.

The Inkjet Color Image, Born in a Rainbow of Tiny Drops

By MATT LAKE

TEN years ago, the computer industry gave desktop computer owners a choice of options when it came to printers. But it was pretty much a Henry Ford kind of choice: you could print in any color as long as it was black.

At the time, most people used typewriter-style daisy-wheel printers or clacking dot-matrix printers to put ink on paper. If you wanted a silent printer, your cheapest option was an inkjet machine — and back then, that would cost you more than $1,000.

Like most computer hardware, inkjet printers have plummeted in price. New models start at less than $200 and are soon discounted below $100. Some computer dealers literally give inkjet printers away as bonuses with the purchase of new computers.

As the costs have gone down, the print quality has improved. Many inkjet printers can produce color pages that are almost indistinguishable from photographs. The quality is so high, in fact, that inkjet printers are widely cited as one of the reasons the United States Treasury has been redesigning its paper currency in the past few years — it was becoming too easy to print passable counterfeit bills with a good inkjet machine.

It is strange that a technology that produces such sophisticated results began with an accident involving a soldering iron. The story goes that in 1977, an engineer at Canon put a hot soldering iron a little too close to an ink-filled syringe. The heat boiled a tiny amount of the ink in the needle, making it expand into a gas. That pushed some ink out of the tip of the needle.

The physics behind this accident became the foundation of what Canon called bubble jet technology, a printing technique in which electrical resistors heat up tiny bubbles of ink, pushing ink droplets through an array of hair-thin nozzles.

Other inkjet printer makers, including Hewlett-Packard, Lexmark and Xerox, have adapted Canon's thermal technology for their own printer lines, adding enhancements (and patents) along the way.

Another printer manufacturer, Epson America, developed its own variation on the theme. Epson printers use piezoelectric crystals to do the job of the bubbles in inkjet printers. Piezo elements change shape when subjected to an electric charge; they are used in loudspeakers to convert signals from a stereo tuner into sound waves. In Epson's line of inkjet printers, microscopic piezo elements subjected to tiny charges are distorted, pushing out tiny droplets of ink.

Of course, there is more to inkjet printing than the firings of microscopic squirt guns. It is a complicated dance of motors that move a print head with 300 to 600 nozzles across a sheet of paper, squirting 5,000 or more droplets of ink per second in a precise minute pattern.

Most color inkjet printers use four colors of ink: cyan (a pale blue), magenta, yellow and black. They are combined to create all the colors the eye sees in a finished printout. To orchestrate this technique, the computer that is attached to the printer uses driver software that determines the exact color for each dot that will appear on the page. It then figures out how many droplets of each ink color must fall on each dot on the page to make that color. Even when printing at 300 dots per inch (a fairly modest resolution), a letter-size page is more than 2,000 dots wide.

Since photographic prints contain the most complicated shades of color, it is a huge mathematical task for the software to figure out the combination of ink colors needed for each dot. And rendering that subtle color with only four inks can be slow. That is one reason manufacturers are producing desktop printers specifically for photographic printing. These printers use six ink colors, adding paler shades of cyan and magenta to help create realistic halftones.

Despite all that precision, inkjet printers must contend with one great unknown that can radically affect the look of a printout: the paper it is printed on.

With the quick-drying inks that Epson, Hewlett-Packard and others use and the tiny ink drops that current printers squirt out, it is possible to get decent text and spot colors even on cheap photocopier paper. But big blocks of color, as in photos, can soak and buckle thin paper stock. The higher contrast and lower absorbency of brilliant white inkjet paper handles color better. And naturally, the best results of all come from glossy photographic paper (which is what you would hope for after paying the premium price of about a dollar a sheet for it).

You can ensure good results by tinkering with the software settings. If inkjet software can calculate how much yellow and magenta and cyan ink to spit onto a dot the thickness of a human hair, you can be sure that it can adjust that calculation for more or less absorbent paper. Most printer driver software provides a setting for different paper types, so when you put inkjet or photo paper into the tray, make sure that you adjust the software to compensate for it. Just don't do it to print facsimiles of currency.

On the Firing Line

All inkjet printers have a print head containing 300 to 600 tiny firing chambers, each filled with either black, magenta, cyan or yellow ink.

PRINT HEAD

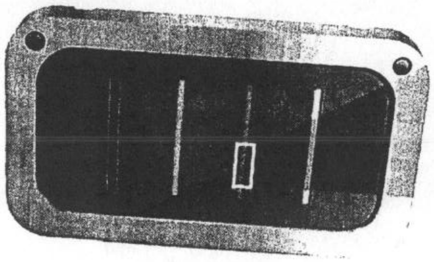

NOZZLES

INKJET PRINTER

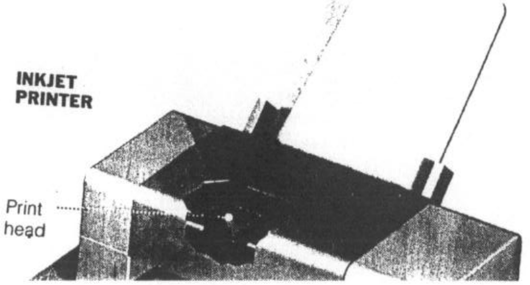

BUBBLE JET PRINTING

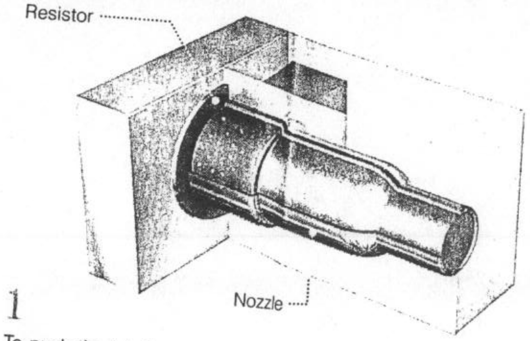

1

To push the ink through the nozzle, bubble jet printers heat minuscule quantities of ink by passing an electrical charge through a resistor, which quickly reaches 900 degrees Fahrenheit.

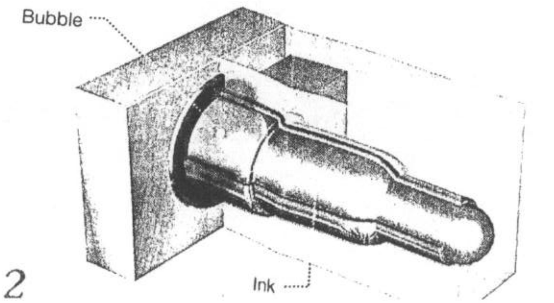

2

PIEZOELECTRIC PRINTING

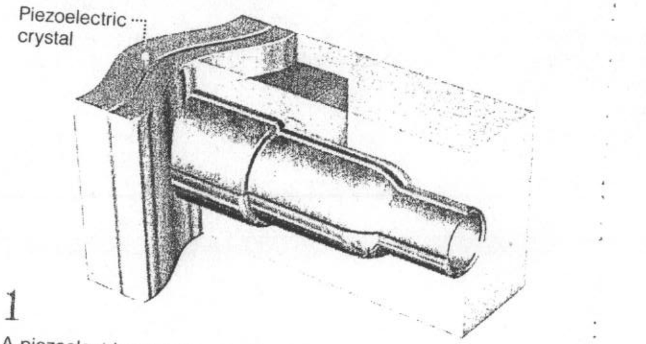

1

A piezoelectric printer works like a squirt gun, but instead of a trigger and plunger, it uses a piezoelectric crystal that changes shape when an electrical charge is applied. A small negative charge deflects the crystal away from the chamber, creating suction.

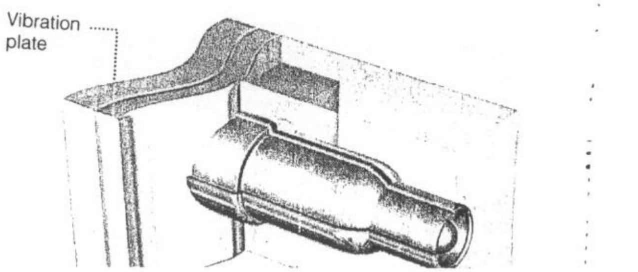

past few years — it was becoming too easy to print

scopic piezo elements subjected to tiny charges are

ink colors, adding paler shades of cyan and magenta to

software to compensate for it. Just don't do it to print facsimiles of currency.

On the Firing Line

All inkjet printers have a print head containing 300 to 600 tiny firing chambers, each filled with either black, magenta, cyan or yellow ink.

PRINT HEAD

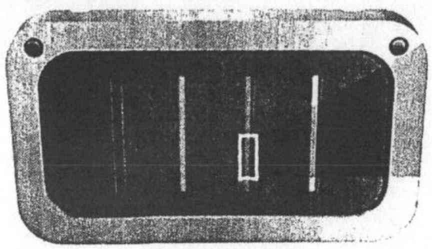

NOZZLES

INKJET PRINTER

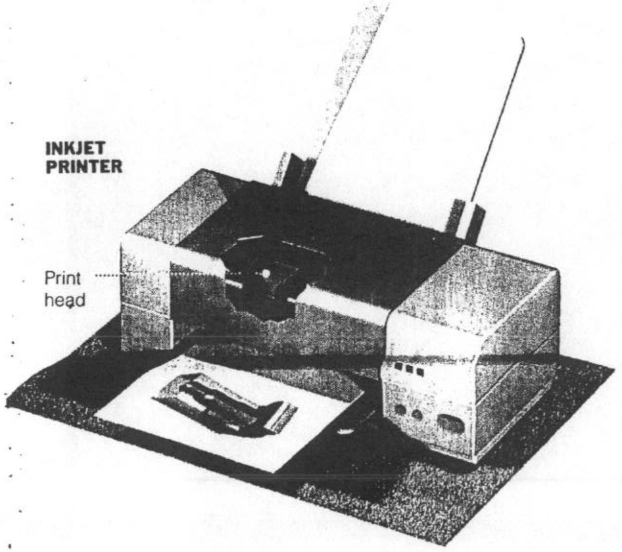

Software drivers translate an image file from the computer into a pattern of microscopic dots of different colors that blend together to look like a continuous image to the naked eye. That pattern is translated into a sequence of firing instructions for the print head as it moves horizontally in the printer.
Most inkjet printers spray the ink on the page by using one of two different mechanisms: the bubble jet (also known as thermal) or piezoelectric techniques.

Sources: Epson America Inc.; Hewlett Packard

BUBBLE JET PRINTING

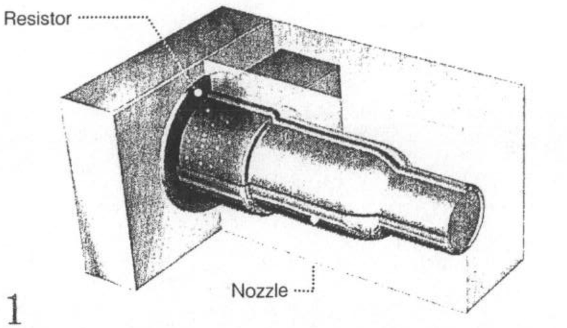

1

To push the ink through the nozzle, bubble jet printers heat minuscule quantities of ink by passing an electrical charge through a resistor, which quickly reaches 900 degrees Fahrenheit.

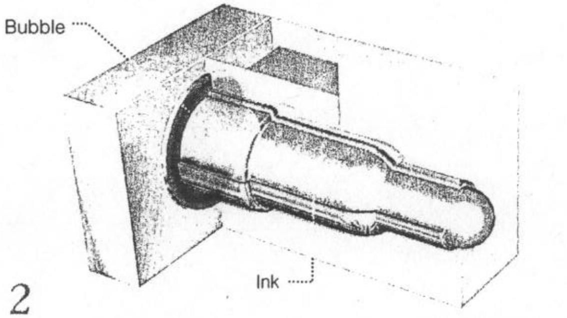

2

The heating element vaporizes a tiny layer of ink at the bottom of *the chamber* for a few millionths of a second, forming a bubble, pushing the ink down the nozzle.

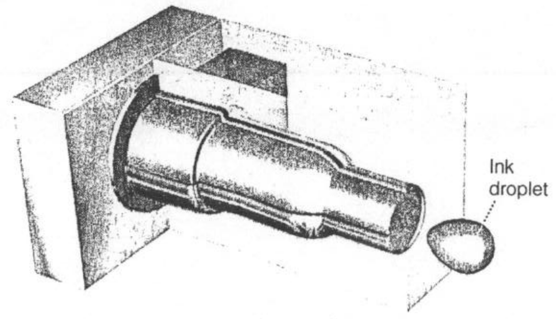

3

The bubble expands and forces a droplet of ink out of the nozzle. Colored droplets are generally so small that a quart could contain 100 billion such drops or more; black droplets are about four times as big. The whole process takes about 10 millionths of a second.
The heating element cools and the bubble collapses, creating suction that draws more ink into the chamber from the ink cartridge.

PIEZOELECTRIC PRINTING

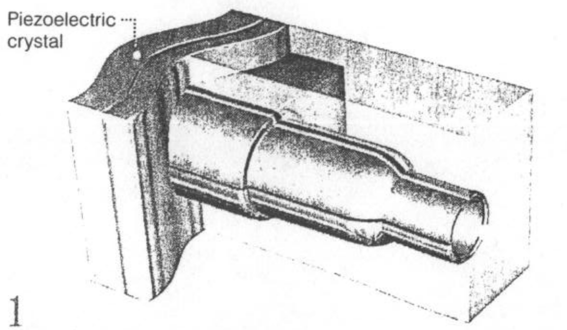

1

A piezoelectric printer works like a squirt gun, but instead of a trigger and plunger, it uses a piezoelectric crystal that changes shape when an electrical charge is applied. A small negative charge deflects the crystal away from the chamber, creating suction.

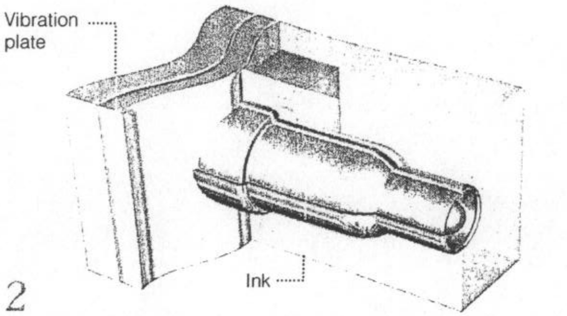

2

A positive charge bends the crystal in the other direction, which pushes a plate into the chamber to create the pressure to *expel a* droplet of ink.

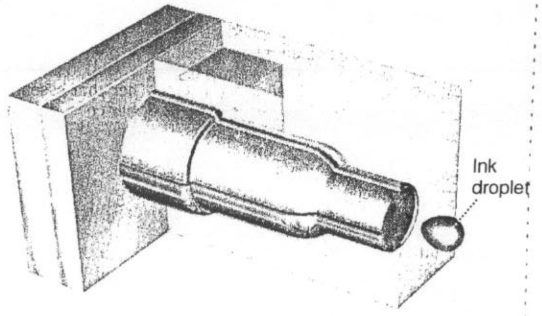

3

An advantage to this approach is that the quantity of ink in the droplet can be precisely controlled. A small charge causes a slight deflection, enough to discharge as little as three-trillionths of a liter of ink through the nozzle. Larger charges can produce larger droplets. The ink is forced out through the nozzle

Illustration by Frank O'Connell/The New York Times

WHAT'S NEXT

The Inkjet Color Image, Born in a Rainbow o

By MATT LAKE

TEN years ago, the computer industry gave desktop computer owners a choice of options when it came to printers. But it was pretty much a Henry Ford kind of choice: you could print in any color as long as it was black.

At the time, most people used typewriter-style daisy-wheel printers or clacking dot-matrix printers to put ink on paper. If you wanted a silent printer, your cheapest option was an inkjet machine — and back then, that would cost you more than $1,000.

Like most computer hardware, inkjet printers have plummeted in price. New models start at less than $200 and are soon discounted below $100. Some computer dealers literally give inkjet printers away as bonuses with the purchase of new computers.

As the costs have gone down, the print quality has improved. Many inkjet printers can produce color pages that are almost indistinguishable from photographs. The quality is so high, in fact, that inkjet printers are widely cited as one of the reasons the United States Treasury has been redesigning its paper currency in the past few years — it was becoming too easy to print passable counterfeit bills with a good inkjet machine.

It is strange that a technology that produces such sophisticated results began with an accident involving a soldering iron. The story goes that in 1977, an engineer at Canon put a hot soldering iron a little too close to an ink-filled syringe. The heat boiled a tiny amount of the ink in the needle, making it expand into a gas. That pushed some ink out of the tip of the needle.

The physics behind this accident became the foundation of what Canon called bubble jet technology, a printing technique in which electrical resistors heat up tiny bubbles of ink, pushing ink droplets through an array of hair-thin nozzles.

Other inkjet printer makers, including Hewlett-Packard, Lexmark and Xerox, have adapted Canon's thermal technology for their own printer lines, adding enhancements (and patents) along the way.

Another printer manufacturer, Epson America, developed its own variation on the theme. Epson printers use piezoelectric crystals to do the job of the bubbles in inkjet printers. Piezo elements change shape when subjected to an electric charge; they are used in loudspeakers to convert signals from a stereo tuner into sound waves. In Epson's line of inkjet printers, microscopic piezo elements subjected to tiny charges are distorted, pushing out tiny droplets of ink.

Of course, there is more to inkjet printing than the firings of microscopic squirt guns. It is a complicated dance of motors that move a print head with 300 to 600 nozzles across a sheet of paper, squirting 5,000 or more droplets of ink per second in a precise minute pattern.

Most color inkjet printers use four colors of ink: cyan (a pale blue), magenta, yellow and black. They are combined to create all the colors the eye sees in a finished printout. To orchestrate this technique, the computer that is attached to the printer uses driver software that determines the exact color for each dot that will appear on the page. It then figures out how many droplets of each ink color must fall on each dot on the page to make that color. Even when printing at 300 dots per inch (a fairly modest resolution), a letter-size page is more than 2,000 dots wide.

Since photographic prints contain the most complicated shades of color, it is a huge mathematical task for the software to figure out the combination of ink colors needed for each dot. And rendering that subtle color with only four inks can be slow. That is one reason manufacturers are producing desktop printers specifically for photographic printing. These printers use six ink colors, adding paler shades of cyan and magenta to

On the Firing Line

All inkjet printers have a print head containing 300 to 600 tiny firing chambers, each filled with either black, magenta, cyan or yellow ink.

PRINT HEAD

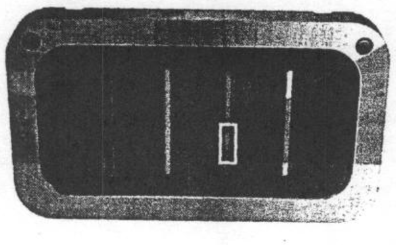

NOZZLES

INKJET PRINTER

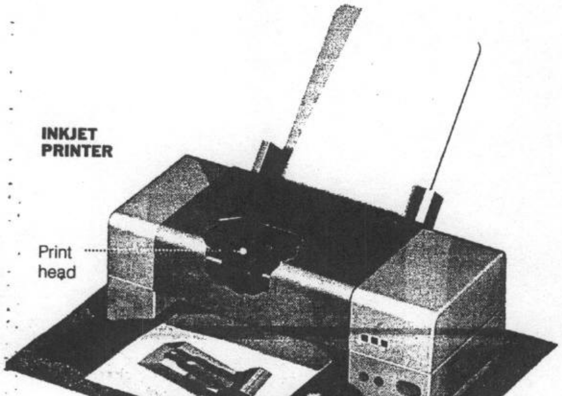

Software drivers translate an image file from the computer into a pattern of microscopic dots of different colors that blend together to look like a continuous image to the naked eye. That pattern is translated into a sequence of firing instructions for the print head as it moves horizontally in the printer.

Most inkjet printers spray the ink on the page by using one of two different mechanisms: the bubble jet (also known as thermal) or piezoelectric techniques.

Sources: Epson America Inc.; Hewlett Packard

BUBBLE JET PRINTING

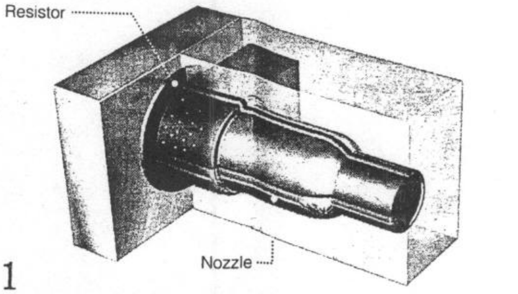

1

To push the ink through the nozzle, bubble jet printers heat minuscule quantities of ink by passing an electrical charge through a resistor, which quickly reaches 900 degrees Fahrenheit.

