

AD-A163 453

OPTICAL DISK TECHNOLOGY FOR LARGE SCALE MASS STORAGE
(U) ROME AIR DEVELOPMENT CENTER GRIFFISS AFB NY
J D PETRUZZELLI DEC 85 RADC-TM-85-17

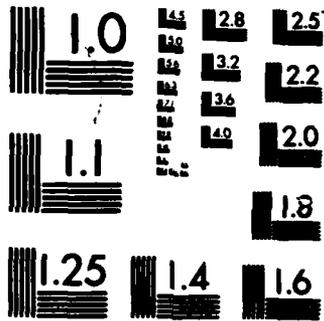
1/1

UNCLASSIFIED

F/G 9/2

NL

							END						
							FILED						
							OPC						



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

12

AD-A163 453

**RADC-TM-85-17
In-House Report
December 1985**



***OPTICAL DISK TECHNOLOGY FOR LARGE
SCALE MASS STORAGE***

Jack D. Petruzelli

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

DTIC FILE COPY

**DTIC
ELECTE
JAN 28 1986
S E D**

**ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700**

86 1 27 102

This report has been reviewed by the RADC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TM-85-17 has been reviewed and is approved for publication.

APPROVED:



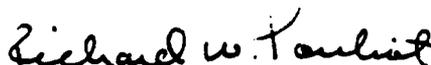
THADEUS J. DOMURAT
Chief, Signal Intelligence Branch
Intelligence & Reconnaissance Division

APPROVED:



WALTER J. SENUS
Technical Director
Intelligence & Reconnaissance Division

FOR THE COMMANDER:



RICHARD W. POULIOT
Plans Office

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

AD-A163453

REPORT DOCUMENTATION PAGE				
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS N/A		
2a. SECURITY CLASSIFICATION AUTHORITY N/A		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A		4. PERFORMING ORGANIZATION REPORT NUMBER(S) RADC-TM-85-17		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) RADC-TM-85-17		5. MONITORING ORGANIZATION REPORT NUMBER(S) N/A		
6a. NAME OF PERFORMING ORGANIZATION Rome Air Development Center		6b. OFFICE SYMBOL (If applicable) IRAP	7a. NAME OF MONITORING ORGANIZATION Rome Air Development Center (IRAP)	
6c. ADDRESS (City, State, and ZIP Code) Griffiss AFB NY 13441-5700		7b. ADDRESS (City, State, and ZIP Code) Griffiss AFB NY 13441-5700		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Rome Air Development Center		8b. OFFICE SYMBOL (If applicable) IRAP	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N/A	
8c. ADDRESS (City, State, and ZIP Code) Griffiss AFB NY 13441-5700		10. SOURCE OF FUNDING NUMBERS		
		PROGRAM ELEMENT NO. 62702F	PROJECT NO. 4594	TASK NO. 15
				WORK UNIT ACCESSION NO. TK
11. TITLE (Include Security Classification) OPTICAL DISK TECHNOLOGY FOR LARGE SCALE MASS STORAGE				
12. PERSONAL AUTHOR(S) Jack D. Petruzelli				
13a. TYPE OF REPORT In-House		13b. TIME COVERED FROM Jan 85 TO Apr 85	14. DATE OF REPORT (Year, Month, Day) December 1985	15. PAGE COUNT 36
16. SUPPLEMENTARY NOTATION None				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	Optical Disk	
09	02		Digital Optical Disk	
09	03		Mass Storage	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Optical disk has been under intense investigation as a digital mass storage mechanism since the mid 1970's. Since that time, even before potential manufacturers had systems in hand, the technology has been touted as the panacea to cure all the supposed ills of magnetic based storage mechanisms. It is very important to gain an understanding of digital optical disk and what it can or cannot do in a particular application scenario before deciding if it is the answer to a storage problem. This report attempts to acquaint the potential user with optical disk technology. In addition, it briefly discusses a few of the techniques employed to overcome some of the perceived shortcomings. <i>References included</i>				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL JACK D. PETRUZELLI		22b. TELEPHONE (Include Area Code) (315) 330-4581	22c. OFFICE SYMBOL RADC (IRAP)	

UNCLASSIFIED

Block 17 (Cont'd)

Cosati Codes

Field	Group
09	04
14	03
15	04

UNCLASSIFIED

ACKNOWLEDGEMENT

This paper was originally submitted at the request of the National Center for Atmospheric Research of the University Corporation for Atmospheric Research (National Science Foundation). The National Center for Atmospheric Research is operated by the University Corporation for Atmospheric Research and is sponsored by the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation.

