Fermi National Accelerator Laboratory

FERMILAB-Pub-91/348

Data Storage Technology; Evaluation Criteria, Current Technology Experiences, Future Technology Expectations and Evaluation

J. Stith
Technology Tracking and Transfer, Computing Division
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

December 1991

* Submitted to *Computers in Physics, Computer Technology Review, Computer Design,* and *Supercomputing Review.*

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CHO3000 with the United States Department of Energy

Disclaimer

Data Storage Technology;

Evaluation Criteria, Current Technology Experiences, Future Technology Expectations and Evaluation

Joseph Stith
Technology Tracking and Transfer
Computing Division
Fermi National Accelerator Laboratory
P.O. Box 500
Batavia, IL 60510
stith@fnal.fnal.gov

December 1991

ABSTRACT

The conventional technology for bulk data storage until just a few years ago was the 2400 foot by half-inch 9-Track tape reel, stored at either 1600 or 6250 bits per inch. With the explosion of data, including within the field of High Energy Physics, it became clear that nine tracks across a half inch magnetic tape was not very efficient use of the media. This paper discusses various characteristics of sequential storage technologies and assist in deciding on a technology for a new project, or for changing an existing project.

INTRODUCTION

Non-standard tape devices were introduced into the computer data storage market, often borrowed from consumer or professional audio or video markets. These allowed more bits to be stored within a smaller form factor (either measured as bits per square inch of media, or gigabytes per square foot of floor space).

Some of the introductions have been the quarter inch cartridge (QIC), 8mm and 4mm helical tapes which rotate heads and write diagonal tracks, and 3480 (half inch) cartridges. With IBM seemingly banking on relatively minor improvements to the 3480 and 3490 drives, the industry is faced with a need for high-capacity storage of which vendors must now design on their own, rather than producing IBM-like drives. With this introduction of multiple new technology formats, it is no longer obvious what technology should be selected for a given application. Since there is no longer a de facto standard, the characteristics of each technology have to be reviewed for each application. Some such characteristics are the cost of the drive, the cost of the media to store some fixed amount of data, the throughput of the drive, reliability of the technology, and the compatibility with other sites or other vendors' equipment.

In today's computing center, it is impractical to attempt to support all of these emerging technologies (a combination of media, transports, and formats). It is best to select a standard technology (either by design or by de facto usage). However, the center is likely to be forced into also supporting both the previous technology, as well as the next up-and-coming technology (supporting even higher densities). Even if a new technology is "better" than the technology you currently use, it must be "significantly better" in order to justify the overhead costs of supporting an additional technology format. Likewise, different technologies may have conflicting benefits (price vs performance) and it might be best to support multiple technologies. It is beneficial, however, to limit the number of technologies with similar characteristics.

Optical disk offerings, both WORM and re-writable, are now more reliable, compatible, and readily available than they were previously. These provide significant cost savings relative to magnetic disks, due to their removability (and, therefore, robotic options), while still retaining random access capabilities. Their speeds remain somewhat slower than magnetic media, however.

The rate of CPU technology improvements is increasing faster than that for storage technologies, creating a large I/O gap making data storage a more important issue. As with most highly technical fields, however, the field of data storage is changing fast enough that it is often difficult to keep up to date. Vendor laboratories are able to produce higher densities of storage, either using the same media or different media. Future possibilities with helical storage (eg, D1, D2, or DX format for either 1/2" or 19mm magnetic tape) and half-inch optical tape prove worthy of watching.

One merely hopes that when it is time to choose a technology, it is not too difficult to determine which technology to use, and that the rapid turnover of new technology does not prematurely obsolete your choice, introducing support problems. The ANSI committee on Information Processing Systems (X3B5) acknowledges this in their Annual report of June 30, 1991, as they state that they plan to "proceed with individual projects. Some may fall by the wayside, and the committee will have wasted its time on those." [1]

Outline

High Energy Physics

Before discussing storage solutions any further, the context in which this information was gathered must first be presented. Typically tape drives are utilized only 5% of the time, and perform backups, most of that time. Within high energy physics tape drives are often used 95% of the time, around the clock, to write data tapes rather than backup. Hence the need for understanding what a High Energy Physics Research Laboratory (such as Fermilab) does.

Fermilab provides an environment for sub-atomic particles to collide with each other or to hit a fixed target. A collaboration of physicists, often from different institutions (i.e. universities), designs a detector to determine what particles are produced out of each collision. These detectors can have up to 150,000 data channels (and this number is growing every year).

Since the rate of collisions is very high and the amount of data per collision is also very high, no tape drive has the speed to support storing all of the data, nor would it be practical to purchase enough media to support it. Thus, *online* computers perform the first line filtering of data so that only the *interesting* collision events are actually recorded onto a media (i.e. tape).

This *raw* data is then passed through a single *track reconstruction* pass, which analyzes the data and calculates particle motions. The resultant tapes, referred to as *Data Summary Tapes* (DSTs), however, often require more data capacity than the input data. This data is then analyzed by several different physicists, both at Fermilab and at the collaboration institutions, each analyzing it in different ways.

The recorded data then has the following characteristics:[2]

1. Large data volume (2-20+ TBytes)
2. Large record sizes (0.1-1.0 MByte)
3. Variable computational intensity vs input/output rates
4. Raw or summary data may have wide distribution for analysis
5. Loss of 5% of the data due to media errors can be tolerated because there is enough raw data some can be lost, and the summary data can be re-constructed.

Since physics is changing at a rapid rate, the data taken today is uninteresting within five years, and does not need to have an archival life beyond ten years (while others often need longer lifetimes).

Storage Hierarchy

The first thing to acknowledge in analyzing storage technologies, as with most subjects, is that the world is not ideal. You can not store lots of data on a high quality inexpensive media using a highly reliable inexpensive fast drive. Thus, it is important to put the data that requires quick access on faster storage devices (which are more expensive), while keeping the lower priority data on the less expensive storage devices (which are slower). Supporting multiple technologies, however, increases certain overhead costs, so each selected technology must provide "significant" improvements over other technologies in use. The storage hierarchy attempts to maintain a balance between data needs and economic capabilities.