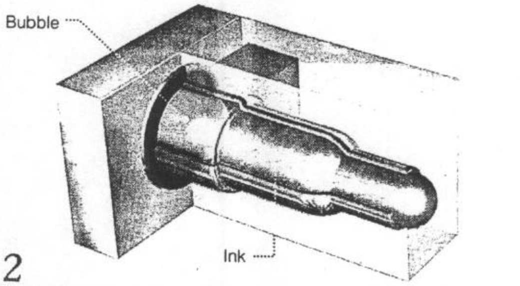

2

The heating element vaporizes a tiny layer of ink at the bottom of the chamber for a few millionths of a second, *forming a bubble, pushing the ink down the nozzle.*

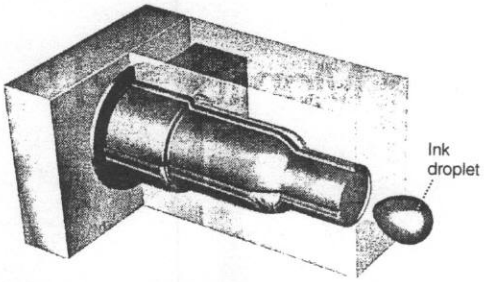

3

The bubble expands and forces a droplet of ink out of the nozzle. Colored droplets are generally so small that a quart could contain 100 billion such drops or more; black droplets are about four times as big. The whole process takes about 10 millionths of a second. The heating element cools and the bubble collapses, creating suction that draws more ink into the chamber from the ink cartridge.

PIEZOELECTRIC PR

1

A piezoelectric print
ger and plunger, it u
when an electrical c
flects the crystal awa

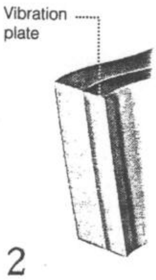

2

A positive charge b
pushes a plate into
droplet of ink.

3

An advantage to this
let can be precisely
flection, enough to d
ink through the nozzl
The ink is forced out

WHAT'S NEXT

As Gadgets Shrink, They May Wind Up in Surr

ıkjet Color Image, Born in a Rainbow of Tiny Drops

ſ LAKE

ɩuter industry gave desktop
ɩice of options when it came
pretty much a Henry Ford
nt in any color as long as it

ɩple used typewriter-style
king dot-matrix printers to
nted a silent printer, your
t machine — and back then,
an $1,000.
dware, inkjet printers have
ɩdels start at less than $200
low $100. Some computer
printers away as bonuses
mputers.
lown, the print quality has
rs can produce color pages
shable from photographs.
ct, that inkjet printers are
reasons the United States
ıg its paper currency in the
ecoming too easy to print passable counterfeit bills with a good inkjet machine.

It is strange that a technology that produces such sophisticated results began with an accident involving a soldering iron. The story goes that in 1977, an engineer at Canon put a hot soldering iron a little too close to an ink-filled syringe. The heat boiled a tiny amount of the ink in the needle, making it expand into a gas. That pushed some ink out of the tip of the needle.

The physics behind this accident became the foundation of what Canon called bubble jet technology, a printing technique in which electrical resistors heat up tiny bubbles of ink, pushing ink droplets through an array of hair-thin nozzles.

Other inkjet printer makers, including Hewlett-Packard, Lexmark and Xerox, have adapted Canon's thermal technology for their own printer lines, adding enhancements (and patents) along the way.

Another printer manufacturer, Epson America, developed its own variation on the theme. Epson printers use piezoelectric crystals to do the job of the bubbles in inkjet printers. Piezo elements change shape when subjected to an electric charge; they are used in loudspeakers to convert signals from a stereo tuner into sound waves. In Epson's line of inkjet printers, microscopic piezo elements subjected to tiny charges are distorted, pushing out tiny droplets of ink.

Of course, there is more to inkjet printing than the firings of microscopic squirt guns. It is a complicated dance of motors that move a print head with 300 to 600 nozzles across a sheet of paper, squirting 5,000 or more droplets of ink per second in a precise minute pattern.

Most color inkjet printers use four colors of ink: cyan (a pale blue), magenta, yellow and black. They are combined to create all the colors the eye sees in a finished printout. To orchestrate this technique, the computer that is attached to the printer uses driver software that determines the exact color for each dot that will appear on the page. It then figures out how many droplets of each ink color must fall on each dot on the page to make that color. Even when printing at 300 dots per inch (a fairly modest resolution), a letter-size page is more than 2,000 dots wide.

Since photographic prints contain the most complicated shades of color, it is a huge mathematical task for the software to figure out the combination of ink colors needed for each dot. And rendering that subtle color with only four inks can be slow. That is one reason manufacturers are producing desktop printers specifically for photographic printing. These printers use six ink colors, adding paler shades of cyan and magenta to help create realistic halftones.

Despite all that precision, inkjet printers must contend with one great unknown that can radically affect the look of a printout: the paper it is printed on.

With the quick-drying inks that Epson, Hewlett-Packard and others use and the tiny ink drops that current printers squirt out, it is possible to get decent text and spot colors even on cheap photocopier paper. But big blocks of color, as in photos, can soak and buckle thin paper stock. The higher contrast and lower absorbency of brilliant white inkjet paper handles color better. And naturally, the best results of all come from glossy photographic paper (which is what you would hope for after paying the premium price of about a dollar a sheet for it).

You can ensure good results by tinkering with the software settings. If inkjet software can calculate how much yellow and magenta and cyan ink to spit onto a dot the thickness of a human hair, you can be sure that it can adjust that calculation for more or less absorbent paper. Most printer driver software provides a setting for different paper types, so when you put inkjet or photo paper into the tray, make sure that you adjust the software to compensate for it. Just don't do it to print facsimiles of currency.

ring Line

ı print head containing 300 to
ers, each filled with either
r yellow ink.

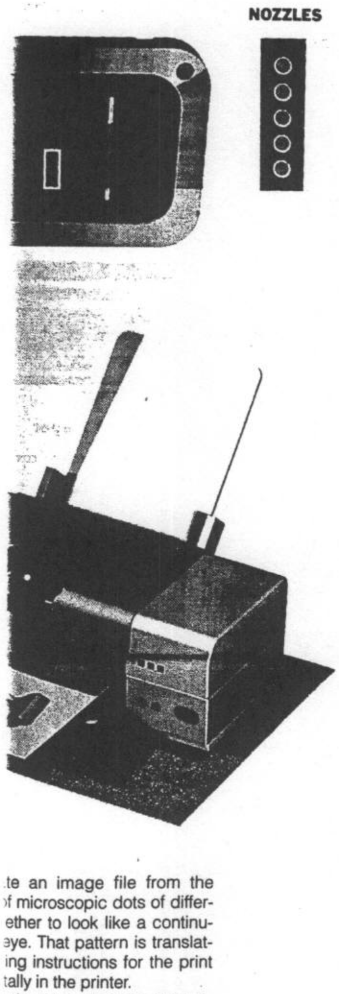

te an image file from the
ıf microscopic dots of differ-
ether to look like a continu-
eye. That pattern is translat-
ing instructions for the print
tally in the printer.
the ink on the page by us-
nechanisms: the bubble jet
ır piezoelectric techniques.

ett Packard

BUBBLE JET PRINTING

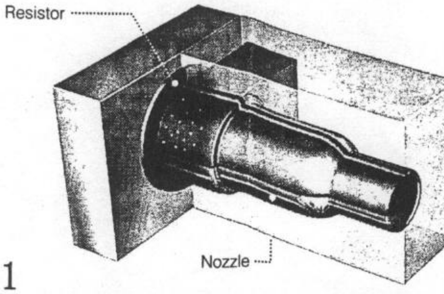

1

To push the ink through the nozzle, bubble jet printers heat minuscule quantities of ink by passing an electrical charge through a resistor, which quickly reaches 900 degrees Fahrenheit.

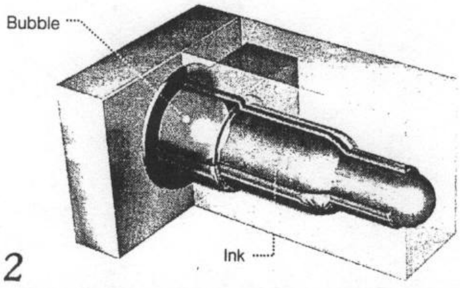

2

The heating element vaporizes a tiny layer of ink at the bottom of the chamber for a few millionths of a second, forming a bubble, pushing the ink down the nozzle.

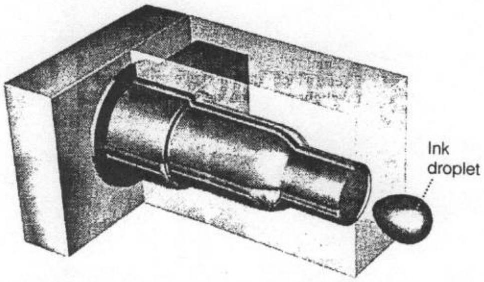

3

The bubble expands and forces a droplet of ink out of the nozzle. Colored droplets are generally so small that a quart could contain 100 billion such drops or more; black droplets are about four times as big. The whole process takes about 10 millionths of a second. The heating element cools and the bubble collapses, creating suction that draws more ink into the chamber from the ink cartridge.

PIEZOELECTRIC PRINTING

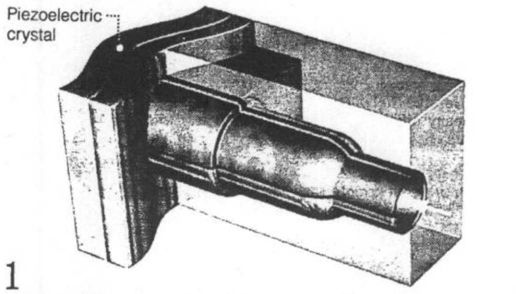

1

A piezoelectric printer works like a squirt gun, but instead of a trigger and plunger, it uses a piezoelectric crystal that changes shape when an electrical charge is applied. A small negative charge deflects the crystal away from the chamber, creating suction.

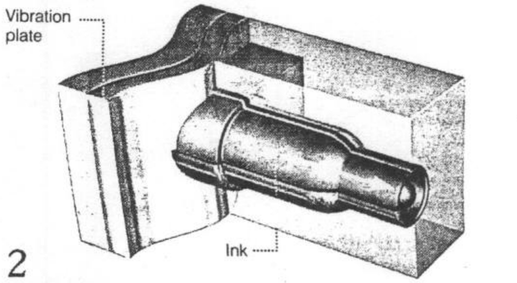

2

A positive charge bends the crystal in the other direction, which pushes a plate into the chamber to create the pressure to expel a droplet of ink.

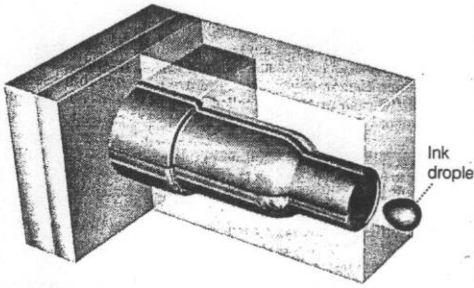

3

An advantage to this approach is that the quantity of ink in the droplet can be precisely controlled. A small charge causes a slight deflection, enough to discharge as little as three-trillionths of a liter of ink through the nozzle. Larger charges can produce larger droplets. The ink is forced out through the nozzle

Illustration by Frank O'Connell/The New York Times

WHAT'S NEXT

Four Lasers, No Waiting

BY M. DAVID STONE

When it comes to color laser printers, the one drawback of true lasers—as compared with LED or solid ink printers—has always been speed.

With an LED printer, for example, it's relatively easy to mount four LED arrays next to one another and print all four colors in one pass. But a laser printer, which has only one laser, needed four passes to print in color and therefore took four times as long to print a color page.

The new **Xerox Phaser 7700DN** changes all that. Inside the chassis are four lasers, each with its own drum, for transferring the cyan, magenta, yellow, and black toner needed for a four-color print (see the diagram). This means the 7700DN can print color output in a single pass at a claimed 22 pages per minute (ppm) for color or monochrome.

Let's get to the big news first: The 7700DN lives up to its claim. We created two 100-page text files with identical text; we formatted one as black text only, then formatted the second with red, green, blue, cyan, yellow, and magenta text on each page. The times for both files were essentially identical, with a 20-ppm tested speed.

The 7700DN printed color photos quickly, taking only about 30 seconds for an 8- by 10-inch photo stored at 200 pixels per inch (ppi) and about 1 minute for a photo stored at 300 ppi.

Just as important, output quality is excellent. Color laser printers generally produce better-looking output than LED printers, and the 7700DN maintains that distinction. Text and lines in graphics are crisp and well formed, color text in graphics stands out well against a color background, and photos printed on standard copier paper are as close to photo quality as we've ever seen from any laser-class printer.

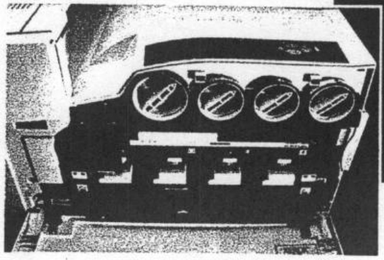

The printer also scores well on setup and network administration. Setup consists of removing the packing materials, putting four toner cartridges in the printer, inserting paper, and connecting a cable and power cord.

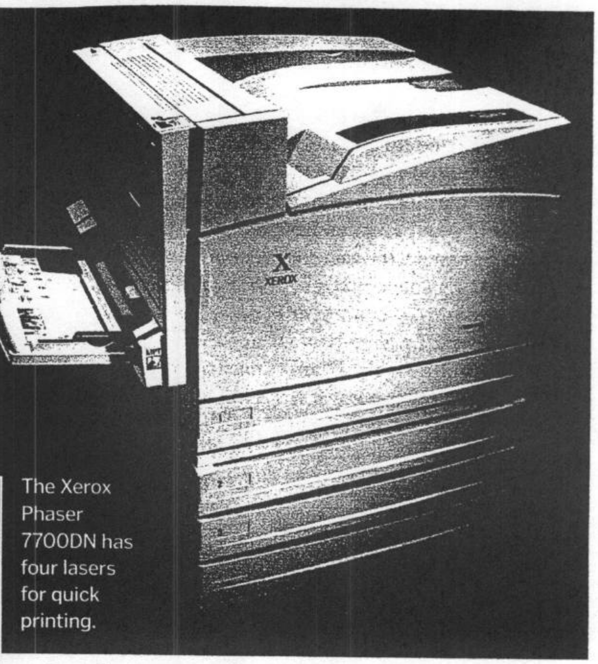

The Xerox Phaser 7700DN has four lasers for quick printing.

Installation on a network is even easier: Insert the CD to start the setup routine and let it find the printer. The installation routine automatically shares the printer and, on our Windows 2000 server, even installed the Windows 95/98 driver to make the system available for easy installation on client workstations.

As you might guess from the speed, the 7700DN is designed for heavy-duty printing. It can print on paper as large as 12 by 18 inches and comes standard with a 650-page input capacity. The 7700DN is the base version and includes a duplexer and network card. Other versions add paper-handling features and memory and range in price from $8,300 to $9,000.

All of this adds up to a printer that can truly serve as an office workhorse and can, by the way, print color output at eyebrow-raising speeds.

Xerox Phaser 7700DN

Street price: With network card, built-in duplexer, 5GB hard drive, $7,000. Xerox Corp., Wilsonville, OR; 877-362-6567, 503-682-7377; www.xerox.com/officeprinting. ●●●●●

Four Colors in a Single Pass

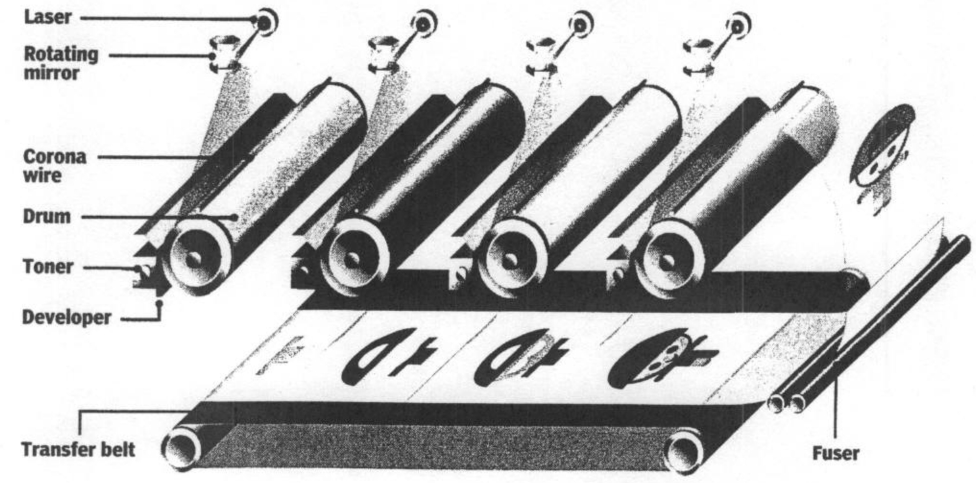

Until now, a color laser printer was able to have only a single laser, which meant the paper had to make four passes around the drum to pick up the cyan, magenta, yellow, and black toner that produces a color printout. Printing a color page took four times as long as black-only.

The groundbreaking Xerox Phaser 7700DN uses a new single-pass engine, made up of four individual lasers and drums, so all four colors can be laid down in a single pass. On our tests, it was just as quick in producing a color text document as it was in delivering an all-black text file (20 pages per minute).

More than Meets the Eye

By Alfred Poor

Understand the technology behind flat-panel displays *before* you buy.

Every time you look at a cathode ray tube in a computer monitor or a television, you're looking at technology that is over 100 years old. Though CRT manufacturers have made some amazing improvements, thin is in. People really want the light weight, small footprint, and sharp image that come with a flat-panel display.

Some types of flat-panel displays, however, are better suited to certain tasks than others. To understand the differences, it helps to know a bit about how flat panels work, how they are manufactured, what flat-panel variations are available now, what technologies are just over the horizon, and how this all can affect your purchasing decisions.

Behind CRTs, LCD technology is the second-most-popular display technology, and it is the mainstay for notebooks and handhelds. It is also playing a central role in data projectors and making inroads into the desktop monitor market. In fact, LCDs are even finding applications in head-mounted and other near-to-eye uses.

LCD technology is older than you might think. The special properties of liquid crystals were discovered by an Austrian botanist, Friedrich Reinitzer, in 1888. It took 85 years for these properties to be commercially exploited. The Japanese company Sharp Electronics released the first LCD product in 1973: an electronic calculator with a digital display.

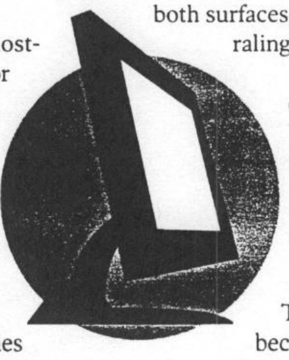

NEMATIC MATERIALS

Liquid crystals get their name from the fact that the molecules are crystals, but they form a liquid instead of a solid. Liquid crystals display two important characteristics: First, if you run an electric current through a layer of liquid crystals, they will orient themselves in relation to the positive and negative poles of the current. Second, without a current, they will try to line up parallel to each other. But if you provide a surface with fine grooves, the layer next to the surface will line up with the grooves.

A third characteristic completes the magic that makes liquid crystals useful: a layer of crystals can bend light waves. A liquid-crystal layer acts as a polarizer (it can filter out all the light waves except those that are oriented in a specific direction). Furthermore, if the crystals in the layer are twisted, the light waves will follow the twist and emerge at the other side of the layer with a different orientation.

Taken together, these three attributes allow liquid crystals to be used as a switch—or *cell*—that can either block or transmit light. The top and bottom surfaces of the cell are made with tiny ridges, which cause the molecules next to these surfaces to line up in parallel. The ridges in the top and bottom surfaces are at an angle to each other. The liquid crystals try to line up with both surfaces and twist slightly in between. The result is a spiraling layer that can twist light passing through it.

If an electrical current is run through the liquid-crystal layer, all the molecules will line up to match the flow of the current, which eliminates the twist of the light. As a result, the presence or absence of the electrical current can determine whether the light will be transmitted or blocked.

PASSIVE-MATRIX LCDs

The first type of high-information-density LCD to become commercially viable was passive-matrix technology. Passive-matrix gets its name from a simple design that lets it switch the liquid-crystal cells on and off.

The individual liquid-crystal cells are sandwiched between two sets of electrodes. The electrodes on the bottom layer run at right angles to the ones on the top layer. As a result, activating one row of electrodes and one column of electrodes will result in a current running through one specific cell.

