InterPARES Project

International Research on Permanent Authentic Records in Electronic Systems

Media

Draft Appendix

October, 2001

P C Hariharan

Media

Historical evidence for written records dates from about the middle of the third millennium BC. The writing is on media[1] like a rock face, cave wall, clay tablets, papyrus scrolls and metallic discs. Writing, which was at first logographic, went through various stages such as ideography, polyphonic syllabary, monophonic syllabary and the very condensed alphabetic systems used by the major European languages today. The choice of the medium on which the writing was done has played a significant part in the development of writing. Thus, the Egyptians used hieroglyphic symbols for monumental and epigraphic writing, but began to adopt the slightly different *hieratic* form of it on papyri where it coexisted with hieroglyphics. Later, *demotic* was derived from hieratic for more popular uses. In writing systems based on the Greek and Roman alphabet, monumental writing made minimal use of *uncials* and there was often no space between words; a soft surface, and a stylus one does not have to hammer on, are conducive to cursive writing.

Early scribes did not have a wide choice of media or writing instruments. Charcoal, pigments derived from mineral ores, awls and chisels have all been used on hard media. Cuneiform writing on clay tablets, and Egyptian hieroglyphic and hieratic writing on papyrus scrolls, permitted the use of a stylus made from reeds. These could be shaped and kept in writing trim by the scribe, and the knowledge and skill needed for their use was a cherished skill often as valuable as the knowledge of writing itself.

Historically, a characteristic of human writing, which distinguishes it from *electronic* recording, is that it involved the transfer of *mass*, either from the writing surface (as in etching and carving), or to it, in the form of ink, from a pen or brush. The writing is perceived as a difference in color, or as a difference in texture, of the written portions of the surface. Both in Rome and Greece, the wax tablet, which was a piece of board with a very shallow depression covering almost its entire area and coated with a thin layer of beeswax, was used for making non-permanent records, and this may be viewed as the predecessor to rewritable media. Edison experimented with recording sound on a wax-coated cylinder, but settled on tinfoil instead. Today, the most widely used distribution medium for audio and computer software, the Compact Disc (CD) and the CD-ROM (Read-Only Memory), are replicated using a technique which is based on the removal of mass from the media; information is recorded on a CD in the form of pits, just as they were on the 70 rpm records, and its successors, the LP records. However, the pit diameter, the track pitch and the pit depth are microscopic, and the fidelity with which millions of copies of such an intricate and complex pattern can be faithfully reproduced is a tribute to modern technology. In the not too distant future, it is likely that there will be further reductions of two or three orders of magnitude in the size of the features or domains, which store the information. The Rosetta HD-ROM (High Density Read-Only Memory) from Norsam Technologies of Santa Fe NM relies on milling using an ion beam, and can reproduce features of the order of 25 nm on a number of substrates, from plastic polymers to silicon to metal. With the exception of CD, DVD, CD-ROM, DVD-ROM and HD-ROM, the electronic recording methods use an *energy transfer* process for information transfer rather than the traditional mass transfer process.

Analog and digital recording in use today are based on either permanent magnetism (*ferromagnetism*), or the use of a laser (*optical recording*) to etch away material, or cause a chemical or physical change. Magnetic media are available as rigid disks, floppy disks, and flexible tape, of various widths, wound on an open reel, or enclosed in a cartridge

[1] Since the 1970's, the custom has been to use the plural *media*, to stand for both singular and plural.

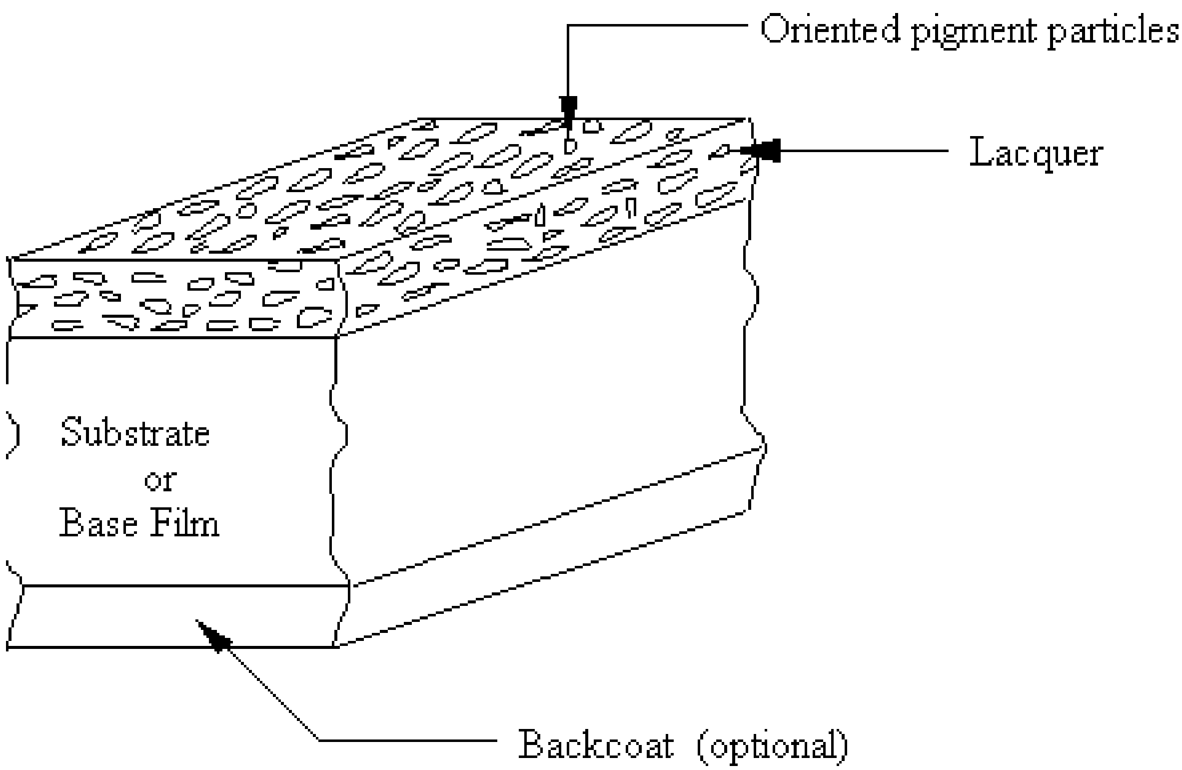

Figure 1-1: Cross-sectional view of magnetic tape

containing a single reel, or in a *cassette* containing two reels (supply and take-up reels). Optical media are available only in the circular disk format, but may enter the market in a tape format soon.

For users of media, a number of characteristics of the technology are important: linear density, areal density, volumetric density, transfer rate, reliability, error correction, stability of medium and longevity of technology. The mass transfer methods have proved, over historic times, the stability of both the technology and the media, which were used. On the other hand, in about a century over which magnetic recording has been employed, its susceptibility to decay has been noted by almost everyone who has used it. The shorter lifetime is not due just to media characteristics, but is also related to the short lifetimes of particular technologies.

While rock carvings have survived for millennia under exposed conditions (*e.g.*, the Behistun inscriptions of Darius I), and frangible material like papyrus was preserved in arid climates such as the Egyptian desert, twentieth century electronic media are less durable, and require storage under controlled conditions. The complex nature of the media is one reason. Table 1 compares two commonly used media of the current decade, magnetic tape and CD (Compact Disk).

	Magnetic tape	CD
Number of layers	2 or 3	3
Substrate	PET	Polycarbonate
Signal carrier	γ-Fe_2O_3, MP	Metallic backing on polycarbonate
Signal encoded in	Magnetic domains	Reflectivity change

Table 1-1.

