Case Study: An Investigation into the Cause of “Error Rate Drift”

Simon Chandler[*], B. Tucker[*], P. Turner[*] and P. Heard[†].

Abstract

During product development some Hewlett-Packard DDS tape drives experienced a gradual increase in their uncorrected error rate when running a particular operating test at low humidity. This was known as “error rate drift”. A combination of system-level testing, signal processing, and atomic force microscopy showed that the degradation was caused by contamination on the surface of the heads. The “stain” thickness increased each time tape was reused, resulting in progressive attenuation of the readback signal amplitude. Using time-of-flight secondary ion mass spectroscopy the contaminant was identified as inorganic Fe, Si and Al, transferred from the media.

Keywords: *tape, head, contamination.*

Introduction

During product development Hewlett-Packard (HP) Digital Data Storage (DDS) helical scan tape drives[1] are put through many operating tests in which their performance is evaluated. Some drives were found to experience a progressive increase in their uncorrected error rate when subjected to a particular test in a low humidity environment. This is colloquially known as “error rate drift”. Although the DDS format’s error correction system protects the end user from being affected by this degradation, HP wanted to understand the phenomenon in order to reduce its impact on the drive’s margin to other impairments.

Figure 1 shows a typical example of error rate drift that occurred when a drive was operated in an environmental chamber at 25°C and 10% relative humidity (RH). Starting at a known point on unused tape a read-after-write (RAW) error rate measurement was made over 768 frames (~9 seconds), with the drive writing a 1T data pattern (transition every bitcell). The tape was then rewound to the starting point and the RAW error rate measure-

---

[*] *Hewlett-Packard Ltd, Computer Peripherals Bristol, Filton Road, Stoke Gifford Bristol, BS12 6QZ, UK.*

[†] *Interface Analysis Centre, University of Bristol, 121 St Michaels Hill, Bristol, BS2 8BS, UK.*

ment was repeated over the same section of tape. This RAW-rewind cycle was repeated 40 times over the same section of media (i.e., 40 "passes" take place) resulting in the first region of drift. The tape was then moved forwards at normal recording speed for three minutes so that unused tape passed over the write and read heads. This resulted in the first drop in error rate. The drift sequence was then restarted on a unused piece of tape to give the second region of drift. This cycle of drifting and recovery could be continued indefinitely.

The size and shape of error rate drift were found to be very variable from drive-to-drive although individual drives showed fairly repeatable behaviour. The size of drift was not correlated with the drive's absolute error rate, and was found to occur with all DDS products and media, which is metal particulate (MP).

Figure 2 shows how error rate drift decreases as the humidity of a drive's environment is increased. The metric plotted is the RAW error rate of pass number 20 divided by the error rate from the first pass. Within the tape drive's specified operating range of 20% to 80% RH (non-condensing) the increase in error rate can be considerable. For this reason an investigation was carried out to determine the root-cause of the error rate drift. Two aspects of the drift phenomenon were considered: what is the impairment that causes the error rate to increase, and why does it get systematically worse when the drive is operated as shown in figure 1?

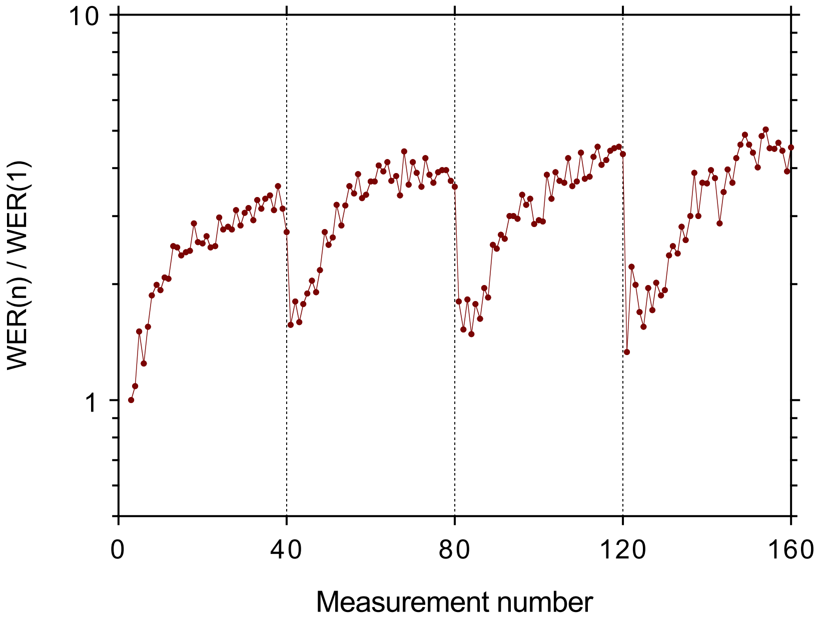

Fig 1 Four consecutive regions of "error rate drift" at 10% RH. The sequence of operations yielding this behaviour is described in the text. The quantity plotted is the uncorrected read-after-write error rate (WER) normalised by the first measured value.