Passive matrixes create an image by activating each row of electrodes in turn and then, while a given row is selected, selecting a column of electrodes to turn on specific pixels in that row. The design is simple and adds relatively little cost to the panel production. If too great a current is run through a cell, however, adjacent cells may also be affected, resulting in ghosting. Too weak a current and they switch too slowly, which reduces contrast and loses detail in moving images.

EXTREMETECH™

LCD TECHNOLOGY

Adapted from **ExtremeTech.com**

TWISTED-NEMATIC DISPLAYS

Early passive-matrix panels placed the top and bottom layers at right angles, and the crystals had just a 90-degree twist between bottom and top. Contrast was very low, however, and cell response was slow. This approach worked fine for low-information-density displays but proved unsuitable for computer displays. Supertwisted-nematic displays achieve more twist with a different chemical composition. The image is improved with a twist of 180 or 270 degrees.

Although higher twist increases polarization, it also can color the light. This factor gave early laptops their characteristic blue and yellow coloring. A *dual supertwisted-nematic* (DSTN) panel cures this by stacking two LCD layers, twisting in opposite directions. DSTN absorbs more light, however, and is more complex to manufacture. A simpler approach is to add compensating films to reduce the coloring effects. Most of today's passive-matrix color panels rely on *film-compensated STN* (FSTN) technology.

One final improvement in passive-matrix is *dual-scan* design, which divides the panel in half horizontally. The top and bottom halves are scanned simultaneously, allowing any particular LCD cell to be scanned more frequently, which improves contrast, color quality, and response time. Dual-scan panels are still widely used in mid- and low-price notebook computers.

ACTIVE-MATRIX LCDs

Like passive-matrix panels, active-matrix LCDs have transparent electrodes made of indium tin oxide running in rows and columns inside the top and bottom surfaces. The difference is that the current is switched on and off at each cell, thanks to a tiny transistor. The transistors are fabricated right on the panel substrate using thin films. This gives rise to the other common name for these displays: *thin-film transistor* (TFT) LCDs.

The transistors can switch the cells on and off more rapidly than the passive-matrix scanning scheme—and with less electrical crosstalk. You can use higher currents, which will avoid ghosting or streaking, providing higher contrast and sharper/brighter images in both color and monochrome.

Unlike LCDs, CRTs and televisions form their images with phosphors inside a glass picture tube. The light is emitted more or less in all directions from the surface of the phosphor layer. As a result, you can view the image from a wide range of angles with negligible loss in brightness and clarity.

LCDs, however, use a bright backlight to shine through the liquid-crystal layer (and polarizing films and other layers) to create the image. The light that emerges is very directional. Imagine looking through a bundle of drinking straws: The light is blocked except when you're perfectly aligned with their openings.

Manufacturers address this problem mechanically by putting a diffusing film on the top layer. The same problem can also be addressed electrically by running the polarizing current from side to side rather than top to bottom, giving a greater horizontal viewing angle, though the vertical angle is still limited.

As an added cure—at a higher cost—manufacturers can divide each liquid-crystal cell into multiple regions (*multidomain* technology) and pretilt the substrate in different directions so that the molecules line up in different ways. The transmitted light follows these varying directions for a wider viewing angle.

CRTs can produce an almost infinite range of colors, because they are analog devices. Vary the strength of the red, green, or blue analog signal and you vary the output brightness for that color. LCD cells also vary in brightness according to the voltage applied, but only active-matrix panels give you the ability to control it on a pixel-by-pixel basis. Each subpixel of an LCD panel is used to produce red, green, or blue light thanks to minute colored filters. The subpixels block or transmit varying amounts of light to the filters. A typical panel uses an 8-bit controller, which can produce 256 shades. With each subpixel capable of 256 shades,

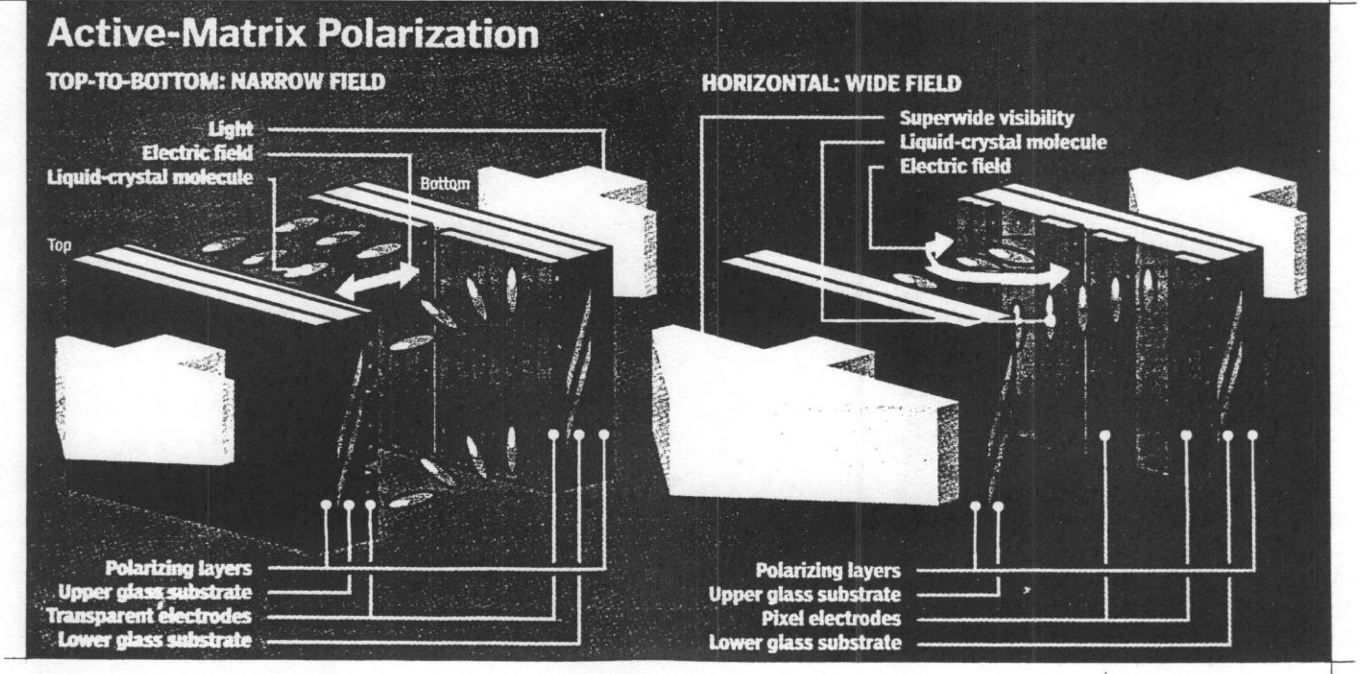
Active-Matrix Polarization

with less elec-
hich will avoid
st and sharper/
.
ir images with
is emitted more
phosphor layer.
range of angles

ne through the
r layers) to cre-
tional. Imagine
light is blocked
penings.
cally by putting
lem can also be
g current from
reater horizon-
ill limited.
cturers can di-
s (*multidomain*
it directions so
he transmitted
viewing angle.
colors, because
e red, green, or
htness for that
g to the voltage
e ability to con-
an LCD panel is
minute colored
mounts of light
ller, which can
of 256 shades,

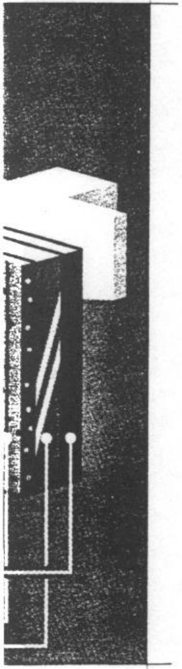

you get 2,563 combinations, or 16,777,216 colors per pixel. Even this 24-bit color depth may not be ideal, because the human eye is nonlinear in its response. LCD engineers adjust the voltages applied at the different intervals so colors appear more uniform.

Engineers also use *frame rate control* (FRC), or temporal dithering, for a greater color scale. This approach refreshes the pixels multiple times per screen image. Instead of using space to display the mix of shades as in color dithering, this method uses time to mix the shades. If the time spent displaying an image were divided into frames, a pixel could be switched between a darker and lighter shade during the frames to produce an intermediate shade. Four frames can produce three intermediate shades. This approach does not reduce the image resolution and is widely used in today's active-matrix panels.

LCDs have a distinct disadvantage compared with CRTs when displaying fast-moving images such as scrolling data or movies. The time it takes a CRT picture tube's gun to emit an electron stream and the excited phosphors to emit light is almost instantaneous. LCDs, however, depend on an electrochemical reaction, which has inherent latency.

The lag time is formally called *response time* and is typically measured in milliseconds. Passive-matrix panels have very slow response times—as much as 150 ms or more—so they are not suitable for moving images such as movies. Standard active-matrix panels typically have a 40-ms response time, which means they are capable of displaying 25 frames per second (fps). In-plane switching increases the viewing angle but results in a slower display; a response time of about 70 ms is typical. Multidomain panels are often faster, with a 25-ms response time.

POWER CONSUMPTION

Active-matrix LCD panels consume little power compared with CRTs. Active-matrix has become the standard type of display for portable applications, but LCD technology is woefully inefficient. For example, if you display a white screen, less than 10 percent of the light shining into the back of the panel from the backlight is transmitted through the front. The rest of the light is absorbed.

The inefficiency forces notebook computer designers to face some tough choices. Brighter backlights require more power, so premium notebooks need larger batteries to sustain longer running times. The backlight represents the largest portion of an LCD's energy budget, which has remained fairly constant at 2 to 5 watts. It consumes about 1.2 watts per tube and so accounts for 1.2 to 2.4 watts of the energy budget, depending on whether the panel uses a single- or dual-tube backlight.

PDAs typically use *transflective* surfaces, either backlit or illuminated by ambient light. This is not particularly efficient, but it drastically cuts the required energy and device weight.

One of the key concerns about LCD panels is their cost. If they were as inexpensive as CRTs, they would likely take over almost all of the display market. Unfortunately, 15- to 17-inch-diagonal LCDs—the sizes used most often for desktop displays—cost roughly three to five times as much as CRTs with a similar viewing area. For larger displays, the cost difference is even greater.

Why do LCDs cost so much? The manufacturing process is complex and unforgiving, and maintaining high yields is a constant struggle.

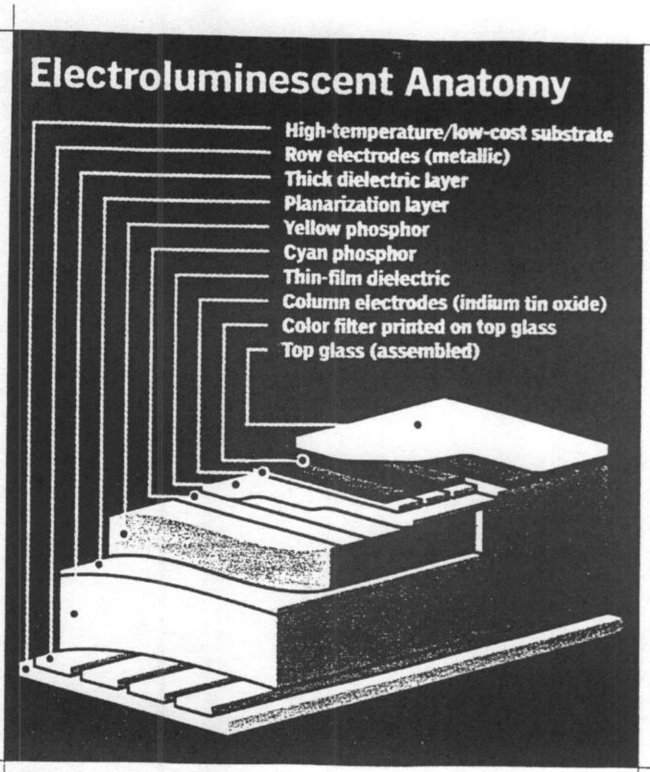
Electroluminescent Anatomy

TRADITIONAL PROCESS

LCD panel production starts with very thin sheets of glass measuring 0.4 to 1.1 mm thick. The bottom sheet is coated with a thin layer of amorphous silicon, which makes it possible to create the semiconductor device for each pixel location. Through a series of photolithography, etch, mask, and deposition steps, a switching transistor is created at each pixel, along with color filters and other components.

A transparent alignment layer is deposited over all the components, and a similar transparent alignment layer is applied to the inside of the top sheet. These layers are then mechanically brushed or etched using a photochemical process that creates tiny grooves in each layer. When the liquid-crystal material is added, the liquid-crystal molecules will align with the grooves.

Tiny spacers are sprayed onto the bottom surface, so that one or two spacers are in each pixel space. These spacers provide a physical barrier to separate the top and bottom glass layers, to create a space for the liquid-crystal material.

Next, a sealant is applied around the edges of each panel on the bottom glass, leaving a small gap at one edge. The top sheet is then fused to the bottom, and the fused sheets are cut into individual panels.

The liquid crystal material is then injected into the small gap using a process that first evacuates all air in the empty space between the sheets and then uses nitrogen gas to force the liquid-crystal material into the panels. The gap is then closed, and the panel is ready for testing to make sure that all the subpixels are working properly. After testing, polarizing films and other layers are added. The final step in creating the display module is adding the electronic driver circuits and interface connectors to the computer or other device.

YIELD ISSUES: DEFECT SOURCES AND SOLUTIONS

One of the key cost factors is yield. It doesn't matter whether you can produce many panels inexpensively if you have to throw them all out at the end of the production line because they don't pass quality control standards. But there are many ways to create substandard panels.

Defects can be optical, electrical, or mechanical. *Optical* defects include pixels that are stuck on or off. Pixels that are stuck on are more noticeable than ones that are always off. In some cases, a bright defect is "repaired" by zapping that cell's circuits, changing it into a dark defect—one that is always off. Other optical defects include artifacts from contamination, inconsistent cell thickness, and a lack of uniformity in the color filter layers.

Electrical problems include flaws in the driver circuitry and other electrical components in the panel. *Mechanical* defects include cracked glass substrates, misalignment of the electrical connectors, and misalignment of the panel in its case.

Most of these defects can be identified only after the liquid-crystal material has been injected and the panel has been sealed. The vast majority of time and materials have already been invested in the unit.

COMPETING FLAT-PANEL TECHNOLOGIES

Like CRTs, which have a commanding lead in the desktop monitor and television markets, LCDs dominate the portable-display market. Several competing technologies hope to gain a share of these markets at the expense of CRT and LCD products.

Most displays are designed to be viewed directly and without magnification. Here is an overview of some of the new flat-panel technologies that hope to make inroads in these markets.

- **Plasma.** A plasma display runs an electrical current through a gas to create a charged plasma, which radiates ultraviolet light toward the back of a panel that is coated with phosphor material. The phosphors, which are arranged in alternating red, green, and blue columns, emit light that is transmitted through the front layer to the viewer. Plasma panels have response times fast enough for TV and movie viewing. They can be large—50 inches diagonally—and yet are relatively thin and lightweight. The range of viewing angles is similar to that of a CRT.

Plasma panels, however, have relatively short lives. Some are rated for 10,000 hours, about 13 months of continuous use, before the phosphors degrade to half their brightness. Some newer designs are rated for 20,000 to 30,000 hours, which is comparable to the life of many CRTs. The displays also are not particularly energy-efficient. Plasma displays remain expensive; a 42-inch unit costs around $10,000.

Major companies are working hard to market plasma panels and reduce production cost. Although plasma displays are serious competition for large CRTs, they are also vulnerable to competition from other technologies.

- **Organic Light-Emitting Diodes (OLEDs).** Technology that uses organic light-emitting diodes (OLEDs) is the current darling of the industry. Many companies are now investigating OLED panels, and others are trying to bring such products to market.

OLEDs rely on special carbon-based molecules that emit light of different colors when an electrical charge is passed through them. Some of the materials can be made into polymers—like plastics—that may be applied by printing. OLED panels require relatively small amounts of power at low voltages. The panels are extremely thin and light, the fabrication is relatively simple, and the materials have good lifetime specifications. And they are emissive, so they have good viewing angle characteristics.

For computer displays and televisions, however, OLEDs require active-matrix switching, causing the same yield problems as active-matrix LCDs.

For the short term, expect OLEDs to be limited to low-density displays, as on car stereos and other gauge applications. It may take some time before they reach the market in computer and video applications.

- **Field Emitter Displays (FEDs).** FED technology uses the electron gun (emitter) of a CRT and moves it up so that it is right behind each phosphor in the display. Instead of three guns covering the entire screen, each subpixel gets dozens of "guns." The result is a very thin panel with a CRT-like picture. Tens of millions of emitters are used for the entire screen. When turned on, the electrons are directed at the phosphors, which emit light.

The biggest advantage of an FED design is that it is emissive and thus has an excellent range of viewing angles while remaining as thin as an LCD. The response times are sufficient for video and movies and significantly faster than with active-matrix LCDs. The panel structure is relatively simple, and the large number of emitters provides redundancy, making higher yields likely. Even if 10 emitters out of 100 fail to work, the loss of brightness for the affected pixel or pixels would likely be unnoticed.

Power consumption, however, makes FEDs less attractive for portable applications. They are also more difficult to produce than manufacturers initially expected. And they are limited in size; the largest demonstration unit shown so far is 15 inches diagonally.

A number of big-name companies have experimented with FEDs and quietly dropped their development programs. Several small companies are still working on the technology, but it remains to be seen whether they can obtain the funding needed to bring their products to market.

- **Electroluminescent (EL).** Electroluminescent materials, like OLEDs, emit light when a charge is passed through them. In the past, they have been used as backlights and low-information-density displays. But in recent years, some companies have tried to develop high-information-density products for use in entertainment and computing. These structures tend to be very simple, using layers of materials that are quite thick by LCD and semiconductor standards.

The thick layers have important advantages. Contamination has little visible impact on the image quality, so the displays cost much less to produce than technologies that require clean rooms. EL materials have fast response times and a wide viewing angle.

The coarse construction makes it difficult, however, to create fine-resolution devices. Therefore, the technology is being applied mostly to segmented displays and television screens.

The most likely application for EL displays in the near future will be limited to televisions. If manufacturing-cost advantages can be realized, this technology will compete well against large LCDs and plasma displays in the thin-TV market.

s. The panels
atively simple,
ons. And they
haracteristics.
ver, OLEDs re-
ield problems

to low-density
ations. It may
computer and

logy uses the
that it is right
ree guns cov-
of "guns." The
. Tens of mil-
en turned on,
ich emit light.
it is emissive
while remain-
cient for video
-matrix LCDs.
rge number of
ds likely. Even
htness for the

attractive for
lt to produce
are limited in
r is 15 inches

imented with
grams. Sever-
logy, but it re-
nding needed

materials, like
h them. In the
-information-
ies have tried
use in enter-
be very sim-
k by LCD and

ontamination
displays cost
clean rooms.
iewing angle.
ever, to create
y is being ap-
screens.
e near future
st advantages
against large

MICRODISPLAYS MAKE THEIR MARK

Microdisplays, at 2 inches or less diagonally, open a whole new world of markets. One of the most familiar applications is the data projector. Up to ten years ago, the only way to project data on a screen was with a large, 80-pound projector that used CRT technology. Such projectors were not very bright, and they required tedious adjustments to produce an acceptable image. Today you could buy a projector with a much higher resolution and a much brighter image for less than half the price. It would weigh about 3 pounds and fit into a coat pocket.

Another potential market for microdisplays is desktop monitors. Imagine shrinking a home projection TV down to about 20 inches diagonally. You'd have a monitor about the same thickness as the base of a typical LCD monitor, at a fraction of the weight of a CRT and with the same screen size.

The third type of application is often referred to as *near-to-eye displays*, which are typically head-mounted for privacy or hands-free viewing. They are already in use as electronic viewfinders for digital cameras. This kind of display can also be used with a PDA or cell phone to create a virtual image of a large, high-density screen—a near necessity if wireless Web browsing is to gain acceptance. Near-to-eye displays may also enable entirely new classes of entertainment and productivity devices.

As described earlier, high-temperature polysilicon is already in wide use in projector applications. The display panels feature high aperture ratios—very little of the light is blocked by the transistors in the individual cells—and the fabrication process is relatively mature. Although this technology is expected to continue as a mainstay in projectors weighing 10 pounds or more, it faces some stiff competition from newer designs.

ELECTROMECHANICAL PANELS

One of the most successful new designs is the electromechanical panel. In this design, the electrical signals are transformed into physical movement of mechanical devices built into the display panel. This may sound impractical or old-fashioned, but it works remarkably well.