The desire to maintain any given data online (or nearly so via robotics) may vary through the life of the data. Thus, a given file may be accessed frequently during a one week period, but then not be accessed again for several weeks. A Hierarchical Storage Management (HSM) system which includes software that is capable of transparently moving the data from one storage technology to another is the best solution to this problem.[3] Unitree is an example of such a product based on the IEEE Storage Reference Model which details the framework and interfaces but does not detail the implementation method.[4]

A well balanced hierarchy allows you to leverage the lower cost magnetic tape media and drives, without significant penalties for the lower speed characteristics typical of magnetic tape drives. A balanced hierarchy may include any combination of the following:

Main memory (RAM and ROM)
Solid-state disk
Cached disk
Raid disk
Standard magnetic disk
Automated WORM and erasable optical disk libraries
Automated magnetic tape libraries
Manual WORM and erasable optical disks
Manual magnetic tape

This report generally refers to those products aimed toward competing with the traditional magnetic tape products and includes all layers past standard magnetic disk.

Additionally, just as a hierarchical storage management scheme can leverage your way around a drive bottleneck, additional software may be needed to leverage your way around a network bottleneck. Today, the computing environment is moving away from mainframes and toward many "smaller boxes" based on RISC processors that can, together, give you more total computing power than a mainframe for less money. This new distributed environment introduces the network file-server issue of getting the data to the correct distributed system in a timely manner. Intelligent network software, such as the Andrew File System (AFS, being renamed to DFS) with its local cache, will help leverage the network. However, a tiered Hierarchical Storage Management system which would migrate files between multiple file-servers and/or the target distributed node based on usage pattern would be the goal. Packaged dedicated file server systems, such as ones offered by Epoch, include hierarchical software support within the package.

Evaluation Criteria

When selecting a storage technology, you must determine what criteria are important to you. Some sample criteria include:

1. Availability
2. Drive Related Costs
    - A. Drive Unit Cost
    - B. Throughput
    - C. Cost/KB/Sec
3. Media Related Costs
    - A. Media Unit Cost
    - B. Capacity
    - C. Cost/GB
4. Functional Requirements
    - A. Sequential vs Random
    - B. Writability
    - C. Search Speed
5. Data Interchange and Compatibility
6. Robotic Options
7. Data Compression Options
8. Level of Integration Within Your Environment
9. Robustness
    - A. MTBF
    - B. Error Rate
    - C. Passes over read/write heads
    - D. Shelf Life
10. Standards and Competition
11. Supply Competition, Availability, Service for Drive and Media

1. Availability

When did (or will) the device become available for purchase on the open market. A device that is not yet shipping may be a great technological achievement in the lab, but it may get delayed or entirely canceled.

2. Drive Related Costs

A. Drive Unit Cost: Purchase cost of one drive unit. A central computing facility might have the luxury of looking at cost of throughput (ie, a drive might be expensive, but its speed makes it worth it), but institutions with smaller budgets (such as universities) have to look at entry level pricing (and have to live with slower throughput). Thus the center may be able to afford the $1/4 million drive, but the university can not. The center might purchase multiple slower, cheaper, drives, so the university can afford to purchase one or two of them for data interchange.

B. Throughput: How fast can you write data onto a single unit.

Drives can run from 200 KB/sec to 10 MB/sec or higher. Running a benchmark against a proposed drive is better than relying on vendor figures, because your application may request data in a fashion that slows down the performance of the drive unit. The typical example is a start/stop device which may have a high throughput rating, but may significantly slow down if you do not request data fast enough due to the time consuming process of stopping and re-positioning the tape.

Also, higher performance is not necessarily better. If you know your processing model and the processing required for your application, then you should *match* I/O with the processor. For example, if you have a 10 MIPs machine and need to perform an average of 100 instructions per byte, then a device that is 100KB/Second is fast enough. Purchasing a device capable of 500KB/sec will gain nothing since the CPU power is not sufficient to support it.

C. Cost/KB/Sec: If your application has to have a single data stream, then raw throughput is important, but if your application can use multiple drives to achieve the required aggregate throughput, then this calculation is important.

3. Media Related Costs

A. Media Unit Cost: Purchase cost of one volume of media. If you expect to put only small amounts of data onto each volume (despite the volume's capacity), this is important.

B. Capacity: How much data you can store on a single volume.

If your typical data file is 1GB in size, then a 200MB tape cartridge is probably not ideal. Likewise, if your typical file size is 150MB, then a 2GB sequential volume is probably not ideal due to the time required to skip over undesired files to get to the file you are interested in. Of course there are relatively easy work-arounds for both of these, but the time and effort to work around this issue, both in up-front development and in long-term day-to-day operations, must be part of your evaluation.

Note that some drives' capacities changes from one volume to the next. The typical scenario is when you use data compression (discussed elsewhere in this paper). Another example is if the file-marks themselves use a lot of real-estate on the media; writing many small files will then reduce the over-all capacity of a volume.

C. Cost/GB: For large capacity users that fill up many tapes, this figure is important. At Fermilab, this determination alone lead to the extensive use of 8mm drives because more science was possible due to higher total storage capacity with the given budget.

4. Functional Requirements

A. Sequential vs Random: The typical functionality question is whether or not your access needs are sequential or random access. If your application need is random, and you can not tolerate delays of migrating a file from a sequential device to a random access device (as in staging files between tape and disk devices), then your options are reduced to random access technologies.

B. Writability: Differences, especially in optical devices, in writability and re-writability. One can write over data previously written by most of these technologies. WORM optical drives, as their name implies (Write Once, Read Many), can only be written to once. This is not a problem for some applications, and is, in fact, desirable in certain applications.

C. Search Speed: File search time is also of concern if you do not always start processing data from the beginning of sequential media. If you wish to read data in the middle of a sequential volume, it takes a certain amount of time to reach the data you are interested in. This time varies from technology to technology.

