Chapter 11. Magnetic Tape Drives.

Magnetic tape storage is often the forgotten or ignored peripheral device until the first disk crash or virus infection. In the early days, when disk storage was small and expensive, magnetic tapes were used for data processing as well as storage. A typical process would include reading data in from a set of punched cards. These input records are sorted and written out to tape. Often the sorting process requires the use of two tape drives to perform the sort merge operations as internal memory could be quite limited. The sorted input data tape is mounted on one drive, the "old" master tape on the next drive, and the updated output master on a third drive.

Today the main applications are for archive and backup, where the cost as well as volume per megabyte is still much lower than any other competing medium.

11.1 Classification and Application of Tape Drives

There are several ways of classifying magnetic tape storage devices. One way is to classify them by the mode of operation, viz. block mode or streaming mode tape drives. Another way is to divide them according to whether the recording is done by a fixed or a rotating head. The chart in Figure 0-1 gives a finer classification.

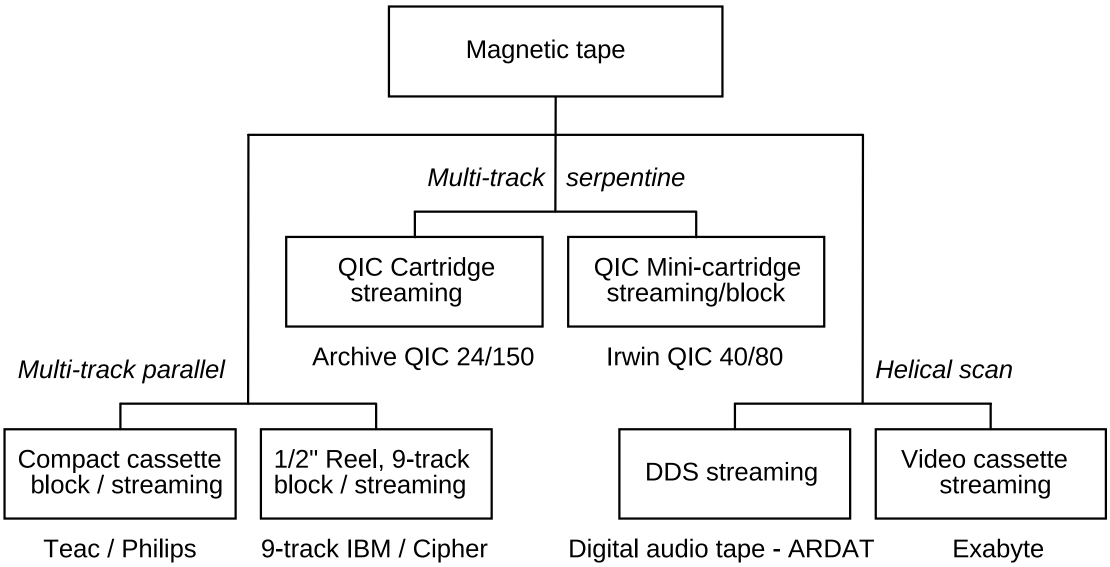

Figure 0-1. Classification of magnetic tape drives.

11.1.1 Start-Stop Blocked Mode

In the early years of data processing using computers, tape drives were used for the storage of data files on a record-by-record basis. For updating a record or block of data is first read into the CPU, calculations performed, and an updated record is written back to the tape. Provided the block length remains the same, the drive would rewind the tape by one block length and the updated block written back over the old record.

In this blocked mode of operation, data are recorded in fixed physical blocks of 512 to 4096 bytes. Between each block there is an Inter-Record Gap (IRG) of up to 0.5 inch. The magnetic tape controller uses the IRG to locate the beginning of each block and can update or over-write data on the tape on a block by block basis.

Obviously blocked mode operation places a higher demand on the tape transport mechanism. Tape movement now consists a sequences forward, stop, reverse, stop operations in quick succession. In addition, it has the capability of searching for a specific block by skipping in the forward or reverse directions.

Although blocked mode operation can be implemented on any kind of tape drives, it is usually found only in the large 1/2-inch fixed head multi-track drives and s few mini-cartridge units. In the case of the mini-cartridge units, additional software is used to format the tape into "tracks" and "sectors" so that it emulates a floppy disk drive with a capacity of 20/40 Mbytes.

Table 1 gives the typical specifications of some popular IBM blocked mode 1/2-inch tape drives.

Table 1. Typical specifications of IBM reel-to-reel tape drives.

