

NASA Conference Publication 3165, Vol. II


INSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications

Volume II

(NASA-CP-3165-Vol-2) NSSDC
CONFERENCE ON MASS STORAGE SYSTEMS
AND TECHNOLOGIES FOR SPACE AND
EARTH SCIENCE APPLICATIONS, VOLUME
2 (NASA) 257 p

N93-16320
--THRU--
N93-16333
Unclas

H1/82 0126634



NASA Conference Publication 3165, Vol. II

NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications

Volume II

*Edited by
Ben Kobler
Goddard Space Flight Center
Greenbelt, Maryland*

*P. C. Hariharan
and L. G. Blasso
STX Corporation
Lanham, Maryland*

Proceedings of a conference held at
NASA Goddard Space Flight Center
Greenbelt, Maryland
July 23-25, 1991

NASA

National Aeronautics and
Space Administration

Office of Management

Scientific and Technical
Information Program

1992

[The page contains extremely faint and illegible text, likely bleed-through from the reverse side of the document. The text is too light to transcribe accurately.]

Preface

The National Space Science Data Center (NSSDC) at NASA's Goddard Space Flight Center has been charged with the archiving of data collected from NASA's scientific spaceflight missions flown over the past 30 years. During this time NSSDC has accumulated an archive of several terabytes of data. In the coming years NASA will be generating this volume of data every few days or less. Thus, data storage media and systems become critically important to NASA if it is to successfully manage this data volume and to have a chance to transform these data into scientific knowledge.

NSSDC will play an important role in NASA's awareness of and exploitation of emerging mass storage systems, both at NSSDC and in the increasingly distributed NASA scientific data environment. For this reason, NSSDC organized a conference at Goddard in the summer of 1991 to review the status of and the outlook for data storage media and systems. Leading experts in each of several areas were invited to make presentations, and a highly informative conference transpired. In order that the record of that conference be preserved, this set of presentations is being published.

The Proceedings of the NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science applications are published in four volumes, with each of the first three volumes containing the talks and presentations for that particular day. Discussions following some of the talks are collected in the fourth volume along with introductory biographical material on the speakers. Despite our best efforts, the questions and answers were sometimes inaudible to the transcriptionist. An effort was made to contact the participants to clarify the transcript, and we are grateful to the speakers who cooperated.

The success of an endeavor of this magnitude depends on the generous help and cooperation of the participants. We would like to record our gratitude to the speakers, the audience, and in particular, to the following individuals and organizations for their assistance:

The Program committee whose membership is listed in the front of each volume

The session chairs who kept the schedule on track:

Professor Bharat Bhushan of Ohio State University
Dr. Barbara Reagor of Bellcore
Dr. Robert Freese of Alphasatronics
Mr. Patric Savage of Shell
Professor John C. Mallinson of Mallinson Magnetics
Dr. Kenneth Thibodeau of the National Archives and Records Administration

The members of the Panel Discussion and Professor Mark Kryder of Carnegie Mellon University who moderated the discussion

The hard-working crew from Westover Consulting

Dr. Dennis E. Spiliotis of Advanced Development Corporation for his help with the transcript of the Panel Discussion

Dr. James L. Green and the National Space Science Data Center

Dr. J. H. King
Dr. P. C. Hariharan
Benjamin Kobler

PRECEDING PAGE BLANK NOT FILMED

NSSDC Mass Storage Conference Committee

Ben Kobler, *NASA/GSFC (Chair)*
John Berbert, *NASA/GSFC*
Sue Kelly, *DOE/Sandia*
Elizabeth Williams, *SRC/Bowie*
Al Dwyer, *STX Corporation **
P. C. Hariharan, *STX Corporation **
Sanjay Ranade, *STX Corporation **

Conference Coordinator

Kim Blackwell, *STX Corporation **

* Hughes STX Corporation as of October 1, 1991

MASS STORAGE CONFERENCE PROCEEDINGS

CONTENTS

VOLUME I - FIRST DAY - TUESDAY, JULY 23, 1991

1. *Introduction*.....1-1
Dr. Milton Halem, Chief
NASA Space Data and Computing Division
2. *Enterprise Storage Report for the 1990s*.....1-3
Fred Moore, Corporate Vice President, Strategic Planning
Storage Technology Corporation
3. *Optical Disk and Tape Technology*.....1-33
Dr. Robert Freese, President, Alphasatronics
4. *Magnetic Disk*.....1-69
Dr. John C. Mallinson, Mallinson Magnetics, Inc.
5. *Magnetic Tape*.....1-79
L. Harriss Robinson, Manager, Mass Storage Systems
Datatape, Inc.
6. *Storage System Software Solutions For High-End User Needs*.....1-109
Carole Hogan, Director of Technical Development, DISCOS
7. *File Servers, Networking and Supercomputers*.....1-145
Dr. Reagan Moore, San Diego Supercomputing Center
8. *Mass Storage System Experiences and Future Needs at National Center
for Atmospheric Research*.....1-163
Bernard T. O'Lear, Manager, Systems Programming,
Scientific Computing Division, NCAR
9. *The Long Hold: Storing Data At The National Archives*.....1-183
Dr. Kenneth Thibodeau, Director, Center for Electronic Records,
National Archives and Records Administration
10. *Banquet Presentation*.....1-187
Dr. John C. Mallinson, former Director of Center for Magnetic Recording
Research, University of California, San Diego

MASS STORAGE CONFERENCE PROCEEDINGS

CONTENTS (Continued)

VOLUME II - SECOND DAY - WEDNESDAY, JULY 24, 1991

1. *Stewardship of Very Large Digital Data Archives*.....2-1
Patric Savage, Shell Development Company
2. *High-Performance Storage Systems (Text Not Included in this Publication)*.....2-7
Robert Coyne, IBM Federal Systems Research Division
3. *An Open, Parallel I/O Computer as the Platform for High-Performance, High-Capacity Mass Storage Systems*.....2-9
Adrian Abineri, APTEC Computer Systems
4. *EMASSTM: An Expandable Solution for NASA Space Data Storage Needs*.....2-21
Anthony L. Peterson, P. Larry Cardwell, E-Systems, Inc.
5. *Data Storage and Retrieval System*.....2-35
Glen Nakamoto, MITRE Corporation
6. *The Challenge of a Data Storage Hierarchy: Data Archiving*.....2-55
Michael Ruderman, Mesa Archival Systems, Inc.
7. *Network Accessible Multi-Terrabyte Archive*.....2-85
Fred Rybczynski, Metrum Information Storage
8. *ICI Optical Data Storage Tape*.....2-98
Robert A. McLean, Joseph F. Duffy
9. *ATL Products Division's Entries into the Computer Mass Storage Marketplace*.....2-110
Fred Zeiler, Odetics Automated Tape Library (ATL)
10. *Panel Discussion*.....2-122
Moderator: Dr. Mark Kryder
Dr. Barbara Reagor, Bellcore
Darlene M. Carlson, 3M National Media Laboratory
Kazuhiro Okamoto, Sony Magnetic Products, Inc.
Allan S. Hadad, Ampex Recording Media Corporation
John W. Corcoran, Consultant/Ampex Corporation
Dr. Dennis E. Speliotis, Advanced Development Corporation
11. *Poster Session*.....2-237
Steve Atkinson, Ampex Recording Media Corporation
Barbara Matheson, STX Corporation
Kirby Collins, CONVEX Storage Systems

MASS STORAGE CONFERENCE PROCEEDINGS

CONTENTS (Continued)

VOLUME III - THIRD DAY - THURSDAY, JULY 25, 1991

1. *Tribology of Magnetic Storage Systems*..... 3-1
Dr. Bharat Bhushan, Ohio State University
2. *Network Issues for Large Mass Storage Requirements*3-73
James "Newt" Perdue, Ultra Network Technologies, Inc.
3. *The Role of HIPPI Switches in Mass Storage Systems; A Five Year Prospective* ...3-97
T. A. Gilbert, Network Systems Corporation
4. *The National Space Science Data Center: An Operational Perspective*.....3-127
Ronald Blitstein, Hughes STX Corporation/NSSDC
Dr. James L. Green, NSSDC
5. *EOSDIS/DADS Requirements*.....3-141
John Berbert, Ben Kobler
6. *The Preservation of Landsat Data by the National Land Remote Sensing
Archive*.....3-153
John E. Boyd, EROS Data Center
7. *Status of Emerging Standards for Removable Computer Storage Media and
Related Contributions of NIST*.....3-163
Fernando L. Podio, National Institute of Standards and Technology Computer
Systems Laboratory
8. *Data Management in the National Oceanic and Atmospheric Administration
(NOAA)*.....3-189
William M. Callicott
9. *Storage Needs in Future Supercomputer Environments*3-211
Dr. Sam Coleman, Lawrence Livermore National Laboratory
10. *Requirements for a Network Storage Service*3-225
Suzanne M. Kelly, Rena A. Haynes, Sandia National Laboratories

MASS STORAGE CONFERENCE PROCEEDINGS

CONTENTS (Continued)

VOLUME III - TRANSCRIPTIONS

Introduction of Fred Moore.....3-237

Introduction and Discussion of Dr. Robert Freese.....3-238

Introduction and Discussion of Dr. John Mallinson.....3-240

Introduction and Discussion of Mr. Harriss Robinson.....3-243

Introduction and Discussion of Ms. Carole Hogan.....3-245

Introduction and Discussion of Dr. Reagan Moore.....3-249

Introduction and Discussion of Mr. Bernard T. O'Lear.....3-251

Introduction and Discussion of Dr. Kenneth Thibodeau.....3-254

Introduction of Mr. Patric Savage.....3-258

Introduction and Discussion of Mr. Adrian Abinerl.....3-259

Introduction and Discussion of Mr. Anthony L. Peterson.....3-262

Introduction and Discussion of Mr. Glen Nakamoto.....3-265

Introduction and Discussion of Mr. Michael Ruderman.....3-269

Introduction and Discussion of Mr. Fred Rybczynski.....3-270

Introduction and Discussion of Mr. Robert McLean.....3-272

Introduction of Mr. Fred Zeiler.....3-275

Introduction and Discussion of Dr. Bharat Bhushan.....3-276

Introduction and Discussion of Mr. James "Newt" Perdue.....3-277

Introduction and Discussion of Mr. Tom Gilbert.....3-278

Introduction and Discussion of Mr. Ronald Blitstein.....3-279

Introduction and Discussion of Mr. John Berbert.....3-281

Introduction and Discussion of Mr. John Boyd.....3-282

Introduction and Discussion of Mr. Fernando Podio.....3-284

STEWARDSHIP OF VERY LARGE DIGITAL DATA ARCHIVES

Patric Savage
Shell Development Company

Call an archive a permanent store.

Some of the largest *digital data archives* are operated by oil exploration organizations. The vast bulk of these archives is seismic data. It is kept forever because of its extremely high cost of acquisition, and because it often cannot be re-acquired (due to cultural buildup, political barriers, or difficult logistical/administrative factors). Western Geophysical operates a seismic data archive in Houston consisting of more than 725,000 reels/cartridges, typical of the industry. Oil companies fondly refer to their seismic data troves as "family jewels."

There are relatively few "very large" digital data archives in existence. Most business records are (gladly) expired within five or ten years depending on statutes of limitations. And many kinds of business records that do have long lives are embedded in data bases that are continually updated and re-issued cyclically. Also a great deal of "permanent" business records are actually archived as microfilm, fiche, or optical disk images - their digital version being an operational convenience rather than an archive.

So there is not really much widely known about operating *digital data archives*, let alone *very large* ones. Even the oil companies have been in a sense overwhelmed by this somewhat unplanned for hugeness.

This paper addresses the problems foreseen by the author in stewarding the very large digital data archives that will accumulate during the mission of the EOS. It focuses on the function of "shepherding" archived digital data into an endless future.

Stewardship entails a great deal more than storing and protecting the archive. It also includes all aspects of providing meaningful service to the community of users (scientists) who will want to access the data. The complete steward will:

1. Provide against loss due to physical phenomena.
2. Assure that data is not "lost" due to storage technology obsolescence.
3. Maintain data in a current formatting methodology. Also, it may be a requirement to be able to reconstitute data to original as-received format.
4. Secure against loss or pollution of data due to accidental, misguided, or willful software intrusion.

5. Prevent unauthorized electronic access to the data, including unauthorized placement of data into the archive.
6. Index the data in a *metadatabase* so that all anticipatable queries can be served without searching through the data itself.
7. Provide responsive access to the metadatabase.
8. Provide appropriately responsive access to the data.
9. Incorporate additions and changes to the archive (and to the metadatabase) in a timely way.
10. Deliver only copies of data to clients - retain physical custody of the "official" data.

Items 4 through 10 are not discussed in this paper. However, the author will answer questions about them at the conference or by email or telephone.

Providing Against Loss Due to Physical Phenomena

Broadly classifying these we have:

1. Site destruction
2. Theft/robbery
3. Sabotage
4. Media unit suffers severe damage
5. Systemic media degradation

The first three can be guarded against, but not absolutely. The fourth is a rare inevitable eventuality (e.g., a mechanically faulty drive "eats" a tape.)

Systemic media degradation is best managed by using only media that are known to have archival properties, by conservatively rewriting media that, when accessed, are found to have an error, by regularly running PM according to vendors' recommended practice (e.g., winding and re-tensioning tape), and copying the entire archive to new-generation media. The last must be planned for, budgeted for, and be resigned to - it is an imperative. Generally speaking, one media generation can be leapfrogged by the copy procedure: for example, when Shell adopted 3480 technology, all of the 1100bpi tapes were copied; when 3490 technology is adopted, all of the 6250bpi tapes will be copied. However, copying can be mandated earlier if media are observed to be systemically degrading faster than anticipated.

Media failure occurs when an uncorrectable bit error is detected. This always causes/implies loss of an entire error correction block - ordinarily a minimum of one kilobyte of archived data. Archivists should be aware that the media vendors' touted "hard error rate" always has a 10,000-fold impact. A badly degraded media unit might have relatively many unreadable error correction blocks; hence even a redundancy array of media units might then (by a little bad luck) have an unrecoverable error correction block.

A practical cost-effective solution to the problem of protecting against physical loss can be tailored around the following concept (which came to me while ruminating about extending the now-familiar RAID idea to striping tape). This is merely the seed of an idea, to which a good deal of systems thought will have to be given.

Some number (in this example, 10) of archiving sites are chosen to participate/cooperate in a redundancy scheme that provides mutual protection against all modes of physical loss.

The sites should be geographically distributed in order to eliminate concern that a calamity (e.g., earthquake/meteor strike) would wipe out multiple archives. Of course, all sites individually should have reasonably good physical security.

All sites must be accessible via state-of-the-art WAN technology.

Each site houses, primarily, its own archive of data. (A variation would have a single archive partitioned and distributed among its own multiple sites.) Clients of an archive would communicate only with the primary site.

Each site also houses either p-parity or q-parity data generated from (in this example) 9 other sites. (Optionally two sites could be dedicated, one for p- and the other for q-parity.)

In the eventuality of a loss at a site (of an error correction block, or a media unit, or the site itself) any 8 of the 9 other sites reconstruct the lost data. This would not be instantaneous, as with RAID, because an extraordinary procedure would have to be executed; but the insurance would be very certain. Clearly, each site should be practicing high quality archiving methodology, so that losses would occur with extreme rarity (say, no oftener than one per month).

The merits of this scheme are first, that the storage overhead for backup can be small (25% for this example); second, that the degree of protection can be high (with both p- and q-parity) or lower (with p-parity only); third, independent archives do not have to create their own backup systems, but can band together in a consortium for mutual protection.

Assuring That Data is Not "Lost" Due to Storage Technology Obsolescence

The 1960 census was archived on the best storage medium known at the time: UNIVAC metal tape. There was a rude awakening some years later when it was discovered that only two drives existed in the world - one in Japan, and the other, dismantled, in the Smithsonian.

We know now that drive technology lifetime, even assuming heroic geriatric care, is scarcely ten years. Vendors drop maintenance after low-level parts' technologies disappear. For a while thereafter, drives can be cannibalized for parts; but ultimately maintenance becomes impossible.

The optical disk vendors, for example, tell of the fine archival qualities of their media. But their technology is evolving quite rapidly - vendors come and go, and recording formats with them. Here we have a single medium that undoubtedly lasts a long time, but the drive and recording technology has a half-life of less than five years. Considering the relatively high cost of optical media, copying an archive every five years seems out of reason.

Archiving demands that digital data on old storage technology be copied to new storage technology periodically. The frequency depends on the media, on how widely the drives were accepted, and on whether the old technology satisfies current access requirements. Keeping too many generations of storage technology in use can cause serious operational problems, even if they are all in good working condition. For example, 556bpi tape would be much too slow for regular use today, so, even though drives are still available, that technology is obsolete.


Maintaining Data in a Current Formatting Methodology

The winds of computation methodology are ever varying. Yesterday there was no C language. Today C-readable records might be a good bet. Tomorrow the fad may be object files. What will come next? The curse of required media copying is really a blessing because it enables us to continually modernize our data language. Cuneiform tablets were certainly archival, but they contain antiquated, almost unreadable language.

Standards for "self-defining" data formats are evolving rapidly and are already very useful. The time has come to abandon schema-less data formats (where programs know implicitly where every field is in a record, and what each field means).

Even fixed (schema'd) formats are passé for scientific data because of the continual change in interest and emphasis in almost every scientific specialty.

Archivists can extract a side benefit when copying to a new media generation. Indeed, the planning for the copy should include deciding which new formatting standard is to be adopted. Migrating from old to new formats is only slightly less important for archiving as migrating from old to new media technology. What's more, it's almost free.




High-Performance Storage Systems

**Robert Coyne
IBM Federal Systems Research Division**

Permission for this talk could not be obtained in time for the Text to be included in these proceedings.

PRECEDING PAGE BLANK NOT FILMED



**An Open, Parallel I/O Computer as the Platform for High-Performance,
High-Capacity Mass Storage Systems**

**Adrian Abinerl, APTEC Computer Systems
Y. P. Chen, APTEC Computer Systems**

For those of you who are not familiar with APTEC Computer Systems, we are a Portland, Oregon based manufacturer of I/O computers. About 400 of our systems are installed today, typically in real-time oriented, high bandwidth environments. Applications have included satellite ground systems, mass storage archival systems, signal and image processing systems, etc.

Much of the discussion here today has focused on mass storage solutions exclusively. That is high density storage media, attached to a general purpose computer, which in turn supports network connections to users.


APTEC's focus in this environment is on programs requiring real-time data capture, with low latency processing and storage requirements. As an example my second introductory slide illustrates the Loral / Space Telescope - Data Archival and Distribution System. This is an existing Loral AeroSys designed system, which utilizes an APTEC I/O Computer.

The key attributes of a system architecture to address these types of requirements include:

- Data acquisition alternatives
- A wide range of supported mass storage devices
- Data processing options
- Data availability through standard network connections
- An overall system architecture (hardware and software designed for high bandwidth and low latency.

The following slides outline APTEC's approach, which is designed to provide flexible, standards based, system solutions.

PRECEDING PAGE BLANK NOT FILMED



Aptec



- **Introduction**
- **Mass storage system attributes**
 - Data acquisition**
 - Mass storage devices**
 - Processing options**
 - Data availability**
 - Architecture**
- **Conclusion**

Introduction / Data Deluge

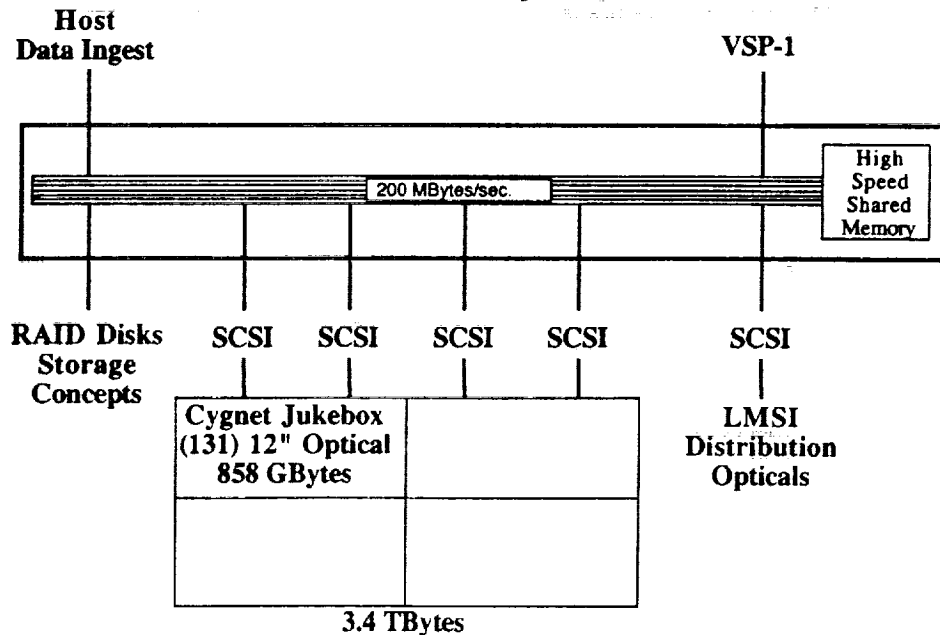


- **ERS-1**
 - Transmits data at 100 Mbits/sec.**
 - During this 30 minute presentation**
 - 160 9-track 6250 bpi tapes would be**
 - filled with data.**
- **EOS**
 - Expected to exceed 1 TByte/day**

Introduction / Loral ST-DADS



Space Telescope Data Archival and Distribution System




Mass Storage System



Attributes

- Data acquisition
- Mass storage devices
- Internal processing capabilities and connections to external processing elements
- Data availability
- Architecture

Data Acquisition Options



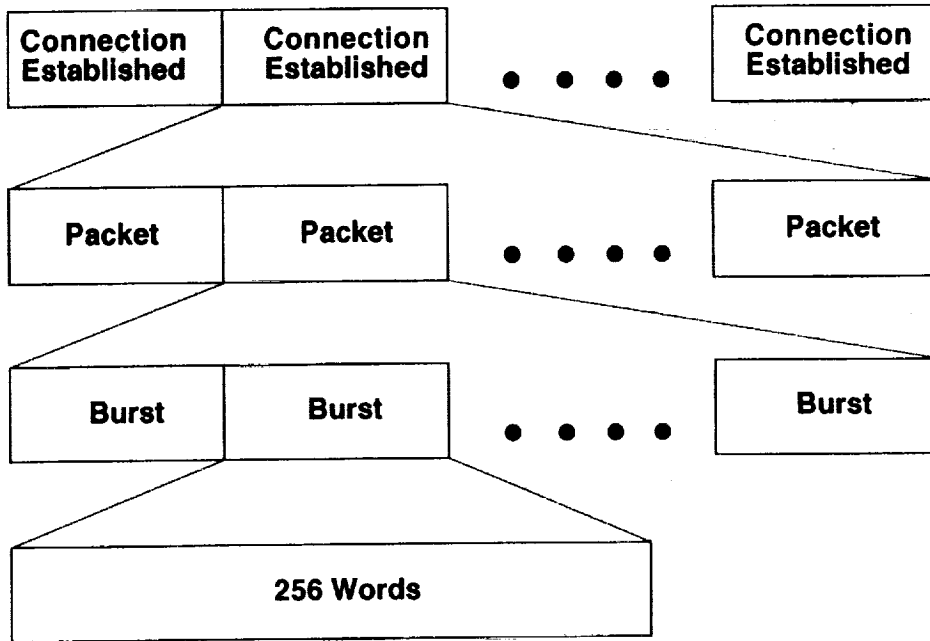
HIPPI




High Performance Parallel Interface

- ANSI Standard (X3T9.3)
- Efficient high speed interconnection optimized for large block transfers
- Point-to-point connection
- 32-bit channel
- 100 MByte/sec simplex channel

HIPPI



HIPPI




**Only HIPPI channel runs at 100 MByte/sec.
How fast the HIPPI channel is fed is not specified.**

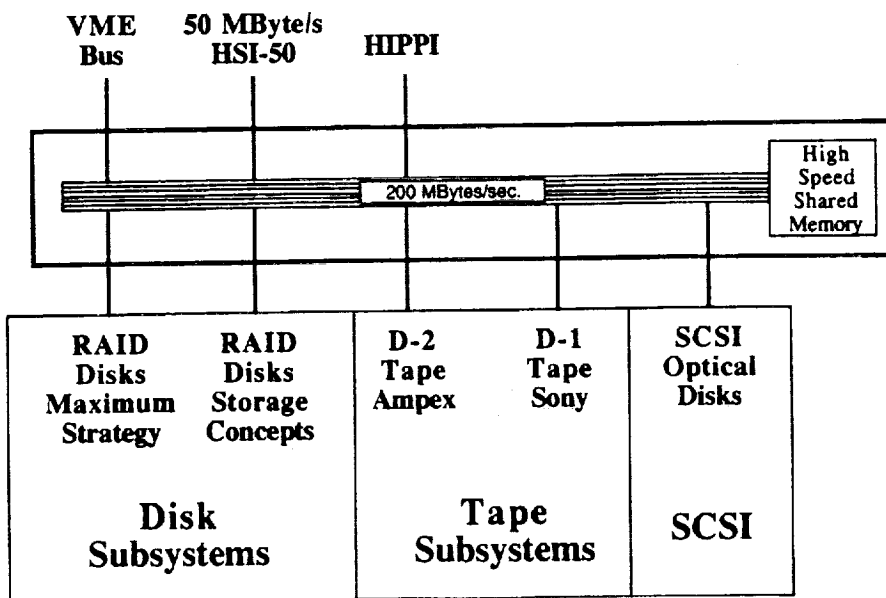


**The Aptec HSI-50 / HIPPI design provides 50 MByte/sec
sustained throughput to/from the HIPPI channel.**

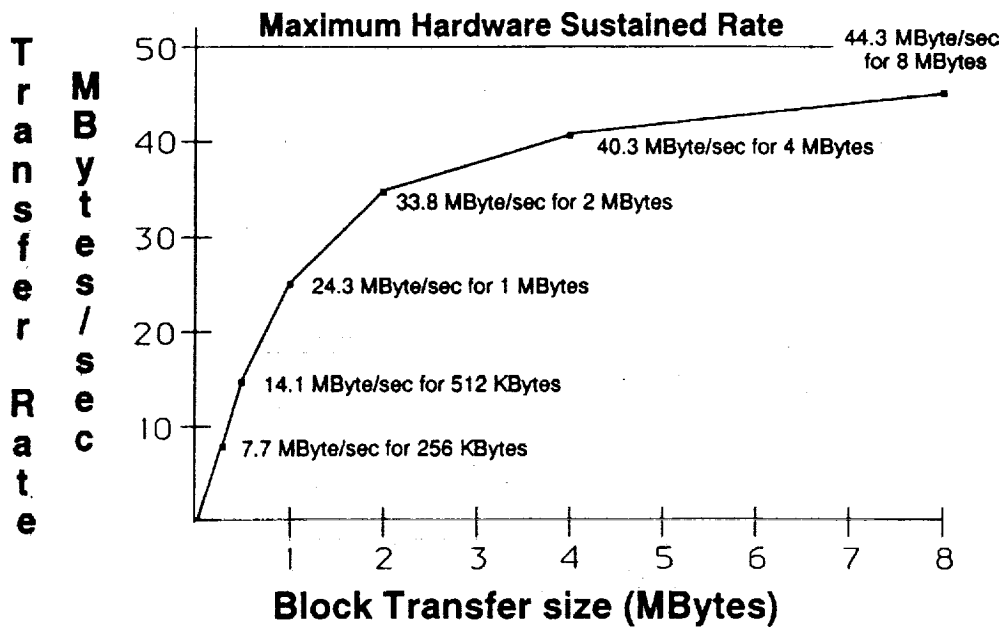
HIPPI (User defined ULP)



Mass Storage Devices



Maximum Strategy Disk



Tape Drives



Ampex DCRSi

11.4 MByte/sec transfer rate
38 GByte capacity per cartridge


Sony DIR 1000 (D-1)

Up to 32 MByte/sec transfer rate
12, 41, or 96 GByte capacity per cartridge

Ampex TeraStore (D-2)

15 MByte/sec transfer rate
25, 75, or 165 GByte capacity per cartridge
Ampex TeraAccess robotic system (6.4 TByte)

Processing Options

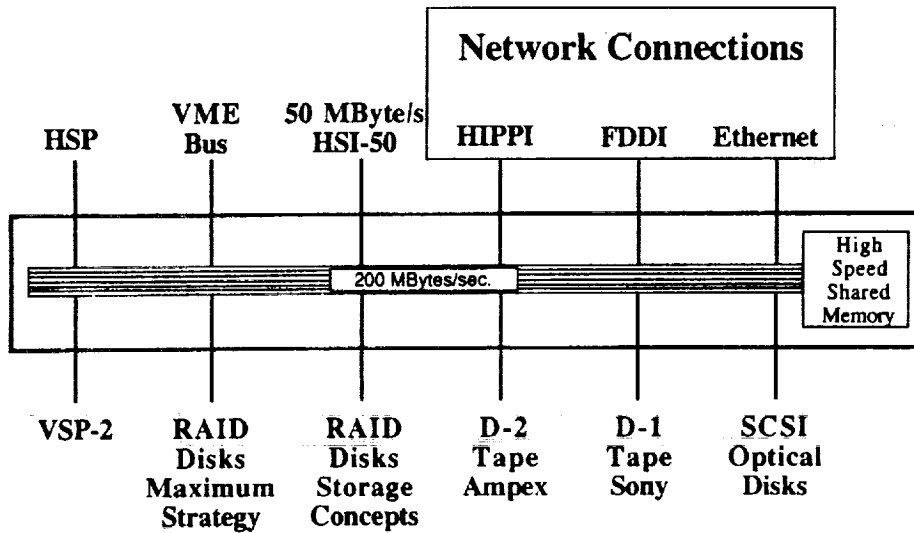


Processing Options



- **HSP - High Speed Scalar Processor**
20 MIP processor with 50 MByte/sec connection to memory.
VxWorks and C.
- **VSP-2 - Vector / Scalar Processor**
150 MFLOP Array Processor with 50 MByte/sec connection to memory.
VxWorks, C, and Math Advantage library of callable vector subroutines.
- **Many external processor links supported.** Convex, Alliant, Sun, Silicon Graphics, AMT/DAP, HIPPI etc.

Data Availability



Data Availability



Client / Server Model

- TCP/IP Access
- Server Software
- NFS Network File Access
- HIPPI
- FDDI
- Ethernet

Architecture



Aptec architecture can sustain multiple concurrent high data rate transfers with predictable repeatable performance.

- **Synchronous bus**
- **Dedicated I/O Processors**
- **Real-time kernel / VxWorks with Aptec's MultiProcessor services**

Conclusion



- **High performance solutions are available today using commercial-off-the-shelf systems and peripherals.**
- **They are cost effective and low risk systems offering flexible, modular architectures.**
- **Standards based.**
 - UNIX development environment**
 - Connectivity / networking**
 - VME, HIPPI, FDDI, Ethernet, TCP/IP**
 - VX/Works real-time kernel**



**EMASS™: AN EXPANDABLE SOLUTION
FOR NASA SPACE DATA STORAGE NEEDS**

Anthony L. Peterson
P. Larry Cardwell

E-Systems, Inc. Garland Division
Dallas, Texas

Abstract

The data acquisition, distribution, processing and archiving requirements of NASA and other U. S. Government data centers present significant data management challenges that must be met in the 1990's. The Earth Observing System (EOS) project alone is expected to generate daily data volumes greater than 2 Terabytes (2×10^{12} Bytes). As the scientific community makes use of this data their work product will result in larger, increasingly complex data sets to be further exploited and managed. The challenge for data storage systems is to satisfy the initial data management requirements with cost effective solutions that provide for planned growth. This paper describes the expandable architecture of the E-Systems Modular Automated Storage System (EMASS™), a mass storage system which is designed to support NASA's data capture, storage, distribution and management requirements into the 21st century.

Introduction

We first discuss NASA's requirements for mass storage with a focus on functional and performance specifications. Next, an overview of the EMASS architecture is presented and evaluated with respect to NASA's requirements. The major EMASS architectural components, hardware, software and interfaces, are then explored with emphasis on the data management capabilities of the EMASS software.

NASA Requirements

Requirements for large volume, mass storage systems have been well established in order to meet the storage needs for NASA's space and Earth science information systems. The use of sophisticated data acquisition instrumentation will continue to evolve, providing large, increasingly complex data sets to be processed, distributed and archived. Therefore, data storage requirements

will continue to grow nonlinearly through the 1990's. For example, the Earth Observing System (EOS) project alone, generating daily volumes greater than 2 Terabytes, will require automated storage libraries with capacities greater than 500 Terabytes by the late 1990's. E-Systems is also currently developing storage systems to meet existing U. S. Government and commercial requirements to be delivered in 1993 having automated data storage library capacities greater than 200 Terabytes.

Data management requirements such as these within NASA and other U. S. Government data centers present significant challenges that must be met in the development of new mass storage systems. These systems must meet increasing performance requirements with cost effective solutions while providing for planned growth. E-Systems is developing the EMASS architecture to address these requirements for extremely large, expandable data storage and data management systems.

As we view NASA's supercomputer-based data management systems we see a need for high bandwidth, high density tape recorder systems having the data quality characteristics of a computer peripheral. As scientific data processing requirements move towards open systems environments, the file management software and server should support a UNIX environment. The file management software structure should provide application specific integrated data management solutions. A file server with high I/O bandwidth is required to accommodate simultaneous data transfers from multiple high bandwidth tape recorder systems. Finally, to keep a perspective on hardware and maintenance costs, the use of commercially available equipment is strongly emphasized.

EMASS Architecture Overview

The EMASS architecture is a family of hardware and software modules which are selected and combined to meet these data storage requirements. Figure 1 illustrates the EMASS architecture. EMASS is a UNIX-based hierarchical file management system utilizing both magnetic disk and tape. The storage capacity ranges from one to several thousand Terabytes depending on the type of storage library used. It has the capability to support both a graphical and metadata interface to the user. The system is user driven by standard UNIX and unique EMASS commands and has user configurable automatic file migration. The EMASS system employs standard protocols for user file transfer, communications and network interfaces.




Figure 1. EMASS Architecture

The EMASS system operates as a large data storage node on a network, servicing client requests over a number of standard interfaces, including Ethernet, FDDI, DECnet, HYPERchannel™ and UltraNet™. The system is a two-level hierarchical data storage system. Magnetic disk is the first level of storage, and magnetic tape is the second. Data is managed via a selectable migration policy based on data class, a method of data segregation addressed in a subsequent section. Two alternative types of magnetic tape for data storage are included in the EMASS design: 3480 tape cartridges and D2 digital tape cassettes. Files are migrated to 3480 or D2 tape depending on the migration policy for the data class to which that file belongs. The physical volume repository (Miller¹) functionality is implemented in three separate types of storage libraries which are selected based on user requirements for performance and expandability.

EMASS Hardware

The storage system file server function is implemented in a CONVEX C3200 series computer. The CONVEX was selected after an extensive survey of available computers. The major evaluation factor leading to the selection of the CONVEX machine is its high I/O throughput performance. The CONVEX supports four channels, each having I/O bandwidths of 80 Megabytes per second peak and 60 Megabytes per second average. Other key evaluation factors included cost, compatibility with a UNIX environment, modularity, expandability, upgradability, reliability, and support.

The file server interfaces with three types of tape libraries, the STORAGETEK (STK) 4400 Automated Cartridge System, the EMASS DataTower™ and the EMASS DataLibrary™. The STK tape cartridge library data interface is implemented through ANSI standard Block Multiplexor Channel interfaces which connect to the STK 4480 drives through a SUN Library Server. The DataTower™ and DataLibrary™ data interfaces are implemented with enhanced ANSI standard IPI-3 tape controllers within the file server connected to E-Systems ER90 digital D2 recorders.

The DataTower and DataLibrary robotic systems provide data archive expandability. The DataTower, with dimensions illustrated in Figure 2, serves as a medium scale storage device, with a capacity of 6 Terabytes on 227 small D2 cassette tapes. This device was implemented by





Figure 2. EMASS DataTower

modifying an existing automated robotics tower currently in volume production for the broadcast industry. The device may be expanded by adding up to three additional expansion storage units, for a total capacity of 25 Terabytes.

The DataLibrary, illustrated in Figure 3, is a modular aisle architecture comprised of a series of modules each four feet in length. This design specifically addresses the needs for a modular, expandable data storage solution required for NASA's large data archives from EOS and Space Station Freedom. Each shelf module contains up to 207 small, or 192 medium, D2 tape cassettes, for a maximum capacity of 14 Terabytes. Shelf units reside in rows on either side of a self-propelled robot and can be added incrementally as the library grows. The row of modules may be expanded to lengths of 80 feet, providing a maximum of 288 Terabytes per row. Further expansion is accomplished by adding additional rows and robots. Cassette access times are specified at 45 seconds maximum for robot travel spanning an 80 foot aisle for cassette retrieval. The DataLibrary configuration will be housed within a sealed watertight structure with interior fire protection using CO₂ supplied on demand.




Figure 3. EMASS DataLibrary

E-Systems selected D2 helical scan tape and recorder system technology to meet high density and bandwidth requirements. As shown in Table 1, the 19mm D2 tape cassette is available in three form factors: small, with a capacity of 25 Gigabytes of user data, medium, with a capacity of 75 Gigabytes and large, with a capacity of 165 Gigabytes.

The suitability of the 19mm D2 helical scan media and recorder for use as a computer peripheral has been reviewed by Wood². The D2 recorder provides a format already in wide use within the broadcast industry. Ampex, SONY, and Hitachi have delivered over 2000 D2 units to the broadcast industry since 1988. The D2 video broadcast recorder has been modified to develop the ER90 digital recorder peripheral. Key features of the ER90 recorder include air guides to minimize tape wear, azimuth recording and automatic scan tracking.

The ER90 provides a sustained data rate of 15 Megabytes per second, with burst rates up to 20 Megabytes per second. Additional error detection and correction coding has been implemented using a three-level interleaved Reed-Solomon code. Resulting error event rates of 1 in 10^{13} bits are being achieved. The recorder absolute positioning velocity is 300 inches per second and logical positioning velocity is 150 inches per second. This results in an average file access time for mounted media of 10 seconds for D2 small cassettes and 30 seconds for D2 medium cassettes. To provide compatibility with existing file management systems the ER90 provides ANSI 9-Track file labeling compatibility.

PARAMETER	SPECIFICATION
Tape Media	19mm - D2
Tape Cassette Capacities	25 GB (S), 75 GB (M), 165 GB (L)
Data Rate	15 MB/sec - Sustained 20 MB/sec - Burst
Error Detection/Correction	3-Level Interleaved R-S Code
Error Event Rate	1 in 10^{13} bits
Tape Positioning Velocity	300 in/sec - Absolute Address
Average File Access Time (Mounted Media)	10 sec - D2 Small 30 sec - D2 Medium
Data Format	Compatible With ANSI 9-Track File Labeling
Peripheral Interface	Enhanced IPI Physical (ANSI X3T9/88-82) IPI-3 Logical (ANSI X3.147-1988)

Table 1. Recorder System Performance

The ER90 drive uses the enhanced IPI physical interface (ANSI X3T9/88-82) and the IPI-3 Magnetic Tape Command Set (ANSI X3.147-1988) at the logical interface level. The enhanced IPI physical interface can sustain transfer rates commensurate with the basic transport performance. A second enhanced IPI interface port can be added to allow a separate master-slave path to another server. A large internal buffer (approximately 60 Megabytes) has been incorporated for rate smoothing to minimize recorder start-stop sequences.

EMASS Software

The EMASS server stores files in an extended UNIX File System (UFS). EMASS software is divided into separate components as depicted in Figure 4. These components are the user interface, the event daemon, the migration manager, the file mover, and the physical device manager. All EMASS software executes as UNIX processes at the application level. All UNIX kernel enhancements/modifications were accomplished by CONVEX and are included in ConvexOS™ 9.0.

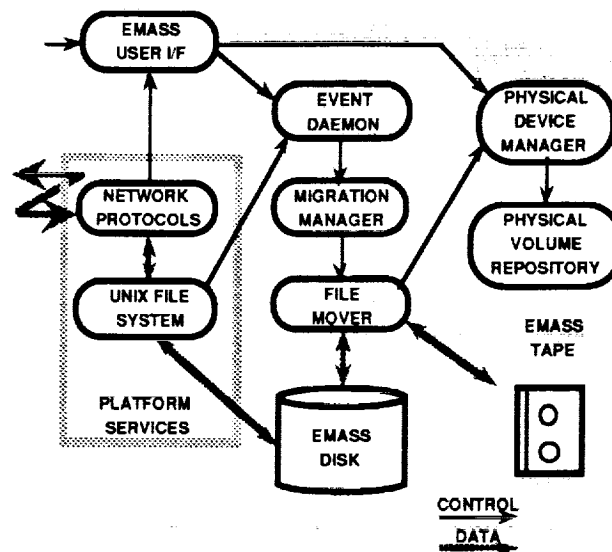


Figure 4. Overview of EMASS Software.

Users have two methods to gain access of EMASS migration services. One method is through a user interface front-end which provides migration override control to end-users. The other method is through direct access of the CONVEX UFS. This second method provides transparent access to EMASS file migration services to both local and networked users.

In order to provide the ability to transparently migrate files, CONVEX has upgraded their ConvexOS to allow the UFS to provide notification of selected critical file system events to a user-level event daemon. This modification is similar to those made by the BRL/USNA Migration Project (BUMP)³, a joint development of the US Army Ballistic Research Laboratory and the US Naval Academy. The EMASS event daemon receives file events and forwards them to the migration manager. The migration manager collects this information. When migration policy is triggered, the migration manager will select files for migration and forward the list of selected files to the file mover.

The movement of files from an EMASS server disk to magnetic tape and from magnetic tape to an EMASS server disk is controlled by the file mover. For each list of files to migrate, a file mover process is created to perform the read and write operations. The file mover design provides for the addition of software routines to support new media types.

The final major EMASS software component is the physical device manager. The physical device manager provides a standard interface for tape movement services to all other EMASS software components. The physical device manager will translate a generic tape movement request into the format required by the target physical volume repository (PVR). The translated request is then sent to the PVR for processing. The physical device manager will later receive the results from the commanded PVR. The results are then placed into a generic format and sent to the process that requested the tape movement. Additional PVRs will be supported by the addition of software modules to the physical device manager.

DataClass™

Before describing EMASS software in more detail, a discussion of the abstraction known as DataClass™ is required. The file systems that are to be provided EMASS migration services are subdivided at specific points in the file system tree structure by identifying those directory point(s) beneath which all files are to be managed alike. These directory point(s), which are referred to as migration directories, are what define each DataClass. When a new directory point is added to a DataClass, the event daemon will request the UFS to associate the directory and all files below it (both present and future) with the event daemon.

Figure 5 depicts a DataClass to migration directory relationship. The directory `/test/dick/special` is the only migration directory in DataClass SPECIAL. All files beneath `/test/dick/special` will be managed together. The DataClass PURPLE contains all files under the directories `/prod/blue` and `/prod/red`, but none under `/prod/green`, showing that some directories at a certain level may be excluded from a DataClass. All files under `/test/jane` and `/test/dick/public` belong to DataClass TESTERS. This illustrates that the assignment of migration directories to DataClass is not restricted to a certain level in the tree structure. In fact, migration directories from different file systems may be in the same DataClass. Also, a file system can be mounted onto a mount point underneath a migration directory, for example `/test/jane/dir1`.

The definition of DataClass is key to site administration. Migration policy parameters are configurable on a DataClass basis, thus providing the EMASS administrator with a great deal of control over the behavior of the EMASS system. Time interval between policy application, time required on disk prior to migration, and desired time for migrated files to remain cached on disk are examples of DataClass based migration policy parameters. Quotas for tape utilization (both

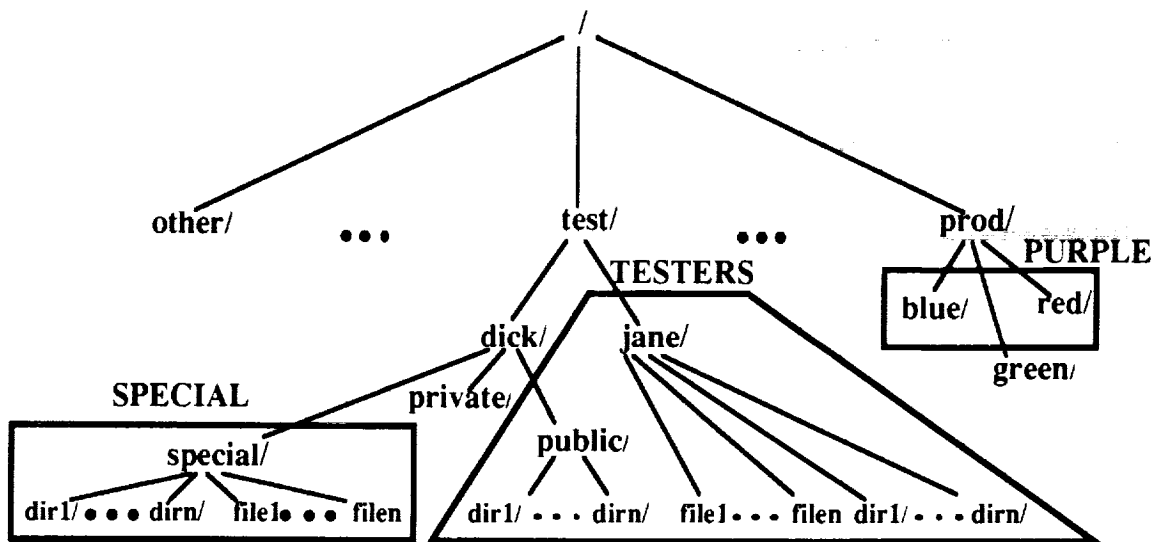


Figure 5. DataClass Example.

a warning limit and a hard limit) are also kept on a DataClass basis. DataClass based parameters are kept in the INGRES database, so tuning can be easily done while the EMASS system is active.

