

1995117710

Optical Storage Media Data Integrity Studies

N95-24130

Fernando L. Podio

National Institute of Standards and Technology

Building 225, Room A61

Gaithersburg, MD 20899

fernando@pegasus.ncsl.nist.gov

301-975-2947

301-216-1369 (fax)

43466

p. 12

Abstract

Optical disk-based information systems are being used in private industry and many Federal Government agencies for on-line and long-term storage of large quantities of data. The storage devices that are part of these systems are designed with powerful, but not unlimited, media error correction capabilities. The integrity of data stored on optical disks does not only depend on the life expectancy specification for the medium. Different factors, including handling and storage conditions, may result in an increase of medium errors in size and frequency. Monitoring the potential data degradation is crucial, especially for long term applications. Efforts are being made by the Association for Information and Image Management Technical Committee C21, Storage Devices and Applications, to specify methods for monitoring and reporting to the user medium errors detected by the storage device while writing, reading or verifying the data stored in that medium. The Computer Systems laboratory (CSL) of the National Institute of Standards and Technology (NIST) has a leadership role in the development of these standard techniques. In addition, CSL is researching other data integrity issues, including the investigation of error-resilient compression algorithms. NIST has conducted care and handling experiments on optical disk media with the objective of identifying possible causes of degradation. NIST work in data integrity and related standards activities is described.

Introduction

Many organizations are using optical disk-based systems for the storage and retrieval of large sets of valuable information. One general indicator for long term storage of data is the optical disk media life expectancy. For this indicator to be of value, a standard method to determine life expectancy is essential. Extrapolated life expectancy values may vary greatly because they depend on the test method used for calculating the quality parameter (e.g. the byte error rate), the measurement approach (including areas on the disk tested, data patterns written, and amount of data tested), the mathematical model used, and the criteria for data analysis (including the statistical analysis used and the confidence levels), Podio [1].

If a standardized test were employed by all media manufacturers, media life expectancy could

be a good parameter to select media for long term applications. However, to determine a life expectancy specification, tests are run with a small sample of disks from a population of manufactured disks. In addition, media technological changes would require running life expectancy tests almost continuously to test newcomers on the market, since the old data obtained on previous life expectancy tests may not apply to new technology. In conclusion, a life expectancy specification is useful only as a general indicator for media selection. Individual disks will still fail at different times.

All storage devices are designed with powerful, but not unlimited, error correction capabilities. Because of different factors which include handling and storage conditions, errors may increase in size and frequency. If the level of errors increases beyond the maximum capacity of the ECC (error correcting codes) in the device, data will be uncorrectable. By not being informed of the level of error correction that is taking place in the optical disk device, users learn about critical error events only when the data is already irretrievable. If these types of critical errors occur on a specific unlinked data structure, this data structure may no longer be recoverable but data degradation may not have caused extensive damage outside the unlinked data structure. However, if data degradation at these critical levels of unrecoverable errors occurs in any linked structure or with a compressed data entity, substantial data losses may result.

Several approaches can be followed for improving data integrity. One approach is to monitor data errors with time. Users can gather information to highlight trends in particular selected disks or their entire data sets. This monitoring capability allows users to make decisions on transferring data to new media in a timely and economic manner before data loss occurs. Another method to increase data integrity is to use layered ECC. Although layered error correction decreases the user data capacity, it adds error resilience.

Compressed data in the presence of errors is especially vulnerable to catastrophic data failure. Woolley [2] emphasizes the importance of robust error control in data compression applications. For compressed data, in addition to using media error monitoring and layered ECC, there are other techniques to improve data integrity in the presence of errors: error correction integrated with data compression Kobler [3], entity reduction and error-resilient compression. NIST is currently investigating the error-resilience of these techniques.

Efforts to develop standard media error monitoring tools and techniques for optical disk drives and NIST's involvement in the development of this standard are described. NIST investigations on media error monitoring tools, data analysis statistical models for error distribution, and media error visualization are also described.

Another aspect of data integrity research at NIST includes an experimental program for the care and handling of optical disks. A series of experiments were performed using different types of optical disks. The experiments included exposure to liquids and vapors, cleaning agents, solvents, fire smoke, food substitutes, paint fumes and paint, temperature and humidity cycles, heat and cold shocks, uniform pressure, static electricity, gamma rays, etc.