IBM Product No.	726	3420	3480
FCS (First customer shipment)	1953	1973	1985
Linear Density (BPI)	100	6250	38,000
Number of Tracks	7	9	18
Reel Capacity (MB)	2.2	156	200
Data Rate (KBytes/sec)	75	1250	3000
Recording Code	NRZI	GCR(0,2)	GCR(0,3)
Tape Transport	Vacuum	Vacuum	Cartridge

Figure 0-2. Half-inch vacuum column magnetic tape transport.shows the mechanism of a vacuum column tape drive. Because of the start-stop operations, keeping the tape well tensioned across the read/write head is non-trivial. Lower cost units, running at low speeds of 45 ips, use spring-loaded arms, but for the higher performance units, a vacuum column is used. The lower pressure within the column draws in excess slack in the tape gently, without causing damage by stretching, or physical contact.

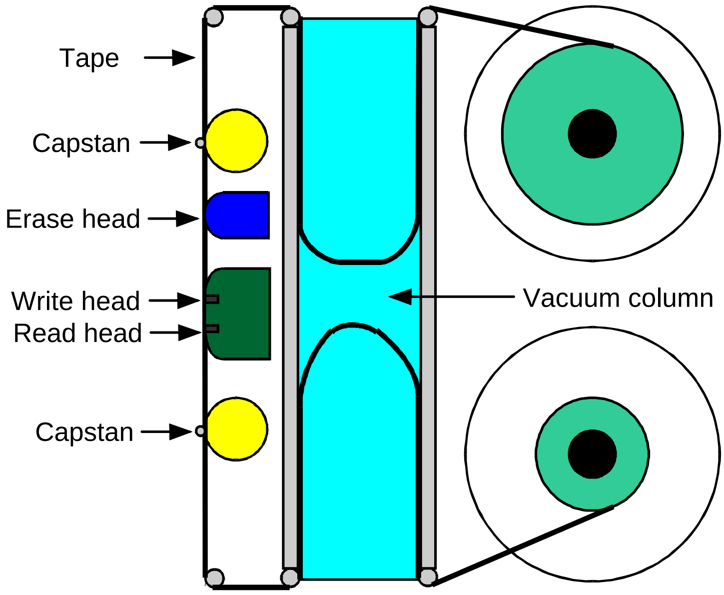

Figure 0-2. Half-inch vacuum column magnetic tape transport.

11.1.2 Streaming Mode

When the primary application of tape drives is for the archiving and backup, the need to operate in the start-stop mode is no longer required. In such applications, a file or set of files are written onto the tape. If these files have changed, the updated versions are normally kept by re-writing the complete files. No attempt is made to update just the affected records within the file. In this case streaming mode tape drives can be used. Streaming tapes write complete files or groups of files as a continuous "stream" of data. Although the actual data may be recorded in physical blocks, there is no support for locating and modifying a particular block on the tape and leaving the portions before and after unaffected.

Tape motion is controlled only in the forward direction for read, write or skip forward to file mark or end of file. Reverse motion is used only to rewind the tape back to the beginning. We will consider in turn the two current tape drive mechanisms, the fixed head and rotary head machines, both operating in the streaming mode.

11.2 Stationary Head Recorders

The ubiquitous walkman is an audio tape recorder with a stationary head. Early digital tape drives were adaptations of audio tape recorders. In fact normal audio recorders have been used for storing digital data by using the digital information to modulate tone signals in the audio frequency range. FSK, PSK and other modulation schemes have been used. However, digital tape drives normally have different designs to optimize data storage density and transfer. In order to achieve high storage efficiency, direct digital recording is used. The tape is moved rapidly across the head as the data transfer rate is a direct function of tape speed. Digital tape drives run the tape at over 45 ips compared the several inches per second for audio units.

As the recording track runs along the length of the tape, several tracks are laid out in parallel. On the multi-track recorders, these tracks are recorded simultaneously using a head block with a number of recording heads mounted in parallel.

11.2.1 The QIC Standard

The QIC (Quarter-Inch Cartridge) industry standard describes a stationary head, streaming mode tape storage which is becoming increasingly popular because of its compact size and large capacity. However it has only two sets of read/write heads, one for each direction of tape movement. Parallel tracks are obtained by physically moving the head perpendicularly across the width of the tape to access a new track. Each track is written sequentially resulting in a *serpentine* pattern as can be seen in Figure 12.1.

QIC defines a number of standards, differing in capacity, as can be seen in Table 2.

Table 2. QIC tape standards.