Compared to metal, papyrus or clay tablets, magnetic tape and CD (Figs 1-1 and 1-2) are made up of layers with different characteristics. Each of them reacts differently to changes in the most basic environmental conditions: temperature and humidity. Manufacturers claim a hundred year life for CDs, but magnetic tape will last barely 15 years unless stored in carefully controlled environments. Polethylene terephthalate (PET), the substrate used in magnetic tapes, has been estimated to have a lifetime of a thousand years [Smith, Brown and Lowry 1982, 1983], but the complex soup which goes to make up the lacquer containing the magnetic pigment is susceptible to hydrolytic degradation of the binder. It is doubtful, however, if the devices required for playback will be available for longer than a decade after introduction of the technology.

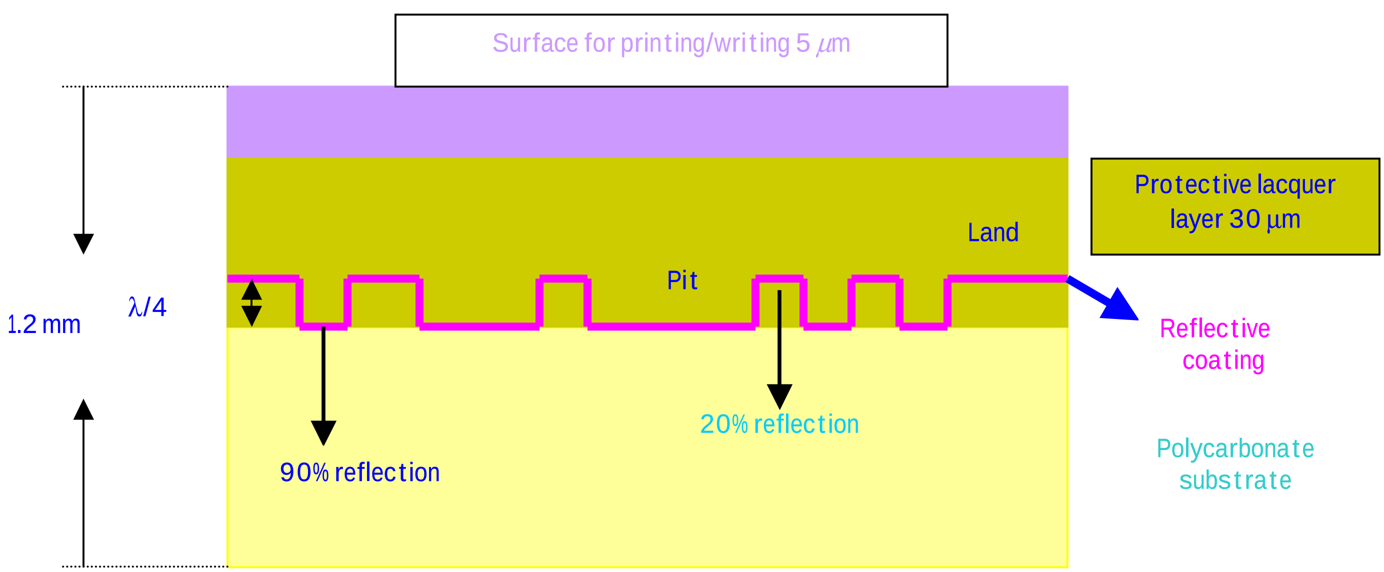

Figure 1-2. Cross-sectional view of Compact Disc

Non-human readable recording can be *analog* or *digital* (Fig #a). Analog recordings can be produced both on magnetic and on optical media. The response of an analog device is generally proportional to the input stimulus; in other words, the output is a *linear* function of the input. The response of a digital device is of the on-off, abrupt kind, compared to the smooth, continuous response from an analog one; stated differently, a digital device responds in a *non-linear* fashion to its input.

Analog recordings lose their fidelity as they are copied, and the loss is more noticeable with the later generations. Washed out colors on a VHS tape, which has been copied from another VHS copy, exemplify this situation. Even the original recording can become corrupted as the recording does not have error correction, or even detection, built in. Digital recordings, on the other hand, can be copied from generation to generation with little loss of data integrity. This is possible because digital data, represented as *bytes* or *words*, can be subjected to arithmetic and logical operations to generate extra *CRC* (cyclic redundancy check) or parity bits or bytes and, when these are stored and retrieved with the original data, the operations can be reversed to determine if an error has occurred during readback and possibly even corrected. Error correcting codes (ECC) are designed to handle both *random* errors (also called *bit* errors) and *burst* errors (when a whole sequence of adjacent bits can either not be read, or the error extends over a sequence of bits). The CD-ROM ECC, for example, can handle data missing due to a 4 mm scratch on the disk surface. Appendix B gives an overview of the calculation and layout of the data and ECC bytes for the Ampex DD-2 tape system.

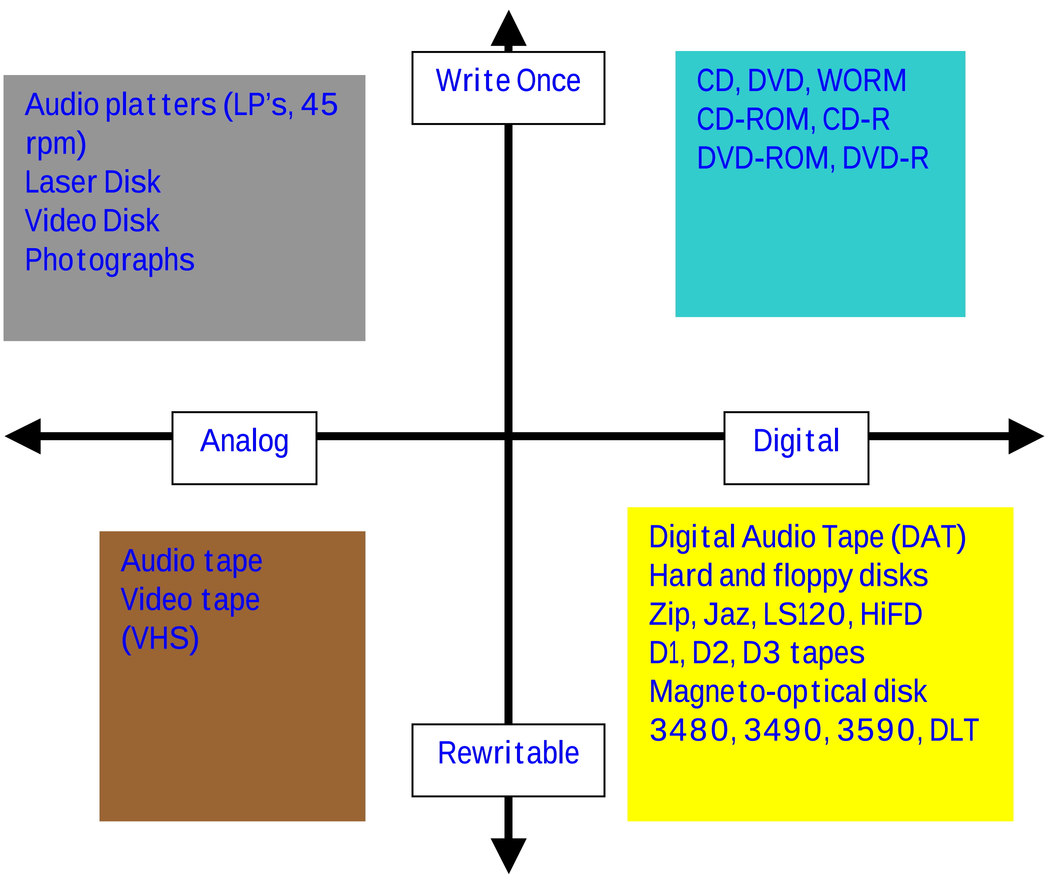

Figure 1-3. Two-dimensional diagram of data recording

Unlike printed books, recording technology available today can be used to store not just text drawings and paintings, but also sound, animation and moving pictures, separately and in combination. These recordings can be made on both analog as well as digital technology, but only the digital format is capable of producing clone-like copies of the original recording. The digital format introduces a number of processing layers.:

1. Size of word and order of bits within the word (*byte length*, *endian order*)
2. Selection of error correction code (ECC): number of levels of ECC, interleaving
3. Channel code (also called modulation code): this is designed to reduce the dc content (over a given sequence of bits, the number of 1's must equal the number of 0's), enable clocking (which tells the detecting system the starting and ending points of the receiving window), and reduce or eliminate the effects of *intersymbol interference* (ISS).

