N94- 33797

Certification of *ICI 1012* Optical Data Storage Tape

J. M. Howell

Senior Engineer, ICI Imagedata
Brantham, Manningtree, Essex CO11 1NL, England.
Phone 044-206-392424 Ext 6432, Fax 044-206-391472

Introduction

ICI has developed a unique and novel method of certifying a Terabyte optical tape. The tape quality is guaranteed as a statistical upper limit on the probability of uncorrectable errors. This is called the Corrected Byte Error Rate or CBER, and is defined below.

We developed this probabilistic method because of two reasons why error rate cannot be measured directly. Firstly, written data is indelible, so one cannot employ write/read tests such as used for magnetic tape. Secondly, the anticipated error rates need impractically large samples to measure accurately (Smythies and Woodley [1]). For example, a rate of 1E-12 implies only one byte in error per tape.

The archivability of *ICI 1012* Data Storage Tape in general is well characterised and understood; see for example Ruddick [2]. Nevertheless, customers expect performance guarantees to be supported by test results on individual tapes. In particular, they need assurance that data is retrievable after decades in archive. This paper describes the mathematical basis, measurement apparatus and applicability of the certification method.

Tape format and error correction

See figure 1. Data is stored as records written transversely across the tape. Each record on the tape is built up from 1024 codewords. Each codeword contains 64 bytes of user data and 16 error correction code (ECC) bytes. The ECC algorithm is a Reed-Solomon code which can completely correct up to 8 defective bytes within a codeword. If there were 9 bytes or more, then the codeword would be flagged as uncorrectable and the regenerated data would very likely contain errors.

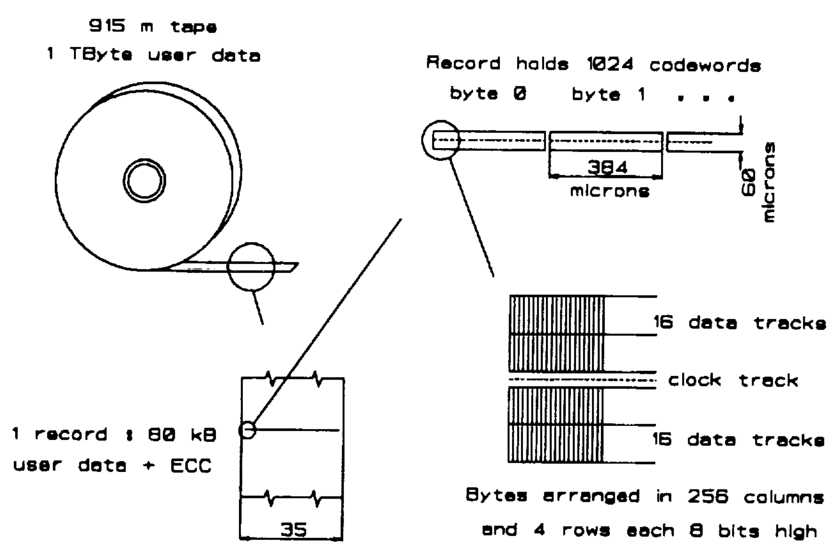

Figure 1. Data format on *ICI 1012* Data Storage Tape

The CBER is the probability that a byte regenerated from the raw data in any codeword is corrupt. Because of the ECC, and the way the data is packed in each record, the number of defective bits in each byte does not affect the ECC; it is thus appropriate to measure errors in bytes rather than bits. Statistically, the number of errors in the data increases with the raw byte error rate, or BER.

The ECC works on expanded codewords 255 bytes long consisting of the 80 user and ECC bytes plus 175 bytes filled with zeroes. The algorithm generates 8 correction bytes and 8 correction addresses in the expanded codeword. With more than 8 corrupt bytes in the original codeword, the corrections and addresses are wrong, and there is a high probability that some of the 175 padding bytes will be toggled non-zero. When this happens ECC breakdown is detected and the recovered data is issued as it was read.

