VOLUME XVIII NUMBER 10 OCTOBER 1998

Computer Technology Review®

TECHNOLOGY SOLUTIONS FOR THE INTEGRATION/RESELLER MARKET

A New Tape Recording Technology For Data Backup

by Juan Rodriguez

Mark Twain once said after reading his own obituary in a newspaper, "reports of my death have been greatly exaggerated." Since shortly after it emerged as the standard for the storage, backup, and retrieval of data, many in the computer industry have maintained that tape storage technology had one foot in the grave. These perennial predictions of tape mortality, like those of Mr. Twain's, have until now proven to be greatly exaggerated.

However, in recent years the pace of tape technology innovation has slowed considerably. If this reduced pace of invention were to continue, tape will become an obsolete technology and reports of its death will sadly become true.

Tape Vs. Disk Technology

As you can see from Fig 1, disk technology has been improving price/performance at a much faster rate than tape and, since early 1998, has actually become cheaper than tape. If the tape industry is to preserve its competitiveness, it must do two things: it must re-establish its price/performance advantage and then, it must preserve it. Since price/performance is a necessary-but-not sufficient condition for competitiveness, there are a number of limiting features and problems that must be improved and/or fixed.

Limitations Of Streaming Tape Technology

Streaming is the basic technique used to transfer data in all of today's linear and helical scan tape drives. Streaming tape drives operate by reading an entire track, thousands of bytes long at a fixed speed. Streaming tape technologies require track-following (or 'signal-following') to maintain proper head to track alignment while the track is read in its entirety at a fixed speed. The drive mechanism and media tolerances must be tightly controlled to maintain a very precise alignment between the path traced by the heads and the written data tracks on a tape. These fundamental features of streaming devices suffer from a number of inherent weaknesses that negatively impact reliability, performance, and product cost.

Reliability Degradation

* **Back-hitching** - Tape drives often send or receive data faster than a host system can receive or send it. A streaming tape drive cannot slow down to match the host's speed; it must come to a complete stop, back up, wait until the host is ready for the next transfer, accelerate to the appropriate speed, and continue the

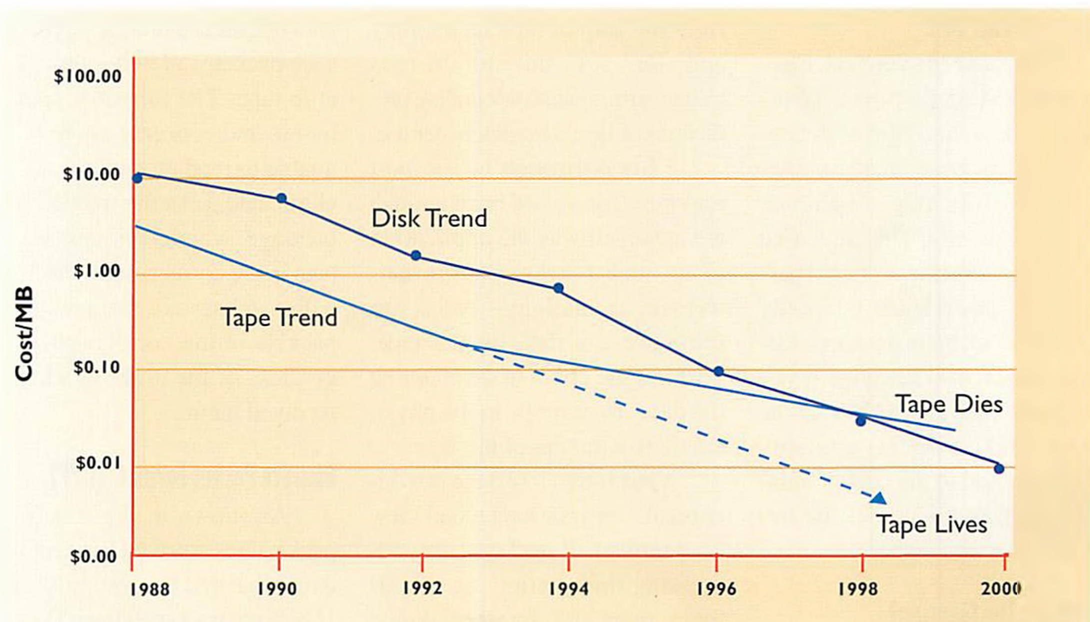

Fig 1 Disk technology has been improving price/performance at a much faster rate than tape and, since early 1998, has actually become cheaper than tape.