These requirements, making up the *physical format*, impose a heavy processing load on the recording system, and the processing itself is accomplished using software (in the form of *firmware* in a Read-Only Memory [ROM] in the recorder). This is the *digital overhead.* The application, which creates the records on the media, may impose a *logical* format, *e.g.*, a file structure. Together, the physical and logical formats add to the complexity and to the difficulty of preserving the data. *Recovery of the bit stream from a media is not sufficient for the recovery of the data committed to the media*.

A simple example will illustrate the complexity. CD-ROMs use an 8-bit byte as the data unit. 680,000,000 user bytes (~650 Mebibytes, see Appendix A on units and prefixes) can be recorded on a CD-ROM. We will denote this by *C.* The ECC overhead is 13%, and increases the number of bytes needed to be recorded to ~770,000,000 (1.13*C*). The channel code is Eight to Fourteen (EFM) and 3 additional merge bits are used; each 8 bit-byte is written as 14+3 = 17 bits on the CD. The total number of bits written to the CD is therefore 770,000,000 × 17/8 ≈ 1,636,630,000. The numbers are approximate, and no attempt has been made to show how the bytes are ordered, nor the steps carried out to generate the ECC bytes. A somewhat longer code walk-through for the generation of the bit-stream for the Ampex DD-2 tape system is outlined in Appendix B.

2. Magnetic Recording

Valdemar Poulsen demonstrated magnetic recording in 1898 in Copenhagen, Denmark. His set-up is described in Jorgensen [Finn Jorgensen: The Complete Handbook of Magnetic Recording]. As with other electronic equipment, recording devices were mostly in studios and laboratories at first. After WW II, electronics increasingly expanded into the consumer entertainment segment, and round-reel recorders for music replay began to be available. In the early 50's, IBM and UNIVAC began selling digital computers in the US. Magnetic tapes were the ubiquitous interchange media in the days before networking. As computer peripherals, they have been around since 1952, and predate the disk drive by five years; thus they also served as the earliest secondary storage media, giving way, gradually, to the magnetic drum and, later, to the magnetic disk. Table 1 lists the various ½-inch tape drives that IBM, once the principal peripheral provider, has introduced.

Year	Product	Package Geometry	Pigment	Linear Density* Bpi	Transfer rate KB/s	Capacity MB	Track density tpi
1952	IBM726	Round	γ-Fe_2O_3	100	7.5	1.4	14
1953	IBM727	Round	γ-Fe_2O_3	200	15	5.8	14
1957	IBM729	Round	γ-Fe_2O_3	800	90	23	14
1965	IBM2401	Round	γ-Fe_2O_3	1600	180	46	18
1968	IBM2420	Round	γ-Fe_2O_3	1600	320	46	18
1973	IBM3420	Round	γ-Fe_2O_3	6250	1250	180	18
1984	IBM3480	Square	CrO_2	38000	3000	200	36
1991	IBM3490E	Square	CrO_2	76000	3000	800	72
1995	IBM3590	Square	MP-II	92000	9000	10000	256

*This is the user bit density; the number of flux reversals per inch (fci), which is not indicated here, will include the channel (modulation) code, ECC (error correction and control) and other overhead is higher.

Table 2-1. From Round Reel to Square – the evolution of tape technology from IBM

This market, however, is not restricted to just ½-inch tapes; other widths in use in the computer segment are: 4 mm (called DDS and derived from the Digital Audio Tape), ¼-inch (called Quarter Inch Cartridge or QIC, and, more recently, Travan), 8 mm (Exabyte, Mammoth, Sony Advanced Intelligent Tape, Ecrix), and 19 mm (ID-1, DD-2). A representative list appears in Table 2.

On any medium, data are written along *tracks.* On disks, the tracks can be either concentric circles, or one long spiral, either right-handed or left-handed, moving out from the center towards the outer edge of the disk. In the case of tapes, there are three methods in use currently for track layout. The simplest is *longitudinal* where the tracks define lines along the length of the tape, see Fig 2-1. This corresponds to the way books are printed in most west- European languages.

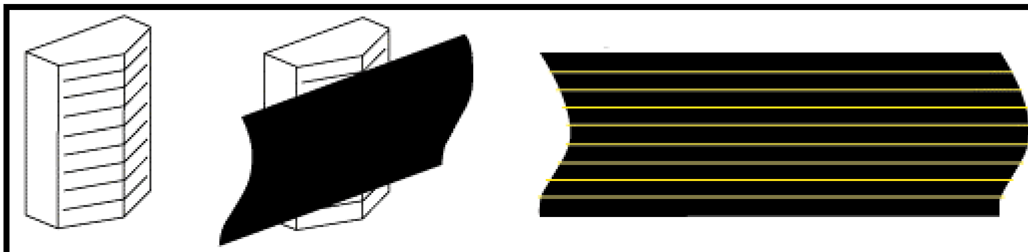

Figure 2-1. Longitudinal tracks

Type	Drive Manufacturer	Product Name	Capacity* (GB)	Transfer Rate (MB/s)	Recording Format
19mm	Ampex	DD-2	50(S)	15	Helical
19mm	Ampex	DD-2	150(M)	15	Helical
19mm	Ampex	DD-2	330(L)	15	Helical
19mm	SONY	ID-1	8.7	2 - 64	Helical
19mm	SONY	ID-1	43	2 - 64	Helical
19mm	SONY	ID-1	96	2 - 64	Helical
1/2 in*	IBM	3590	10	9	Serpentine
1/2 in*	STK	Redwood	50	11	Helical
1/2 in*	STK	Eagle 9840	20	10	Serpentine
1/2 in*	Quantum	DLT2000	15	1.25	Serpentine
1/2 in*	Quantum	DLT4000	20	1.5	Serpentine
1/2 in*	Quantum	DLT7000	35	5	Serpentine
1/2 in*	SONY	DTF-1	42	12	Helical

1/2 in*	SONY	DTF-2	100	24	Helical
1/2 in*	SONY	DIR-120	38(S)	15	Helical
1/2 in*	SONY	DIR-240	38(S)	30	Helical
1/2 in*	SONY	DIR-120	125(L)	15	Helical
1/2 in*	SONY	DIR-240	125(L)	30	Helical
1/2 in*	LMS	NCTP	18	10	longitudinal
1/2 in*	HP, IBM, SEAGATE	LTO			
8 mm	EXABYTE	Mammoth	20	3	Helical
8 mm	SONY	AIT-1	25	3	Helical
QIC^	TANDBERG, SEAGATE,HP	QIC-4	4	0.8	Serpentine
QIC^	TANDBERG,SEAGATE,HP	QIC-5010	16	1.5	Serpentine
QIC^	TANDBERG, SEAGATE,HP	QIC-5020	25	4.5	Serpentine
4 mm	HP, SONY, EXABYTE	DDS-1	2	0.183	Helical
4 mm	HP, SONY, EXABYTE	DDS-2	4	0.366	Helical
4 mm	HP,SONY, EXABYTE	DDS-3	12	1.5	Helical

*S, M, L – Small, Medium, Large denote cassette sizes

Table 2-2. The variety of tape formats

The video industry introduced the *helical scan tracking* format, Fig 2-2. Instead of the tracks running parallel to the long edges of the tape, they subtend an angle of ~5° to the edge. The chief advantage of this format is its ability to use most of the tape area for data recording, *i.e.*, an increased *areal efficiency*. DAT, VHS and Exabyte tapes are examples of technologies using helical scan recording. The recording/playback heads are mounted on a drum which rotates at about 3600 rpm while the tape itself moves forward at a few meters per minute. The relative velocity between the head and the tapes, however, can reach values as high as 50 km/hr (about 30 mph). It can be seen from the figure that the tape is wrapped around the drum at a shallow angle, and the wrap angle (the angle between the entrance and exit points of the tape wrap) is between 90° and 180° depending on the particular technology. Contact between tape and heads is extremely close on helical scan drives, and the high relative velocity can cause abrasive wear in both head and media.