EMASS software also uses DataClass as the means to segregate files on tape. All files on a tape will be of the same DataClass. This provides a level of physical security for those sites which might require it. This segregation also ensures that retrieval of files from different user groups (as defined by DataClass) will not collide trying to access the same physical tape.

EMASS Interfaces

A key concept of the EMASS architecture is that it will provide multiple archival file storage choices to the users of various networked client systems. EMASS supports connectivity over industry standard interfaces including Ethernet, DECNet, FDDI, HYPERchannel, and UltraNet. This provides connectivity from the smallest workstations to the largest supercomputers. The industry standard transfer protocols available to the user include the File Transfer Protocol (FTP)

from the TCP/IP family, the Network File System (NFS™) as defined by Sun Microsystems, and the UNIX utilities UNIX to UNIX copy (UUCP) and remote copy (rcp). This support is provided by placing the EMASS interface under the UFS. The EMASS system will receive notification of all managed file system events for files in every DataClass. Migration services are therefore provided for any connectivity available on the EMASS server system to the UFS. Thus, as new connectivity options are offered by the EMASS server vendor, the EMASS system will automatically also provide support.

Hierarchical Data Storage

EMASS provides three levels of data storage. These three levels of hierarchical storage are EMASS server disk, robotically managed tape, and human managed tape. When files are placed in a DataClass, the residency is on server disk. The EMASS migration policy will schedule placing the files onto tape based on the migration rules defined for its DataClass. The user can preempt the migration policy by giving the EMASS system a directive to migrate specific files to tape immediately.

When the migration policy is executed for a DataClass, all files in that DataClass which are not solely on tape are examined. If the time since last modification (or the time since retrieval from tape if unmodified) of a file is greater than that specified in the policy for that DataClass, then that file's data is placed in the staging directory. If an up-to-date copy of the file is not on tape, the file is added to a list of files to be migrated. When all files in the DataClass have been examined, the list of files to migrate is forwarded to the file mover.

To store files on tape, the file mover first allocates tape(s) through the physical device manager. The files in the list are next migrated to tape. The file mover will record the location of the tape copy of each file in the INGRES database as they are successfully written to tape. The disk copy is left intact as a cached copy which will later be removed from disk when either 1) the length of time since migration has exceeded its DataClass defined limit or 2) the file system requires additional disk data blocks and it is the oldest staged file on disk. When the file mover has completed migrating files to tape, the tape and drive are released back to the physical device manager.

Not removing disk data blocks from the staging directory on strictly a first-in-first-out basis provides additional flexibility. A file system can be divided into DataClasses which can have different access requirements. Each DataClass will have its own disk retention period, so one

portion of a file system can have data that is not cached on disk as long as another portion. One DataClass thus can be set up to never free disk data blocks except when required to provide needed free disk space for the file system.

When the UFS receives a request for data blocks for a file which is not currently on disk, the requesting process is suspended and the EMASS event daemon is notified. The event daemon forwards that notification to the migration manager which immediately instructs the file mover to retrieve the requested file to disk. The file mover requests the tape containing the file be mounted and copies the file to disk. The requesting process is now allowed to continue processing. The EMASS system maintains knowledge of the tape copy. If the disk copy is unchanged, the file will not be re-migrated by the migration policy.

The retrieval of a user-directed range of bytes from a file is also available to the EMASS user. This is accomplished much like the retrieval of a complete file. However, the file mover will copy the specified range of bytes into a UNIX file of a different name as specified by the user. Thus, the user can retrieve only the portion of the file of interest, reducing the amount of data brought back.

The EMASS system will also manage tapes that are not under robotic control. This will allow sites to have EMASS management of many more tapes than the robotic system can support. When access is requested for a file that is only on a tape that is not under robotic control, the EMASS system will request the operator to return the tape to active service so that the file may be copied onto disk. The effect to the user is only a longer delay waiting for a tape mount.

Infinite File Life

A mass storage system must provide for the integrity of its client's data. In order to insure that the client's data is always available, the EMASS system has several features to provide safeguards against data loss. These features include automatic Error Detection and Correction (EDAC) monitoring and secondary file copy maintenance.

For every file segment written by an ER90 drive, the drive will automatically perform a read while write comparison. If the data written is not recoverable or to successfully recover the data written required more than a minimum threshold of correction, that segment is automatically re-written to tape by the drive without any action required by the host system. This provides positive assurance that the data written is retrievable without much stress on the EDAC at the time it is recorded.

For every file read by an ER90 drive, the EMASS system will request the drive to return EDAC statistics and if the level of correction was excessive, that tape will be placed in a "suspect" list for system administrator action. The system administrator can then at a later time request the EMASS system to move all files off of the old tape. This provides for refreshing the EDAC encoding for all files that were on the old tape.

To ensure the health of D2 tapes that have not been accessed for a while, the EMASS system provides a tape sniffing service. Tape sniffing is the process of periodic monitoring of tapes that have not been accessed for a length of time defined by the data center. The EMASS system will schedule the reading of sample files from the tape and then examine the EDAC statistics to determine if the level of correction was excessive. If excessive, the tape is placed in the "suspect" list for system administrator action.

As an added measure of protection for tape-based files, secondary file copying is provided. If enabled for its DataClass, files will automatically have a secondary copy maintained on a separate tape. This DataClass feature can be overridden on a file basis, thus allowing the user to request a secondary copy be created when the DataClass default is to not maintain a copy. The user can conversely request the EMASS system not to maintain a secondary copy of a file when its DataClass default is to maintain a copy.

Through the use of automatic EDAC monitoring and secondary copies, the EMASS system provides for the integrity of its client's data. The life of the EMASS client's data can in fact be prolonged well beyond that of any one type of storage media, as the file sniffing service will promote data from one media onto another.

Summary

EMASS software provides a UNIX-based data storage solution with automatic and transparent file migration and retrieval. Data archive centers can be provided with very large (up to Petabytes), expandable, automated data storage systems. These data storage systems connect to high speed networks, providing 24 hour per day accessibility for rapid delivery of requested data in the Space Station and EOS era. File access can be provided to networked users through standard file transfer protocols. A graphical user interface can also be provided. Thus, the client is not required to have special networking software. The implementation of DataClass provides a flexible method for tuning the behavior of the system at each installed center.

Trademarks

EMASS, DataTower, DataLibrary and DataClass are trademarks of E-Systems, Inc.

CONVEX and ConvexOS are trademarks of CONVEX Computer Corporation.

UNIX is a trademark of AT&T.

INGRES is a trademark of ASK Computer Systems.

NFS is a trademark of Sun Microsystems, Inc.

HYPERchannel is a trademark of Network Systems Corporation.

UltraNet is a trademark of Ultra Network Technologies, Inc.

References

¹ S.W. Miller, "A Reference Model for Mass Storage Systems", *Advances in Computers*, Vol. 27, Academic Press, 1988, pp. 157-210.

² Tracy G. Wood, "A Survey of DCRSi and D2 Technology", *Digest of Papers, Proc. Tenth IEEE Symposium on Mass Storage Systems*, May 1990, p. 46 (1990).


³ Michael John Muuss, Terry Slattery, and Donald F. Merritt, "BUMP, the BRL/USNA Migration Project", *Unix and Supercomputers*, 1988.

**Data Storage
and
Retrieval System**

24 July 1991

Glen Nakamoto

MITRE Corporation
Bedford, MA 01730
(617) 271-3032



Data Storage and Retrieval SystemBackground

The Data Storage and Retrieval System (DSRS) consists of off-the-shelf system components integrated as a file server supporting very large files. These files are on the order of one gigabyte of data per file, although smaller files on the order of one megabyte can be accommodated as well. For instance, one gigabyte of data occupies approximately six 9 track tape reels (recorded at 6250 bpi). Due to this large volume of media, it was desirable to "shrink" the size of the proposed media to a single portable cassette. In addition to large size, a key requirement was that the data needs to be transferred to a (VME based) workstation at very high data rates. One gigabyte (GB) of data needed to be transferred from an archiveable media on a file server to a workstation in less than 5 minutes. Equivalent size, on-line data needed to be transferred in less than 3 minutes. These requirements imply effective transfer rates on the order of four to eight megabytes per second (4-8 MB/s). The DSRS also needed to be able to send and receive data from a variety of other sources accessible from an Ethernet local area network.

System Configuration

In order to meet these requirements, a system was configured using Aptec's Input/Output Computer (IOC-24) with Storage Concepts C51 disk array and Honeywell's Very Large Data Store (VLDS) tape drive (dual channel unit) as the basic components for this file server. The IOC-24 has eight megabytes of shared memory and was hosted on a VAX 11/750 which, in turn, was connected to an Ethernet local area network (LAN). The interface to the VME based workstation was accomplished via Aptec's VME Gateway Controller. The specific (and initial) VME based workstation that needed to be interfaced for this project was the Sun 3/260 workstation containing a Vicom II-9 image computer and the Vicom Fast Disk (Maximum Strategy's parallel disk array). The Sun workstation also contained a 16 megabyte high speed (32 MB/s) memory card from Micro Memory. This memory card was used as a high speed receiver (or transmitter) of data during the initial period (prior to the Vicom Fast Disk becoming available) for debugging purposes. Data needed to be transferred to/from the DSRS file server (C51 disk or VLDS tape) to the parallel disk array under the control of the Sun workstation operating at the data rates previously discussed. The specific configuration is illustrated in Figure 1.

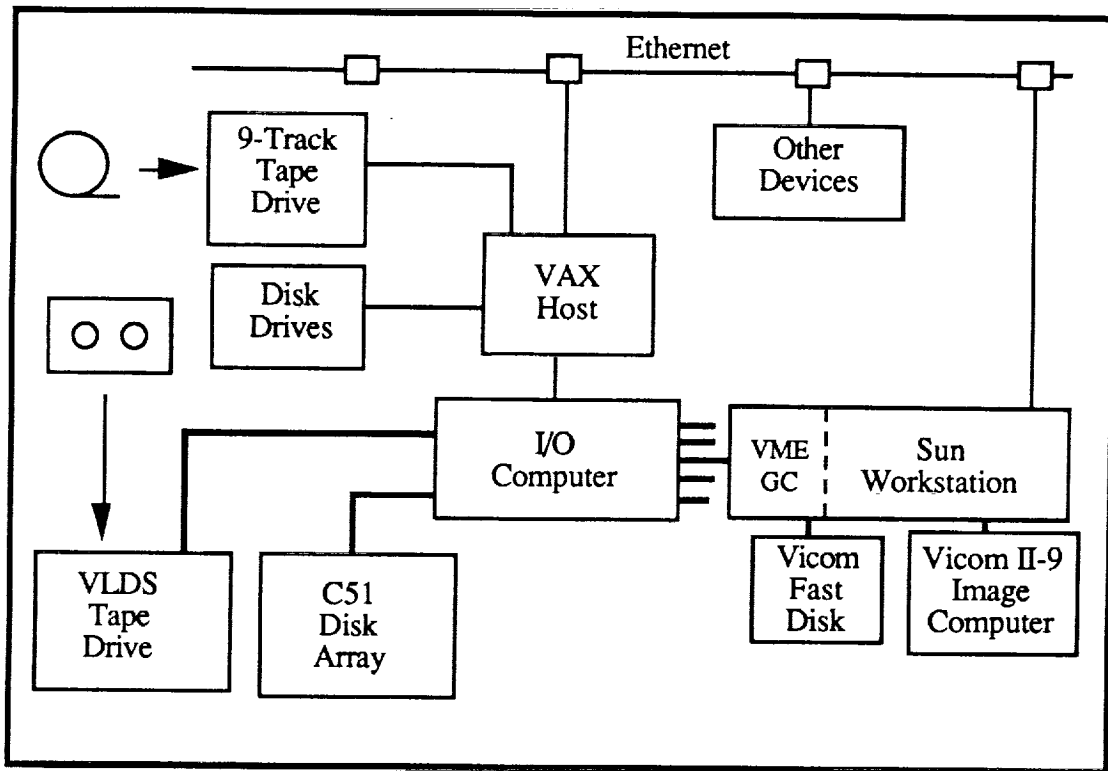


Figure 1. Data Storage and Retrieval System

Component Performance

From a performance standpoint, the C51 disk is the fastest peripheral on the DSRS. It has been measured as transferring (across the VMEbus to high speed memory) in excess of 9 MB/s using transfer block size of one megabyte. The transfer block size is critical in determining the effective transfer rate. Figure 2 shows how the effective transfer rate varies with the block size used when transferring data from the C51 to the Vicom Fast Disk (C512FD) or vice versa (FD2C51). These transfers were done using files on the order of 500 M bytes of data. For the DSRS application, a block size of two megabytes (512 K words) was chosen. This final size was dictated more by the Vicom Fast Disk rather than the C51 disk. For this configuration, the C51 disk array has 2.5 gigabytes of formatted disk space.

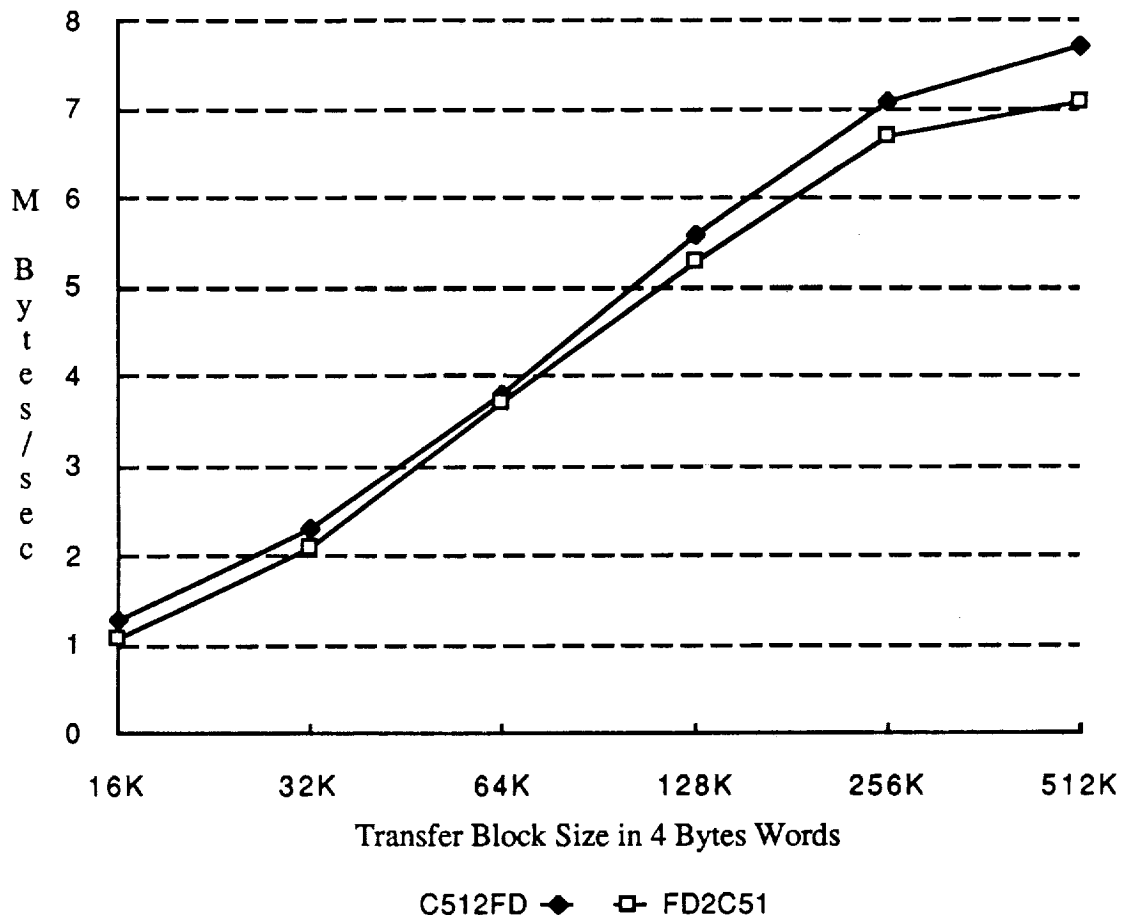


Figure 2. Transfer Rate as a Function of Block Size

The VLDS is a streaming tape drive and streams data at approximately 4 MB/s. Since it cannot start/stop like a conventional tape drive, it is imperative that it operate at its full transfer rate. In the read/playback mode, the VLDS can stop/restart taking approximately eight seconds to restart. Any significant mismatch in transfer rates between the VLDS and another device could slow the overall transfer rate down to eight kilobytes per second (KB/s). On the write/record side of the transfer, the VLDS will write "padding blocks" while continuing to stream. However, there is a maximum number of padding blocks that is (user) specified prior to the system halting (i.e., no restart). VLDS tapes that have padding blocks will have a natural degradation in transfer rate as well as in tape capacity. With no padding blocks, one VLDS cassette (super VHS T-120 cassette) will store approximately 5.2 gigabytes.

The VME gateway controller can effectively move data from the IOC's shared memory to the VMEbus on the workstation at rates in excess of 11 MB/s. The transfer rate also varies as a function of block size as shown in Figure 3. These transfer rates were




Figure 3. VME Transfer Rate as a Function of Block Size

obtained using the VME gateway controller in the Master mode. In the final configuration of the DSRS, the VME gateway controller is used in the Slave mode. Experience to date indicates that Slave mode transfer performance is similar to Master mode operation. For Sun 3/2XX workstations, the CPU exhibits a 190 microsecond bus timeout. This means that the VME GC cannot hold the bus longer than 190 microseconds after the CPU asks for it else it times out and the Sun operating system (SunOS) crashes. Since the VME GC operates in a release-when-done (RWD) mode versus release-on-request (ROR), it holds

the bus for the entire transfer. This practically limits the maximum block size to under one kilowords in Master mode operation (in order to stay under the 190 microsecond time limit). This limitation was overcome by using the VME GC in Slave mode and using a fast master controller (Maximum Strategy disk controller) operating in a ROR mode to effect the high speed transfer.

While not part of the DSRS, the Vicom Fast Disk is the other key device when examining end-to-end file data transfer. The Vicom Fast Disk is actually a disk array made by Maximum Strategy, Inc. (Strategy One Controller). In this configuration, the disk array contained approximately 6 GB of formatted space. The hardware is unchanged from the original product. However, Vicom has written a device driver and a file management system for it. The Vicom Fast Disk was rated at running at eight MB/s. While Vicom literature indicates 12 MB/s burst and eight MB/s sustained, in reality, 12 MB/s burst only occurs during the first megabyte of data (since it comes from memory and not the disk array). The eight MB/s sustained rate applies to the transfer once the disk heads are in position and is "streaming" data. In reality, large block sizes were needed to maintain transfer rates near eight MB/s.

Key Challenges

While all the hardware components were available off-the-shelf, no software was available to allow the components to function as a system. The key integration task was to develop the software and ensure that device performance would not be impeded by this software. A key challenge was to ensure that the VLDS tape unit operated at its full 4 MB/s capacity since it functions as a streaming tape drive. As indicated earlier, running the VLDS at slower than 4 MB/s would significantly slow down the total transfer time. Paramount to the development effort was to keep all overhead to an absolute minimum. Another challenge was to keep all software development at a high level - not develop any assembly language or microcode. Efficient use of library routines was essential. In order to promote a high level of portability, all routines developed on the Sun workstation had to be written in the C language and interfaced to existing Vicom device drivers. No kernel modifications were allowed as well. Since the software effectively controlled key hardware components directly, it was extremely difficult to debug since the typical error message was "bus error" (followed by a system crash). The use of a bus logic analyzer was essential to do problem identification and debugging. During this development effort, several key problems were discovered and fixed on the VME gateway controller. All problems dealt

with the use of the VME GC in slave mode operating at high (~12 MB/sec) speeds. Software modifications at the microcode level (in Aptec's software) were also made to get the system operating properly. These changes have now been incorporated as part of Aptec's baseline.

VLDS Tape Unit

The VLDS was the first peripheral to be interfaced to the IOC-24. By using a dual buffering scheme to keep data transferring between the VLDS and the shared memory, the effective transfer rate of 4 MB/s was easily maintained when sending data to the IOC's memory. When the VME gateway controller (VME GC) was added to the IOC, data was then made to flow from the VLDS, to the shared memory, to the VME GC Input/Output Processor (IOP), to the VME GC, and then on to a high speed memory card on Sun 3/260 workstation with no delays. Later when the Vicom Fast Disk was added to the Sun workstation, it became apparent that the Fast Disk needed large (one megabyte or larger) transfer blocks in order to maintain high throughput. Since the VLDS reads and writes in principal block increments of 65,536 bytes (64 KB), a buffer size mismatch needed to be fixed. The dual buffering scheme had to be modified to accommodate 64 KB buffers on the VLDS side and 2 MB buffer size on the VME/Vicom side. This was accomplished by using multiple VLDS buffers adding up to the (two MB) VME side buffer, then taking into account partial buffers and last buffer anomalies. With this approach, it became possible to transfer files from slow devices such as the Sun local disk to/from the VLDS at high data rates (as fast as the slowest device) for files up to 2 megabytes in size with no degradation in performance. Larger files could also be sent but the VLDS start/stop action would cause a degradation in performance.

C51 Disk Array

Interfacing the C51 disk array into the DSRS involved making a key decision regarding its use. The C51 could be used in a dedicated manner, i.e., used by a single process until completed, or shared like a disk server. Used in a dedicated manner, the performance would be optimized and the software would be easier to develop. The major drawbacks were that the disk array would not be shareable between processes and only contiguous files would be supported. This last condition was quite restrictive when there is 2.5 GB of disk space and file sizes of one to two GB could be expected. The contiguous requirement would prevent files from being written even though the space may exist due to

disk fragmentation or possibly bad disk blocks/sectors. The DSRS configuration uses the more complex VAX/VMS file management system (QIO) to create and manipulate files on the disk array. This allows files to be written out with multiple extents if needed. Another advantage of this approach is that the disk subsystem is shareable by different processes. Thus, a lengthy transfer that may take 3 minutes can function at the same time that another user is accessing a small file on the same disk, without having to wait for the first transfer to complete. From a VAX user viewpoint, the C51 appears as a standard disk (though non-system bootable) and functions, such as file transfer (ftp) and copy, can be used to transfer data from VAX based peripherals or other peripherals attached via the Ethernet LAN. Transfers involving the C51 also use a similar dual buffering scheme to maximize the transfer to/from the IOC's shared memory then on to the destination device.

VME Gateway Controller

The VME GC was, by far, the most challenging piece of equipment to understand and integrate into the system. A series of protocols were developed to allow the DSRS to communicate with the target workstation. The first protocol (command protocol) involved sending a command and appropriate parameters from the Sun workstation indicating what transfer needed to be done with what files. The second protocol (information protocol) involved communicating information regarding file size, transfer block size, and size of the last transfer. This information was critical to ensure both sides (DSRS and workstation) knew exactly what and how much data was being sent. Finally, the transfer protocol involved the low level "handshaking" needed to keep the data transferred in proper synchronization, i.e., ensure that source and destination devices were ready to receive the appropriate blocks of data. From an ISO networking viewpoint, these protocols fall into the application layer. They effectively establish a means of communications at a high level between a workstation and the file server (DSRS) for the transfer of a file. From a logical viewpoint, the VME GC is set up like a data structure featuring a first-in-first-out (FIFO) location, a mailbox area for small messages, and a set of 16 registers. The base address to this logical structure is user programmable and can exist anywhere within the workstation's memory space. Physically, the VME GC is also user programmable and can exist anywhere within the VMEbus 32-bit address space (barring conflicts with existing devices). The FIFO is used to transfer large blocks of data between the workstation and the DSRS. The mailbox region is used to pass file names, file sizes, and other miscellaneous pieces of information. The registers perform all control functions including setting the direction of transfer, number of bits per word, FIFO full/empty indication,

semaphore acknowledgements, etc. The DSRS is set up with a predefined data structure. The Sun workstation also sees this same data structure and communicates with the DSRS by reading and writing to these memory locations as if the DSRS was local to it. By using this approach, different VME based workstations can be integrated with the DSRS with minimal difficulty. All that is needed is an understanding of how to memory map to a specific memory location and a VME memory device driver (both commonly standard with any workstation/operating system). An interface guideline document describing both the hardware interface requirements (for the VME GC board set) as well as a detailed description of the above protocols (for a software interface) has been published.

Typical Data Flow (C51 to Vicom Fast Disk)

Once a user on the Sun workstation "launches" a transfer function, the VME IOP spawns a request to the C51 IOP to initiate the file transfer. Four megabytes of shared memory (organized as two 2 MB buffers) are allocated in the IOC. The C51 IOP's task is to fill each 2 MB buffer and mark the empty semaphore flags as "full" upon completion of writing a buffer. At the same time, the VME IOP reads the same buffers (when the semaphore indicates that the buffer is "full") and flags the buffer "empty" upon completion of the read function. The two processes run concurrently switching buffers as necessary to keep a steady flow of data moving. When the VME IOP reads the 2 MB buffer, the data is simultaneously transferred to the VME GC which has previously initiated "handshaking" with the Sun workstation CPU. In the meantime, the Sun CPU has transferred control to the Vicom Fast Disk controller which masters control of the VMEbus and extracts the 2 MB of data from the (FIFO address location of the) VME GC and transfers the data to its input buffer and subsequently to the disk array.

This single cycle which involves transferring data across three busses (DIB, OPENbus, and VMEbus) with "handshaking" and resource arbitration is executed 500 times (to transfer a gigabyte of data) and runs at 96% of the maximum burst speed of the slowest link.

Software Architecture

The software on the VAX was set up as a single program running as a server listening to commands (via a semaphore register) that it might receive from the Sun workstation. In actuality, the VME IOP has software constantly checking one of the

registers to determine if a workstation wants to deposit a command into the mailbox. Once a command is received, the VME IOP is allocated and cannot be used by another device until the command has been completed. The VME IOP "downloads" the appropriate code and executes all transfer routines or spawns appropriate routines to accomplish the commands. Upon completion, the semaphore register is setup for the next request. During the transfer of information, if any substantial delays occur, either side will timeout and revert back to listening mode. This allows the server software to recover from errors that may occur on the workstation. Using this approach with multiple VME IOPs and VME GCs would allow multiple VME based workstations to access the DSRS resources in parallel, limited only by the device speeds, the amount of shared memory, or bus bandwidth within the IOC-24. Over 60 STAPLE routines/procedures were written to support this server software.

The software on the Sun workstation was written as a series of short C routines that effect a transfer from one device to another. A typical routine name was vl ds2fd implying (in this case) transfer of data from the VLDS to the Vicom Fast Disk (FD). Following the command, the source and destination file names would be included as parameters. Ten of these routines were developed to transfer data in all conceivable directions. Over 30 routines/procedures were written to support the Sun based software.

Due to the development environment and the tools that were available, the entire system integration and software development effort for the DSRS took less than four months (with a staff of two): Interfacing the Sun/Vicom system took an additional two months (although hardware problems precluded use of the system for almost half of that time).

Conclusion

In the end, the DSRS successfully transferred one gigabyte of data from the VLDS to the Sun/Vicom Fast Disk in 4 min. 13 sec. + 25 sec. for tape setup and rewind. This translates to an effective transfer rate of 4 MB/s during the transfer (which is the streaming rate for the VLDS). The file transfer from the C51 disk array to the Sun/Vicom Fast Disk took 2 min. 15 sec. using 2 MB transfer blocks. This translates to an effective transfer rate of 7.7 MB/s (out of a theoretical maximum rate of 8 MB/s as constrained by the Vicom Fast Disk). File transfer was also performed between the C51 disk array and a device (a local disk attached to another Sun workstation) over Ethernet while, at the same time,

transferring a file from the Vicom Fast Disk to the C51. Although the performance was slightly degraded due to the sharing of the C51 disk, it was not noticeable using small files (less than 10 megabytes).

All requirements were fulfilled using commercially available off-the-shelf components with a relatively small software development effort. The system is now operational and is being used to store and retrieve large files on VHS cassette tapes and can load the Sun workstation (Vicom Fast Disk) in minutes versus the many hours it used to take when using 9 track tapes.

Data Storage and Retrieval System

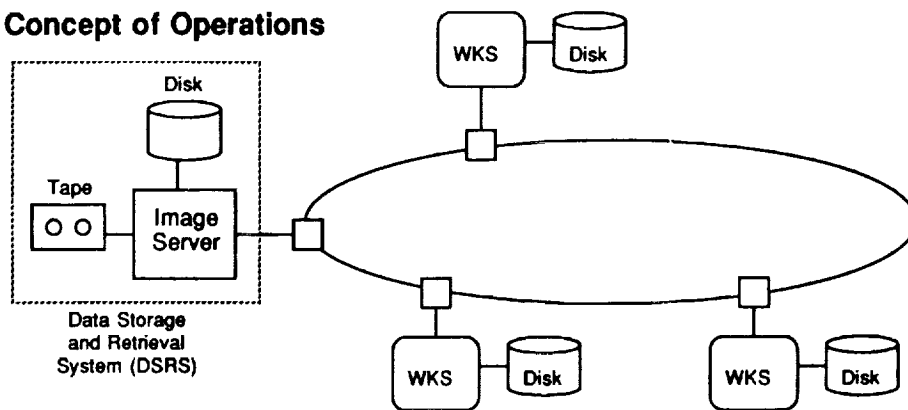
Glen Nakamoto

24 July 1991

MITRE

Background


Concept of Operations



- Support evaluation of operationally-oriented softcopy imagery exploitation
- Two sessions per day; four hours per session
- Preload images into workstation prior to session
- Ad Hoc access to any image stored on the server

MITRE

A Key Problem



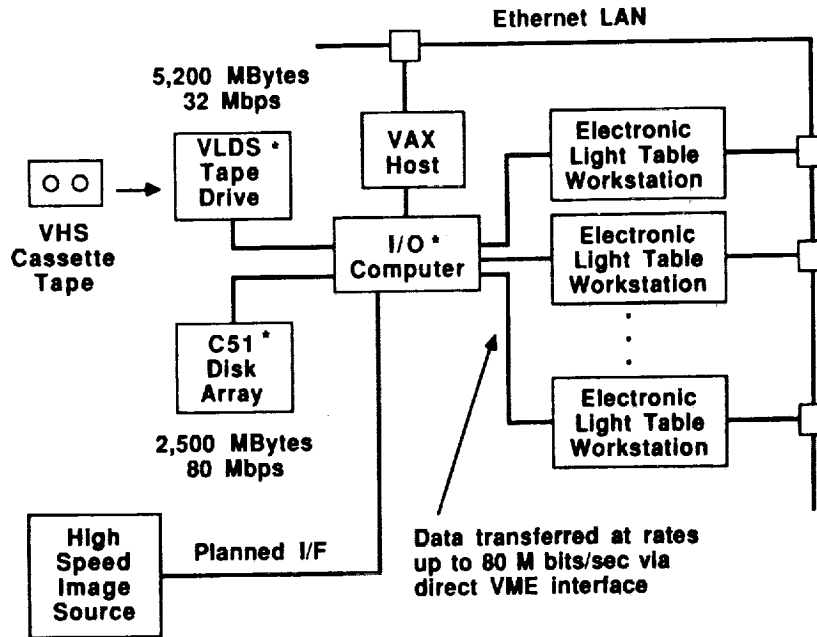
MITRE

FY90 Goal

- **Develop a prototype data storage and retrieval system**
 - Support image files 20 - 30 times larger using single portable media
 - Transfer files at rates 90 times current capability
- **Establish interface guidelines for future workstations**
 - Multiple standards (Physical, networking, application)
 - Portable media (Tape, format)
- **Interface Initial commercial workstation**

MITRE

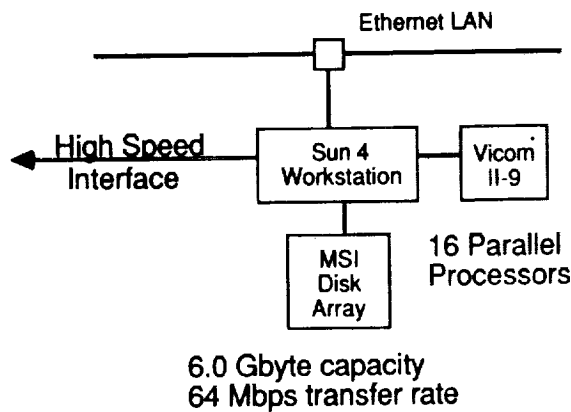
DSRS Architecture



* These items fit in a rack space of 27 inches.

MITRE

Electronic Light Table Workstation

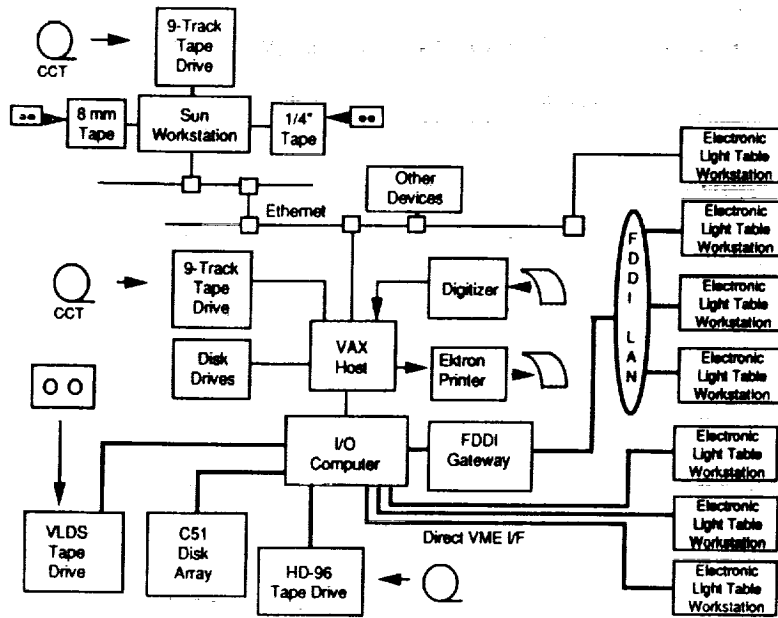


- Electronic Light Table Functions
- Integrated Database -
 - Text, Graphics, Imagery
- Large Image File Manipulation
- Real-time Decompression
- 2K by 2K display size (search)
- High Speed Image Transfer
 - (64 Mbps)
- Off-the-Shelf Hardware

Previously Pixar

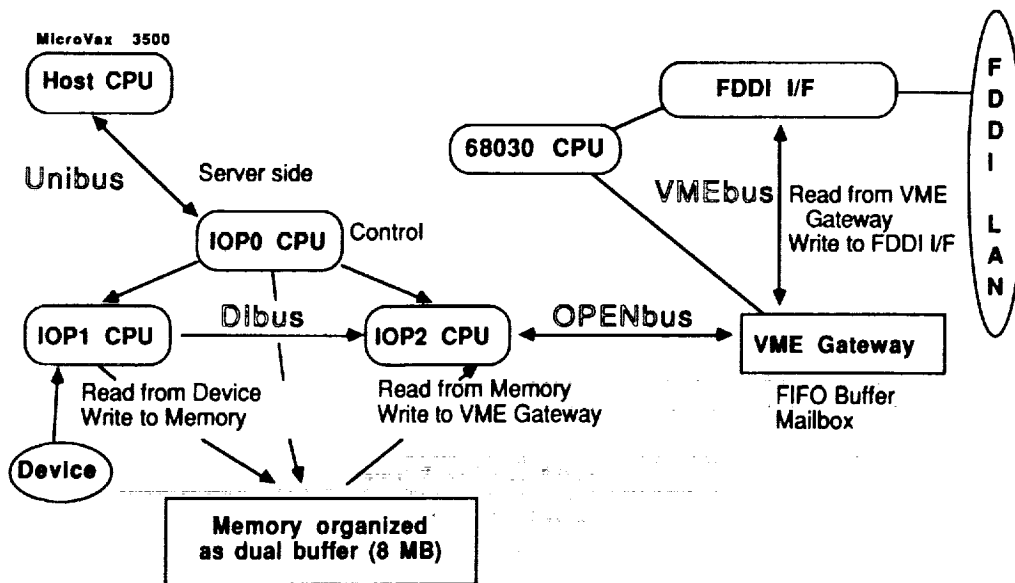
MITRE

Image Server Configuration



MITRE

Software Architecture



MITRE

Measured Transfer Rates

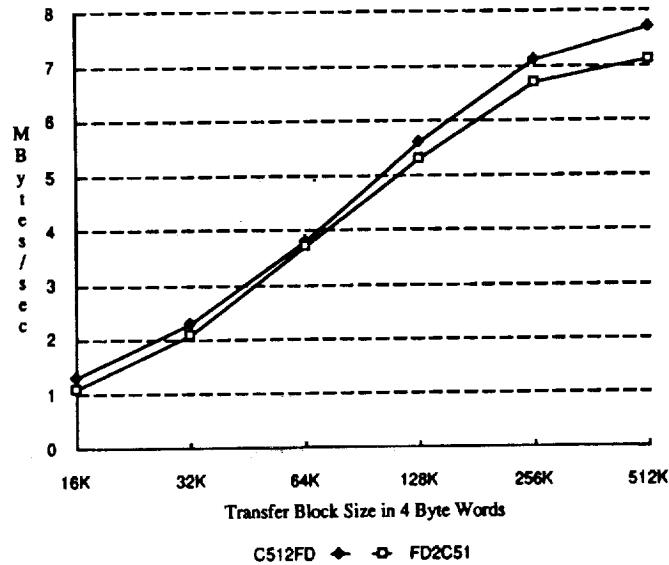
<u>Network</u>	<u>Burst Rate</u>	<u>Effective Throughput</u>
Ethernet	10 Mbps	2.4 Mbps
FDDI	100 Mbps	2.5 Mbps
UltraNet	125 Mbps	20.0 Mbps
DSRS	96 Mbps	93.6 Mbps
IOC based FDDI	100 Mbps	48.0 Mbps (not measured)

Notes:

Sun memory to Sun memory transfer with no network contention.
 SunOS 4.1 (Sun 3/260 with 8 MB memory).
 Sun provided TCP/IP software.
 UltraNet used TP4/IP with board level protocol processing.


MITRE

C51-PFD: Transfer Rate vs. Block Size



MITRE

VME-MM Transfer Rate



MITRE

Summary


- Requirements for DSRS finalized in September 1989
- IOC based system purchased in October 1989
- System delivered on 18 January 1990
- The DSRS was developed and delivered 18 June 1990
- Integration of the Sun/Pixar workstation was completed by 9 August 1990
- Improves transfer times 90:1
- Improves storage approximately 100:1
- Allows search oriented experiments to be conducted
- Improves the management of an image library
- Promote standards and interface guidelines

MITRE

FY91 Goal

- Integrate an FDDI network into the DSRS
 - Develop an FDDI gateway for the DSRS
 - Initially support TCP/IP protocols
 - Provide capability to install other protocols
 - Provide capability to support multiple gateways per IOC
 - Maintain maximum performance end-to-end
- Upgrade IOC-24 to IOC-100
- Upgrade 2.5 GB Disk Array to 7.5 GB capacity
 - Provide means to address greater than 32 bits

MITRE



THE CHALLENGE OF A DATA STORAGE HIERARCHY

**Mr. Michael Ruderman
Vice President for Marketing
Mesa Archival Systems, Inc.**

MR. RUDERMAN: Good morning. I have the pleasure of standing between you and lunch. So, I will try and speak quickly and get through as much of this material as we can. I'll tell you a little bit about Mesa Archival Systems first.

We are a small private company located in Boulder, Colorado. We were founded as a technology spin-off of software originally developed at NCAR, the National Center for Atmospheric Research. Currently, their version of this software is managing 19 terabytes of data on over 102,000 3480 cartridges. They are in the process of moving from 3480 to 3490 tape cartridges, but it is clear that this represents a large amount of data from a variety of systems and is composed of fairly large set of bitfiles.

Something else you ought to know about Mesa Archival is that we are part of the winning team of the recent NASA Goddard mass storage award that hopefully, barring any obstacles--legal or otherwise--you will all be able to see running here at NASA Goddard in the next couple of months. We are looking forward to that. (Showing of viewgraphs)
MR. RUDERMAN: Now, before going into the storage hierarchy aspect of the system, I will put it in some perspective with a brief overview. What is the system in general?

It is a data archiving system; it fits into the kinds of environments that all of you have or are working towards, whether you are vendors or users, with multiple heterogeneous client systems networked to a central archive server.

That is the simple explanation.

(Change of viewgraph)

MR. RUDERMAN: We can picture our system, as illustrated. Obviously, any kind of client system on the network can be connected through the Unix interface of our system. We are strictly a software system sitting on a mainframe, managing the data into permanent file storage. We do not have time today to go into all the details of how we do everything. But suffice it to say, it is a central archive manager, with standard network access and it is a standard commercial product--that is very important.

We came from a lab environment; however the version that is now available is not lab code. It is commercial product code. It sits in a high-performance computer, accessible from multiple processors.

(Change of viewgraph)

MR. RUDERMAN: And just to reiterate, when we took this code from NCAR and designed it for commerciality, it had to be designed for change. Much of what you are hearing during this conference is a lot of information, from both users and suppliers, about a lot changes taking place in the area of mass storage devices.

We believe one of the most important ideas to keep in mind is that data must outlive any device or any media; the data is more important than any of those things.

The hardware is going to keep changing; and so, you have to have a software control system that is independent of any hardware system. We believe in that.

(Change of viewgraph)

MR. RUDERMAN: Let's take a look at the next level into the system. As I said, we support standard networks, TCP/IP with FTP access and NETEX with User Access client support. We put no code on client systems. The network access server provides the access; the data library access manager provides all the client and system communication dialogue. As demand warrants, support of additional new network protocols will be added easily.

The archival object database is the heart of the system, and here we have implemented our own object oriented analysis. At the back end of the system there are multiple storage servers, and of course, system administration. This is where we will focus on a little more detail.

Before getting into that, I want to mention that we are an active member of the IEEE Mass Storage Reference Committee and Storage Subsystem Working Group, - we have been involved with it from the very, very beginning. In fact, NCAR was one of the originators of the whole idea back in the mid-1980s.

(Change of viewgraph)

MR. RUDERMAN: This is our version of the IEEE components model, and what I want to talk about today is the storage system side of it. As you can see, we have separate data paths between bitfile movers--the bit movers from the client side and the server side--and here we have multiple storage servers, which is the area that we want to look at. The issue is how do we implement that?

(Change of viewgraph)

MR. RUDERMAN: With the variety of devices available--and you have all been hearing about a whole bunch of them in the last couple of days--the variations cover access speed, permanent versus temporary storage, capacity, cost--cost is a major factor.

I think half of you could go home--those of you who are hardware vendors--if there were infinite rotating disk storage available at a very low price, we probably wouldn't be worrying about this conference too much. Life would be simple, but it's not that way.

And products keep changing. What may seem like simple changes from 3480 to 3490 technology have implications--big implications. New products--D-1, D-2, optical disk, optical tape, etc.. Our objective on the software side is to keep the data independent of any of these hardware changes.

(Change of viewgraph)

MR. RUDERMAN: Now, this morning, you heard Bob Coyne from IBM talk about the storage hierarchy and the issues and problems of a static storage hierarchy. Because there are now multiple devices being implemented in single archives, both for cost reasons, technical reasons, and user reasons, the traditional, what we call "static hierarchy" of primary and secondary storage alternatives, is just not sufficient.

The problem surfaces as soon as you try to figure out what you do with the third one you want to add--the third type of storage device. Where do you put it in the hierarchy? For different kinds of data it belongs in different places.

We have examined this; we have spent a fair amount of time on this over the past year or so. To say that we have solved it, is easy. What we will try and show you is how--the conceptual approach as to how we solved it. This is, in fact, what is available from Mesa Archival Systems today.

(Change of viewgraph)

MR. RUDERMAN: To give you a feeling for what I'm talking about a little more specifically, we can see that from the data library system, we can support magnetic disk, optical disk, cartridge tape, helical scan magnetic tape--all of these kinds of storage devices, multiple devices, heterogeneous devices--simultaneously.

You cannot do that very efficiently with the traditional static hierarchy of a fixed physical system. So, we have developed a **structured hierarchy**, which gives multiple views of those physical storage options. It is dynamic, and it must be able to be varied by user, by bitfile--down to the bitfile level--by class of data, by accounts, by any combination of categories to be derived for each system.

Now, that is an easy thing to say, but it's not so easy to implement. Let me just make a note. I see some of you looking through the proceedings and not necessarily finding all the slides identical. We recently updated some of these and created some new ones for today; and I apologize for any confusion.

We didn't get them sent in, but we will get a new set available for the followon proceedings book. So, you might as well just ignore what's in there at the moment. (Change of viewgraph)

MR. RUDERMAN: It's easy to say, not so easy to do.

The first thing we did is separate the archive system from the storage system conceptually. Bitfiles come into the system. There can be one or more images of them, depending upon the attributes of the bitfile. You may want to keep a copy on disk, put a copy on optical; you may want to keep copies on multiple devices for various reasons.