Origin of Errors

Optical tape is written by short bursts of intense light from a laser which reduces the dye-polymer recording layer thickness, making it appear dark. The data is read by another lower power laser which discriminates the dark and light regions. Errors arise if either a written region appears bright or an unwritten region appears dark. There are many potential sources of errors, but for media certification we are only interested in those from physical anomalies in the recording layer.

Bright spots occur when there is a break in the dye-polymer coating so that the underlying aluminium alloy reflector is exposed. Alternatively, the coating is so thick that the write laser cannot 'punch through' to sufficient depth. Dark spots are generally due to scratches and debris which scatter the read laser light. They can also be seen if the dye-polymer layer is thinner than normal, so that it is comparable in thickness to a written area.

Calculating the CBER from measurable properties

For reasons given above, one must use statistical probability to compute CBER. This section starts with the general case and develops a practical tool for deriving the CBER from the background defects and observable point defects.

1. General case.

The probability of a given codeword being corrupt is the sum of the probabilities of that codeword containing y=9, 10, .. 80 corrupt bytes. The CBER for this codeword is the sum of the products of

i) the probability $p_{y,80}$ of a codeword having y corrupt bytes before correction, and
ii) the fraction y/80 of corrupt bytes in the codeword (since correction is inhibited).

$$CBER = \sum_{y=9}^{80} p_{y,80} \cdot \frac{y}{80}$$

2. Effect of background defects.

If the defective bytes are completely random, then the binomial distribution can be used to compute the $p_{y,80}$ from the probability p of any individual byte in the codeword being corrupt. The fraction p is the byte error rate, or BER.

$$p_{k,n} = \frac{n!}{k!.(n-k)!} \cdot p^k.(1-p)^{n-k}$$

It was shown in Howell [3] that the CBER calculated in this way is the most pessimistic value; if the defective bytes are not random, i.e. they are clustered in the record, then the CBER will be less. It represents an upper bound on CBER when the raw BER in a record measured by the drive is used directly in the binomial distribution.

In practice the CBER estimated from data from large scale experimental trials [3] is substantially less than the upper bound. A lot of perfectly serviceable material would fail if this simple test was applied indiscriminately. The reason is clear when the tape surface is examined more closely. A relatively small peak in BER would be amplified by the non-linear binomial expression so it dominates the CBER result. However we find that these peaks are not showers of defects spread over the tape width. They are compact clusters caused by well-defined circular or 'point' defects. Such defects corrupt nearly every byte within their boundary, but do not affect those outside.

A practical certification method must take the morphology of these point defects into account. As the tape is not preformatted it is impossible to know exactly which codewords will be affected. We can still obtain an important statistical CBER estimate from the size alone. This will now be derived.

3. Effect of individual 'point' defects.

Consider codewords within a single record containing a defect which is exactly one byte block (384 microns) wide. This defect is superimposed on the random background errors. Every codeword in the record will have at least one corrupt byte from the defect. The probability of 9 corrupt bytes in total is then the binomial probability of 8 out of the remaining 79. For a defect which is 2 byte blocks wide, we must find the binomial probability of 7 out of 78, and so forth. In general, consider a 'point' defect which is (w+w') blocks wide, where w' is the fraction and w is the integer part. The CBER for a record containing this defect is given by the following formula:

$$\begin{aligned}\mathrm{CBER} &= \sum_{y=9}^{80}\left(w'.p_{y-w-1,80-w-1} + (1-w').p_{y-w,80-w}\right).\frac{y}{80}\\ &= w'.\sum_{y=9}^{80} p_{y-w-1,80-w-1}.\frac{y}{80} + (1-w')\sum_{y=9}^{80} p_{y-w,80-w}.\frac{y}{80}\\ &= w'.B_{w+1} + (1-w').B_w\end{aligned}$$