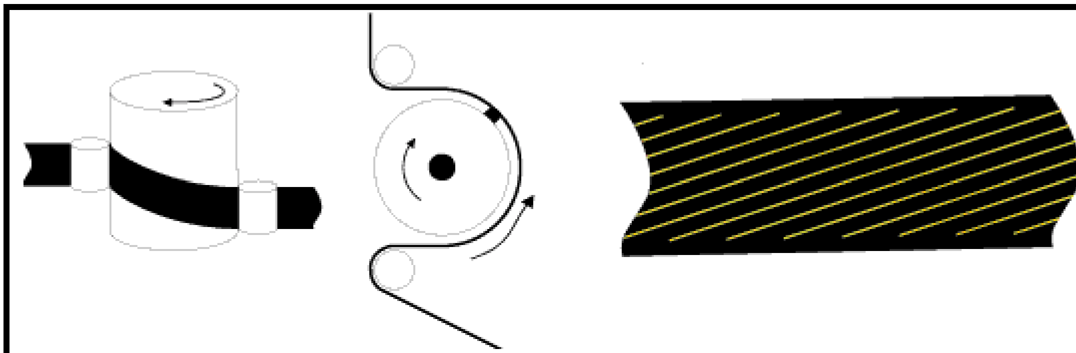

Figure 2-2. Helical scan recording

The third recording format is *serpentine* which is known as *boustrophedon* when applied to human writing. The serpentine format, illustrated in Fig 2-3, is very much like longitudinal except that the tracks change direction when they reach the end of the tape.

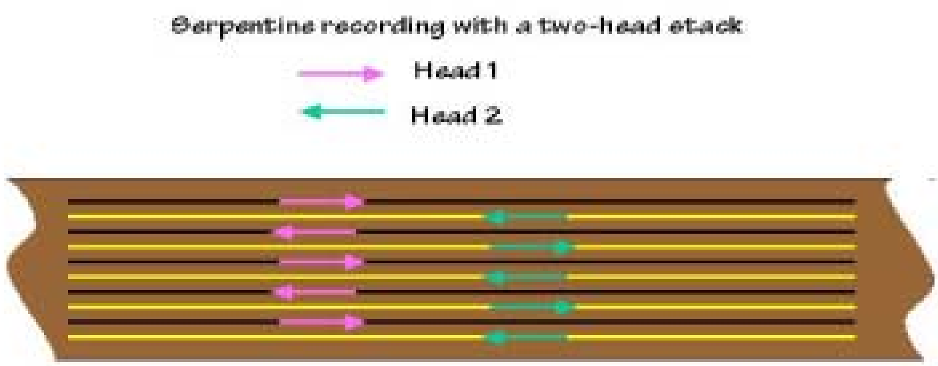

Figure 2-3. Serpentine Recording

There is one recorder, the DCSRi, from Ampex, which uses *transverse* recording. The DCSRi uses 1-inch wide tape, which passes over a canoe while in contact with a cylindrical read/write head (see Fig 2-3). This results in a set of tracks being written which are perpendicular to the long edges of the tape.

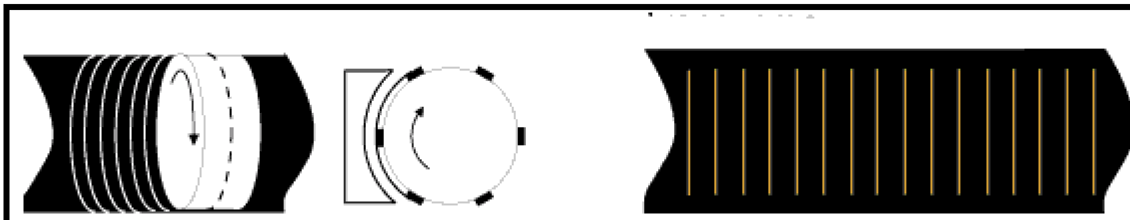

Figure 2-4. Transverse Recording

Arctuate track recording has been tried out in the laboratory, but has never been brought to market in a product. Figure 2-5 shows an example of arctuate recording. The read/write heads are mounted on a circular disc rotating in a plane parallel to that of the moving tape. The rotational motion of the disc, combined with the linear movement of the tape, results in tracks in the shape of an arc being written on the tape.

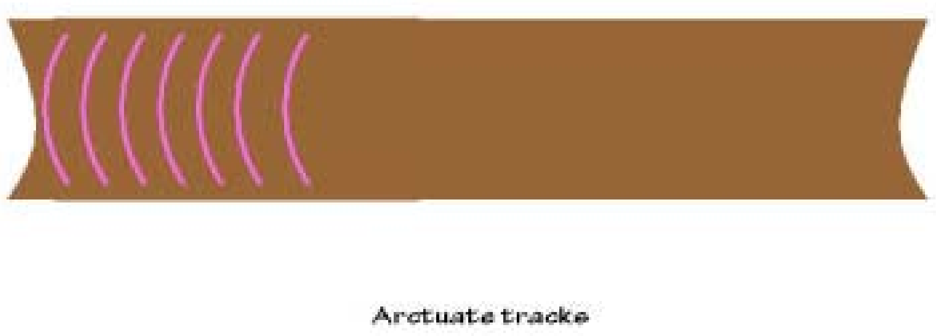

Figure 2-5. Arctuate Recording

Tape drive prices range from just above a hundred dollars for low-end QIC and Travan systems with EIDE interfaces to well over a quarter of a million dollars for high transfer rate 19 mm helical scan recorders. Media costs average from a few dollars for 3480-type cartridges to a couple of hundred dollars for high-capacity DD-2 cassettes.

Inspection of Table 2-1 shows that both linear bit density (bpi) and capacity (MB) have increased by a factor of 10^4 between 1952 and 1995, a period of 43 years. In the same time span, track density has increased only modestly, from 14 tpi to 256 tpi. Helical scan recording packs tracks more closely together, and the track density is almost an order of magnitude larger than in longitudinal recording. Transfer rates have lagged capacity and linear density increases, and has grown only by a factor of 10^3. Large DD-2 cassettes are already approaching a TB in capacity; if the drives can only write to them at 15 MB/sec, it takes about 66 sec to write a GB if data can be streamed to the drive at this rate in a sustained fashion. A full cassette will then require about 6600 sec to be filled, or read completely.

In the semiconductor industry, Moore's Law (named after Gordon Moore, one of the founders of Intel Corporation) states that the number of components on an integrated circuit (IC) doubles every 18 to 24 months. Since 1990, areal density of magnetic disks has been doubling every 18 months. More recently, magnetic tapes have been able to adapt to this pace too. In September 1999, IBM announced a laboratory demonstration of 35.5 Gb/in^2 for disks, and products at this density are likely to be available in the market within the next 3 years. This raises an interesting question: what is the ultimate areal density, which can be achieved in magnetic recording. The answer is not straightforward. Using current technology, it is very likely that the *superparamagnetic* limit will restrict densities to be below ~100 Gb/in^2. The industry is already using single-domain grains; as their size decreases, thermal effects can cause orientational instability in the domains and this leads to the loss of recorded data. With the exception of magneto-optical recording, all other forms of magnetic recording result in the creation of magnetic regions, which lie in the plane of the media. Work has been going on in the last two decades, mostly in Japan, on *perpendicular* recording. In this technique, the magnetic dipoles are oriented at right angles to the media surface, and it may permit the density to go beyond the superparamagnetic limit for in-plane magnetization. However, ring heads can no longer be used for writing; the two poles of the inductive write head need to be on opposite sides of the media being recorded. Alternatively, a magnetically soft material may be used as an underlay to complete the flux circuit for a modified ring head. Barium Ferrite ($BaO.6Fe_2O_3$), a green pigment often used in the lining on refrigerator doors to keep them closed, is an ideal material for perpendicular recording.