A brief description of this work is also included.

Standards for Media Error Monitoring and Reporting Techniques

In 1991, the Computer Systems Laboratory of NIST sponsored a workshop to identify the state of the art on media error monitoring approaches for optical disks and to identify the user's needs, Podio [4]. As a result of the workshop, a working group was formed to identify media error monitoring techniques. The working group developed a set of procedures for monitoring and reporting media error correction levels on optical disk devices. The results of this activity are being used as a basis for formal standard work.

With NIST leadership, the Association for Information and Image Management (AIIM) Committee C21 (Storage Devices and Applications) is developing the American National Standard ANSI/AIIM MS59. ANSI/AIIM MS59 specifies media error monitoring and reporting techniques for the verification of information stored on optical digital data disks¹, AIIM C21 [5].

Parallel efforts are taking place in developing an international standard under the auspices of the International Standards Organization (ISO) Technical Committee TC 171, Micrographics and Optical Memories for Document and Image Recording, Storage and Use. The current content of the proposed ISO standard is based on ANSI/AIIM MS59.

ANSI/AIIM MS59 provides a toolkit of media error monitoring and reporting techniques, any combination of which may be employed. The standard provides two levels of media error monitoring and reporting techniques, a functional approach and an implementation of a selected set of Small Computer Interface (SCSI-2) commands.

The high level approach (a set of functional commands) is independent of the host operating system (e.g. DOS, Unix, OS/2, etc) and the interface that connects the optical disk device with the host (e.g. SCSI-2, IPI, LAN, etc). This high level interface approach is media type and size independent. That is, it can be used with systems that use WORM (write-once read many), rewritable or partially read-only media and optical disk drives for different media sizes from 90 mm to 356 mm media. The implementation of a selected set of SCSI-2 commands enables media error monitoring and reporting techniques at the device level providing direct communication with an optical disk drive that uses the SCSI-2 interface.

¹

The U. S. National Archives and Records Administration (NARA) has recently published a Technical Information paper NARA [6]. NARA's publication provides recommendations on long-term access strategies for Federal Agencies using digital-imaging and optical digital disk storage systems. One of NARA's recommendations on data integrity states that users should "require that equipment conform to the proposed national standard ANSI/AIIM MS59".

The media error information that can be obtained using the tools specified in the standard include:

- A list of reallocated sectors.
- Corrections that exceed some specified media error levels.
- Warning on specified verify media error levels.
- Total number of bytes in error, number of bytes in error per sector and maximum number of bytes in error in any sector codeword.
- The uncorrected or corrected sector content.
- Errors encountered reading header information such as the sector address, sector marks, and synchronization signals.
- The maximum length of contiguous defective bytes.

From the user's perspective, the purpose of ANSI/AIIM MS59 is to allow users of the standard:

- To have a better understanding of the status of their data stored on optical disks.
- To obtain media error information as directed by the system administrator.
- To enable data recovery with tools of the desired level of sophistication.
- To provide media error information allowing the user to make decisions about the media at the present time, and also provide error information which will highlight trends in particular selected disks or in their entire data sets.
- To make decisions about how long the media can be used without an unacceptable risk of data loss.
- To develop more cost effective backup, copying and data transfer policies.

The user or implementor of ANSI/AIIM MS 59 will be able to:

- Format the optical digital data disks with or without certification.
- Reallocate sectors when specified media error levels are exceeded.
- Obtain information about all the reallocated sectors and/or a defect list of initial media defects.
- Set media error level values to obtain early warning information about the status of the data and/or interrogate the drive to obtain the values of those set media error levels.

The **media error levels** are what the optical disk drive will use for error recovery. If the ECC level of correction exceeds one or more of the set levels and reallocation is enabled, the sector that exceeded the media error level(s) is reallocated to a spare sector. Whether reallocation is enabled or not, the optical disk drive reports to the host that a set level was exceeded, indicating which one was exceeded, and whether or not the data was recovered.