5. Data Interchange and Compatibility

If your work is self-contained within one site, then interchange may not be an issue. However, within high energy physics, it is often the case that copies of the data are made to send back to various home institutions of the collaboration. Thus, it is important to use a technology that they are likely to be able to read on existing equipment or be able to purchase a drive at reasonable cost. Alternatively, you must deal with the issues of copying data from one technology to another, different media capacities, and different error handling techniques.

6. Robotic Options

This capability allows for the use of library systems which give you automated access to large number of volumes of data, with significantly fewer actual drive devices. This data is neither online or offline, and is often called sideline. Thus, a library unit may hold 100 volumes, but have only 4 actual drives. This allows for less expensive access to all 100 volumes, but also introduces the contention issues of trying to simultaneously read more volumes than the number of drives in the robot. If the drive is fast enough, this wait time is reduced.[5] [6]

Storage Technology Corporation's Automated Cartridge System (ACS) unit is a commonly used 3480 form factor high-end robotic example. This silo supports multiple tape drives and 6,000 tapes of 3480 form factor (200 MB each) for a total of over 1 TeraByte of data. It initially supported IBM channel connection, but now also supports SCSI connections. Future developments in non-3480 formats on media of similar physical size are under way by STC.

7. Data Compression Options

Data compression for visual data (i.e. pictures) will often be "loss-full" as it reduces the size of the data representation, resulting in an inability to recover the original figure, but maintains a figure of lower quality. In most data storage, however, data compression is "loss-less" and can be a method of encoding redundant copies of data to reduce the number of bits required to store the data. When a match on previous data is found, a pointer back to the previous data is required, rather than a repeat of the data itself. Thus, the original data can be recovered completely.

Typical compression rations of 3:1 (3-to-1) are achievable, which means that if you have a 300MB file, you only need a media capable of storing 100MB in order to store it. The compression ratio is determine by the repetitive nature of the data, however, and changes with each file. The amount of data which can be stored on a given media volume, therefore, will change from one volume to the next.

Although traditionally data compression has slowed down throughput, with today's chip technology, data compression is now very quick. In fact, if you take a drive capable of writing at only 300KB per second, add data compression logic to the front end that is capable of 3:1 compression, you can achieve an *effective* data rate of 900KB per second.

Because compression ratios change with each file compressed, data compression has had a quicker adaptation to sequential tape media with variable record size capability), rather than into the disk market (which uses fixed length sectors, tracks, etc). Due to the lower speeds of optical disk, however, increased research is being performed on compressing disk data. The technology will then be utilized on both optical and magnetic disk.[7]

Data compression can be performed in software, in the controller, or in the drive itself. Also, modules are available for some drives which plug between your interface and the drive and provide transparent compression. If your software already uses bits efficiently, as generally is the case with high energy physics, then your compression ratios will be less than the "average" quoted by the vendor. Fermilab's tests with IBM's IDRC compression has been roughly a 1.5:1 ratio. Note, also, that if your software or hardware (ie, controller) already compresses your data, compressing it again could actually *increase* the storage requirement for the data.

There is no standard compression algorithm, although the DCLZ (Data Compression, Lempel Ziv) method is being used by many 4mm DAT manufacturers.[8] IBM's IDRC (Improved Data Recording Capability) is the next closest thing to a standard, but must be licensed from IBM. Due to the current lack of a public standard any compressed data will often have to be kept in-house, although drives which offer compression also allow for the writing of the standard uncompressed format when it is needed for compatibility purposes. Standards committees are currently working on this issue, which will help the situation, but not eliminate the issue due to the fact that there will remain multiple standards.

8. Level of Integration Within Your Environment

The device must be attachable to your computer system in a reasonable fashion. This usually means making sure the hardware interface, and software device drivers, are available and robust for your system. Typical interfaces are the various flavors of SCSI, ESDI, TurboChannel, EISA (Advanced Computing Environment), Futurebus+, VME, IPI (Intelligent Peripheral Interface).

9. Robustness

A. MTBF: The device must have a reasonable Mean Time Between Failure, MTBF, or it could not survive in the market. However, you must determine what is sufficient for your application. Be careful to understand the *duty cycle* of the drive. If a drive is quoted at 50,000 hours MTBF, that figure could be accurate for someone performing backups and using the drive 10% of the time (10% duty cycle), but if you want to use the drive 50% of the time (50% duty cycle), then your MTBF is only 10,000 hours.

B. Error Rate: The error rate of the storage technology must also be reasonable. Are error correction schemes, such as Reed Solomon Dual Error Correction Code (ECC) in use? What about read-after-write verification? The important final question is "what is the uncorrectable bit error rate." In HEP a relatively high bit error rate is acceptable because each event is separate and no one event is enough to prove a new discovery (it takes a dozen or so). Thus HEP can withstand a higher bit error rate than, say, a financial application.

C. Passes over read/write heads: The number of passes of the media before error rates significantly increase may depend partially on the tensioning mechanism used in the tape transport[9] and whether air bearings (rather than capstans) are used. Lower cost drives will have the media touch the heads and capstans, while more expensive drives will have air bearings to prevent physical contact, thus increasing the life of the media.

D. Shelf Life: Shelf life of the media is important if you desire to save the data more than a couple years. Also, the environmental conditions required to obtain the full shelf life may differ from one media to another. In HEP, the data is usually not read after five years but we keep it around for about ten. In Astronomy, it is generally kept "indefinitely" which results in issues of copying old data from old media to a new media after some number of years.

10. Standards and Competition

All else being equal, it is best to have a device which satisfies an independent standard such as an ANSI or IEEE standard. This improves your chances that the technology will remain in the market place longer, helping you maintain support for years to come. Be sure, however, that (1) it is a standard for information processing systems (eg, not just for instrument recording), and (2) it is a complete standard that includes not only media but recording formats on the media itself, and (3) the standard is in fact in use by multiple vendors.