3. Optical recording

Optical recording is a term which is used for a number of technologies whose common characteristic is the use of a laser, principally for reading. It embraces both the mass-transfer and energy-transfer methods. The compact disc (CD), introduced in 1982, can be mass-produced because, once a glass master has been made, copies can be reliably and rapidly manufactured by injection molding. On a CD, a spiral track, ~5 km long, holds data in the form of *pits* of various lengths, with pieces of intervening *land* of different lengths. The pit depth is $\lambda/4$, where λ is the wavelength of the laser used for reading (for *CD circa* 780 nm, which lies in the infra-red). The light reflected from the bottom of a pit is 180° out of phase with the incident radiation, and interferes destructively with it. The net result is a substantial decrease in the intensity if light reflected from the bottom of a pit; about 90% of the laser light is reflected from a *land* area, whereas the reflectivity of a pit is only about 20%. DVD (Digital Versatile Disk) uses the same principle, although the wavelength of the laser light is smaller. Table 3-1 summarizes the principal characteristics of CD and DVD.

	CD	DVD
Disc diameter (mm)	80, 120	120
Number of sides	1	2
Thickness per side (mm)	1.2	0.6
Track pitch (μm)	1.6	0.74
Minimum pit length (μm)	0.834	0.4
Laser wavelength (nm)	780	640
Data capacity per layer (GB)	0.68	4.7
Number of layers	1	1, 2 or 4
Numerical aperture	0.6	0.38 to 0.45
Modulation code	8/14 (8/17 with merge bits)	8/16
User data rate, 1X (Mbps)	1.41	11.08
Error correction	CIRC	RS-PC
Error correction overhead %	34	13
Format overhead %	252	136
Scanning speed (m/s)	1.2 to 1.4	3.49 to 3.84
Rotation speed (1X) (rpm)	200 to 500	570 to 1600

Table 3-1. Principal characteristics of CD and DVD

Although the CD and DVD disks have the same diameter and thickness, the storage capacity of DVD, on a single layer, is almost seven times that of a CD. This is achieved through the use of laser with a shorter wavelength, a more efficient channel code and stronger error correction code. The diameter of the spot which can be produced using a source of wavelength λ is given by the expression

$$d = C\lambda/NA$$

where *NA* is the numerical aperture of the lens and *C* depends on the distribution of the energy of the beam at the focal plane. *C* is 0.6 for lenses used in CD and DVD drives.

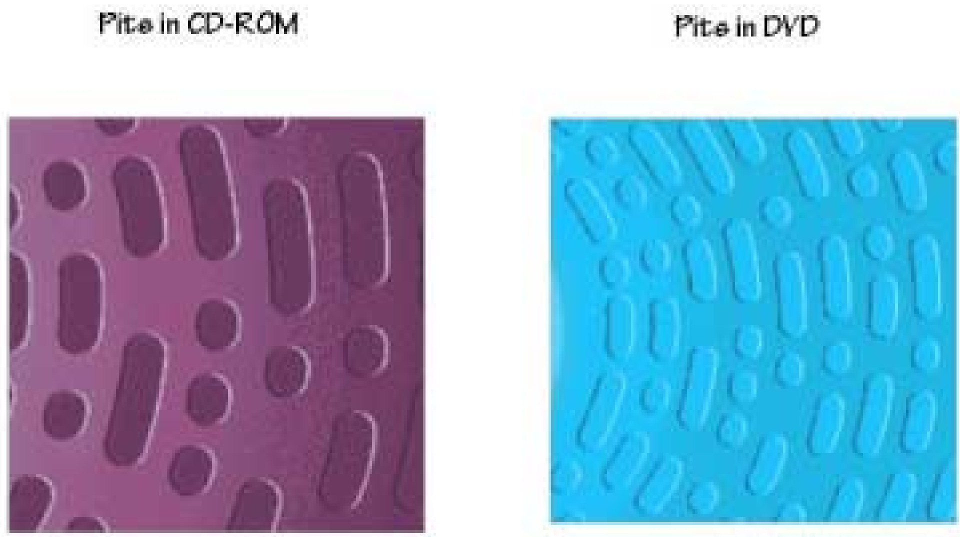

Figure 3-1. Pit size in CD and DVD

It is possible to have reflectivity differences created in other ways than by removal of mass. In CD-R (CD-Recordable), this is achieved by a laser-induced reaction in a polymer dye; the ones generally used are cyanine and phthalocyanine dyes. The *irreversible chemical* reaction, which gives the user a Write Once Read Many (WORM) media, has regions of low reflectivity where the dye has undergone a reaction. CD-RW (Rewritable) media use a coating of an alloy which can change from an amorphous to a crystalline phase when illuminated by a laser. The crystalline regions have higher reflectivity than the amorphous areas. The amorphous to crystalline, and the reverse, *phase transformations* are *reversible.* Manufacturers, however, caution that material fatigue constrains their use to a maximum of about 1000 write cycles. Since the laser power used for reading is below the threshold for initiating a phase transformation, a CD-RW disc can be read more times than it can be written on.

Calimetrics, of Alameda CA, has developed a *pit-depth modulation* technique which is based on an *M-ary* code and promises to increase both the storage capacity and transfer rate by a factor of 3 or more. The technique can be used with both CD and DVD technologies.

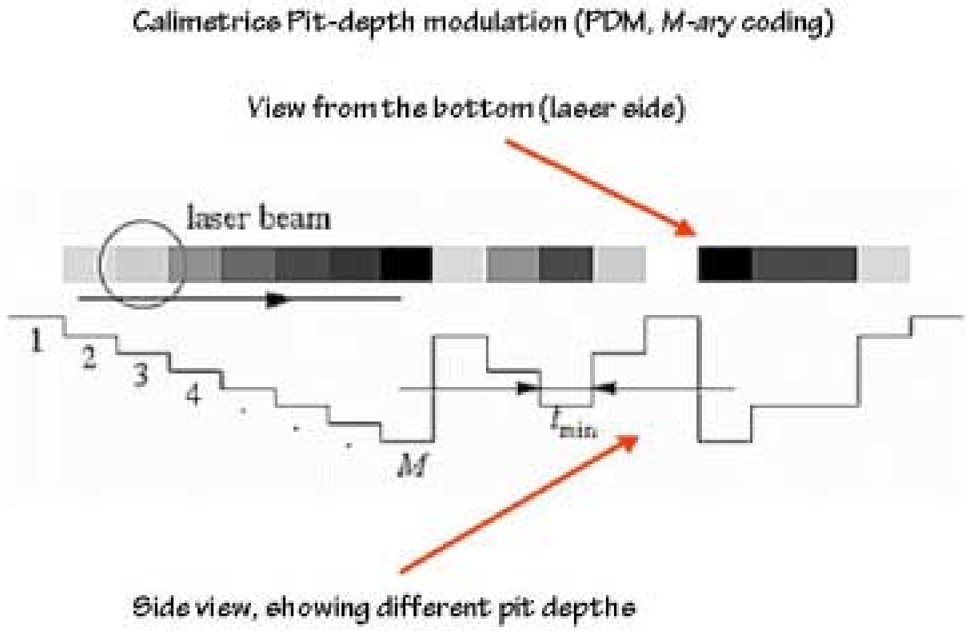

Figure 3-2. Pit depth modulation

Magnetic hard disks, and removable media like Jaz, Zip and magneto-optical disks, record data along a series of *concentric* circular tracks on the surface of the disk. Track density depends on the particular technology, but is about 20,000 tpi for hard disks. Hard disks rotate at a constant rate; it used to be 3600 rpm (revolutions per minute), but more recent products operate at 5400, 7200 or even 10,000 rpm. The length l of a track is related to its radius r by the expression