Each image of a bitfile, unless it is very small, will have multiple fragments; and these fragments reside on specific types of media. The media gets broken down further into packages and surfaces. I'll show you later how that all fits together.

But basically, the first thing to do is to separate the archive system from the storage system and implement the whole concept from the point of view of object-oriented analysis and design.

That makes the use of it simple; it makes the coding of it not so simple, but it has been accomplished.

(Change of viewgraph)

MR. RUDERMAN: Now, let's take a look at this from an object orientation. We have a client file in the system. This client file talks to the bitfile server, requesting services.

The bitfile to be archived will request storage from various storage options. Here we are just showing two; there can be any number of storage options based upon what the physical system has installed.

(Change of viewgraph)

MR. RUDERMAN: As we take a look at this, we see the kinds of objects that sit under each of these categories. In the client file system we have the directories and files and users and groups. The bitfile archive has the bitfiles and accounts that they belong to, if that is appropriate, and the templates.

The templates are very important. The templates describe which of the storage options are available for a particular bitfile or, in other words, which of the options that

are physically on the system are the permitted options in this instance. In this way, the storage hierarchy is constructed for each different need.

All of the communication between the various tasks is very standard client server protocol, make a request, receive a reply; it makes use of the system very simple, as well as modification and administration of the system--very, very simple. (Change of viewgraph)

MR. RUDERMAN: This is the simplicity of the storage option view from the bitfile. The bitfile doesn't see much more than that.

However, below this point, we have analyzed and organized all of the information that we believe we need to know. The head of development doesn't like me to say "all" about anything; he's very sensitive about that, but I believe it's all of it that we need to know, and it works!

We have broken this down to fixed and removable media. Disk is obviously fixed, with a one-to-one media/device relationship. Most of the removable media that we are dealing with, on the systems that we support today, use 3480 tape interface. Obviously, a 3480 cartridge tape is a one-to-one; however, the implementations of both helical scan and optical disk have multiple logical 3480 volumes on each physical volume.

The reason for this analysis and breakdown is to make it simple, to be able to both move data and make requests of the system, such as change devices in and out. We are not aware of any devices that we cannot put into this organization, and with it, we will make the subject of managing them very, very simple.

Now let me try and give you an example of what we have been able to do with this approach.

(Change of viewgraph)

MR. RUDERMAN: If we take a look at a particular system, let's say this system has these options available to it. This includes the disks, helical tape, optical tape, and standard cartridge tape.

A particular bitfile in the system may have attributes such that it can reside on disk or helical tape or optical tape--it doesn't matter. If there is an image sitting on disk at a point in time and this disk reaches its threshold and an event is triggered such that it needs to be migrated or "scrubbed," as we call it, the option is: Do I put it here on helical tape or do I put it here on optical disk?

We have built the system and designed the system in such a way that, as we can get more information feedback from the hardware system, we can make that selection much more sophisticated. What I'm driving at here is that the system manager may have said: Because helical is faster, the next best choice for this file, when you migrate from the disk is to go to helical tape.

If at a particular moment in time the helical system has large cues and is overloaded, from a performance point of view, if the optical tape were available, it might make more sense at that point in time to go there.

The kind of information required to be able to make those selections automatically, transparent obviously to the user and transparent to the bitfile, will require some additional information that tends not to be available today from most robotic suppliers.

And that gets us into the challenge that we see in dealing with robotic storage devices, and we have dealt with quite a few of them. When there are multiple and/or shared robotic systems in a single archive, there are some potential problems that we think most

of the robotic manufacturers have not anticipated. They tend to think they are the only devices on the system and they have not anticipated the need for programmatic dialogue about their own status.

But I believe that you will find that the user community is going to demand multiple heterogeneous storage systems to be installed, and the ability to install and remove and add and change any device at any time without taking down the archive demands the approach we have taken.

So, what we have found is that the ability to do reads and writes is easy; anybody can do that. But the lack of a standard programmatic interface between software archive systems and the storage devices is a problem. The one that we have implemented, and we are pushing to become IEEE standard is the standard client server protocol. It's the same as the ISO managed object interface for networks. It's the same protocol we use internally between all of our task communications.

(Change of viewgraph)

MR. RUDERMAN: This is something we have all seen in any client server relationship. With a storage server in this case, the client would make a request to the server, and get a reply and/or be triggered with an event. (Change of viewgraph)

MR. RUDERMAN: For instance, if an automated media library is shared, then how do we know if a package we need is not being used elsewhere or has been ejected from the device?

Because we are dealing with high performance requirements and massive amounts of data, we would like to know that before we issue a read or a write so that we don't hang up the archive system. The objective is to improve performance, not issue commands to any device when a piece of media or the device itself is not available.

This is what we are looking to see. We have a design for it; we have implemented it with certain devices, and others of you in the vendor world need to think about how you are going to coexist in this environment.

(Change of viewgraph)

MR. RUDERMAN: Just summarizing. The devices currently supported, that we have experience with, the 3480 cartridge, the STK silo, the Memorex tape library, Dataware optical disk, and Masstor helical scan tape.

We see in the future adding additional archival devices. We are very interested in lots of the new devices, in D-1 and D-2 areas especially. We also will be expanding network connectivity to new standards, as they emerge.

We intend to be very oriented towards standards. Additional operating system support is probably one that a few of you might be interested in. For those who don't know, right now we are MVS-based. We definitely have plans for expanding beyond MVS to Unix. A fundamental objective must be that all changes need to be transparent to the client systems.

(Change of viewgraph)

MR. RUDERMAN: When we talk about adding new operating system support, what we are really saying is that we don't see the world of data archiving suddenly and totally abandoning MVS for Unix. MVS systems are very, very powerful for moving massive amounts of data, simultaneously.

We see the next version of our software to be able to have multiple distributed servers, both for bitfile serving and for storage servers and multiple processor hosting, including both MVS and Unix.

That's what we wanted to communicate with you today, and I'll think we'll make it to lunch. You may ask some questions if you'd like.

DR. FREESE: Thank you, Michael.

(Applause)

DR. FREESE: Questions, comments, discussion from the floor?

MR. SAVAGE: I do have one question. (Inaudible) DR. FREESE: Could you paraphrase that?

MR. RUDERMAN: Yes. The question had to do with, I believe, an interpretation of the NCAR system, which was that the archive manager was really just directing the requesting system or telling it where the data was, as opposed to actually shipping it to it.

Yes, what you are referring to is the fact that there is direct data transfer between the IBM disks and the Cray at NCAR. This is done for performance reasons, so the data does not have to be sent through the IBM mainframe.


MR. SAVAGE: The IBM machine would look it up, find out what disk it was, and ship that information to the Cray. The Cray would then create a channel program and send it down over that channel to actually directly read the disk.

DR. FREESE: Any other questions or comments?

Data Archiving.
Michael Ruderman
Mesa Archival Systems, Inc.

MESA
Archival Systems, Inc.

Computing Environment



MESA
Archival Systems, Inc.

NATIONAL CENTER FOR ATMOSPHERIC RESEARCH (NCAR)

Atmospheric and oceanographic research

Initiator of IEEE Storage Model

Status

- **Operational since 1986**
- **2,000 users**
- **102,000 3480 cartridges 5/91**
- **~19 TB, growing at 6 TB/year with Y/MP**

MESA
Archival Systems, Inc.

User Needs

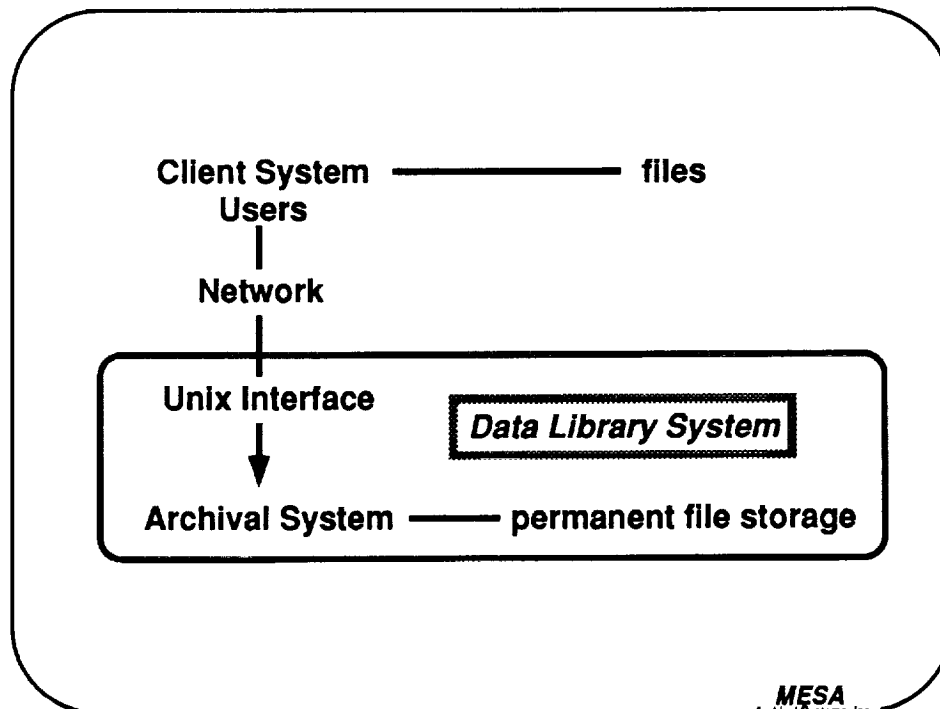
- **Data Integrity**
- **Consistent, familiar interface**
- **No system dependencies**
- **Accessible from everywhere**
- **Reduced local storage mgmt**

MESA
Archival Systems, Inc.

System Manager Needs

- Data integrity
- Storage hierarchy
- Mass storage alternatives
- Ability to deal with change
- Accountability by user
- Performance
- Reduced operations cost

MESA
Archival Systems, Inc.



MESA
Archival Systems, Inc.

Data Library System

- **Central Archival Data Management Facility**
- **Standard Commercial Product**
- **High Performance Computer**
- **Access from Multiple CPUs**
- **Expandable, Device Independent Architecture**
- **Standard Operating System**
- **Standard Network Software**


MESA
Archival Systems, Inc.

System Environment

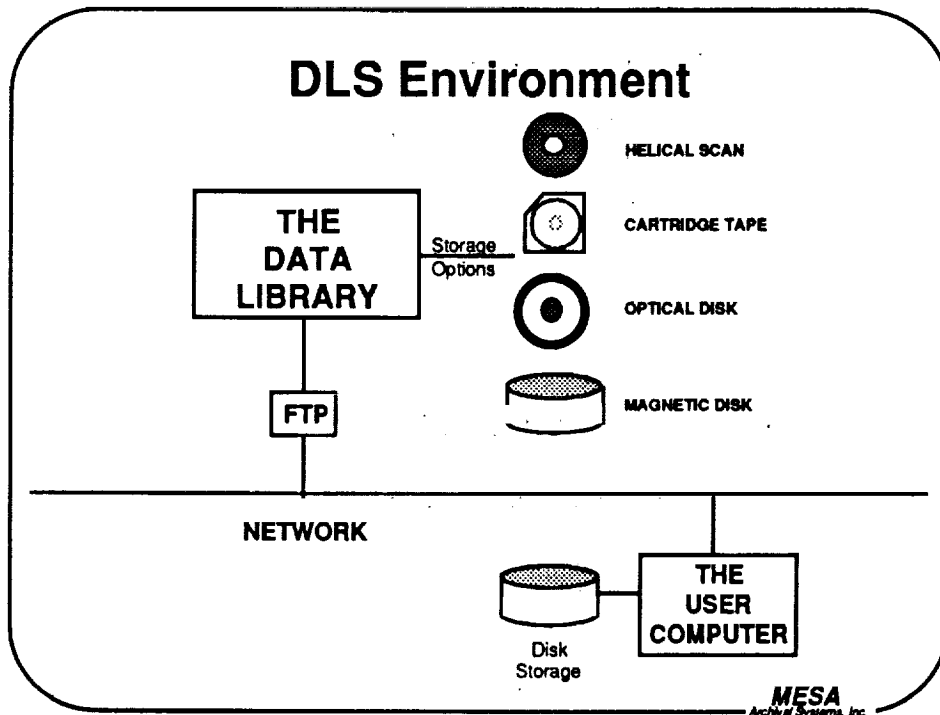
- **MVS/XA**
- **One dynamic user SVC**
- **SMP/E Installation**
- **Security software interface**
- **Tape management system interface**

MESA
Archival Systems, Inc.

Data Library System



DLS Environment



Data Library System

- **Network Access Servers**
 - **FTP or User Access**
 - **Unix File System appearance**
 - **Gateway to DLS**
- **Archival Object Data Base (AODB)**
 - **Powerful Facilities**
 - **Object Orientation**
- **Storage Servers**
 - **Uniquely Mounted Media**
 - **Variably Mounted Media**
- **System Administration**

MESA
Archival Systems, Inc.

DLS Features

- **Modular Implementation**
- **Client applications**
(Volume backup)
- **Resource accounting**
- **Security**
(Client - POSIX, System - MVS)

MESA
Archival Systems, Inc.

Networks

Protocols

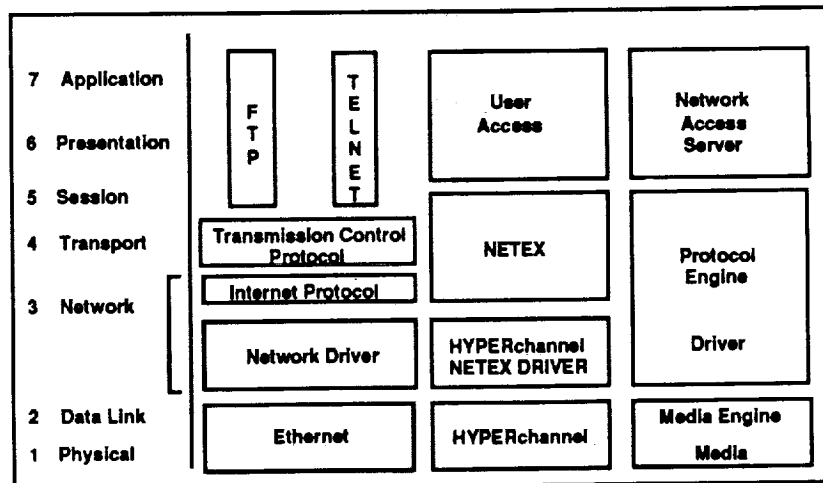
- FTP, User Access
- TCP/IP, NETEX

Networks

- Ethernet
- Ultranet
- HYPERchannel

MESA
Archival Systems, Inc.

OSI Model

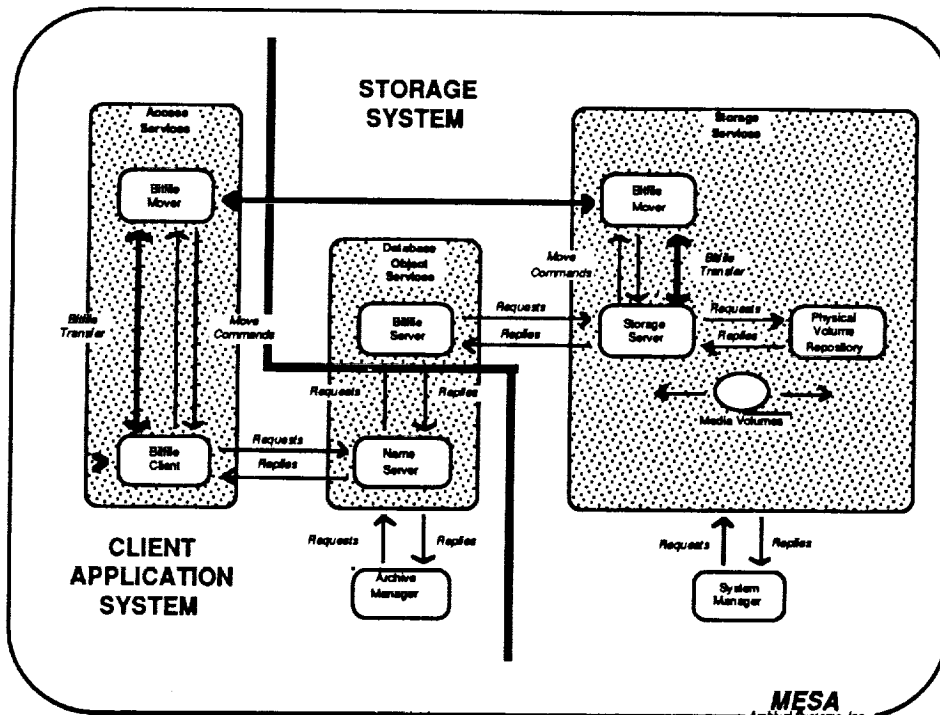


MESA
Archival Systems, Inc.

IEEE Mass Storage Reference Model

- Deals with Named Files - "Bitfiles"
- File Structure Insensitive
- Modularity of Design
 - Application Client
 - Bitfile Server
 - Storage Server
 - Physical Volume Repository
 - Bitfile Mover
 - Name Server
 - Site Manager

MESA
Archival Systems, Inc.



MESA
Archival Systems, Inc.

DLS and UNIX

- File Naming Conventions
- Directory Structure
- Command Syntax
- Security

MESA
Archival Systems, Inc.

Client System Examples

- APOLLO AEGIS
- CDC NOS, NOS/BE, NOS/VE
- CRAY COS, UNICOS
- DEC VMS, MICRO VMS, ULTRIX
- IBM MVS, VM, AIX
- PRIME PRIMOS
- SUN UNIX BSD 4.3
- UNISYS OS/1100

MESA
Archival Systems, Inc.

The Users View

- **User Capabilities**
- **Directory Organization**
- **File Security**

MESA
Archival Systems, Inc.

User Capabilities

- **Store a file**
- **Retrieve a file**
- **Examine the directory**
- **Other**

MESA
Archival Systems, Inc.

Client System Commands

FTP Interface

- *GET*
- *PUT*
- *DIR*
- *LS*
- *RENAME*

Other

- *USER ACCESS*
- *IMPORT / EXPORT*

MESA
Archival Systems, Inc.

Other User Capabilities

- **Add** Adds a DLS directory
- **Delete** Deletes a DLS file or directory
- **Copy** Creates a copy of a DLS file within the DLS
- **Help** Invokes DLS help facility

MESA
Archival Systems, Inc.

Ways to Organize User Files

- Simple (flat)
- Hierarchical (tree-structured)

MESA
Archival Systems, Inc.

Security

- DLS User Validation Password
- Owner, Group, World read / write access
- File read / write Password
- Account group, security & accounting

MESA
Archival Systems, Inc.

Multiple Mass Storage Devices

Differing User Storage Needs

Access Speed

Permanent Retention

Cost/MB

Interchangeability

Continuing Product Evolution

3480 ' 3490

New Mass Storage Devices


D1/D2

Optical tape


MESA
Archival Systems, Inc.

Archive / Storage System

archive system



storage system

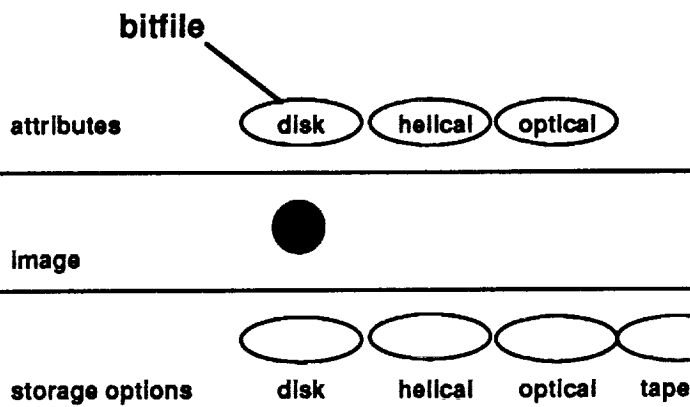


MESA
Archival Systems, Inc.

Constructed Hierarchy

- Multiple views
- Dynamic
- Vary by bitfile or user

MESA
Archival Systems, Inc.




MESA
Archival Systems, Inc.

Robotic Challenge

- Multiple / shared robotic systems
- 3480 service interface not sufficient
- No standard programmatic interface
- Client / server protocol recommended
(ISO managed object)

MESA
Archival Systems, Inc.



package

injected - in the robot
injectable - out of the robot

surface

mounted - in a drive
mountable - out of a drive

media

opened - has an active file
openable - has no active file

MESA
Archival Systems, Inc.

Archival Devices

- **IBM 3480/90 Cartridge Tape**
- **STK 4400 Cartridge Tape Robot**
- **MTC 5400 Automated Tape Library**
- **DataWare Optical Disk**
- **Masstor Helical Scan Tape**

MESA
Archival Systems, Inc.

System Implications

Media Orientation

Media/Device Relationships

Robot Awareness

User Profiles

MESA
Archival Systems, Inc.

Vendor Implications

Storage Management System

- **Must be unbound from File System**
- **Bitfile "Instances"**
- **Must be aware of different media, devices, robots**

Robotic Systems

- **Must be able to interact with SMS**
- **Media content and status**
- **Must notify "clients" of changes**
- **Emerging ISO standard for managed objects**

MESA
Archival Systems, Inc.

DLS / IEEE Storage Model

Active committee member

General compliance

Organization of physical volume repository

- **Media sets and pools**

Multiple archival devices

Direct I/O capability

Commitment to continued compliance

MESA
Archival Systems, Inc.

System Growth

- **Additional Archival Devices**
- **Additional Connectivity Products**
- **Additional Operating System Support**
- **Transparent to Client System**

MESA
Archival Systems, Inc.

Future Direction

Version 3

- **MVS or Unix**
- **Multiple processor hosting**
- **Support for Version 2 PVR**

MESA
Archival Systems, Inc.

User Factors

- **Prevalence of Unix**
- **Mass storage devices**
- **Software and people expense**
- **Integrity and performance**
- **Standard network support**

MESA
Archival Systems, Inc.

MVS

System integrity/availability

Proven I/O throughput

Wide range of I/O devices

Security

UNIX

Standard, familiar interfaces

Standard development platform

Standard networks

MESA
Archival Systems, Inc.

Implementation Factors

- **Multiple flavors of Unix**
- **Mass storage driver support**
- **Unix vendor commitment to performance**
- **Relationship with Unix vendor**

MESA
Archival Systems, Inc.

Benefits

Client System

- **Reduced disk and tape drive expenditures**
- **Reduced operational expense**
- **Reduced media expense**

MESA
Archival Systems, Inc.

Benefits

- **Improved data integrity**
- **Higher rate of data backup**
- **Increased data reliability**
- **Control of organizational data**
- **Ability to deal with change**

MESA
Archival Systems, Inc.

Implementing the DLS

- **Hardware Configuration Planning**
- **System Customization**
- **Installation Planning**
- **Training and Support**

MESA
Archival Systems, Inc.

DLS is a Central Archival Data Management Facility

- **I/O Server Computer - MVS**
- **Standardized Library Access from Multiple CPUs**
- **Standard High Speed Networks**
- **Expandable Device Independent Architecture**

MESA
Archival Systems, Inc.

Archiving Facilities

- **User / system initiated transfers**
- **Simple user interface**
- **Archival device independence**

MESA
Archival Systems, Inc.

Backup Facilities

- **List-driven backup**
- **Media clustering by expiration date**
- **VAX / Unix backup utility**

MESA
Archival Systems, Inc.

Mesa Archival Provides ...

- **Data Archival Product**
- **Data Archiving Expertise**
- **Archival System Integration Capability**

MESA
Archival Systems, Inc.

N 9 3 - 1 6 3 2 6

NETWORK ACCESSIBLE MULTI-TERABYTE ARCHIVE

NSSDC CONFERENCE
GODDARD SPACE FLIGHT CENTER
JULY 24, 1991

Fred Rybczynski
Metrum Information Storage

ROTARY STORAGE SYSTEM (RSS)

- T-120 Super-VHS media with 10.4 GB per cartridge
- Very Large Data Store (VLDS) tape drives
 - Sustained 1.0 MB/S with Read-After-Write
 - 1E-13 Bit Error Rate w/ error mgt software
 - 45 sec average file access time
- Robots with on-line capacity to 6.2 TB
- Bar code reader
- Data management computer
 - Data referenced by file name
 - Files organized by hierarchical directories
 - File and directory access protections
- Network accessible

Metrum
Information Storage
021891 drj

METRUM INFORMATION STORAGE SYSTEMS BASED ON STANDARDS

MEDIA

- World-wide VHS media standard
- T-120 half-inch tape cartridge
- Billions of cartridges sold

TAPE DRIVE

- SMPTE format VHS transport
- Millions of transports sold
- Commercial grade

ROBOT

- Commercial broadcast industry

Metrum
Information Storage
022481 drj

METRUM INFORMATION STORAGE STANDARDS - concluded

BAR CODE READER

- Standard formats

INTELLIGENT DATA MANAGER (COMPUTER)

- UNIX based
- POSIX compliance

NETWORK CONNECTIVITY

- Ethernet TCP/IP and DECNET
- HYPERCHANNEL
- FDDI
- GOSIP
- NFS

Metrum
Information Storage
[022401 .sh]


MANAGED DATA STORAGE GOALS

- Maximize data per unit volume, balanced by
 - Cost/MB of Media
 - Cost/MB of Technology
 - Required data rates
- Optimize media access
 - Unattended media access
 - Media ID verification
 - Speed of loading
- Provide ease of access to the data
 - Easily identify data to be stored/retrieved
 - Concurrent access from multiple computer systems

Metrum
Information Storage
[031801 .sh]

STORAGE HIERARCHY

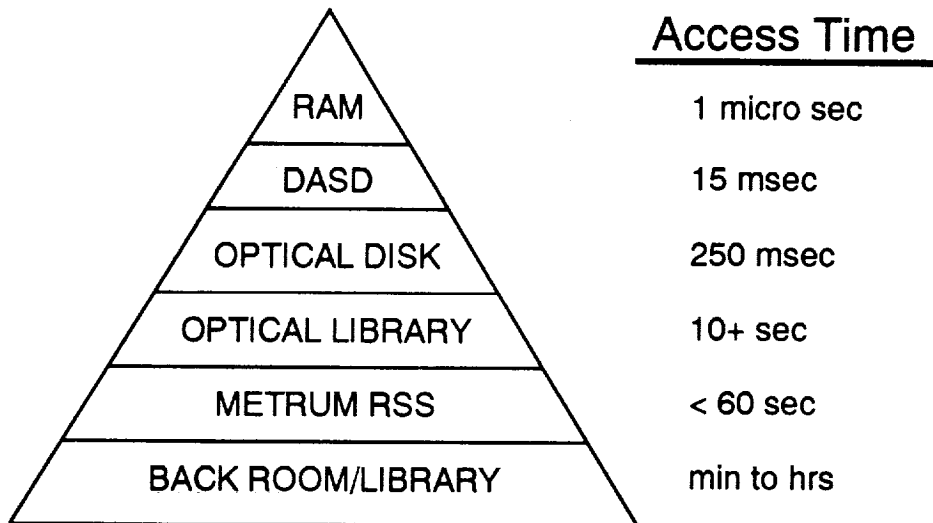
TYPICAL MEDIA PRICE/MB



Metrum
Information Storage
(110180-44)

STORAGE HIERARCHY


TYPICAL ACCESS TIMES




Metrum
Information Storage
(110180-44)

VOLUME COMPARISON

6 Terabytes



9-Track, 6250 bpi
2400 Ft, 170 MB
35,000+ Reels
850+ Sq Ft



T-120 Cartridge
10.4 GB/Cartridge
600 Cartridges
4 Sq Ft

Metrum
Information Storage
[p40191 sh]

VOLUME COMPARISON

One T-120 Cartridge (10.4 GB) Equals

- 61 nine-track reels at 2400 ft/reel, 6250 bpi and 170 MB/reel
- 50 cartridges of 3480 at 200 MB/cartridge

One Cubic Foot of Storage Can Contain

- 12-15 nine-track reels
- 60 cartridges of 3480 (equals 70 nine-track reels)
- 30 T-120 cartridges (equals 1,800 nine-track reels)

Metrum
Information Storage
[p40891 sh]

RSS FEATURES

- 6.2 TB in < 20 sq ft
- System cost < 8.5 cents per MB
(RSS-600, 2 drives, 600 cartridges = \$530,000)
- Media cost < 0.3 cents per MB
(\$30/cleaned and certified cartridge)
- Unattended operation
- Network accessibility
- Read-After-Write reliability (1E-13 BER)

Metrum
Information Storage
{031861 sh}

RSS DATA ACCESS

- Archive system accessed via network
- Access authorization screening (UserID & Password)
- Entire archive looks like a single huge disk
 - Data accessed by file name
 - Multi-level hierarchical directories
 - Selective directory & file protection
- File access is random (cartridge rewind not required)
- Avg file access is 45 sec for loaded cartridge
- Cartridge load in < 8 seconds
- Intelligent queue management

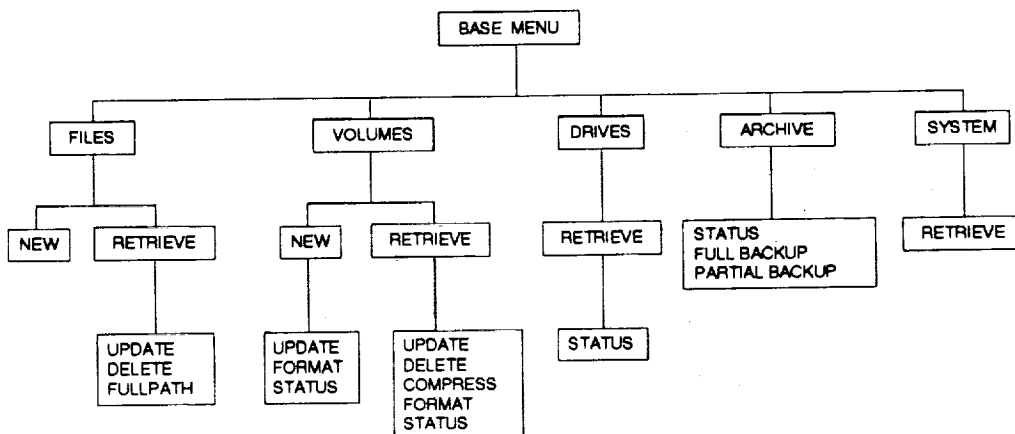
Metrum
Information Storage
{040161 sh}

RSS DATA MANAGEMENT

- Database Management System
 - 90 million files
 - 65,000 cartridges
 - Maintains system performance log
 - Comprehensive report generator
- Monitors system events
 - Errors are categorized by severity
 - Multiple notification targets, including e-mail
 - Incorporates recovery algorithms
- User data can be directed to specific cartridge(s)
- Admin functions available at console & via network

Metrum
Information Storage
[031801-ah]

RSS ADMINISTRATION MENU MAP



Metrum
Information Storage
[052191-ah]

RSS ADMINISTRATION BASE MENU

AMASS Administration

Version: AMASS/2.4

Files(1) Volumes(2) Drives(3) Archive(4) System(5)
Wipe(0) Help(PF2) End(PF3) Quit(PF4)>

Metrum
Information Storage
(031801-01)

RSS ADMINISTRATION FORMS INTERFACE

FILES / RETRIEVE MENU

AMASS Administration

Name:	1990_wp_files	Record:	32
User:	100	Parent's Record:	15
Group:	101	Volume:	
Mode:	rwx----	Volume Group:	3
Type:	Directory	Starting Block:	
Size:		Bad Blocks:	
Archived:	01-13-91 13:18:22	Nr of Errors:	
Accessed:	01-13-91 13:18:22		

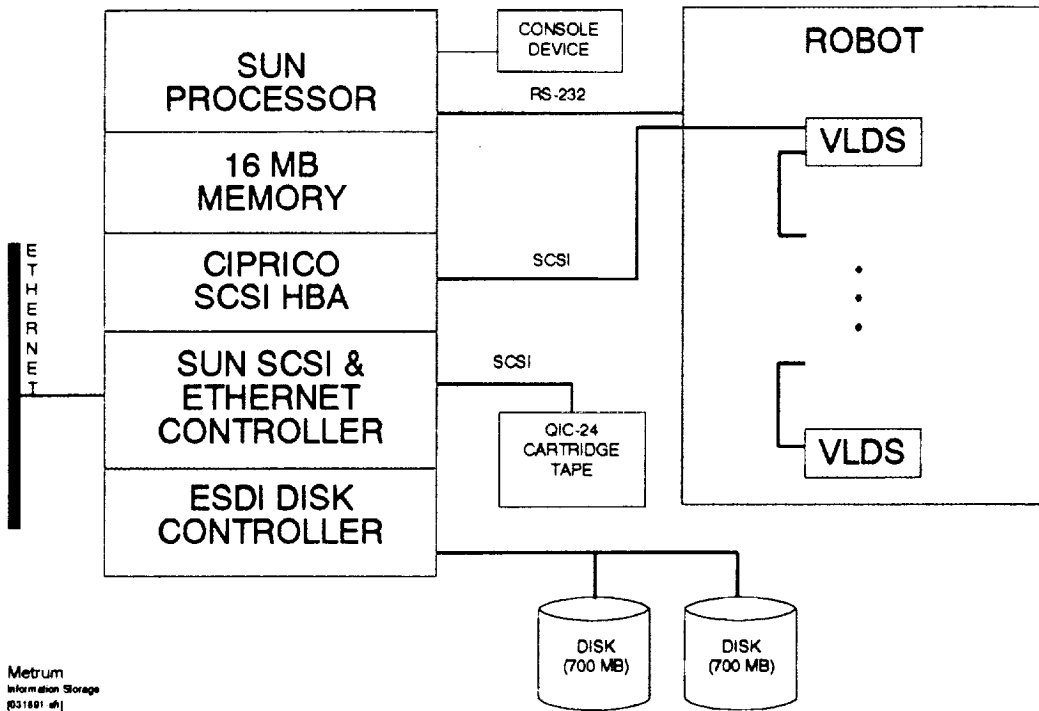
3 of 13

Files/Retrieve

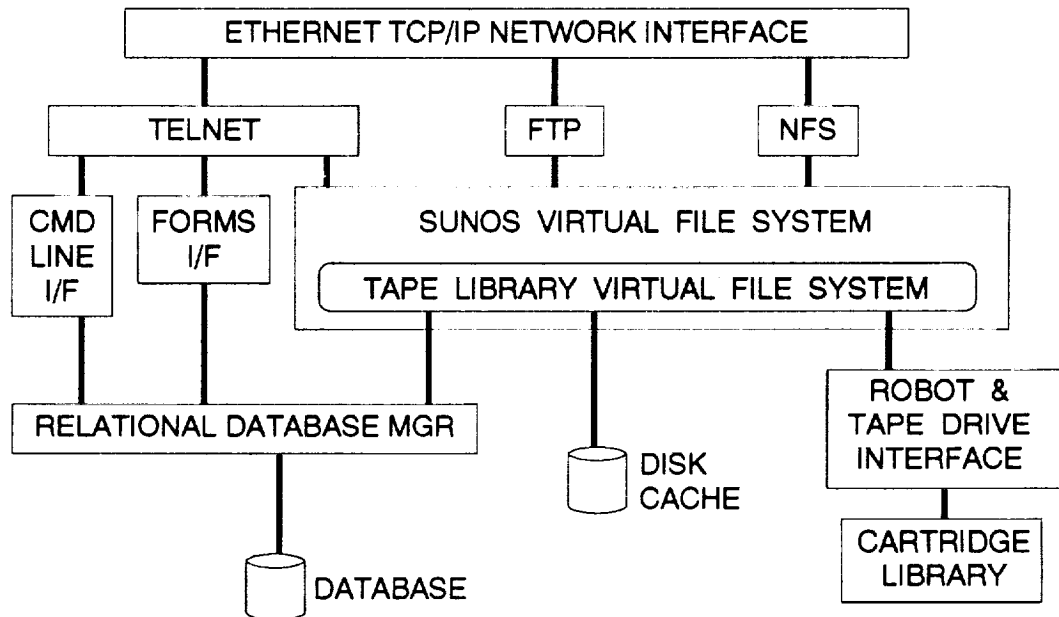
Update Delete Fullpath

Metrum
Information Storage
(031801-01)

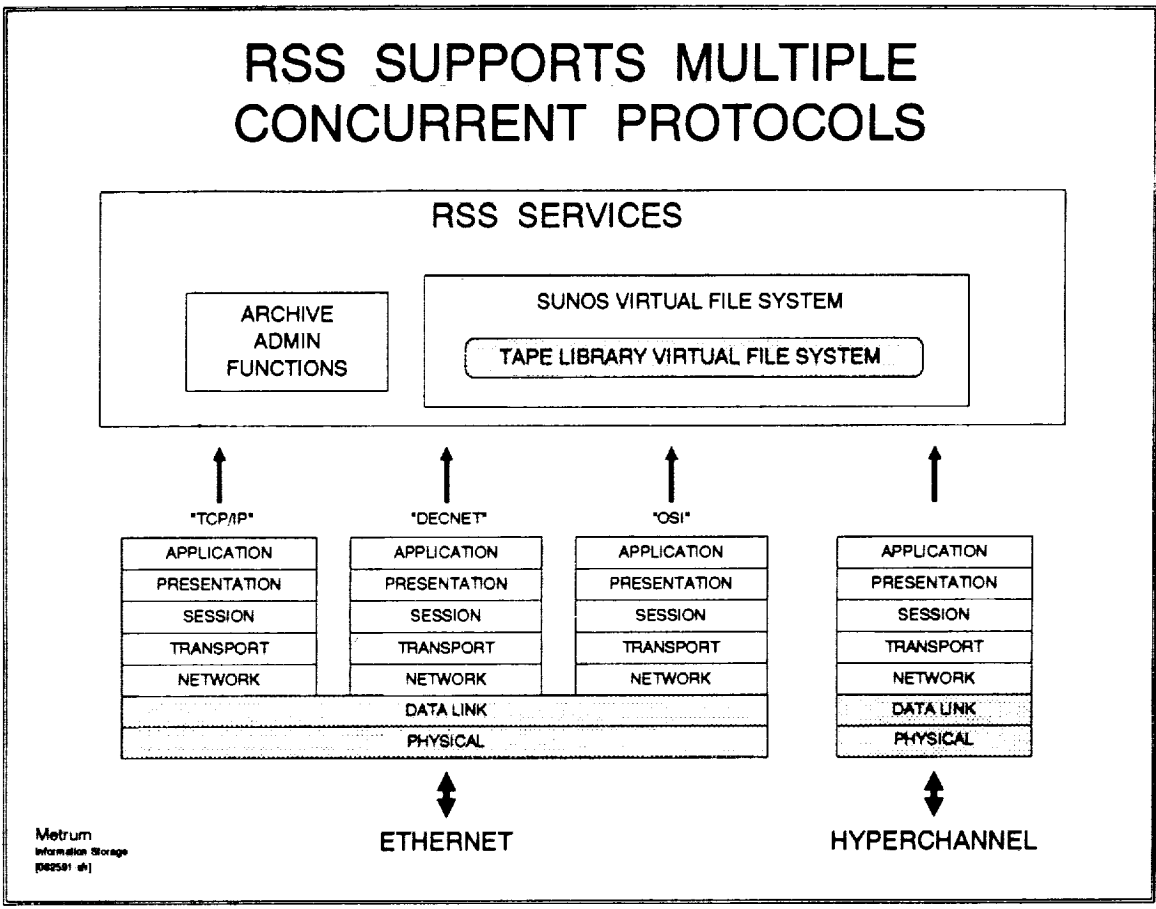
RSS HARDWARE ARCHITECTURE



RSS SOFTWARE ARCHITECTURE




RSS SUPPORTS MULTIPLE CONCURRENT PROTOCOLS



DATA SOURCES


- Data acquisition systems
- Data transcription systems
- Local computer file systems
- Remote computer file systems



TAPE DRIVE CONFIGURATIONS READ-AFTER-WRITE VLDS

- 1 MB/S fixed rate, read while write
- $1E-13$ BER with error management software
- 10.4 GB cartridge capacity (\$.002/MB media cost)
- 6.2 TB in RSS-600 (\$.085/MB system cost)


Metrum
Information Storage
[072291-eh]



TAPE DRIVE CONFIGURATIONS VLDS FAMILY

- 1 - 4 MB/S fixed rate
- $1E-8$ to $1E-10$ BER, config dependent
- Variety of interfaces
- Up to 10.4 GB per cartridge (\$.003/MB media cost)
- Up to 6.2 TB in RSS-600 (\$.085/MB system cost)


Metrum
Information Storage
[072291-eh]



TAPE DRIVE CONFIGURATIONS BUFFERED VLDS

- 0 - 4 MB/S variable rate
- 4 MB/S maximum sustained rate
- 20 MB/S burst (16 MB buffer)
- Variety of interfaces
- 1E-8 Bit Error Rate
- 10.4 GB cartridge capacity (\$0.003/MB media cost)
- 6.2 TB RSS-600 capacity (\$0.09/MB system cost)


Metrum
Information Storage
(772291-01)



TAPE DRIVE CONFIGURATIONS RSP-2150

- 0 - 2 MB/S variable rate
- 2 MB/S maximum sustained rate
- 4 MB/S burst rate (8 MB buffer)
- SCSI or PERTEC interface
- 1E-14 bit error rate (auto re-read/write)
- Up to 14.5 GB per cartridge (\$0.001/MB media cost)
- Up to 8.7 TB in RSS-600 (\$0.06/MB system cost)

Metrum
Information Storage
(772291-01)



OPTICAL TAPE DRIVE HD-1000

- 0 - 3 MB/S variable (3 MB/S sustained)
- SCSI interface
- 1E-12 bit error rate (auto re-read/write)
- 1 TB per reel (\$.005/MB media cost)
- 28 sec avg access time (60 sec end-to-end)
- \$250K drive cost
- Future adaptation to RSS-600 chassis

Metrum
Information Storage
[37291-01]

RSS ARCHIVE ARCHITECTURE SUMMARY

- UNIX-based operating system
- Standard network protocols
- User-directed archiving
- Supports automatic archiving
- Managed data
- Fast, easy access

Metrum
Information Storage
[33101-01]



ICI Optical Data Storage Tape

Robert A McLean

Joseph F Duffy

1. Introduction to flexible optical media

Optical data storage tape is now a commercial reality. The world's first successful development of a digital optical tape system is complete. This is based on the *Creo 1003* optical tape recorder with *ICI 1012* write-once optical tape media.

In order to understand the significance and potential of this step change in recording technology, it is useful to review the historical progress of optical storage. This has been slow to encroach on magnetic storage, and has not made any serious dent on the world's mountains of paper and microfilm. There are numerous reasons for this, the most of important of which are:

- The long time needed for applications developers, systems integrators and end users to take advantage of the potential storage capacity.
- Access time and data transfer rate have traditionally been too slow for high-performance applications.
- Optical disk media has been expensive compared with the competition, eg magnetic tape.

As one of the world's major international chemical companies, ICI's strategy in response to these concerns has been to concentrate its efforts on flexible optical media; in particular optical tape.



ICI Imagedata

2. Manufacturing achievements

Flexible optical media offers many benefits in terms of manufacture; for a given capital investment, continuous, web-coating techniques produce more square meters of media than batch coating. The coated layers consist of a backcoat on the non-active side; on the active side there is a subbing layer, then reflector, dye/polymer and transparent protective overcoat. All these layers have been tailored for ease of manufacture, and specific functional characteristics.

3. Media characteristics

The media has the ability to deliver high system performance and functionality over a very wide range of system capacities, depending on the drive design and media format. In addition to low cost per MegaByte, this can also be achieved with the archivability and indelibility that is vital for many storage applications. Consequently, the media permits the development of drives and systems that provide a unique set of features:

- **Low on-line cost/MB:** 10¢ to 40¢/MB depending on the format.
- **Low media cost/MB:** $\frac{1}{2}$ ¢ to 1¢/MB at first, falling with time.
- **High performance in access time:** the low track spacing and high longitudinal speed of fixed head optical tape drives allows search rates between 2 GB/s and 20 GB/s. This compares very favorably with helical scan magnetic drives.
- **High data rate:** a single laser channel can achieve 4 MB/s with 60 mW on the media surface, and this can readily be increased with the use of multiple lasers.
- **High volumetric efficiency:** flexible optical media offers efficiencies a factor of 10 higher than advanced helical magnetic recording.
- **Indelible media:** the original information cannot be erased and altered.



ICI Imagedata

- **Unlimited read cycles:** non-contact recording means that the data will survive in excess of 30,000,000 read cycles. The number of tape wind cycles is very high.
- **Long media lifetime:** the use of chemically stable materials has allowed us to predict a lifetime in excess of 15 years.

4. Media lifetime

Optical media based on organic, dye/polymer materials is well-known for its chemical stability. An unfortunate consequence of this is the difficulty in predicting a lifetime. Several accelerated aging techniques are in common usage: steady-state elevated temperature and humidity with Arrhenius extrapolation; elevated levels of temperature plus corrosive gases (the "Battelle" test); and chemical stability investigations.

We have used all of these extensively, both to attempt to predict an absolute lifetime limit, and in comparisons with magnetic tapes: iron oxide, cobalt modified iron oxide, chromium dioxide and metal particle formulations. We have also compared our media with rigid optical disks, including stability to uv light.