	QIC-24	QIC-150	QIC-525	QIC-1350
Capacity (formatted) MB	45 or 60	125 or 150	320 or 525	1.35 GB
Track Format	9	18	26	30
Flux Density	10,000 ftpi	12,500 ftpi	20,000 ftpi	38,750 ftpi
Data Density	8,000 bpi	10,000 bpi	16,000 ftpi	51,667 bpi
Tape Speed	90 ips	90 ips	120 ips	120 ips
Data Transfer Rate KBytes/Sec	90	112.5	240	600
Recording Code	GCR (0,2)	GCR (0,2)	GCR (0,2)	RLL(1,7)
Track Width (in)	0.0135	0.0056	0.0070	0.0070
Tape Length (ft)	450 or 600	600	600 or 1000	750
Soft Error Rate	1 in 10^8	1 in 10^8	1 in 10^8	1 in 10^8
Hard Error Rate	1 in 10^{10}	1 in 10^{10}	1 in 10^{10}	1 in 10^{10}

11.2.2 QIC-24 Cartridge Tape Drive

We shall use QIC-24 standard to look at the various aspects of the QIC tape drive. Referring again to Figure 0-3 the read/write operation is explained in the following section.

11.2.2.1 Read/Write Operation

1. When the tape is loaded, rewinding takes place. Tape is then forwarded to Beginning of Tape (BOT), and head is positioned for Track 0.

2. On WRITE command, tape is forwarded to the Load Point (LP). This is the beginning of the recording zone and drive begins writing to tape on Track 0. Data is written bit serial on one track. In addition, during the writing on Track 0, the erase head is enabled. Full track erase is performed by applying AC signal to the erase head ahead of the write head.

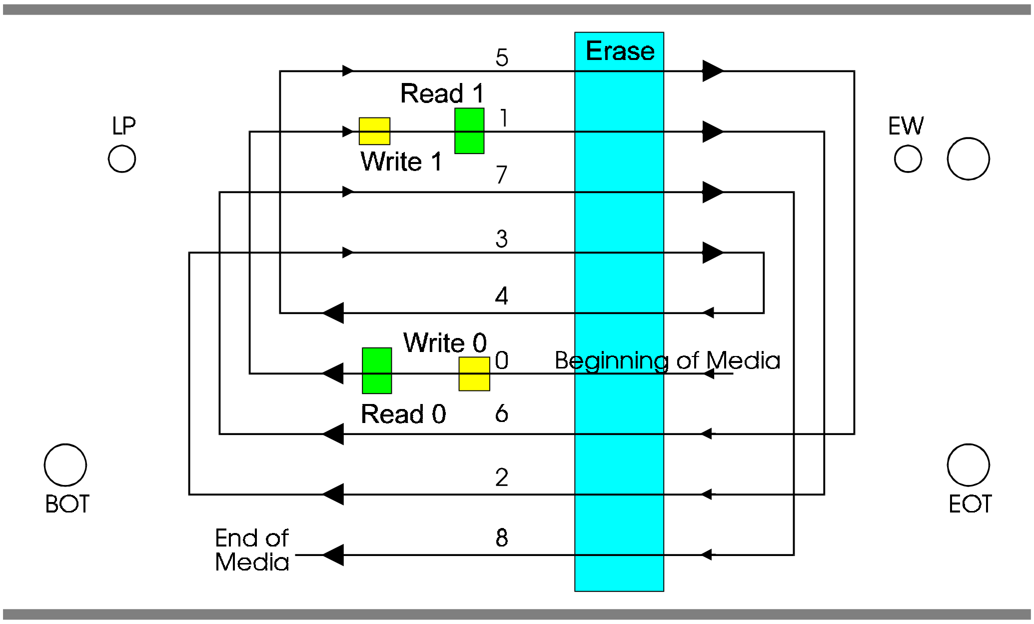

Figure 0-3. Serpentine tracks in the QIC standard.

3. Writing on Track 0 continues till the Early Warning (EW) hole is detected. Drive stops accepting data from host. Recording stops when all remaining data in buffers are written to tape.

4. Tape movement continues in the forward direction until End of Tape (EOT) is reached. Erase head is disabled. Read / Write head for Track 1 in reverse direction is selected.

5. Tape is now moved in the reverse direction and recording begins when EW hole detected. Recording continues until the LP hole is detected.

6. When end of Track 1 is reached at BOT, tape head is moved (down 0.048 inch from Track 0 to Track 2), and positioned at the next track. Head is moved across width of tape to read/write the various data tracks by attaching it to a screw driven by a stepper motor, with resolution of 0.001 inch.

7. Tape movement is changed back to the forward direction and Track 2 is written.

8. Similar procedures are used, alternating between forward and reverse directions and shifting the head after every pair (F/R) of tracks.