transfer. This process, called back-hitching, induces extremely high transient forces in the drive mechanism, reducing the drive's mechanical reliability and increasing tape wear—ultimately impacting the reliability of data resident on tape. Hardware compression, used by many streaming devices to increase effective tape capacity, only increases the frequency of back-hitching events.

* **Track-Following** - Through use and over time, tracks tilt and bend for a variety of reasons including temperature, humidity, tape wear, and drive tension variations. The servo control of track-following technologies is limited in its ability to follow distorted tracks. This causes errors when attempting to restore data and when attempting to interchange media among tape drives.

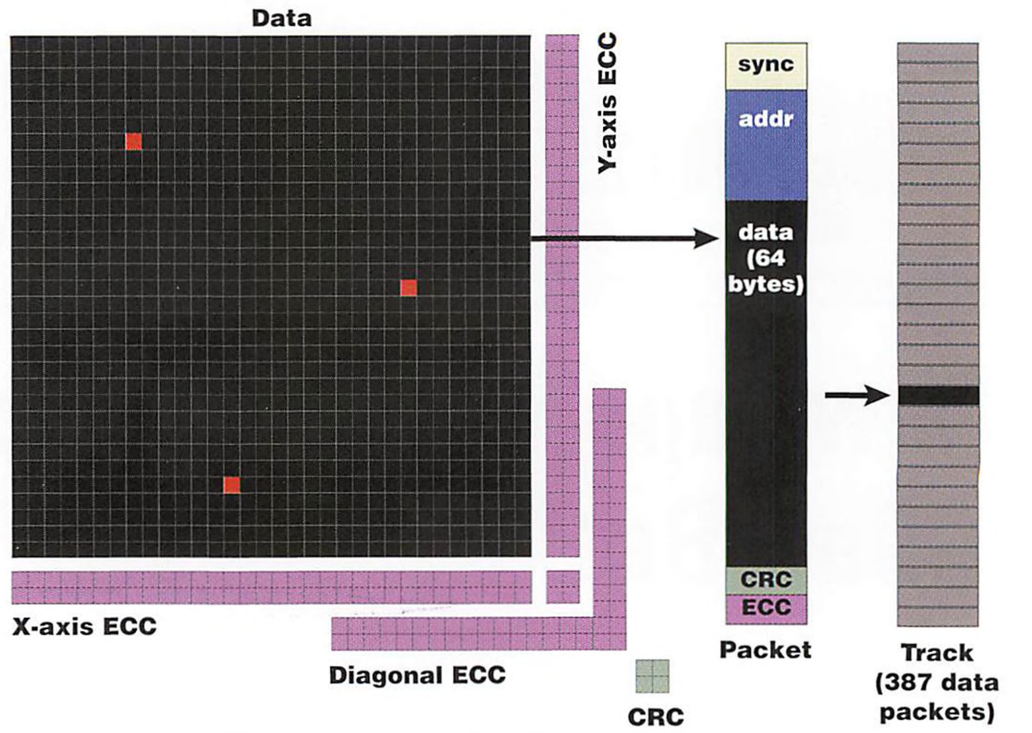

Data packets are read into the buffer segment and reassembled to produce a complete segment. Imagine the buffer as a checkerboard in which the black squares are "good" data and the red squares are "missing" data. ECC corrects the data until all the squares are black.