Accession For	
NCIS GRA&I	<input checked="" type="checkbox"/>
ERIC TRS	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Identification	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



TABLE OF CONTENTS

	Page
Introduction	1
Background	2
Data Density	4
Data Encoding Techniques	7
Error Management	9
Media Archival Characteristics	14
Advantages of Non-Rewrite Media	16
Disadvantages of Write Once Media	18
Formatted Versus Unformatted Media	20
Other Issues	22
Summary	23
Conclusion	26

INTRODUCTION

Optical disk has been under intense investigation as a digital mass storage mechanism since the mid 1970's. Since that time, even before potential manufacturers had systems in hand, the technology has been touted as the panacea to cure all the supposed ills of magnetic-based storage mechanisms. It is very important to gain an understanding of digital optical disk and what it can or cannot do in a particular application scenario before deciding if it is the answer to a storage problem. This report will attempt to acquaint the potential user with optical disk technology. In addition, it will briefly discuss a few of the techniques employed to overcome some of its perceived shortcomings. It is hoped this will allow the reader to gain sufficient insight so that well-informed decisions can be formulated.

BACKGROUND

The techniques to employ optical disk as a digital storage medium are rooted in the video industry. Serious investigations into using a disk format for the storage, dissemination and retrieval of video program material began in the early 1970's at laboratories such as RCA, 3M, Xerox, Kodak, Philips and various Japanese firms. The basic goals were to devise a very high density, inexpensive, reliable, easily replicated method for the storage of video program material such as full-length motion pictures, concerts, instructional material, etc.

These substantial investments in both media and systems developments resulted in a very large array of product introductions. Unfortunately, these products were brought to the marketplace coincident with the advent of the home video tape recorder. The home video tape machine offered a lower cost, more reliable alternative to the video disk systems. In addition, and most importantly, they offered a media that could be recorded and reused by the consumer. By contrast, the video disk systems were read-only.

Sales of the video tape machines immediately outstripped those of the disk systems. As a result, there are very few video disk systems for sale to the consumer today. However, there is a substantial institutional and instructional market for custom video disk program material, and industrial grade players. At least one manufacturer, 3M, operates a very successful video disk mastering facility boasting an 8-hour turnaround for small quantities of fully certified video disk program material.

As a result of this unexpected competition from the video tape recorder, many would-be video disk manufacturers either abandoned their consumer projects, switched to concentrate exclusively on the institutional market or shifted emphasis to digital data systems. The reasons for the shift from video to digital were compelling. The optical disk had many potential advantages as a digital storage and retrieval device. A single video disk had approximately 54,000 user addressable tracks per side on a 12 inch diameter format. Digital storage on such a device could easily

approach 1 to 2 gigabytes of user data with some 14 inch diameter formats providing between 4 and 10 gigabytes on a single removable platter. The then current magnetic disk drives offered 300 megabytes on 9 platters (18 surfaces) of non-removable media. However, there were disadvantages with a digital system. It had to be able to write as well as read data as opposed to the read-only function of the video systems. The increased complexity and performance inherent in such designs had to be balanced against system cost, particularly since they had to compete in the marketplace against conventional magnetic storage peripherals. Also, the devices had to have very good random access times that were not necessary in a purely video device. Typical consumer video disk players had no random access capability, only a serial, frame-by-frame search mode. The industrial grade players had a random search capability giving access to any frame in seconds. In order to be viable, a digital device required access times of tens of milliseconds. In addition, there were and still are no standards for digital data disk to follow in terms of disk size, center hole size, rotation rates, data formats, sector size, encoding scheme, media interchangeability etc.

Finally, but most importantly, the digital optical disk must possess extremely good error detection and correction schemes to compensate for media defects. Such defects can occur both at the time the data is written and can continue to propagate as the media ages. If digital optical disk is to compete with the computer peripheral industry, it must demonstrate data reliability of 1 bit in 10^{12} bits detectable and correctable and better than 1 in 10^{13} hard error rate. This performance must be obtained from a media exhibiting a raw Bit Error Rate (BER) in the 10^{-5} range.

DATA DENSITY

One of the most impressive and often quoted facts about digital optical disk is the currently demonstrated areal density. Bit densities of 10^8 bits per square inch are commonly recognized as an achievable production number. By contrast, an IBM 3380 magnetic disk drive possesses bit areal densities in the 10^7 bits per square inch regime.