11. Supply Competition, Availability, Service for Drive and Media

All else being equal, it is better to have a product in which there is competition in the marketplace among different vendors for both the drive and the media. This introduces competitive pricing, complementary products, as well as avoiding dependency on a single source company. Remember that a product that is great, but new, also will not have significant market acceptance yet, and if that acceptance never develops, the vendor may be forced to drop support for the product regardless of the technical merits of the product.

Current Technology Experiences

Half-Inch 9-Track

The traditional 2400 foot by one-half inch tape reel had 9 tracks running longitudinal along the tape using fixed heads. The longitudinal density was initially 800 bits per inch but current use is generally 1600 or 6250 bits per inch (bpi).

3480/3490

The 3480 cartridge was introduced by IBM in the early 1980's, had 18 tracks and can store 200MB per cartridge at a rate of 2-3MB/Second.

Recent enhancements on the technology have included data compression (IDRC) which increased throughput and cartridge storage capacity from 1.5 to 5 times, depending on the data, and 36 track bi-directional (3490) for double storage capacity.

Future development in thinner tape could allow a tape of double length to be used to double capacity again.

Adding all these together, one can expect the traditional 200MB cartridge to store an average of 2400 MB.

QIC

The Quarter-Inch cartridge is generally a lower performer than other technologies and is intended for the PC/Workstation backup market. There are a wide variety of products which are often incompatible. Developments are under way, especially by 3M, which target 100GB/cartridge by 1999 which may make this technology more of a contender in the future.[10]

8mm

8mm involves the Sony helical tape transport to which Exabyte has sole rights for computer data storage use within the United States. After Exabyte adds the circuitry to connect the drive to SCSI, various vendors re-sell the drives as-is, with driver software, and/or with added controller hardware to connect the drive to another interface.

The original device, called the 8200, has throughput capabilities of 220KB/Sec and has a tape cartridge capacity of 2.3GB. Notable disadvantages are the slow throughput and the slow file search time of only 10 times the nominal speed. Thus, it can take 3 hours to write a tape, and afterward it could take 20 minutes to find a single file if it is located near the end of the tape.

Current augmentations include the 8200SX, which incorporates a file search capability of 75 times nominal speed (which is the full rewind speed). Another augmentation is the 8500 model, which is purported to have 5GB per tape capacity with double transfer rates (440KB/Sec), and also be compatible with tapes written on the 8200 model.

Future augmentations include data compression. Should the standard 3:1 ratio be achieved, then the achievable tape capacity could be 15GB with transfer rates of 1200KB/Sec.

Exabyte is also conducting development projects for further enhancing the capacity of the media.

4mm DAT (Digital Audio Tape)

Traditional product: 1.3GB on one 60 meter tape at 180KB/Sec. File search time has always been faster than 8mm and runs at about 200 times nominal speed.

Multiple vendors support the DAT format and have converged on the DDS (Digital Data Storage) format (to the near demise of the DATA/DAT format).[11] Some care still needs to be made in insuring compatibility between different models of DDS 4mm drives.

Current augmentations include transports designed for the start/stop and fast file search needs of computing (rather than the long-play needs of audio, which is where the technology was borrowed from). Now products are appearing with data compression and/or supporting 90 meter tapes that purport to hold 8GB of data with an effective throughput of 733KB/Sec.

Although 8mm drives were available before 4mm drives, and are more popular today, several industry analysts expect 4mm to beat out 8mm in the long run.[12] [13]

CD-ROM

This technology has a relatively high cost for off-site writing of the first copy from another media, but significantly decreased cost in making copies. Thus, the typical markets here are for large quantity copy operations such as software distribution or data bases. Drives conform to a standard of 650MB/Disk with 150KB/Sec transfer rate.[14]

Optical WORM Disks

WORM was the second optical technology to be introduced, but enough vendors produced the product early enough that there is no standard.[15]

In WORM technology, one laser burns holes in the media surface exposing a non-reflective sub-layer. A second, less powerful, laser then reads the reflectivity of the spot, ie, is it reflective (not burned out) or is it not reflective (ie, burned out).

This device has the definite advantage of random access while still being cheaper than magnetic disk drives. However, their performance is lower than the magnetic disk as they generally rate between 500-1000KB/sec reading and half that for writing.

WORM's biggest disadvantage is its write *once* characteristic.

Optical library units are now generally available to provide automated access to many platters without requiring one drive for each platter or operator mounts.

Future augmentations expected are drives that can access both sides of the platter without manual flipping of the cartridge and dual function (WORM and erasable) drives.[16] Data compression research is under way to provide higher data densities and higher *effective* data transfer rates.

Optical, Rewritable Disks

Rewritable optical drives provide the re-usability characteristic of most media. However their performance remains at WORM levels of 500-1000KB/sec reading and 250-500KB/sec writing (due to the three-pass Erase/Write/Verify procedure). While throughput is expected to increase, it is expected to remain slower than magnetic technologies. The ANSI/ISO standard 5 1/4" disk hold 650 MBytes although there are plenty of drives in production that do not support any standard.

The older method of recording, referred to as Magneto-Optical, involves heating up the media and applying a magnetic field to change the magnetic orientation of the media. This method has proven media life, although one drawback is that re-writes involve a two step process, one to erase, the second to write. The newer technology, refereed to as Phase-Change, involves changing the surface from a dull to a reflective crystal structure, which involves only one pass to re-write, but has not proven itself to the large number of rewrites that Magneto-Optical has.[17]

Some devices now support almost double capacities by increasing the number of sectors per track as the tracks get closer to the outer edge, thus keeping the actual bit densities near constant throughout the disk.[18] Unfortunately these disks do not satisfy an ANSI/ISO standard.

Future augmentations are the same as listed under WORM.

VHS

Some drives are available which will write to standard VHS casettes. Metrum (formerly Honeywell) introduced a unit which stored 5 GB on one cartridge a few years ago and has upgraded it to 10.4 GB and now 14.5 GB at a rate of 2MB/second, sustained with a file search time of 90 seconds or less. Metrum expects to reach 12-15MB/second with 100GB/cartridge within a few years.