The Arrhenius test is performed at various constant conditions of elevated temperature and humidity (Ref. 1). The test assumes the following relationship, where failure is dominated by one chemical process with a given activation energy: $1/\text{time} = A\{e^{-E/RT}\}$. We defined failure as a change in reflectivity (Fig. 2) or a drop-off in the Carrier-to-Noise Ratio (CNR) (Fig. 3). In fact, within the measurement accuracy, failure was not detected for any of the samples, and this test is continuing. A better way of detecting failure is to measure the raw BER on a drive; we are now doing this on the Creo drive. However, by assuming that failure did occur, it is possible to estimate a lower limit on the lifetime (Fig. 1). The mildest conditions of 55 °C and 60 %RH do not correlate with a straight line through the three data points at more severe conditions. The slope of the graph gives an activation energy of 1.12 eV. A comparison with other forms of optical media is shown in Table 1.



ICI Imagedata

This implies that in comparison with the sample that failed after 121 days at 75 °C and 60 %RH, a sample kept at conditions of 20 °C and 60 %RH would survive for an extra 393 years!

We have used two standard uv stability tests. The British Standard *Blue Wool Test* was developed for dyes used in the clothing industry. It assesses light fastness on a scale of 1 (low) to 6 (high). ICI's data storage dye measures > 4 with the test still continuing. By comparison, the average textile dye measures < 4.

The British Standard *Accelerated Weathering Test* exposes a sample to 120 hours under UVA light at 55 °C and 60 %RH. This is equivalent to 72 hours of sunshine. Under this test we saw 1% drop in reflectivity. These results are similar to that for *Sony Century Media* in an identical test.

We also compared our media to metal particle and metal oxide tapes. In the case of metal particle, we compared changes in recording characteristics, and for metal oxide, we compared the chemical stability.

The metal powder tapes were stored under accelerated aging conditions, and the recording characteristics were assessed by looking for any changes in the magnetic properties. We bought two tapes; tape "A" is consumer R-DAT and tape "B" is 8 mm video, and both were different major Japanese brands.

The test method consisted of storing the tapes at a constant 60 °C and 80 %RH for 3 and 6 weeks, plus keeping a Control tape at room conditions. The magnetic properties were then measured in a high saturation field VSM. This work was done under contract for ICI by the Fulmer Research Company in the United Kingdom, and is summarized in Table 2. This shows drops in signal strength of up to 22% after 6 weeks.

We performed an analogous test on the ICI optical tape out to 9 weeks (Fig. 4). The signal strength was unchanged, and the CNR curves did not change within the



ICI Imagedata

measurement accuracy. In addition, we used a Time Interval Analyzer (TIA) to look for an increase in the intrinsic error rate by measuring a Geometric Error Rate (GER) on unwritten media, and looking for changes. The GER is directly proportional to the sum of the defective areas. Figure 5 displays this against the detector threshold level, normalized to unwritten media at 100%. Points below 100% are dark defects, and those above are light defects.

ICI media darkens upon writing and a typical detector threshold level is 60%, so the GER appears to have increased by less than a factor of 2.

Temperature cycling tests can be used to test not only the chemical stability of the media, but also the mechanical integrity of the structure. It is not possible to estimate a lifetime from cycling tests. We subjected the optical tape media to 20 temperature cycles defined by Fig. 6. We again looked for changes in the recording characteristics. There was a 1 or 2 dB drop in the CNR (Fig. 7), which is just greater than the measurement accuracy. The time interval results presented in Fig. 8 showed a corresponding increase in the signal jitter (standard deviation) by 3 ns. The "(+/-)" refers to triggering on the rising and falling slopes of the signal, and "(-/+)" is the other half cycle.

A range of oxide magnetic tapes were stored under accelerated aging conditions, and any degradation was assessed by looking for any decomposition products. This work was also done under contract for ICI by the Fulmer Research Company in the United Kingdom.

We tested three tapes; tape "C" is iron oxide instrumentation tape, tape "D" is cobalt modified iron oxide VHS consumer video tape, and tape "E" is chromium dioxide *IBM 3480* type data cartridge tape.

The test method was to store the tapes at a constant 60 °C and 80 %RH for 3 and 6 weeks, while keeping a Control tape at room conditions. Then standard solvent extraction techniques were used to look for any decomposition products. One meter



ICI Imagedata

length samples were immersed for two hours in a Soxhlet extraction by Delifrene, then the soluble extract was weighed. This was followed similarly by a further two hours in acetone, then weighing the extract.

The results are presented in Table 3. The initial extract is a baseline, and the percentage increase above this is a measure of chemical decomposition in the binders. These results can be put into perspective through a paper by Bertram and Cuddihy (ref. 2). This states that after aging, an increase in extract of 1.4 % by weight corresponds with a degradation in tape performance. When our results are compared with Bertram and Cuddihy's, there was very good agreement for the Tape "C".

We performed a similar chemical stability test on the optical tape. Again, the test method was to store constant 60 °C and 80 %RH, plus 80 °C and 80 %RH for 1, 2 and 3 weeks, plus a Control tape at room conditions. We were unable to detect any decomposition products at a constant 60 °C and 80 %RH, and so the test was repeated at the more severe conditions of 80 °C and 80 %RH for three weeks.

We devised an extraction technique suitable to the solubility of the dye/polymer material. This involved immersing 500 cm² samples for 72 hours in a Soxhlet extraction by ethanol. Samples and extract were then weighed and analyzed by sensitive FTIR techniques. Using FTIR, we could detect binder degradation only at the harsher conditions of 80 °C and 80 %RH, for a minor component of the protective overcoat, not the recording layer. The quantities were too small to be weighed.

Another test we used is the Battelle Class II accelerated aging test, performed under contract at the Battelle Institute in Ohio. This test has been correlated to aging of materials in cities and other locations where combustion byproducts form a mix of corrosive gasses. A fully assembled open reel *ICI 1012* Optical Tape, without the protective sealed storage and shipping box, was kept at a constant 23 °C and 70 %RH in a flowing mixed gas environment consisting of 10 ppb H₂S, 10 ppb Cl₂ and 200 ppb NO₂ for 30 days. This environment has been correlated to 15 years lifetime. We detected no damage in terms of corrosion, reflectivity or modulus.



ICI Imagedata

5. Conclusion

Optical tape systems can offer a unique set of attributes to potential end-users. The use of chemically stable recording materials, indelible write-once technology, plus the advantages of non-contact reading and writing, yield a very robust and archival medium.

Our lifetime studies have concentrated on measurement techniques and comparisons with magnetic technologies that are independent of the drive and read/write channel. The comparisons with magnetic tape technologies are favorable:

- No deterioration in optical recording characteristics after 9 weeks at 60 °C and 80 %RH.
- Metal particle magnetic formats show up to 20 % deterioration in magnetic capability after only 6 weeks.
- No binder hydrolysis in the recording layer after severe environmental exposure and extraction process.
- Minor hydrolysis present in overcoat after 3 weeks at the extreme conditions of 80 °C and 80 %RH.
- Clear evidence of deterioration in magnetic oxide tapes under an industry standard test.

We plan to extend these studies to raise the lifetime prediction for *ICI 1012* optical tape from 15 to 30 years, and to characterize any changes in the raw BER through the read/write channel on the *Creo 1003* optical tape recorder.

References

1. A B Marchant, *Optical recording: a technical overview*, Addison-Wesley Publishing Company, pp 380-382, 1990.
2. H N Bertram and E F Cuddihy, "Kinetics of the Humid Aging of Magnetic Recording Tape", *IEEE Trans Magn*, vol MAG-18, no 5, pp 993-999, Sept 1982.



Table 1

Company	Media Type	Test Method	Activation Energy in eV
Sony	Metal Alloy	BER	1.5
ICI	Dye/polymer	%R and CNR	1.12+
OITDA*	Multi WORM	%R	1.0
Hitachi	12" WORM	BER	1.0
NEC	Magneto-optic	DER	0.97

*OITDA = Japanese Standard Committee for the Optical Data Disk

Table 2

Tape	Time	M _g % fall	M _r % fall	Sq % fall
"A"	Control	0	0	0
	3 weeks	10.64	11.30	0.74
	6 weeks	18.84	20.34	1.84
"B"	Control	0	0	0
	3 weeks	4.20	4.38	0.19
	6 weeks	22.35	22.31	-0.015

Table 3

Tape	Time	Total Extract in %	Increase in %
"C"	Control	1.20	0
	3 weeks	2.35	1.15
	6 weeks	3.49	2.29
"D"	Control	1.24	0
	3 weeks	1.45	0.21
	6 weeks	1.86	0.62
"E"	Control	1.53	0
	3 weeks	2.04	0.51
	6 weeks	2.55	1.02



Fig. 1





Fig. 2



Fig. 3





Fig. 4



Fig. 5





Fig. 6



Fig. 7

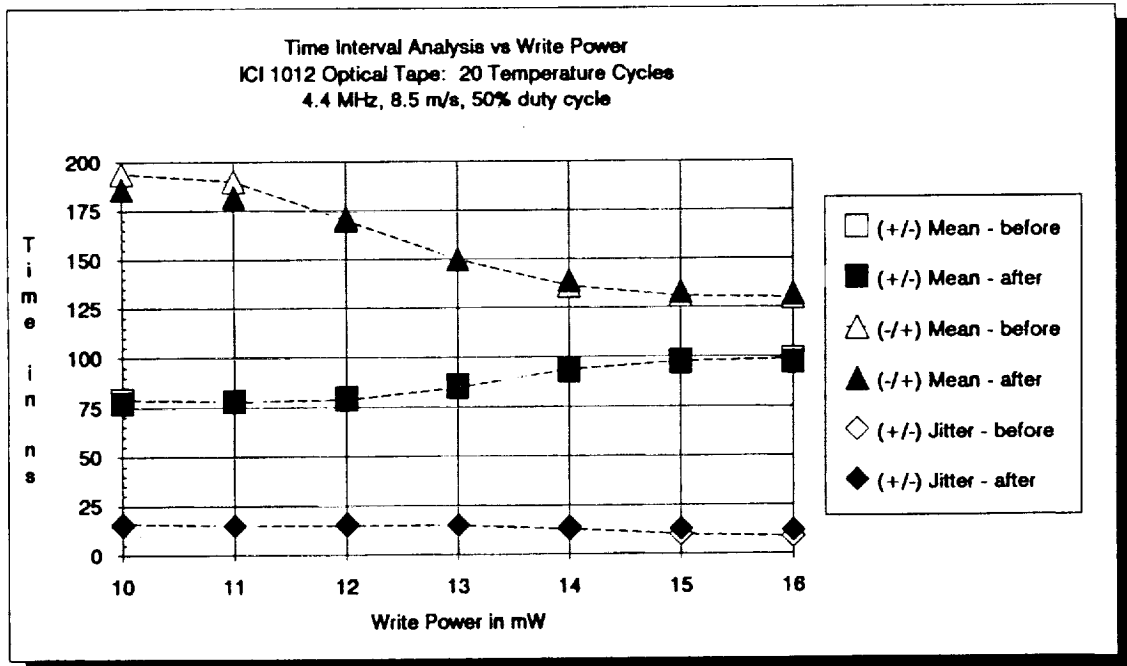
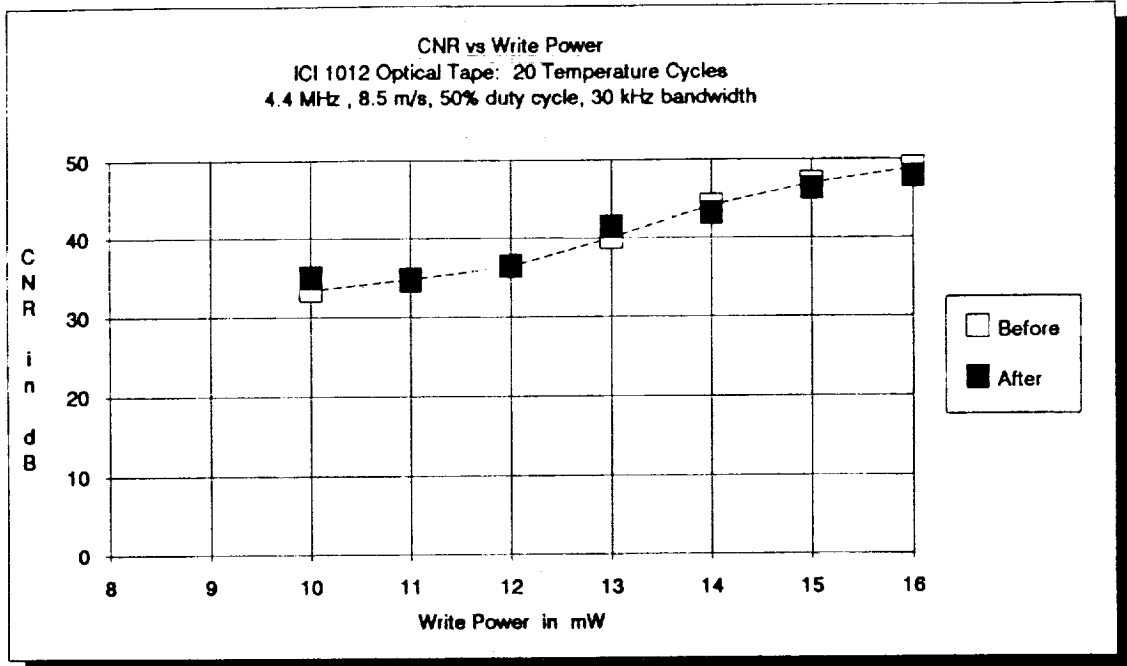


Fig. 8

Presentation for NSSDC Conference

**Mass Storage Systems and Technologies
for Space and Earth Science Applications**

**ATL Products Division's Entries Into the
Computer Mass Storage Marketplace**

July 24, 1991

Fred Zeller
Odetics, Inc.
ATL Products Division
240 E. Palmdale Road
Anahem, California 92805
714-774-6900

DOCUMENT NO. 9011002

ATL Products Division

Odetics

Odetics Background


- **High Tech Company Founded In 1969, Publicly Traded**
 - Serving Well-Defined Niche Markets
 - Through Variety of Product Groups
- **Roots are In Space Borne Recorders**
 - Own 80% of Marketplace
- **Evolved Into Robotics In Mid-70s**
 - AIM, Broadcast, DMS, ATL Products Divisions
 - 35% of Revenues and Growing
 - Technology and People Move From One Division to Another

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Spaceborne Tape Recorder



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Product Evolution


- **Robotics Genesis - AIM**
- **Company's High Technology Group**
 - 1979 Committed to a Six Legged Robotic System
 - 18 Months Later Demonstrated ODEX I
 - Symbol of the Corporate Commitment to Robotics
 - Demonstrates High Strength to Weight Ratios
 - All Electric, Compact, Extremely High Performance
 - Six Units Built - Three Generations of Technology
 - Predominantly for Nuclear Plant Maintenance
- **Evolution to Other Robotic Subsystems**
 - Arms, Hands, and Effectors
-

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Three Generations of ODEX



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Product Evolution Technologies and Markets Served

- Innovators In "Small Package" Handling
- Do Not Serve General Purpose Robotics Handling Market
- Design Intent of Our Products
 - Move "Small Light Weight Objects" Very Quickly
 - Accent On Longevity of "Object" Being Moved
 - High Degree of Reliability
- Necessitates
 - Expertise in Low Mass, Light Weight, High Speed Systems
 - Requires Unique Robotic Handlers, Arms, End Effectors
 - Products Designed for Niche Markets
 - Aperture Card Storage Module Systems
 - Tape Cassettes and Cartridges
 - Optical Disks

DOCUMENT NO. 9011002-2
JULY, 1991

ATL Products Division

Odetics

Product Evolution Infodetics' Aperture Storage Module Library

- **First Linear Servo Based Expandable System**
- **Modules: 10 Ft. Long By 3 Ft. Deep By 7 Ft. High**
 - 2000 Cartridges Per Two Rows or Module
 - 100 Aperture Cards Per Cartridge
 - Robotic Handler in Aisle Between Rows Within Module
- **Large System With Multiple Modules and Pass-Through**
- **5 Seconds Average "Pick and Place"**
 - Access Cartridge and Load Into Aperture Card Reader
- **Document Management System "Storage Server"**
 - Cache Microfilm Images to Disk
 - Transmitted to Work Stations for Viewing

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Product Evolution Broadcast Division's TCS2000 Video Cart


- **First "Tower" Based Expandable System Introduced '86**
 - Designed as a TV Station or Network Automation System
- **Built as Part of a Joint Venture With RCA In 18 Months**
 - RCA Dropped Out, Odetics Entered End User Market
- **System Consists of:**
 - Robotics and Up to 6 Tape Recorders Per Tower
 - 225 to 300 Tapes Per Tower Depending On Formats
 - Switchers, Sequencers, Monitor and PC Based Work Station
 - Hierarchical Software
 - Real-Time Controller/Operating System, Relational DB, Playlist
- **Supports VHS, Beta, D-2 Formats**
- **Robust, Redundant and Extremely Reliable**

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

TCS2000 Cart Machine With Library Expansion Module



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Broadcast Division's Newest Product TCS90 Videocart System


- **Bookshelf Design**
 - X-Y High Speed Linear Servo System
 - One Armed Gripper With Holding Tray
- **Accommodates Combination of Cassette Sizes and Format**
 - Beta SP, VHS, or D-2
 - Small and Medium Sizes
- **Tape Recorders are Standard: Non-Modified**
- **Autoloader Accommodates Up to 8 Cassettes**
- **Fixed Size**
 - No Expansion Capabilities
- **Software From TCS2000 Migrates Directly to This Product**
 - 50 Man Years of Development

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

TCS90 Cart Machine



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

ATL Products Division Marketing Strategy

- **Serve the Evolving Computer Based Mass Storage Market**
- **Develop Library Automation Subsystems**
 - Robotics, Control, Storage and Computer Interfaces
 - Support a Broad Range of Tape Sizes and Formats
 - Interface to a Variety of Tape Drives in Each Size
 - Provide Low Level Library Control and Management Software
- **Sell Through Distribution Channels That:**
 - Integrate Tape Drives
 - Add High Level Mass Storage Management Software
 - Service a Broad Range of User Markets
 - Provide "Private Labelled" and "General Purpose" Products
- **Pursue Major Market Shares**
 - High Density/High Capacity Storage Market
 - General Commercial Market By Supplying a Range of Solutions

DOCUMENT NO. 9011002-2
JULY, 1991

ATL Products Division

Odetics


High Density Systems Business Product Lines

- **Developing Two Basic 19mm ATL Storage Subsystems**
 - To Serve High Capacity Markets: Terabytes and Petabytes
 - Support Small and Medium Size D-2 Cassettes
- **Expandable "Tower" Based System**
 - Broadcast System as Platform
 - Auxiliary Towers for Expansion
 - Delivered March, 1991
- **New Linear Aisle Based Expandable System**
 - 30 Months in Development: Delivery August, 1991
 - Most Advanced Robotic System On the Market

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division
Odetics


19mm Aisle Based ATL



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division
Odetics

19mm ATL Cut Away Side View of Robots



DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

High Density Systems Business Product Lines

- **10 Minute Video On Two Technologies**
 - Copies of Video Available Upon Request
 - Text of Narration Follows Presentation

DOCUMENT NO. 9011002.2
JULY, 1991

ATL Products Division

Odetics

E-Systems Business Relationship

- **Exclusive Supplier of 19mm ATL Subsystems**
 - Computer Mass Storage Marketplace Only
- **Can Market Other Odetics ATL Subsystems**
- **E-Systems is the Integrator**
 - Providing Systems Expertise - ATLs, Tape Drives, Computer Integration
 - Library Management Software for the ATLs
 - Supplying Storage Server Software

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

3480 "Medium Size" Library First Commercial Product Offering

- **"300" Cartridge Baseline System**
 - Expandable in Increments of Approximately 300 Cartridges
 - "1500" Cartridge Maximum
- **Up to 2 Tape Drives Available in Baseline**
 - Up to 4 Additional Drives as System Expands
 - Supports All Low Cost 3480 Tape Drives
- **Small Footprint**
 - Fits Standard 28 Inches Wide By 45 Inches Deep By 78 Inches High Cabinet
 - Very High Density Storage
- **Cartridge Autoloader and Bulk Loading**
- **RS-232C or SCSI-II Interface**
- **Serve the Distributed Computing and File Server Markets**

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Storage and Library Management

- **ATLs are Driven With "Low Level" Commands**
 - Pick From Bin and Move to Tape Drive
 - Status Provided Back Through Sensors
 - Electrical Interface: RS-232C, Ethernet, SCSI-II
 - New Software Interface: SCSI-II, Chapter 16 Jukebox Commands
- **Library Management: Physical Volume Repository**
 - Input PVS and Provide Level of Intelligence
 - Management Resource and Allocation of Bins and Drives
 - Automatic Error Recovery
- **Servers and Applications Provide Next Level**
 - Storage Servers and Bit File/Client Servers
 - Backup

DOCUMENT NO. 9011002
JULY, 1991

ATL Products Division

Odetics

Summary What Odetics Can Offer the Market

- **We Are and Will Be a Major Supplier of Robotic Libraries**
 - Advancing Technologies
- **By Year End, From Broadcast and ATL Products Divisions**
 - Four Different ATL Technologies and Five Products
 - Cross "Breeding" of Technologies Across Divisions
- **In the Future**
 - Broader Reach of Products and Markets Using Robotics
 - Further Transfer of Technologies at Component Levels

DOCUMENT NO. 9011002.2
JULY, 1991

ATL Products Division

Odetics

ATL PRODUCTS VIDEO

19mm High Density Tower & Aisle Technologies

Narration

- The information explosion is here. The amount of data being collected is growing at an ever increasing pace. As a result, the task of managing that data presents new challenges for data processing professionals.
- How do you effectively manage your data while providing rapid availability to your users, reduce costs for labor and storage, eliminate misplaced tapes and ensure security?
- Odetics, a leader in robotic storage and retrieval systems has a solution for you. A variety of automated tape library, or ATL systems are being developed for use in a broad range of applications.
- A culmination of twenty years experience, these ATL products will utilize robotic technologies found in hundreds of installed, field-proven aperture card information and broadcast video cart systems.
- This presentation introduces two of our mass storage subsystems -- We'll start with the aisle-based automated tape library from Odetics.
- This expandable automated tape cassette library represents the highest density mass storage subsystem available. It's capable of storing terabytes or even petabytes of data. This system was specifically designed to meet the needs of the federal government which is the largest user of data in the United States.
- Its unique architecture moves tapes through the library rapidly and efficiently. And it's completely modular. As a result, it can easily be configured to meet your specific floor plan requirements while maximizing storage density and the system can be easily expanded to keep pace with your growing requirements.
- This animation utilizes a combination of solid and wire-frame images. The outer skins of the system have been removed to provide the best perspective of the ATL in operation.
- For this presentation, a system configuration of 2 aisles and 3 rows is shown.
- Small 25 gigabyte and medium sized 75 gigabyte 19mm D-2 cassettes are stored vertically in the library. Each medium D-2 cassette has the same storage capacity as 520 reels of 9 track, 6250 bpi magnetic tape.
- The cassettes are kept in rows of modular bins called cassette storage modules. These modules can be joined to form rows up to 80 feet in length.
- Robots operate in the aisles between the rows. They move cassettes between bins and tape drives located at either end of the library.
- Maximum speeds are nine feet per second in the horizontal axis and 2 feet per second in the vertical.
- The robots are gantry-based. They travel on tracks mounted along the top of each row of storage bins.
- As a result, there's sufficient clearance between the robot and the floor to protect any cassettes which might inadvertently fall from a bin.

DOCUMENT NO. 9011002-2

JULY, 1991

ATL Products Division

Odetics

A network of sensors can halt the robot within a half second in response to any object that would interfere with its movement.

Each robot is controlled by electronics which interfaces with the system's host computer through an Ethernet Local Area Network. Commands to move the robot are sent from the host to the control electronics which are not shown in this presentation.

Servo drive electronics power the robot's components.

Each robot has two tape accessors which consist of a sliding arm and a uniquely designed gripper which can grip the different sized cassettes.

The arms and grippers enable the robot to move cassettes to and from cassette storage modules on either side of the aisle.

The two accessors are linked to one another by a belt which controls their vertical movement.

A barcode scanner is mounted on the accessor arm. This scanner reads the barcode on the cassette to verify its identity prior to removing it from the bin.

Up to four tape drives are housed in a recorder mounting module stacked vertically in compartments. The end unit houses a corresponding number of cassette loader modules which move tapes between the robots and the tape drive.

The robot places a cassette in the bin on the platform of the cassette loader module.

Since tape drives can be positioned on both ends of the system, facing inward, the platforms on one end, in this case, the left rotate 180 degrees to position cassettes properly for placement in the drive.

The bin is mounted on a vertical slide that lifts the cassette to where it's engaged by a gripper.

The gripper swings the cassette 90 degrees into the horizontal position and inserts the cassette into the tape drive.

For library management, a bulk loader for loading and unloading up to 10 cassettes from the library at a time is also provided. Although they're not shown in this presentation, the bulk loader and the control electronics each replace a tape drive and cassette loader module in an end unit.

Let's watch the subsystem operate at actual speed as we summarize its key features.

As you have seen, the ATL is extremely versatile and redundant. Cassettes can be accessed by either of two robots.

Performance and efficiency are exceptional. Maximum time to do a cassette pick and place from one end of an 80 foot row to the other, including tape drive insertion, takes less than 23 seconds.

The ATL is completely modular. Components can be arranged in combinations that fit any floor plan.

It's also expandable. Rows and aisles can be added to the library without interrupting operation. A user can start with a small system and as storage requirements grow, it can be easily expanded over time.

Finally, the system architecture is reliable. Redundancy built into the system provides maximum uptime. And individual components are designed to last.

DOCUMENT NO. 9011002-2

JULY, 1991

ATL Products Division

Odetics

- Another application of our mass storage technology, the TCS2000 Odetics Broadcast Videocart Machine, has established a new industry standard for efficiency and economy in playing both television station and network programming and commercials to air.
- In the system's tower, robotics control the programs and commercials which are recorded on 1/2" analog VHS, Beta or 19mm D-2 videocassettes.
- These cassettes are stored in up to 280 bins within the storage tower.
- A 4 armed robot moves the cassettes from the bins to tape drives which are stacked in the rear.
- There is a variety of other hardware including video switchers and monitors as well as a menu driven interactive touch panel which is used to make changes to the play list rapidly and without mistakes.
- A powerful microcomputer controls the system.
- Commands are sent from the playlist to a sequencer which manages the robotics and controls the tape drive electronics so events are played to air in the correct order.
- Software maintains a database of up to 65,000 commercials and program segments on cassettes stored inside and outside the tower.
- With over 50 man years of development, the software also has a myriad of user oriented features for on-line editing of playlists that make the TCS2000 not only easy to use but allows error-free operation.
- Over 130 TCS2000 systems are installed worldwide, and it has become the broadcast industry standard for videocarts.
- It's proven to be a highly effective solution that reduces costs and dramatically improves accuracy and efficiency.
- And now, the tower technology has been adapted to serve as an automated tape library for mass storage of computer data. And users can expect the same outstanding results.
- With high density digital D-2 cassettes, the tower has the capacity to archive 5.7 terabytes of data.
- This makes the tower ATL a more cost-effective storage alternative than magnetic storage or even optical disks for "near" line storage.
- It can store up to 227 small D-2 cassettes with each cassette storing 25 gigabytes of data. That's equal to the capacity of 175 9 track 6250 bits per inch magnetic tapes.
- Like the broadcast system, an automatic cassette loader, mounted in the door, allows up to ten cassettes to be loaded or removed from the library quickly and easily.
- As cassettes enter or exit through the autoloader, a barcode reader scans their barcode labels. The information is automatically entered into a database residing in host computer software which tracks cassettes and generates reports on usage.
- The drive electronics, mounted in the rear of the tower, accept pick and place commands which direct the robot to the proper bin and tape drive locations as it moves cassettes through the system.
- The rapid vertical movement of the 4 armed robot is controlled by a digital stepping motor mounted at its base, geared to a nylon belt.
- Each of the four arms includes a custom-designed, highly reliable end effector which grips the cassette with a controlled force to protect it from damage.
- The cassette loader module is the transport mechanism that passes cassettes between the robot and its corresponding tape drive.
- Once the cassette reaches the loader module, the robot places the cassette into the module's input port. The cassette then moves along a conveyor which is aligned to the throat of the drive.
- These tape drives are mounted in a recorder mounting module at the rear of the tower.
- When unloading, cassettes move through the cassette loader module to the robot arm in a similar manner.
- Simple cassette pick and place commands are transmitted through an RS232C hardware interface.
- The maximum time for moving a cassette from its storage bin to a tape drive is less than 7 seconds.
- As with the broadcast system, the robotics are extremely reliable. With over 200 man years of maintenance logged on TCS2000 systems in the field, it has a proven record of 99.999 percent availability. As a result, the system has dramatically improved the efficiency and accuracy of broadcast play to air operations for our customers. That means significantly lower costs and greater profitability. And data storage users can expect similar benefits.
- The modular tower design facilitates ease of maintenance. The entire robot can be quickly removed and replaced in a maximum of 30 minutes.
- And like the TCS2000, the system is easily expanded.
- Up to six towers can be linked together via a specially designed mobile storage bin which transfers cassettes between units within six seconds. Now, user's have an extremely cost-effective means of increasing capacity as storage needs expand.
- As companies grow, the problem of finding an efficient, reliable, cost-effective means of archiving large quantities of digital data becomes more and more challenging.
- The mass storage systems manufactured by Odetics meet these challenges by providing longer hours of reliable operation, with significantly improved efficiency, improved security and lower costs.
- Their proven technology and intelligent, modular design make them easily adapted to meet your archiving needs.
- As with all of our automated storage systems, Odetics developed these two powerful technologies to solve your company's data explosion by providing the easiest, most productive and cost-effective means of managing and utilizing large quantities of data.

**PANEL DISCUSSION:
STABILITY OF HIGH DENSITY MAGNETIC TAPE**

Moderator: Professor Mark Kryder

Participants: Bellcore, Ampex, Sony, Advanced Development Corp., NML-3M

MR. SAVAGE: We have come to the panel section of the afternoon session. Our panel moderator this afternoon is Professor Mark Kryder from Carnegie Mellon University in Pittsburgh, Pennsylvania. Mark has a respectable biography here, which I think I will read because they are still making a few preparations there.

Mark received a B.S. Degree in Electrical Engineering from Stanford in 1965, an M.S. in Electrical Engineering from Cal Tech in 1966, and a Ph.D. in Electrical Engineering and Physics from Cal Tech in 1970.

From 1969 to 1971, he was a research fellow at Cal Tech; and from 1971 to 1973, he was a visiting scientist at the University of Regensburg, West Germany. From 1973 to 1978, he was a research staff member and manager of exploratory bubble device research at the IBM T.J. Watson Research Center. In 1978, he joined Carnegie-Mellon University, where he is now Professor of Electrical and Computer Engineering and Director of the Engineering Research Center in Data Storage Systems, which he founded.

Professor Kryder has over 150 papers and 10 patents in the field of magnetic materials and devices. Currently, he is actively involved in research on magnetic and magneto-optical recording technology. Mark, it's all yours.

DR. KRYDER: Thank you very much. I hope we are going to have some fun this afternoon. The success of this panel in this afternoon's activity really depends on the audience. You know, the speakers are going to get up here and make a few statements; but it really depends upon you, the audience, to challenge them and to provoke some controversy and get some good discussions going.

So, I hope you are all primed and thinking about getting some of those controversial questions out there.

What we are going to try to do is to have a panel discussion on the archival characteristics of various tape media. And we have representatives here from a large number of organizations; and we will actually have an opportunity to hear from people talking about chromium dioxide, cobalt gamma type materials, as well as metal particle and metal evaporated type materials, as well as barium ferrite.

And we will hope to see some comparisons of those; I hope to maybe find out what some of the problems are. We may not come to any solutions here, but the hope is that we can identify some of the problems; and maybe we can conclude that some of the problems that people thought were there really aren't problems.

Anyway, we would like to get the issues out on the table.

What we are going to do is this: We have a number of people to bring up here, and I have asked each group to go no longer than 15 minutes--absolute maximum--in their presentation. There are really five groups; so, that means we will have an hour and 15 minutes of presentations maximum.

And then, following that, we will have an open panel discussion.

Now, in order to make room for the viewgraphs and so forth, although later on we will ask all the panel members to come up when we have the discussion session, I have decided to have them here in the first row at the beginning; and then, we will let each one give the individual presentation and then take their seat. And then, when it is time for the discussion, everybody will come forward.

What I have tried to do is organize the session from the historical perspective forward in terms of when various media were introduced.

So, the first person this afternoon whom I am going to have speak is Dr. Barbara Reagor from Bellcore. She is Division Manager of Chemistry and Materials Science Research. She has 21 years experience in contamination and materials reliability in telecommunications and data center environments.

From 1970 to 1983, she was at Bell Laboratories. From 1984 to date, she has been with Bellcore. She has a B.S. in Chemistry from Monmouth College in New Jersey, an M.S. in Organic Catalysis and a Ph.D. in Chemistry, doing work on pico-second laser spectroscopy, from Seton Hall University in New Jersey.

Barbara will talk about the environmental impacts on mass storage systems, with particular emphasis on chromium dioxide in the 3480 cartridge. Barbara?

ENVIRONMENTAL IMPACTS ON MASS STORAGE SYSTEMS

Dr. Barbara Reagor
Bellcore

DR. REAGOR: To give you an overview, back in 1986, what brought my Bellcore group into involvement with this was the appearance of an article of this nature saying that chromium dioxide tapes would pose a toxic hazard for employees as well as environmental disposal of the units.

(Showing of viewgraphs)

DR. REAGOR: The companies that own Bellcore are all of the regional Bell operating companies; and our use of data is quite extensive in data storage.

We basically must store for greater than ten years all of the telephone call transactions that go through this country, all of the billing information, all of the operational parameters of the switching equipment. And we were very concerned about the reliability of tapes, in particular if there are questions being raised about their environmental robustness or in terms of their toxic characteristics.

The types of long-term stability we are looking for are the impacts from indoor air pollutants, heat and humidity, and also again this question of toxic removal.

We have heard a lot today about different tapes. (Change of viewgraph)

DR. REAGOR: And I threw one of these in very quickly just to give an overview of what the magnetic metal oxide tapes were like, where we have your metal oxide particles--iron oxide or chromium dioxide--embedded in a polymer-based binder. That binder can be an elastomer of polyurethane; it may have lubricants, dispersants or curing agents, as well as other additives in it to give it its various properties.

And then, we have our polymer base, which is usually PET, polyethylene terephthalate; PVC, polyester, or cellulose acetate material.

MAGNETIC TAPES -

● Information storage (10+ years)
telephone calls, billing records, etc...

● Long term stability
- pollutants, heat, humidity,...

● Disposal
- no toxic components

In doing our study, we approached this from the sense of trying to learn as much as we could about the new chromium tapes. In 1987, we purchased seven different vendor products, all chromium tapes, all 3480 IBM compatible materials.

And these are the tests we subjected them to: differential scanning calorimetry to see how it responds to heat cycling and heat tests; thermal/mechanical analysis, where we look at the stretching parameters as we heat and move this tape in use; thermogravimetric analysis.

And we actually developed a magneto- thermogravimetric analysis technique, where we could look at the effect of heating in the presence of a 150 Gauss magnetic field. We did the magneto-hysteresis to see the impact of aging and chemical stressing on these tapes; and other typical chemical analyses, with FTIR analysis, extractions, X-ray spectroscopy, as well as macroscopic investigations.

(Change of viewgraph)

DR. REAGOR: In terms of some of the thermomechanical studies, what we would look at is a glassy type material at room temperature. You could slowly extend the materials, strain them as you start to heat. You would then see stress relaxation of the coatings on the physical binders.

And finally, if we got to about 200^o C, we would start to see melting of the actual binder and substrate materials. And we looked at the effects this had on the magnetism.

(Change of viewgraph)

DR. REAGOR: If we ran a TGA analysis, we would start to see some very interesting things. With the TGA, we can find the Curie temperature, where we lose all magnetism. That is where the chromium tapes do have a disadvantage over iron.

In a chromium tape, as you scan the temperature range, when you reach about 120^oC, you actually lose magnetism--at the Curie point, about 125^oC for these tapes. We could actually monitor this with and without magnetic fields.

We would watch the deterioration of the actual substrate; comparing that to the case of an iron oxide tape, where the Curie temperature is 600^oC, you have no loss of magnetism with that type of system.

(Change of viewgraph)

DR. REAGOR: To correlate it a little bit better. We could even do some interesting things where we could identify the actual degradation of the materials using these types of techniques.

We can see the effect of alignment of the chromium particles at the Curie temperature; we actually start to see a disalignment of the particles, a loss of magnetism, as they become mobile within the matrix of the binder.

As you reach the Curie point, we can look at the actual degradation. And a lot of work was put into using this material to actually analyze the tapes.

One of the things that we were able to learn from this--and it can be shown here --

(Change of viewgraph)

DR. REAGOR: Is we would look on when the alignment onset changed--at what temperature. At what temperature do we actually see the particles on the tape start to move and to lose alignment from a magnetic field? And in all of the tapes, the binder onset for alignment were showing up in about the 83^o C range. We were able to obtain, like I said, through this

method the Curie temperatures; and it shows you right here that T- 5 was actually an iron oxide tape and not a chromium dioxide tape.

From these types of information, from the weight gain and the loss of weight, depending on whether we have magnetic fields in place, we could determine the actual percent of substrate, the percent of metal, in the systems. We could take these all the way to total combustion; and in doing that, what we actually came up with was almost all of the tapes really were between the 30 and 35 mm thick PET films; and they had between 2 and 4 mm of coating in terms of the metal oxide. Yesterday, you heard a lot about the coercive force, about applied fields.

(Change of viewgraph)

DR. REAGOR: I put this up just as a way of explaining some of the data I'm going to talk to you about. When we do magnetohysteresis measurements, what we gain out of this hysteresis loop is that we can then get our remanent magnetization measurement. We can see our saturation magnetization; we can develop and measure the tape, and the surface roughness drastically changes the frictional drive.

PARTICIPANT: How long was the aging?

DR. REAGOR: 15 to 30 days. 85 degrees. This is not unreasonable for being in the back seat of your car on a humid day, and transportation was an issue. It is not unreasonable if you are in an environment that is in a disaster situation; and we have to maintain the possibility of that.

We also did, just looking at a variety of tapes, and we did studies of this nature, where we kept the temperature low and high.

(Change of viewgraph)

DR. REAGOR: We changed the humidity; we actually measured this magnetic weight. It gives you a feeling for how much magnetism is still retained in the tape and looked at the percent losses that we would see.

And again, situations that developed--massive changes--are shown here. At low temperature, basically just a little bit above room temperature, the humidity alone and the low temperature alone don't really hurt us unless it is an acidic medium.


If we have acidic conditions at 30°, yes, we will see 20% loss of magnetic ability. If we are on the other side, if we are in high pH conditions --basic conditions--we actually only see 5 percent for the same situation.

Once we go into the 85° areas, we are starting to see losses in all of these tapes for the conditions. These are short-term exposures--two hours--for the tapes.

If we look at, again, three different tapes--this is the iron tape--again T-5; and that one behaved very, very well no matter what pH. It was actually the best performer of the group, which still says something for the iron oxides.

(Change of viewgraph)

DR. REAGOR: To give you a feeling for the analysis of this and just what we did in terms of looking at the tapes, in doing the elemental analysis of the materials, we found all of the chromium tapes still contain iron. That may be a trade secret; I'm not sure. But basically, there are still iron magnetic materials in all of these tapes except, like I said, for the T-5 which did not have any chromium.



Magnetic metal oxide particles

**Anisotropic
Iron oxide and/or Chromium dioxide**

Polymeric base tape

**PET, PVC, Polyester
Cellulose triacetate**

Magnetic Layer

**Metal oxide attached
via elastomeric polyurethane
with lubricants, dispersants, curing agents
and other additives.**

DSC

TMA

TGA

MAGNETO-HYSTERESIS

FTIR / ATR

EXTRACTIONS

X-RAY, MICROSCOPY

THERMOMECHANICAL STUDIES

Below room temp.

flat TMA trace = glassy material

30-85 deg C

extends slowly reaching a max.
strain of 0.002 at 70-80 deg C

85-200 deg C

contractions begin at 80-90 deg due to
stress relaxation. Above 140-150 deg C, it
contracts at an increasing rate until 200 deg C


Above 200 deg C

Melting point approached, PET flows and
extends, eventually breaking at 230-240 deg C


(2.98WT ± .18MAG)


WT: 3.8000 mg RATE: 20.00 deg/min

T4 - CHROMIUM DIOXIDE TAPE -



TEMPERATURE (C)





Frictional Drag Results

TAPE	Frictional Drag (pounds)	
	new	<u>aged 85C/85%RH</u>
IBT1	0.27	0.52
IBT2	0.34	0.55
IBT3	0.54	(a)
IBT4	0.61	0.49
IBT5	0.49	0.64
IBT6	0.38	0.38
IBT7	0.42	0.53

(a) insufficient aged material for testing.

W_m = weight increase due to magnetic field
= defined as magnetic weight

W_{mp} = magnetic weight at peak maximum
during 100-120 deg C magnetic transition

W_{m3} = maximum magnetic weight obtained by
cooling from 175-200 deg C in
the presence of the magnetic field
(150 Gauss).

Magneto-Hysteresis Results


(aged = 1 month in 85°C-85%RH)

TAPE	Coercive Force, H_c (Oersted)		Magnetization				Percent change (a)
	new	aged	R		S		
			new	aged	new	aged	
C1	570	520	11.1	5.99	13.6	7.68	-45%
C2	513	487	9.67	6.16	11.2	7.56	-35%
C3	537	515	12.6	7.56	14.8	9.05	-40%
C4	530	500	9.70	5.36	11.5	6.64	-43%
F5	687	681	10.1	9.91	11.9	11.80	- 2%
C6	539	500	11.0	6.23	12.9	7.72	-42%
C7	536	497	8.59	4.74	10.3	5.99	-44%

(a) percent change = average of changes in R and in S.

Magneto-TGA Analysis

Tape	W_m/M x 100	Magnetic Transition		Curie Temp. (°C)	900°C Residue (% of M)	Composition	
		onset (°C)	W_{mp}/W_m			PET %	Coating %
C1	7.5	85	4	122-5	18	76	24
C2	8.0	83	3	120-7	16	77	23
C3	8.6	89	4	122-7	20	70	30
C4	6.7	83	4	122-9	14	80	20
F5	9.2	90	1.5-2	585	18	75	25
C6	7.0	92	4	122-6	18	70	30
C7	6.8	91	4	120-7	17	75	25
PET	0	-	-	-	0	100	0



 30-35 μ m PET
 2-4 μ m coating

W_{m3} Values after Aging in 85°C-85%RH Chamber

TAPE	W _{m3} /M		percent decrease
	new	aged	
C1	0.50	0.25	50
C2	0.48	0.40	20
C3	0.59	0.3	50
C4	0.47	0.27	40
F5	0.12	0.13	0
C6	0.44	0.28	40
C7	0.37	0.19	50

We looked at the effect of iron-to-chromium ratio and its effect on being changed as we dulled the tapes. Dulling was caused when one put these through the environmental stressing situations.

And again, you can see the ratio of these peak areas was changing slightly as we go in terms of the shiny and the dull; the ratio is changing because we were basically corroding the metal, forcing it deeper into the surface. And with the technique we were using, we were trying to do a surface ratio penetrating very lightly into the surface of the tape structure. And this gave us a good feeling for the actual surface composition being changed by the actual study.

(Change of viewgraph)

DR. REAGOR: All of this work, putting together, basically brought us to a proposal for what the degradation mechanism was and that there were two mechanisms that could impact the tape. One was a chemical one, where we could have an acid catalyzed hydrolysis; it was the binder in all cases that was allowing the loss of magnetism as well as an interaction of the particle surface with the binder that could reduce some of the magnetism.

We were seeing in a physical sense that thermally induced disordering where the particles could misalign and we would lose their magnetic effect were showing some entropic relaxation.

What we saw in these cases was that these two types of degradation mechanisms would reduce the adhesion of the magnetic particles to the tape. And if you do that, then your alignment and your magnetic fields would have less effect for read/write. We would see it disrupt the particle array, and we would see orientation changes in the tape; and that was critical if we were in a longitudinal tape reading mode and all the CrO₂ particles are going off axis.

And the thing with the chromium that we liked so much was these elongated needles that gave us our best magnetic reading capability. And now, all of a sudden, I'm changing the configuration of that needle, and I'm not using its long length; I'm now using its width because it's turning away from the longitudinal axis, that we can see that that can disrupt our ability to read/write.

And then, we were seeing that we could actually create defect sites, that as you read across the tape, these disorders--the small amount of corrosion--could lead to problems.

This work was actually published in a paper, and I'll just give you the source. It was published in the book by Mittal called *Polymers in Information Storage Technology*. * The key team leader here was Trevor Bowmer, and he published it as "Characterization and Hydrolysis of Magnetic Tapes."

Our lawyers would not allow us to use the term "3480" anywhere in the publication because they thought that was a trade name and that would look as if we were doing a product analysis. So, there is this paper; it was published in 1989 based on a very large amount of work that was done on these tape materials.

Additionally, we did solvation studies and extraction studies; and we looked at the solvent impacts on all of these different tape materials. The one material that affected all of the tapes in terms of long-term exposure, we had five 15-minute exposures to different types of solvent systems and then greater than 24.

* K. L. Mittal, ed: *Polymers in Information Storage Technology*
(New York: Plenum Press, 1990)

Aging Results for C2,C3 and F5.