In order to achieve packing densities of 10^8 bits per square inch, optical disk systems designers rely heavily on the ability to precisely and repetitively locate extremely small features (spots) on the surface of the disk. Typically optical disk systems fall into two categories based upon the choice of the laser source employed. The highest capacity and highest data transfer rate systems, those demonstrating areal densities well in excess of 10^8 bits per square inch and transfer rates in excess of 4 megabytes per second are largely based upon the use of Argon gas lasers. Such a system can form diffraction limited spots as small as 0.4 microns in diameter on 1.0 micron centers with track-to-track spacing, or pitch, of 1.25 microns. The second class of systems, and the scheme most commonly employed in the industry, use semiconductor laser diodes as the light source. The major difference between the Argon and the laser diode systems is data capacity. Since the wavelength of the laser diode is twice that of Argon, the smallest achievable spots are also twice that of Argon, in the 0.8 to 1.0 micron range. Consequently, density in both the along-track and cross-track directions suffer by approximately 2:1. This means that an Argon-based system can provide a maximum of 10 gigabytes of storage per disk while a similar diode-based system would provide 4.0 gigabytes per disk.

While Argon systems offer maximum capacity and data rate, they suffer from some important drawbacks. An Argon laser typically has an optical head 3 to 4 feet long. With its attendant power supply, an additional 1 to 2 cubic feet, the laser system alone weight 50 to 100 pounds. Such lasers require several kilowatts of 3-Phase electrical service and 2 to 5 gallons per minute of cooling water. Cost is another factor. They usually are in the \$5 to 10 thousand range. These factors relegate Argon systems to use in

relatively large, very high transfer rate and large capacity systems where its cost, size and weight are but a fraction of the total system.

The greater majority of optical disk systems are designed around laser diode sources. They are extremely small, are relatively inexpensive, consume just a few watts of power and rely on ambient cooling. Aside from their 2:1 disadvantage to Argon in spot size, the only real drawback to their use is that they have limited light output power. Consequently, system data transfer rate, disk sensitivity and optical system efficiency must be carefully designed to operate effectively within the constraints of the light power available. Other potential problems from the use of diodes are their lifetime, long term pointing stability and mode stability. Nonetheless, the use of laser diode sources gives the system designer the opportunity to arrive at a relatively compact, inexpensive design that can compare favorably to conventional magnetic peripherals in all respects.

While the single order of magnitude density difference between the best optical and magnetic systems is not impressive standing alone, it translates to a substantial volumetric data density difference. For instance, an IBM 3380 consists of 2 non-removable, multiple platter (10 platters each) stacks. Each stack holds 2×10^{10} user bits (2.5 gigabytes). The stacks are 14 inches in diameter and are approximately 12 inches tall for a volume displacement of approximately 1850 cubic inches. The stacks are non-removable. By contrast, a Storage Technology Corporation (STC) data cartridge for the Model 7640 Optical Storage Unit consists of a single disk platter housed in a cartridge approximately 15 x 15 x 0.5 inches (approximately 112.5 cubic inches) and holds 4.0 gigabytes (3.2×10^{10} user bits). In addition, it is a removable cartridge. Put in other terms to store 10 gigabytes of data would require 2.14 cubic feet in the IBM case and 0.16 cubic feet for STC, FOR THE MEDIA ALONE. This is complicated by the fact that the IBM media is not removable; the STC is. This difference in volumetric efficiency provides very strong impetus for the utilization of digital optical disk as a mass storage device. From a floor space standpoint, the 3380 requires 100 square feet of floor space to store 10

gigabytes on-line, while an STC 7640 would require approximately 70 square feet, including service area. Also, since the optical media is removable, it facilitates automated handling for off-line storage of much larger amounts of data. This option is unavailable to the IBM user. To store larger amounts requires the use of additional components and/or systems, or other types of storage media, such as the IBM 3480 tape system.

NOTE

In the previous example, STC was chosen as representing the highest performance optical disk unit that is in or near commercial availability. In light of the recent Chapter 11 proceedings at STC, the future of their optical storage product is in doubt. Nonetheless, the STC 7640 remains the highest performance commercial development to date.

If one were to consider special developments, such as those being put forth by RCA Corp., Moorestown, N.J., a single 14 inch diameter optical disk platter holds a very impressive 10 gigabytes (8×10^{10}) user bits. While not currently in commercial development, this technology has been demonstrated to the U.S. Government and has resulted in the delivery of at least two test bed mass storage systems dubbed "Jukeboxes". These systems offer on-line access to 1.25 terabytes through the use of robot accessing mechanisms. This allows automated access to any of 128 optical disks in 6 to 10 seconds in a footprint of approximately 30 square feet. Off-line access to even larger amounts of data is facilitated either automatically or manually. The point to be made is technology has demonstrated the ability to increase the volumetric efficiency of digital optical disk. Although, mitigating factors such as lack of a commercial development base, nonrecurring development costs, custom hardware, specific interfaces and custom software device handlers must be tolerated.