Magnetic Disk

This media should not be ruled out as 1GB SCSI or ESDI disks are now common and 2GB disks are now making it to market. The number of MB in a square inch has increased from .02MB in 1970 to 60MB in shippable products today to over 1GB in the laboratories.

Futures involve increased capacity at lower cost and faster spinning disks to increase throughput somewhat, but the best gains will be from efficiently cached disks because the mechanical delays will remain the biggest obstacle.

RAID (Redundant Array of Inexpensive Disks) disks are also being developed to raise the reliability and performance of the 5.25" SCSI disks to those expected from high-end systems. By grouping multiple SCSI disks together to make them look like one big disk, *striping* achieves higher throughput by reading data off one disk while another is moving its actuator. Redundancy, and, thus, the reliability, is achieved by using error correction codes in such a manner that if any one disk in the group were to fail, the disk could be replaced and the data re-constructed from the remaining disks while the disk set remains live.[19]

Future Technology Expectations

Helical Scan

Helical scan technology, in which heads which rotate at an angle to the media read and write angled tracks onto the tape, was borrowed from the video market and adapted for the computer industry. 8mm and 4mm, from the consumer video and audio markets, have proven themselves as high-capacity drives, but with high error rate, and slow speeds.

19mm drives (both DD1 and DD2), from the broadcast industry, are now emerging as higher capacity, faster, more reliable, and also more expensive drives. The "smaller" D2 cartridge, which is roughly the size of a VHS cartridge, holds 25GB of data, while the "Medium" and "Large" sizes hold more. Current players are Ampex, Sony, and Datatape. Sony's model has an advantage in that they intend to sell multiple models, at varying prices, that have full media interchangeability.

Future developments in helical scan technology are expected from storage manufacturers who are designing helical scan devices, from the ground up, for the computer industry.

Optical Tape

Optical tape is a WORM technology based on a plastic and metal sandwich developed by ICI Imagedata and named “Digital Paper” and which Dow has now developed their own product which is partially compatible.[20] The major differences with the traditional optical disk technology is that “Digital Paper” involves altering the polymer coating rather than the metal layer, resulting in higher densities and transfer rates, as well as the fact that the media is flexible and, thus, can be used in various formats, including tape.

Creo developed an optical tape product based on Digital Paper but it is targeted at the high end market as the drive cost is prohibitive for many as the drive costs roughly $250,000.

The optical tape product under development by LaserTape will provide 50GB of WORM data on one 3480 size cartridge with 541 feet of optical tape, at a rate of 200Mb/inch^2 with SCSI and SCSI-2 interfaces at up to 3MB/second with autoloader capabilities. The cost is estimated at $0.005 per MB. Seek time is expected to be 15 seconds over the entire 50GB. However, this product is not expected out until late 1992 or early 1993.

Rough estimates on the cost of the drive is $25,000 with media costing $250 each.[21]

Magnetic Media

The media itself will change as Metal Particle tape is being replaced by Metal Evaporative tape, as well as barium-ferrite and cobalt-chrome work underway at 3M. Each new technology allows for a higher bit density capability.

Evaluation

At the risk of putting down in one table lots of data which can be easily disputed, I list here my understanding of the situation as it exists today.[22] [23] [24]

Ratings are H=High (ideal device), M=Medium, and L=Low (undesired attribute of the device). These ratings are only relative, not linear (ie, a rating of M does not mean the device is twice as good as a rating of L, nor is an H twice as good as an M). Granularity is intentionally kept to a minimum.

Please note that the figure which graphs the cost of throughput (1 KB/Sec) and media (1 GB) has logarithmic scales.

Key for Table

Criteria:	H=Ideal product
1. Availability	H=Product shipping non-beta units for 6 months or longer
2. Drive Related Costs	
A. Drive Unit Cost	H=Low cost
B. Throughput	H=Fast Throughput
C. Cost/KB/Sec	H=Low Cost/KB/Sec
3. Media Related Costs	
A. Media Unit Cost	H=Low media unit cost
B. Capacity	H=High Capacity
C. Cost/GB	H=Low Cost/GB
4. Functional Requirements	
A. Sequential vs Random	S=Sequential only, R=Random
B. Writability	W=Writable (and re-writable), WO=Writ only, R=Read only
C. Search Speed	H=Fast file search capabilities
5. Data Interchange and Compatibility	H=Well-used standard media
6. Robotic Options	H=Multiple complementary offerings by multiple vendors
7. Data Compression Options	H=Data Compression is available
8. Level of Integration Within Your Environment	H=Standard interfaces
9. Robustness	
A. MTBF	H=High MTBF with reasonable duty cycle
B. Error Rate	H=Low uncorrectable error rates
C. Passes	H=Many (ie, millions) of passes
D. Shelf Life	H=Long shelf life (ie, 20 years or more)
10. Standards and Competition	H=Recognized by ANSI and/or IEEE and used by many vendors
11. Supply Competition, Availability, Service	H=High availability for Drive, Media, and Service