Performance Degradation

* **Data Transfer Rate** - Each back-hitching operation delays tape motion by 1–2 seconds. Multiple back-hitching events can reduce the effective data throughput to a fraction of nominal. Even more important, this degradation is data dependent and consequently impossible to optimize for the variety of jobs normally present in a typical computer environment.

Higher Product Cost

***Stringent Mechanical, Electrical and Manufacturing Tolerances** - The architecture of streaming tape drives grows more complex as data transfer rates and linear densities increase. This increased complexity comes at a price. There is a high dependence on costly state-of-the-art, high-precision electrical and mechanical components. Similarly, the complexity of the design and stringent assembly tolerances add to the cost of manufacturing, and increases the frequency of maintenance.

So What's The Solution?

Since the beginning, tape engineers have been adding complexity in order to increase recording density. The added complexity has finally resulted in diminishing returns in cost-per-megabyte.

We concluded two years ago that a new approach was needed to get tape back on the technology curve. This article describes a new approach to tape architecture called VXA. We believe VXA will right the ship of tape technology and allow it to sail into the next millennium, again confounding predictions of tape's imminent demise.

* **Breakthrough 1** - The most important conceptual breakthrough we achieved was the application of an idea borrowed from networking technology—namely to packetize the data on the tape. This has the effect of de-coupling the data structure from the physical track structure of the tape.

Your Internet service provider breaks up your long email message into small packets, adds addressing information, and sends them over the Internet. Some packets of your email may actually take different routes and arrive at their destination at different times and in a different order than your original message. At the email destination, the packets are reassembled in the right order, addressing is stripped off, and your recipient reads a nice long email, unaware of the whole process.

Similarly, VXA Technology takes a data file, breaks it up into packets, adds a unique address to each packet, and writes that packet to tape. The format is packet-based, and concerns about being unable to read an entire track are eliminated. Like the Internet mail message, which reassembled the data in the correct order, the VXA buffer architecture reassembles the packets in the correct order, regardless of the order in which it received them.

Discrete Packet Format™ (DPF)

As shown in Fig 2, a VXA packet consists of 64 bytes of user data preceded by sync bits and a 16bit local packet address. Following the data in each packet is CRC and a Reed-Solomon error correction code.

VXA implements a 4-layer Reed-Solomon error correction and corrects at both the packet level and the 32 x 32 packet buffer level. The error correction scheme means that there can be as many as two lost packets in each row, two in each column, and two in each diagonal, and VXA recovers the correct data. This results in an error rate equivalent to reading all the books in the Library of Congress 100 times with only a single error.

* **Breakthrough 2** - The next important conceptual breakthrough was variable speed tape. To understand what this means, picture the head-tape interface (Fig 3). It consists of a drum containing the read and write heads rotating at a fixed speed of 4,345rpm. The axis of the drum's rotation is perpendicular to the path of the tape. As the drum rotates, the heads move from the lower to the upper edge of the tape.

Now, consider what would happen if the tape were to slow

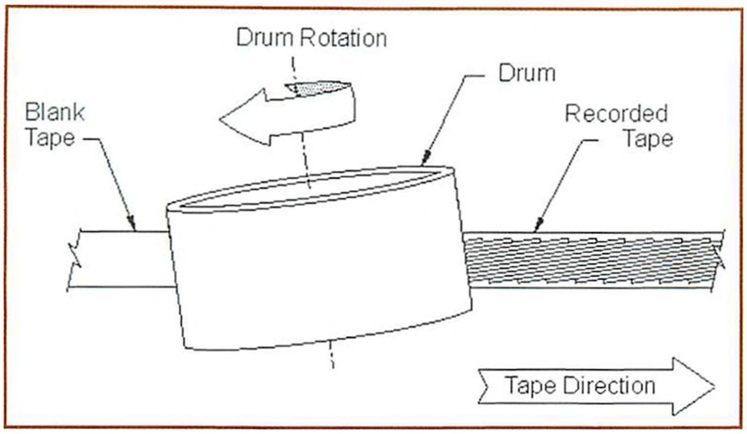

Fig 3 The head-tape interface consists of a drum containing the read and write heads rotating at a fixed speed of 4,345rpm. The axis of the drum's rotation is perpendicular to the path of the tape. As the drum rotates, the heads move from the lower to the upper edge of the tape.

down during a read operation. In a standard helical-scan streaming tape drive, a change of the tape speed would cause the track and head to be out of alignment and consequently the tape drive could no longer read the tape. In linear-track streaming tape drives, a reduction in tape speed would reduce the quality of the data signal and increase the error rate.