DATA ENCODING TECHNIQUES

A major factor in the ability of optical disk designers to pack data so tightly on the surface of a platter is the sub-micron sized features created by the laser beam. However, this is not the only technique used. In all successful digital recording and storage schemes, whether optical or magnetic, disk or tape, some form of sophisticated data encoding algorithm is employed. They have names such as Randomized-Non-Return-to-Zero (RNRZ), Enhanced-Non-Return-to-Zero (ENRZ), 8, 10 Group Code invented by Independent Broadcast Authority, known simply as IBA, and Run Length Limited Codes (RLL) such as 3-Phase, 5-Phase, 7-Phase, and finally, Bi-Phase and Delay Modulation (DM), known also as Miller Code. There is also a further derivation of Miller known as Miller Squared (M^2). This is by no means an all inclusive list but merely some of the more popular coding algorithms employed by the industry. Of the codes mentioned above, the two that seem to enjoy the greatest popularity among optical disk manufacturers are Delay Modulation, or Miller, and 3-Phase.

When selecting a code for a particular application, a system designer obviously looks for many things that are system dependent. Among the characteristics considered are: its minimum and maximum run length, required channel bandwidth for a given data rate, the size of the decision window in terms of fraction of bit periods or nanoseconds, whether it is self-clocking and/or self-synchronizing code, the rate conversion factor and whether there is a DC component to the code. Finally, what effects the code selection will have on overall system Bit Error Rate and if it will work effectively with the error detection and correction scheme selected.

Encoding provides a certain measure of recording efficiency in that each code offers the ability to effectively place more than a single user bit per recorded feature. The measure of this efficiency varies from code to code. The frequency content of the various codes can limit the required system frequency response for a specific digital bandwidth. This can both limit the digital data rate as well as significantly affect circuit design criteria and consequently system cost.

Whether or not a particular code is self-clocking and/or self-synchronizing can significantly reduce system overhead. A self clocking code does not require the introduction of a separate timing/synch track or surface to sustain the system. The amount of DC content a code possesses will limit the string length of a succession of 1's or 0's, or other baseline shifting patterns that can occur before the system begins to pick up a DC bias or offset. This can result in decision errors that propagate through the system. Usually, a code with little or no DC content is preferred. The choice of a given code seriously affects overall parts count in the system. Implementation of certain codes can vary the required semiconductor parts count by as much as 30 percent. The effect on overall system complexity and cost rises accordingly. However, there are situations where the advantages to the use of such a code outweigh the negative factors.

It would require volumes to treat the topic of code selection and to give complete derivations of the coding rules. It is sufficient to say that the various encoding techniques provide the system's designer with several advantages.

ERROR MANAGEMENT

The strategy for handling errors in digital optical disk systems is very similar to that used in the magnetic industry. There are some basic differences since most of the optical systems must assume the use of a non-rewriteable media. Obviously, the first step in controlling errors is characterization of the system. The word system here refers to the hardware, software and the recording medium. Of these, the medium is the hardest to fully characterize due to its variable nature. It is not the intention to dismiss the characterization of the hardware and software for error management purposes. The point being that a good system designer can prove both in design and implementation phases that the hardware and software function as specified from an error budget standpoint. The media, on the other hand, is subject to a large number of variables that are outside the control of the system designer.

The most difficult step in managing the error budget in an optical disk system is obtaining a high quality media. An optical disk consists of a rigid or semi-rigid substrate such as aluminum, glass, plexiglass or mylar, with a series of extremely thin layers, in some cases less than 50 nanometers thick, built up in a series of individual processing steps. The composition of these layers vary from metals to polymers to dyes. The coating methods range from RF sputtering to spin coating. They must all reside in intimate and benign contact with one another if the media formulation is to succeed. The physics, chemistry, metallurgy and in some cases, the magic that goes into this process has utilized millions of dollars and is the subject of volumes of technical information. The ultimate goal is to provide a defect-free, uniform, stable, durable and reproducible platter within such systems constraints as available laser power, spot size, data rate and laser wavelength. In addition, it must be able to survive certain environmental impacts such as temperature and humidity cycling, dust contamination and operator handling.

Defects can manifest themselves as particulate contamination from airborne dust, non-uniform surfaces, poor surface preparation, the use of

incompatible materials or poor coating techniques. Any or all of these can result in poor individual layer uniformity, gaps or voids in layers or non-uniform sensitivity. Collectively these are called drop-outs. Drop-outs can appear as small as a single bit or as large as several thousand bits in both the along-track and cross-track directions.

The reason a disk manufacturer must pay such extreme attention to obtaining a defect-free media is due to the very nature of write-once recording. One cannot fully characterize each and every disk as is done in the magnetic disk industry. To do so would require writing on it. Once written over with test data, the disk cannot be used by the customer for valid data recording. In order to circumvent this difficulty in media testing, manufacturers, after assuring themselves that their formulation is valid and their process control is well in hand, resort to several proven techniques for quality control testing such as batch sampling or visual inspection.