where B_W is the cumulative binomial expansion from background errors for the CBER of a codeword which has w bytes within a 'point' defect:

$$B_w = \sum_{y=9}^{80} p_{y-w,80-w}\frac{y}{80}$$

4. Effect of multiple 'point' defects.

If several 'point' defects occur in the same record, their combined effect is calculated from the probability of the defects coinciding in any codeword. Consider two defects each less than one block wide of sizes w' and v' blocks. Assume they are randomly placed in the record. The total CBER will be the sum of the contributions from codewords with none, one or two bytes affected, as follows:

$$\begin{aligned}\mathrm{CBER} = &(1-w').(1-v').B_0\\ &+((1-w').v'+w'.(1-v')).B_1\\ &+(v'.w').B_2\end{aligned}$$

The calculation can be developed iteratively to include any number of such defects, of any width (w+w') blocks. Let d_i^n be the probability of a codeword having exactly i corrupt bytes from n 'point' defects. Imagine a codeword which is initially free of point defects. Here $d_0^0 = 1$ and $d_1^0 \ldots d_{80}^0 = 0$. We include the effect of each defect which might affect the record in turn. For the k'th defect the vector $\mathbf{d}^k$ is generated from its predecessor $\mathbf{d}^{k-1}$ and defect width (w+w') thus:

$$d_i^k = 0, \text{ for } i < w$$

$$d_i^k = (1 - w').d_{i-w}^{k-1}, \text{ for } i = w$$

$$d_i^k = (1 - w').d_{i-w}^{k-1} + w'.d_{i-w-1}^{k-1}, \text{ for } i > w$$

The expression for the CBER in a record with n point defects is then the sum of the contributions from the codewords with 0, 1, 2, ... 80 bytes corrupted by 'point' defects.

$$\text{CBER} = \sum_{i=0}^{80} d_i.B_i$$

5. Practical implementation.

The vector $\mathbf{d} = (d_0, d_1, \ldots d_{80})$ is a complete description of the effect of 'point' defects in a record on the CBER for that record. The vector $\mathbf{B} = (B_0, B_1, \ldots B_{80})$ is only a function of the background BER. One can therefore compute **d** as an arithmetic mean over all records in a region of tape (say 1 metre) and use it in the above expression for an average CBER. If the BER is invariant along the tape the vector **B** will be constant for all regions. Thus one has a computationally efficient method of predicting the CBER profile for a tape from the sequence of **d** vectors and the background BER.

Measuring Point Errors

'Point' defects are regions of the tape which introduce substantial clusters of errors, and are due to irregularities in the active layer. A fully automatic system is required which will measure these in line for every tape pancake (figure 2).