The following are the media error levels specified in ANSI/AIIM MS59:

- Number of bytes in error per codeword
- Number of bytes in error per sector
- Number of bad sector IDs
- Number of missing resync bytes (when the drive uses resync bytes)

AIIM C21 is also developing an accompanying ANSI Technical Report, AIIM C21 [7] describing guidelines for the use of the media error monitoring techniques documented in ANSI/AIIM MS59. The current outline for the guidelines includes:

- A description and use of the media error monitoring and reporting techniques documented in ANSI/AIIM MS59.
- Discussion of error management strategies.
- Methods of visualization of media error information using the techniques specified in MS59.
- Methods to estimate the data integrity including:
 - use of sampling methods
 - use of baseline media error parameters and distributions
 - use of statistical models

NIST Investigation of Media Error Monitoring Tools for Optical Disks and Media Error Visualization

Concurrently with the standardization efforts, NIST has also been conducting laboratory work in investigating media error monitoring and reporting (MEMR) techniques, statistical models for error distribution, and methods for error data visualization.

The **MEMR techniques** are used in optical disk drives for the verification of information stored on the optical disks. These techniques allow users to obtain timely information about the status of their data. NIST has investigated some of the MEMR techniques available in commercial drives, and has researched possible new implementations. All of this work has contributed to the content of the proposed ANSI/AIIM MS59 standard and the parallel proposed ISO standard.

NIST has also developed guidelines for the use of the MEMR tools. The guidelines include procedures that end users or system integrators can use to monitor the status of data stored on optical disks. These MEMR techniques for optical disk drives may be the basis of similar MEMR techniques for other types of high density/high capacity mass storage technologies, such as magnetic media disk/tape drives, optical tape drives and devices based on new page-oriented memories.

NIST has looked at **statistical models for media error distributions** on optical disks. One model is the modified Gilbert model, Marchant [8]. This model is based on two different classes of defects and has been found to give an excellent fit to defect statistics on media that

uses the magneto-optical recording Takeda, Saito, and Itao [9]. This model takes into account modeling non random errors (long burst defects). The Gilbert model requires two byte error rate (BER) values and two average burst lengths. One BER is derived from microscopic defects, the other from larger media substrate damage. The model output is the burst-length probability distribution.

A simpler statistical model has also been developed at NIST. This model is the basis for predicting the maximum number of errors in a sector codeword. One version of this model assumes a uniformly random binomial distribution of errors, and uses only one byte error rate (BER) and the number of sector interleaves as input. It produces a baseline for comparison at different times on the status of an optical disk every time the disk is tested. The output of this model can provide a basis for comparison with reported disk error statistics so that the user can identify abnormal changes in the media error distribution. Figure 1 shows the model output, a distribution of bytes in error per codeword.




Figure 1. For this particular disk type, the sector data field, which includes the user data bytes, the ECC bytes, and the CRC bytes, is divided into five codewords. The maximum number of bytes in error of the five codewords per sector is determined, and the probability of occurrence of one, two or more bytes in error is plotted. The maximum number of errors per codeword that can normally be corrected in drives that use this type of media is eight.

Depending on the different values of the expected or empirical BER, the model can also indicate that the drive's maximum error correcting capabilities are being approached or exceeded. The modified Gilbert model describes only error length distributions, but does not make the connection to error correcting capabilities. This simpler model makes this link. It uses the real disk data structure (sectors and interleaves) and assumes only a simple uniformly random distribution of errors, such as binomial or Poisson.

Given a measured BER, the model can be used to generate a predicted distribution of the

number of bytes in error per codeword. If the predicted distribution is not acceptable (number of bytes in error in any codeword exceeds certain user established number), the user may consider retiring the media at this point. If the predicted distribution is acceptable, the user should determine the real distribution of the number of bytes in error per codeword (using a ANSI/AIIM MS59 compliant drive or any other drive that provides this type of media error information). If the measured distribution shows an excess in the number of bytes in error, the user may consider retiring the disk. Because the model assumes a random distribution of errors, if the predicted and measured distributions are significantly different, a non random error distribution might be suspected and the user may use other MEMR tools to investigate the level and the distribution of these errors further.

NIST has also developed **media error visualization tools** for the byte-error statistics retrieved via the media error reporting tools documented in the ANSI/AIIM MS59 standard. The media error visualization tools include:

a. Line graphs depicting:

- Sector reallocations over time.
- Bad sector ID's over time.
- Byte error rate over time.
- Maximum bytes in error per codeword and per sector over time for a given disk as shown in Figure 2.




Figure 2. This line graph shows how the maximum, or worst-case, number of bytes in error per codeword may change over time for a given disk. The normal ECC correction limit capability is eight bytes per codeword.