Table: Evaluation Chart

Device	1 Time Avail	2a Cost of Drive	2b Speed KB/Sec	2c Cost of 1KB/Sec	3a Cost of Media	3b Capacity of Media	3c Cost of 1GB	1	2a	2b	2c	3a	3b	3c	4a	4b	4c	5	6	7	8	9a	9b	9c	9d	10	11
800 BPI 9Trk	1950	4000	400	10.00	7	20	350.00	H	L	L	L	H	L	L	S	W	L	L	L	L	H	H	M	M	M	M	M
1600 BPI 9Trk	1960	6000	800	7.50	8	40	200.00	H	L	L	M	H	L	L	S	W	L	H	M	L	H	H	M	M	M	H	H
6250 BPI 9Trk	1970	15000	1000	15.00	8	160	50.00	H	M	M	L	H	L	L	S	W	M	H	M	L	H	H	M	M	M	H	H
3480	1981	20000	3000	6.67	5	200	25.00	H	M	H	M	H	L	L	S	W	H	H	H	L	H	H	H	H	M	H	M
QIC	1986	1000	240	4.17	12	600	20.00	H	L	L	H	M	L	L	S	W	L	M	M	L	M	L	L	L	M	M	M
8mm 8200	1987	2000	246	8.13	5	2300	2.17	H	L	L	M	H	M	M	S	W	L	M	M	L	H	L	L	L	M	L	L
3490	1987.2	50000	3000	16.67	5	400	12.50	H	H	M	L	H	L	L	S	W	H	M	H	L	H	H	H	H	M	H	M
4mm	1988	2000	183	10.93	7	1300	5.38	H	L	L	L	H	M	M	S	W	M	M	H	L	H	L	L	L	M	M	M
VHS VLDS	1989	40000	1000	40.00	20	5000	4.00		H	M	L	L	H	M	S		H	L	M	L	M	H	H	M	M	L	L
4mm DC	1990	3500	550	6.36	7	4000	1.75	M	L	L	M	H	M	M	S	W	M	L	H	H	H	L	M	L	M	M	M
3490E	1991	75000	9000	8.33	7	1200	5.83	M	H	H	M	H	M	M	S	W	H	M	H	H	H	H	H	H	M	M	L
VHS RSP	1991	40000	2000	20.00	22	14500	1.52	M	H	M	L	L	H	M	S	W	H	L	M	L	M	H	H	H	M	L	L
8mm 8500	1991	2500	440	5.68	7	5000	1.40	L	L	L	M	H	H	H	S	W	M	L	H	L	H	M	M	L	M	L	L
19mm Sony	1991	125000	32000	3.91	35	25000	1.40	L	H	H	M	L	H	H	S	W	H	L	M	L	M	H	H	H	M	L	L
19mm Ampex	1991.4	125000	15000	8.33	35	25000	1.40	L	H	H	M	L	H	H	S	W	H	L	M	L	H	H	H	H	M	L	L
8mm 8500 DC	1992.1	3500	1200	2.92	7	15000	0.47	L	L	M	H	H	H	H	S	W	M	L	H	H	H	M	M	L	M	L	L
3490E+thin	1992.2	90000	9000	10.00	8	2400	3.33	L	H	H	L	H	M	M	S	W	H	M	H	H	H	H	H	H	M	M	L
Optical Tape	1992.3	25000	3000	8.33	25	50000	0.50	L	M	H	M	L	H	H	S	WO	H	L	M	L	M	M	M	H	H	L	L
Future Helical	1993.4	30000	12000	2.50	25	25000	1.00	L	M	H	H	L	H	H	S	W	H	L	M	L	M	H	H	H	M	L	L
VHS Future	1994	60000	12000	5.00	25	100000	0.25	L	H	H	H	L	H	H	S		H	L	M	L	M	H	H	H	M	L	L
Future LaserTape	1998	50000	12000	4.17	20	100000	0.20	L	H	H	H	L	H	H	S		H	L	M	L	M	H	H	H	H	L	L
Future 8mm	1999	5000	6000	0.83	20	65000	0.31	L	L	H	H	L	H	H	S	W	H	L	H	H	H	H	M	M	M	L	L
Future QIC	1999	5000	1000	5.00	20	100000	0.20	L	L	M	H	L	H	H	S	W	H	M	H	H	H	M	M	M	M	M	M
Magnetic Disk		3000	1500	2.00	3000	1000	3000.00	H	L	H	H	L	M	L	R	W	na	na	na	N	H	H	H	H	M	na	H
CD-ROM	1987	1000	150	6.67	7	650	10.77	H	L	L	M	H	L	L	R	R	na	H	H	N	H	H	H	H	H	H	H
WORM Disk	1987.2	50000	800	62.50	600	700	857.14	H	H	L	L	L	L	L	R	WO	na	M	H	N	M	H	H	H	H	M	H
Rewritable	1989	60000	1000	60.00	600	800	750.00	H	H	M	L	L	L	L	R	W	na	M	H	N	M	H	H	H	H	M	H
Dual function	1991	75000	1000	75.00	600	800	750.00	M	H	M	L	L	L	L	R	WO	na	M	H	N	M	H	H	H	H	M	M

Shaded areas are less accurate due to dependancies on pricing, vendor, compression, configuration, and future product uncertainty or merely unconfirmed information.