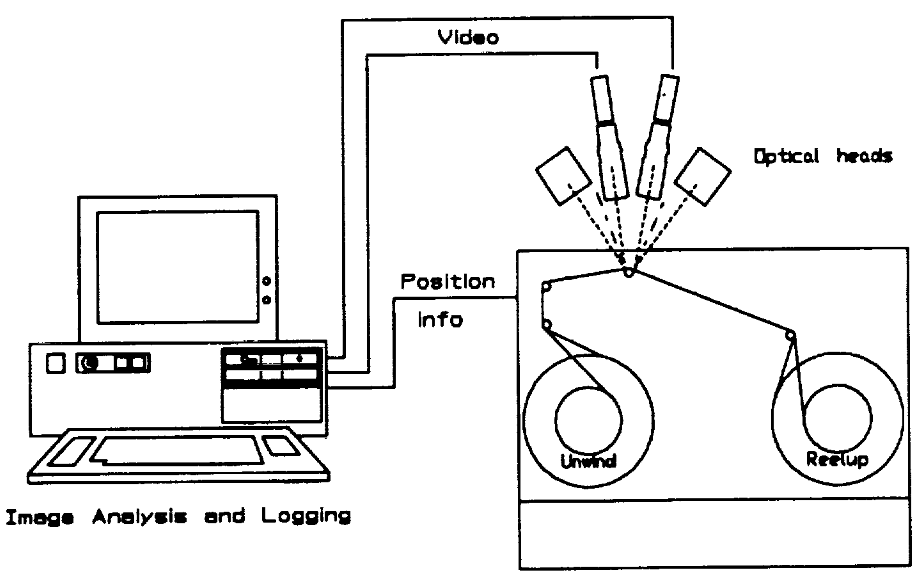

Figure 2. Automatic inspection for *ICI 1012* optical tape

The system uses a pair of line scan cameras which observe the tape in the near infra-red as it passes at constant speed. Defects show up in both cameras, but the wavelengths have been carefully chosen to measure the reflectivity in the unwritten and written states. The latter measurement is clearly an indirect one. The optics incorporate many innovative features to reduce sensitivity to other disturbances, such as tape movement. This means that the system can be made very sensitive to variations in coating which are known to affect read/write performance.

The cameras will resolve to 20 microns at the tape surface. This means that any defect can be measured to 5% of a byte block width. The software analyses the image in real time, including correlation of the two camera images to avoid double-counting of defects. The codeword defect vector **d** is built up for each transverse scan which approximates to a record width. The cumulative codeword defect vector may then be used to compute the CBER directly for each metre run of tape, as described above. This is also a traceable record for quality assurance.

The software further analyses the image to classify the physical form of the defects. Thus, scratches, thin and thick dye-polymer, debris and exposed reflector are individually logged and displayed. This map then becomes a permanent record of every tape which is produced, so that in the unlikely event of problems the product is fully supportable.

Conclusions

The method described is a non destructive, statistically valid prediction of the CBER for certification of individual tapes. It is flexible in three respects.

Firstly, it is not a simple pass or fail quality check. Different applications will have different tolerances to uncorrectable errors. One inspection provides all the information needed to determine the suitability of a tape. It is then possible to grade tapes if the need arises.

Secondly, the CBER is known along the whole length of the tape pancake. This means that cutting positions can be optimised to minimise waste when removing any substandard regions of the pancake.

Finally, the defect vector is a compact record which is independent of subsequent changes in service[1]. The CBER can be recalculated easily and accurately for any background BER. This could be BER measured by the tape drive after writing. Equally it could be BER at some point in the future as estimated by models from accelerated ageing experiments.

References

1. Smythies, DC and Woodley, BR. Bit Rate qualification Criteria for ICI Optical Media, Creo Products Inc. internal report, May 1991.

2. Ruddick, AJ. ICI Optical Data Storage Tape - an archival mass storage medium, Goddard Conf. on Mass Storage Systems and Technologies, NASA Goddard Space Flight Centre, September 1992.

3. Howell, JM. The Effect of Defect Distribution on Error Correction in 1012 Tape, ICI Imagedata internal report, September 1992.

[1] This is because in practice the point defects themselves are invariant. Any changes during the life of the tape occur inside the measured defect boundary which the theory already assumes is unusable.

The IEEE Mass Storage System Reference Model: Update on Version 5

Bob Coyne

IBM Federal Systems Company
3700 Bay Area Blvd., Houston, TX 77058

Overview

- Work in Progress
- Key Concepts
- The Big Picture(s)
- Fundamental Abstractions
- The Mover
- The Physical Volume Repository
- The Physical Volume Library
- The Storage Server
- The Bitfile Server
- Environmental Services

Work in Progress

Latest document integration and editing
by Rich Garrison of Martin Marietta
Version 5 unapproved draft 1.2, Oct 18, 1993
Current draft available on request after October meeting
Collecting public comments, send to
ieee+mss@larc.nasa.gov and/or
coyne@houvmscc.vnet.ibm.com