The Texas Instruments Digital Light Processor (DLP) is a small chip covered with tiny metal mirrors. When a signal is sent to turn on a pixel, a mirror rotates to reflect light through the lens toward the screen. When the *off* signal is sent, the mirror tilts the other way and reflects light away. In this design, mask, etch, and strip cycles create microscopic mechanical structures. A projector may use one DLP and a color wheel (field-sequential, or one color at a time) or three DLPs and filters to make a full-color image.

Electromechanical panels have rapid response times, which makes field-sequential displays possible. One panel can switch to the red, blue, and green (and sometimes white) components of an image frame as it is illuminated by a rapidly spinning wheel in the time allotted for a single frame. The separate fields are combined in the human brain, where they are perceived as a single, full-color image.

For higher quality, three DLP panels can be used. Each is illuminated with colored light—red, green, and blue—and the images are then combined optically to create a single, full-color image. This approach is used in certain large digital movie projectors.

DLP panels have a high resolution and have proved extremely reliable. They also have roughly twice the contrast of LCD projectors, making them more effective in bright rooms.

DLP panels have a few drawbacks. They are more expensive than LCDs. Field-sequential images can appear to break up into different colors when you look from one spot on the screen to another, especially if the images have small white objects on a black background. In projector applications, the motor used to spin the color wheel can be noisy, but new, solid-state color filter systems coming to market eliminate this problem.

The prospects for DLP applications are excellent, as small

Potential Manufacturing Breakthrough

Alien Technology's fluidic assembly could improve yields.

The chances are slim that the traditional LCD manufacturing process can be improved enough to cut production costs to one-half or one-quarter of their current levels. But this is what it would take for the price of LCDs to come close to the cost of CRT monitors.

Researchers have been looking for new approaches that could have a major impact on panel prices. One company that has the potential to revolutionize LCD production is Alien Technology. The company has developed a *fluidic self-assembly* (FSA) process that could cause a dramatic drop in display prices.

The central concept is that an LCD panel can be seen as a large area that is mostly empty. At precisely spaced locations, the panel needs switches in the form of transistors. The traditional process prepares the entire glass substrate surface for creating semiconductors by coating the glass with amorphous silicon but uses only small regions of this coating for the cells' transistors. Amorphous silicon is a less-than-optimal substrate for this process, and you can end up with bad transistors that result in pixel defects.

Alien Technology takes the view that you should put semiconductors only in the precise locations where you need them and not waste time and effort preparing the rest of the surface. In addition, the company wants to use high-quality semiconductors that are known to be good before they are used in a display. Alien uses several clever techniques to create tiny transistors, separate the good ones from the defectives, and position them on the LCD substrate.

The company has a pilot production plant and is building a full-scale production facility, which is scheduled to be up and running by the end of 2001. If Alien Technology can deliver on the promise offered by its FSA process, it could revolutionize the display industry.—*AP* projectors move into the home entertainment market. There is also a good outlook for the use of DLP panels in desktop projection monitors.

LIQUID CRYSTAL ON SILICON (LCOS)

Another class of microdisplay uses liquid crystals with a reflective backplane, fabricated directly on a silicon chip. Some such displays use liquid-crystal material, and others use exotic alternatives such as the ferroelectric liquid crystal used by microdisplay innovator Displaytech. Some are transmissive; others are reflective. They are suitable for both near-to-eye and projection applications.

An LCOS display is similar to a standard LCD panel except that the backplane is a silicon chip with a reflective coating. The control structures can be tiny, so you can get very high resolutions on tiny panels. Also, driver circuitry can be built on the edge of the chip, reducing component counts and physical connections. A top glass layer holds the liquid crystal material in place. The gap between the layers is typically 1 to 4 microns thick.

Such displays use CMOS fabrication procedures similar to those that chip makers use. The resulting panels are moderately priced and have relatively fast response times. The prices of these panels could be reduced greatly, especially if demand permits production in large volumes.

LCOS displays do not have the high contrast of DLP images. Some are designed for field-sequential illumination—such as those used with single-panel DLP applications—but others either rely on combining the image of three separate panels or use color filters and triple the number of pixels for a given resolution. One source lists more than 20 companies actively developing or producing these panels, so they will play an increasingly important role in mobile communications, computing, and entertainment applications. These products already have started to appear in the viewfinders of digital cameras and camcorders and in projection television systems.

RETINAL SCANNING LASERS

One of the most novel display techniques is the retinal scanning laser. The design is primarily for head-mounted displays, such as those created for military pilots. A single, low-power laser beam is modulated by circuitry and then scanned in rows to create a complete image. The image is projected directly onto the viewer's retina, and field-sequential color filters are used to create full-color images.

The image can be bright, which makes it visible even when you're viewing daylight scenery. It can also be dimmed for nighttime use. Although the display has a number of fail-safe design features, some analysts are concerned that a failure of the control mechanism might leave the laser beam stationary on one spot of the retina long enough to cause irreparable eye damage.

Given the current expense and bulk of this technology, it is unlikely to make the jump from defense research to consumer products in the near future.

VISIBLE PROGRESS

It's all too easy to take the displays in our lives for granted. We look at them for hours a day, yet we often don't see them.

Flat-panel displays in general—and liquid crystal displays in particular—have made enormous strides in the past ten years. The advances in the next ten years will be even greater. The CRT remains king of the desktop displays, but there are developments in the labs and on the market that may result in a whole new way of seeing your computer information, movies, and television entertainment. Keep watching.

See www.extremetech.com *for the full version of this article.*

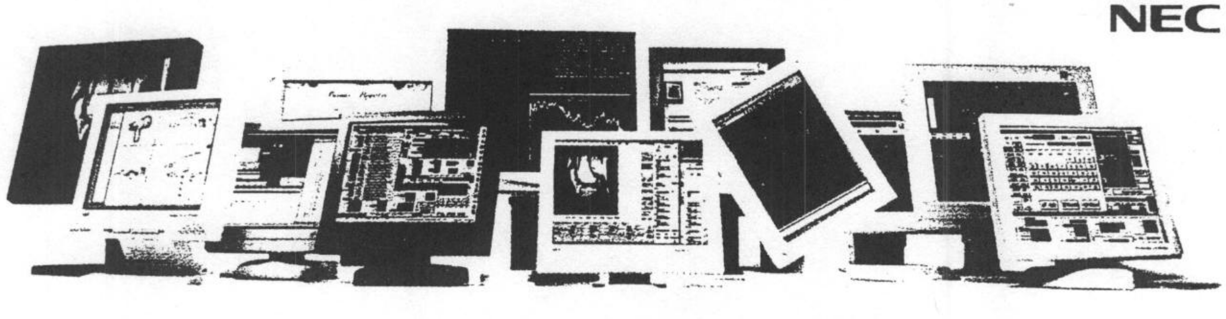

Lifting the Cover of Those Ubiquitous Touch Screens

By MATT LAKE

SOMETIMES a computer mouse just isn't easy enough to use. When bank customers are taking money from automated teller machines or restaurant workers are tallying bills, grabbing an electronic device encased in a handful of plastic doesn't come naturally at all.

That's when people find the original pointing device — the index finger or its surrogate, the stylus — much more efficient and easier to use. That is why touch-screen technology was developed and why it is now used for everything from palm-size computers to controllers in airplane cockpits.

Touch screens combine two separate technologies, one to display items on a screen and the other to figure out what you are pointing at. The displays are either cathode-ray tubes, like those used in televisions, or flat panels, like those in laptops.

Every part of a touch-screen system must work fast. Without a visible change or audible cue, most people touch the screen repeatedly (which can overload the system) or assume that the device is broken.

Touch screens use one of four different technologies to detect touch, and electronic controllers and software drivers to communicate that information to a computer. The touch sensor system must be calibrated so a touch can be correlated exactly to the right spot. Two of the technologies, called resistive overlay and capacitive overlay, can drift out of synchronization and need to be recalibrated periodically. The other two use ultrasound or infrared signals and tend to stay aligned.

The earliest touch-screen technology, which is also the most widespread and least expensive, is the resistive overlay. Resistive overlays were first developed by Elographics, a company in Oak Ridge, Tenn. Elographics (which has since changed its name to Elo TouchSystems) came up with the idea of overlaying two conductive sheets separated by a mesh of tiny raised dots, then passing electricity through one layer. When the screen was touched, it would make contact between the two layers, letting the electricity take a shortcut. That interruption of the flow of electricity is at a measurable point — one that can be determined precisely enough to register the touch of a stylus tip.

The resistive overlay system has a few drawbacks, but it is the system used in Palm and Pocket PC devices. The overlay is susceptible to abrasion, which rules out its use in public kiosks. And the thin conductive layers do block some light from the display screen. At first, up to 35 percent of the light was blocked; that's now down to about 18 percent.

More light makes its way through capacitive overlay touch screens, which has helped that technology catch on, especially for gambling machines, where graphics are important. When the screen is touched, the user draws a small amount of electricity to the point of contact, so that touch can be detected. The capacitive coating is thinner and harder than the glass it is applied to, which makes it more durable for public uses. But the user cannot wear gloves or point at the screen with a stylus or pen.

That gave rise to another type of touch sensor system, called surface acoustic wave technology, that combines durability with sensitivity. Surface wave touch screens do not require any opaque overlay so their displays are bright — and they do not require periodic recalibration. These screens detect the absorption of ultrasonic waves at the point of touch.

Like capacitive screens, acoustic wave screens are robust. That's why they are used in machines dispensing Metrocards for riding the subways and buses in New York City. Many banks use the technology in automated teller machines. The screen can become quite dirty and dusty without giving false readings.

The most recent addition to the touch-screen pantheon, infrared screens, use a "Mission Impossible"-style grid of infrared light and sensors to detect the point of touch. Infrared screens tolerate a range of temperatures and do not require frequent recalibration, which makes them suitable for use in airplane cockpits and A.T.M.'s alike.

A Sandwich of Layers Gives a Screen Just the Right Touch

Touch screens have to perform two separate tasks: they must display a screen full of items to point at, then must detect which ones you're pointing at. The display screens are like television sets or computer monitors, but there are four different technologies for detecting the touch of a finger or stylus.

Resistive overlay screens

Used for point-of-sale terminals in stores, hotels and restaurants, and, as shown here, in hand-held computers like the Visor.

A resistive screen is a sandwich, with two conductive layers that are kept apart, usually by a grid of tiny insulating dots. Pressure from a finger or stylus pushes the two layers together, creating an electrical contact.

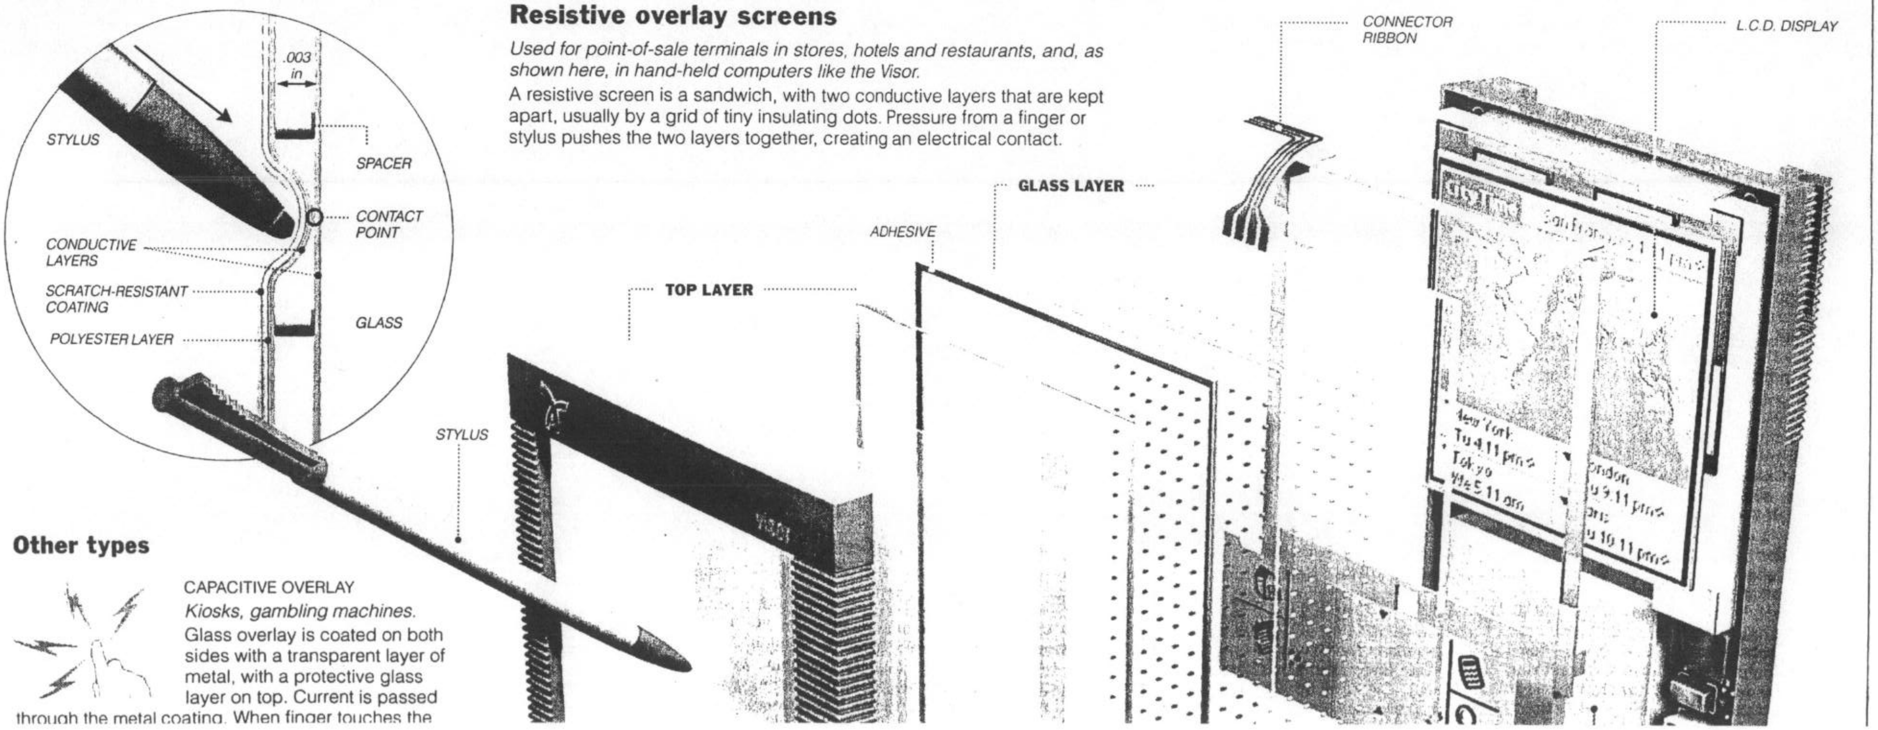

Other types

CAPACITIVE OVERLAY

Kiosks, gambling machines.

Glass overlay is coated on both sides with a transparent layer of metal, with a protective glass layer on top. Current is passed through the metal coating. When finger touches the

A Sandwich of Layers Gives a Screen Just the Right Touch

Touch screens have to perform two separate tasks: they must display a screen full of items to point at, then must detect which ones you're pointing at. The display screens are like television sets or computer monitors, but there are four different technologies for detecting the touch of a finger or stylus.

Resistive overlay screens

Used for point-of-sale terminals in stores, hotels and restaurants, and, as shown here, in hand-held computers like the Visor.

A resistive screen is a sandwich, with two conductive layers that are kept apart, usually by a grid of tiny insulating dots. Pressure from a finger or stylus pushes the two layers together, creating an electrical contact.

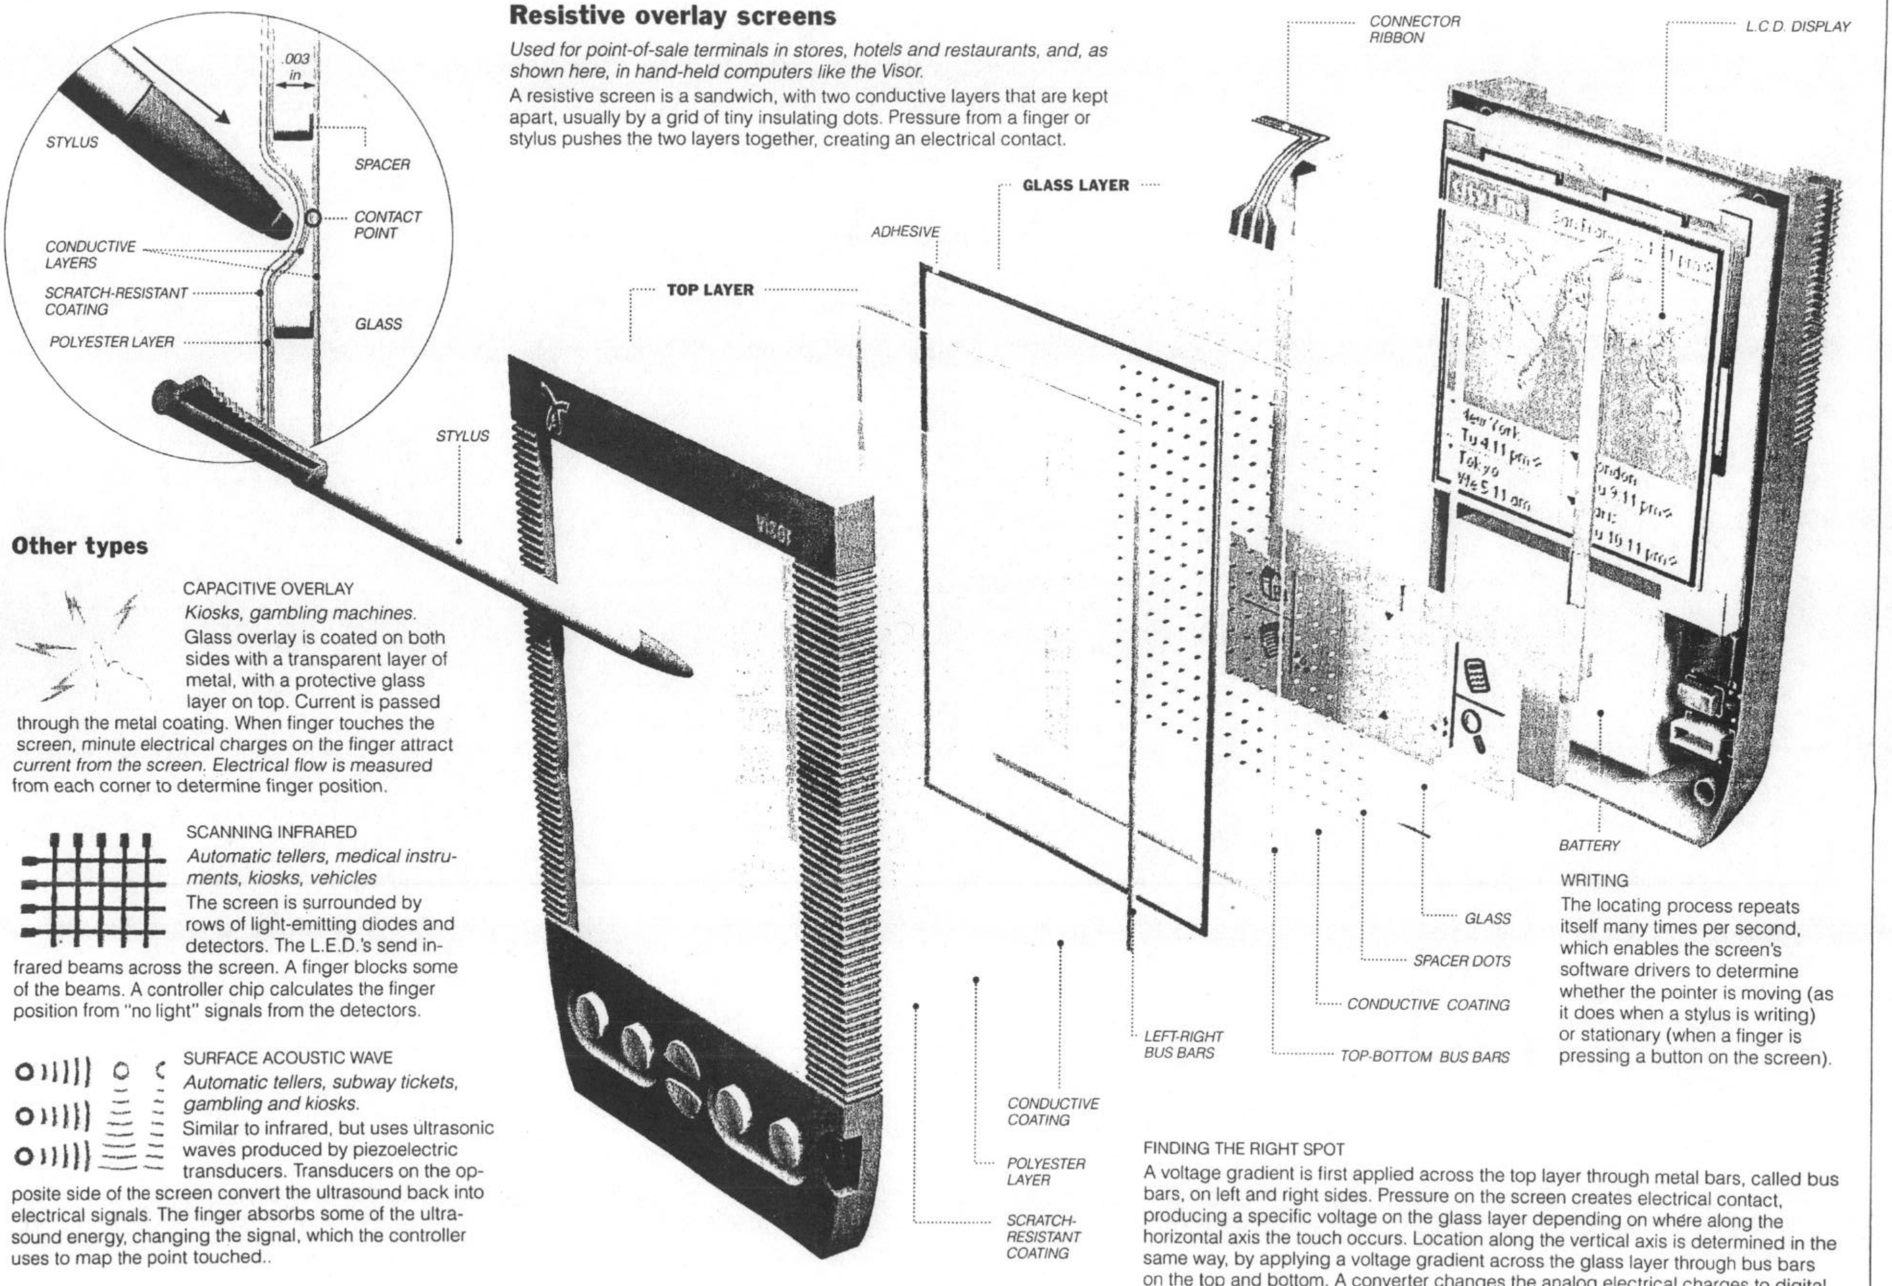

WRITING

The locating process repeats itself many times per second, which enables the screen's software drivers to determine whether the pointer is moving (as it does when a stylus is writing) or stationary (when a finger is pressing a button on the screen).