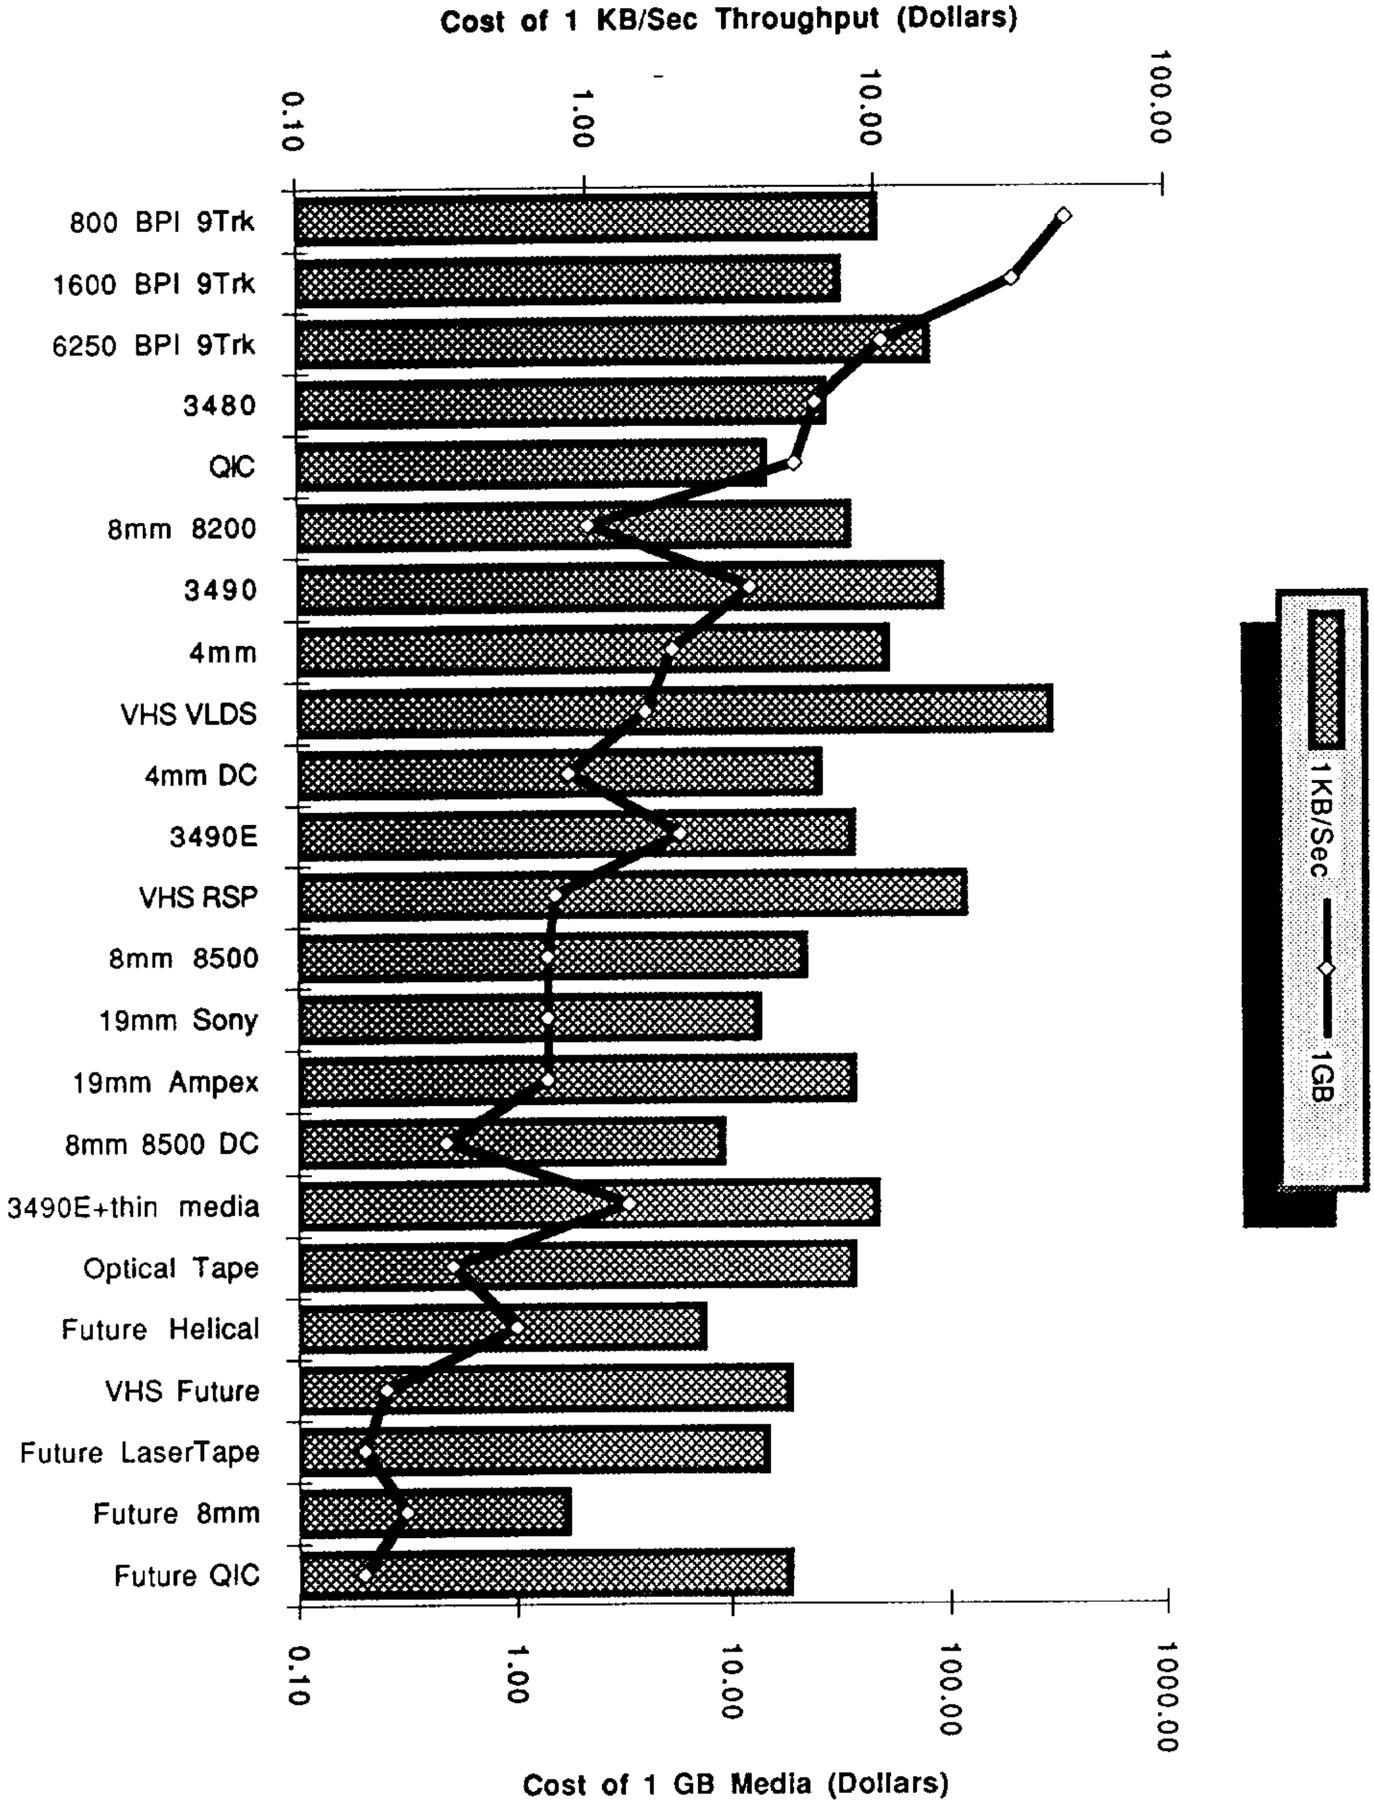

Figure: Cost of Throughput and Media for Various Technologies

Conclusion

In my opinion, IBM seems to be relying on relatively small increases to its 3480 and 3490 technology for the next five years. This will not satisfy the needs for much of high energy physics, as well as others requiring high capacities. Thus, a vacuum will result in which storage vendors will be forced to produce non-IBM-like drives. Whether or not these drives become accepted by the market will depend on:

