MAGNETIC TAPE STORAGE

Magnetic tape has been used for **data storage** for over 50 years. In this time, many advances in tape formulation, packaging, and data density have been made. Modern magnetic tape is most commonly packaged in cartridges and cassettes. The device that performs actual writing or reading of data is a tape drive. Autoloaders and tape libraries are frequently used to automate cartridge handling.

When storing large amounts of data, tape can be substantially less expensive than disk or other data storage options. Tape storage has always been used with large computer systems. Modern usage is primarily as a high capacity medium for backups and archives. As of 2008, the highest capacity tape cartridges (Sun StorageTek T10000B, IBM TS1130) can store 1 TB of data without using compression.

Open reels

10.5 inch reel of 9 track tape

Initially, magnetic tape for data storage was wound on large (10.5 in/26.67 cm) reels. This defacto standard for large computer systems persisted through the late 1980s. Tape cartridges and cassettes were available as early as the mid 1970s and were frequently used with small computer systems. With the introduction of the IBM 3480 cartridge in 1984, large computer systems started to move away from open reel tapes and towards cartridges.

[edit] UNIVAC

Magnetic tape was first used to record computer data in 1951 on the Eckert-Mauchly UNIVAC I. The UNISERVO drive recording medium was a thin metal strip of ½″ wide(12.7 mm) nickel-plated phosphor bronze. Recording density was 128 characters per inch (198 micrometre/character) on eight tracks at a linear speed of 100 in/s (2.54 m/s), yielding a data rate of 12,800 characters per second. Of the eight tracks, six were data, one was a parity track, and one was a clock, or timing track. Making allowance for the empty space between tape blocks, the actual transfer rate was around 7,200 characters per second. A small reel of mylar tape provided separation from the metal tape and the read/write head.

[edit] IBM formats

IBM computers from the 1950s used ferrous-oxide coated tape similar to that used in audio recording. IBM's technology soon became the de facto industry standard. Magnetic tape dimensions were 0.5" (12.7 mm) wide and wound on removable reels of up to 10.5 inches (267 mm) in diameter. Different tape lengths were available with 1200', 2400' on mil and one half thickness being somewhat standard. Later during the '80s, longer tape lengths such as 3600' became available, but only with a much thinner PET film. Most tape drives could support a maximum reel size of 10.5"

Early IBM tape drives, such as the IBM 727 and IBM 729, were mechanically sophisticated floor-standing drives that used vacuum columns to buffer long u-shaped loops of tape. Between active control of powerful reel motors and vacuum control of these u-shaped tape loops, extremely rapid start and stop of the tape at the tape-to-head interface could be achieved. (1.5ms from stopped tape to full speed of up to 112.5 IPS) When active, the two tape reels thus fed tape into or pulled tape out of the vacuum columns, intermittently spinning in rapid, unsynchronized bursts resulting in visually-striking action. Stock shots of such vacuum-column tape drives in motion were widely used to represent "the computer" in movies and television.

Early half-inch tape had 7 parallel tracks of data along the length of the tape allowing, six-bit characters plus one bit of parity written across the tape. This was known as 7-track tape. With the introduction of the IBM System 360 mainframe, 9 track tapes were developed to support the new 8-bit characters that it used. Effective recording density increased over time. Common 7-track densities started at 200, then 556, and finally 800 cpi and 9-track tapes had densities of 800, 1600, and 6250 cpi. This translates into about 5 MB to 140 MB per standard length (2400 ft) reel of tape. At least partly due to the success of the S/360, 9-track tapes were widely used throughout the industry through the 1980s. End of file was designated by a tape mark and end of tape by two tape marks.

[edit] DEC format

LINCtape, and its derivative, DECtape, were variations on this "round tape." They were essentially a personal storage medium. The tape was ¾ inch wide and featured a fixed formatting track which, unlike standard tape, made it feasible to read and rewrite blocks repeatedly in place. LINCtapes and DECtapes had similar capacity and data transfer rate to the diskettes that displaced them, but their "seek times" were on the order of thirty seconds to a minute.

[edit] Cartridges and cassettes

Quarter-Inch cartridges.

In the context of magnetic tape, the term *cassette* usually refers to an enclosure that holds two reels with a single span of magnetic tape. The term *cartridge* is more generic, but frequently means a single reel of tape in a plastic enclosure.

The type of packaging is a large determinant of the load and unload times as well as the length of tape that can be held. A tape drive that uses a single reel cartridge has a take up reel in the drive while cassettes have the take up reel in the cassette. A tape drive (or "transport" or "deck") uses precisely-controlled motors to wind the tape from one reel to the other, passing a read/write head as it does.

A different type of tape cartridge has a continuous loop of tape wound on a special reel that allows tape to be withdrawn from the center of the reel and then wrapped up around the edge. This type is similar to a cassette in that there is no take-up reel inside the tape drive.

In the 1970s and 1980s, audio Compact Cassettes were frequently used as an inexpensive data storage system for home computers. Compact cassettes were logically, as well as physically, sequential; they had to be rewound and read from the start to load data. Early cartridges were available before personal computers had affordable disk drives, and could be used as random access devices, automatically winding and positioning the tape, albeit with access times of many seconds.

Most modern magnetic tape systems use reels that are fixed inside a cartridge to protect the tape and facilitate handling. Modern cartridge formats include DAT/DDS, DLT and LTO with capacities in the tens to hundreds of gigabytes.

[edit] Technical details

[edit] Tape width

Medium width is the primary classification criterion for tape technologies. Half inch has historically been the most common width of tape for high capacity data storage. Many other sizes exist and most were developed to either have smaller packaging or higher capacity.