Tape	Conditions			W _{m3} /M after 1 mth	percent loss
	Temp.	%RH	pH		
C2	30	20	-	0.48	0-5
	30	70	-	0.46	0-5
	30	100	4	0.38	20
	30	100	10	0.45	5
	85	< 5	-	0.42	15
	85	20	-	0.39	20
	85	85	-	0.40	20
	C3	30	20	-	0.64
30		70	-	0.63	0-5
30		100	-	0.62	0-5
30		100	4	0.44	33
30		100	7	0.62	0-5
30		100	10	0.64	0-5
85		< 5	-	0.49	25
85		20	-	0.38	40
85		85	-	0.30	55
100		< 5	-	0.35	45
F5		30	< 5	4	0.13
	↓ 100	↓ 100	↓ 10		



X-ray Analysis.

TAPE	Cr/Fe Peak Area Ratios (a)				Elements found on shiny side
	new		aged 85C/85%RH		
	shiny	dull	shiny	dull	
IBT1	21.5	2.5	22.7	2.9	Fe, Cr, O, no C
IBT2	56.3	2.7	56.9	2.5	Fe, Cr, O, no C
IBT3	55.2	2.8	51.2	2.6	Fe, Cr, O, no C
IBT4	52.7	2.6	52.9	2.4	Fe, Cr, O, Cl, no C
IBT5	0	0	0	0	Fe, no Cr, O, Al, Si, no C
IBT6	54.8	2.7	51.8	2.5	no Fe, Cr, O, no C
IBT7	46.9	2.5	40.9	2.4	Fe, Cr, O, no C

DEGRADATION MECHANISM

CHEMICAL

Acid catalysed hydrolysis

PHYSICAL

Thermal-induced
disordering
Entropic relaxation

- (A) reduces adhesion of magnetic particles to tape
- (B) disrupts particle array/orientation on tape
- (C) creates defect sites - disrupting magnetic domains

We saw either no impact, or we saw a small amount of deterioration of the binder. The one product that did impact them all was halon-211, in a long-term exposure; five 15 minutes, no change; 24 hour soak time, a noticeable deterioration in the binder of the tapes.

I raise this right now because halon-211 is leaving us, and we are going to see some new halon products; and it is going to be vital that the studies of any tape medium be done with these new halon replacements to ensure that we don't have short-term exposure changes for those materials.

That was the basis for the talk I had provided for the conference; and one additional thing I wanted to bring up is a concept for you.

(Change of viewgraph)

DR. REAGOR: At this point in time, we use the chromium-based tape in all of our data centers. I have 75 maxcenters across this country with more than 10,000 DASDs in operation at this time; and those 10,000 3380 DASDs are dumping to the 3480 magnetic tape. And then, that goes into our storage.

The numbers, I can't tell you. I can't tell you the number of tapes, but it is an ungodly large number. And we are required by law to store this magnetic media minimally ten years. And it has everything you can imagine on it in terms of a telephone call placed in this country: date, time, to/from, what carrier it went through, and the whole bit.

And that is an awful lot of data that needs to be stored and retrieved; and so, from a user standpoint, we are very concerned.

The difference I see is that this is the concept of a typical data center, like our typical data center, like your typical data center. A controlled environment, filtered from a major system, and then uniquely filtered with high change rates inside.

We know that they are not highly contaminated areas yet.

(Change of viewgraph)

DR. REAGOR: Sometimes, this is typical of what they look like in terms of contamination.

The difference between a data center and a telephone switching office is basically in the particulate level and the temperature windows. From a particulate standpoint, this gives you a good feeling of what a low activity and a high activity data center--D.C.--versus a telephone center office, in terms of low activity.

Our data centers in those spaces stay under the 10,000 particle per cubic foot at .5 micron; so, basically, a 10,000 foot clean room.

High activity in a normal central office is still below the 100,000 that we set for computer standards. I bring this up since we are now moving the processing and storage capability of the computer center, and it is being merged with the electronic switch. And this is happening on a daily basis.

When you pick up your phone today and want to look at caller I.D. or have the various calling features, that is being done on the switching product with the 3480 line, some products with a 340 megabyte drive system, on-line storage of data right there in the switching office.

The difference that we see for the future is that that move to have storage within the central office is going to be greater, but we are going to have a higher need for data storage and data manipulation within the management of the system.


Corrosion Chamber

(17days > 10 years in the field)

- 30 deg C
- 70% relative Humidity (RH)
- hydrogen sulfide - 10 ppb
- chlorine - 10 ppb
- nitrous oxide - 200 ppb

in an air stream

MULTI-GAS ACCELERATED CORROSION CHAMBER



(C) COMPUTER INTERFACE
 - - - - - ELECTRICAL SIGNALS
 → GAS FLOW

We are going to have user influenced telephone systems, that as a user I can, with the touch of a dial, change what I get for information or change where I go in terms of pathways that I can select one carrier one day and another carrier another day. To do that, you now have to maintain a computer system in this environment because these are the standards for telephones.

The operating environment is from 40°F to 100, a minimum of 35 and a maximum of 120. The nominal temperature will swing 12 degrees an hour on any given day and will cycle in that range.

There is a humidity operating range of from 20 to 55% humidity; and it will float and is allowed to go as low as 10 and as high as 80, and the projections are now to go as high as 90. And in that environment, computing equipment and storage media must work. And so, we are looking at some interesting changes.

(Change of viewgraph)

DR. REAGOR: In addition, I have brought the environmental standards for gases and for contaminants. In both our data centers and our central offices, we only control particulates. The rest of it--whatever is outside --comes inside.

From the standpoint of robustness, we look at three ingredients that we recommend vendors must be able to withstand: 25 micrograms of dust per cubic meter of air; volatile organic compounds on the order of 1,200 parts per million; and ammonia to the level of 500 parts per billion. Ammonia is generated by people, and it's in the products we use to clean our floors; and it will always be there. The rest of these gases are typical 95% level, short-term exposure possibilities.

And all of these gases that are outside a data center or central office will be pulled inside that data center and central office, and their resident lifetimes are a whole lot higher than people realize.

I was very happy to hear the comment on ozone because we have measured the ozone; and depending on the outside air strategy of a building, you can have 80 percent of the outside level indoors as a constant level during the day. And that is a critical amount if your outdoor level happens to be 200 parts per billion.

And ozone is something that we actually have in our standard test now. And for our purposes, we have ozone levels as even potentially having a maximum of 500, which is everything outside plus the fact that I have laser printers and a variety of other products inside contributing to that type of level.

So, my words to you at this point are: We see chromium as a little bit less sensitive or less robust than the iron, but it is still good in our environment; but it is more of a word of wisdom to you that, for the future, we are going to need some very robust materials because we are going to be living in two environments.

(Applause)

DR. KRYDER: Write down your questions and make a note of them to yourselves because we are going to proceed all the way through the presentations before we entertain questions.

The next speaker will be Darlene Carlson. Darlene is a chemist working as Manager of Operations Support in the 3M National Media Laboratory. She has worked in magnetic media for 18 years with broad-based experience in media formulation and applications development.

Darlene was the principal investigator in the development and introduction of 3M metal audio tape in 1978. In recent years, Darlene has developed high-density digital recording tape for Government applications.

Darlene will be talking about the stability of cobalt-doped gammaferric oxide tape.

STABILITY OF COBALT-DOPED GAMMAFERRIC OXIDE

Darlene Carlson
3M National Media Laboratory

MS. CARLSON: Good afternoon. I would like to talk to you today about cobalt-doped gammaferric oxide; and you have heard a lot about a lot of different products out on the market today.

(Showing of viewgraphs)

MS. CARLSON: The question is: Why are we so interested in cobalt-doped gammaferric oxide?

If you consider the entire recording industry, cobalt-doped gammaferric oxide is basically an order of magnitude more in production than any other medium that is in the market place. It is a work horse; it is the cost driver. It is what makes the media manufacturers, both here and in Japan, churn. This is the driving force.

So, you really have to understand why people are involved in cobalt-doped materials.

The question is always: What is the stability of the materials?

(Change of viewgraph)

MS. CARLSON: What I've done is given you an overview of the various types of recording media. And what we are talking about is basically a consumer's report type thing, looking at temperature, humidity, the pollutants in the Class II Battelle environment, and then a combination of those two.

And what you can see is that the normal iron oxide, which is your typical instrumentation tape, is better than the cobalt-doped iron oxide. Now, I'll get into the reason why that is; but it is also better than the MP and the metal evaporated and the cobalt chromium that is coming down the line.


(Change of viewgraph)

MS. CARLSON: What I'm talking about here is the surface-modified or surface-adsorbed cobalt material; and what we have is an interior material of g-ferric oxide in there, which is the normal 300 Oersted instrumentation tape, that has been surface-modified with cobalt, which gives you a very interesting engineering material.

And if you remember John Mallinson, who was talking about the paranoia that people have in the recording industry because they are concerned about the next generation of thing that will overcome the recording media, my contention is that the reason why we are still with magnetic recording media is the versatility of this particular species.

We have a coercivity range of 450 to 900 Oersteds; and we have some control over the aspect ratio and the surface area. So, this is a very interesting engineering material, which has allowed us to 'fend off' all the different media recording technologies that are out there.

Co- γ -Fe₂O₃ Pigment



Surface Adsorbed Co

Hc (oe) - 450-900

σ s (emu/gm) - 84

l/d - 5:1 to 10:1

Storage Systems Using Co-γ-Fe₂O₃ Tape




System ID	Head	Media Format	Thickness (μ)	Capacity/Pkg (Gbits)	Durability (passes)
QIC 525	Fixed	1/4" Cart.	11	4.2	10,000
QIC 1350	Fixed	1/4" Cart.	13	10.8	10,000
Interferometrics	Fixed	1"x14" reel	16	5500	1,000
DCRsi	Trans.	1" Cart.	25	380	200
VHS	Helical	1/2" Cart.	20	32	NI
Digidata	Helical	1/2" Cart.	20	43	NI
VLDS	Helical	1/2" Cart.	16	80	NI
ID1	Helical	19mm Cart.	16	300	300
DCR	Helical	1" Cart.	25	240	200


* NI - No Information

Recording Bit Density Comparison

Linear, Low Density: 2060 bits per inch, 20 mil width; 165X.



Helical, High Density: 45,000 bits per inch, 1.2 mil width; 165X.



(Change of viewgraph)

MS. CARLSON: It is used in a wide variety of systems, both in fixed transverse and helical scan material and a lot of different media formats. And it is coated on a lot of different thicknesses of materials.

And you can see that it has a wide range of capacity for package; and again, this is in gigabits, not gigabytes. And, depending upon the application, the durability has a wide range.

Now, one question is: How long can you archive this material?

(Change of viewgraph)

MS. CARLSON: Unfortunately, this particular viewgraph isn't very clear; but what I would like to do is do a bit of a comparison. This is a 2K bit; this is just one bit right here. And this is a -- technique of a search track on a digital recording system.

These bits are 20 mils wide, and these represent probably about three times as large as your 6250 bpi material that everybody is used to. But you can see that this particular -- shows you a 45K bits. So, when you ask: what is the archivability of cobalt-doped material?, it really depends upon your bit density; and the implementations and archivability of a 2K material or 6250 material is going to be much, much different than a 45 Kfci or as an 8 millimeter, we are going up to 75 Kfci.

So, you can see that these bits are actually getting smaller as you evolve into different products.

(Change of viewgraph)

MS. CARLSON: Similar to what Barbara was talking about, magnetic tape components have a lot of different contributions that could impact archivability. We have a coating layer of polyester, and some people have backcoating

And each of these components can be the weakest link in any archive system. Now, what I would like to do is take a broad-brush approach on what can happen to the various individual components in a tape system.

(Change of viewgraph)

MS. CARLSON: If you look at the substrate, what you are interested in is the dimensional stability and how that affects crossplay.

If you look at any ANSI standard or SMPTE standard or MIL standard, what we normally do is define the tape format on the tape. That's the system in which the head lays down your track of information.

Because of either tape shrinkage or tape creep, which is involved with tension on the tape, you can actually have the polyester changing so that your recorded track is outside of that tape standard. And therefore, you will not get crossplay; and sometimes, it can be so bad that you can actually lose self play.

So, dimensional stability is well known, and it is a concern in any archive environment.

(Change of viewgraph)

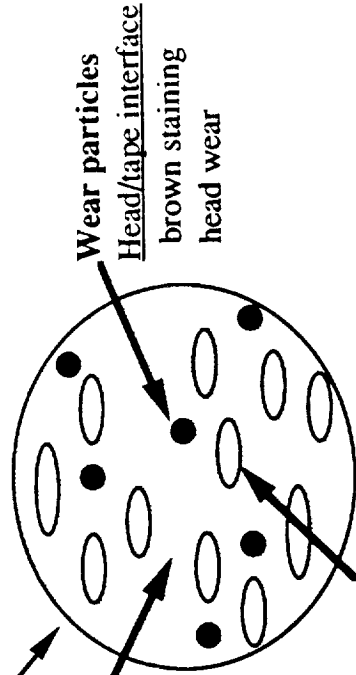


Magnetic Tape Components:
Environmental Sensitivities of
Co-γ -Fe2O3 Tape

Substrate:
 -Dimensional Stability:
 directional modulus
 temp./humidity coeff.
 stress relaxation

Magnetic Coating

**Organic Binder/
 Lubricant System**
Surface Tribology Changes
 stick-slip
 swelling...
Composition Changes
 hydrolysis
 molecular wt. change
 lubricant loss




Loss of Magnetic Moment
 diffusion (Co- Fe2O3)
 stress demagnetization

Backcoat:
Backcoat/Magcoat
 Adhesion (blocking)

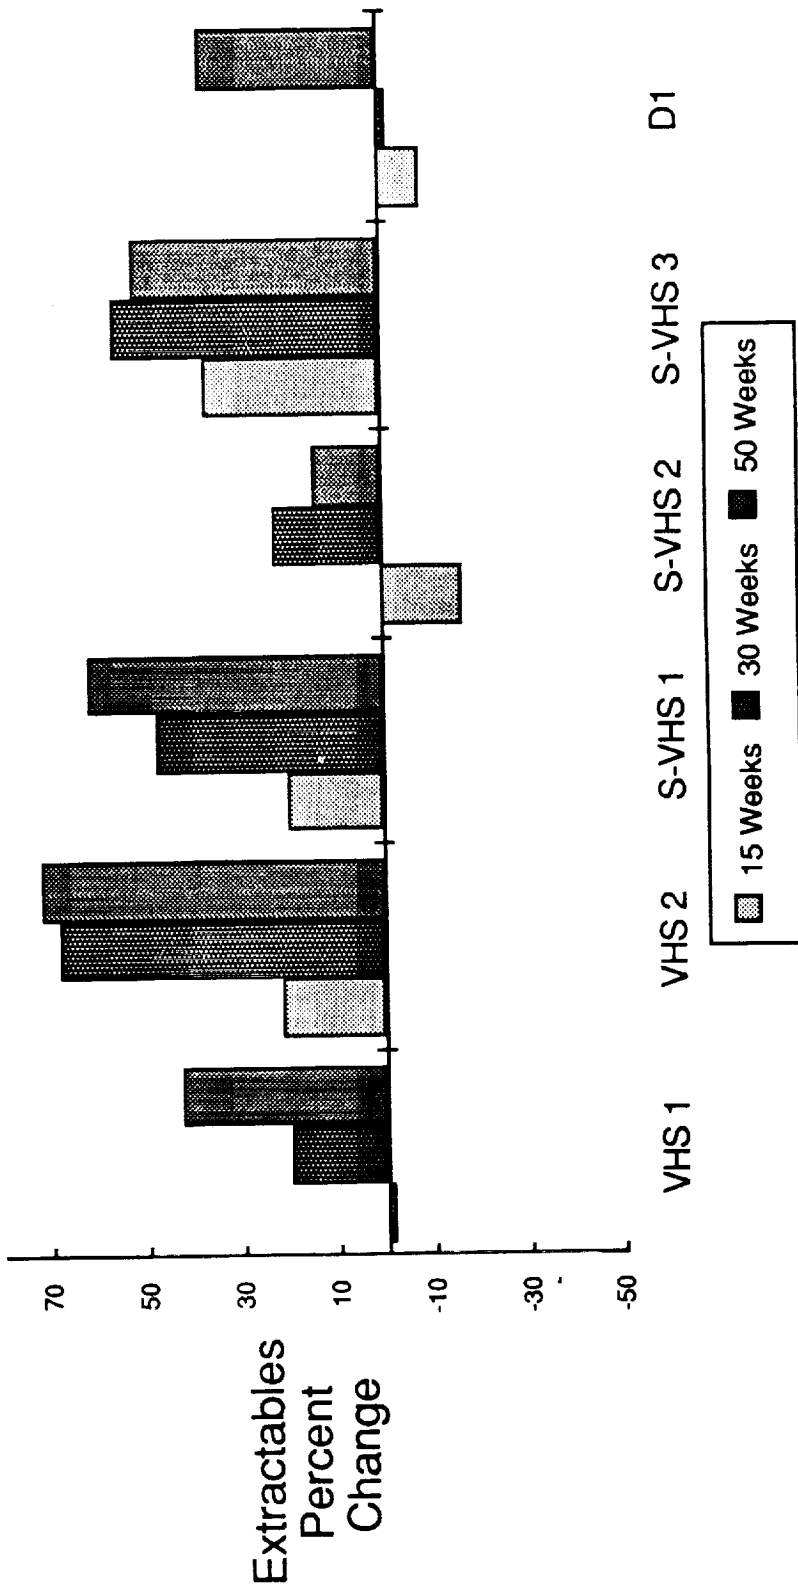



Dimensional Stability and Crossplay






58°C - 90% RH Acetone Extraction





58°C - 90% RH Coercivity



NML DMC 7/24/91

MS. CARLSON: Another thing that is of importance is the hydrolysis. Barbara talked earlier about hydrolysis of the binder and compared instrumentation tape with various tapes; and this is Co-doped material; and what you can see is that hydrolysis is still with us--it is alive and well.

And this is a solvent extraction at 58° C/90% RH; and we are looking at two different vendors of VHS tapes, three different vendors of S-VHS and D-1.

And you can see that, after 50 weeks--which is an extremely long period of time--that we still have hydrolysis. Now, the purpose of this study is that we have a family of curves that we are developing, and we are developing various equations from them; and we will have models to predict for all the various materials.

But you can see that in the cobalt-doped materials, we do have hydrolysis; and we will have to manage for that change.

(Change of viewgraph)

MS. CARLSON: If you remember the earlier slide, I told you that the cobalt-doped material is not as stable as other things like g-ferric oxide. And the reason for that is because of the coercivity change.

What happens is that the cobalt actually diffuses into the g-Fe₂O₃; and what you do is you lose coercivity. And again, what we are plotting here is the coercivity change at 58° C/90% RH at up to 50 weeks.

And you can see that there is a variation from different types of cobalt to how fast or what the mechanism of the cobalt migrating into the g-Fe₂O₃ is.

(Change of viewgraph)

MS. CARLSON: Now, in contrast, the f_R , which is the measure of the remanent magnetization, is very, very stable; and I don't know if you remember this from Barbara's talk, but chromium is much worse than this. And this is basically within the measurement error.

(Change of viewgraph)

MS. CARLSON: In summary, we have to be careful about how we specify any archive requirements. You can archive cobalt-doped gamma materials for up to ten years. Now, the caution that I have is that you have to retension and clean the tape.

In most mass storage situations recertification is not an issue because what you are doing is you are recording on one tape or reading from another or whatever. But if you are reusing a tape, it should be recertified.


The archive requirements are system-specific, specifically because there is such a wide range of requirements out there; and each system has its own peculiar way of being implemented.

And many times, a tape vendor recommends a certain range of temperature and humidity controls; but the user environment, as Barbara was saying, is outside of that. We recommend never putting a tape outside of a 70% RH environment.

And as we heard today, they realize that we will store at 95% RH. So, the question is: what implementation procedures should we utilize if we are storing at 95% RH? And that is what we are in the process of developing.



58°C - 90% RH Phi R



NML DMC 7/24/91


Overview of Archival Stability of Recording Media

Environment	Magnetic Pigment						
	γ -Fe ₂ O ₃	CO- γ -Fe ₂ O ₃	BaFe	CrO ₂	MP	CoNi	CoCr
Temp \geq 50°C	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Humidity \geq 75%	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Pollutants \geq CI II*	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Temp & Humid	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Temp & Poll	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Humid & Poll	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
T & H & Poll	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>

PRODUCT:	Audio Cassette	Video	4 MByte	Video	Hi-8mm	Next
	9T Comp. Tape Lo-Density Diskettes	D-1 Hi Bias Audio Hi-Density Diskettes	Diskettes Hi-8mm	3480 Cart. Hi-Bias Audio	R-DAT D-2 Audio	Generation
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

KEY: = GOOD No corrosion or signal loss problems expected
 = FAIR May be suitable if some signal loss and/or bit errors can be tolerated
 = POOR Unsuitable for storage under this condition

NOTES: *Interim recommendations, based on results of NML and others, as of January, 1991.*
 Does not include possible binder and substrate problems not specific to media type.
 *Battelle Class II Environment



Summary

Co- γ -Fe₂O₃ Tape Stability

- Archive Stability: 10 yrs in a dry, controlled environment
- Retensioning
- Cleaning
- Recertification
- Archive requirements are system specific
- Wide range of system performance
- User requirements outside recommended controlled environments
- Distributed Storage Systems require Shipping
- Anomalous events

I know that there are going to be some network managers here who would like to throw tomatoes at me at this particular moment ; but even with local area networks and wide-area networks, media will require shipping. And the analogy I always cite is: Remember when they said that we were going to have a paperless office? the same thing here. Media will require shipping; and therefore, we have to be concerned about how it is being shipped.

And we also have to be concerned about anomalies such as fire, water damage as a result of people spilling coffee or whatever, and various other shipping impacts such as DC motors, erasing the tape.

So, we have to be able to do a recovery from these types of anomalies and events. And we have been involved with some tape recovery operations.

That concludes my talk.

(Applause)

DR. KRYDER: We will hold the questions until the discussion. Thank you, Darlene.

Our next speaker is Kazuhiro Okamoto. He obtained his M.S. Degree in Physics from Hokaido University. He has been working at Sony for the past eight years. He is a materials engineer working in the Magnetic Particle Development Group.

The title of his presentation is "Stability of Metal Particles/Metal Particulate Media."

STABILITY OF METAL PARTICLES/METAL PARTICULATE MEDIA

Kazuhiro Okamoto
Sony

MR. OKAMOTO: Thank you. Today, I would like to talk about the study of metal particles and the metal particulate media.

Stability of Metal Particle and Metal Particulate Media

Kazuhiro Okamoto

Sony Magnetic Products Inc., Tagajo, Miyagi, Japan

1. Introduction

Metal particulate (MP) video tape was launched for 8 mm video tape in 1985. Since then MP tapes have been applied to several consumer formats and instrumental formats because of its superior electrical performance. Recently data storage media, such as DDS, and D - 8, have started employing MP tape.

However, there are serious concerns with archival stability of MP tape particularly in the case of data storage use, as metal particles essentially have problems with chemical instability and are susceptible to oxidation and corrosion. Although there have been some studies about the archival stability of metal particles or MP tapes (1) ~ (3), a clear understanding has yet to be reached.

In this paper, we report the stability of magnetic properties of current metal particles, and then discuss the new technologies to improve the stability further.

2. Stability of Current Metal Particles

A metal particle is composed of three layers. Pure iron core with $\sigma_s=222$ emu/g is covered with a passivated iron oxide layer, and it is covered again with a ceramic layer composed of Al and/or Si compounds. The outer two layers prevent the oxidation of the iron core. The important role of the surface ceramic layer to the stability of metal particles is shown in Fig. 1.

The stability of metal particles is improved still more when they become part of the tape, since they are covered uniformly with binder molecules during the mixing process (Fig. 2). Fig. 3 shows the archival stability of the RF-output of Betacam tape. Little degradation is observed even after long term storage in a 45°C, 80%RH environment. So, from a practical viewpoint the current MP tape has sufficient stability due to the multilayer protection of the iron core.

3. Newly Developing Technologies

3-1. Dense Oxide Layer

The morphology of pure iron core is thought to be measured with X-ray crystallite size D_x . The aging effects of σ_s and D_x of current metal particles under 60°C, 90%RH environment is shown in Fig. 4. The initial degradation of σ_s occurs without damage to the morphology of iron core. This fact indicates that the degradation of σ_s begins with the destruction of the oxide layer. Therefore the formation of a uniform and dense oxide layer would promise improvement in the stability of metal particles.

Fig. 5 shows the comparison of the degradation ratios of flux density between current MP tape and improved MP tape which utilizes new metal particles with a dense oxide layer under the storage conditions of 65°C, 90%RH. The improved MP tape is twice as stable as the current one which is already stable.

3-2. Anticorrosive Agent

In order to improve the archival stability to the highest degree possible, we are researching anticorrosive agents. A low molecular organic compound would be selected as an anticorrosive agent to minimize the decrease of magnetization by the use of non-magnetic materials. The organic anticorrosive agent which is adsorbed onto the surface of the metal particles changes the hydrophilic surface into a hydrophobic surface and cuts off direct contact with oxygen and moisture in the atmosphere.

The results of the aging test of the currently used metal particles, the improved ones (with a dense oxide layer), and the anticorrosive agent treated particles are shown in Fig. 6. The stability of metal particles are amazingly improved through the application of the anticorrosive agents. The degradation ratio of treated metal particles is comparable to that of magetite particles which have the same size.

4. Stability of Coercivity

The coercivity of metal particles usually shows almost no change or a little increase from the aging test. The degradation of coercivity, however, is observed only when the magnetization decreases extremely.

5. Conclusion

Currently used metal particles are stable, as they have protection layers of ceramic coating and passivated oxide. Therefore the MP tape has sufficient stability for practical use. In order to improve the stability further, new metal particles which have a uniform and dense oxide layer and an anticorrosive agent are under development.

References

- (1) T. D. Lee, S. Hu and N. Madulid; IEEE Trans. on Magn. MAG-23 (1987) 2880
- (2) D. E. Spiliotis; IEEE Trans. on Magn. MAG-26 (1990) 124
- (3) Y. Yamamoto, K. Sumiya, A. Miyake, M. Kishimoto and T. Taniguchi; IEEE Trans. on Magn. MAG-26 (1990) 2098

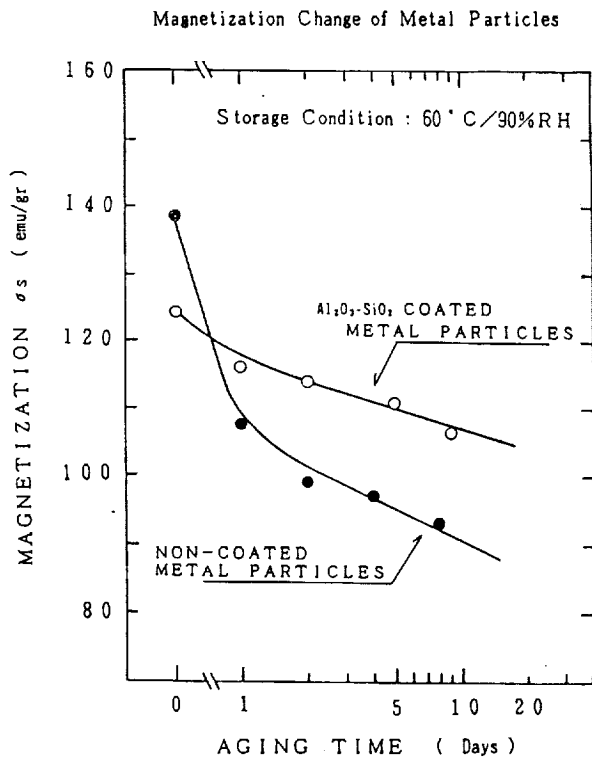


Fig. 1

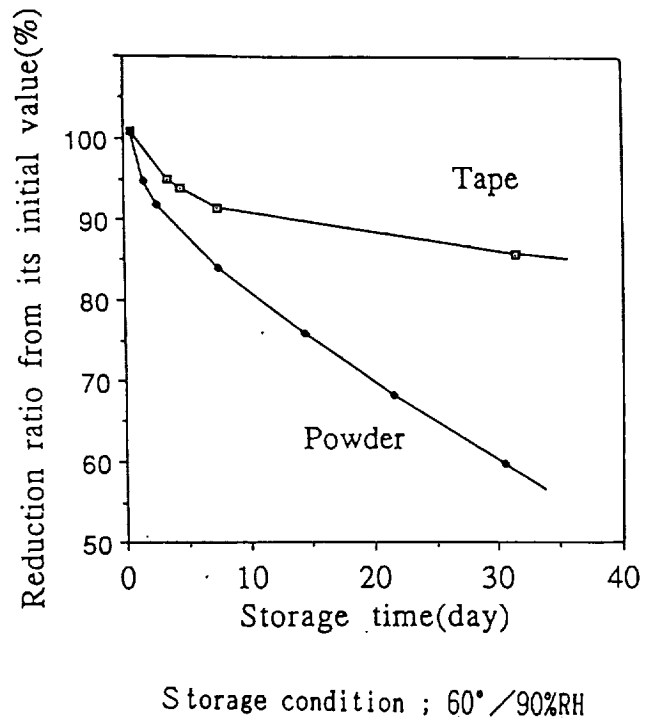


Fig. 2

Video S/N stability of Betacam
Metal and Co-oxide tape

SONY


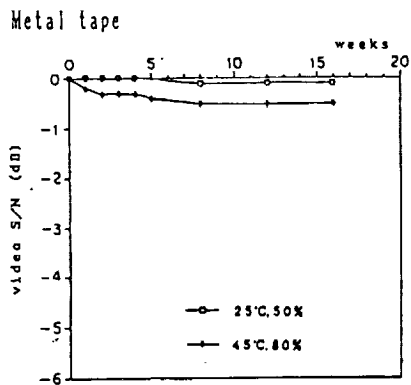


Fig. 3

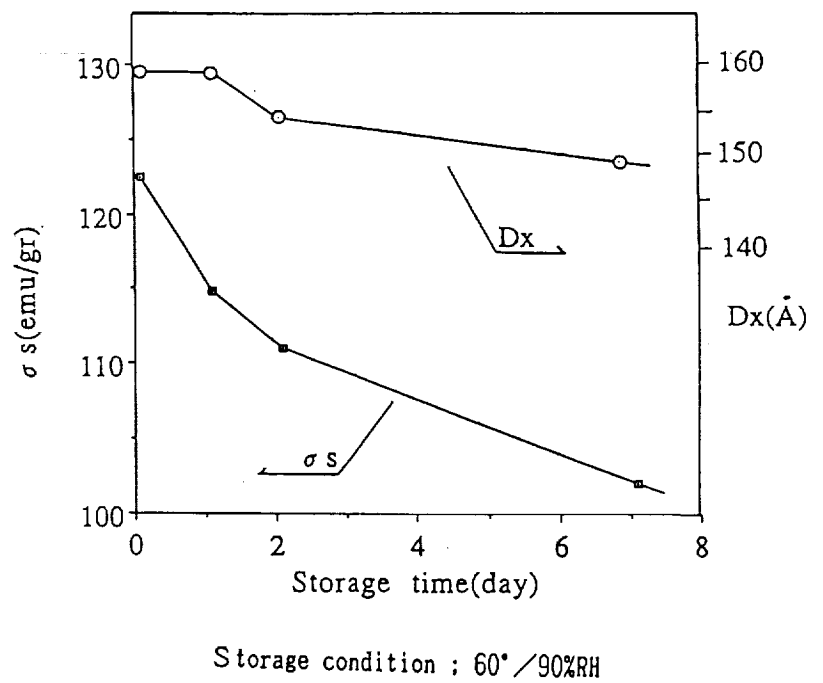
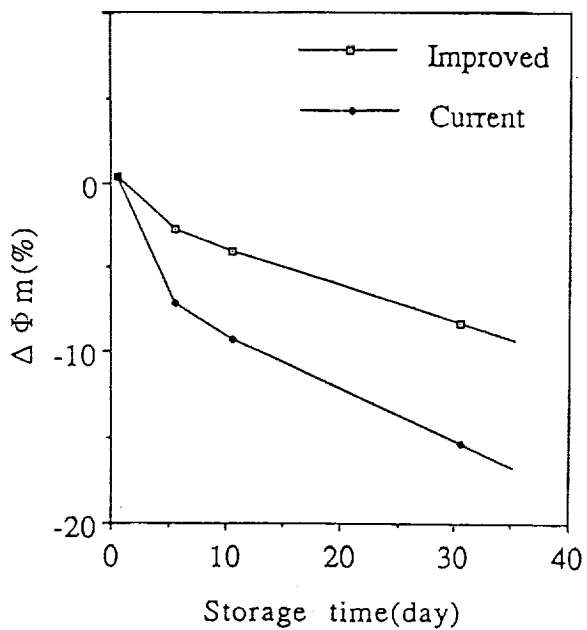



Fig. 4



Storage condition ; 65°/90%RH

Fig. 5



Storage condition ; 60°/90%RH

Fig. 6




Fig. 7

DR. KRYDER: Next, we have a two for one deal. Ampex is going to split their time between two speakers; they promise to keep both their talks within the 15-minute time limit.

The first talk will be by Allan Hadad. Allan has a B.S. Degree in Chemistry and an M.S. Degree in Materials Engineering. He started his professional career as a synthetic organic chemist working in the application of high-performance polymers in the wire and cable industry.

Allan began working in the recording industry in 1974 as a formulation development chemist. He has developed various tape formulations for both consumer and professional use. He has participated in the application and introduction of products based on the evolution of new magnetic materials, from cobalt-adsorbed iron oxides to iron particles.

Allan has also managed Ampex's metal-evaporated tape program. Allan's current assignment is Manager of Formulation Development for Ampex Recording Media Corporation in Redwood City, California and Program Manager for an all metal particle development effort in both Redwood City and Opalatica, Alabama.

Allan will be speaking on an empirical approach to predicting archival stability and primarily discussing metal particles.

AN EMPIRICAL APPROACH TO PREDICTING ARCHIVAL STABILITY

Mr. Allan Hadad
Ampex Corporation

MR. HADAD: Good afternoon. Of course, John and I are going to be splitting the time; but we never did come to an agreement about where the split would be. So, since I'm here first --

(Laughter)

MR. HADAD: I would also like to thank Dr. Hariharan for inviting me. He called me about a month ago; he asked me what I thought about the concept of free speech. I said I agreed with it; and he said: Good. I'd like you to give one at this conference. So, here I am.

(Laughter)

(Showing of viewgraphs)

MR. HADAD: I would like to talk about the stability of high-density media, particularly of course metal particles, that I have been involved with pretty much since 1978.

I like to separate environmental stability from archival stability. I feel that most of the testing that we have done so far--and much of the data that you have maybe even seen today--is environmental stability, discrete conditions in which tapes have been exposed. You have seen some data. I call that a user condition, perhaps during acquisition, as opposed to long-term storage.

Archival stability, I think, has to be determined by predictive techniques, mathematical modeling, taking a lot of data from a variety of environmental exposure conditions, sort of lumping it together and building an expression, to be able to predict what will happen at any temperature and humidity condition.

What will the environmental stability of a tape be? How long will it really last? What's the lifetime of tape?

So, what I am going to present here is some data that we have acquired on environmental stability, basically to give the users confidence that under a certain variety of conditions, we feel very confident ourselves that metal particles do produce a viable recording medium for high-density information.

I will follow that then by an empirical approach that we have developed and are trying to propose as really the means or the vehicle to predict what the lifetime of tape will be.

(Change of viewgraph)

MR. HADAD: When we first got into the metal particles, we worked initially a great deal with the particle manufacturers themselves. They have made a lot of changes for us, as well as for all the other magnetic tape manufacturers, to produce much more stable particles, more passive particles.

Just some data that we've got. It was collected between 1987 and 1990. Three different particles that have been evaluated in our laboratory since 1987; and basically, the conditions are essentially ambient: 22° C/40% relative humidity.

As you can see, obviously there are a variety of materials available from the particle manufacturers; but in general, what we are seeing here is that after three and a half years--and this is naked tape; it is not protected by a cassette case

(Change of viewgraph)

MR. HADAD: Tapes exposed to the ambient environment that, after three and a half years, have experienced losses of about or less than 2 percent in B_m .

(Change of viewgraph)

MR. HADAD: Going a little further then, looking at some of those environmental exposure conditions, I've got a chart here. We are looking at four different tapes, commercially available tapes--some of our own, some of our competitors--at three different environmental conditions.

Ambient again, 30° C, 70% relative humidity, and 54° C, 85% relative humidity--fairly extreme conditions. Although I believe that certain users may for short periods of time have exposed the tapes to these conditions, especially during initial acquisition of data, or perhaps through transient activities, as Darlene said--a fire, a sprinkler system going off--that often tapes can be exposed to some of these more extreme conditions.

The bottom line: the loss in B_m . Obviously, for ambient conditions again, we see some variation; but basically, around a 1 to maybe a 3 percent loss, after nearly six months exposure.

Moving up to the 30° C /70% relative humidity, losses on the order now--it depends on the manufacturer--around 2 to 3 percent loss in B_m again.

And then finally, the very extreme case, where now we start to see very rapid or much more devastating losses in magnetic performance, although again we wouldn't certainly call that an archival or storage condition.

Now, perhaps for certain short amounts of time, and it could be tolerated maybe on up to six months, or maybe a couple of weeks, in an environment like that, that data could be collected; and acceptable losses in magnetic performance would occur.

(Change of viewgraph)

MR. HADAD: Going out a little bit further, some D-2 tape, again under ambient conditions, six months worth of data. And we see somewhere around a 1 dB loss after 24 months--two-year data now.

In this case now, we are looking at tapes that are protected by the cassette. We are moving a little bit further along now where, on the tapes themselves, there is an extra level of shielding. Of course, the particles degrade the fastest; we saw some of that in the paper given by the gentleman from Sony.

We put it into the tape system; the binders give it a degree of shielding. Now, it's inside a cassette. There is an additional degree of shielding. Two-year data then, around a dB loss under ambient conditions for this particular sample.

Of course, there is one thing about all of this. If you look at two and three-year data, that's the old stuff now--right? Improvements have been made since.

(Change of viewgraph)

MR. HADAD: A real quick look at errors. This seems to be the real key. Again, about six-month data and essentially with some variability of three of these samples at ambient conditions, 22° C/50% relative humidity, with no dramatic change in the errors per frame here with three of the manufacturers.

One of the tapes here apparently had some debris on it that was actually being removed with the repetitive testing over a period of time. Of course, we have seen that before.

(Change of viewgraph)

MR. HADAD: Under more extreme conditions there is an increase in errors, again cassettes. We've got one here--actually, it was the same one that was cleaning up-- but fairly ragged out here at the end.

But in general again, no real change in the error count at 30° C/70% relative humidity for a period of almost six months.

Okay. That's environmental stability. That's not really kinetics. As a chemist and as a materials scientist, I like to reduce things down to mathematical expression and predict: How long will these tapes really last under normal storage conditions?

We have to come up with a coefficient that relates; you can sort of additively combine all the exposure conditions and determine how much time is really left on this tape.

(Change of viewgraph)

MR. HADAD: My favorite chart is the psychrometric chart. For those of you who have not seen it, it basically relates temperature to water content in air. Relative humidity are the curved lines. Absolute humidity are the horizontal lines.

What we have done is through an experimental design, these points are actual conditions that we have exposed powders to. Now, you start at the front end, and you adjust the powders themselves; there are about 12 conditions here. You have to do a multitude of testing.

Powders are subjected to these temperature/ humidity combinations for periods extending up to several months. Data is collected periodically; a curve is drawn of the degradation behavior and the slope of that line is determined. That's the degradation rate.

Those degradation rates at each of these points we will put into the computer, crunch it, and a mathematical equation was generated, describing -

(Change of viewgraph)

MR. HADAD: degradation rate as a function of, in this case, the silicon content of the powder that we are looking at as a variable, relative humidity and temperature, cross products between those and the square terms.

(Change of viewgraph)

MR. HADAD: We have a nonlinear equation that allowed us to generate lines of constant degradation rate, if you will, as a function of temperature and humidity for a given particle and gives us the ability now to predict long-term behavior at any temperature/humidity combination within the experimental region.

What these numbers are, again, are lines of constant degradation rate. One thing it shows is, that the relative humidity is the big driving factor in degradation of iron particles; temperature is not as great a factor until you get below this critical value here whereas say, you increase temperature, degradation rate does increase.

N93-16330

AN EMPIRICAL APPROACH TO PREDICTING LONG TERM BEHAVIOR OF
METAL PARTICLE BASED RECORDING MEDIA

By


Allan S. Hadad
Ampex Recording Media Corporation
Redwood City, California

Narrative

Submitted for the

National Space Science Data Center's
Conference on
Mass Storage Systems and Technologies for Space and Earth
Science Applications

Goddard Space Flight Center
Greenbelt, Maryland
July 23-25, 1991



**AN EMPIRICAL APPROACH TO PREDICTING LONG TERM BEHAVIOR OF
METAL PARTICLE BASED RECORDING MEDIA**

Alpha iron particles used for magnetic recording are prepared through a series of dehydration and reduction steps of $\alpha\text{-Fe}_2\text{O}_3\text{-H}_2\text{O}$ resulting in acicular, polycrystalline, body centered cubic (bcc) $\alpha\text{-Fe}$ particles that are single magnetic domains. Since fine iron particles are pyrophoric by nature, stabilization processes had to be developed in order for iron particles to be considered as a viable recording medium for long term archival (i.e. 25+ years) information storage. The primary means of establishing stability is through passivation or controlled oxidation of the iron particle's surface.

The usual technique of producing the protective layer is through re-oxidation of the iron particle's surface after synthesis starting with a mixture of 0.1% O_2 and 99.9% N_2 . The oxygen content is slowly increased to 20% (the composition of air) so as to maintain the reaction at room temperature. This results in a particle that is stable in air provided it is not subjected to any form of 1) mechanical abuse that could disturb the outer layer, or 2) source of heat that would initiate combustion.

PRECEDING PAGE BLANK NOT FILMED

The nature of the passive layer on iron particles has been found to consist of either Fe_3O_4 , $\gamma\text{-Fe}_2\text{O}_3$ or a mixture of the two. The thickness of the passivation (or total) oxide layer is typically about 3.0 nm. The condition of passivity slows down the oxidation of iron, but, the formation of the iron-oxide layer around metallic iron particles does not preclude further oxidation. The continued stability of iron particles is controlled by the integrity of the oxide layer and the kinetics of diffusion of iron ions through it.

Since iron particles used for magnetic recording are small, additional oxidation has a direct impact on performance especially where archival storage of recorded information for long periods of time is important. Further stabilization chemistry/processes had to be developed to guarantee that iron particles could be considered as a viable long term recording medium.

In an effort to retard the diffusion of iron ions through the oxide layer, other elements such as silicon, aluminum and chromium have been added to the base iron to promote more dense scale formation or to alleviate some of the non-stoichiometric behavior of the oxide or both.

The presence of water vapor has been shown to disrupt the passive layer, subsequently increasing the oxidation rate of the iron.

A study was undertaken to examine the degradation in magnetic properties as a function of both temperature and humidity on silicon-containing iron particles between 50-120°C and 3-89% relative humidity. The methodology to which experimental data was collected and analyzed leading to predictive capability is discussed.

EXPERIMENT

To study the effect of temperature and humidity on the stability of passivated iron particles as a function of silicon content, three particle batches were prepared from the same precursor containing 0.81, 0.95 and 1.12% silicon by weight.

As stated earlier, the oxidation process of iron particles is diffusion controlled. The increase in oxide thickness is determined by how fast iron ions can migrate to the surface (passivation kinetics). It is assumed that the reduction in σ_s is proportional to the oxide thickness; therefore, Fick's second law of diffusion for concentration dependent systems can be applied:

$$Y^2 \approx D t$$

where:

Y = diffusion distance in cm

D = diffusivity in cm^2/sec

t = diffusion time in sec.

For this analysis, the data are plotted as specific magnetic moment versus the square root of time. The resultant slope is therefore the rate of degradation expressed as EMU/gram-sec^{0.5} for a given temperature and humidity condition.

Analysis of the interactive effects of temperature, humidity and percent silicon on the degradation of magnetic properties of iron particles were determined using a statistical experimental design package run on an IBM AT compatible computer. The characteristic measured was the log₁₀ of α_s loss per second^{0.5} (rate of degradation).

The procedure of data analysis consists of selecting an initial mathematical model (usually a quadratic equation to analyze interaction and curvature effects of the form $Y = C_0 + C_1X_1 + C_2X_2 + C_{13}X_1X_3 \dots + C_{11}X_{12} + C_{22}X_{22} \dots$) where C_n is a constant and X_n is the level setting of the respective factor. The model is then used to "fit" the data.

RESULTS AND DISCUSSION

Dry Conditions. The first series of oxidation rate experiments were run at the "dry" conditions for the selected temperatures. Oxidation at various temperatures allowed generation of α_s vs. time curves for iron particles containing 0.81, 0.95 and 1.21% silicon sample at various temperatures.

Figure 1 shows the change in specific magnetic moment with time for the 1.12% silicon sample at various temperatures. As temperature increased, the rate of degradation of σ_s also increased. Differences in the rate of degradation at a given temperature between the three samples studied containing different amounts of silicon are slight, indicating that the iron particles all behave the same at low humidity, independent of Si content or initial σ_s value.