FINDING THE RIGHT SPOT

A voltage gradient is first applied across the top layer through metal bars, called bus bars, on left and right sides. Pressure on the screen creates electrical contact, producing a specific voltage on the glass layer depending on where along the horizontal axis the touch occurs. Location along the vertical axis is determined in the same way, by applying a voltage gradient across the glass layer through bus bars on the top and bottom. A converter changes the analog electrical charges to digital code, and a controller chip determines the horizontal and vertical coordinates

Other types

CAPACITIVE OVERLAY

Kiosks, gambling machines.

Glass overlay is coated on both sides with a transparent layer of metal, with a protective glass layer on top. Current is passed through the metal coating. When finger touches the screen, minute electrical charges on the finger attract *current from the screen. Electrical flow is measured* from each corner to determine finger position.

SCANNING INFRARED

Automatic tellers, medical instruments, kiosks, vehicles

The screen is surrounded by rows of light-emitting diodes and detectors. The L.E.D.'s send infrared beams across the screen. A finger blocks some of the beams. A controller chip calculates the finger position from "no light" signals from the detectors.

SURFACE ACOUSTIC WAVE

Automatic tellers, subway tickets, gambling and kiosks.

Similar to infrared, but uses ultrasonic waves produced by piezoelectric transducers. Transducers on the opposite side of the screen convert the ultrasound back into electrical signals. The finger absorbs some of the ultrasound energy, changing the signal, which the controller uses to map the point touched..

Sources: Handspring; MicroTouch Systems Inc.

Frank O'Connell/The New York Times

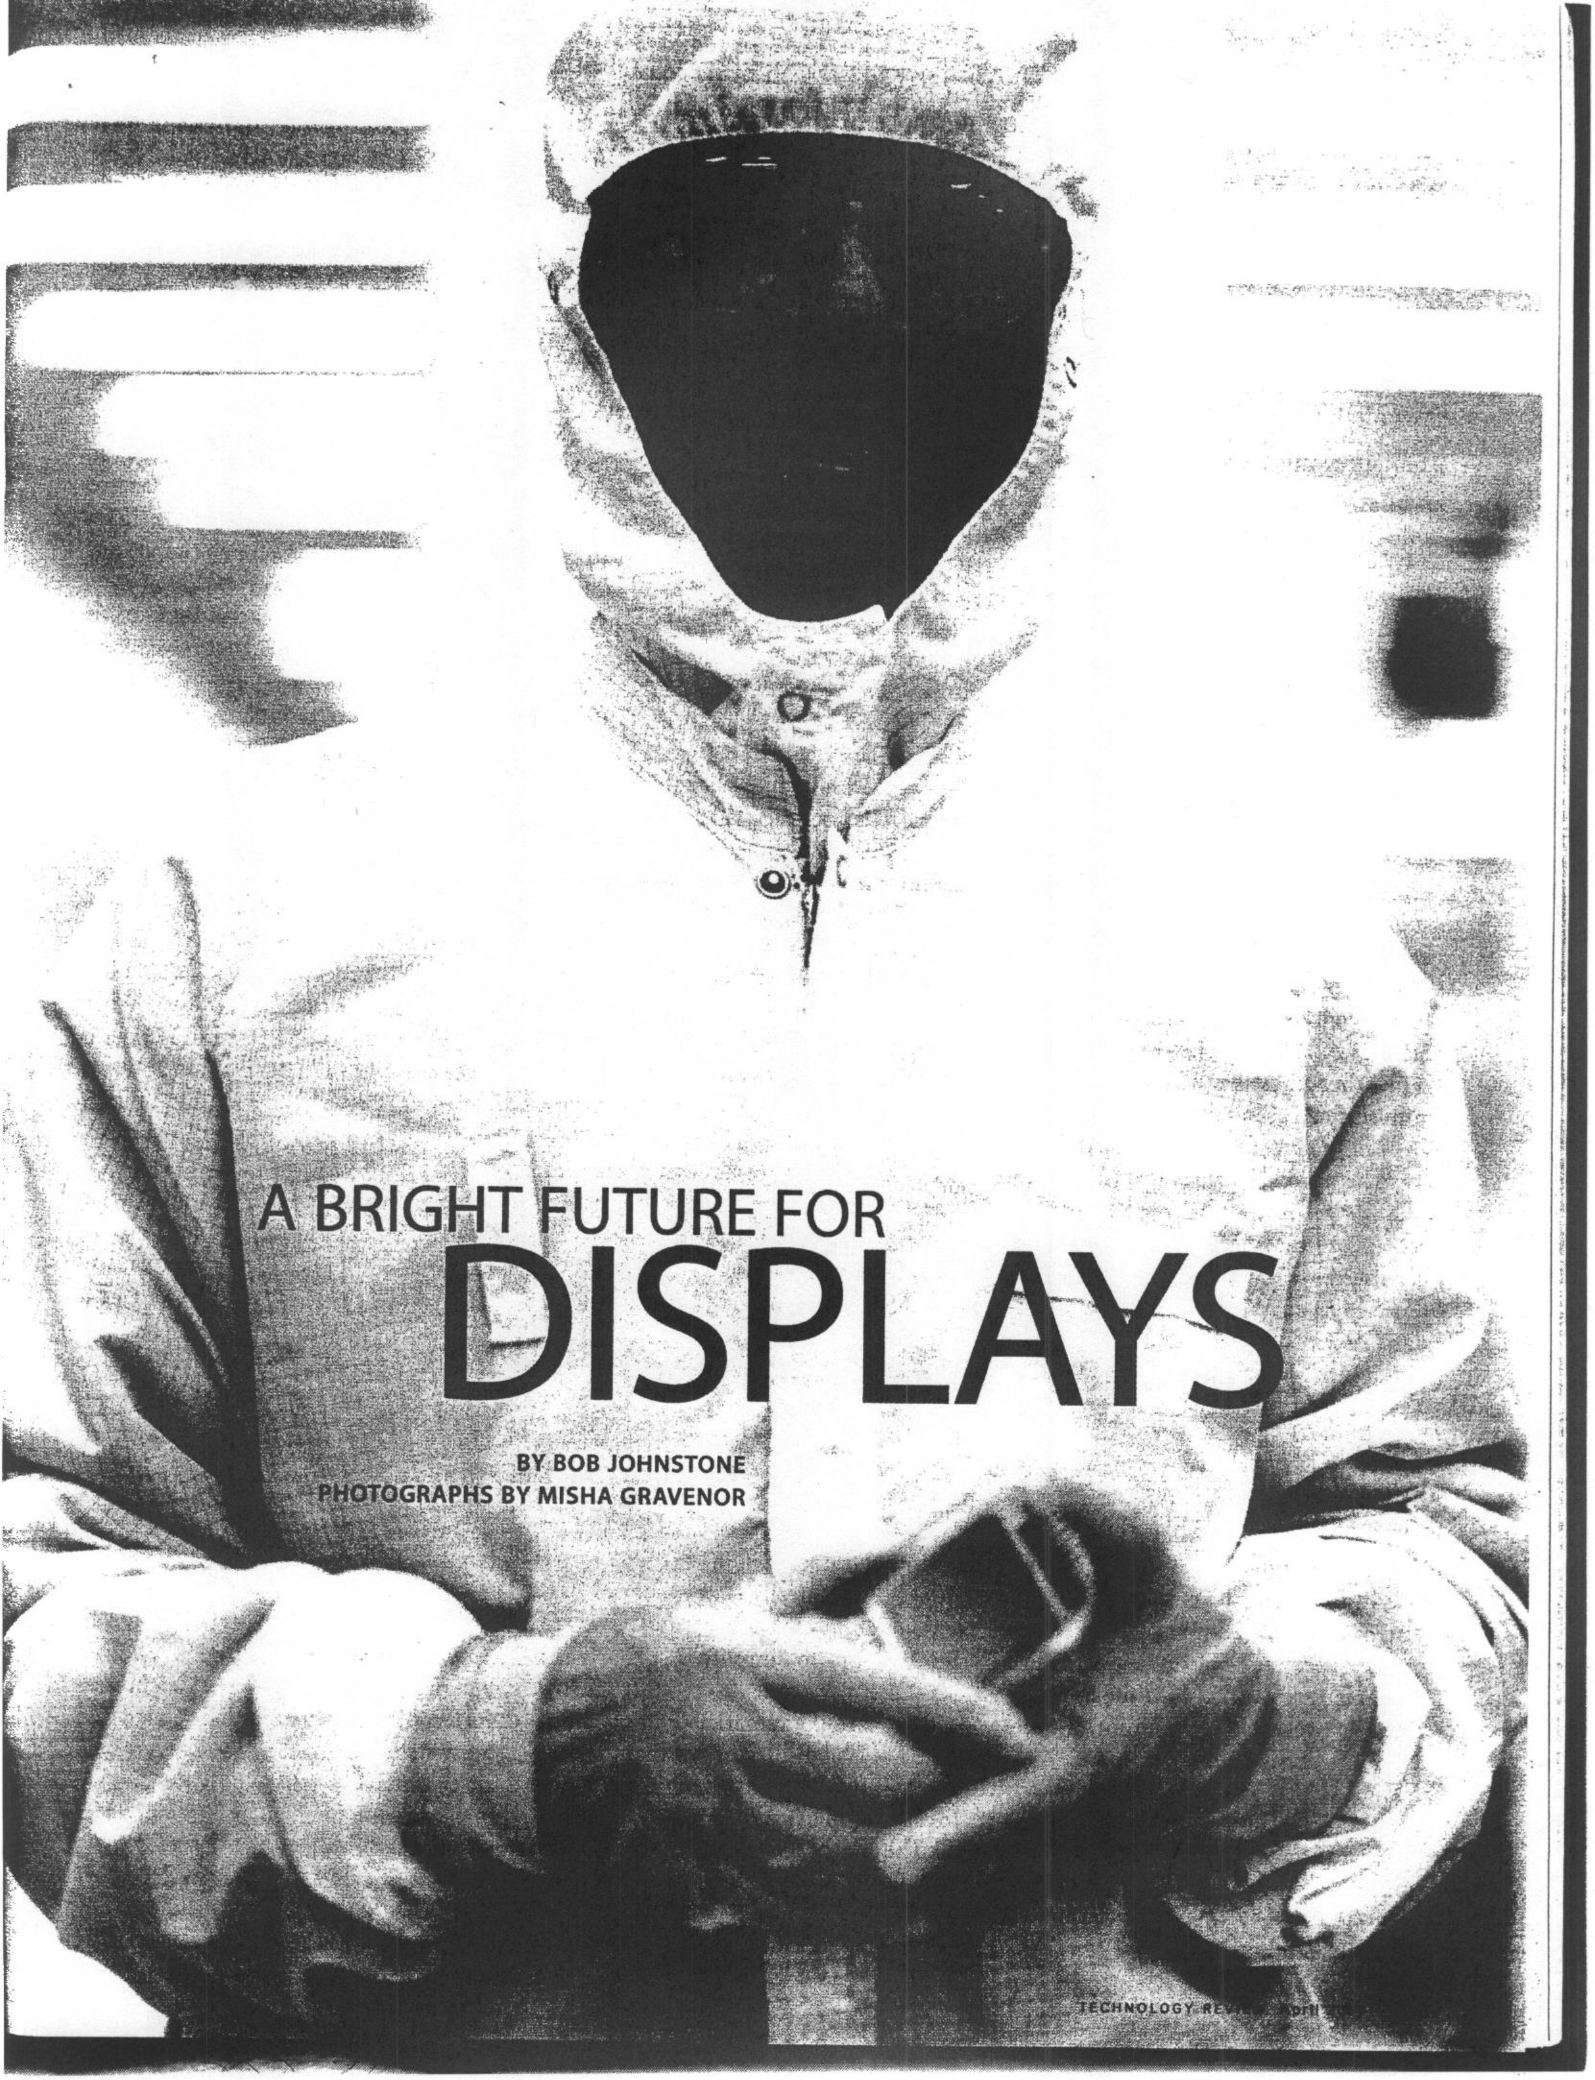

A BRIGHT FUTURE FOR DISPLAYS

BY BOB JOHNSTONE
PHOTOGRAPHS BY MISHA GRAVENOR

Every year, thousands of tourists flock to the California coastal city of Long Beach, the last resting place of the ocean liner Queen Mary—a notable showcase of what was once state-of-the-art technology. But for a select group of visitors last May, the city's main attraction was not a memento of the past but a technology of the future: a dime-thin sheet of glass 14 centimeters along the diagonal whose unparalleled ability to exhibit ultrabright colors and process high-clarity video images holds the potential to have far greater impact on the world than any single ship, no matter how splendid.

Based on a technology called organic light-emitting diodes, the prototype screen was unveiled by Eastman Kodak and Sanyo Electric at the annual conference of the Society of Information Display, the industry's top professional group. As the screen was put through its paces, running images from video cassette, DVD and digital tape, even grizzled veterans of the flat-panel industry who packed into the Kodak booth came away goggle-eyed. Little wonder. Organic light-emitting diodes are shaping up as a superdisplay: brighter, thinner, lighter and faster than liquid crystal displays. They also take less power to run, offer higher contrast, look equally bright from all angles and have the potential to be much cheaper to manufacture than their conventional counterparts.

These advantages, especially the ability to handle video, give the upstart technology the inside track to become the screen of choice for the coming third generation of mobile phones. About to debut in Japan, the third-generation standard seeks to spur the production of phones that are aimed at eyes as well as ears, by giving them the ability to handle high-speed video over the Internet. These wireless Web phones are expected to quickly become a multibillion-dollar global business. But that may be only the start for organic light-emitting diodes, which are threatening to challenge the 30-year hegemony of liquid crystal displays in a broad range of portable electronics.

This promise has fired the imaginations of scientists and engineers and spurred a worldwide race to develop the technology that pits startups against heavyweights such as Kodak and Sanyo. Difficult problems remain to be solved before the promise can be realized. But the potential is too great for some savvy technology companies to ignore. Notes Dalen Keys, chief technology officer of DuPont Displays, the chemical giant's spinoff that is out to win a big share of this emerging market, "We are trying to achieve a complete change of the paradigm of what is a display, and the cost of the display."

A STRANGE BLUE GLOW

If ever a technology has begged to be disrupted, it is liquid crystal displays. Invented in 1963 and originally envisioned as a slimmed-down replacement for bulky cathode-ray tubes or as screens for wall-mounted televisions—a use never realized due to problems scaling up to large surfaces—liquid crystal displays have instead become the standard for everything from watches to laptop computers.

In spite of its spread, however, this ubiquitous technology has an Achilles' heel: the screens are hard to make and therefore expensive—especially when it comes to high-end versions used in color displays. Indeed, they account for as much as a third of the cost of a laptop, and the failure to bring down the price of portables as dramatically as the price of personal computers has been due largely to the inability to simplify display-screen production.

Perhaps the remarkable thing about liquid crystal displays is not that they are so expensive but that, given their technological complexity, they are affordable at all. The core of the screen is a sandwich of two flat sheets of glass a few microns apart, with the liquid crystals that form the display medium poured in between. Since the crystals themselves cannot produce light, it's necessary to provide a source—a backlight. A diffuser is then needed to distribute the light evenly across the crystals, as well as front-and-back polarizers to orient the light. And that's just for monochromatic screens. Full-color displays also require expensive red/green/blue filters made of dichromated gelatin—fish glue. To make things fast and bright for, say, a laptop requires another pricey addition: an active-matrix backplane that puts a thin-film transistor behind every pixel. Finally, to manufacture this Byzantine monster takes a superclean factory that won't leave much change out of a billion-dollar bill.

The result is that for the dozen or so firms, mostly Japanese, that have overcome these obstacles to produce active-matrix liquid crystal displays—only to see them become a commodity item, with wafer-thin margins to match—victory has often been Pyrrhic.

Into this chasm of opportunity, with the potential for cheaper but higher-margin and more versatile displays made of common materials like the dyes used in photocopying and photographic paper, come organic light-emitting diodes. Their story begins in 1979, with a scene straight out of a low-budget sci-fi movie. That's when Ching Tang, a Hong Kong-born chemist at Kodak's research laboratories in Rochester, NY, noticed that one of the organic solar cells he was working on was giving off...well, a strange, blue glow. Curiosity aroused, the Kodak scientist launched a long investigation into this phenomenon, known as "organic electroluminescence." His seminal work, reported in *Applied Physics Letters* in 1987, showed that organic materials were efficient converters of electricity into light that could be switched on and off quickly—especially crucial for showing video, where images are updated 50 or 60 times a second and get blurred if the screen can't keep up. Furthermore, Tang noted, these effects could be obtained with low voltage. In short, organic light-emitting

Organic light-emitting diodes represent the only display technology poised to meet third-generation mobile phone standards.

SLIM FILMS

diodes had all the makings of a sensational display technology.

Another breakthrough occurred a year later, when Jeremy Burroughes, a doctoral student at the University of Cambridge's famous Cavendish Laboratory, showed that electroluminescence was characteristic not just of the small molecules studied by Tang but also of far larger polymer molecules *(see "Displaying a Winning Glow,"* TR *January/February 1999).* This was important because it's much easier in principle to make displays out of polymers than out of small molecules. The smaller materials must be vaporized in a vacuum, then patterned through a perforated metal foil or shadow mask—a costly and delicate process. At least for monochrome displays, however, polymers can be deposited by inexpensive "spin-coating," in which they are simply squirted at a rotating target to achieve a uniform surface.