$$l = 2\pi r$$

Inner tracks, being shorter, hold less data than outer tracks. Hard disks, and most high-density removable disks, therefore use *Zoned Bit Recording* (ZBR), which subdivides the disk surface into zones of different linear densities. A 3.5" hard disk drive for a PC, for example, typically has about 15 zones. One consequence of ZBR is that the transfer rate to/from the disk is dependent on the location of the data. The highest transfer rates are achieved for the outer tracks, but the innermost tracks have the highest data densities. Vendors have been known to load and locate files on the outer tracks to boost perceived transfer rates in benchmark tests [I Bird, R Chambers, M Davis, A Kowalski, B Lukens, S Philpott and R Whitney, *Evaluating RAID in the Real World*, Sixth Goddard Conference on Mass Storage Systems and Technologies, GSFC/CP-1998-206850, pp 345-354]. CD and DVD, on the other hand, record along a *single, continuous, spiral* track, which starts from the inner region of the media and runs outward in a clockwise fashion. One advantage of starting the track at

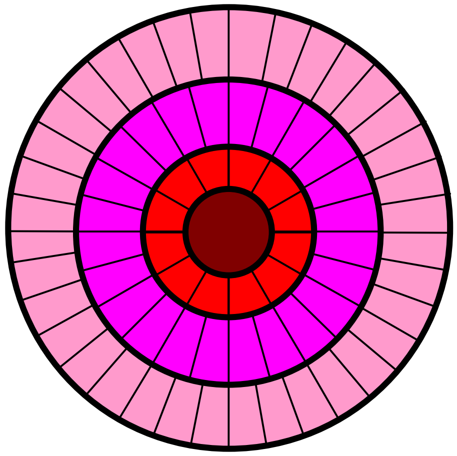

Figure 3-3. Zoned Bit Recording

the inner edge is that media of more than one diameter can be designed. Small CD disks, with a diameter of 80 mm, can be played on the same drives which are used for the 120 mm discs. Both CD and DVD drives use a technique called *constant linear velocity* (CLV) to read data from the tracks. The track passes under the read head at the same linear velocity irrespective of where on the disk surface the track segment is located; since outer tracks are longer than inner tracks, this means that the rotational speed of the CD is highest at the innermost track, and smallest at the outermost one and varies continuously from 200 to 500 rpm (for 1X drives) as the head moves in from the outer tracks. Rectangular CD disks, whose dimensions are close to those of a business card, with a capacity of 80 MB, are being used in place of traditional business cards by some organizations. The printed area contains the usual information, but these cards can be played on most multi-read CD-ROM drives to display multimedia images or play sound clips.

Recordable versions of CD exist in both WORM and rewritable formats; these are CD-R and CD-RW. Neither CD-R, nor CD-RW is designed for mass production. In CD-Rs, the signal carrier is a dye (cyanine or phthalocyanine, see Fig 3-1), which is polymerized by a high-power laser. Areas, which have undergone reaction, have a smaller reflectivity than the unwritten areas. However, the reflectivity values, and differences, are not the same as those for CD. Hence, CD-Rs and CD-RWs can be read only by those CD drives which have *multi-read* capability. Unlike the regular edges found for the pits on stamped CDs, the peripheries of the marks on both CD-R and CD-RW disks are more irregular. Another factor affecting the shape of the marks is the speed of the recorders; CD-R and CD-RW recorders are available at various recording speeds: 1X, 2X, 4X, 6X and 8X.

There are four versions of recordable DVD: DVD-R, DVD-RAM, DVD-RW and DVD+RW. The last three are also writable and are competing formats. DVD-RAM (1.0) divides the disk surface into 24 zones, and uses *Zoned Constant Angular Velocity* (ZCAV) for reading and writing. Within a zone, the angular velocity is fixed, but it changes when a new zone is accessed.

In 1989, seven years after the introduction of CD, the installed base exceeded 25 million drives in the US. By comparison, 70 million Sony playstations were sold worldwide in less than five years.

CD-ROM drives have become standard peripherals on all Personal Computers. Most software is sold on CD-ROMs, and there is a standard called *El Torito* which permits booting from a CD-ROM. The CD has become the analog of the phonograph record in the multimedia era; it has enjoyed a long and successful life in the entertainment market, and has managed to move into computer software distribution and file backup.

CD and CD-ROM media are among the most stable. The data is recorded on them as pits, and a thin layer of a reflective metallic layer (Al) is deposited on it before it is sealed in with another protective layer. The substrate, polycarbonate, is less susceptible to hydrolytic degradation than polymethylmethacrylate (PMMA) used in laser discs. *Laser rot*, oxidation of the reflective metallic layer, has been reported in some cases, but material control and modern manufacturing practices have reduced its incidence. In rewritable media, the signal carrier, a phase change alloy of Te, Sb, Ge and/or Sn, is in a finely-divided state and therefore susceptible to oxidation. The chances of this occurring are minimized by using a number of protective layers: ZnS/SiO_2 and UV-hardened plastic. It is more likely that the number of write cycles will be exceeded before the disc becomes unusable due to other causes. CD-R, and DVD-R, use a dye polymer. The dye is generally cyanine or Metalized phthalocyanine and there were early anecdotal reports of its being bleached by UV and blue lights. It is best to keep all CD disks away from direct light and heat. Manufacturers quote a lifetime of 100 years for CD and CD-ROM disks, and 30 years for CD-R.

Summary

Electronic recording is now just over a hundred years old, starting out as magnetic recording. Optical recording is only about 25 years old, and had to wait for the development and availability of lasers. Magnetic media are available in fixed and removable disks, and flexible tape loaded in cartridges and cassettes. Optical media, on the other hand, can be had today only as disks. While media, when stored under controlled conditions, can last for decades, the lifetime of a particular technology, especially in the magnetic recording world, is less than 10 years. Charts 4-1 and 4-2 summarize the growth and development of the two technologies.