Batch sampling simply refers to removing a statistical sample of disks or test targets from each batch or production run. These samples are subject to a battery of quality checks. This includes placement in environmental chambers to evaluate accelerated life test performance, and fully writing the disks to evaluate their recording and playback performance. This is essentially destructive testing.

There are a number of optical inspections that can be performed on candidate optical media to give an indication of its performance in a recording system. Measurements of reflectivity, absorption, layer uniformity and thickness, defect counts, defect mapping and others are routinely performed. This provides a good indication of the condition of individual samples.

Even with all the care that goes into the characterization, design, manufacture and quality control testing of optical disk media, the best disks available are far from error-free. Typical disks seen in the industry exhibit raw Bit Error Rates (BER) of 1 bit in 10^5 (10^{-5}). Systems designers

would like to see much better. Raw error rates in the 10^{-10} range would be considered acceptable. The reasons why are quite simple. The higher the raw (uncorrected) BER of the disk, the larger the amount of space that must be left available for re-writes of incorrect data blocks when recording. Also, the system designer must include additional levels of error detection and correction to bring the corrected BER performance of the system into an acceptable range for the computer industry. As the amount of re-write space and mathematical error detection increases, the higher the complexity of the circuitry and the greater the total overhead figure of the system. This lowers user capacity on the disk and raises the system cost.

For the computer industry, a BER of 10^{-5} is totally unacceptable. Even 10^{-10} is not good enough for computer peripheral quality hardware. Most magnetic peripheral hardware has specified BER's in the 10^{-12} range. That is the detectable, correctable BER. The hard error rate for uncorrectable errors must be 10^{-13} or better. That's 1 bit in every 100 to 1000 disks or so.

To make a 10^{-5} media perform at 10^{-12} or better, requires the introduction of both hardware and software Error Detection and Correction (EDAC) schemes. The introduction of EDAC to optical systems is not without precedent. Exactly the same techniques are used in the magnetic industry. EDAC involves breaking the incoming data up into pieces ranging in size from several bits to several thousand or more and adding a series of parity or check bits. These additional bits are recorded along with the incoming data at a corresponding increase in data volume and data rate. The increase could total as much as 20 to 25 percent. Whenever the data is read off the disk, the parity bits are stripped out of the data stream. They are then used to check for the presence of errors and to try to correct them.

The most widely used scheme of EDAC in the industry is the Reed-Solomon code. It can be used in a series of interleaved embedded "layers" to provide successively more powerful levels of correction on larger data blocks. It is important to state that the number of layers and consequently the more powerful the EDAC employed are directly related to the amount of

overhead recorded with the data and the system complexity. The size of the blocks that the EDAC will correct must be matched closely to the error distribution, burst length and burst distribution of the media. The lower the number of residual errors left after a given layer of EDAC has done its work, the greater the percentage of error reduction the next layer will provide. In other words, good characterization of the defect profile of the media is critical to the successful implementation of an efficient and effective EDAC scheme.

For example, if one assumes a typical 10^{-5} raw BER disk, a single layer of Reed-Solomon EDAC applied to a properly characterized media could give a corrected BER of 10^{-8} to 10^{-9} . Adding a successive layer could move the BER to 10^{-10} or so. The net increase in recorded data, and consequently the data rate, could be in the 25 percent range for this level of improvement. In essence, for every user bit recorded the system would actually lay down 1.25 bits on the disk. If the incoming data rate were 3 megabytes per second, the system would require a recording bandwidth of 3.75 megabytes per second.

It is important to stress that EDAC is most effective when employed to correct errors in the data that occur as a result of aging and handling effects. The first line of defense in most error management schemes is referred to as read-after-write. It too is a take-off from the magnetic industry. Simply stated, read-after-write means that the data is read immediately, within a few bit periods, after it is written and compared on a bit-for-bit basis with the original. If even a single bit is in error, that block of data is rewritten over and over until an error-free readback is obtained. This is done either on-the-fly or on a block or sector basis. Space is left, depending on the scheme employed, by the system designer in the recording format for these rewrites. If the system reaches a pre-determined limit in the number of rewrites, it notifies the operator that intervention is required. The sole aim of this technique is to write an error-free disk so that the EDAC must only compensate for errors due to time and environment. Of course, this function also adds to system overhead in both circuitry and lost real estate on the disk.

To summarize, the error management scheme of a "typical" digital optical disk is as follows: first and foremost, a high quality, low defect, very uniform and stable blank platter is required, second, read-after-write is employed to write an error-free disk, third, embedded, multiple layer EDAC, typically Reed-Solomon, finally, reasonable care, storage and handling.