This finding helped catapult the technology into prime time. Burroughes and his professor, Richard Friend, formed Cambridge Display Technology, which established itself alongside Kodak as a key player in the race to commercialize organic light-emitting diodes. Although these pioneers had the field to themselves briefly, they're not alone anymore. David E. Mentley, a senior vice president with San Jose, CA-based market research firm Stanford Resources, estimates that some 90 other companies have since joined the fray. These include giants Philips, DuPont and NEC, as well as startups eMagin and Uniax, a Santa Barbara, CA, firm founded on technology from Nobel laureate chemist Alan Heeger (DuPont bought the company last year). Some license the Cambridge polymer technology; some follow Kodak's small-molecule lead; some pursue their own variants. Whatever the strategy, a savage battle to commercialize organic light-emitting diodes is underway *(see "Global Race for a Better Display," p. 85).*

The reason for all this activity is straightforward: the more they are studied, the more organic light-emitting diodes look to be just about everything their liquid-crystal counterparts are not. For starters, their structure is about as simple as one could imagine: an electrode, some organic stuff, then another electrode. Hook it up to a voltage and, presto, out comes light. There's no backlight, no diffuser, no polarizers or any of the other baggage that goes with liquid crystals *(see infographic below).*

Such simplicity should translate into a manufacturing process between 20 and 50 percent cheaper than liquid crystal display processes. It also means a thinner and lighter screen with far lower power consumption: backlights in conventional screens are a major drain on laptop batteries. In addition, organic light-emitting diodes shine much brighter than their conventional rivals and are visible even in daylight. In short, enthuses analyst Mentley, "They have all the ideal features you look for in a display."

These impressive qualities have sparked a flood of gushy predictions about potential applications for organic light-emitting diodes that range from a

Cleaner Design, Brighter Image

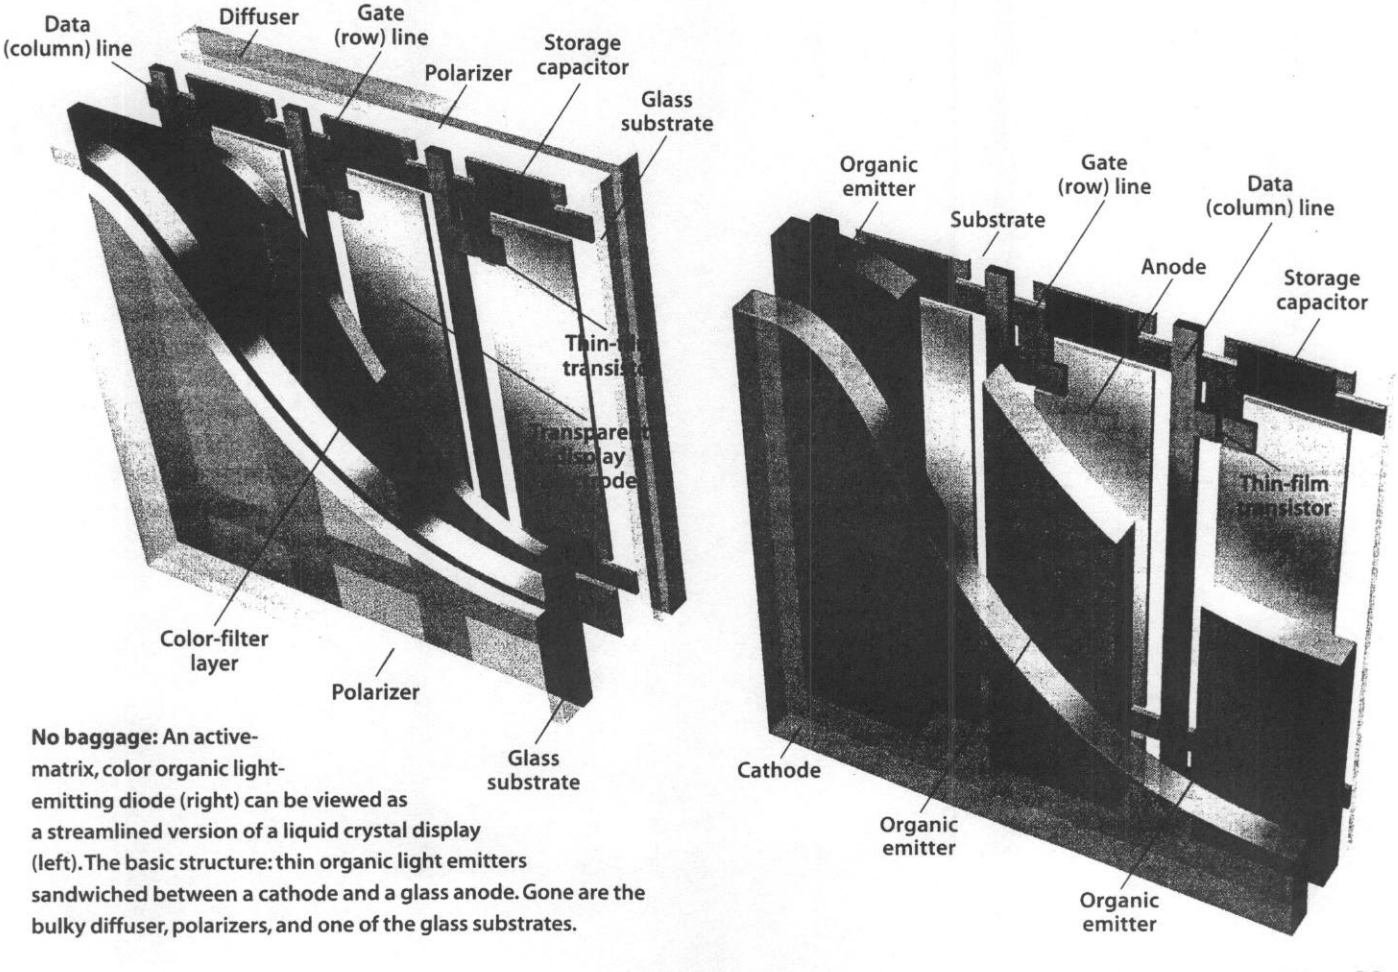

No baggage: An active-matrix, color organic light-emitting diode (right) can be viewed as a streamlined version of a liquid crystal display (left). The basic structure: thin organic light emitters sandwiched between a cathode and a glass anode. Gone are the bulky diffuser, polarizers, and one of the glass substrates.

SLIM FILMS

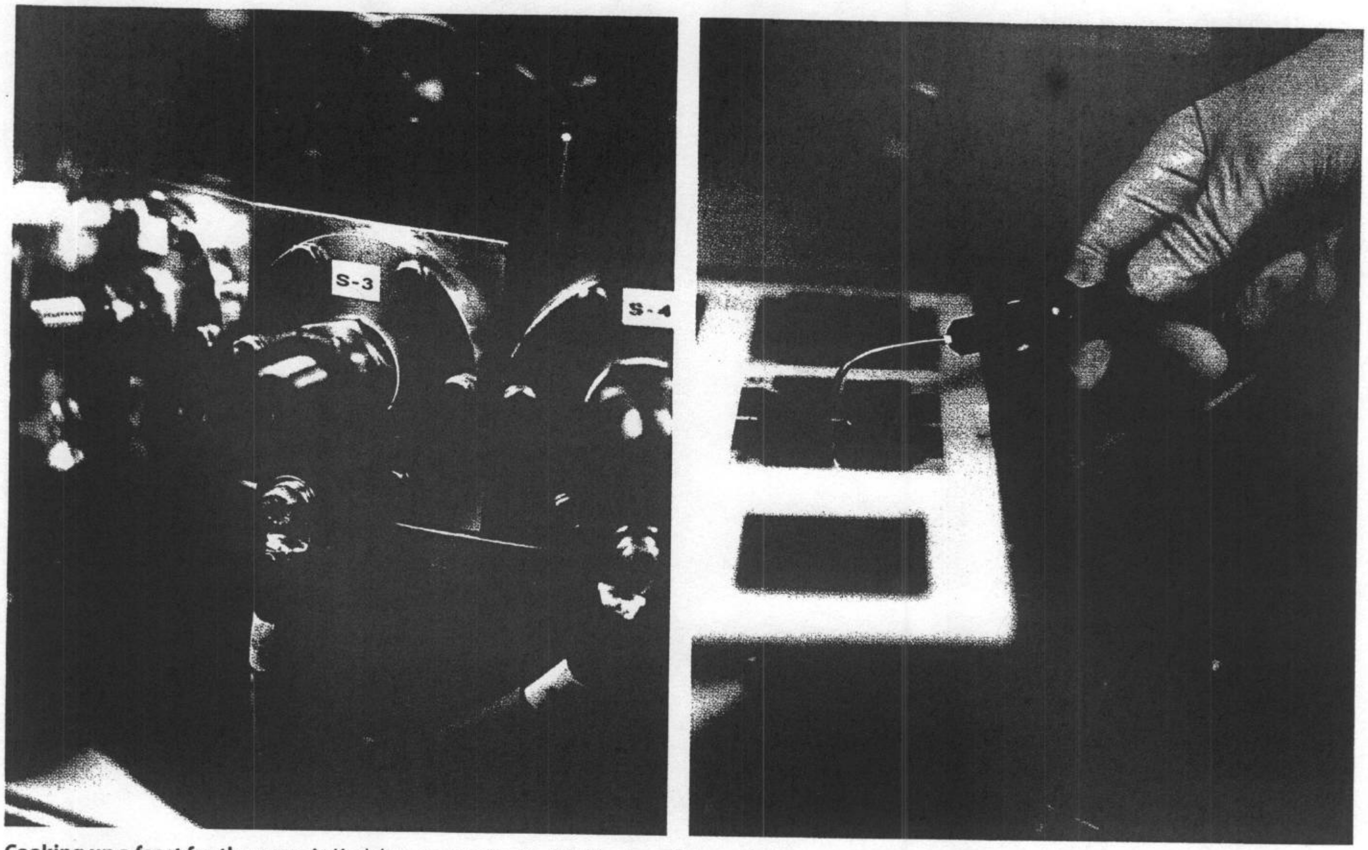

Cooking up a feast for the eyes: At Kodak, a vacuum "coater" (left), where light-emitting organic materials are turned into the ultrathin films used in displays. A researcher wields a pickup tool (right) to load a glass substrate used in device fabrication into the coater.

new generation of affordable wall-hung TVs to highly flexible displays that can be rolled up and carried around like newspapers. It's still too early to evaluate most of these uses, which hinge on clearing formidable engineering hurdles. But there is one arena where the technology is ready to have immediate impact: cell phones.

On May 1, Japanese telecom giant NTT DoCoMo will launch the world's first third-generation mobile phone service. Organic light-emitting diodes represent the only display technology poised to meet all the requirements of the new standard, and the prize for winning the market will be significant. Overall sales of organic light-emitting diodes should grow from \$3 million in 1999 to \$2.7 billion in 2005, as they capture as much as 40 percent of the third-generation market, according to DisplaySearch, a consulting group based in Austin, TX. Even that figure may be conservative. "This is just a vast business, and there is an appetite for a better display," asserts Richard Friend of the University of Cambridge. In the third-generation marketplace, he predicts, the display will become the key differentiator between products.

> "It's paper thin and really beautiful. They put it next to an active-matrix liquid crystal display, and it just blew the LCD away."

The first phone to hit the market with an organic light-emitting diode display is Motorola's \$300 Timeport P8767, which went on sale last September. Manufactured by Pioneer of Japan and based on small-molecule technology licensed from Kodak, it has a passive matrix display that is somewhat limited in its ability to show colors. Still, when stacked up against an equivalent LCD phone—the \$250 Motorola P8167—its brightness and clarity are startling, and perhaps enough to justify the higher price tag. "There's no comparison as to which one people would rather have," claims Daniel Gisser, director of strategic marketing for Kodak's display business unit. "The OLED [organic light-emitting diode] looks so much better."

These are some of the first shots in the battle to dominate the third-generation market. More firepower is on the way. Last July, Seiko Epson demonstrated a cell phone-sized passive matrix screen—this time using polymer technology licensed from Cambridge Display Technology—capable of handling full-color video; production is scheduled to begin in mid-2002. Later this year, Samsung NEC Mobile Display, a joint venture formed last December, expects to begin producing 700,000 full-color, five-centimeter passive-matrix organic light-emitting diode displays a month.

INTO THE RED/GREEN/BLUE YONDER

All this activity represents the first wave of what many insiders believe will be a revolution in displays for devices ranging from small-screen applications such as digital camera viewfinders to handheld computers and laptops. Here, the stakes are far bigger than with cell phones, which DisplaySearch estimates will represent only about a tenth of the total \$75 billion liquid crystal display market in 2005.

Last December, eMagin, a startup in Hopewell Junction, NY, made a first strike on this vast market by garnering the Society for Information Display's Display of the Year Gold Award for its advances in organic light-emitting diode technology, including a prototype 1.5-centimeter active matrix display it hopes will set a new standard for viewfinders. Then there's the hit of the prototypes, the Kodak-Sanyo 14-centimeter active matrix panel exhibited in Long Beach as the future of handheld devices like the Palm Pilot. "It's an absolutely fantastic-looking display," enthuses Nick Colaneri, director of new technology at Uniax, DuPont's just-acquired organic light-emitting diode subsidiary, which has its own offering in the works. "Paper thin, really beautiful. They put it next to an active matrix LCD, and it just blew the LCD away."

Indeed, if the industry buzz is to be believed, organic light-emitting diodes hold the potential to blow away the entire economics of display making. One approach is to dramatically lower manufacturing costs by printing displays directly from an inkjet printer onto a substrate. Last July, Seiko Epson demonstrated a full-color 6.3-centimeter screen made using this method. A next step might be to replace glass with plastic as the screen's substrate. That will produce a display cheaper, lighter and more rugged than today's organic offerings. But the critical point is that once the substrate is plastic, it is easy to imagine a transition from the current, hands-on batch production to an automated rolling process where displays are churned out more like newspapers than chips.

Beyond new forms of manufacturing comes a potentially even bigger step, in which polymers replace silicon in the thin-film transistors that form the active matrix backbone. Philips Electronics, Lucent Technologies, and Plastic Logic, formed last year by Richard Friend, are among those exhibiting prototype polymer transistors. Combined with a plastic substrate, this advance could enable a further milestone: electronic paper *(see "Electronic Paper Turns the Page,"* TR *March 2001).*

Of course, talk about e-paper and other advanced applications is getting ahead of the story, since organic light-emitting diodes face obstacles in nearly every application. One serious current problem is color. Leave the Kodak screens on for a month or so, and the color becomes very nonuniform. Reds and blues die first, leaving a very green display. Cambridge Display Technology has done better with its polymer displays, achieving a working life of 100,000 hours for red and 30,000 hours for green—but just 1,000 hours for blue.

Both technologies are probably good enough for cell phones, which are typically used 200 hours a year and would likely be replaced before the colors start to fade. But such performance is not adequate for handheld or laptop displays, for which several thousand hours of life are required. "Whether the material technology can make it to that level has yet to be proven," admits Kodak's Tang. And the list of technical obstacles grows longer the farther out one looks.

Ultimately, however, the biggest challenge that organic light-emitting diodes face may not be so much technological as commercial. That is, while liquid crystal displays will probably fail to match the attractiveness or performance of organic light-emitting diodes, they will continue to be reliable and affordable—and their manufacturers will no doubt find ingenious ways of further lowering costs and improving capabilities. As analyst Mentley warns, "I haven't heard of any of these LCD guys saying they're just going to fold up and go away. It's going to be a battle."

Still, the battle is joined. And, as Richard Friend notes, few technologies last forever. "I mean, really new things will happen," he asserts. It looks as though organic displays could be one of those really startling new things that come barreling over the technology horizon.

Global Race for a Better Display

Although organic light-emitting diode technology originated in the United States and Britain, firms around the globe are hot on the trail of commercial applications. Following are some of the most prominent players.

SMALL-MOLECULE DISPLAYS

Kodak began licensing passive matrix versions of its organic light-emitting diode technology in 1995. Long-standing licensees include TDK, Nippon Seiki and Pioneer, which has already commercialized the technology for use in car stereo displays and as an option on Motorola's Timeport cell phones. A recent licensee is Lite Array, a California startup with a production base in China.

Last October, Kodak set up its own business unit to pursue active matrix displays, establishing a partnership with Sanyo, a leader in the polysilicon transistors that might form the screens' backplanes. Polysilicon is faster than conventional amorphous-silicon thin-film transistors and therefore a potentially better match for organic light-emitting diode displays. Kodak has also licensed its technology to eMagin, a Hopewell Junction, NY, startup pursuing active matrix microdisplays for near-eye applications such as camera viewfinders.

The Motorola Timeport

COURTESY OF MOTOROLA

POLYMER DISPLAYS

Cambridge Display Technology began licensing its polymer technology in 1997. Its closest relationship is with Seiko Epson, which runs a lab next door to CDT. But the first license went to Philips, a player in cathode-ray tubes, plasma panels, active matrix LCDs and other display technologies. Philips's first product will be a backlight for liquid crystal displays.

Philips is also collaborating with DuPont, which in March acquired Uniax, a Santa Barbara, CA, startup founded on polymer technology from Nobel Laureate Alan Heeger.

The first commercial polymer organic light-emitting diode display—an alphanumeric screen—was to be launched this spring by Taiwan's Delta Optoelectronics, which works under a CDT license and uses materials from Dow Chemical.

CDT's only other publicly announced agreements are with the German chemicals firm Hoechst and Agilent Technologies, a leader in inorganic light-emitting diodes.

Prominent independent makers of small-molecule organic light-emitting diodes include Samsung NEC Mobile Display and Princeton University spinoff Universal Display.

David Ranada

The Ghost of a Machine

Despite an onslaught of new displays, CRTs just won't die

EARLY IN AUGUST, *The New York Times* ran an article proclaiming the end of the picture tube. Cathode-ray tubes (CRTs) are "heading for the dustbin of history much faster than anyone expected," wrote our colleague and contributor Eric Taub. But unlike a dead-end technology like the SACD — whose death I proclaimed in my last column — many characteristics of the CRT live on in HDTV systems.

Way back in 1940, Vladimir Zworykin, a pioneer in the development and commercialization of television, said in his early textbook on TV that the CRT "will continue to keep step with the inevitable advance of television standards." He would probably be as impressed with the technology's 66-year-plus longevity as I am. Zworykin would also be the first to recognize those capabilities and limitations of picture-tube technology that have survived into the HDTV era.

Take scanning, for instance — the process by which an image is disassembled in a camera and reassembled on a TV tube. The concept of line-by-line scanning was fundamental to early television because it was used in both the image sensors (such as Zworykin's Iconoscope, a sort of picture tube in reverse) and in display CRTs. We still speak of pictures being built up of scan lines, even though solid-state CCD image sensors in video cameras, as well as all fixed-pixel display technologies (LCD, DLP, LCoS, plasma), don't have to scan the image line by line — they all sense or display an entire frame all at once. In principle, you could directly hook up a 1080-line CCD camera to a 1080-line fixed pixel display on a pixel-by-pixel basis, though it would take a couple of million cables! Other CRT-derived systems similarly survive into HDTV, including interlaced video encoding and the precise shades of the red, green, and blue primary colors for HDTV (derived from CRT-phosphor colors).

While CRTs won't be missed by anyone with a small apartment, there are people already very concerned about the eventual shutdown of CRT production lines — those whose profession requires them to monitor video quality for transmission or mastering purposes. Almost all of these folks are now making decisions about video quality using relatively small CRT monitors. This isn't a bad thing, since no fixed-pixel technology can reproduce the whole video dynamic range quite as smoothly as a CRT with its truly black blacks.

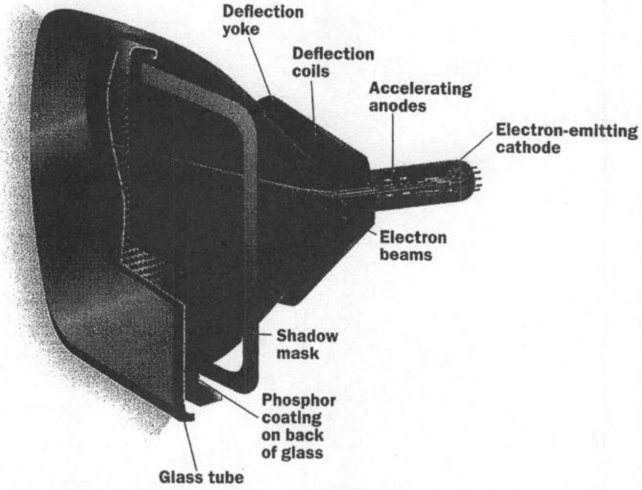

Going but not forgotten: CRTs and their electron guns, shadow masks, phosphor coatings, and massive picture tubes.

But even the largest of the studio-grade CRT-based direct-view HDTV monitors, Sony's BVMA32, measures only 32 inches diagonally, far smaller than a great many of the HDTV fixed-pixel displays already in homes. And while the image it can generate when reproducing HD video can be gloriously realistic — like looking through an open window — its relatively small size can hide defects in the image that may become obvious when enlarged in a home theater. Furthermore, since the BVMA32 costs $42,850, not all studios working on DVDs or HD video can afford one, with the result that many video-quality decisions are made with far less capable devices, all CRT monitors.