[edit] Recording method

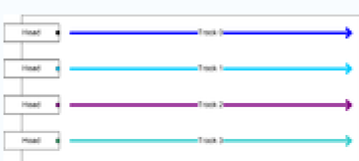

Linear

Recording method is also an important way to classify tape technologies, generally falling into two categories:

[edit] Linear

The **linear** method arranges data in long parallel tracks that span the length of the tape. Multiple tape heads simultaneously write parallel tape tracks on a single medium. This method was used in early tape drives. It is the simplest recording method, but has the lowest data density.

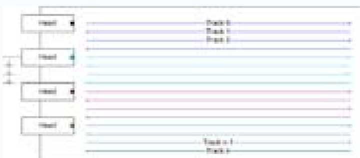

Linear serpentine

A variation on linear technology is **linear serpentine** recording, which uses more tracks than tape heads. Each head still writes one track at a time. After making a pass over the whole length of the tape, all heads shift slightly and make another pass in the reverse direction, writing another set of tracks. This procedure is repeated until all tracks have been read or written. By using the linear serpentine method, the tape medium can have many more tracks than read/write heads. Compared to simple linear recording, using the same tape length and the same number of heads, the data storage capacity is substantially higher.

[edit] Scanning

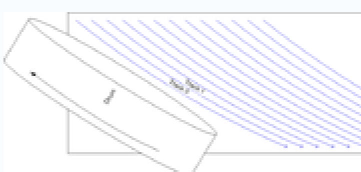

Helical

Scanning recording methods write short dense tracks across the width of the tape medium, not along the length. Tape heads are placed on a drum or disk which rapidly rotates while the relatively slowly moving tape passes it.

An early method used to get a higher data rate than the prevailing linear method was **transverse scan**. In this method a spinning disk, with the tape heads embedded in the outer edge, is placed perpendicular to the path of the tape. This method is used in Ampex's DCRsi instrumentation data recorders and the old 2 inch Quadruplex videotape system. Another early method was **arcuate scan**. In this method, the heads are on the face of a spinning disk which is laid flat against the tape. The path of the tape heads makes an arc.

Helical scan recording writes short dense tracks in *diagonal* manner. This recording method is used by virtually all videotape systems and several data tape formats.

[edit] Block layout

In a typical format, data is written to tape in blocks with inter-block gaps between them, and each block is written in a single operation with the tape running continuously during the write. However, since the rate at which data is written or read to the tape drive is not deterministic, a tape drive usually has to cope with a difference between the rate at which data goes on and off the tape and the rate at which data is supplied or demanded by its host.

Various methods have been used alone and in combination to cope with this difference. The tape drive can be stopped, backed up, and restarted (known as shoe-shining, because of increased wear of both medium and head). A large memory buffer can be used to queue the data. The host can assist this process by choosing appropriate block sizes to send to the tape drive. There is a complex tradeoff between block size, the size of the data buffer in the record/playback deck, the percentage of tape lost on inter-block gaps, and read/write throughput.

Finally modern tape drives offer speed matching feature, where drive can dynamically decrease physical tape speed as much as 50% to avoid shoe-shining.

[edit] Sequential access to data

From user perspective the primary difference between tape data storage and disk data storage is that tape is a sequential access medium while disk is a random access medium. Hence tape uses a very trivial filesystem in which files are addressed by number not by filename. Metadata such as file name or modification time is typically not stored at all. Over time some tools (i.e. tar) were introduced to enable storing metadata by introducing richer formats of packing multiple files in a single large 'tape file'.

Another difference to hard disk storage is that data is generally added by appending a file to the end of the recording, not by overwriting a particular file (or part of file) in the middle of tape.

[edit] Access time

Tape has quite a long latency for random accesses since the deck must wind an average of one-third the tape length to move from one arbitrary data block to another. Most tape systems attempt to alleviate the intrinsic long latency, either using indexing, where a separate lookup table (*tape directory*) is maintained which gives the physical tape location for a given data block number (a must for serpentine drives), or by marking blocks with a tape mark that can be detected while winding the tape at high speed.

[edit] Data compression

Most tape drives now include some kind of data compression. There are several algorithms which provide similar results: LZ (most), IDRC (Exabyte), ALDC (IBM, QIC) and DLZ1 (DLT). Embedded in tape drive hardware, these compress a relatively small buffer of data at a time, so cannot achieve extremely high compression even of highly redundant data. A ratio of 2:1 is typical, with some vendors claiming 2.6:1 or 3:1. The ratio actually obtained with real data is often less than the stated figure; the compression ratio cannot be relied upon when specifying the capacity of equipment, e.g., a drive claiming a compressed capacity of 500GB may not be adequate to back up 500GB of real data. Software compression can achieve much better results with sparse data, but uses the host computer's processor, and can slow the backup if it is unable to compress as fast as the data is written.

Some enterprise tape drives can encrypt data (this must be done after compression, as encrypted data cannot be compressed effectively). Symmetric streaming encryption algorithms are also implemented to provide high performance.

The compression algorithms used in low-end products are not the most effective known today, and better results can usually be obtained by turning off hardware compression, using software compression (and encryption if desired) instead.

[edit] Viability

Tape remains a viable alternative to disk due to its higher bit density and lower cost per bit.[*citation needed*] Tape has historically offered enough advantage in these two areas above disk storage to make it a viable product, particularly for backup. Tapes are more robust and often smaller than disk drives, and suitable for on- and off-site storage of several generations of backups.[1] The rapid improvement in disk storage density and price, coupled with arguably less-vigorous innovation in tape storage, has reduced the market share of tape storage.