In Figure 2, the user-specified error level is set at a maximum of four bytes in error per codeword.

Using ANSI/AIIM MS59 compliant drives the user can check the default level of bytes in error per codeword that, when exceeded, would enable the reallocation of the sector. Using this type of drive, the user can also change this error level.

b. Bar charts depicting:

- Relative frequency of maximum bytes in error per codeword and per sector.
- Maximum bytes in error per codeword per radial area of a disk.
- Maximum bytes in error per codeword per band of tracks.

The bar chart in Figure 3 shows the maximum number of bytes in error per codeword per band of tracks for a given disk.




Figure 3. The maximum number of bytes in error per codeword is shown in this chart in a different way than in Figure 2. The disk has been divided into bands of n tracks and the maximum number of bytes in error per code word is shown for all the sectors within the tracks specified in a particular band. Dividing the disk into bands of n tracks enables visualization of the entire disk from the inner area to the outer area.

c. Three-dimensional histograms depicting:

- Maximum bytes in error per sector over the disk.
- Maximum bytes in error per codeword over the disk as shown in Figure 4 in a three-dimensional histogram.

The media error information conveyed in this Figure may be, in many cases, sufficient for most users. As in Figure 3, the disk has been divided into bands of n number of tracks (31 bands in this Figure).




Figure 4. The maximum number of bytes in error per codeword over the disk is shown in this Figure using a three-dimensional representation. The maximum number of bytes in error per codeword is shown for all the sectors within the tracks specified in a particular band. Only one band shows sectors that have codewords with more than one byte in error. The ECC correction capability for drives compatible with this type of disk is eight bytes in error per codeword. The Figure shows a fairly healthy disk.

By using ANSI/AIIM MS59 compliant drives or other drives that provide similar media error information, the user can typically obtain from the drive information on media errors with the desired level of detail and sophistication.

When the user wants to analyze media errors in specific disks, media error visualization charts can be used. However, when users want to apply media error monitoring tools to a large number of disks, plotting error distributions or other statistics might be impractical. In this case, the user should access the required information numerically and take the appropriate action. For example, if the users set the media error levels, they may decide to transfer data to another disk when these levels are exceeded. More information about these procedures or how to use the visualization tools is provided in AIIM C21 [5] and [7].

Care and Handling Experiments for Optical Digital Data Disks

NIST has conducted care and handling experiments with the objective of identifying possible causes of data degradation in optical disk media. In order to conduct these experiments NIST developed optical disk media measurement systems capable of determining data degradation parameters such as the byte error rate, and error distributions. Figure 4 was based on one of these error distributions. A measurement system for mechanical characteristics of optical disk media was developed by Nevenzel and Voogel [10].

Optical disks may deteriorate if subjected to some unusual conditions such as extreme temperature and humidity, temperature and humidity cycles, and high energy radiation. Some office cleaning substances and other components like tobacco smoke, liquids and food may also produce data degradation. Such degradation effects were investigated through care and handling experiments. The approach that was followed consisted of writing a selected number of sectors of the disks and reading them back checking the bytes in error and the error distribution. The information that is derived is the byte error rate (BER), which gives an average measure of the number of bytes in error, the defect distributions, and the location of burst errors. Some mechanical measurements were also performed. For CD media testing, which included CD-ROMs and CD-R's (CD recordable media), a commercially available tester was used. NIST is currently analyzing the test results.

The experiments included:

- Cleaning agents immersion tests and vapor and gas exposure.
- Fire smoke exposure and exposure to chemicals used in fire extinguishers.
- Exposure to food substitutes.
- Exposure to paint and wax fumes.
- Temperature and humidity cycles; cold and heat shocks.
- Mechanical experiments such as impact and uniform pressure.
- Human interaction such as scratches, permanent inks, hand creams and bending experiments.
- Electromagnetic exposure such as magnetic fields, gamma-rays, X-rays and electrostatic discharge and sun light.
- Exposure to possible harmful liquids such as gasoline and diesel.
- Read, write and erase cycles.

A complete description of the measurement procedures, the experiments conducted and the test results are to be included in a NIST report that is currently being prepared for publication.

Current and Future Plans

NIST will continue to investigate techniques to improve data integrity in the presence of

media errors. Work will continue in the development and analysis of statistical models for error distribution and visualization tools. The work will also include the investigation of data integrity of storage media using layered ECC, and the investigation of emerging techniques for data compression including entity reduction and error-resilient compression. In addition, there is interest from both users and industry in extending the work done on data integrity for optical disks to other emerging storage device technologies.