The research and development of all the major video-codec systems (MPEG-2, MPEG-4/H.264, and VC-1) used with the DVD or high-definition media were conducted with images monitored on CRTs. Crucial decisions as to how these systems work, especially in how they're structured to minimize the visibility of encoding artifacts, were made using CRTs. So while it's hard to find an explicit expression of CRT technology in these video codec systems, their very architectures and operation were influenced by CRT characteristics, for good or ill.

This only starts to become a problem when you watch media at home. To the extent that your monitor or TV deviates from CRT-like size and visual behavior, you may see increasing problems with an image that looked fine when it left the mastering studio. I've found encoding artifacts such as mosquito noise and blocking to be more visible with fixed-pixel technologies than they are with CRTs of the same size, at least with material of DVD resolution.

What the pro-video folks desperately need is a fixed-pixel technology that produces very large, CRT-like images. Zworykin seems to have been prescient when he wrote 66 years ago: "Today, owing largely to the advent of cathode-ray television systems, a recognizable picture can be taken for granted, and the problem is one of obtaining a high-definition picture." **S&V**

ILLUSTRATION BY DIMITRY SCHIDLOWSKY

Inner Workings

David Katzmaier

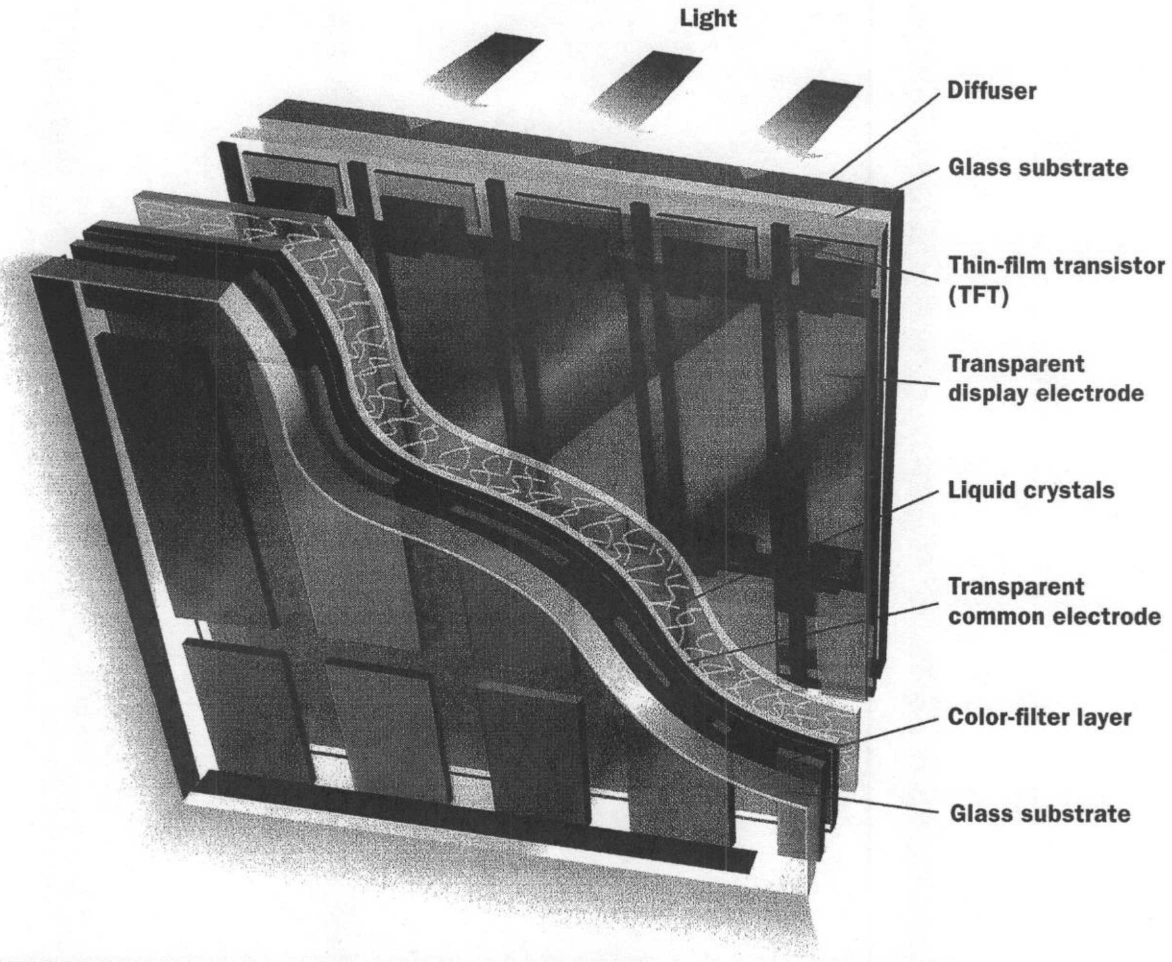

Inside a liquid-crystal display (LCD)

The granddaddy of fixed-pixel technologies, LCDs first appeared in pocket calculators in the early 1970s. LCD technology is amazingly versatile, able to power front projectors, rear-projection TVs, and flat-panel displays.

In an LCD, a matrix of thin-film transistors (TFTs) supplies voltage to liquid crystal-filled cells sandwiched between two sheets of glass. As with plasma panels, a trio of red, green, and blue cells makes up one pixel. When hit with an electrical charge, the crystals "untwist" to filter light generated by a lamp behind the screen (for flat-panel TVs) or a lamp shining through a small LCD chip (for projection TVs).

LCDs have a lot of advantages. Direct-view models are only a few inches deep, and are available in screen sizes as small as 15 inches. Compared to plasmas, LCDs usually have higher resolution at comparable screen sizes, and relatively affordable 1080p LCDs are becoming available. LCD panels weigh less and consume less power than plasmas. And, unlike DLP models, LCD projectors aren't subject to the rainbow effect — flashes of multicolored light sometimes seen around bright objects in motion against a dark background.

But of the fixed-pixel technologies (which also include plasma, DLP, and LCoS), LCD has historically had the most trouble producing deep blacks. Some light always passes through when the crystals untwist, so the best black is usually a very dark gray. The newest LCDs, however, are much better with this, and many rival the best plasmas at producing deep blacks. Also, because of the way light goes through LCD cells, flat-panel LCDs still have narrower viewing angles than plasma TVs.

LCD has always dominated the flat-panel market under 42 inches, but until the last year or so high prices have slowed widespread adoption of bigger models. In 2006, however, 40- and 42-inch LCDs have dropped to plasma price levels, and they're experiencing a surge in popularity. Rear-projection LCDs compete with DLP sets and sell for about the same prices, while front-projection LCDs are generally less expensive for the same resolution than their DLP competition. **S&V**

ILLUSTRATION BY DIMITRY SCHIDLOVSKY

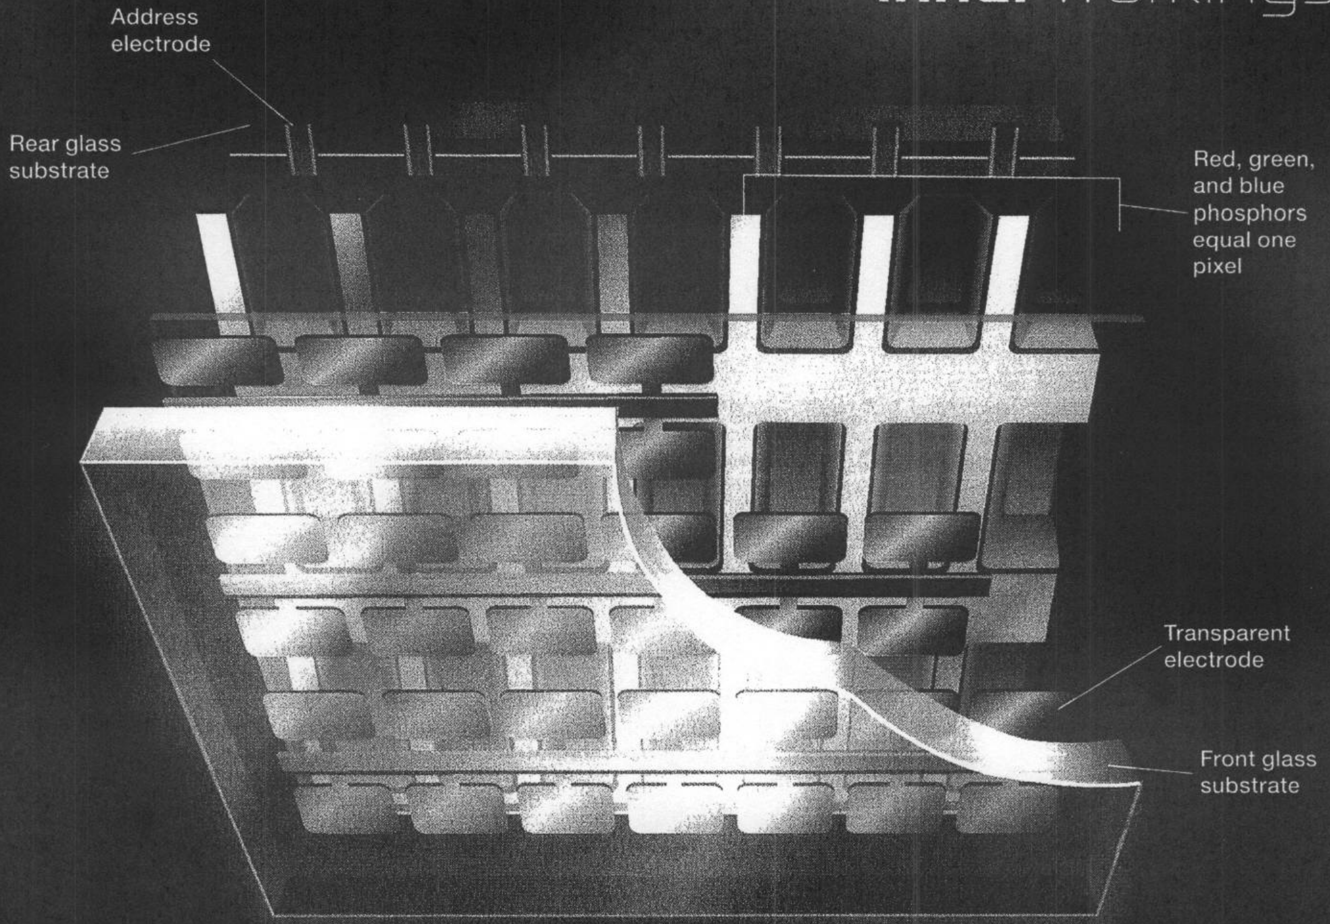

inside a plasma screen

BY DAVID KATZMAIER

The "plasma" in a plasma TV isn't the clear part of blood, or the phase of matter that makes up the billions of stars in the universe. It's an inert gas — usually a mixture of neon and xenon — and it's a big part of what allows these TVs to measure just 3 to 6 inches thick.

Unlike a bulky tube TV, a plasma display uses the combination of thousands or even millions of discrete pixels to form a picture. A plasma with a resolution of 1,024 x 768, for example, has 786,432 (1,024 times 768) pixels. If that sounds like a lot of dots, it's really just a third of the story. Each pixel is actually made up of three separate sub-pixels: one each for red, green, and blue. (Next time you're near a plasma, put your nose up to a white area — you'll be able to see all three sub-pixels, and smell that new-plasma smell.)

The sub-pixels, also known as cells, are basically tiny pockets filled with gaseous plasma. The inside of each cell is coated with red, green, or blue phosphors, like the ones that line the inside of a traditional picture tube. Using phosphors allows plasma TVs to closely mimic the color palette of tube TVs — but it also makes them vulnerable to burn-in, which is basically uneven phosphor wear.

As you can see from the illustration, each cell has a pair of electrodes associated with it — one in the bottom of the cell, and one above — arranged in two layers of glass substrate. The top electrode layer is transparent — you actually look through it to see the picture. When current is run through the electrode layers, the gas inside each cell is stimulated, releasing ultraviolet photons. Those photons react with the phosphors, producing reds, greens, or blues of various intensities. The combination of colors in each complete pixel can produce millions of real-world colors, from white to black and everything in between. **S&V**

ILLUSTRATIONS ON PAGES 94 & 95 BY DMITRY SCHIDLOVSKY

Third-Generation AIT Debuts

BY M. DAVID STONE

As hard drive capacity continues to grow at a dizzying rate, tape drive manufacturers are working hard to stay ahead. Among the latest tape formats is third-generation Advanced Intelligent Tape technology (AIT-3), from Sony Electronics, which you'll find in the **Sony SDX-D700C** tape drive.

The drive boasts fast throughput, 100GB native capacity per cartridge, and a claimed 260GB compressed capacity. That's an enormous amount of storage for a single cartridge, especially one that is about the size of a thick cassette tape.

AIT-3 gets its high capacity in part from helical-scan recording, which is known for offering high data densities. The technology gets another boost from using Advanced Lossless Data Compression (originally developed by IBM), which offers a claimed 2.6-to-1 compression ratio instead of the more common 2-to-1.

Another welcome touch is a 64K chip in each cartridge that stores the tape log, so the drive can find the location of files by reading the chip. The payoff is a claimed average access time of under 27 seconds.

Setup is simple: Just plug the cable in and go. The SDX-D700C comes without backup software but includes a utility for copying from non-AIT drives to the SDX-D700C, which makes an upgrade more convenient. (The drive is compatible with both earlier AIT formats.) We used an Adaptec 29160 SCSI card and Computer Associate's BrightStor ARCserve 2000 backup software running on Windows 2000 Server.

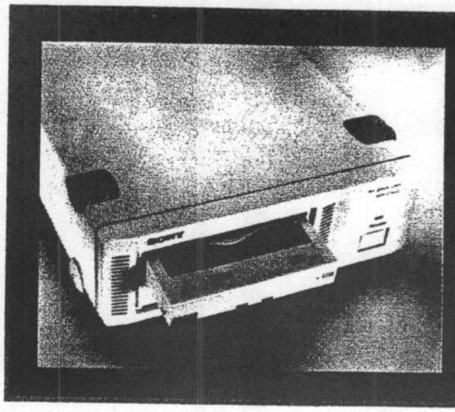

The Sony SDX-D700C, which supports the new AIT-3 tape format, is a good choice for server backup.

On our tests, data throughput came out to 13.2 MBps, or nearly 800MB per minute, with the drive backing up roughly 41GB in less than 52 minutes. This is on the low end of the claimed throughput of 12 MBps to 31.2 MBps, indicating that there was relatively little compression for our test data.

In a nutshell, the Sony SDX-D700C is a practical and useful upgrade that is well worth considering if you're deciding on a tape format.

Sony SDX-D700C

Direct price: $4,225. Requires: SCSI-2 connection (Ultra 160 Wide SCSI LVD/SE recommended), backup software. Sony Electronics Inc., 800-352-7669, 408-432-1600, www.storagebysony.com. ●●●●○

Telecommuter Security Appliance

BY BRUCE AND MARGE BROWN

Installing a hardware firewall for your home network is a great idea. But if you use your home network and broadband Internet access for both personal and corporate-connected computing, other users on your home network can inadvertently threaten your business network's security. The **SonicWall Tele3 TZ** ($550 street) establishes a "trusted zone" to isolate corporate-connected PCs from your home network while still providing firewall protection and VPN access for the entire configuration.

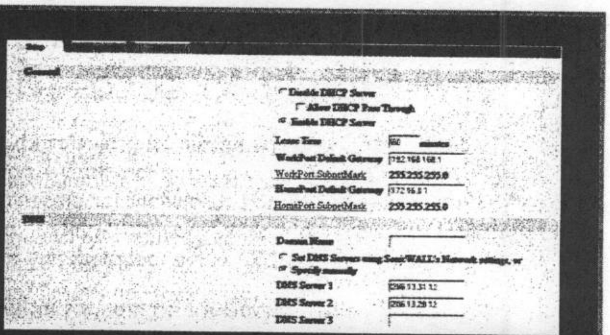

○ Disable DHCP Server

☐ Allow DHCP Pass Through

● Enable DHCP Server

Lease Time 60 minutes

WorkPort Default Gateway 192.168.168.1

WorkPort SubnetMask 255.255.255.0

HomePort Default Gateway 172.16.8.1

HomePort SubnetMask 255.255.255.0

Domain Name

○ Set DNS Servers using SonicWALL's Network settings, or

● Specify manually

DNS Server 1 206.13.31.12

DNS Server 2 206.13.28.12

DNS Server 3

When you configure the SonicWall Tele3 TZ for both home-network and protected-work PCs, separate network subnets are established. You can manage all the subnets through the SonicWall Administration application.

Setting up the TZ for basic operation is straightforward. The unit establishes separate, parallel connections to the Internet for home and corporate use. You can configure what traffic (if any) is allowed between the security zones. The TZ has three RJ-45 ports: The WAN port connects directly to a broadband cable or DSL modem or router. The HomePort connects to your home network via a hub or switch (or you can directly connect a single PC), and the WorkPort connects to a single business-related PC (or a hub or switch for those who use multiple PCs to access corporate systems).

The TZ's security specifications are impressive, including an ICSA-certified stateful packet inspection firewall, 3DES (168-bit) encryption, and licenses for five IPsec VPN tunnels and five firewall users. You can allocate the bandwidth between the two network zones, and if your corporation manages firewalls centrally, the TZ is Global Management System v.2.5–compatible.

We connected the TZ's WAN port to our home network's DSL modem via a D-Link DI-714 combination wireless access point, router, switch, and DHCP server. We plugged a network switch into the HomePort and a single notebook PC into the TZ's WorkPort. When we ran the browser-based configuration program, the TZ configured one subnet for the computers connected to the HomePort via the switch and a separate subnet for the PC connected to the WorkPort (which we used to connect to the Ziff Davis Media server-based e-mail system).

DHCP functions in the device automatically managed and assigned IP addresses to the PCs on each subnet. We also had wireless PCs accessing the Web via the D-Link DI-714, so we had three subnets working simultaneously. The networks were invisible to one another. All were able to get to the Internet, proving that the two-layered DHCP function was working correctly.

The SonicWall Tele3 TZ isn't cheap. But for companies with remote employees (or frequent work-at-homers), the cost is well worth the extra protection.

SonicWall Tele3 TZ

Street price: $550. Requires: Cable, DSL, or network WAN Internet connection, Microsoft Windows 95 or later. SonicWall Inc., 408-745-9600, www.sonicwall.com. ●●●●○

Wise Drives

Manufacturers may soon be unlocking the computing power hidden in your hard-disk drive

BY GORDON F. HUGHES
UNIVERSITY OF CALIFORNIA
(SAN DIEGO)

If disk drives were enabled to perform a host of tasks that have traditionally been the work of CPUs, dramatic low-cost improvements could be made in computing environments ranging from laptops to storage networks and drive makers would gain new profit centers. But this will only happen if computer and drive manufacturers seize the opportunity and create new interfaces for these so-called intelligent disk drives.

Disk drives currently come equipped with 32-bit internal microprocessors, several megabytes of internal RAM, and tens of megabytes of disk capacity reserved for internal drive purposes. This onboard computing power has increased alongside improvements in the speed and capacity of disk drives, and has already allowed them to take on several intelligent tasks beyond basic data storage and retrieval, such as detecting imminent drive failures. This power could be exploited for even more tasks, and the addition of some relatively inexpensive extra processing power would open the door to a huge array of functions, such as searching and encryption, and so shift some of the processing burden from CPUs and networks to disk drives.

Modern drives are nearly all standardized: they have common physical dimensions, electrical and data connections, and operational commands—all defined by ANSI standards. Interfaces between computers on one side and drives on the other constitute an agreed contract on the rules for storing and retrieving data. Adding intelligent drive features may require changing these entrenched standards, and may also require changes in operating system and application software. High-level control of intelligent features may best reside in user applications, allowing them to tune a drive's performance to their specific needs.