Activation energies for degradation of the 0.81, 0.95 and 1.12% silicon-in-iron powder samples were determined to be 0.383, 0.263, 0.333 eV, respectively when the rates of degradation were plotted against the absolute temperature (Arrhenius plot). Thus, the energy requirement to initiate the oxidation reaction for the various Si-contents is essentially the same and all degrade at essentially the same rate as a function of temperature. However, these activation energies are only ~10% of the value for diffusion of iron through any of the possible iron oxides. The implication is that the primary mechanism for the loss in magnetic properties is not due to lattice diffusion but rather diffusion along some short circuit path such as linear or planer defects present in the oxide shell.

Moisture Effects. Magnetic moment loss is shown as a function of humidity at constant temperature (70°C) for the 0.81, 0.95 and 1.12% silicon samples in Figures 2, 3 and 4 respectively. As the humidity increases, the rate at which the magnetic moment degrades increases markedly with decreasing Si content.

Figure 1.




Figure 2.




Figure 3.





Figure 4.



Degradation is most rapid for the 0.81% Si-containing iron particles. The sample containing 0.95% silicon is more resistant to degradation until the relative humidity is above 59% at 70°C. Finally, the sample containing 1.12% silicon is the most resistant to degradation exhibiting almost the same rate of loss at 7% relative humidity as it does at 74% relative humidity.

There appears to be a threshold value of relative humidity that is observed at the 0.95% silicon level and possible the 1.12% level, above which the degradation rate increases very rapidly.

As silicon content increases, the critical humidity (the humidity where rapid degradation occurs, at a given temperature) also increases. However, the degradation ceases at σ_s value of 50-60 EMU/gr for the 0.81% silicon sample (Figure 2). The trend appears to be the same (a saturation limit) for the 0.95 and 1.12% Si particles also.

Interaction of Temperature, Humidity and Silicon. The three parameters of this study (temperature, percent relative humidity and percent silicon content) interactively influence the degradation of the iron particles. Thus, statistical analysis was performed to quantify the contribution of each variable and present it in a manner that could be easily visualized. The measured characteristic was the \log_{10} of the degradation rate in specific magnetic moment (EMU/gram-second^{0.5}) as derived from the results of the experimental matrix.

The design predictor equation determined using XSTAT is as follows:

$$\begin{aligned}\log_{10}(\text{Rate}) = & -14.76 + 25.26(X_{Si}) + 0.04176(X_{RH}) \\ & + 0.0015(X_T * X_{RH}) - 0.05963(X_{Si} * X_{RH}) \\ & + 0.000068(X_T)^2 - 12.49(X_{Si})^2 \\ & + 0.000219(X_{RH})^2\end{aligned}$$

where:

X_{Si} = concentration of silicon in weight percent

X_T = temperature in degrees Celsius

X_{RH} = percent relative humidity

The percent variance explained (how well the regression equation predicts the data) was 96.03% so the model can be considered as quite good. Solving the predictor equation via the computer produces two dimensional contour plots of percent relative humidity versus temperature at constant silicon content. The values of constant degradation rate were transferred to a psychrometric chart (Figures 5, 6 and 7) so the relationships between absolute humidity (expressed as pounds of water per pound of dry air), percent relative humidity and temperature as a function of silicon content could be observed. The non-linear degradation behavior becomes very obvious when presented in this manner. A rate value can be taken from the contour plots for a given temperature, humidity and percent silicon content and be used to determine how long it would take for the magnetic moment to degrade to some predetermined value.

Figure 5.




Figure 6.





Figure 7.



For a constant absolute humidity (above the critical value, depending on silicon content) increasing temperature will cause the degradation rate to decrease. It appears that relative humidity is the controlling factor. The corrosion rate appears proportional to the thickness of the adsorbed water layer on the surface of the test specimens. The higher the relative humidity, the thicker the adsorbed layer, hence the faster the corrosion rate.

The behavior observed in the contour plots follows an expression of the form $\text{Rate} = A_e (b\%RH)e^{-(Q/RT)}$, where at low humidity temperature is the controlling factor, then a transition occurs where humidity effects prevail. It is obvious that special precautions should be taken when attempting to predict long term, low temperature behavior based on high temperature data with a system that is humidity sensitive. Actual archival stability could be less than expected unless humidity control is considered. The value of statistical analysis to the interpretation of oxidation behavior is clearly evident.

CONCLUSIONS

This study has shown the effects of percent silicon content, temperature and percent relative humidity on pre-passivated fine iron particles used for magnetic recording.

When the iron particles were exposed to various temperatures between 50 and 120°C at very low humidity it was observed that the degradation rates were not affected by either the silicon content or the initial value of the magnetic moment. TEM micrographs revealed that the oxide layer grew thicker leading to a condition of passivity.

It was shown that for a given temperature there exists a critical relative humidity value above which the degradation rate of magnetic moment increases markedly. The presence of silicon appears to increase the critical humidity value at which rapid degradation occurs. When the magnetic moment degrades to 50-60 EMU/gram it remains constant for additional exposure time.

A parametric expression has been proposed to relate silicon content, temperature and humidity to the initial rate of specific magnetic moment degradation for the iron particles used in this study.

Up here, if you increase temperature and maintain constant absolute humidity in the air, the degradation rate actually slows down, which means you might make a prediction conservatively that you might have better lifetime than you thought.

What these numbers correspond to is that they are EMUs per gram per root second of degradation rate; and what I can really translate it to is that to sustain a 10 percent loss in the EMUs per gram of the powder--the magnetic signal--this -1 value corresponds to 3 minutes that will lose 10 percent of its signal.

By the time we get down here to the .01 rate, 13.5 days to lose 10 percent; and then, the next contour line, which would be down here, .001, which is a five-year rate. Five years under more normal operating conditions based on data from much accelerated conditions. Five-year lifetime for the particles to lose 10 percent of their signal.

Again, remember, if we get it into tape, there is an additional layer of shielding; put it in a cassette, another layer.

What we are doing is proposing this methodology to do testing to really be able to do empirically derived predictive testing for determining what is the lifetime of tape--what is acceptable. How long will that really take? You have to be able to quantify what these rates are--these rate constants--as a function of both temperature and humidity.

So, to summarize, our feeling is that metal particle media based on the data I have shown you, both for environmental exposure conditions--we saw numbers on the order of 1 to 2 percent over three years--some of this methodology here.

We feel that metal particles are a viable vehicle for high-density recording. They are not as stable as the gamma iron oxides. They are not as stable as the chromium dioxide. But we feel that, within acceptable limits, it is a viable vehicle and that, for whatever lifetime you need before it's recopied to a new format or rerecorded, we think that the metal particles themselves do provide a good means--an excellent means of recording at high densities.

I have also then shown you that the data must be collected and crunched to give true archival predictions as to what the lifetime of the tapes will be. The manufacturers will be able to substantiate their claims with this kind of data as to how long the tape will last.

Thank you.

(Applause)

DR. KRYDER: Okay. The second speaker from Ampex is John Corcoran. John just recently retired from Ampex, last March, but is now working as a consultant for them.

John spent 30 years in optical and magnetic recording. John will be talking on the archival stability of metal particle tape. John?

ARCHIVAL STABILITY OF METAL PARTICLE TAPE

Mr. John Corcoran
Consultant/Ampex Corporation

MR. CORCORAN: I think most of the other speakers all come from the tape manufacturing community; I actually come from the recording community.

About five years ago, SMPTE, a group of 31 companies, representing broadcasters, tape manufacturers, and equipment manufacturers, started to put together what is now referred to as the D-2 standard.

Actually, it should be called the proposed D-2 standard; it takes forever to get these things done.

The question of tape stability was a key issue then.

(Showing of viewgraphs)

MR. CORCORAN: I happen to have here a paper, and I have copies of it if anybody wants, prepared by a chairman of the committee, Robert D. Thomas of ABC. And one of the key conclusions of the group at that time was that metal particle tape will be satisfactory "under rigorous broadcast conditions and the use of barium ferrite would not be beneficial before at least 1990."

(Change of viewgraph)

MR. CORCORAN: Now, four tape manufacturers report their earliest experimental tapes--and Allan made one, and Darlene has made another; and Sony has separately reported on this, and I believe Hitachi has--tapes that they made 12 or 13 years ago, all of which still play back the original recording signals in an acceptable manner.

Three digital recorder manufacturers in the last year have reported that tapes protected in the cassette will not degrade. Ampex presented a paper on this; Digital Equipment Corporation has presented a paper on this. Dr. Speliotis was a co-author on that paper; he did the magnetic work. And Exabyte has also looked at that.

The effect of the cassette on protecting the tape is much larger than is generally realized. In spite of these facts, the question of stability of tapes is still with us; and I don't know how long it is going to remain with us.

There is considerable evidence; I would like to point out, though, that tapes will survive very well under what I will call archival conditions. Now, archival conditions do not include continuous storage at high temperature and humidity.

There is a considerable number of papers that have been presented before--Bertram, Cuddihy--on what happens if you try to do that. And it won't work with the metal particle tape, and it won't work with iron particle cobalt gammaferric oxide tape either. (Change of viewgraph)

MR. CORCORAN: Here's a chart comparing the uses of various types of tape and various applications over the years. You will see that metal particle tape here has become more common. The most recent one I've added here is a metal particle in a 3480 cartridge.

(Change of viewgraph)

MR. CORCORAN: In our studies, we have taken samples of tape from four different manufacturers, all that were available last year: Ampex, Sony, Fuji, and Hitachi. And we ran them in this Battelle environment, we measured the error rates of those tapes beforehand; we measured them after.

What I have done here is plot the before versus after error rate measurements on three or four cassettes of each of the four brands. Now, this data happens to be for what I call "non-burst errors," where you have isolated single errors. That does not necessarily mean random, but they are isolated. They are not clumps.

AMPEX

ERROR RATE MEASUREMENT

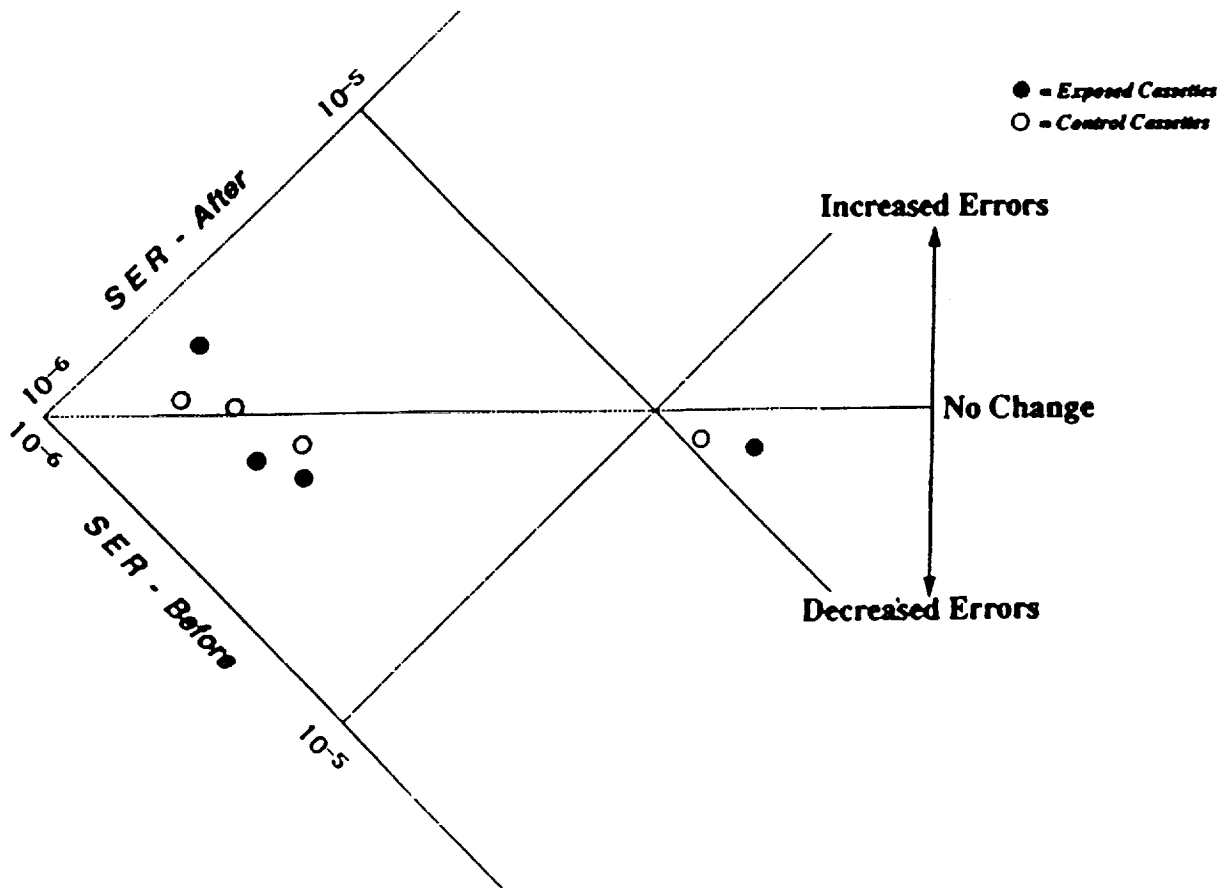
ALL ERROR DATA PRESENTED IS "RAW" TO PERMIT DETECTION OF ANY CHANGE IN THE MEDIA. (WITH ERROR CORRECTION APPLIED, MOST CASSETTES WOULD RUN ERROR-FREE, AND CORROSION EFFECTS WOULD BE UNDETECTABLE).

SYMBOL (OR BYTE) ERROR RATE (SER) IS USED FOR D-2 DATA BECAUSE THE REED-SOLOMON ERROR CODING IS SYMBOL ORIENTED.


BURST AND NON-BURST ERROR DATA ARE PRESENTED. NON-BURST DATA INDICATE ERRORS IN INDIVIDUAL BITS DUE TO LOSS OF SNR, MEASLES, ETC. BURST ERRORS ARE PRIMARILY DUE TO SURFACE DEFECTS.

POST TEST DATA ARE PLOTTED AS A FUNCTION OF PRE-TEST FOR THE SAME CASSETTE. POINTS BELOW THE "NO CHANGE" LINE INDICATE IMPROVEMENT.

Correlation of Non-Burst Error Rates Before and After Exposure Period of D-2 Tapes From Four Manufacturers



Correlation of Burst Error Rates Before and After Exposure Period of D-2 Tapes From Four Manufacturers



The data generally shows the error rates are lower after exposure. The black dots represent the exposed data; the clear dots represent controls. We kept controls with them.

They both go down, indicating probably our transport improved. (Change of viewgraph)

MR. CORCORAN: And we did the same thing with looking not at the isolated errors, but at clumps of errors; and some of us feel that clumps are more likely to be produced by corrosion than single ones. We get exactly the same effect. The tapes did not degrade.

(Change of viewgraph)

MR. CORCORAN: Now, I would like to mention barium ferrite. Barium ferrite is considered to be one of the more stable materials that we can have. Unfortunately, we don't have very much of it.

Ampex had one sample, which we have run on recorders; and we had trouble with shedding. And this is not uncommon when we try a new tape on a recorder. It turns to dust, and then we have to tell the chemist: Go back and fix up the formulation.

We would like to have other samples in the D-2 format. Some day it may prove to be an interesting media; but at the present time, it just isn't available. Thank you.

(Applause)

DR. KRYDER: Okay. Thank you, John. The final speaker this afternoon, before we break into the discussion will be Dennis Speliotis. Dennis was born in Greece and obtained his university education in the United States, receiving a Ph.D. in Solid State Physics in 1961. He worked at IBM from 1961 to 1967 and became Manager of Recording Physics.

Dennis subsequently taught at the University of Minnesota from 1967 to 1971, where he established the Magnetics Research Laboratory. He was a co-founder and vice president of Microbit Corporation, where he worked on electron beam memories.

Since 1977, he has been President of Advanced Development Corporation in Burlington, Massachusetts, doing consulting in advanced magnetic materials and magnetic recording. He co-founded in 1984, and is President of, Digital Measurement Systems, Inc., also located in Burlington, which manufactures advanced vibrating sample and torque magnetometers.

Dennis will talk about corrosion and barium ferrite, and I'm sure he will give us some comparisons to other media.

CORROSION AND BARIUM FERRITE


Dr. Dennis Speliotis
Advanced Development Corporation

DR. SPELIOTIS: Thank you, Mark. I have had two pieces of advice, one from the chairman of our session and one from Dr. Hariharan. One was: Be very brief. And the other one was: Be very nice. And interspersed between those two also was: Be controversial.


This advice is rather conflicting, but I will try to satisfy all the suggestions if I can help it.

(Laughter)


Comparison of a Bit Cell and a Measle



BLOCK DIGRAM OF AN ARCHIVAL STORAGE SYSTEM



Concentration of Agents } Potential
 Flow Rates } Current
 Reaction Rates }
 Time Constants
 Attenuation Factors



WORLDWIDE INDOOR POLLUTANT DISTRIBUTIONS FOR
ELECTRICAL AND ELECTRONIC EQUIPMENT

ADSORPTION IN BATTTELLE TEST CHAMBERS

Gas Adsorption Rate

$$\frac{dQ}{dt} = \{ [G]_{in} - [G]_{out} \} Vf$$


$$= k A [G]_{Avg}$$

$$\frac{2}{V} \sum k_i A_i = \frac{1 \cdot \eta}{f}$$

Adsorption Coefficients

k_i cm/sec

Gas	Chamber Walls	Circuit Boards
Cl ₂	0.20	0.093
H ₂ S	0.086	0.035



GENERAL CORROSION REACTION



RATE FOR SECOND ORDER REACTION (Brackets indicate concentration)

$$\frac{d[M]}{dt} = \frac{d[G]}{dt} = -k [M] [G]$$

For [G] constant (as in Battelle chamber)

FOR FIRST ORDER REACTION

$$\mathbf{M} = \mathbf{M}_0 \exp - t/\tau$$

$$\tau = \frac{1}{k [G]}$$

For low gas concentrations, the time constant is very long. Also, for $t \leq \tau$ the decay rate is constant.

Time Constants for Vaults Without Air Circulation

For Diffusion
in a tunnel of length
 d :

$$\tau = \frac{d^2}{3\Delta}$$

For chlorine in air, $\Delta \approx 0.1 \text{ cm}^2/\text{sec}$ and for $d = 10 \text{ m}$, $\tau = 10^3 \text{ hr}$ or 40 days.

For Thermal Pumping

A daily temperature variation of $3^\circ/\text{k}$ will cause air to flow in and out of any unsealed container or space, thereby carrying corrosion agents into the container. Using perfect gas law, the daily mass flow is

$$\frac{1}{M_c} \frac{\Delta M}{\Delta t} = \frac{\Delta T}{\bar{T} \Delta T}$$

where M_c is the mass of air in the container and \bar{T} is the average absolute temperature. The time constant is the reciprocal of these terms, about 1200 hours, 50 days.

For Barometric Pumping

Similarly

$$\tau = \frac{\bar{P} \Delta t}{\Delta P}$$


For $\bar{P} = 30 \text{ in. Hg}$, $\Delta P = 1 \text{ in}$ $\Delta t = 3 \text{ days}$, $\tau = 90 \text{ days}$.

SHIELDING TEST


STANDARD D-2 CASSETTES (SMALL SIZE) WERE FITTED WITH BATTELLE METAL COUPON SAMPLERS AND EXPOSED FOR 28 DAYS. A COMPARISON OF RESULTS OF THREE INTERNAL SAMPLERS WITH AN EXPOSED ONE INDICATES THE CASSETTE HAS ATTENUATED THE **CORROSIVE** ENVIRONMENT TO NEGLIGIBLE PROPORTIONS.

THE DIFFERENCE BETWEEN THE PRESENT RESULTS AND THOSE REPORTED BY SPELIOTIS FOR "NAKED" TAPE ARE ATTRIBUTABLE TO THE CASSETTE SHIELDING FACTOR.

Sampler Locations on D-2 Cassette




Silver Sampler Film Measurements for Various Locations on D-2 Cassette



CASSETTE IMPEDANCE MODEL

Diffusion Resistance $R_D = \frac{\ell}{\Delta A_x}$



Wall Adsorption Resistance $\frac{1}{K_w A_w}$

$$\frac{[C]}{[G]} = \frac{1}{\frac{K_w A_w \ell}{\Delta A_x} + 1} \text{ Attenuation Factor}$$

DIFFUSION INTO REELS




Fig. 6.10 Moisture Gradients During Conditioning of Aerial Film Rolls at 70°F -5% RH

Data from "Manual of Physical Properties"
- Eastman Kodak Co.

DETERMINATION OF TIME CONSTANTS FOR MOISTURE DIFFUSION




Fig. 6.11 Difference in relative humidity between center to edge of a 70 mm reel following a step change from 70°F - 50% RH

$$\tau = T \left\{ \frac{W}{70 \text{ mm}} \right\}^2$$

FROM BERTRAM & ESHEL

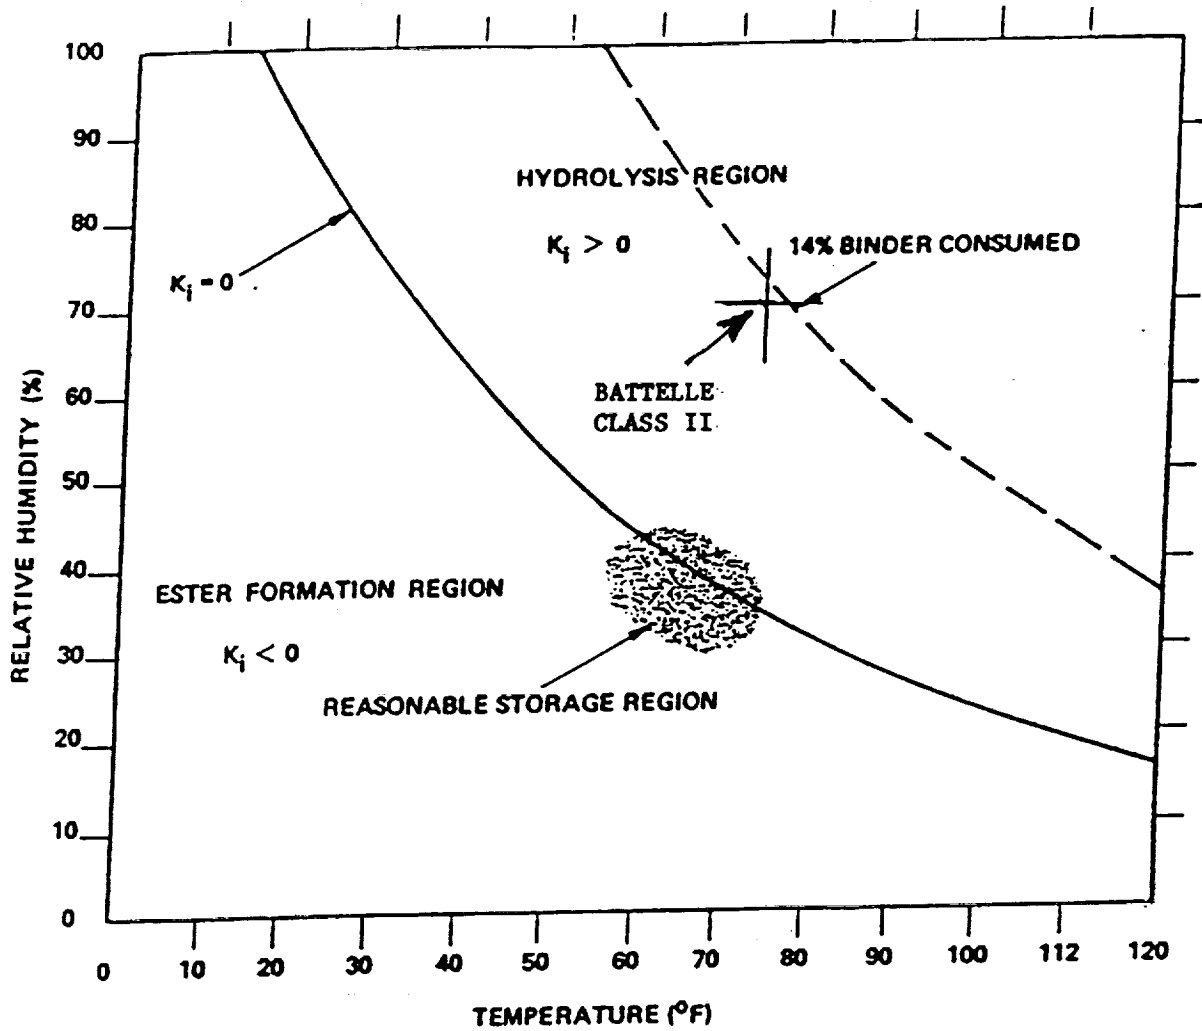


Figure 37 Hydrolysis equilibrium curves vs humidity and temperature. $K_i = 0$ indicates initial equilibrium. Dashed curve denotes 14% binder hydrolysed from an initial 6.7%.

Conclusions

1. The archival life of four brands of D-2 tape have been shown to exceed 14 years in the Battelle Class II environment. No evidence of corrosion was found.
2. The cassette is a necessary element for achieving this life. Dangling tape out of a cassette invites failure.
3. Extended exposure of any type of tape to high temperature and humidity causes binder degradation and coating failure. Archival storage is not possible in such environments.
4. An archival storage system includes: the tape, its cassette, other protective enclosures (if used), the storage vault, and its material parameters and the environment. Methods of determining time constants and attenuation factors for estimating storage life have been suggested.
5. There is a need for research -- perhaps here at the University of Alabama -- to determine adsorption coefficients of tapes, cassettes and other materials used in archival storage to use in life predictions.

REFERENCES

1. Abbott, William G. - Development and Performance Characteristics of Mixed Flowing Gas Test Environment - IEEE Trans on Components Etc., VII, p. 22 Mar. 1988.
2. Cuddihy, Edward F. - Aging Correlation (RH + T) Relative Humidity (%) + Temperature (C) - Corrosion Science, V27, No. 5, p 463-474, 1989.
3. Cuddihy, Edward F. - Aging of Magnetic Recording Tape, IEEE Trans on Mag V16, No. 4, p 558-568, July 1980.
4. Bertram, Neal and Eskel, A. - Recording Media Archival Attributes (Magnetic) Contract No. F 30602-78-C-0181, Final Technical Report to Rome Air Development Center, Ampex Report RR-80-01, 30 Nov. 1979.
5. Speliotis, Dennis E., - Corrosion of Particulate and Thin Film Media, IEEE Trans Mag V26, p 124-6, Jan. 1990.
6. Wolf, Irving - Environmental Stability of D-2 Tape, THIC, Mar. 1990.
7. Thomas, Robert G. - SMPTE Study Group on New Magnetic Media, SMPTE Journal, p 1242, Dec. 1986.
8. Hatai, A. et al. - Archival Stability of Metal Video Tape/Betacam SP, Symposium on Archival Stability, Hull, Canada, May 1990.
9. Yamamoto, Y. et al. - Study of Corrosion Stability in Metal Particle Tape, IEEE Trans Mag V26, p 098-100, 1990.
10. Mukaida, Y., Practical Stability of Metal Powder Magnetic Tape, THIC, June 2, 1990
11. Metal Tape Stability, Sony, THIC, June 19, 1990...
12. Djalali, A.; Judge, J. S.; Speliotis, D., et al. Study of the Stability of Metal Particle Data Recording Tapes - Fall Meeting of Electrochemical Society, Seattle, Oct. 14, 1990.

DR. SPELIOTIS: Well, we already heard about various kinds of tapes and their stability. There is one important issue, however, that I would like to remind everybody; and that is, we didn't hear very much about gammaferric oxide *per se*. Why? Primarily because the performance of γ -ferric oxide is not adequate to meet the demands of high density data storage.

So, really stability and archivability are important, but performance is also important. We are talking about high performance media. We talk about terabytes of data; some people even talk about petabytes of data, 10^{15} .

So, obviously, high density is very important. What we would really like to have is very stable media that will also have high performance.

If we keep this in mind, I submit to you that most of today's highest usage media, such as Co-doped γ -Fe₂O₃ and chromium dioxide are not extendable to very high densities. Therefore, they will fall by the wayside in the future--no doubt about it.

Co-doped γ -Fe₂O₃ is not extendable because if we add more Cobalt, the temperature dependence of the coercivity becomes excessive. This is a reversible change--not irreversible. It is just a reversible temperature change; coercivity drops as the temperature goes up. Chromium dioxide is not extendable because we cannot go up in coercivity very much. There is only one way that has been found to extend the coercivity of chromium dioxide to very high values, which is required for high density storage, and that is to dope it with Iron; but the total amount of Iron in the world would not be enough to satisfy even a minuscule percent of the requirement.

Iron doping of chromium dioxide is another possibility, but that is limited. Maybe the maximum coercivity that can be obtained is about 1,000 Oersted, which is not good enough.

So, I am going to scratch off the list for the future chromium dioxide and cobalt-modified iron oxide. That leaves metal particles and ME. Nobody has talked about ME, and I will say a few words. Metal-evaporated tape is a very natural extension of the technology to go to high density data storage. Why? Because we can get very high coercivity in metal particles and very high magnetization at the same time, which is beautiful.

On ME, of course, we get high coercivity, very high magnetization, and thinness, which is also very important to go to high density.

So, it is only logical that the Japanese and us, later adopted the use of metal particle and are beginning to adopt ME for high density storage media.

Now, there is a new material; John said it is not available, but it is becoming available: barium ferrite. Barium ferrite is a very stable material; Darlene did not mention it, but on her chart of stability, there are only two materials that are really stable. And I don't think anybody argues with this.

Plain γ -Fe₂O₃ and barium ferrite. On Darlene's chart, those were the only ones that had red dots, meaning stable under all conditions--no change. Okay.

Now, barium ferrite also is a material which allows you to attain any coercivity you want--any-- including higher coercivities than metal particles. If you want 5,000 Oersteds, if you can write on it, you can have it.

Now, there is one problem with barium ferrite; and that is that the magnetization is relatively low, about 60 EMU per gram. However, as you saw, on many of the metal particles, the magnetization is dropping and is dropping very rapidly. Pure alpha iron has a magnetization of 210 EMU per gram; today's metal particle tapes use particles which because of the oxide required to protect them, plus the silicon fortification--what is called armor

fortification--silicon alumina, ceramic fortifier- the effective magnetization with the particles is about half of that of a-Fe.

Now, with these particles that are used today in metal particle media, the size of the particles typically is about 1,500 Angstroms in length and about 1/8th to 1/10th of that in diameter.

They have a protective coating roughly on the order of 3 nanometers, about 30 Angstroms.

Now, let us make a simple calculation. If you take this particle and make it smaller, and you must make it smaller because you cannot extend the high-density capacity of metal particles without using smaller particles (the particles are already too large as they are today for real high density). If you make it smaller, like 1,000 angstroms in length, and about 100 angstroms in diameter and have roughly not 30 angstroms of protective coating, but maybe 25, you calculate the magnetization of that particle to be below that of barium ferrite.

So, metal particles are not extendable into the real high-density requirements of the future. I can take the best metal particle tape that is available today, and compare it with barium ferrite experimental tapes; the barium ferrite, in spite of its low magnetization, outperforms the metal particle tape in absolute signal--not in noise--it outperforms it in noise at much lower density--but in absolute signal, for any density over 75 Kfci.

At densities above 75 Kfci the barium ferrite tape has higher absolute signal than metal particle tapes. The reasons for this are beyond the scope of this conference; there are certain fundamental demagnetization processes that are very different in the case of barium ferrite than in the case of metal particles.

Now, many of us have also heard--and I think I should say a few words--about ME tape--the metal-evaporated tape that in some quarters is supposedly going to be the ultimate successor of all of these tapes. It is going to be the high-density, high-definition media of the future. And the question is: How, for example, does barium ferrite tape compare with ME tape recording-wise? The performance is identical within half a dB at anything above 130 Kfci.

So, if you really want to talk about high density, I think the real possible contenders are barium ferrite and ME tape. MP tape is not extendable.

Now, let's go to the subject of our seminar here, which is called instability. There are some very important issues; Allan made a very important distinction between long-term archivability versus environmental stability. It is obviously a complex issue.

What I will show you are some error rates that were measured using a Media Logic ML-4500 tape tester, which uses an Exabyte drive. The measurements were done on 8 millimeter metal particle, ME, and some experimental barium ferrite tapes, all of them 8 millimeter. We did immediately read after write.

These tapes were exposed to a corrosion test which was only temperature/humidity; 7.5 weeks, at 50°C/80% relative humidity.

When I put the tapes in the chamber, they did not have the protective plastic box; I just put the cassette in, the way you mount it on a recorder, simply because I didn't want to be limited by diffusion effects, if any, in that period of time.

(Showing of viewgraphs)

DR. SPELIOTIS: Here are some typical results and error maps, before and after corrosion, for 40,000 tracks, on the Exabyte drive. These are not single bit errors; the machine

Corrosion Test: 7.5 Weeks 50C, 80% RH Dropout vs. Threshold (20/28/80)

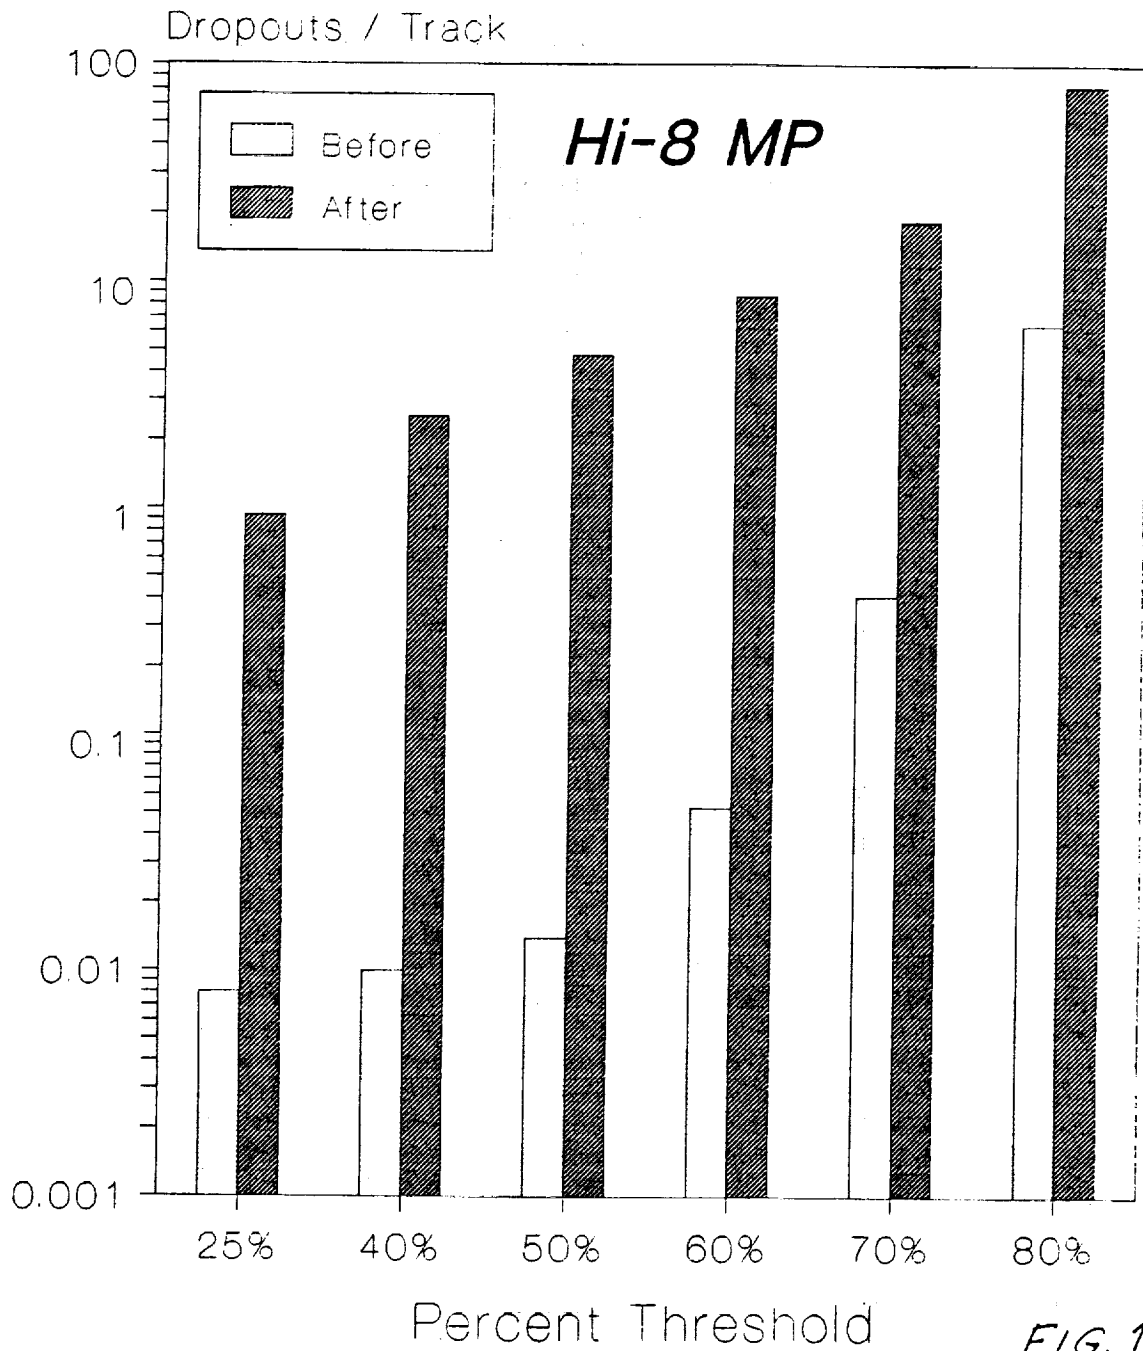


FIG. 1
Hi 8 MP Tape
before and after
corrosion

Corrosion Test: 7.5 Weeks 50C, 80% RH
Dropout Size Distribution (75% TH., 28G)

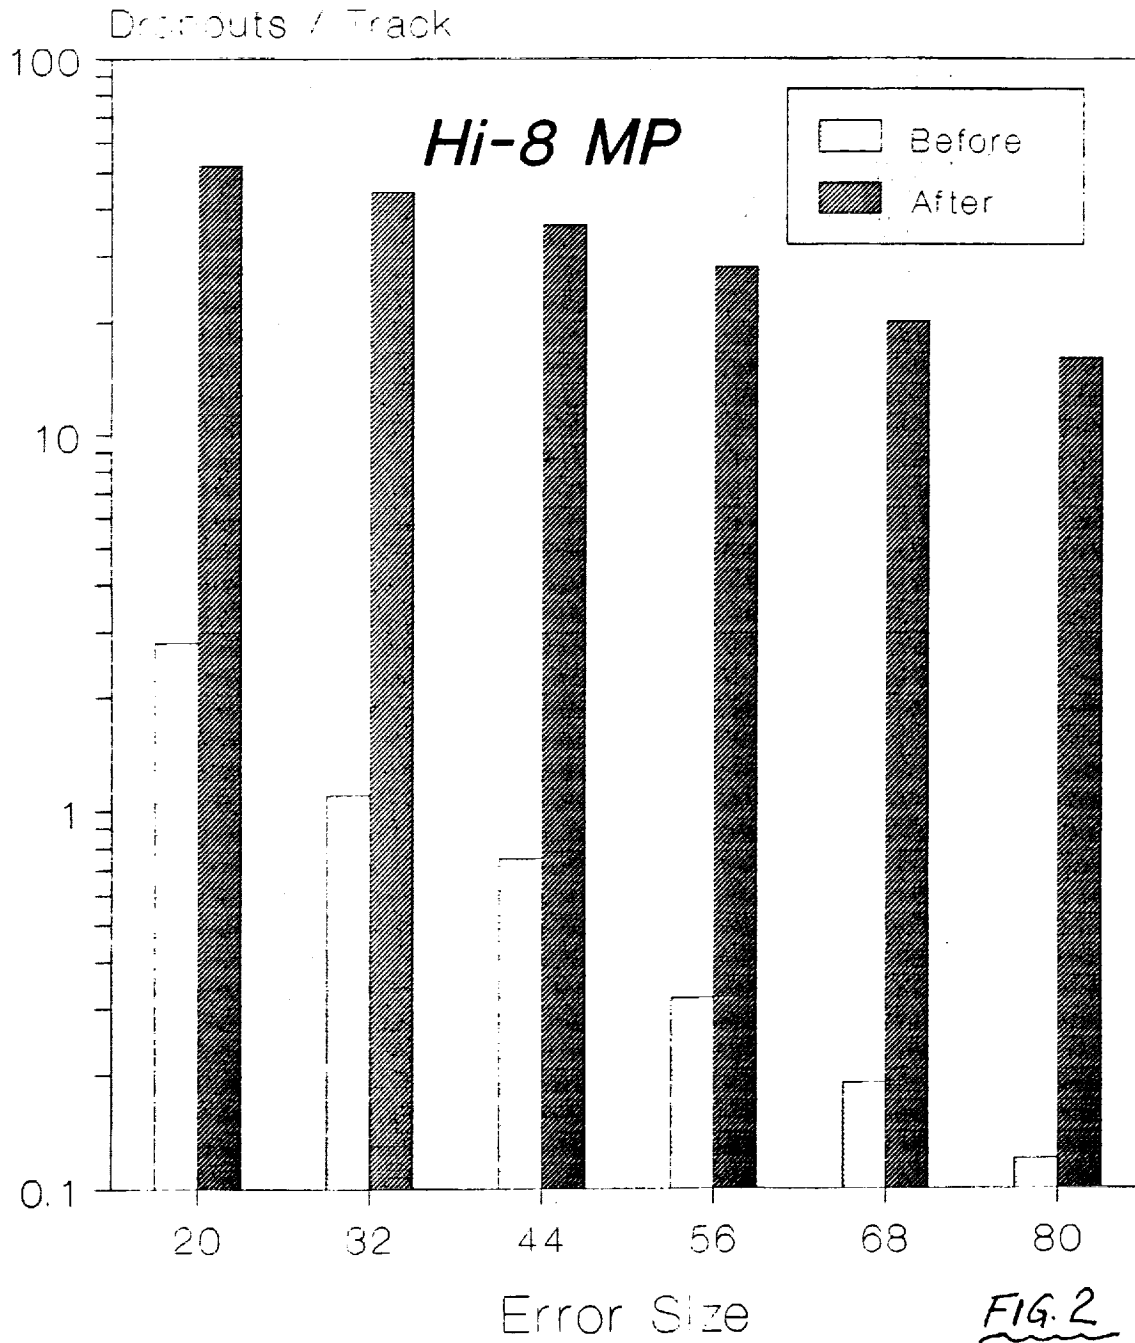


FIG. 2
Hi8MP Tape
before and after
Corrosion

Corrosion Test: 7.5 Weeks 50C, 80% RH Dropout vs. Threshold (20/28/80)

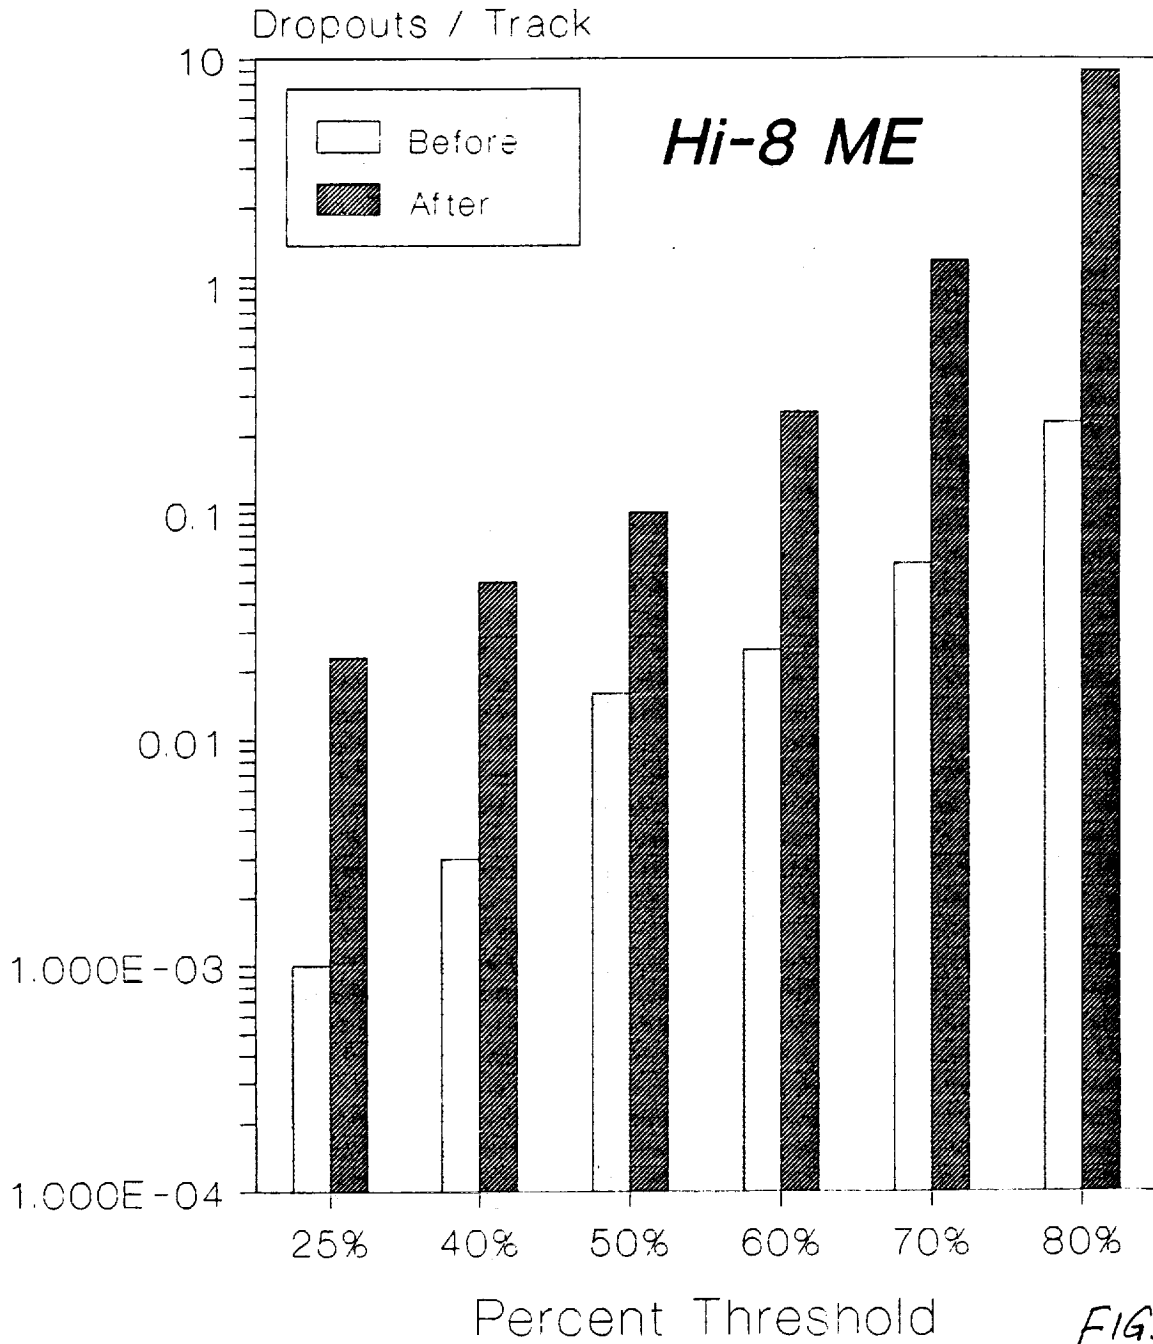


FIG. 3
Hi8ME Tape
before and after
corrosion

Corrosion Test: 7.5 Weeks 50C, 80% RH Dropout Size Distribution (75% TH., 28G)




FIG. 4
Hi8ME Tape
before and after
Corrosion

can certainly count single bit errors, but they are larger errors because the data correction in the Exabyte drive can handle blocks of errors.