Key Concepts and Features

Abstract model for open storage systems interconnection (OSSI)
Modularity
Transparency
Separation of policy and mechanism
Logical separation of control and data flows
Third-party transfers
Layered object naming via name services
Enable automated storage hierarchy management
No scalability limits
OSI system management model
General security model

Fundamental Abstractions

Sets/Containment/Groups
used for managing sets of storage system objects
many types of homogeneous and heterogeneous sets

Stores
structured address space with operations
fundamental object manipulated by the Model
contain all storage data
includes physical stores, virtual stores, virtual volumes

Physical Volume directly maps to storage media

Cartridge contains physical volumes

Device contains read/write interfaces and mount points

SOID typed name for all Model-visible objects

Mover

Media Access Point

a read/write interface with state (position, etc.)

Dual Role:

as *Device Manager*
controls media access points
as *Data Transfer Manager*
controls peer data flow between two movers

Commands: copy, load/unload, position

Physical Volume Repository (PVR)

Contains cartridges and device mount points
Mounts cartridges onto device mount points
Supports location-independent access to contained cartridges and devices
Transfer mechanism may be human, robotic, or a combination
Commands: mount/unmount, stage, inject/eject

Physical Volume Library (PVL)

Location-independent mounting of a set of cartridges resident in various PVRs
Secure and reliable mounting of removable media
Single uniform physical volume name space
Global resource allocation and compatible device selection
Lifecycle management for cartridges and physical volumes (assigned, scratch, maintenance)
Commands: mount/unmount, stage, import/export

Storage Server

Composes physical and virtual stores into other virtual stores
Supports concatenation, replication, RAID, etc.
Virtual stores are managed and exported via storage groups
Range of access semantics:
  fine to coarse grained access & allocation
  shared (locking) to unshared access
Commands: copy, create/delete/reconfigure, mount/unmount, import/export

Environmental Services

Communications
Location by SOID Name
Security
Layered Name Services (Directory Services)