Under VXA Technology, we are not reading entire tracks—we are reading packets. Slowing the tape speed simply instructs the heads to make more than a single pass over a given track, which means VXA will read some packets more than once and out of order. But since packets have unique addresses, VXA doesn't care if it sees a packet more than once or that the last packets VXA read were from the beginning of the track. VXA only keeps one copy of each packet in the packet buffer. On writing tracks at a slower than nominal rate, the tape speed must move at some integer fraction of the nominal tape speed. At half speed, the heads will write on every other rotation of the drum.

With these two breakthroughs, we can eliminate all problems associated with helical and linear streaming tape drives and achieve our lower cost-per-megabyte goal as well. Following are some of the operational benefits of these two breakthroughs.

Variable Speed Operation™ (VSO)

VSO adjusts the tape speed to match the host transfer rate in real time, eliminating the problems common to all streaming drives' use of back-hitching. Unlike streaming tape technologies, VXA dynamically adjusts tape speed to instantaneously match the tape drive to the host's actual transfer rate. This speed-matching capability eliminates the delays and reliability problems associated with back-hitching. VXA can synchronize with slower or inconsistent hosts or bursty network transfer rates. Job times are reduced and time-to-data during restore operations is accelerated. In addition, variable speed for read and write modes significantly reduces the total amount of data buffering. Eliminating unnecessary and costly tape-path components reduces circuitry cost, further reducing the total cost of VXA drives and media.

OverScan Operation™ (OSO)

OSO is a technique for reading packets independent of track shape or the speed at which the tracks were recorded. The area scanned by all four heads is greater than the recorded area, ensuring that all packets are read at least once. This operation further optimizes data availability and enhances recovery possibilities for damaged or hard-to-read data.

VXA technology is the new, revolutionary approach to tape storage, one that offers profound advantages over conventional streaming devices. It is the first completely new approach to tape storage in more than a decade. It succeeds in a number of areas where streaming tape has struggled.

VXA is a mechanically simpler and more inherently reliable design, due to a reduced number of components. Also, because there is no back-hitching, less stress is placed on the mechanics and the media. And since it removes the need to place data in the precise sequence of time and space, it desensitizes tape guidance and improves tape wear, thereby increasing tape life and media interchangeability.

VXA's smaller packet size makes it more tolerant of variations occurring among drives and media. VXA's ability to vary tape speed and to overscan data increases its ability to recover hard-to-read data.

Matching the speed of the tape to the data rate of the host eliminates back-hitching delays, optimizing backup and restore operations. The combination of VXA's packet-based format, dynamic variable speed capability, and overscanning mode provides unprecedented data availability and media interchangeability.

VXA's revolutionary architecture means fewer mechanical and electronic components, as well as a reduced design burden from earlier tape guide and control technologies. VXA is inherently less costly to manufacture, and reestablishes tape's historical price/performance position compared to disk.

VXA eschews incremental improvements and establishes a truly new paradigm for tape storage.

Juan Rodriguez *is chairman and CEO of Ecrix Corporation (Boulder, CO).*

WestWorld Productions

Reprinted by permission from the publisher of the October 1998 edition of Computer Technology Review®.

For FREE subscription information, please call 310/276-9500, ext. 36 or reply via the World Wide Web http://www.wwpi.com.

VXA™

The new shape in tape™

Ecrix Corporation
5500 Central Avenue
Boulder, Colorado 80301

Telephone: 303.402.9262
Fax: 303.402.9266

Ecrix