9. The serpentine recording pattern results.

10. The READ Operation is similar.

11.2.2.2 12.2.1.2. Data Block Format

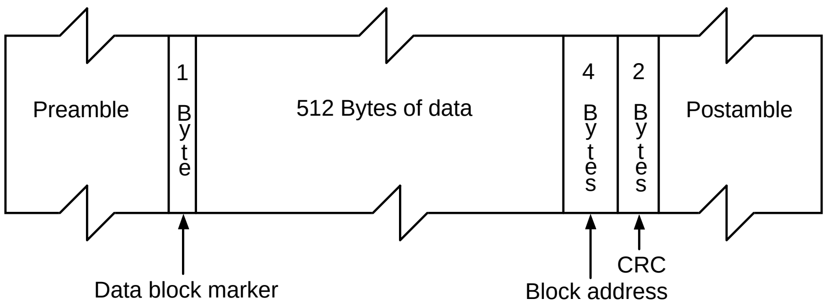

Figure 0-4. QIC data block format.

Figure 0-4 shows the format of the data block used in the QIC-24 tape drives. The various components comprising the block are described below:

a. Preamble

This is a sequence of bytes used to synchronize the PLL to the data frequency. An underrun occurs when during the write process, the data rate from the CPU is slower than the write data rate, so that the write buffer in the tape is empty. Rather than stopping the tape, preamble bytes are written into this area while waiting for the write buffer to fill. A long preamble is used at beginning of each track also.

b. Data Block Marker

GCR code '11111 00111'

c. Data Block

The 512 data bytes encoded as GCR 4/5.

d. Block Address

The block address consists of a 4-Byte block address with track number, a control nibble and block number starting with 1 on each track. This is illustrated in Figure 12.2.

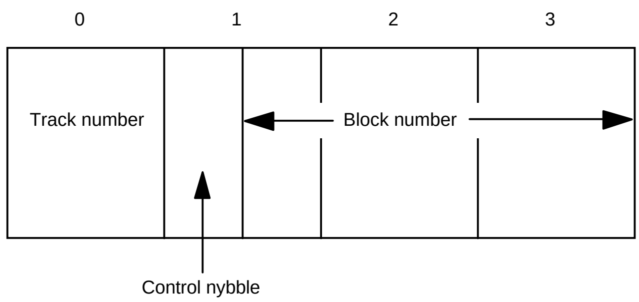

Figure 0-5. QIC data block format.

e. Cyclic Redundancy Check Character

A 2-Byte CRC character is calculated over the 512 bytes of data and 4 bytes of block address using the following polynomial:

$$CRC = x^{16} + x^{12} + x^{5} + 1$$

f. Postamble

A normal post amble of 5 - 20 flux transitions are recorded at maximum density to act as a guard band. Elongated post ambles of 3500 - 7000 flux transitions are appended when an underrun occurs.

11.2.2.3 Underrun

If data from the host are interrupted, or if data transfers from the host drop below 90 Kbytes/sec, underrun occurs and duplicate blocks maybe recorded. These duplicate blocks are transparent to the host.

11.2.2.4 Reliability

The QIC standard has error processing and recovery software incorporated in the firmware. Their operation is transparent to host and user. For each block of 512 bytes of data plus control data, a cyclic redundancy check (CRC) character is computed and recorded. The CRC generating polynomial used is:

$$x^{16} + x^{12} + x^{5} + 1$$

A read check performed on each block of data as it is written. The read head is positioned 0.3 in. as shown in Figure 12.3 after the write head and is used for normal reading and for the read-after-write checking. Based on the tape speed, this distance is equivalent to a delay of approximately 300 bytes.

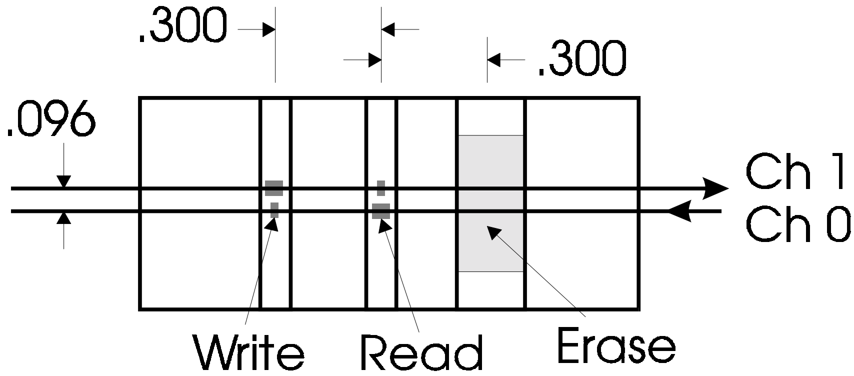

Figure 0-6. QIC-24 tape and head layout.

Thus when the controller is reading and checking block N, the next block, N+1 is already being written to tape. The following, together with Figure 0-7 describes the protocol used in the error processing and recovery procedure:

1. Drive begins writing block N (about 530 bytes long).
2. When block N reaches read head (300 bytes later), read checking begins.
3. Writing of block N completes with Block address and CRC.
4. After inter-block gap, block N+1 is written.
5. Read checking finds CRC error in block N.