In any case, there was a very large increase in error rates in the MP tapes.

(Change of viewgraph)

DR. SPELIOTIS: ME tape has a much higher signal- to-noise ratio than MP; and when you look at it before corrosion, it is very clean. So, there is nothing to show before; there are hardly any errors.

But there is a lot of difference between different ME tapes; and I'll show two of them, both of them after corrosion, for manufacturer A and manufacturer B. There are relatively a lot of errors, a much larger increase percent-wise than on MP tape.

And there are many, many errors at the front part of the tape on which I had done a lot of testing for short blocks; and the rest of the tape, I hadn't looked at, but there was some corrosion.

There is clearly some diffusion limitation; there is edge corrosion. This is one edge of the tape; this is the other edge. And this is down track. And there is some significant edge corrosion limited by diffusion apparently and a tremendous difference between manufacturers.

(Change of viewgraph)

DR. SPELIOTIS: Of course, the same thing applies to MP tape. There is a very big difference between the unfortified armor protected particles and the plain oxide protected particles.

(Change of viewgraph)

DR. SPELIOTIS: Here is a metal particle tape, again with excessive testing on the front: once every day for 20 minutes and then it was put back into its cassette enclosure (this was not exposed to an accelerated corrosion test but was in my office for one month).

The tape developed a pretty large number of errors in one month. In the beginning, this was very clean; when I opened the cassette from its plastic sealed case, it was very clean. It was just sitting in the office which is typical of most offices.

DR. SPELIOTIS: Here is a barium ferrite tape measured similarly to the MP tape--no effect. So, magnetically and error-wise--and of course, the extent of the errors is the real criterion, I think-- there is a very significant difference for these tapes, even in a normal office environment, without doing any accelerated corrosion testing.

Now, to summarize: Is there a disagreement between what the other speakers said and what I have said? I don't believe there is any disagreement at all. Let me state the view as I perceive it including the tests that John Corcoran referred to: the published data, a very extensive study by many people over many months--that was done on behalf of Digital Equipment Corporation.

This is going to be published this August in the *Journal of the Electrochemical Society*,* and it is a study of the stability of MP tapes.

* *J Electrochem. Soc.*, **138** (1991) 2504-2509

Corrosion Test: 7.5 Weeks 50C, 80% RH Dropout vs. Threshold (20/28/80)




FIG. 5
Two different
Hi8 MP Tapes
after corrosion

Corrosion Test: 7.5 Weeks 50C, 80% RH
Dropout Size Distribution (75% TH., 28G)




FIG. 6
Two Different
Hi8MP Tapes
after corrosion

Corrosion Test: 7.5 Weeks 50C, 80% RH Dropout vs. Threshold (20/28/80)

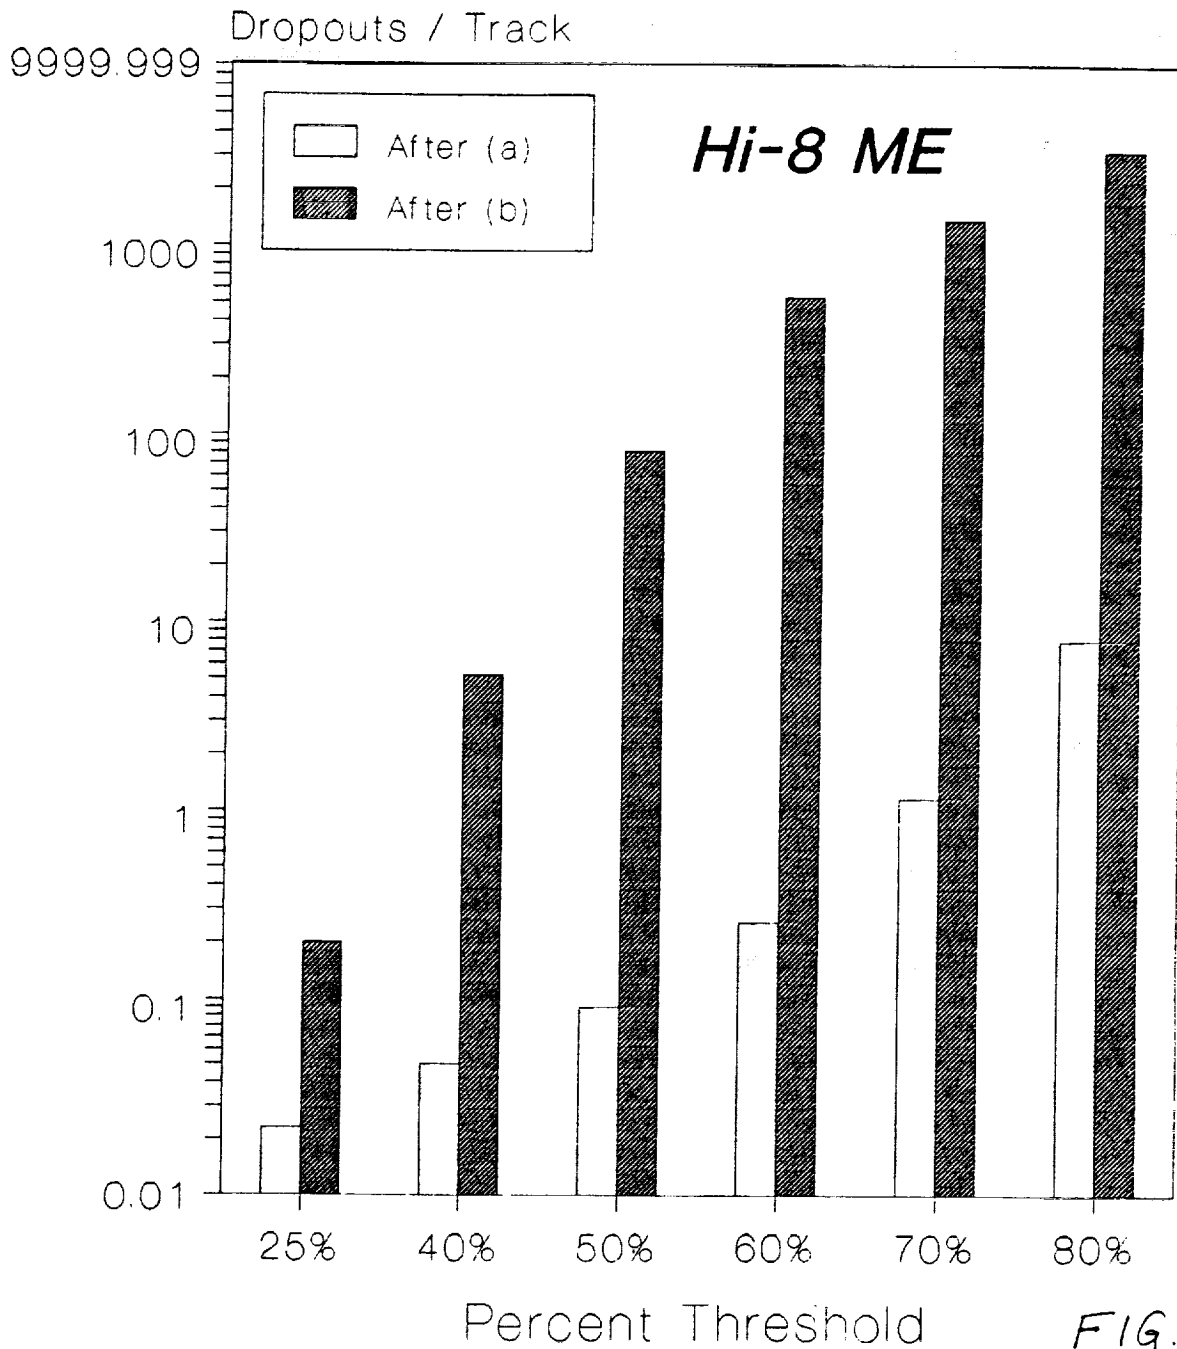
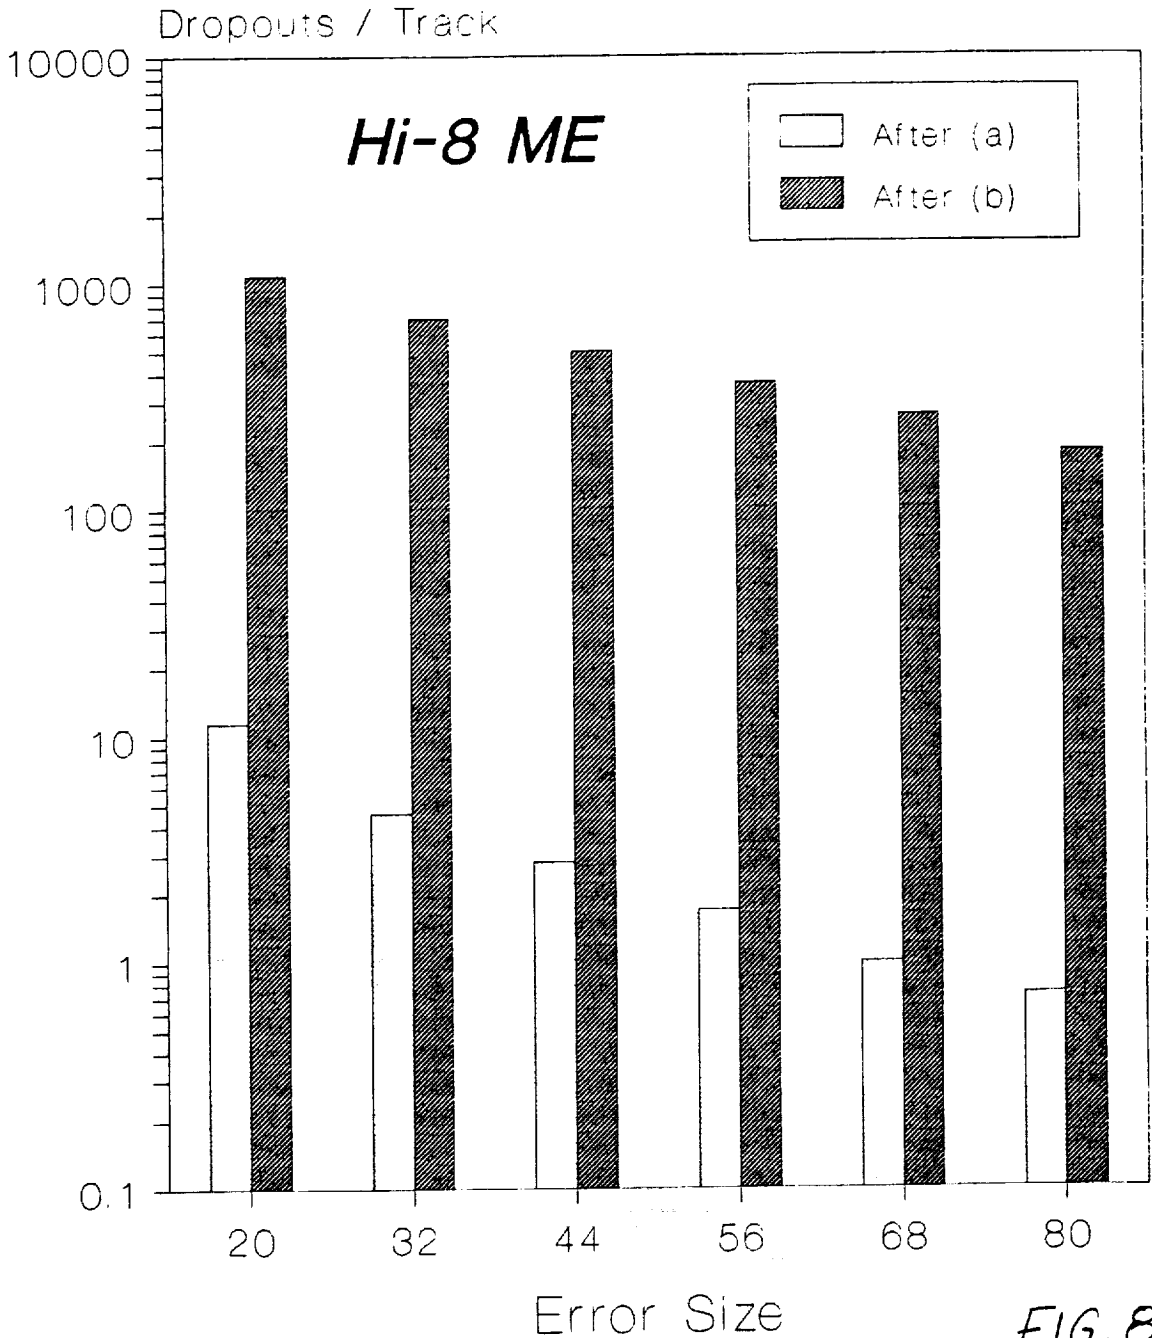


FIG. 7
Two different
Hi8ME Tapes
after Corrosion


**Corrosion Test: 7.5 Weeks 50C, 80% RH
Dropout Size Distribution (75% TH., 28G)**



*FIG. 8
Two different
Hi8ME Tapes
after Corrosion*

DROPOUT MAP

Unit: 1 8MM EXABYTE ROTARY HEAD 10:11:52 12/11/90
Operator: KP Lot: corrosion test Cartridge:
Current: 15.22 mA Frequency: 4.0000 MHz Location: 5.00%
Threshold: 25.0 Bad/Good/Max: 8/28/80 Tracks: 40000 8mm MP




Before

FIG. 9
Error maps of
8mm MP Tape
before and after
Corrosion

DROPOUT MAP

Unit: 3 8MM EXABYTE ROTARY HEAD 12:45:53 02/19/91
Operator: KP Lot: 7.5 WKS TH Cartridge:
Current: 14.28 mA Frequency: 4.0000 MHz Location: 5.00%
Threshold: 25.0 Bad/Good/Max: 8/28/80 Tracks: 40000 8mm MP




After

FIG. 10
Error maps of
two different
Hi8 ME tapes
after conversion

DROPOUT MAP

Unit: 1 8MM EXABYTE ROTARY HEAD 08:07:45
Operator: DES Lot: 7.5 URS TH Cartridge:
Current: 13.40 mA Frequency: 4.0000 MHz Location: 0.50%
Threshold: 75.0 Bad/Good/Max: 20/28/80 Tracks: 40000 **HI-8 ME**




After (a)

0 Total Errors: 58540 Track Number 40000
(c) 1990 MediaLogic, Inc.

DROPOUT MAP


Unit: 1 8MM EXABYTE ROTARY HEAD 11:05:25 **HI-8 ME**
Operator: DES Lot: 7.5 URS TH Cartridge:
Current: 11.93 mA Frequency: 4.0000 MHz Location: 0.50%
Threshold: 75.0 Bad/Good/Max: 20/28/80 Tracks: 40000



After (b)

0 Total Errors: 388968 Track Number 40000
(c) 1990 MediaLogic, Inc.


Operator: DES Lot: Hi8MP K-C Cartridge: 4H50097
 Current: 14.00 mA Frequency: 4.0000 MHz Location: 0.50%
 Threshold: 25.0 Bad/Good/Max: 20/28/80 Tracks: 40000



Total Errors: 107862 Track Number 40000 (c) 1998 MediaLogic, Inc.

DROPOUT MAP

Unit: 1 9MM EXABYTE ROTARY HEAD Q1:14:44
 Operator: DES Lot: EF/I/MAY91 Cartridge: BT1C-02
 Current: 14.00 mA Frequency: 4.0000 MHz Location: 0.50%
 Threshold: 25.0 Bad/Good/Max: 20/28/80 Tracks: 40000




Total Errors: 2871 Track Number 40000 (c) 1998 MediaLogic, Inc.

FIG. 11
 Error maps of
 Hi8MP and Bat
 Tapes after 1 ms
 testing.

DROPOUT MAP


Unit: 1 8MM EXABYTE ROTARY HEAD 16:58:41 06/05/91
Operator: DES Lot: HI8MP K-G Cartridge: 4H30897
Current: 14.00 mA Frequency: 4.0000 MHz Location: 0.50%
Threshold: 50.0 Bad/Good/Max: 20/28/80 Tracks: 40000



Total Errors: 236298 Track Number (c) 1990 MediaLogic, Inc.

DROPOUT MAP

Unit: 1 8MM EXABYTE ROTARY HEAD 16:15:47 06/05/91
Operator: DES Lot: BF/T/MAY91 Cartridge: BTIC-02
Current: 14.00 mA Frequency: 4.0000 MHz Location: 0.50%
Threshold: 50.0 Bad/Good/Max: 20/28/80 Tracks: 40000



Total Errors: 3154 Track Number (c) 1990 MediaLogic, Inc.

FIG. 12
Error maps of
Hi8MP and BaFe
Tapes after 1 mon
testing.

What we found was: If you take these tapes inside the cassettes and expose them to a Battelle environment, which is a very moderate temperature/humidity but with polluting gases present, there is some increase in errors, but no problem at all.

Our opinion is that the cassette protects the tapes primarily because there is not enough chlorine, including diffusion, but even if there were diffusion, there is not enough chlorine to really get in and do the damage. Now, if we expose prerecorded tapes to a Battelle test outside that protected case, (open tape with data) all the tapes, including ceramic protected, were destroyed; after seven days they were unreadable.

And then, after rewriting them, they were unreadable--all gone!

Now, the same test on prerecorded tapes inside their cassettes, exposed to a temperature/humidity test for anywhere from 10 days to 70 days. We found that after 30 days or 40 days, there was some increase in error rates, but not a very big problem.

After 70 days, most of the tapes had data that was unrecoverable, on a temperature/humidity test, inside the cassettes.

We interpret this to mean that there is plenty of humidity around and is doing a job on these tapes. On the Battelle test, the humidity is very low--relatively--and the temperature is moderate; and there is not enough chlorine, which is presumably the primary culprit to damage the tapes.

I don't think these findings are in disagreement with what other people presented; but if there are some opposite opinions, they should come out at the following discussion.

One very worrisome thing about the metal particle tapes that were studied at DEC was that on the Battelle test, we kept taking the tapes out of the chamber every couple of days to determine what was the change in error rates.

This particular Battelle test was a 12-day test, and we took the prerecorded tapes and measured the change in error rates every two days. Normally what you would expect to happen due to corrosion would be that you would have initially some rapid corrosion because of surface effects, and then, it would level off because the corroded surface would become an inhibitor for further corrosion. What we found was exactly the opposite.

For several days of exposure, corresponding to years of projected life, not much change; and then an exponential, dramatic increase in error rates near the end of the test. This really needs to be tried over again and understood.

If it is true, it could mean that we are sitting on a time bomb. It would have these tapes looking fairly stable and fairly good; but after some time, they would fall apart. Thank you.

(Applause)

DR. KRYDER: Could I get all the panel members to come forward, please? I will ask them to sit in the order that they spoke, but leave some extra space for Mike Younker. Mike is an interpreter here for Sony.

(Pause)

DR. KRYDER: I invite questions from the audience to get us started. Go ahead.

PARTICIPANT: Yes. 18 months from now, I will have to start shipping a system with high density tape. What tape should I plan to use? (Inaudible)

DR. SPELIOTIS: You know my answer.

DR. KRYDER: Well, what is your answer? He has got to ship it in 18 months.

DR. SPELIOTIS: 18 months? I think it will. Barium ferrite tape will be available in adequate supply definitely in 18 months.

PARTICIPANT: But it has to be tested.

DR. SPELIOTIS: Right. You have to test it; so, it has to be available much earlier than that.

PARTICIPANT: There are now five companies that have D-2 tape. You will now have your option to look at various sources. I don't think there is any hope that that is going to be available -- (Inaudible)

DR. KRYDER: Could I ask the people to pass the microphone around for those who are speaking?

PARTICIPANT: 8 millimeter, 4 millimeter -- (Inaudible)

(Inaudible discussion)

DR. KRYDER: John, can you repeat the question?

PARTICIPANT: (Inaudible)

MR. CORCORAN: For instance, the other questions like: Do you want it all in one cassette? Do you want many cassettes to a lot of different people? What data rate do you want?

PARTICIPANT: (Inaudible)

MR. CORCORAN: Okay. Now, you can do that in four small D-2's or two big D-2's. If you do it in Exabyte, they are about 2 or 5 gigabytes, depending on which one. So, you could do that with 20.

PARTICIPANT: (Inaudible)

MR. CORCORAN: Right. Now, what data rates do you want them to read at?

PARTICIPANT: (Inaudible)

MR. CORCORAN: I think if you will use the metal particle tape, under the "normal commercial practice" and you don't expose it to high humidity for an extended period of time, you will not have any trouble with it.

DR. SPELIOTIS: But you cannot guarantee this for shipment.

MR. CORCORAN: The tapes that we used and I showed data on were shipped back and forth, from Redwood City to Battelle and back.

PARTICIPANT: With the 85 degree temperatures, I could easily achieve those -- (Inaudible) -- if they were sitting out on an aircraft tarmac --

MR. CORCORAN: Well, I'd be a little surprised if you got 85° Celsius on a tarmac. I know it will get up to 55 or 60, but not 85.

DR. REAGOR: The tentative recommendation we are having with regard to chromium dioxide and metal particle is that you are in various zones of safety. You are very safe if you are between 30 and 50% RH and less than 50° C.

If you go beyond that, you are going into a zone of more risk. If you get up to 70-80% RH for any extended period of time, you are taking a risk.

Now, we are talking about several different risks, not only pigment, but we are also talking about changes in the hydrolysis of the binder and also in the stability of the polyester.

So, again, what we are talking about is the degree of risk that you want to take as you are implementing these types of systems.

DR. KRYDER: A question in the back?

PARTICIPANT: (Inaudible)

MR. HADAD: What kind of problems?

PARTICIPANT: (Inaudible)

MS. CARLSON: There are probably two different mechanisms that could occur; and therefore, it is very difficult for us to make a recommendation without knowing what the mechanism of failure is. But there are probably two different mechanisms.

One: Remember the slide I talked to you about -- If it was sitting for a long period of time in a hot and humid environment, you could get outside of the track, could get outside of the format that the recorder could play back. Or else, you could have signal degradation.

Now, unless we see what you actually have, it is very hard to know what implementation procedures we would recommend; but we can talk about it. There are two obvious mechanisms that can be going on.

MR. HADAD: Yes. I would say it would be more of a physical phenomenon, the -- changing, which is common to all tapes, as opposed to a magnetic change. I showed the data, of course, where we get the magnetics; and in periods of up to six months in those high temperature/humidity conditions, there would only be a loss in magnetization of 1 or 2 percent.

PARTICIPANT: (Inaudible)

MR. HADAD: But again, you said some of those tapes had been shipped from some long distance; and that was the result of receiving it from Germany?

PARTICIPANT: Yes. (Inaudible)

MR. HADAD: Okay.

DR. KRYDER: Another question?

PARTICIPANT: (Inaudible)

(Inaudible discussion)

DR. KRYDER: Speak up, please.

PARTICIPANT: (Inaudible)

MR. HADAD: It would be awfully difficult to monitor. Foreign matter getting on them is very randomized; and I think it would be very difficult to model that. So, the idea would be just to keep the tapes in as clean an environment as possible, although we recognize that --

MR. CORCORAN: It would show an error rate like I showed. If you are getting additional dirt, the error rate is going to degrade.

PARTICIPANT: I would think it would be a higher random error rate.

MR. CORCORAN: Yes.

PARTICIPANT: (Inaudible)

MR. CORCORAN: We did not see it. Okay? Now, there is undoubtedly some condition you might find it in, but we did not find it in the Battelle test.

DR. HADAD: That would be independent of format, of course.

DR. REAGOR: Let me go into that. From a user's standpoint in tape storage, we do the analysis of the particulate impact; and we can actually predict particulate deposition rates onto surfaces.

From the models we have used and from actual examination of the tape storage areas of particular environment, particulate contamination contributing to future degradation is not an issue.

The shielding from the containers is so high that the actual deposition velocity, since there is no movement in the cartridge, there is no impetus to pull particles from the environment in, as their natural grounding movement brings a very small amount just to the openings that are exposed.

The biggest concern we have had with the magnetic tapes, in particular the 3480s, are localized contaminants from the people that operate the systems.

And we have been able to show that, even in the cleanest of a data center, the localized input of a person walking up and removing a cartridge and putting the next one in, it is the dirt on that person that leads to the problem and the failure. And we can actually tell you the color of their hair and who used it last and things of that nature.

(Laughter)

DR. REAGOR: And that is the issue that we have started to address in terms of the air flow dynamics in the data center around a tape reader and the fact that you have to change it in order not to strip off the dirt of a person--and these are things we are not actually looking at--and actually have it come up in the air and flow inside the tape reader at the point of change.

DR. KRYDER: Another question?

PARTICIPANT: (Inaudible)

MS. CARLSON: I think that the point that Dennis made earlier was that for systems that used 2000 fci or 1600 fci--you know, nine-track tape--you should stay with nine-track tape because that is a very stable environment. However, you cannot get up to 45 Kfci or 100 Kfci; you just cannot get the performance from the 9-track tape system

PARTICIPANT: (Inaudible)

DR. KRYDER: Let me just briefly summarize that. I was asked to repeat the question. The second one was: Is there an effect of horizontal versus vertical storage basically on the tape? Go ahead, John.

MR. CORCORAN: I think in general the rule still is that you should store the reels with the axis horizontal. Now, we tend to have smaller cassettes today; I think perhaps the problem is less intense than when we had big 12 and 14 inch reels. But you will still do better if you keep them horizontal.

PARTICIPANT: Why? (Inaudible)

MR. CORCORAN: You don't have slippage of the reel packs.

PARTICIPANT: (Inaudible)

MR. CORCORAN: The smaller cassettes do have less tape on them. If the ratio of O.D. to I.D. exceeds 3, cinching can occur regardless of how large or small the tape is.

PARTICIPANT: (Inaudible)

MS. CARLSON: First of all, I would always recommend rewinding to the beginning of tape or to the end of the tape, specifically because polyester does deform; and if you have it exposed, you are going to get crease marks or deformations on any of the bending radii that you are storing it on.

So, in archives, that is one of the things that was mentioned early today. Some of the systems unload, but don't rewind the tapes. That's fine as long as you are not having those tapes stay in that condition for more than three months.

After three months, you should really have it stored at beginning of tape or end of tape because you will lose data at those crease marks.

DR. KRYDER: What was the first question? I'm sorry; I missed it.

PARTICIPANT: (Inaudible)

DR. KRYDER: Don't you really mean: How realistic are their tests and the extrapolations from those?

MR. HADAD: Personally, I feel the temperature/ humidity are the major issues. And in terms of smoke or ozone, no, we have not addressed those. Looking at it just as a background level, leaving it constant, and then we are more concerned in our own modeling. And again, those models that I showed were only the particles; we haven't gotten to the tape yet. That's proposed.

But all of that other was sort of background; and we felt the temperature/humidity was the most critical.

DR. KRYDER: Dennis, do you agree with that?

DR. SPELIOTIS: Yes, I think temperature/humidity is the most damaging. I believe the gases are not so severe; there are not so much of them around. For metal particles, there is a lot of iron in a D-2 cassette, a lot of iron and there is not enough chlorine to affect it.

PARTICIPANT: (Inaudible)

MS. CARLSON: It's 45. That's cobalt-doped; that's not iron oxide. (Inaudible discussion)

MR. HADAD: It's a cobalt modified --

MS. CARLSON: Yes.

MR. HADAD: Cobalt modified γ -Fe₂O₃.

(Inaudible discussion)

MS. CARLSON: Yes. Co-doped γ -Fe₂O₃ is a fairly stable material. The only problem that you have is at high temperature and humidities, the cobalt will diffuse into the γ -Fe₂O₃; and therefore, you lose coercivity.

And in a normal controlled environment, that is very good. However, the point that Dennis made is that at about 50 KFCI, cobalt-doped is as high as you are going to be able to get with cobalt-doped material.

PARTICIPANT: (Inaudible)

DR. REAGOR: From our standpoint in the Bell system--the Bell Companies--we went from the open reel iron oxide tapes to the chromium dioxide tapes, not because of looking for better quality materials; it reduced the size of our storage by about 80 percent. We could put the same amount of information, and now that bulk to protect our things for archiving is a lot smaller. And so, when we went to that mode, it was not because--and I think you have been hearing this--it is not done because the material may be that much better; it is because it is giving the end user a better storage capacity per inch.

And depending on what that end user needs and what those requirements are for archiving and how much data they project they are going to have, that may govern what you want and what you are going to use in the future.

But in terms of my area, it was strictly space. We were running out of space to keep the tapes; and you know, it costs money to maintain that space.

PARTICIPANT: (Inaudible)

DR. KRYDER: Another question?

MR. CROSBY: I'm Jim Crosby, and I would like to make a comment about -- (Inaudible)

DR. KRYDER: That's cobalt damage.

MR. CROSBY: As a matter of fact, it has a coercivity of 600 -- (Inaudible)

DR. KRYDER: No, it's cobalt.

DR. REAGOR: Okay.

DR. KRYDER: All right. We had a question over here and then another.

PARTICIPANT: (Inaudible)

DR. KRYDER: Let me briefly repeat the question. Let me paraphrase the question: He's pointing out that it is not unusual in some parts of the world, particularly in Washington, D.C. for instance yesterday -- (Laughter)

DR. KRYDER: To sit at 85 or 90% relative humidity for days on end. And how do these tests relate to that sort of an environment?

MR. HADAD: The data has shown the lowest temperature on that was 50° C; that is still considerably above that, but we don't have worldwide data yet from, say, shipping tapes over there and letting them sit in that environment and correlating it back to our environmental chamber being at those types of temperatures. But what we are doing with the experimental matrix is coming down to what I would call the lower temperatures, which are more in line with those kinds of locations in the world. And eventually, we will have to develop that kind of correlation.

Right now, we are just using it to fit the model. We will establish a series of conditions where enough change can be observed in a short enough time to build a model and then extrapolate it further.

I know you run into danger when you extrapolate outside your experimental region; you always run into a danger as to how good the model fits once you have left your original boundaries. But this is something that is going to take some time obviously; it is just sort of in the process now.

DR. KRYDER: Dennis?

DR. SPELIOTIS: By the way, making a model on the basis of magnetic data can be very deceiving because most people measure the entire change of magnetic moment of a piece of material; and typically, the coating thickness is about 3 micrometers. That does not matter at all in recording.

At high density recording, it is only the top .2 or .3 micrometers that has anything to do with the signal that you read out of the tape. Therefore, you must measure of the change in surface magnetization, and not in the average bulk magnetization.

We did some such measurements and found that for all the tapes regardless of whether they were double protected with ceramic coating or just the oxide, roughly, the average surface magnetization change was three times greater than the total magnetization change.

If you take this three-time multiplier for the loss in magnetization, you will find out that you are in trouble under some conditions. So, you must look only at the surface and not inside.

In that respect perhaps, the error rate change is the most realistic measurement, except that error rates are composites. You cannot blame the particles alone; it could be the particles, or it could be the binder; or, it could be the surface; or it could be any combination of the above, and you have to sort of separate the effects.

DR. KRYDER: Okay. Another question?

PARTICIPANT: Somebody earlier said that there was really a difference between archival and environmental. And most of what has been talked about really strikes me as being environmental -- (Inaudible)

DR. KRYDER: Right.

PARTICIPANT: If you protect the tape against these things, aren't there any archival problems? Will it be here 100 years from now? Will it -- (Inaudible)

DR. KRYDER: Okay. Let me repeat the question. The question was: We have been having a lot of comments about having tapes in the typical environments--shipping to Saudi Arabia

and so forth; and what if they are properly stored in the "archival environment"? How long will these various tapes really last?

MS. CARLSON: Maybe I can address that. The bottom line of that is recorder obsolescence.

Historically, in the last 20 years, the weakest link in the chain has been either binder hydrolysis or recorder obsolescence. That's the balancing act that has been going on.

Binder hydrolysis--the chemistries that have been coming on board have really improved a lot. So, the biggest question is: What is your recorder obsolescence or your recorder half-life?

MR. HADAD: Yes. One thing that Pat said is that it would be around forever; but then, I sort of wrote down the next thing he said. "**Make sure we will be around long enough to make the copy.**"

(Laughter)

MR. CORCORAN: Could I give you a criterion? The earliest video recorder made at Ampex in about 1955 was a quad recorder. Those things were used in the studios until somewhere around 1975. They have since disappeared into museums; there are a few places you can find them.

They had these little glass things in them called vacuum tubes. You can't find the replacement parts.

Somebody gave a talk about TBM and finding the replacement parts on that; and that was built from 1968 to 1972. That will give you a kind of rough criterion as to what the typical life of these systems are.

DR. KRYDER: They are probably going to get shorter, though, I would argue. Technology tends to change faster.

PARTICIPANT: There are ways around that. There is the obsolescence of chips, and people now make -- chips that they can put -- (Inaudible)

DR. SPELIOTIS: $\gamma\text{-Fe}_2\text{O}_3$. (Inaudible discussion)

DR. KRYDER: I guess another comment could be made. I mean, if you are willing to spend the money, the tapes still exist. Even if the drive doesn't exist, we could undoubtedly recover it; but you have got to pay for the development of that drive again.

PARTICIPANT: (Inaudible)

DR. KRYDER: We have a question down here?

PARTICIPANT: (Inaudible)

DR. SPELIOTIS: Seven and a half weeks was the test.

PARTICIPANT: Okay. (Inaudible)

DR. SPELIOTIS: There have been some correlations. This was an accelerated test; 50° C/80% relative humidity. This is not a standard operating environment. The correlations show that seven and a half weeks correspond to seven years in real life; those correlations have been done by other laboratories.

DR. KRYDER: A question in the far corner?

PARTICIPANT: (Inaudible)

(Inaudible discussion)

MS. CARLSON: You must be in a situation where budgets aren't of concern; but boy, in the last five years, we have had to do more with less.

PARTICIPANT: (Inaudible)

DR. KRYDER: Okay. We have another question?

PARTICIPANT: What's the wicked truth about -- (Inaudible)

(Laughter)

MS. CARLSON: In some conditions, yes.

DR. SPELIOTIS: The problem is much worse obviously with 4 millimeter than with 8 millimeter. I would think the reliability of 8 millimeter is much better than 4 millimeter. 4 millimeter is probably a real problem spot.

The other question is about the corrosion stability of these tapes. If you keep them in a real archival environment; low temperature, humidity, there is no problem. But 8 millimeter is now finding its way as backup storage for many small computer systems--home computers, office computers.

Those tapes are going to be trouble . (Inaudible)

So, it depends on what environment and how you use it.

PARTICIPANT: (Inaudible)

DR. KRYDER: Can he expect a problem in that environment?

MS. CARLSON: Yes. Normal tape packs--and I think Dr Bharat Bhushan has done a little bit of this tape pack analysis, where what happens is that you do get stress relaxation in the tape pack; and therefore, you have a differential in tension. And so, that is the reason why you recommend retensioning.

Unfortunately, with 8 millimeter, I don't know when and how much at this particular moment in time; but you will have to retension most cassettes.

PARTICIPANT: (Inaudible)

MS. CARLSON: Retension once a year.

MR. HADAD: Once a year is what we recommend.

PARTICIPANT: (Inaudible)

DR. REAGOR: One of the other things that we noticed also in our studies, and all studies we did on the tapes was on virgin tape, it turns out when you manufacture the tape, it was stress-induced during the manufacture and the initial first run; and actually if you take it through one low temperature cycle, you can actually remove that initial stress that sometimes later relaxes when you are using it.

PARTICIPANT: (Inaudible)

DR. REAGOR: Yes.

DR. KRYDER: Other questions?

PARTICIPANT: One issue that we have come across is how we got started in this entire discussion; and that is, we wanted to -- (Inaudible) --

DR. SPELIOTIS: On typical MP tapes in the market, as much as two orders of magnitude difference in error rates under some conditions from one tape to another, the best to the worst.

DR. KRYDER: A question?

PARTICIPANT: (Inaudible)

DR. KRYDER: Yes. So, I guess to sum the last two comments, there are variations from vendor to vendor; and there are also variations even within one good vendor that are significant.

PARTICIPANT: (Inaudible)

DR. KRYDER: Okay.

PARTICIPANT: (Inaudible)

DR. SPELIOTIS: For a long time, from the beginning of our modern era until five or six or seven years ago, most of the data was stored on plain gamma--no chromium, no metal, nothing--and gamma is stable.

So, I think that is one of the main reasons we have not lost the data. Now, we are moving into new media, and these are not as stable as the old ones; they are not.

MR. HADAD: The data I presented, again, under ambient conditions, it is routinely maybe a dB change after three years. I think that is pretty stable. I don't know that there is that much reason to be concerned with those kinds of numbers. I imagine there are 2 percent changes in some cases, exposed to 30° C/70% relative humidity. Again, maybe we are getting too excited, or we see the curves of the downward slope; but yet, we ought to rescale the axis.

With a 0 to 100% loss, our line is way up here, of course we know with our Lotus and Excel spreadsheets that we can take just a little bit of that and blow it up; and it looks like the curve goes straight down. So, three-year data, 1% , somewhere around there.

PARTICIPANT: (Inaudible)

(Inaudible discussion)

MR. HADAD: Yes, it is much worse there. Now, you are starting to exceed the material. You know, basic material properties, at 85° C, you have pushed the glass transition point of the base back.

At the moment, we can't do anything about that; you just can't recommend tape to be in that environment. You can't recommend something that will be that high.

MS. CARLSON: The point that I would like to make on that is that's the reason why archive systems are system-specific. In that particular case, there are workarounds. If you are in a situation where you are at high temperature and humidity, what you would do is you would

identify those tapes and then copy them over onto an archival media. All of these systems can be used in an archive way if it's implemented properly.

Is there any reason why that high temperature and humidity tape has to be kept forever?

PARTICIPANT: No, I think the issue is that -- (Inaudible)

MR. CORCORAN: Let me comment. We have a number of DCRS¹ recorders out there in that desert. And I heard no trouble about the tapes. In fact, the one complaint we got was that they kept reusing them, and they never bought any more.

(Laughter)

MR. CORCORAN: Now, they were rerecording them almost every night and then rerunning them. But I'm pretty sure they didn't go through the glass transition temperature, or I don't think they would have got any data.

PARTICIPANT: (Inaudible)

MR. CORCORAN: I hear what you are saying, but I think what you have to do is distinguish between this "archival condition," when you are trying to go for long life, and the extreme environment for maybe a week.

And it is routine, for instance in electronic news gathering cameras, to have trouble with humidity and so forth on site. But that doesn't generally destroy an archival; it will muck it up in a week, though.

DR. REAGOR: And in your case, what you are trying to do is sending a package data system --

To be honest with you, from a user's standpoint, I would look at it as: What I send is dependent upon what the person who is going to receive it can do with it. So, that format and the type and the volume is going to be based on who you are sending it to.

And the real key is the transportation issue and the protection. No matter what material you are hearing right now, the key seems to be that we don't leave it sitting on the dock in the hot sun on a humid day, unless it is in some type of a protective container.

And I think the issue you are raising, that is: I had to call them and say "Hey, guys, don't leave this in the back seat of your car four days before you read the tape." This is really the thing that you are going to have to do.

DR. KRYDER: Did you have a question?

PARTICIPANT: (Inaudible)

DR. KRYDER: I'll repeat the question. If we wait two years for barium ferrite to be developed, do we have the perfect medium; or are there some down sides to that?

DR. SPELIOTIS: Let me say one thing about that. The Japanese industry has already developed barium ferrite as a copy tape for DAT audio. In other words, they record on metal particle masters, classical music, for example, and for multiple distribution, they copy by anhysteretic transfer onto Barium ferrite tapes.

They are not marketing them yet; I don't know why, but they have been developed.

DR. KRYDER: But that is not his question. The question: What is the down side to it?

MR. HADAD: The chemistries are probably going to be very similar to the binder system. Still, we haven't gotten entirely around all the environmental problems with the polyester urethanes that everybody is using. They are subject to hydrolysis, mobile lubricants--all the things that we need also had some -- to it; and they are probably going to be present in the next two years.

DR. KRYDER: Could I get the panel to agree that if the world does go barium ferrite, the stuff would be as stable as gamma was?

DR. SPELIOTIS: Absolutely -- no change.

MS. CARLSON: Barbara brings up a point, and it depends upon-- Barium is toxic; it is a heavy metal; it is a toxic metal. And it depends upon the pigment -- if that material is going to be a toxic waste issue. In the research that we have done, we have seen a wide variety of experimental products that could be toxic. The design engineers have modified it so it won't be. I am anticipating that, by the time it is on the market, barium will not be a toxic waste issue.

DR. REAGOR: As a user, I can address this. If there is a toxic question concerning any of the materials, from a large user's stance, that is something we would not want to get involved in --

To give you an example of that right now, the EPA is just about ready to make telephone poles toxic waste.

(Laughter)

DR. REAGOR: That would cost my own company about \$500 million a year from now on if that happens. So, we are doing a lot of work right now on telephone poles.

MR. CORCORAN: Could I make one other point about -- From a system viewpoint, it has a drawback on low frequency signals. (Inaudible) -- where we have to put in things like audio and control signals. The barium ferrite is not very good at that. It works great at the high frequencies, but not in high syncopation --

DR. SPELIOTIS: But John, all high density future systems, including DAT, digital video, etc., are going to be PCM or band-limited digital.

MR. CORCORAN: We have to build a whole new family of recorders to do that.

PARTICIPANT: (Inaudible)

MR. CORCORAN: Well, this will be another family. (Laughter)

DR. KRYDER: A question over here?

PARTICIPANT: Is that the only down side for the -- (Inaudible)

MR. CORCORAN: Every time we tried changing the tape--and I do this routinely--trying to bring in some new tapes and try them on the machines, we encounter another problem. And then, we have to find a way to resolve that. Now, for instance, this could be solved by changing the control track electronics. But these are the kinds of things that you stumble through as you try to move forward.

MR. HADAD: It's not a direct swap out to any existing -- right now.

MR. CORCORAN: Yes, but the down side is that it's hard to predict what they are going to be until you actually get it on the machine and try it.

DR. KRYDER: Any time you change part of a recording system, you have to change other parts--that's my experience.

(Laughter)

DR. REAGOR: Earlier, you had referred to the metal evaporated tape as having sort of like this linear corrosion and then, all of a sudden, it went exponential. And you said it was opposite from what you expected.

DR. SPELIOTIS: This was on MP, and not ME tapes.

DR. REAGOR: But from the work we have done on metal surfaces, that is what you would expect.

DR. SPELIOTIS: Yes?

DR. REAGOR: When you are doing gas phase corrosion of a metal surface, whether it is particle size or whether it is a full layer, and you store that with a very smooth, clean surface, then you have to basically have an initiation period for what they call "mound formation"; that's the corrosion term.

And so, you have to start to develop a spot that finally starts to corrode, and that initiation period can be very long. And then those mounds finally grow together. I noticed on your data published, you almost look like in the early stages that you have clusters of corrosion, which would be exactly what you would expect in mound corrosion of a clean metal surface.