Conclusions

A life expectancy specification for optical disks (and other types of storage media) is useful only as a general indicator for media selection. But it cannot be the only indicator for assuring data integrity in optical disks. Individual disks may still fail at different periods of time. Work on standardizing media error monitoring and reporting techniques for the verification of information stored on optical disks systems is ongoing. These techniques provide users with a better understanding of the status of the data stored on optical disks. Users can then make decisions about the media at the present time, identify trends and develop more cost-effective backup, recopying and data transfer policies. Without access to media error information, the level of media errors may increase beyond the maximum capacity of the error correcting codes in the device. In this case, data will be uncorrectable.

The level of errors may increase because of factors such as improper handling or storage conditions. Understanding the data integrity of data structures on a disks is important. Compressed data, in the presence of uncorrectable errors, is especially vulnerable to catastrophic data failure. Error-resilient algorithms and other techniques such as layered ECC may increase the chances of data recovery in the presence of uncorrectable media errors. The need to investigate media error monitoring and reporting techniques to verify data integrity on optical disks and other emerging storage media technologies is apparent. Investigation of error-resilient compression techniques is also needed.

Acknowledgements

The work described in this paper is based on the efforts of several people at NIST. I want to acknowledge my colleagues for their input, support and work in this area. In particular, I want to acknowledge the following individuals: Sandra Woolley and Thierry Gouverneur for their work in media error monitoring, data visualization and sampling plans, Sam Silberstein for the development of a statistical model for error distributions, and Stefan Leigh, James Filliben, and other personnel from the Statistical Engineering Division of the Computing and Applied Mathematics Laboratory of NIST for their useful advise and cooperation. Many individuals participated in the care and handling experiments and I want to acknowledge their participation and support in that work. I also want to acknowledge the members of AIIM C21 for their contribution to the ANSI/AIIM MS59 standard and the guidelines. Special thanks are due to Charles Obermeyer from the U.S. National Archives and Records Administration (NARA) and David Patton from IBM, Tucson, AZ. The work described in this paper is

carried out at NIST, under Interagency agreements with several U.S. Federal Government agencies including the U.S. National Archives and Records Administration, the Social Security Administration and the Federal Bureau of Investigation.

References

- [1] Fernando L. Podio, "Development of a Testing Methodology to Predict Optical Disk Life Expectancy Values", NIST Special Publication 500-200, December 1991.
- [2] Sandra I. Woolley, "The Importance of Robust Error Control in Data Compression Applications", Third NASA Goddard Conference on Mass Storage Systems and Technologies, 1993.
- [3] Ben Kobler, "Techniques for Containing Error Propagation in Compression/Decompression Schemes, Space and Earth Science Data Compression Workshop, 1991.
- [4] Fernando L. Podio, "Monitoring and Reporting Techniques for Error Rate and Error Distribution in Optical Disk Systems", Results of a Workshop, NIST Special Publication 500-198, August 1991.
- [5] AIIM Technical Committee C21, Storage Devices and Applications; Proposed ANSI Standard ANSI/AIIM MS59: "Media Error Monitoring and Reporting techniques for Verification of the Information Stored on Optical Digital Data Disks", Fifth Draft, November 1994.
- [6] The National Archives and Records Administration, "Digital-Imaging and Optical Digital Data Disk Storage Systems, Long-Term Access Strategies for Federal Agencies", Technical Information Paper No. 12, July 1994.
- [7] AIIM Technical Committee C21, Storage Devices and Applications; Proposed ANSI Technical Report ANSI/AIIM TR39: Guidelines for the Use of Media Error Monitoring and Reporting techniques for Verification of the Information Stored on Optical Digital Data Disks", Second Draft, November 1994.
- [8] Alan B. Marchant, "Optical Recording, A Technical Overview", Addison-Wesley Publishing Company, 1990.
- [9] Takeda, T., M. Saito, and K. Itao, "System Design of 90 mm Optical Microdisk Subsystem", SPIE Proceedings, 899, 16, 1988.
- [10] Gerhard Nevenzel, Martin Voogel; "Mechanical Deformation Measurements for Optical Data Disks", NIST IR 5208, February 1993.