6. Block N must be re-written, but block N+1 is already half-written.
7. Drive completes writing block N+1 and begins writing block N for second time. A second copy of block N+1 is also written.
8. If there are no errors after second writing of block N during the read check, the tape drive continues with writing block N+1 and the normal sequence resumes.
9. If and error is found at the second write attempt, a third copy of block N (and N+1) is written. This repeated up to a total of 16 attempts.
10. After 16 unsuccessful attempts, the write operation is abandoned.
11. In the READ mode, the last copy of each block is taken as the valid copy.

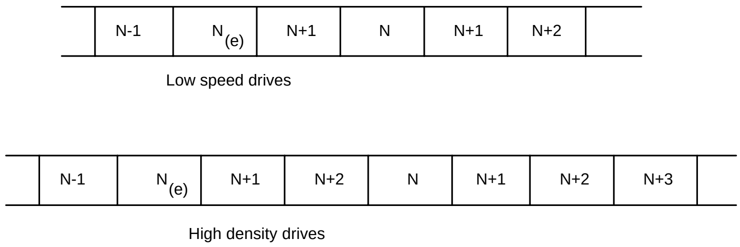

Figure 0-7. QIC data block format.

11.2.3 QIC-02 Standard Electrical Interface

Table 3. The QIC-02 standard command set.

Op Code	Command	Op Code	Command
01	Select Drive, Soft Lock OFF	80	Read
11	Select Drive, Soft Lock ON	81	Space Forward
21	Position to Beginning of Tape	89	Space Reverse
22	Erase Entire Tape	A0	Read File Mark
24	Retension Cartridge	A3	Seek End of Data
40	Write	Bn	Read n File Marks
41	Write without Underrun	C0	Read Status
60	Write file Mark (FM)	C2	Run Self Test 1
80	Read	CA	Run Self Test 2

In addition to the various recording capacity/format standards QIC also defined a standard electrical interface and a set of commands, designated QIC-02. These are shown in Table 3 and Figure 0-8 respectively. This means the same tape controller and device driver is usable across the whole range of tape drives whatever the capacity. The only change that may be required is to modify the capacity parameter so that the software can determine when the next volume of tape needs to be mounted.

11.2.4 QIC-36 Standard Drive Interface

A low level electrical interface for the tape drive has also been defined as the QIC-36 standard. This is basically the electrical and electromechanical interface and does not include any intelligence in the tape drive electronics. Essentially this specifies the interface between the read/write head, the motors and sensors on the basic tape drive and the drive controller.

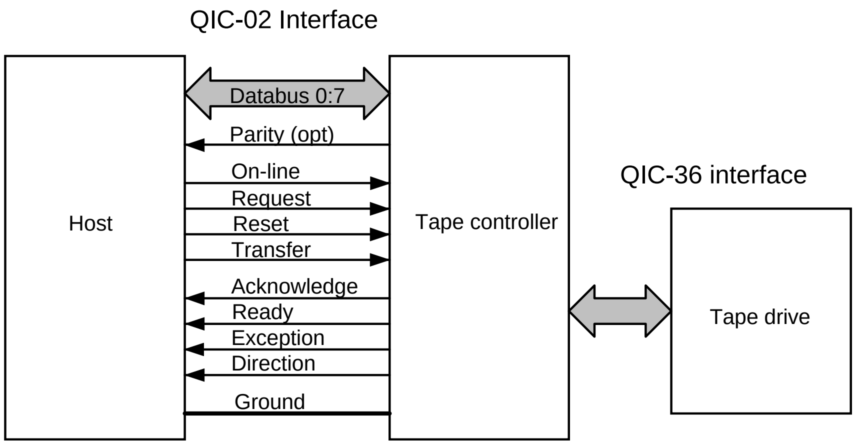

Figure 0-8. The QIC-02 standard electrical interface.

11.2.5 Tape Drive Electronics

Figure 12.4 shows the functional block diagram of a typical QIC tape drive with a SCSI interface. The 8-bit microprocessor together with the custom VLSI, provides complete control of the drive. Memory consists of 16 Kbytes of programme in Eprom and 64 Kbytes of dynamic RAM for buffering. The standard 5080 SCSI controller device operates under the control of the microprocessor, negotiating handshake protocols, commands and data transfer with the host computer. The VLSI brings together a number of functions in one reliable package by reducing component count. These are:

- DMA Controller, to handle data transfer from the host through the SCSI controller.
- Interrupt Controller
- Memory Access Controller, provides multiplexing of memory and address buses, dynamic RAM refresh and parity checking.
- Read/Write Controller, in the write mode performs the read-after-write function to verify correct writing of data.
- Clock Generator, provides the various clocks for the microprocessor, SCSI interface and read/write phase clocks.
- Data Separator; when the data is read from tape, it is a mixture of data and clock signals. The read circuit conditions and detects these signals which are then passed to the VLSI to be separated, making use of the PLL.