So, what we are seeing here is--and I don't want to scare anyone, because I don't think it is going to accelerate after ten years and finally fall off the edge-- but I think it tells us that the metal particles and the metal surfaces are behaving like the metals do. And so, you can now extrapolate even a little further and look at these metal systems and say: Okay, what are the corrosion dynamics of these materials in the real world?

And then, you can go to a Battelle test that does copper surfaces; and if the copper metal surface behaves in this way and is accelerated this way, it is applicable to a metal particle surface or to a metal evaporated tape surface; and so, there is some correlation here.

And what you were saying and describing basically fits into that whole field.

DR. SPELIOTIS: By the way, the corrosion of chromium dioxide and cobalt-modified iron oxide, as far as I can understand, does not involve any significant change in the volume of the particle. There are other mechanisms, like the diffusion of the cobalt that lowers the coercivity in the case of cobalt modified. But in the case of metal particles, the particle grows very significantly in size when it oxidizes--20 or 30 percent growth.

Now, what is going to happen? It's going to distort the surface; it is going to do some damage -- and it's a different mechanism of corrosion.

MR. HADAD: I think that some of those charts that I showed, though, that under normal operating conditions, you are not going to get that kind of corrosion. I want to mention something regarding metal particles starting off as a clean surface; they are not. There is a passive layer. They already are corroded to some degree.

Passivity, just by definition, is an oxide layer that protects further inside the core. And so, I don't know; I've never observed it with metal particle tapes where there is some long

induction period and then they take off. They already come pretty corroded, as Dennis said; the EMUs per gram of pure iron is 210. The material you use for making magnetic tape is 115 or so.

They are precorroded, but they give the outward appearance that people are interested in. The sheer fact alone that that oxide is present is actually protecting further those particles. They have already started corroding and stopped.

Of course, their passivation kinetics determine how long these things are going to last; and what we see is that they are going to last many, many years. Of course, the process will continue; but it is way beyond, I'm sure, our lifetime or our expectations of the media.

DR. KRYDER: A question down here?

MR. ROBINSON: Since I've been invited to make a presentation on tape recording, I thought I might say a word or two. I doubt that the coercivity change of the tape will make a difference, you know, going from Germany to here.

If I were you, I would do what I've done in the lab many times, trying to test machines: I would check my tracking.

(Inaudible discussion)

MR. ROBINSON: Let me just say another word. If I were using metal particle tape, I think I would clean the heads also because of brownstaining. Exabyte, I think, recommends once a day. (Inaudible) So, I would be very concerned about that.

I would also be concerned about using 8 millimeter without my track -- If you recall my presentation, the ID-1 Working Group, including Sony, DataTape, Ampex, and others, selected a 1.6 mm. And we also said, in answer to many, as was reported, that we would make our track as powerful as possible.

I believe we need some advice from manufacturers --from you folks--about whether we are doing the wrong thing when we go with a half mil track width with metal particle tape just because that gives us a hell of a lot more data on tape. Do our users really want just a hell of a lot more data on tape, or do they want a reliable system? Do they really want to crossplay that stuff when they ship it in from the Gulf? Or what do they want to do?

So, in the ID-1, we try to be as reasonable as we can and build in the tolerances that it takes to make the damned things work. But we need advice from you guys. One of the biggest problems I ever had with tape recorders was making the damned media work.

(Laughter)

MR. ROBINSON: On the system 600, we went through two or three tapes to get the quality we needed -- finally solving all of our problems. The media is important; and tolerances are important on tracking--very important.

And I don't know where we are going to stop, you know; maybe with barium ferrite, we can do half-mil write tracks.

MS. CARLSON: You are still going to be on polyester; you know, that's the driving thing.

(Inaudible discussion)

MR. ROBINSON: So, I think you media people should advise the recorder people to make the tracks wide enough so these guys won't have the problem.

PARTICIPANT: (Inaudible)

DR. REAGOR: And from my user's standpoint, we would like to have as much data as you could put in a square inch of media that will be there 20 years from now; and I don't care if I put it under water; I don't care if I burn the building around it or whatever--it's there.

With the phone system we have right now, if there is lightning and your power goes out, you pick up your phone and it works. So, the phone company has this attitude that what they use to work with their system should be the same.

So, from a user's standpoint, I agree with what you said. I can't tell you what to do; I'm just telling you what I want.

MR. ROBINSON: What we should do is be conservative in design.

PARTICIPANT: (Inaudible)

MR. HADAD: Your statement is absolutely correct. 1 percent is something like 8/10ths of a dB, and that doesn't bother an ECC system at all.

DR. SPELIOTIS: But on the other hand, most of these tapes are for real high density environments; they operate at an error rate of 10^{-5} , with no correction. And 10^{-4} is about the limit for error correction to get you to 10^{-12} . In other words, you have an order of magnitude roughly and degradation narrows the safety margin.

PARTICIPANT: I think that is an incorrect statement.

DR. SPELIOTIS: Why?

PARTICIPANT: Why?

DR. SPELIOTIS: Well, of the typical Reed-Solomon error correctors used require tolerate input error levels of about 10^{-4} , or maybe seven times 10^{-3} , before they fall off the edge.

MR. CORCORAN: If you allow us enough overhead, we will bust it through to any level that you want. Okay? On a terabit memory system, they used double recording and then record over it. And I believe on one of the audio systems, they do the same thing. They double record and then they record over that.

And they can get 10^{-15} or 10^{-16} . In fact, it becomes very difficult to measure it.

DR. SPELIOTIS: Right, but what do you start with and what is the limit?

MR. CORCORAN: I just gave you two examples, and I know that we can get up to almost any number you want; it is just going to take more overhead. Most of the time, we are not allowed 100 percent overhead.

We have seen recorders where we were only allowed 10% overhead, 5% for ECC. The lady in the back said: Well, you have all these different problems; and it winds up being a trade-off between the different alternatives, and you are squeezed in almost every direction. And you have to make a good judgment.

PARTICIPANT: (Inaudible)

MR. CORCORAN: Yes.

DR. KRYDER: Other questions?

PARTICIPANT: I would just like to make a statement. I'm surprised somebody up here hasn't already done so; and that is that the care and feeding of magnetic tape is not a new subject.

I know that Ampex has brochures that they make available to customers that say: Here is how you treat magnetic tape. And there is no mystery to that. You don't mistreat it, folks; that's the basic problem.

I'm afraid the old iron oxide tape has been forgiving enough that you may have been misled into thinking that you can mistreat them.

PARTICIPANT: I would suggest that you go back to your friendly dealer and pick up their recommendations on how to treat the stuff. (Inaudible)

When IBM first came out with their 3480 tape, they shipped that stuff in refrigerated cars

MR. HADAD: Well, we are pushing the edge. People's requirements are greater; that's true, and there are materials out there to make more robust tapes. But I don't think anybody wants to pay for them; they are between \$1,000 and \$10,000 a reel. Yeah, we can probably make a tape that is stable to everything; but nobody is going to buy it.

DR. MALLINSON: I have a question for Dennis. (Inaudible)

DR. SPELIOTIS: The recording performance is equivalent?

DR. MALLINSON: Yes.

DR. SPELIOTIS: The recording performance is superior. Barium ferrite is far superior to MP at very high densities (above 80 KFCI)

(Inaudible discussion)

DR. KRYDER: Does anyone have anything else?

MS. CARLSON: To expand on what Dennis is saying, just think about barium ferrite particulates. If you go into Target or K-Mart, you can see VHS all over the place. The industry is scaled up for particulate media. So, the cost per package is going to be much lower just because there is an industry that is all scaled up to manufacture, at very reasonable cost, particulate media. With ME, you have to put the capital investment in in order to get your manufacturing up to even 1% of the worldwide capability that we have now in particulates. So, that is going to be a huge investment for the manufacturers, and they don't want to go out of business. It has got to make very good business sense in order for that to happen.

(Inaudible discussion)

DR. KRYDER: Any other questions? Comments?

(No response)

DR. KRYDER: Okay. I want to thank the panel and thank the audience for all your participation.

(Applause)

(Whereupon, at 6:25 p.m., the conference was recessed, to be reconvened on Thursday, July 25, 1991 at 9:00 a.m.)

N93-16331

**19 mm DATA RECORDERS
SIMILARITIES AND DIFFERENCES**

by

**STEVE ATKINSON
AMPEX CORPORATION**

PRECEDING PAGE BLANK NOT FILMED

Vertical text on the left edge of the page, possibly a page number or margin note.

Vertical text on the right edge of the page, possibly a page number or margin note.

Main body of the document containing several paragraphs of text, which is extremely faint and illegible.

19 mm DATA RECORDING FORMATS

Confusion over the use of non-video 19 mm data recorders is becoming more pronounced as we enter the world of high performance computing. What is the difference between ID-1, ID-2, MIL-STD-2179 and DD-2? What is the proper machine for my application? How do I integrate it into my environment? These are all questions the user community needs answered. This paper attempts to address these issues and clear up any misconceptions there might be about 19 mm tape recorders.

HISTORY

The MIL-STD-2179 and ID-1 tape formats are modifications of the D-1 tape format standard used in the television and motion picture industry. The development of D1 based instrumentation standards was driven by the need of military and government users for high speed digital recorders to capture data. The Military Standard 2179 (MIL-STD-2179) was supported by a group of manufacturers and military representatives (primarily the United States Navy) to provide such a recorder. The first prototypes of these recorders are becoming available now for Beta testing in a number of military programs. Loral, Sony, Datatape, Honeywell, and Schlumberger are all working to deliver these machines. An ANSI standard known as ID-1 also exists which is being implemented by same manufacturers. It differs slightly from the 2179 implementation but is not interchangeable due to the use of azimuth recording, which is not used in the D1 TV format or the MIL-STD-2179 format.

Products to support the television D1 format became available in early 1987 and were followed by D2 format products in late 1988. The two TV digital formats differ in that D1 supports component recording while D2 supports composite recording. In laymen's terms, this means that D1 records video as individual data streams and D2 records it as a single stream. The Ampex ID-2 product is based upon the unmodified D2 NTSC video format, but is supplemented with user interfaces specific for instrumentation applications. As such, the Ampex ID-2 products can serve as either a video or instrumentation recorder. D1 recorders are produced for the video industry by Sony and BTS and D2 recorders are manufactured by Ampex with Sony and Hitachi. The D2 recorder is the dominant 19 mm machine in the video industry constituting over 75% of the digital market today and is expected to increase in market share to over 80% by the early 90's.

The need to support high speed data storage and retrieval for the supercomputer industry was identified by Ampex as a major market opportunity as data creation and storage grew exponentially in the 1980's. Hence, the Data D2 (DD-2) development was undertaken as a joint development by Ampex and E-Systems. The basic difference between the DD-2 recorder and instrumentation recorders is that DD-2 is designed as a computer peripheral with computer interfaces and file structures as opposed to an instrumentation recorder which records input as a single uninterrupted stream of data. Each class of machine has its strengths and weaknesses depending on the mission the user wants to accomplish.

COMMON FEATURES

DD-2, MIL-STD-2179, ID-2, and ID-1 physically record data in very similar manners. All implementations use a recording technique know as helical scan. Helical scan recorders lay down data at an angle on tape and provide improved storage density over longitudinal formats and also provide powerful error correction capabilities. Ampex was a pioneer in the helical scan technology used in today's home VCR's.

Each of these scans of recorded data are called tracks. They are grouped together to form track sets, the lowest addressable unit of data from a physical standpoint. The

helical scan data tracks are accompanied by three longitudinal tracks. In the video version, these tracks contain control, time code and audio information. The instrumentation and data recorders use these tracks for control, timecode, and for file labels in the case of DD-2.

All four 19 mm formats are implemented in high performance machines with several versions tailored for specific missions. Features commonly supported include high rate data storage and retrieval, robotic compatibility, and positioning to data at speeds of 30 to 60 times playback speed. For DD-2 at 60 times playback speed, this translates to over 800 megabytes per second (MB/s) or the equivalent of four 3480 cartridges every second.

With these similarities in mind, it is easier to discuss how the DD-2 computer storage peripheral and 19 mm instrumentation recorders use the 19 mm platform to accomplish different tasks with different derivations of the same basic technology.

INTERFACES

The applications for instrumentation recorders such as ID-1, and computer peripherals such as DD-2, drive the design of their interfaces. Instrumentation machines are commonly used for data capture which means that data will come in a continuous stream with the recorder turned on when the data starts and off when the data stops. This alleviates the need for any buffering on the recorder since the data should arrive in a continuous mode at a steady rate. The standard implementation is a 16-bit wide data front-end which operates in a synchronous mode. By implication, this means that the recorder operates as the master with the input or output source acting as the slave. This is because the recorder expects to record and retrieve data at a steady, uninterrupted stream. Typically, it handles different data rates by the ability to record/retrieve data at selectable, binary rates of up to approximately 30 MB/s in some implementations. Therefore, it is up to the source to buffer the data for smoothness within a small percentage of the recorder's data rate needs. While this is a good implementation for the field and downlink data recording it was designed for, it poses severe problems in a computer environment.

Most computer operating systems simultaneously execute multiple programs in an interactive, time sharing environment. Input and output to a peripheral is not supported as a steady, uninterrupted stream, but instead is dependent upon the dynamics of system loading, computational rates, and application mix. These operating environments do not include the dedicated data buffering and strict command scheduling required to match rates with a data peripheral. The logic of these operating systems is designed to use compute cycles to serve applications -- not peripherals.

Instrumentation recording on a computing platform requires a dedicated computer, executing only from a predetermined set of controlled applications. Development of instrumentation interfaces, controllers, and special device drivers are required to connect a given peripheral. Memory must be dedicated to support the strict input/output rates of the peripheral.

The DD-2 recorder was designed with these interface shortcomings in mind. The DD-2 recorder from Ampex uses the ANSI standard Intelligent Peripheral Interface Level 3 (IPI-3) interface to send data to and retrieve data from the recorder. This standard computer interface defines command, control, and status for both disk and tape peripherals. The subset of the IPI-3 command set dealing with tape devices is used for the DD-2. A large buffer is included as part of the interface electronics. This buffer performs the rate smoothing to provide a sustained data rate of 15 MB/s with a burst rate of 20 MB/s. Using the IPI-3 command set and the buffering capability allows the host to request data units as small as a single byte so that the host computer does not have to retrieve an entire physical block. The IPI-3 implementation also allows up to 8 recorders to be "daisy-chained" on a single control interface. While one recorder is transferring data, the others can be given positioning commands to minimize the

time spent waiting for the subsequent file transfers. Other standard interfaces will be supported in the future. In summary, the DD-2 looks like a standard tape peripheral to the host computer.

DATA FORMAT

As stated earlier, both DD-2 and 19 mm instrumentation recorders physically record data on tape in very much the same way, helical scans grouped as track sets. Once again, the differences lie in the fact that the DD-2 is a computer peripheral recorder. Both machines can retrieve data by track sets if that is the search parameter they are given. The DD-2 also makes use of the longitudinal tracks to implement ANSI 9-track file labeling. Files on the DD-2 can be retrieved by using file labels rather than track set ID's. This allows software already developed on the host computer to manipulate 9-track tape and 3480 cartridges to also make use of DD-2 cassettes. Minimizing the impact of incorporation of DD-2 into existing computer is one of the highest priorities of the DD-2 development.

The traditional concept of tape volumes is incorporated into the DD-2 design while not in 19 mm instrumentation formats. Each DD-2 physical cassette looks like one or more logical volumes to the host computer. These logical volumes can be edited and appended to look like traditional 9-track volumes. Another somewhat related innovation is that the cassettes can be loaded and unloaded without rewinding to beginning of tape or end of tape. This greatly decreases the amount of search time spent positioning to the beginning of a file.

Inter-record gaps are another format difference between 19 mm instrumentation recorders and DD-2. Because most 19 mm instrumentation machines record data at binary rates, the gaps between recorded data are of variable size. This leads to less efficient use of tape as it is indeterminate how much tape is left unrecorded in between data records. The DD-2 uses a consistent gap to further improve upon the substantial data density advantage DD-2 holds over the ID-1/MIL-STD-2179 class of recorder.

ERROR RATE

Image capture and retrieval is one of the most common uses of ID-1 and MIL-STD-2179 recorders. When manipulating images on a host computer, errors on tape can be recovered from by reconstructing the lost part of an image from the remaining parts which surround it. For that reason, the Bit Error Rate (BER) of 10^{-10} specified for MIL-STD-2179 and ANSI ID-1 recorders is sufficient for most of the applications in which it is used but may require tape certification to achieve. This BER is not always adequate for computer applications. DD-2 will deliver an error rate approximately three orders of magnitude better by guaranteeing less than one error event in every 10^{12} bytes read.

All commercial 19 mm recorders have the capability to read what has been written on tape immediately after it is written. However, the implementation of a data buffer in the DD-2 interface makes verification of the data written on tape possible. When an uncorrectable error is detected by the read-after-write verification function, the data will be rewritten to tape a selectable number of times by the DD-2 controller. When retrieved, invalid data is ignored with only the corrected data returned to the host. This capability is coupled with three levels of powerful Reed-Solomon coding to achieve the improved error rate over 19 mm instrumentation formats. For example, the extended burst correcting capability of DD-2 provides a 50-times improvement over ID-1, minimizing errors due to tape defects.

In addition, the DD-2 will monitor the stress level of the error detection and correction (EDAC) equipment and store the statistics in the recorder. When requested from the host computer, it will be returned as status. In this manner, the host can monitor the "health" of an individual cassette and record the data contained on it to a new cassette when appropriate.

MEDIA

The media used by data storage peripherals and instrumentation recorders is the same as used in the video industry. The cassettes used by all 19 mm recorders contain tape which is 19 millimeters wide (about 3/4 inch). The main difference is that DD-2 and ID-2 use a metal oxide based tape with a total thickness of 13 micrometers, while the ID-1 and MIL-STD-2179 use the older iron oxide formulation with a total thickness of 16 micrometers. 19 mm cassettes are very similar in appearance to home VHS cassettes except that VHS uses 1/2 inch tape. D2 tape is manufactured by Ampex, Sony, Fuji, and Maxell with 3M expected soon to enter the market. D1 tape is available from Ampex, Fuji, and Sony. The important point is that both are supported by more than one source which should result in both a reliable future supply and low user costs.

The storage density of D1 and D2 based recorders is a function of the recorder and tape used. A storage density comparison is contained in this section for simplicity's sake. Both DD-2 and ID-1/MIL-STD-2179 have three sizes of cassettes: small, medium and large. The small D2 cassette contains 25 gigabytes (GB) of user data with a medium containing 75 GB and a large holding 165 GB. This compares to 14 GB, 44 GB, 92 GB for the small, medium, and large D1 cassettes respectively.

A concern has been raised on the shelf life of metal oxide tape due to the high content of iron particles in the tape. Metal particle tape has been in existence over 15 years and has been used in commercial products for over 10 years. All 8 mm Camcorders, as well as other video and data storage products, are based on the metal oxide tape. Ampex has recently completed accelerated life tests which simulate a 14-year archive. No degradation in BER or magnetics was detected with the tape in the cassette. Details were presented by Ampex at the Tape-Head Interface Conference (THIC) proceedings on 9 January 1991. Copies of the presentation are available from Ampex. Tests by other manufacturers also support these results.

SUMMARY

DD-2 and 19 mm instrumentation recorders have missions for which each is well designed. While the differences may appear subtle, understanding the difference between the two is the key to picking the right recorder for your particular application.

N 9 3 - 1 6 3 3 2

**DATA STORAGE AND RETRIEVAL
SYSTEM ABSTRACT**

by

**BARBARA MATHESON
STX CORPORATION**



DATA STORAGE AND RETRIEVAL SYSTEM ABSTRACT

The STX mass storage system design is intended for environments requiring high speed access to large volumes of data (terabyte and greater). Prior to commitment to a product design plan, STX conducted an exhaustive study of the commercially available off-the-shelf hardware and software. STX also conducted research into the area of emerging technologies in networks and storage media so that the our design could easily accommodate new interfaces and peripherals as they came on the market.

All the selected system elements were brought together in a demo suite sponsored jointly by STX and ALLIANT where the system elements were evaluated based on actual operation using a client-server mirror image configuration. Testing was conducted to assess the various component overheads and results were compared against vendor data claims.

The resultant system, while adequate to meet our capacity requirements, fell short of transfer speed expectations. A product team lead by STX was assembled and chartered with solving the bottleneck issues. Optimization efforts yielded a 60% improvement in throughput performance.

The ALLIANT computer platform provided the I/O flexibility needed to accommodate a multitude of peripheral interfaces including:

- Up to twelve 25MB/s VME I/O channels
- Up to five HiPPI I/O full duplex channels
- IPI-2, SCSI, SMD and RAID disk array support
- Standard networking software support for TCP/IP, NFS, FTP.
- Open architecture based on standard RISC processors
- V.4/POSIX-based operating system (Concentrix)

Standard 1 Terabyte System:

- Alliant SRM-1 computer with approximately 100 GBytes of user online storage (capable of expansion to over 1 TeraByte) in the maximum strategy arrays and one point-to-point link for a Cray supercomputer using the UltraNet link to provide up to 50 Mbits/s transfer and a STK 4400 ACS with one LSM and a 2-controller CAS with eight transports (four transports for

archiving and four transports for backup to provide a backup transport for each function).

This configuration distributes I/O over four VME data paths with each bus providing a sustainable transfer rate of 40 MBytes to memory and 24 MBytes from memory. Three VME chassis are used to provide four data paths by subdividing one chassis into two paths, yielding a total of 29 device slots. The described configuration uses 14 slots, leaving 15 space slots that can be used for a mix of additional disk arrays, network interfaces or archival read/write stations.

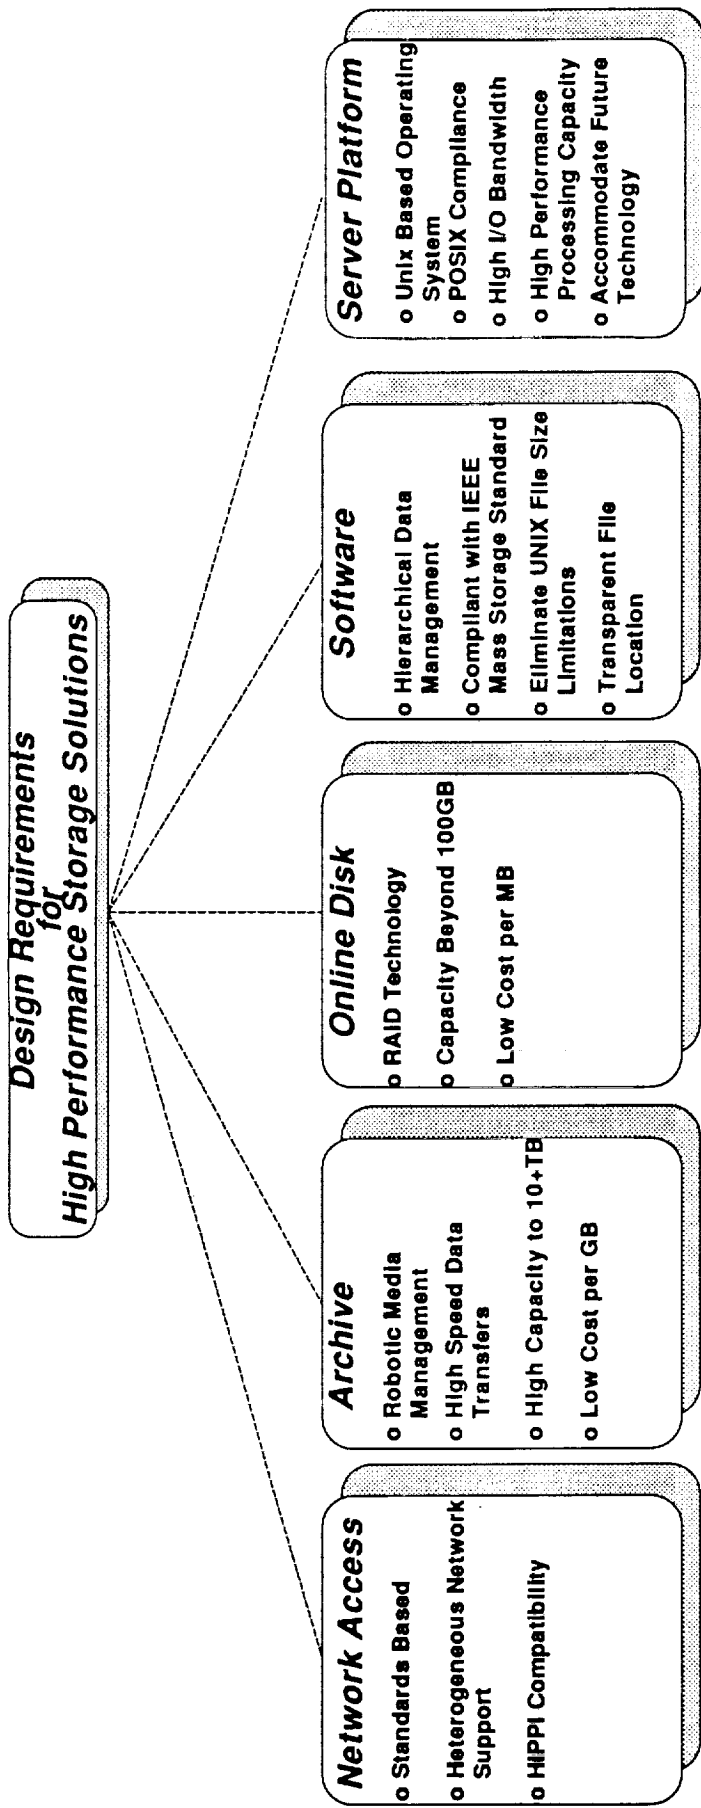
Aside from performance requirements, the companies surveyed many candidate users to determine other features and benefits that would be valuable to an end-user. The most common concern was that a system put in place today still be viable in 5 years. Our product team then turned its attention to growth path issues to ensure that newer, higher speed network architectures could be accommodated and that new storage technologies could be added as they became available.

All components including the software are modular in design and can be reconfigured as needs and system uses change. Users can begin with a small system and add modules as needed in the field. Most add-ons can be accomplished seamlessly without revision, recompilation or re-linking of software.

Since the software is device-dependent only at the lowest layer, new interfaces and peripherals can be easily accommodated.

Ongoing development is focused on:

- Implementation of HiPPI capability
- Selection and interface of D-1, D-2 tape systems
- Interface to Metrum RSS-600
- Interface to SUMMUS//MAGNUS tape jukebox



Design Requirements for High Performance Storage Solutions

Network Access

- o Standards Based
- o Heterogeneous Network Support
- o HIPPI Compatibility

Archive

- o Robotic Media Management
- o High Speed Data Transfers
- o High Capacity to 10+TB
- o Low Cost per GB

Online Disk

- o RAID Technology
- o Capacity Beyond 100GB
- o Low Cost per MB

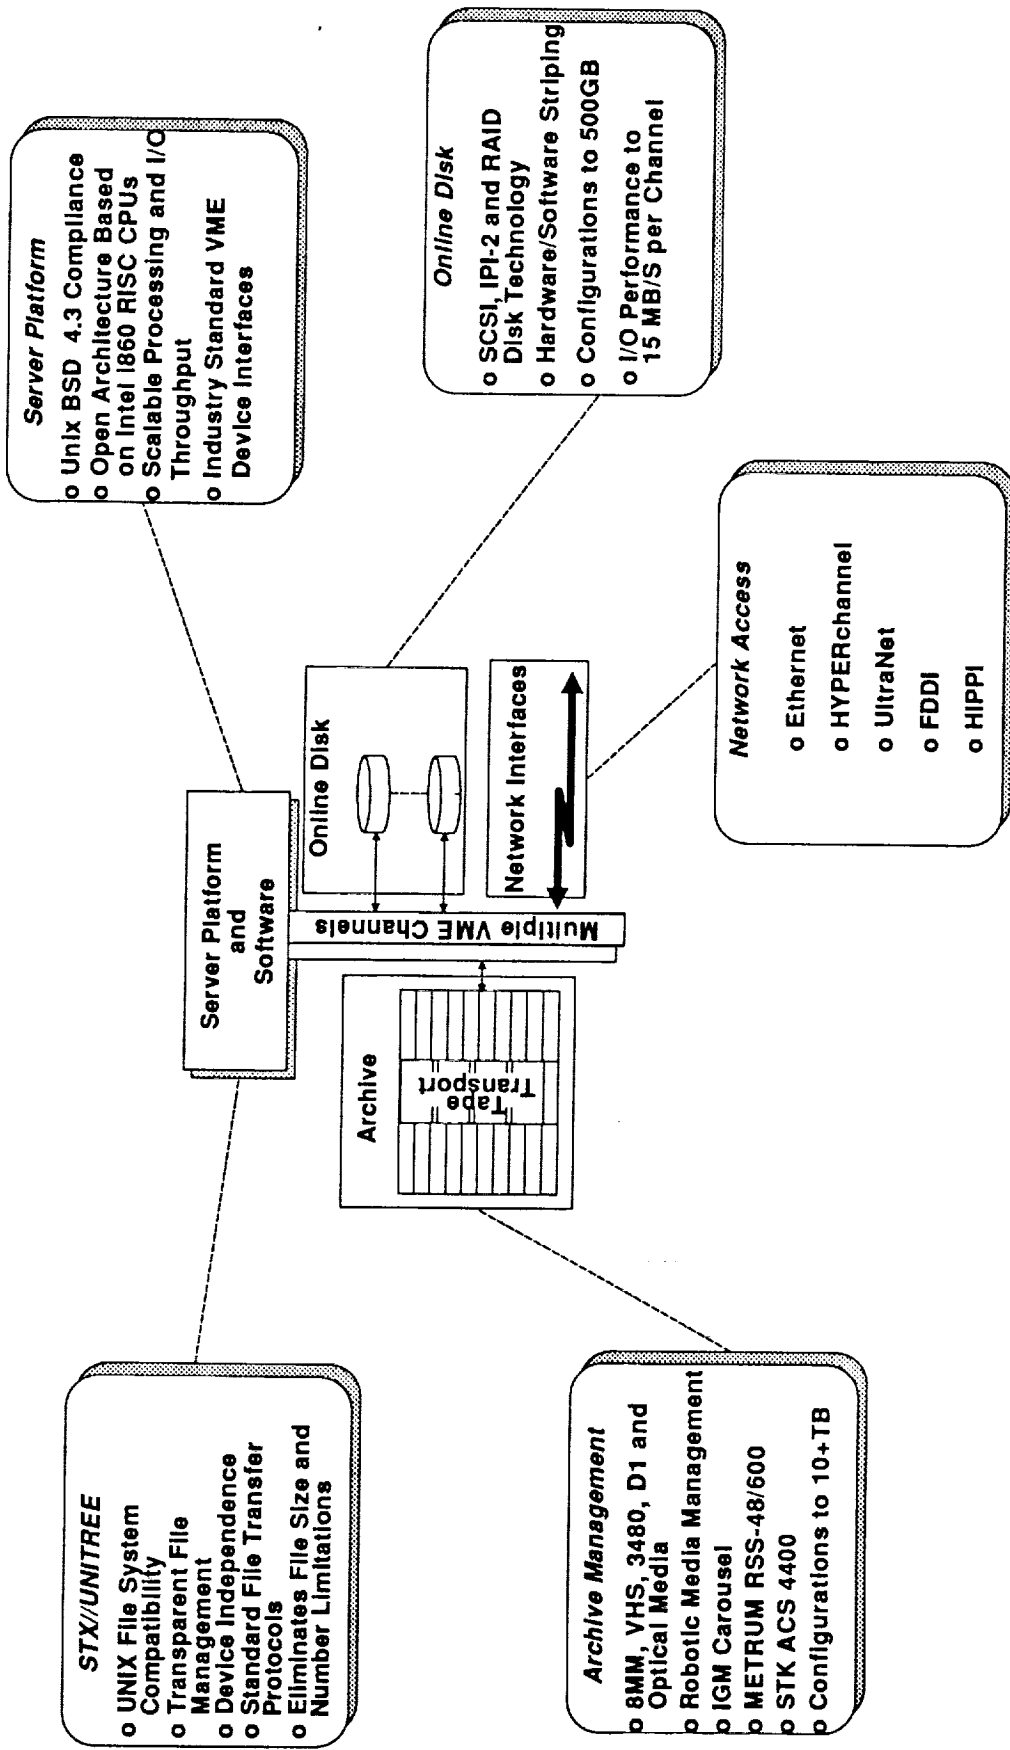
Software

- o Hierarchical Data Management
- o Compliant with IEEE Mass Storage Standard
- o Eliminate UNIX File Size Limitations
- o Transparent File Location

Server Platform

- o Unix Based Operating System
- o POSIX Compliance
- o High I/O Bandwidth
- o High Performance Processing Capacity
- o Accommodate Future Technology

STX//Storage Solutions High Performance File Management Systems



STX//UNITREE

- o UNIX File System Compatibility
- o Transparent File Management
- o Device Independence
- o Standard File Transfer Protocols
- o Eliminates File Size and Number Limitations

Fully Optimized Storage Solutions

- o Easy Expansion
- o High Reliability
- o Low Cost
- o Improved Performance
- o Small footprint
- o Cost Effective Growth
- o Maximized Media Usage


N 9 3 - 1 6 3 3 3

OPEN SYSTEMS STORAGE PLATFORMS

by

**KIRBY COLLINS
CONVEX STORAGE SYSTEMS**

PRECEDING PAGE BLANK NOT FILMED



Open Systems Storage Platforms



Storage Applications

ConvexOS System Software

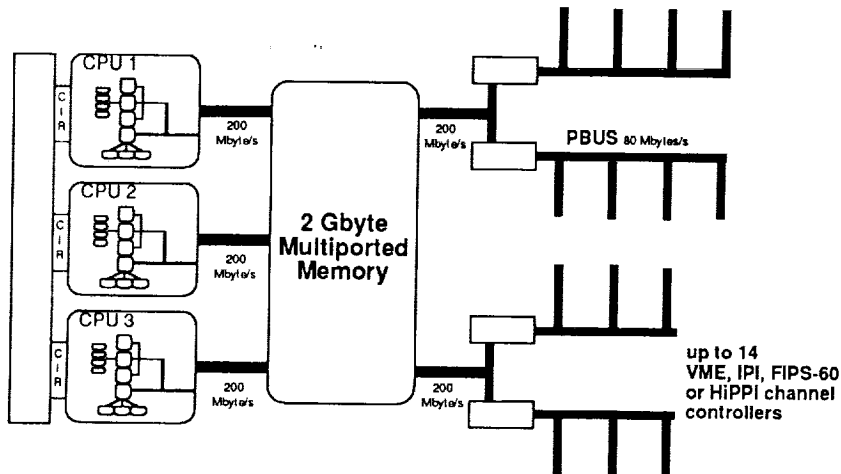
High Performance Peripherals

DS Series Data Server Systems

The building blocks for an open storage system includes a system platform, a selection of storage devices and interfaces, system software, and storage applications

CONVEX storage systems are based on the DS Series Data Server systems. These systems are a variant of the C3200 supercomputer with expanded I/O capabilities. These systems support a variety of medium and high speed interfaces to networks and peripherals. System software is provided in the form of ConvexOS, a POSIX compliant derivative of 4.3BSD UNIX. Storage applications include products such as UNITREE and EMASS.


DS Series Systems



The performance of the DS Series is driven by the main memory system, a multiported arrangement with up to 2 Gigabytes of RAM and a total bandwidth of 800 megabytes/second. A crossbar connects this memory system to five 64bit wide ports, each capable of transferring data at up to 200 megabytes/second. In a C3200 supercomputer, four of these ports are dedicated to CPUs, with the fifth attached to the I/O system. In the DS Series version of the architecture, two memory ports are allotted to I/O and only three CPUs are supported.

The I/O system is designed around a high speed I/O bus called the PBUS. Each memory port supports two PBUS's, for a total of four in the DS Series. The PBUS in turn supports intelligent channels, each of which supports one or more VME, IPI, FIPS-60, or other types of standard I/O busses.

Peripherals - Tape Systems





A range of tape systems is available as well - again based on industry standard interfaces.

A channel controller will soon be available which generates four IPI-3 channels, which can be used to interface such devices as the Ampex R90 DD2 recorder. This device uses a tape format based on video broadcast technology to store from 25 to 150 gigabytes per tape, depending on cartridge size, with transfer rates of up to 15 gigabytes/second.

The TLI generates two FIPS-60 (IBM Block Mux Channel) interfaces, and can be used to connect to such devices as the Storage Technologies ACS (popularly known as the silo) or high duty cycle 3480 compatible tape transports.

Lastly, VME cards can be used to interface to lower speed devices such as nine track tapes, or SCSI based devices such as DAT or medium duty cycle rack mount 3480 compatible drives.

Peripherals - Disk Systems

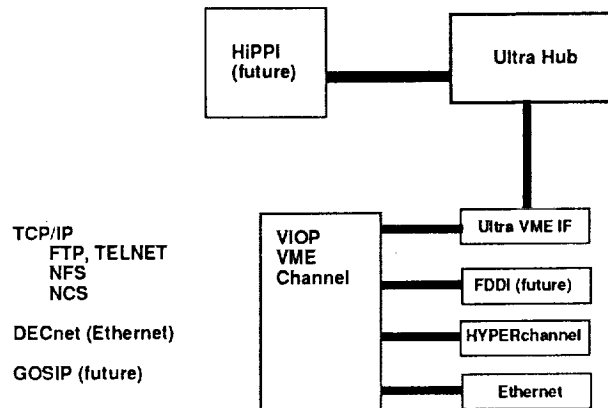


CONVEX supports three basic kinds of disk systems.

High performance IPI-2 disk drives are supported via the IDC, an 88000 based integrated channel controller which generates four IPI-2 interfaces, each capable of transferring data at up to 10 megabytes/second. Each of those interfaces can support up to 8 disk drives, for a total of 32 per IDC. With multiple IDC's more drives can be attached, up to the system maximum of 255. The system is capable of transferring data into these drives at aggregate rates in excess of 50 megabytes per second.

Lower speed drives are supported via the VIOP VME channel, which supports up to two VME busses with 10 megabytes/second of bandwidth each. An SMD-E interface is available along with 1 gigabyte winchester drives, and an ESDI interface is also supported for low cost 780 megabyte drives.

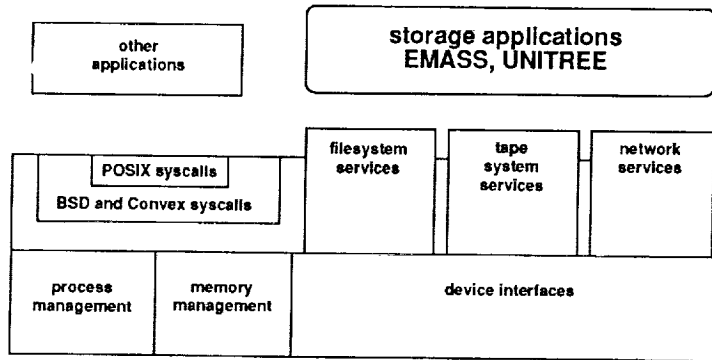
Peripherals - Networks



Network interfaces are provided primarily through VME based interface cards. Controllers are available for Ethernet, HYPERchannel, and Ultra, with FDDI available in the near future. A HiPPI channel is also planned in the near future, to interface to the Ultra hub at high bandwidths, as well as other HiPPI based devices and switches.

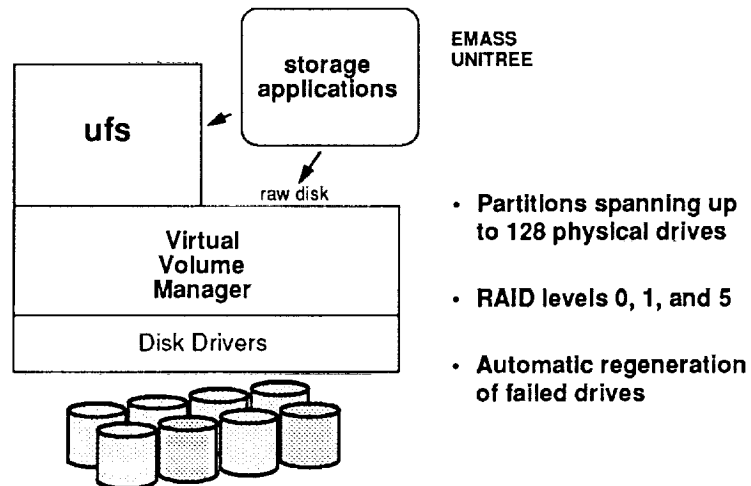
The operating system supports the most commonly used protocols, such as TCP/IP, DECnet (Ethernet only), and GOSIP (in the future).

ConvexOS



The heart of any open systems strategy is the system software, since it is compatibility at the system interface that allows a high degree of interoperability and portability of storage applications. ConvexOS is a derivative of 4.2BSD UNIX, updated with 4.3BSD enhancements, and modified to be compliant with the POSIX.1 standard. Semaphore for symmetric multiprocessing, it uses the parallel architecture of the C3200/DS architecture to provide high throughput under a heavy I/O load. Extensions to the basic UNIX filesystem and tape services provide support for storage applications such as EMASS and UNITREE.

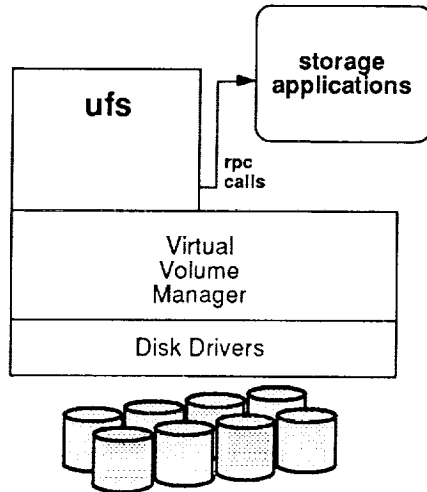
Disk Management



One clear trend in storage is towards large numbers of magnetic disks. To ensure that the connection of large numbers of disk drives does not reduce the overall reliability of a storage system, ConvexOS provides a driver layer called the Virtual Volume Manager (VVM).

VVM takes up to 128 physical partitions on different disks, and combines them into one large partition to the file system and other higher level applications. VVM interleaves successive blocks across different drives to allow parallel access for sequential disk I/O, and can generate mirror or parity blocks so that a failure of any one disk does not cause loss of data. If a drive does fail, VVM can reconstruct data blocks on a spare drive, so that the system can tolerate a drive failure and return to redundant operation, all without interrupting applications that are accessing the partition or requiring operator intervention.

File Fault Interface



- filesystem events can trigger rpc calls to an application
 - file open, read/write
 - metadata update
 - block allocation
 - space watermarks
- allows storage applications to manipulate files transparently to user applications

One of the most difficult challenges in supporting a variety of storage applications is how to tightly integrate these applications with the host system's own file system. Since the host filesystem is usually highly tuned to the architecture, it may make sense to use it for local buffering of files, especially if the host is also used for compute intensive applications as well as storage. Many storage applications have client code that is inserted in the kernel of the host operating system. Providing support for client code for very many storage applications is problematic, since there is little uniformity in these interfaces.

To resolve this ConvexOS provides an interface which allows a storage application to trap many filesystem events, with notification via RPC calls to a daemon process outside the kernel. This allows a storage application to be notified of file creation and access, and migrate data blocks into and out of the filesystem on demand. This activity is completely transparent to applications accessing files through the ConvexOS filesystem, and allows hierarchical storage management to be implemented without modification to applications.


Convex Storage Servers



- **Assembled from off-the shelf components**
- **Industry standard interfaces**
- **Convex engineering for performance and reliability**

With the DS Series of storage systems, Convex has developed a set of products which provide open systems solutions for storage management applications. The systems are highly modular, assembled from off the shelf components with industry standard interfaces. The C Series system architecture provides a stable base, with the performance and reliability of a general purpose platform.

This combination of a proven system architecture with a variety of choices in peripherals and application software allows wide flexibility in configurations, and delivers the benefits of open systems to the mass storage world.



REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 1992	3. REPORT TYPE AND DATES COVERED Conference Publication, July 23-25, 1991		
4. TITLE AND SUBTITLE NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications <i>Volume II</i>			5. FUNDING NUMBERS JON 933-656-80-02-26	
Ben Kobler, P. C. Hariharan, and L. G. Blasso				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA-Goddard Space Flight Center Greenbelt, Maryland 20771			8. PERFORMING ORGANIZATION REPORT NUMBER 92B00093 Code 933	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, D.C. 20546-0001			10. SPONSORING/MONITORING AGENCY REPORT NUMBER CP-3165, Vol. II	
11. SUPPLEMENTARY NOTES Kobler: Goddard Space Flight Center, Greenbelt, MD; Hariharan and Blasso: STX Corporation*, Lanham, MD. *Hughes STX Corporation as of October 1, 1991				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 82			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report contains copies of nearly all of the technical papers and viewgraphs presented at the NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications. This conference served as a broad forum for the discussion of a number of important issues in the field of mass storage systems. Topics include magnetic disk and tape technologies, optical disk and tape, software storage and file management systems, and experiences with the use of a large, distributed storage system. The technical presentations describe, among other things, integrated mass storage systems that are expected to be available commercially. Also included is a series of presentations from Federal Government organizations and research institutions covering their mass storage requirements for the 1990s.				
14. SUBJECT TERMS Magnetic tape, magnetic disk, optical disk, mass storage, software storage			15. NUMBER OF PAGES 269	
			18. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	16. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	