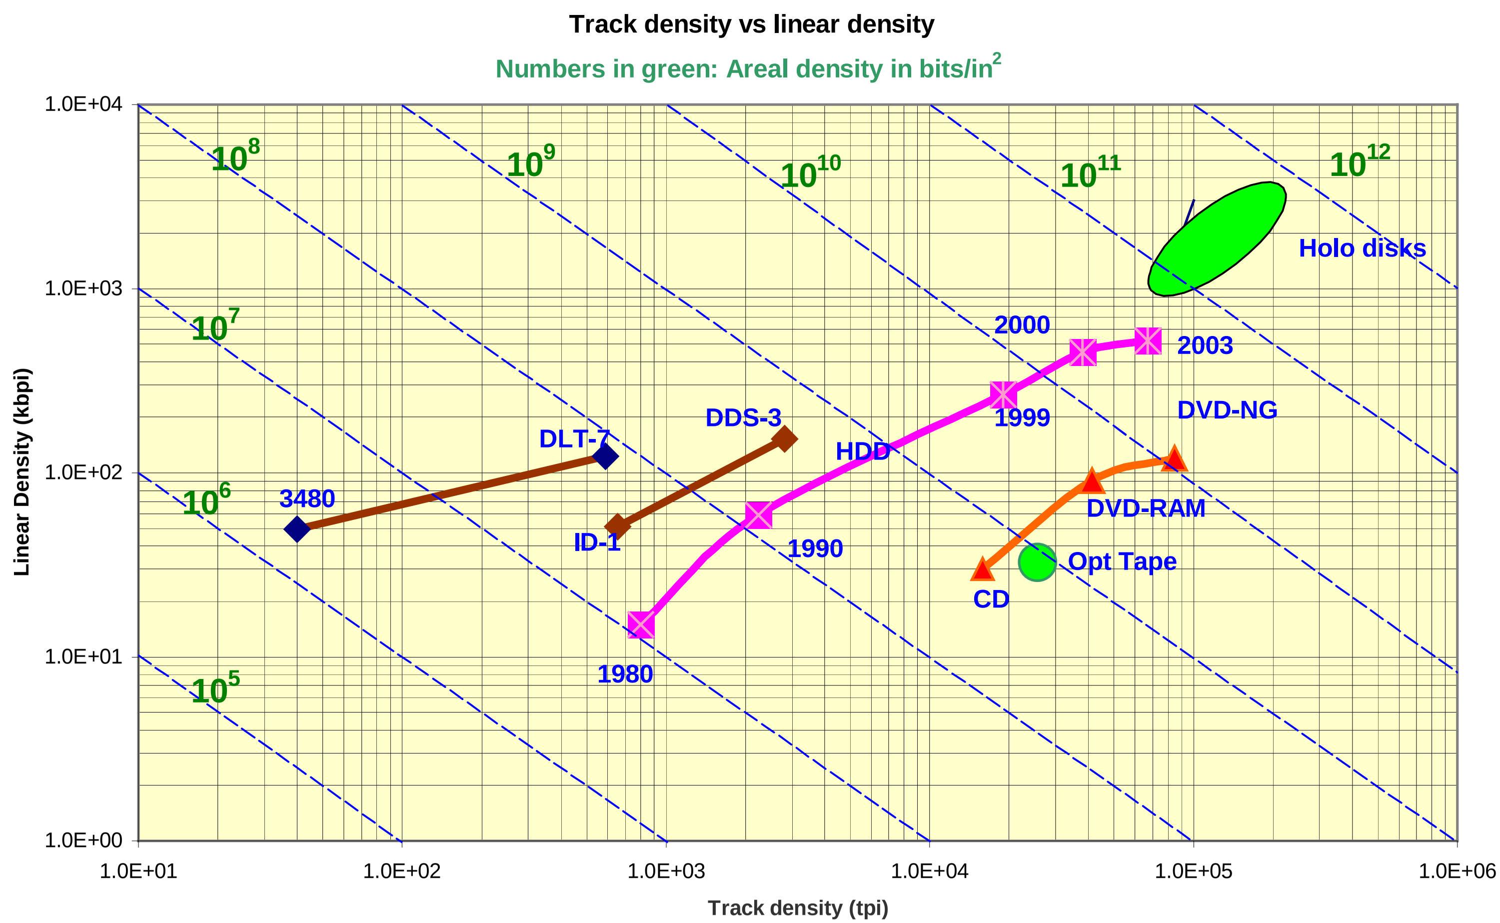

Figure 4-4. Growth of linear and track densities for various recording technologies

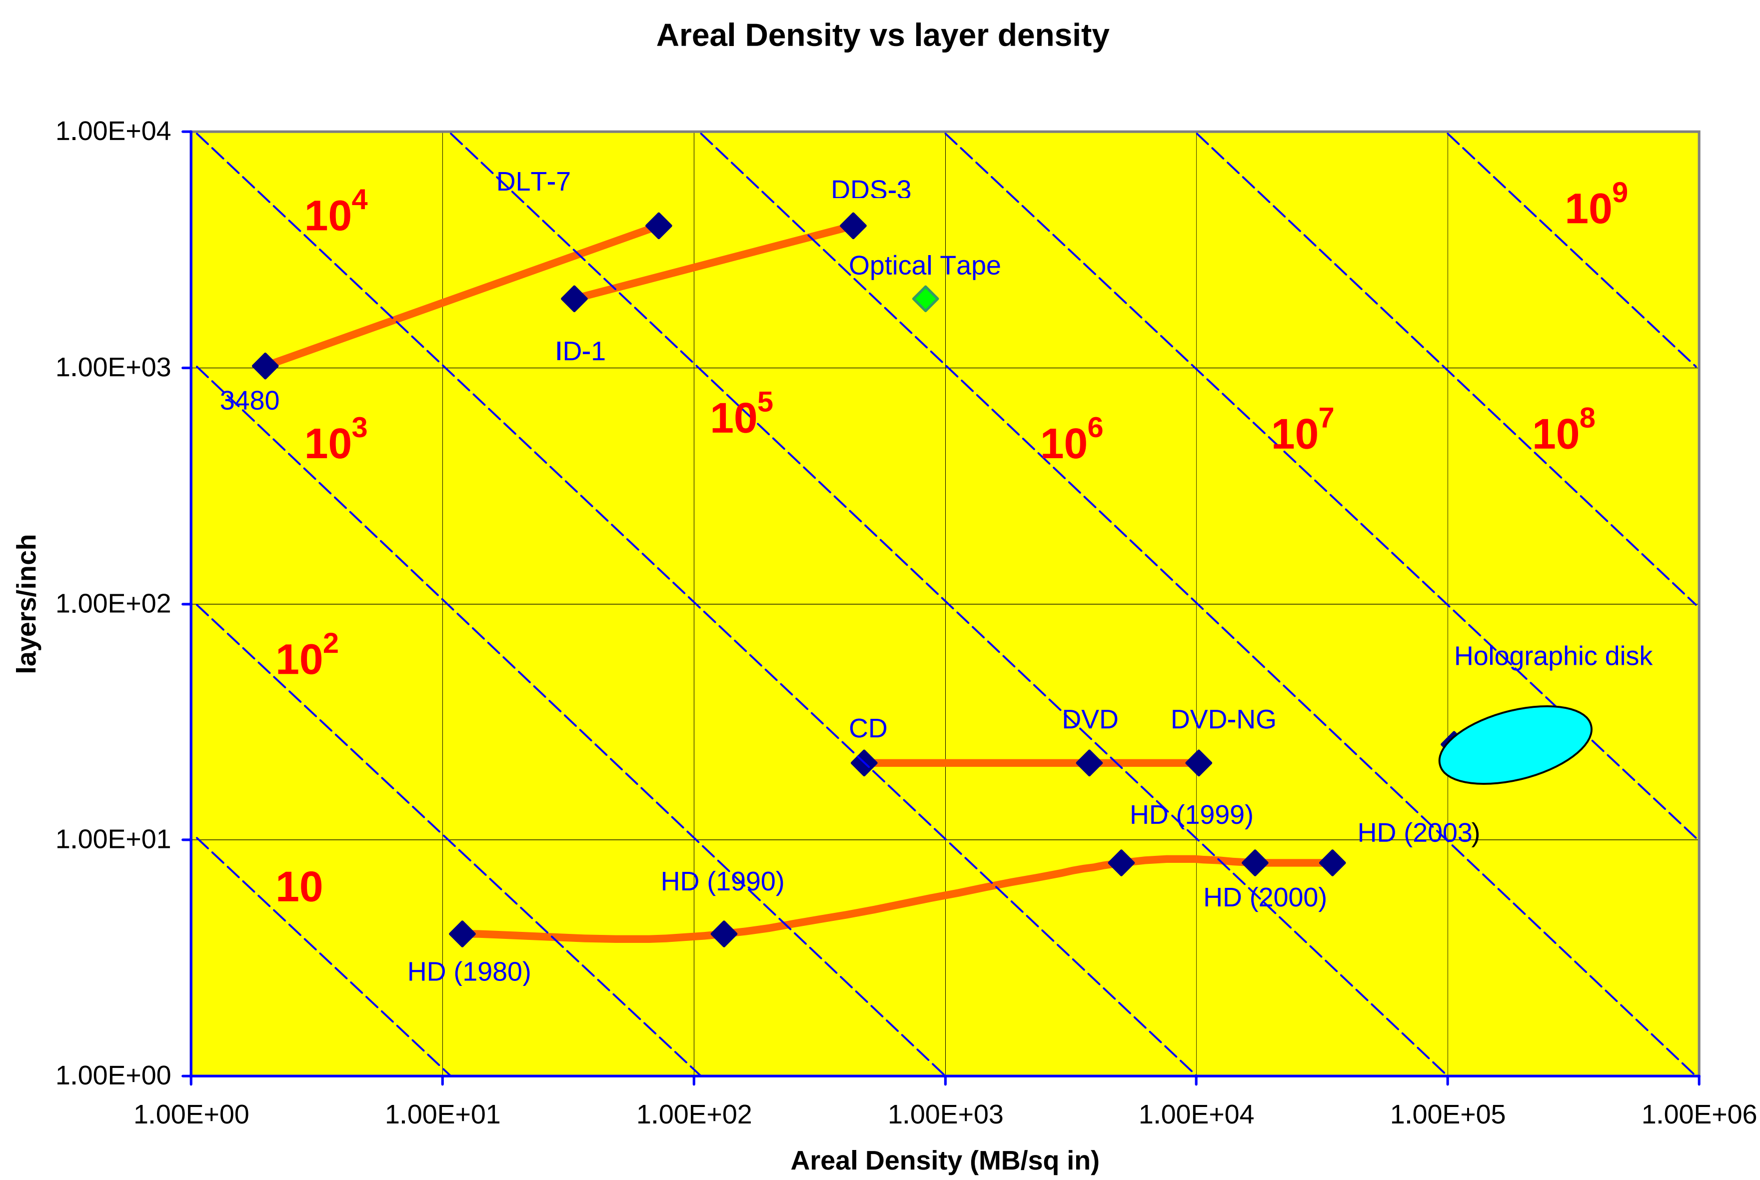
Areal Density vs layer density
layers/inch
Areal Density (MB/sq in)
1.00E+04
1.00E+03
1.00E+02
1.00E+01
1.00E+00
1.00E+01
1.00E+02
1.00E+03
1.00E+04
1.00E+05
1.00E+06
10^4
10^3
10^2
10
10^5
10^6
10^7
10^8
10^9
DLT-7
DDS-3
Optical Tape
ID-1
3480
CD
DVD
DVD-NG
Holographic disk
HD (1999)
HD (2003)
HD (1990)
HD (2000)
HD (1980)