During the write operation, the erase circuit is active during the write of track 0, when a complete erase is performed on the whole tape. The write circuit provides accurate control of the write current to the head to cause the flux transitions on the tape.

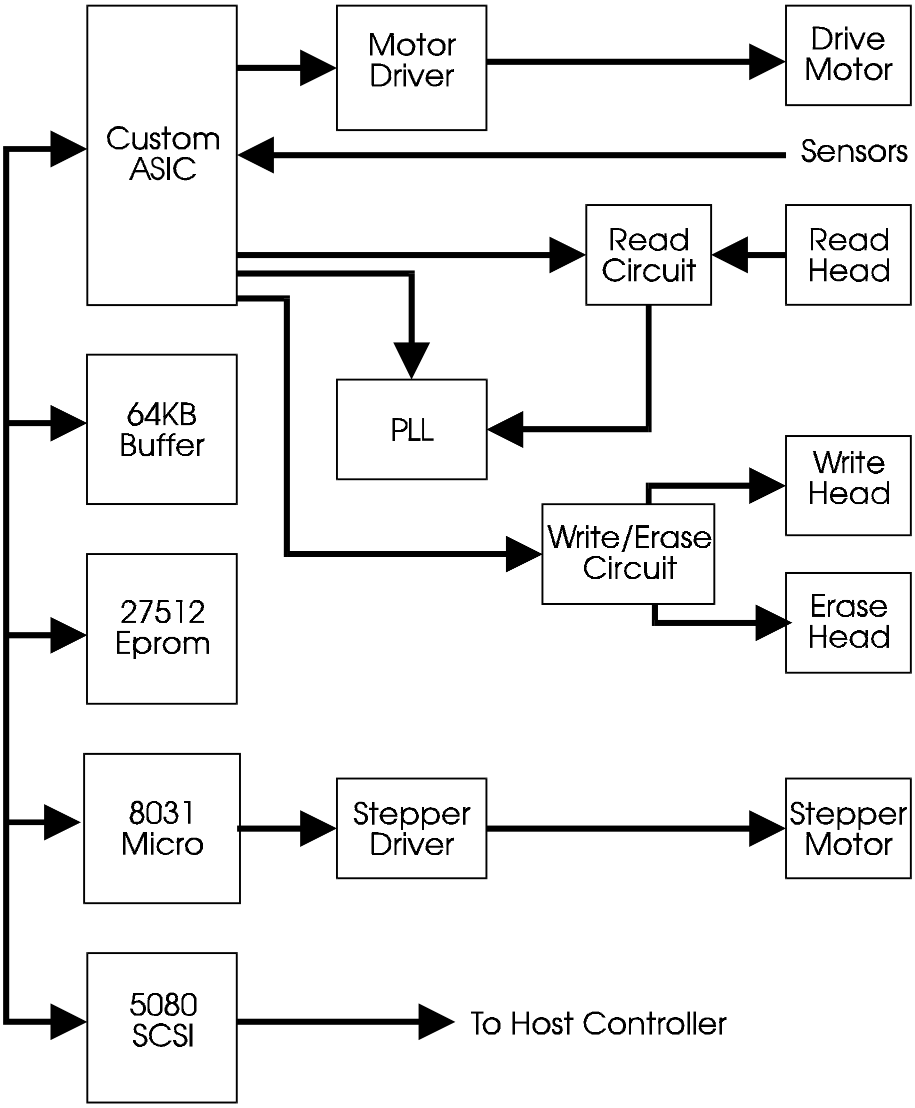

Figure 0-9. Functional block diagram of QIC drive.

11.2.6 Conclusions

Disk capacities are increasing rapidly; and backup devices with sufficient capacity are needed to keep pace with this increase. Current Winchester disk capacity is about 1.3 GB, reaching out to 2 GB. The two competing candidates for backup devices, are the cartridge tape and optical disks (WORM). Presently tape drives have the price advantage, and the QIC-1350 tape, with a capacity of 1.35 GB has been designed for this application. Development continues to continuing increasing this capacity.

11.3 Rotary Head Recorders

Whilst the stationary head recorders are similar on the common audio recorders, the rotary head recorders works on the same principles as the video cassette recorders. The two main storage devices in this group are based on the video cassette recorders and the newer rotary-head digital audio tape (RDAT). With the multi-track stationary head recorders, tape utilisation is low as there has to be sufficient separation between the recording tracks to prevent cross-talk interference. Also for high data rates, the tape has to be moved rapidly across the heads.