It is my hope that after reading this article, developers of such software would become involved in the interface committees of organisations such as the

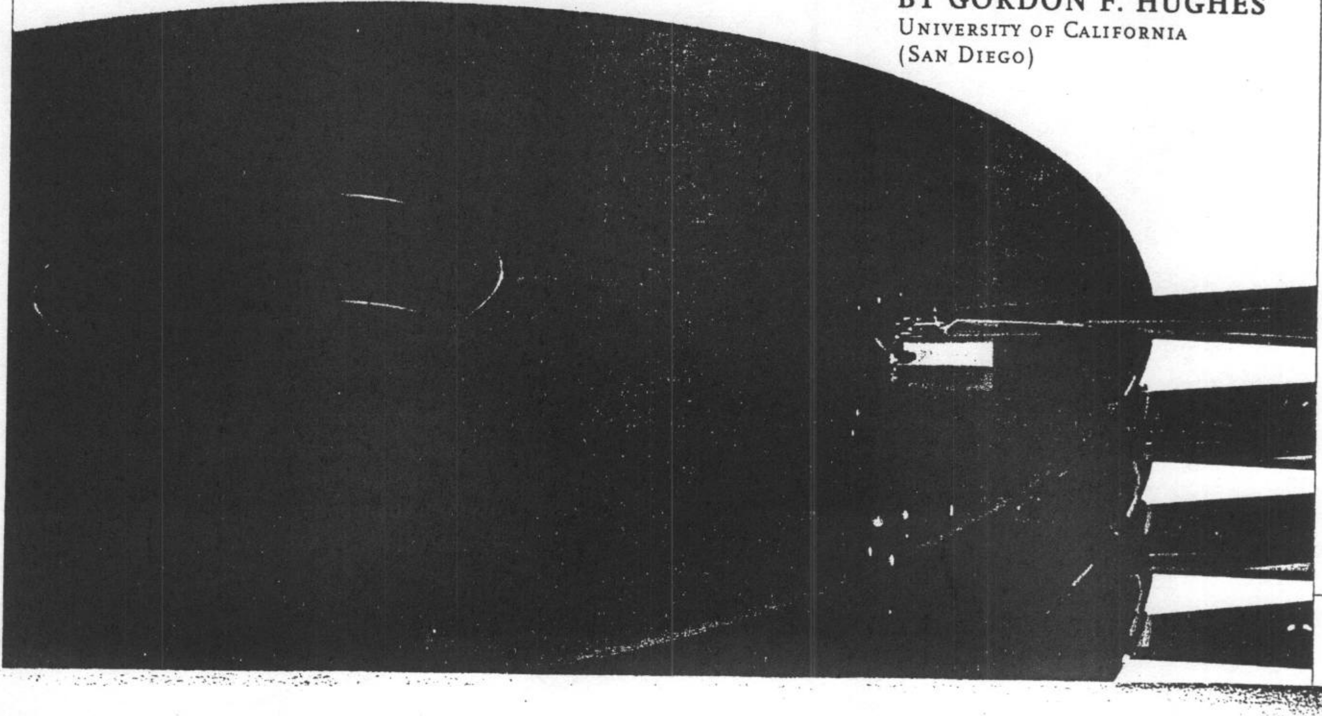

LARRY GILPIN/STONE

Storage Network Industry Association (SNIA), which currently consist primarily of drive makers and drive buyers. But why should intelligent features appear now, when for decades disk drive development has been fixated on faster access to more data?

For 50 years hard disks have acted as relatively dumb devices, simply serving and storing blocks of data. The first generation of hard-disk drives were the size of refrigerators and stored a few megabytes for mainframe computers. Since 1992, drive storage capacity has been increasing at an annual rate of 60 percent, and over 100 percent in recent years—a rate exceeding the 18-month doubling of Moore's law for IC complexity [see graph below]. Single drives today can store up to 200 GB and standard 3.5-inch disk drives are in everything from PCs to supercomputers (laptop computers use smaller 2.5-inch drives).

But the drive manufacturers have to some extent become victims of their own success. The industry has matured and is delivering highly reliable low-cost storage, at capacities that exceed the real needs of the majority of today's end users, typically desktop PC users. All drive manufacturers offer similar-performance mass-market products, which are seen as commodity items.

In my position as associate director of the Center for Magnetic Recording Research (CMRR) at the University of California, San Diego, I see how this commoditization shrinks the number of drive manufacturers and drive component makers, and reduces the number of research jobs for our students—hampering the very development that produced the remarkable drive technology advances in the first place. Customizable intelligent features would allow drive makers to distinguish their products and provide a competitive edge and an impetus for innovation.

The most common drive interface in consumer desktop computers is the advanced technology attachment (ATA) interface, also known as integrated drive electronics (IDE). The small computer system interface (SCSI) is prevalent in high-performance systems. Each interface consists of an adapter integrated into the host at one end connected via a cable to the disk drive(s) at the other end.

An Embarrassment of Riches

The bit density of disk drives is growing even faster than the transistor population of microprocessor chips (which still obeys Moore's law). With disk capacity no longer an issue for most users, drive manufacturers would do well to add features to stay ahead of the competition.

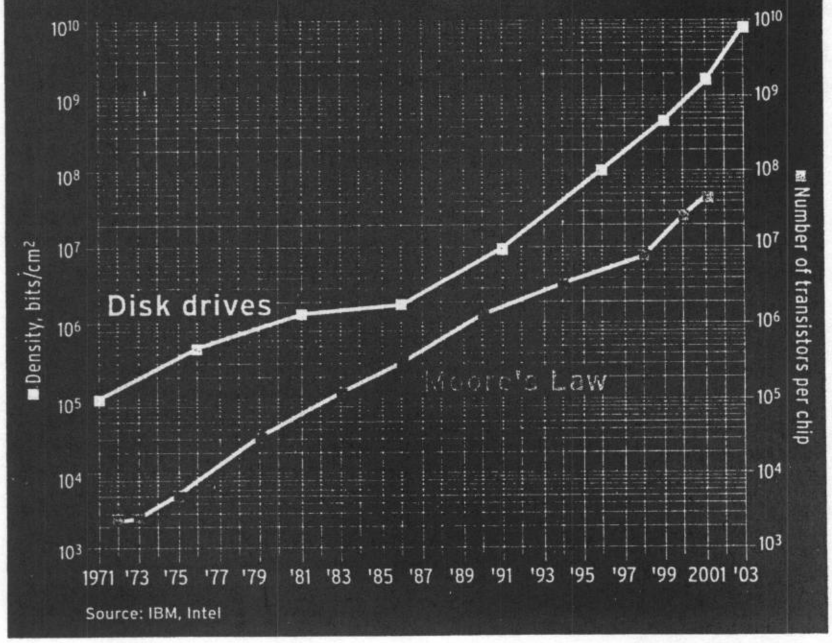

Source: IBM, Intel

In both ATA- and SCSI-based systems, computers commonly divide a data file into a set of so-called logical blocks whose size ranges (depending on the platform and operating system) from 512 bytes to 4096 bytes and up. The logical blocks are passed over the interface to the drive, which stores them as physical blocks, usually of the same byte size, by magnetically recording bits on a cobalt alloy magnetic film coated on the disk surfaces. The physical blocks are stored in concentric tracks subdivided into angular sectors on the disks, which are constantly rotating during operation.

Conversely, in accessing stored data, the drive reads back each physical block magnetically as a noisy analog electronic signal and translates it into digital bits to reconstitute the physical block.

Already pretty smart

To do all this, disk drives already perform complex tasks hidden from the host computer. Errors can be detected and corrected. An impending crash can be anticipated. An array of disk drives can even recover data on their own from a disk that has crashed, without involving the system at large.

Data integrity is ensured by sophisticated algorithms that clean up the raw data from the disk using digital error checking and correction codes (ECC). ECC code bits are recorded along with user data in the physical blocks. A first level of error correction is built into the drive control electronics, to deliver corrected user data without delays while decreasing the probability of an incorrect bit from about one in a million at the playback magnetic bit detector to less than one in 10 trillion for the delivered logical block.

If this first line of error correction fails, additional levels of off-line error recovery are invoked. In disk drives, off-line means that the drive stops delivering data at full speed as time is needed to correct a data error. Typical off-line error-recovery techniques include: attempting to reread the data on a second disk revolution; rereading the data by positioning the read/write head slightly off track; and using a higher-level ECC algorithm in the drive microprocessor firmware. Hard-to-read phys-

PHOTODISC

ical blocks are reassigned (re-recorded) to alternative disk locations and the original locations are not reused.

The philosophy here is that returning error-free user data is so important that delays are acceptable. For some applications, though, this assumption may not be correct—users of a multimedia player showing a digital video clip may find a pause more distracting than the occasional dropped frame. Currently there is no good way to allow a drive to distinguish between different types of user data.

Another intelligent feature of today's disk drives reduces the chances that a hard-disk crash will leave a user without data. Disk drives have failure rates of only about 1 percent per year, but the loss of a user's data can easily be more costly than the drive itself. If they were forewarned of an impending drive failure, users could take steps to back up their data onto another storage device.

PHOTODISC

"Commodization [of disk drives] shrinks the number of manufacturers and reduces the number of research jobs"

To this end, in 1994 the drive industry adopted a standardized specification for such failure warnings called Smart (self-monitoring and reporting technology) at the request of Compaq Computer Corp. Smart is based on monitoring for excessive internal drive errors, such as bit-read errors and track-seek errors. A failure-warning algorithm running in the drive's microprocessor checks whether error rates exceed a threshold value and if they do, produces a warning that is sent over the drive interface to the host CPU.

Virtually all disk drives today have this failure-warning system built in. All the same, some computers simply ignore the warning. A research program on Smart here at the Center indicates considerably higher Smart warning accuracy is possible with improved drive firmware.

A third intelligent drive task today helps high-end storage systems, those using a RAID-5 disk organization (redundant array of independent disks). These systems store data blocks over multiple drives to gain the speed of parallel access in such a way that drives can fail and be replaced without loss of user data. Such computer systems mostly use the SCSI interface, which allows drives to recover the data autonomously from the failed drive, without burdening the host computer system with the work. So drives are already networking without computer help.

Better, faster than before

But if more intelligence is to be added to disk drives in the form of new firmware and/or extra onboard processing power, computer manufacturers must be convinced of the benefits of the new features to them and to application developers and end users. An example would be in-drive tools for handling data delays, integrity, and corruption control. These would allow so-called disk-drive quality-of-service (QoS) features.

In retrieving user data, drives have variable delays due to track-seek and disk-rotation times. A drive may have to wait several milliseconds for its read/write heads to travel back and forth over a span of tracks, and then wait a few more milliseconds for disk rotation to bring the desired data block under a head. When beginning an access, drives can internally determine these motion delays, but there is currently no way of informing the application using the disk drive—even though applications like streaming video cannot use data that arrives too late.

Interface QoS commands could be added to provide precise information about delays and allow applications to abandon a data access if the delays were deemed excessive, or to optimize a group of data accesses. Today, drive QoS refers to nothing more precise than the average data-access speed of a group of drives.

Off-line error correction also introduces unexpected data delays. Currently, drives sold for video streaming applications may simply have off-line error correction disabled to guarantee throughput, but QoS commands could allow computers to turn the off-line correction on and off depending on the file type—off for MP3 music data files and on for application software files, for example.

Sending data over a network to a drive or array of drives can also cause data errors, which must be detected. Currently, drives calculate a data-error detection code (called a CRC polynomial code) before sending blocks across the interface to the computer, intended to allow checking of storage device connectors and cables. The CRC system could be extended to assist storage network administrators, allowing fault tracing down to individual hardware such as routers, disk and tape interfaces, and drives, and to environmental factors such as the power supply and ambient temperature.

Orderly storage

Storage systems such as file or database servers offer a fertile ground for other improvements. Many of them cache logical blocks in their RAM during drive accesses. This helps because consecutive logical block operations are often made on the same or the next logical block in the sequence, forming a complete file or database record. Simultaneously, drives perform physical block caching (such as reading all the sectors on a track into the drive's on-board RAM when a single sector is read) to speed access because the next logical block needed is often also the next physical block.

These two caching operations are unaware of each other and may be redundant. Worse, they may even mutually interfere, as when the storage system demands a chain of logical blocks to cache, which may be scattered over several disk tracks, while the disk drive is trying to cache all the physical blocks in each track. Intelligent drives could allow servers to control drive caching on the level of entire files. To this end, an Object-Based Storage Devices (OSD) standard has been proposed for SCSI systems by the SNIA.

A storage object may be any number of things. It can be a file consisting of an ordered set of logical data blocks, a database containing many such files, or just a single application record such as a database record of one transaction. Information about the data is also stored in an object, which can include QoS, security, caching, and backup requirements. OSD disk drives could perform data preaccess, which could put sequential logical blocks in the drive's cache memory, instead of sequential physical blocks. Drives could transmit them at the full interface transfer rate without waiting for disk rotation or track seeks. And once drives start dealing with data objects instead of their individual logical blocks, even more possibilities open up. Storage systems, which must handle such tasks as data mining, backup, and compression, could be implemented and scaled up much more efficiently.

Object-aware drives could speed up the process of making backups by maintaining internal lists of which object blocks have been rewritten since the last backup (currently a housekeeping task managed by host computer backup software). Only altered blocks would need to be transferred to an incremental backup system, instead of entire files.

Data compression on individual files could also be performed by OSD-based systems. Without OSD, compressing data at the drive level can generate a block consisting of data from multiple files. If any errors occur during compression, all those files are at risk. Another risk is storage capacity overflow, because the size of a compressed file isn't known until its compression is finished.

Spreading the load

But data mining has perhaps the most to gain from object-aware drives. Data mining searches are a major bottleneck for the data farms that provide storage for such applications as e-commerce and scientific analysis of very large data sets such as the human genome.

Designers have long known that these intensive data mining searches benefit from parallel or distributed architectures. Dealing with such applications has given rise to many solutions in which storage intelligence has been moved as near to the drive as possible, but ultimately still resides on the computer side of the drive interface.

Since the 1980s, specialized database systems have been built that use a main computer to divide search jobs into independent tasks that are sent to subprocessors for execution, each having its own memory and own local disk drives. The parallelism allowed searches of terabyte-sized databases with an order of magnitude performance increase over normal serial computer systems. But these specialized systems were expensive—and they still just used standard disk drives.

Today database computers more commonly use storage systems based on RAID systems using fiber-channel SCSI drives. Damaged disk drives in these redundant arrays can be replaced without losing data. Up to 127 drives can be daisy-chained in a loop that can accommodate two simultaneous drive commands. Because many drives are running in parallel, an array is also fast. Yet the maximum data transfer rate of the array can be far below the combined

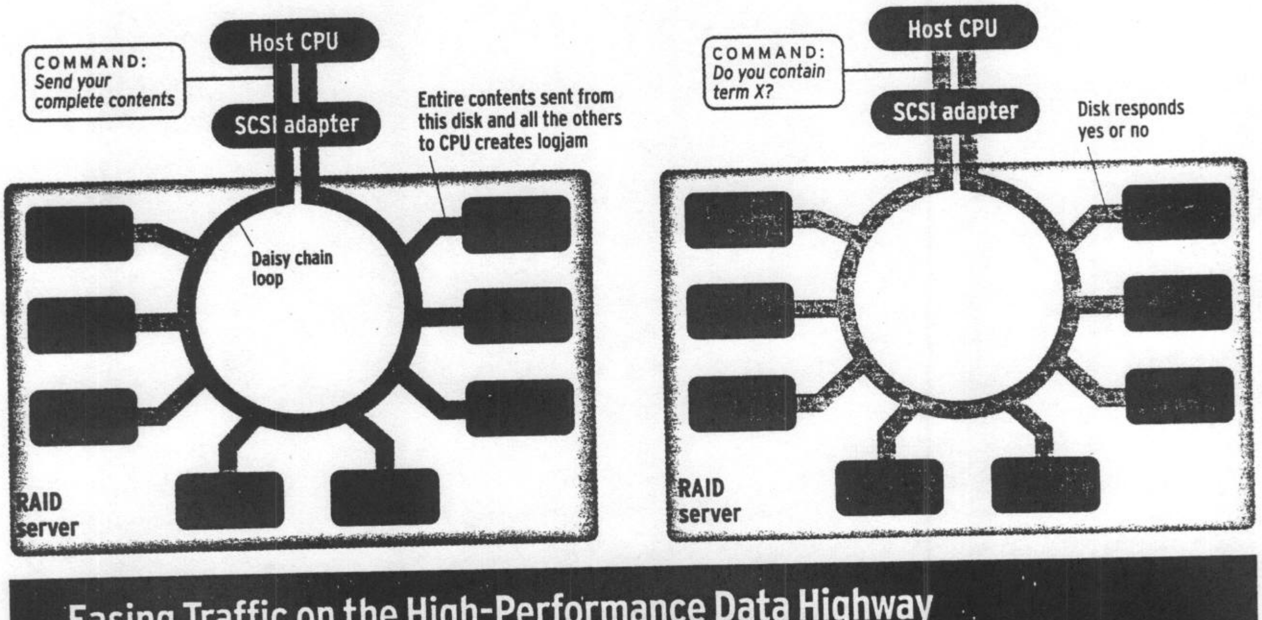

Easing Traffic on the High-Performance Data Highway

Currently, data searching on a RAID server is inefficient. To find records containing a given term all the disk drives must send their entire contents back to the host CPU for examination, causing data congestion in the daisy chain loop [left]. Intelligent drives could be sent the search term and autonomously return just those records that contain it, allowing faster searches with less congestion [right].

transfer rate of all the drives in a loop—because all the drives share the bandwidth of a single interface connection. For example, a database query requiring searching through all the drives could be as little as 2/127, or 1.6 percent, as fast as the maximum data transfer rate of all the drives operating independently.

As a consequence, in storage systems where speed is a concern, loops are limited to five or six drives each, trading off cost for performance. A technique known as fiber-channel switched fabric can also address the bottleneck by establishing direct independent connections between the storage computer and the drives, but again at increased cost.

Putting low-level search intelligence inside drives could be faster and far less expensive. The database application would break searches into individual commands, which would be sent out simultaneously to all drives. Only successful search results from drives would be returned. [See illustration, opposite page.]

"The encrypted contents of a disk could be destroyed in an instant ... imagine such a feature onboard the EP-3 spyplane forced to make an emergency landing in China"

Inner security

Intelligent disk drives could also make computers much more secure and alleviate some of the concerns that have arisen in recent years in military, governmental, and business circles when laptops or hard disks themselves have gone missing.

With current systems, normal erasure of disk data does not necessarily prevent later retrieval of that data. Indeed, drives are designed to resist accidental erasure by a user. Special programs are widely available for retrieving apparently erased files, forcing security-conscious organizations to take precautions to ensure that disk drives have been completely wiped before being disposed of.

Commands that aid in this disk sanitizing were added to both the ATA and SCSI drive specifications at CMRR's request. The secure-erase process is similar to whole-drive formatting because it overwrites all user-accessible data areas on a drive, including any disk flaw reassignments. Still, it can take over an hour to repeatedly overwrite all the blocks on a modern 180-GB drive—if data is only overwritten once, it may be recovered by scanning the disk platter with a magnetic force microscope. Also, secure erase is still an optional feature in the specs and not always implemented. Nor does secure erase eliminate the problem of securing data on disks in active use.

TO PROBE FURTHER, SEE PAGE 63

To protect data in active use, operating systems such as Microsoft Windows XP and Linux provide file encryption. However, encryption at the software level is complex to manage. Implementing file encryption keys and algorithms in disk drives would be easier, with the drives transparently encrypting the data. (Decrypting data from drives would remain an external software task.)

The long execution time penalty of secure erase would vanish, because the encrypted contents of an entire disk could be effectively destroyed in an instant by simply deleting the external software decryption keys (securely, of course). One can imagine the utility of such a feature onboard, for example, the EP-3 spy plane that was forced to make an emergency landing in Hainan last year as the crew hastily destroyed disk drives and other sensitive equipment in a bid to prevent information falling into Chinese hands.

Contingent arrival

While intelligent drive features can be added without interface specification changes by using the "vendor reserved commands" already defined in ATA and SCSI, this limits the features to just those agreed on between a single pair of drive and computer manufacturers. Such an arrangement may be a useful proving ground for a new feature, but for widespread use its input/output and command requirements need to appear in the interface specification.

In short, there must be an industry consensus that the task is of general appeal and offers market opportunities for multiple computer and drive companies. A recent example is the addition of the specialized requirements for flash memory PC data-storage cards to the ATA interface specs.

Computer application developers, such as large database search engine designers, may need to first see potential benefits from QoS features demonstrated. University research projects may be the initial step (RAID storage systems were first studied at the University of California, Berkeley). Computer simulation and monitoring of storage networks can also allow assessing performance improvements offered by intelligent features in advance of actual deployment.

Object storage devices have been studied at Carnegie Mellon University, and are under study by the SCSI spec committee. Interested members of the SNIA trade group are moving this proposal forward by working on the definition of an OSD standard.

Customers always determine the acceptance of new technologies, and many of these intelligent drive tasks are of primary benefit to final computer users, who are not the direct buyers of disk drives. For example, although Smart is advertised by PC makers Dell and Compaq as a principal hard-disk feature, PC makers are not currently asking for the higher-accuracy Smart warning discussed above, although it adds no drive cost. This may offer a market opportunity window for the first drive and computer makers to offer it. When developers and users become aware of potential new intelligent features, and ask computer and storage system makers for them, then drive makers will design them in. ●

Stephen Cass, *Editor*