MEDIA ARCHIVAL CHARACTERISTICS

The preceding systems discussion has assumed the use of what is termed a write-once or archival media. In other words, once written the media is immediately available to be read. Overwriting will obliterate previously written data. It cannot be erased or re-written. The media requires no development, fixing or other post processing before data is available to be read. If properly fabricated and cared for, disks should have a designed usable life of 10 to 15 years. Critical to this storage or archival life is the stability of the total disk package. Most important are the stability of the active recording layer, the impermeability of any intervening hermetic barriers and any protective overcoats. These three, the recording layer, the hermetic layers(s) and the overcoat are critical elements in providing a truly long term storage media.

The greatest concern in designing an optical disk package for long term storage life is the stability of the active recording layer. This is in addition to the concern for its uniformity and low defect count for low BER performance. The constituency of the active layer ranges widely. Early work used such materials as titanium, which had good environmental stability but very low sensitivity. Later work used such metals as tellurium, tellurium alloys, dye polymers, gold and others. These had very good sensitivity but poor environmental characteristics when left unprotected by intervening layers for hermetic protection. Humidity is the greatest culprit in degrading the optical recording properties of these materials. Some materials have usable storage lifetimes as short as a few hours when unprotected. By introducing low permeability layers such as silicon dioxide to inhibit corrosion of the active layer, the storage lifetime could be extended dramatically. In addition to hermetic sealing layers, most manufacturers also include a final relatively thick overcoat layer. The primary purpose of this is to place dust, fingerprints, etc. well out of the plane of focus of the final objective lens. This allows the system to accommodate relatively large foreign particles and normal handling without a loss of recording or playback signal from the disk.

While no manufacturer has stepped up to a 10 year warranty, 1 year and some 4 year usable lifetime guarantees are available. Consensus in the community is that as production and user experience and consequently the confidence level with the materials grows, so will the length of factory warranties.

ADVANTAGES OF NON-REWRITE MEDIA

The use of non-rewriteable media offers important advantages to certain segments of the mass storage community. Potential advocates cut across the full spectrum of commerce, industry, government and the legal profession. An "archival" media offers a user the ability to record important, unchanging or unchangeable data and be reasonably assured that there need not be any concern for the availability and integrity of the data for the warranty period of the media.

Examples of such customers are the banking industry, the Social Security Administration, the U.S. Patent Office, Library of Congress, personnel records organizations of all types, the medical community and the legal community. In some court cases, the use of non-rewriteable optical disk media has been ruled admissible as evidence. Judges have essentially ruled that the integrity of data on optical disk is sufficient to enable detection of and assure freedom from tampering. This presents an interesting array of possibilities from a product liability standpoint. Test data recorded on optical disk cannot be tampered with or altered. Therefore, it would be available for presentation as evidence years after a product has been fielded. The implications for medical case histories, financial disputes, etc. become obvious with this growing acceptance by the legal hierarchy.

The steps to be taken to protect data on optical disk from either intentional or unintentional alteration are relatively simple. Perhaps the most straightforward means is to have archival data, once recorded, housed in mechanisms possessing only playback capability. This concept has two advantages. Firstly, it is not a selectable function. There are no write-protect switches, rings, etc. to be selected or installed by an operator. The chances for accidental "erasure" are minimal. Secondly, the playback-only hardware could be less expensive and physically smaller than dual function record/play hardware.

Another method used to prevent data loss and one commonly employed by systems designers of record/playback hardware is that of a directory or look

up table. That is, when a disk is loaded or staged in a particular drive unit, the machine automatically goes to the disk table of contents or directory area and reads the address of the previously written areas. The controller simply refuses write commands to areas previously written. Conversely, before a disk is unloaded from the drive, the directory information is updated with the latest track information and the directory is rewritten to the disk. In essence, this scenario removes responsibility for write enable decisions from the operator.

DISADVANTAGES OF WRITE ONCE MEDIA

Perhaps the greatest disadvantage of using a media and a device whose properties do not allow rewrite of data into a given space is from the software standpoint. Nearly every piece of software written for digital data storage assumes the ability to overwrite updated, corrected or otherwise changed data into a previously written and currently unused space. Therefore, the introduction of a non-rewriteable media is a serious consideration. This is one of the greatest stumbling blocks to the serious implementation of optical storage products by major mainframe manufacturers.

Optical disk manufacturers have gone to great lengths in their architectures to overcome this perceived shortcoming. The most straightforward of these schemes involves the directory resident on the disk. Each time a disk is mounted the system immediately locates the directory tracks and reads them into a buffer. As a session proceeds, this directory is constantly updated. If a previously written sector or file is to be rewritten or changed, the system writes it into the next available space. It then updates the links and pointers in the directory. This effectively nulls or erases references to the previous version of the file. When a session is complete, one of the last things the system does before unloading the disk, is to write an updated directory.