Rotary head recorders overcome this two problems. Instead of moving the tape, the head is rotating at a high speed across the tape. The tape itself moves slowly at a rate sufficient to feed a new portion to the head as it rotates. In fact, the tracks are allowed to overlap, and cross-track interference is reduced by offsetting the azimuths of the multiple heads.

The Exabyte EXB-8200 drive uses the video cassette format with 8 mm tape. It has a rated storage capacity of 2.5 GB and a data transfer rate of about 246 Kbyte/sec. There are a number of vendors using the DAT cartridge format with 60 m of 4 mm tape. These have a capacity of 1.3 GB and a sustained data transfer rate of 183 Kbyte/sec. Our discussion will focus on the DAT devices.

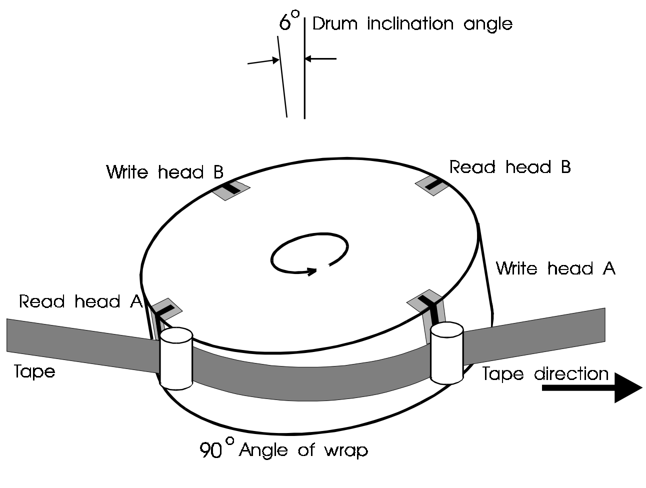

Figure 0-10. Helical scan rotary-head recorder.

Figure 0-10 show the general arrangement of the rotary-head recorder. Four heads are located on the rotating drum, setup as two pairs of write/read heads. Differing from the standard VCR, the tape is only wrapped around a 90° quadrant of the drum, which is rotating at 2000 rpm. The tape itself moves forward at about 1 inch in 3 seconds. It will be seen that each head traces out a diagonal track on the tape. By proper positioning of the read head relative to the associated write head, a read-after write operation can be performed on the same track. The two sets of head and the tape speed are set up in such a way that each set of heads writes adjacent tracks onto the tape. The ANSI Digital Data Storage (DDS) standard format has been proposed for the use of DAT for data storage. Typical specifications of a currently available DAT product is given in Table 4.

Table 4. Typical specifications of a DDS DAT drive.

Product:	Archive Python 4330XT
Capacity:	1.3 GBytes with 60m tape.
Sustained transfer rate:	183 Kbytes/sec, sustained.
Average access time:	20 sec. seek time.
Small form factor:	3 1/2"
Standard recording format:	ANSI DDS
Low cost:	Currently US$0.01 / Mbyte.
Interface:	SCSI-1 and SCSI-2
Media	4 mm. DAT Cartridge, 60/90m length.
Packing density	1869 tracks/in.
Areal density	114 Mbits/sq. in.
Uncorrectable error rate	Using ECC, 1 in 10^{15} bits.
Drum rotation speed:	2000 RPM
Tape speed:	0.32 in/sec.
Search/rewind speed	200 X normal speed
Head-to-tape speed:	123 in/sec, Helical scan (RDAT)

11.3.1 DDS DAT Format

DDS DAT is the **D**igital **D**ata **S**torage format for **D**igital **A**udio **T**ape, a standard format definition which uses the DAT in a streaming mode. There is a second format defined by Hitachi, Data/DAT, which supports random access. However Data/Dat has a lower data transfer rate and less storage capacity.

Referring to Figure 0-11 it is observed that areas of the tape is allocated for device, reference and vendor information. These are device and vendor specific. There is usually one or two data areas and only one is shown here.

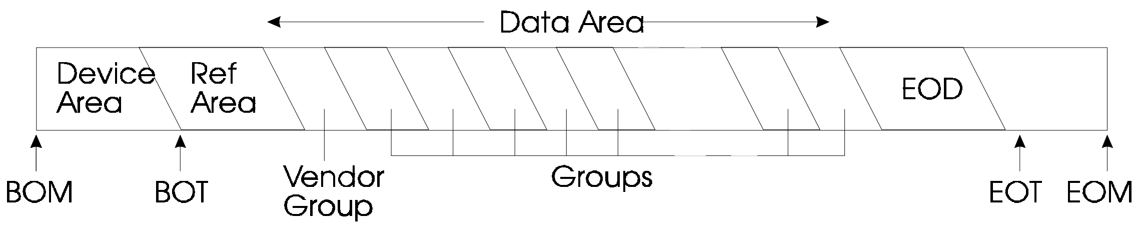

Figure 0-11. Layout of DDS DAT with one data area.