Appendix A. Prefixes for magnitudes

Prefixes for magnitudes are defined by the International Conference on Weights and Measures ***ICBP***, and published by the ***BIPM*** (Bureau International des Poids et Mesures).
These prefixes all refer to *powers of 10*. In the computer industry, without further explanation, these prefixes are also *sometimes* used as *powers of 2*. The prefixes are usually four letters, and end in *–o* for the negative powers (but see below for exceptions).

yocto	y	10^{-24}
zepto	z	10^{-21}
atto	a	10^{-18}
femto	f	10^{-15}
pico	p	10^{-12}
nano	n	10^{-9}
micro	μ	10^{-6}
mill*i*	m	10^{-3}
cent*i*	c	10^{-2}
dec*i*	d	10^{-1}
deca	da	10^{1}
hecto	h	10^{2}
kil*o*	k	10^{3}
mega	M	10^{6}
giga	G	10^{9}
tera	T	10^{12}
peta	P	10^{15}
exa	E	10^{18}
zetta	Z	10^{21}
yotta	Y	10^{24}

More recently, it has been proposed that, for positive powers based on the radix 2, the prefixes should be modified so that the last two letters, or the last letter if the prefix has only three letters, are changed to *bi*; thus a me*bi*byte is 2^{20} = 1048756 bytes, but a megabyte is 10^{6} = 1000000 bytes.
Always watch the units; they can be tricky and treacherous, *e.g.*, ream.

1 Ream of paper = 480 sheets, usually
= 472 sheets for handmade and drawing paper
= 500 sheets for book paper and for flat-plate newsprint
= 516 sheets in a perfect ream

Appendix B. Code walk-through for DD-2

1. Accept a set of 1,199,824 bytes (source bytes). After encoding for error-correction, section identification, etc., this will grow to a size of 1,584,000 bytes and will be written out as a physical block on 32 adjacent helical-scan (HS) tracks.
2. Subdivide this set into 8 blocks of 149,978 bytes each.
3. Add a 2-byte checksum to each block, increasing its size to 149,980 bytes. The input byte stream has now grown to 1,199,840 bytes.
4. Combine the eight blocks into four data streams of 299,960 bytes each.
5. Augment each stream with 8 bytes of error detection information, increasing its size to 299,968 bytes and that of the original input stream to 1,199,872 bytes.
6. Encapsulate the data in the C3 (outermost) error correction code. Each group of 86 bytes is replaced by a 96-byte C3 codeword. This code, denoted (96,86,5) is capable of correcting 5 bytes in an incorrect codeword when the location of the errors is not known. If the positions of the erroneous bits are known, upto 10 bytes can be corrected.
7. The C3-corrected stream has 13952 C3 codewords; the original data stream, at this stage, has increased to 13952 X 96 = 1,339,392 bytes. Shuffle the bytes among the 32 tracks on which they will be ultimately written. This is accomplished as follows:

byte 1 $\rightarrow$ track 1	byte 33 $\rightarrow$ track 1
byte 2 $\rightarrow$ track 2	byte 34 $\rightarrow$ track 2
.	.
.	.
.	.
byte 32 $\rightarrow$ track 32	byte 64 $\rightarrow$ track 32

assigning 41856 (=1,339,392/32) bytes to each track.
8. Insert 192 bytes of logical format information in each track whose size consequently increases to 42048 bytes (total 1,345,536 bytes).
9. Encapsulate the data stream from step 8 in the middle, C2, error correction code. This (106,96,5) code replaces each group of 96 bytes with a 106-byte C2 codeword; each of the 32 tracks, at this stage, has assigned to it, 438 C2 codewords (438X106 = 46428 bytes). The total number of bytes, in the 32 tracks, is 1,485,696.
10. Interleave the C2 codewords to reduce the effect of burst errors. The data is arranged in 438 columns, and is read out row-wise to fill 32 rows, each of which will occupy a helical scan track.
11. At the end of each group of 219 bytes, insert a 1-byte identifier. This is essentially the ordinal number of the group on the helical scan track. There are 212 groups on each track. The track now holds, with the addition of the 212 1-byte identifiers, 46,640 bytes. The total number, for all 32 tracks, is 1,492,480 bytes.
12. Encapsulate the C2-encoded data stream in the inner, C1, code. This (228,220,4) code replaces each group of 220 bytes(the group from step 11) by a 228-byte C1 codeword.
13. Each C1 codeword is augmented by 3 bytes of synchronizing information, leading to a 231-byte sync block. The number of bytes assigned to each track at this stage is $212 \times 231 =$ 48336.
14. The data is now ready to be written out on tape. Each track starts with a 174-byte preamble which is followed by a 18-byte track sync block, the 212 sync blocks from step 13, and is terminated by a 174-byte postamble. Each helical scan track now has 49500 user bytes.

15. At the outset, the user data stream of 1,199,824 bytes, when spread out over 32 helical scan tracks, averages to only 37494.5 bytes per track. However, after encapsulation in error-correcting codes, each track has 48336 bytes. The overhead for error correction is, therefore, 48336 - 37495 = 10841 bytes. This amounts to an overhead of 28.9%.
16. The choice of Miller squared code for the channel facilitates the identification of error locations. This code has a lower rate, 0.5, than the 8/9 code used in ID-1, but the lower rate provides additional error-detecting capability. The 49,500 bytes in the data, when converted to the Miller-squared modulation code, expand to 99,000 bytes.
17. The total overhead in the DD-2 system is 128.9%. One thousand bits of user data are committed to tape as 2289 bits.