Problems with this can occur when a heavily utilized disk undergoes many updates. The system literally runs out of space on which to put the updated and/or new information and returns a "disk-full" status. The disk must be recopied to a new platter in order to "free-up" blank space on the disk. This is similar to "garbage collection" routines performed on magnetic disks. There is an additional requirement of a fresh disk on which to place the good data. However, it should be pointed out that a disk that is truly archival in nature should not have to undergo this procedure frequently, if at all, during its usable life. The key here is that the optical disk is best utilized when used in an archival, relatively stable data environment.

Variations on this theme of directory update differ depending on whether a disk is or is not preformatted, what type of preformatting is employed, i.e., on a track/sector basis or just track banding, etc. The techniques are all basically the same. Critical is the ability to locate, update and rewrite the directory information as required. Another variation is the use of interim storage methods for the data such as magnetic disk or tape until it is ascertained that the data is in its best or final form. At this point, it is committed to optical disk with a reasonable assurance that future updates will be minimal in number.

FORMATTED VERSUS UNFORMATTED MEDIA

As it leaves the manufacturer, the optical disk platter is typically unformatted. There are no track, sector or other types of hard or soft reference marks on the disk. Some systems designers prefer to use the disk as received with no preformatting step. This involves the design of a disk recorder to much higher tolerances than required for a preformatted disk. It also means a more expensive machine to manufacture. In spite of these potential drawbacks, the use of such a recording system offers the advantage of being able to use blank media from any manufacturer that meets its physical constraints. It also offers some gains in media capacity, since the space necessary for formatting can be used for writing data.

The majority of optical disk systems designers opt for the use of some type of preformatting. Typically, this is accomplished by some form of mastering machine. This is a very precise, very expensive piece of hardware whose sole function is to place permanent marks on the disk that will be recognized by the hardware. The marks can be optically "burned" into the surface or physical marks pressed or stamped onto the blank disk. These marks can take the form of relatively wide bands, perhaps the equivalent of 10 or more tracks in width, placed at intervals of 50 to 100 tracks. These form a series of concentric bands over the usable surface area of the disk. The user data tracks and sectors would then be recorded between these bands. The bands serve two primary purposes. First, they provide the coarse level tracking mechanism for the hardware. The opto-mechanical system locates and follows successive bands as it reads/writes data on the surface. This relieves some of the precision requirements from the unformatted case since that hardware must only locate and track a relatively wide tracking band rather than a single track. By contrast, in unformatted systems, the mechanism operates in an open loop mode where it must make extremely precise single track jumps to record successive tracks. This requires the mechanisms to move with extreme, repeatable precision from track to track to avoid bumping into or overwriting previous data or otherwise degrade system performance. Secondly, the track bands aid in the address function of the system directory. Rather than merely placing a wide

track band on the disk surface, many designers opt to include complete track and sector address information on the disk. The system then uses this information to determine its absolute and relative position. This aids in both access time and in the complexity of the accessing scheme. It also alleviates the requirements on the tracking mechanisms even further at the expense of the complexity of the mastering machine and the manufacturing cost of the disk.

The issue of preformatted versus unformatted disks boils down to a question of cost. An unformatted platter will be less expensive to the customer in both purchase price and consumption. A formatted platter will be more expensive to produce and consumption will be higher, however, it will operate in a less complex and less expensive machine.

OTHER ISSUES

One final issue both proponents and opponents of digital optical disk point to is that of compatibility machine-to-machine. This is primarily a standards issue. As with most things pertaining to optical disk, there are absolutely no standards recognized nationally or internationally. There are merely conventions. Media sizes range from 5.25 to 8, 12 and 14 inch diameters. Substrates are glass, aluminum, rigid plastic or "floppy" mylar. Center hole sizes range from 0.6 inch to 6.625 inch diameters. Some use captive cartridges, some removable sleeves, some none at all. There are no standards for recording code, error detection and correction, bit error rate, data rate, capacity or disk formatting. There are no interface, architecture or protocol standards. These are just the major issues. There are a myriad of other less important, but certainly related points that stand in the way of a truly standardized product becoming available to the data storage consumer. Until this issue of standards is resolved and products appear that adhere to them, it is the opinion of many in the community that optical disk will remain a fledgling industry and will not attain the status of a serious contender to the magnetic industry.