Data are organised into groups, with each group containing about 128 Kbytes divided into 22 data frames as shown on Figure 0-12. DDS DAT devices usually has a large buffer so that N groups of data are written at a time.

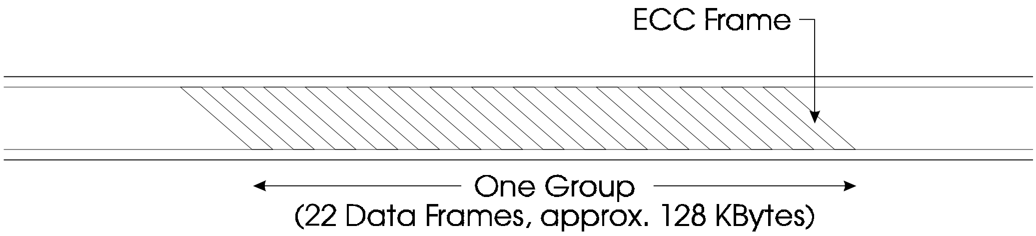

Figure 0-12. One DDS DAT data group

Examination of Figure 0-13 shows that each frame is made up of two tracks; track A is written/read by head A and track B by head B respectively. It can be noted that there is no inter-track gap, and each track overlaps the previous track slightly. At the overlap, signals from one track will interfere with the signal from the other. To eliminate this source of cross-talk interference, the heads are given a small azimuth offset of 20° from the normal. In this way, the component of noise from the interfering track is strongly attenuated as there is a total disorientation of 40° between the write head A and read head B.

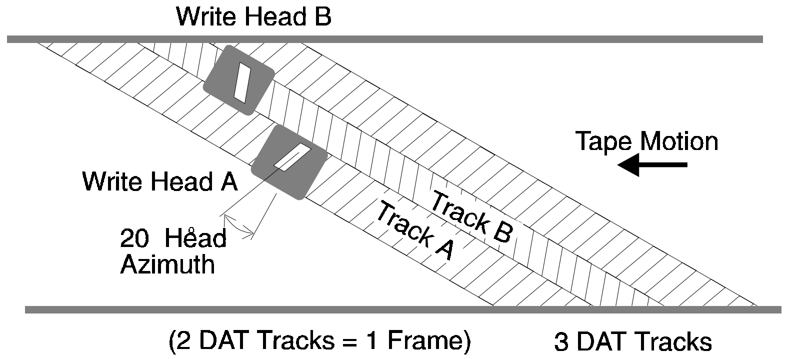

Figure 0-13. Alternate azimuth angles.

Figure 0-14 illustrates the format used in each DDS track. The 8/10 code is used and additional processing like interleaving and Reed-Solomon redundancy is used for error correction. Tone bursts are recorded onto the ATF section. On read-back these signals are fed into the track following circuits. In operation, when head B is following track B, the pilot tones from the two adjacent A tracks are sensed and the amplitude difference is use to more accuracy position head B over the track.

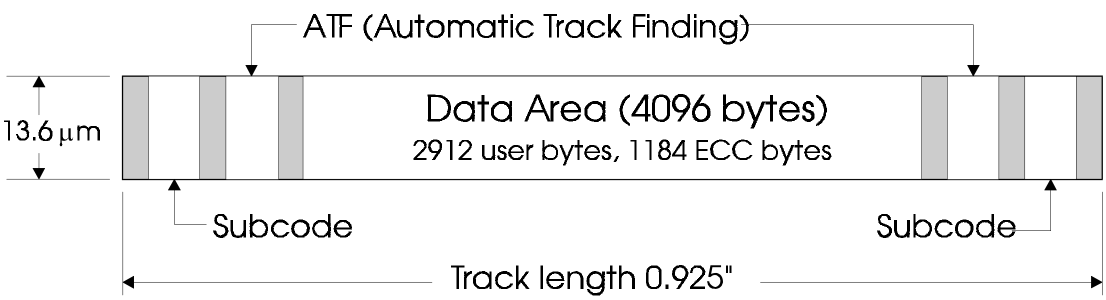

Figure 0-14. DDS DAT track format.

11.3.2 Conclusions

DDS DAT, because of it compactness and high storage capability is rapidly gaining popularity as a backup storage device. Some of the latest products use tape lengths of 90 m. in a cartridge. Data compression is built into the drive controller resulting in a device that can store up to 8 GB per cartridge. If that is still insufficient, there are "juke boxes" with autoloaders that can manage a library of cartridges.

11.4 Additional Reading Guide

John Watkinson, *RDAT*, Focal Press, 1991