Storage System Reference Model, Version 5

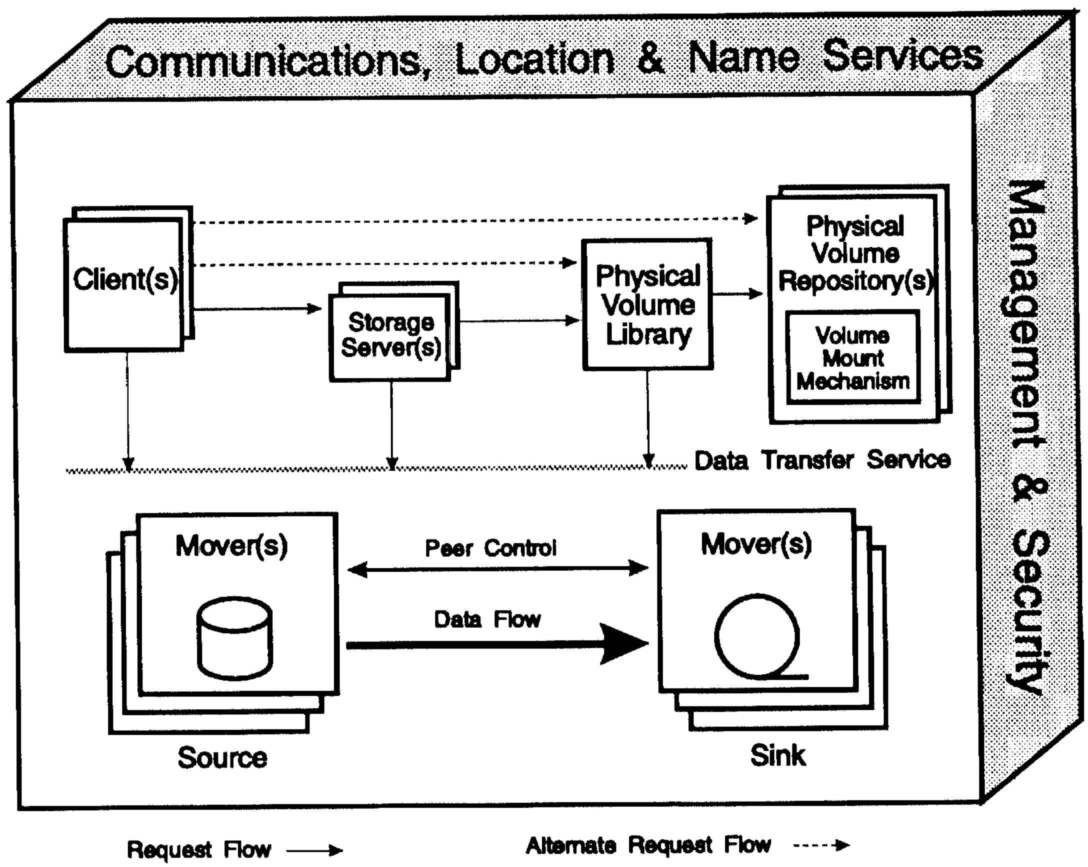

Mapping the Reference Model to Existing Components

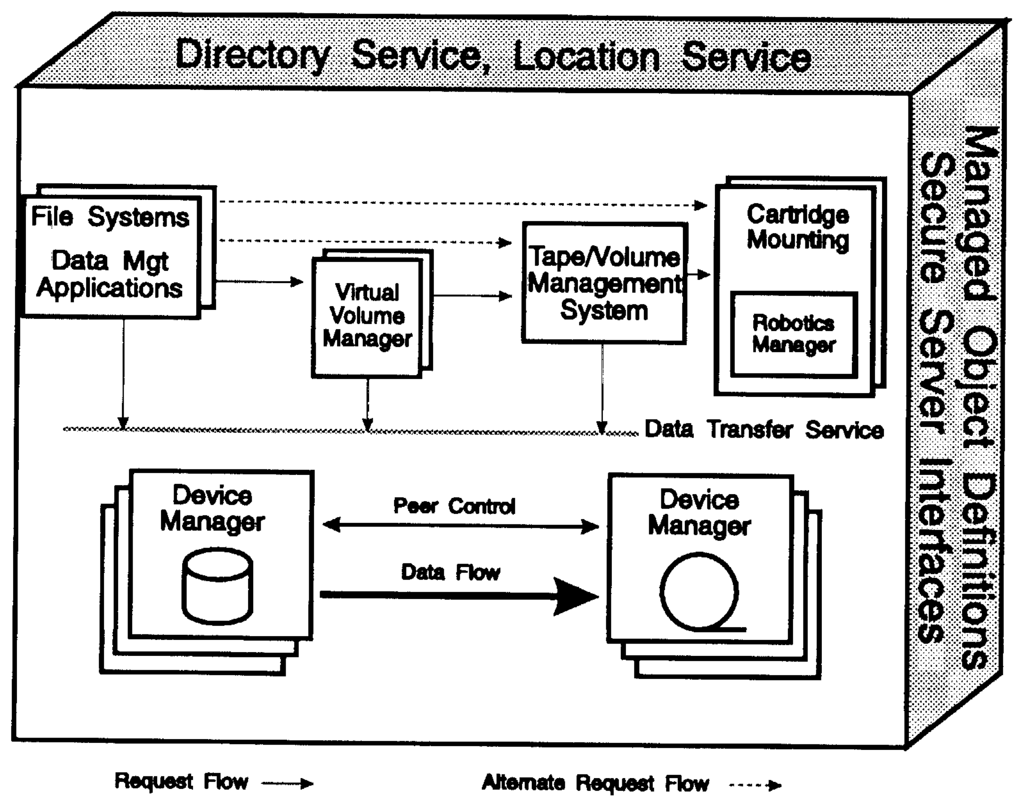

Authorized Standards Projects in the SSSWG