1. Past performance of the company, and its installed base.
2. If the product is truly a technically superior product in the market in areas such as cost/GB or cost/MB/sec.
3. If the product can be accepted as a standard (note that to be approved as a standard is one thing, but to have multiple vendors accept the standard and *use* it in competitive and/or complementary products is critical).

The technologies I believe are worth watching are:

1. 8mm with its dual density 8500 coming out now and future developments, with a target of remaining a low-cost drive.
2. Other helical scan technology, such as 19mm by Ampex or Sony, or new helical formats under development by other organizations including Storage Technology Corporation.which plans to have 20 to 25 GigaBytes on one tape cartridge, raising their 6,000 volume silo unit from 1 TeraByte to 120 TeraBytes by 1994[25] [26]. Developments in VHS storage also are of interest.
3. Optical tape which promises to have high data capacity, and throughput.
4. Robotics is an important factor in considering any device as *sideline* data, which is cheap yet accessible without operator intervention.
5. Software that will make the data available to the user in an easy and timely manner. This includes simple catalogs for robotic library units as well as Hierarchical Storage Management (HSM) systems, hopefully with multiple server locations throughout the network. Also important in providing this data to a more distributed environment is intelligent network software which is capable of caching data locally to reduce the load on the network.
6. Magnetic media chemical composition changes which will allow for higher bit density on the media.

Thus, it is clear that data sotorage technology is neither a simple topic, nor is it a static one. It is important for anyone intending to store large volumes of data to familiarize themselves with the technologies and to continually review them.

References

[1]ANSI document X3B5/91-225A, "Technical Committee X3B5 Digital Magnetic Tape Annual Report July 1, 1990 - June 30, 1991", Richard Steinbrenner, Chair, June 30, 1991.

[2]"Recent Experiences and Future Expectations in Data Storage Technology" By Jack Pfister, talk presented at the 1990 Conference on Computing in High Energy Physics, Fermilab paper FERMILAB-Conf-90/70.

[3]"Hierarchical Storage Carves Out Signpost To Future" by Ted Smalley-Bowen, Digital Review, July 8, 1991, p 41.

[4]"Mass Storage System Reference Model: Version 4," Edited by Sam Coleman and Steve Miller, May 1990.

[5]"Tape Libraries--Good Reading For The Masses" by Mark Ferelli, Computer Technology Review, Sep 1991, Vol XI, No 11, p 1.

[6]"Automated Tape Libraries Are Impacting The Mass Storage Marketplace" by Fred Zeiler, Computer Technology Review, Summer 1991, Vol XI, No 9, p 95.

[7]"Data Compression Breaks Through To Disk Memory Technology" by Robert A. Monsour and Douglas L. Whiting, Ph.D., Computer Technology Review, Spring 1991, Vol XI, No 6.

[8]"DCLZ Emerges As An Open Data Compression Standard" by Mike Franek, Computer Technology Review, Summer 1991, Vol XI, No 9, p 89.

[9]"Top-Rated Tape Technologies Put To Proof" by Jim Tierney, Computer Technology Review, May 1991, Vol XI, No 5, p 18.

[10]"Tape Drive Firms Shift to Ultra-Dens Recording" by Stuart Zipper, Electronic News, Sep 23, 1991, Vol 37, No 1879, p 16.

[11]"Will Interchangeability Derail DAT?" by Mark Ferelli, Computer Technology Review, July 1991, Vol XI, No 9, p 1.

[12]"DAT expected to bypass 8mm tape rival" by Ted Smalley Bowen, Digital Review, September 23, 1991, p 3, based on Freeman Associates report "Computer tape outlook: Helical scan products".

[13]"DAT May Take The Lead," Digital News, May 13, 1991, p 5, based on a report by Data Storage Concepts of Boulder Colorado.

[14]"CD-ROM, by David Harvey, Byte, Sep 1991, p268.

[15]"WORM Recording Keeps Technology Green" by Ken Goodnow and Karen Walker, Computer Technology Review, Sep 1991, Vol XI, No 11, p 8.

[16]"Optical Drives' New Reading Habits" by Dwight B. Davis, Datamation, Sep 15, 1991, p 71.

[17]"Optical Cauldron Is Bubbling Hot" by Mark Ferelli, Computer Technology Review, July 1991, Vol XI, No 8.

[18]"ZCAV Format Delivers One Gbyte Capacity Magneto Optical Drives" by John Freeman, Computer Technology Review, July 1991, Vol XI, No 9, p61-64.

[19]"Arrays More For Bandwidth Than Capacity" by Mark Ferelli, Computer Technology Review, May 1991, Vol XI, No 5, p 20.

[20]"Flexible Optical Media Boost Data Density" by Chris Terry, EDN, April 25, 1991, p 77.

[21]"Tape Backup with an Optical Twist," Digital Review, July 22, 1991, p 50.

[22]"Tape Backup: Finding The Magic in the Medium" by Ted Smalley Bowen with accompanying table based on information from Freeman Associates, Digital Review, July 22, 1991, p 47.

[23]Data Sources: Hardware, Volume 11, No 1, 2nd Edition 1991.

[24]Digital Storage Technology Handbook, Digital Equipment Corp, 1989.

[25]"Storage Tek's Storage Plans" by Thomas Hoffman, Information Week, October 21, 1991, p 57.

[26]"StorageTek to Enter Helical Scan," Electronic News, October 14, 1991, p 13.