SUMMARY

In the optical disk mass storage community, there are several current offerings that should be mentioned. The lower performance end of the spectrum is represented by such firms as Filenet, Integrated Automation Inc., Cygnet Systems and a few of the Japanese. These are typically designed around the document storage requirements of a large office environment. Their principle design criteria is a data file the size of the information content of an 8.5 x 11 inch page. It is either digitized character-by-character or image-scanned and digitized. They are usually a disk handling mechanism interfaced to a low performance, vendor-supplied optical disk drive such as those from Thompson CSF or Shugart-Optimem. The drives hold 1.0 gigabyte and have effective data transfer rates of 1.0 megabyte per second. Effective system transfer rates can be as low as 0.5 megabytes per second. Total on-line storage is in the 100 gigabyte range. The major drawback to these systems, aside from their low performance is that several of them do not have standard interfaces available to IBM or any other large mainframes. Although several are touted as being in the works, their main interface is to terminals or other similar device handlers, hence the reason for their low data performance. There are statements regarding the soon-to-be-available all digital versions of such systems but none are available to date.

At the very high performance end of the spectrum is the RCA Corporation "Jukebox." As mentioned previously, there have been two deliveries of such systems to test sites within the U.S. Government. One is at the U.S. Air Force Rome Air Development Center (RADC), Griffiss Air Force Base, New York. The other resides at NASA's George Marshall Space Flight Center (SFC), Huntsville, Alabama. These units are designed to hold 1.25 terabytes of data on a total of 128 optical disks. Each disk holds 10 gigabytes of user data. The systems possess worst case access times in the 6-10 second range, depending on access scheme, to any point on any disk in the data store. The units have a footprint of approximately 30 to 35 square feet. The RADC unit employs a Network Systems Corporation (NSC) Hyperchannel (TM) interface, allowing the unit to be interfaced to any of a long list of mainframes

individually or in a network. The NSC hardware and software are all commercially available and supported. The NASA unit has a custom interface to a very high speed fiber optic data bus. Transfer rates on the units are asynchronous up to 6 and 12 megabytes per second for the RADC and NASA units respectively. The drawback to the RCA units is that they are not currently in volume production. To do so would require significant investment in production design and scale up which RCA has openly professed an interest in pursuing. However, it is strictly dependent on market interest. The units were built by a Government Systems Group, so they would have to be transitioned to a commercial part of the corporation, which has been done in the past. It should be noted that RCA does not have a computer peripheral hardware line of any type. However, they do have an excellent service company to maintain company products in the field.

If the requirement for a mass storage system was immediate, for instance in a 1 year time frame, the most serious recommendation would be the IBM 3480 system. It offers perhaps the best available solution to an existing mass storage problem. The 3480 is a current IBM product. It represents probably the most modern magnetic tape technology available in production. The media is a metallic thin film encased in compact, removable cartridges giving the unit very good area data density. It should soon be available as an automated library offering impressive on-line storage capability. It is fully compatible with IBM and other large mainframes and is fully supported by IBM field service. The technology and parameters are such that higher data density improvements should become available in the near future with the upgrades such that they would be installed by field engineers and be transparent to the user. The system is modularly expandable to encompass a very large storage capacity. The mean-time-between-failure (MTBF) is very good with a very short mean-time-to-repair (MTTR) and a short regular preventive maintenance schedule.

There are, however, two major drawbacks to the IBM 3480. First, it is a tape-based system. This means that both average and worst case access times will be significantly longer than a disk-based system. The user must wait for the tape to get from one end of a staged cartridge to the desired

file's physical location. Cartridge load times should be similar to a disk-based system. The other drawback to the 3480 is availability. Delivery schedules are fairly long owing to the success of and subsequent demand for it.

CONCLUSION

Digital optical disk is a viable candidate for mass data storage and retrieval. However, in order to become accepted by the community, several issues must be resolved. First and foremost, there must be a successful, large scale commercial development and product offering. Storage Technology Corporation (STC), until its apparent demise, offered the greatest promise in this area. A commercial offering of this type should and must provide the following: an IBM channel compatible hardware and software protocol, essentially IBM "plug compatible," it must be fully staffed and supported by field engineering, both hardware and software, there must be at least one and perhaps more second source media vendors. From an equipment standpoint, the hardware must support at least 4 gigabytes of user storage per platter, more would be better. Platters must be removable and error rates must meet or surpass the 10^{-12} soft and 10^{-13} hard bit error rates of the magnetic industry over a usable media lifetime of 5 to 10 years. There must be a hardware development in the automated library function area. There is a large and increasing requirement in the community for huge amounts of on-line storage. Storage requirements in the 10^{15} to 10^{20} bits range, on-line or automatically accessible are not farfetched. Access to this data need not be in milliseconds, times as long as seconds or in some cases minutes are acceptable.

*MISSION
of
Rome Air Development Center*

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control, Communications and Intelligence (C³I) activities. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C³I systems. The areas of technical competence include communications, command and control, battle management, information processing, surveillance sensors, intelligence data collection and handling, solid state sciences, electromagnetics, and propagation, and electronic, maintainability, and compatibility.

END

FILMED

3-86

DTIC