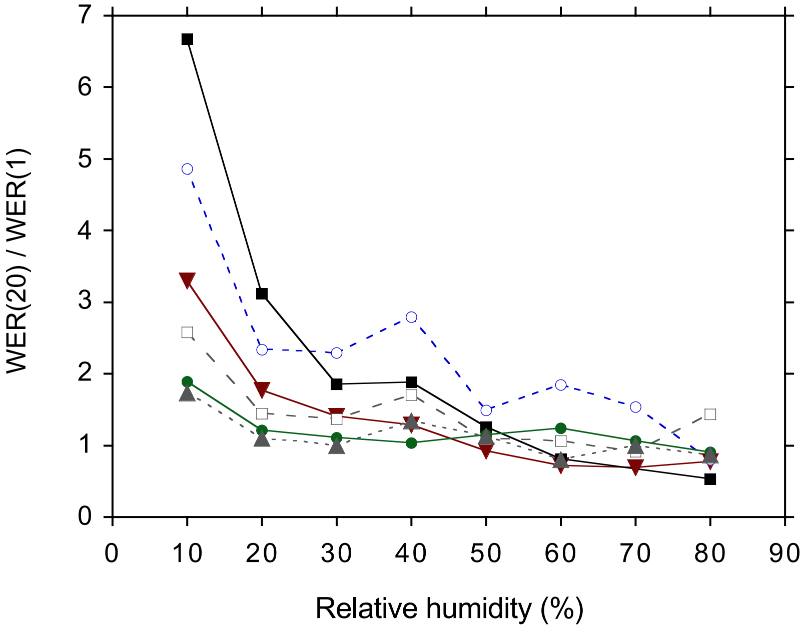

Fig 2 Amplitude of error rate drift vs. humidity of six randomly selected DDS-2 tape drives.

Method

Experiments were performed on ten first generation HP DDS-2 drives, five prototype second generation DDS-2 drives, and five prototype DDS-3 drives. These were operated in an environmental chamber at 25°C and at a controlled relative humidity, usually 10%. Throughout the experiments these drives used only a single batch of HP DDS-2 tape. Initially the drives were "run-in" to the long-term steady state head profile for this tape by operating for 120 hours at 80% RH. Subsequently, no other type of tape was used to avoid any head reprofiling.

To compare the size of two occurrences of error rate drift the starting conditions must be the same. This was not the case in figure 1, because the error rate after moving onto new media before measurements 40, 80 and 120 was not as low as the initial value; the size of the first drift cannot be compared with the other three. By moving unused tape forwards at four times the normal recording speed for 10 minutes at 80% RH a drive's error rate could be returned repeatedly to approximately the same level - usually the lowest achievable with that drive. This "recovery process" was used before all experiments in which the size of the drift was important.

The drives can be controlled using a simple assembly language program stored on their internal EEPROM. Using this a sequence of operating steps could be performed, and error rate measurements made and recorded without connecting the drive to a host computer. This was how the data shown in figure 1 was collected. By making changes to this initial program and observing the affect on error rate drift it was possible to ascertain much about the cause of drift. These “system-level experiments” helped to identify both the underlying impairment, and to explain the systematic changes that occur during the tests.

By connecting to a host computer and using an instrumentation control program it was possible to control the drive as in the system-level experiments, measure error rate, and simultaneously record the analog signal from the read heads using a digital storage oscilloscope. Signal processing techniques were then used to analyse the captured waveforms, and in this way measurements were made of the concurrent behaviour of the error rate and lower-level metrics, such as the coherent signal amplitude, channel frequency response and signal-to-noise ratios.

A Digital Instruments Dimension 3000 Atomic Force Microscope (AFM) was used to monitor the physical condition of the drive’s heads at intervals during error rate drifting. To do this it was necessary to remove the drum from the drive’s mechanism and then replace it after imaging, but this did not appear to disturb the error rate or the drifting behaviour. To maintain a similar tip-sample interaction throughout the sequence of measurements the system was operated using the same tapping mode cantilever and approximately the same AFM system settings.

Surface specific chemical analysis was performed on clean and contaminated heads, on multiply-used tape, and on an adjacent unused region of the same tape. This work was carried out by the University of Bristol using a Vacuum Generators Time-of-Flight Secondary Ion Mass Spectroscopy (ToF-SIMS) system. Before measurement, all sample surfaces were sputter cleaned to remove atmospheric contamination.

Results

This section will describe four groups of experiments that aimed to identify the nature of the impairment that underlies error rate drift: system-level experiments; signal processing-based experiments; microscopy; and chemical analysis. It will then describe a further set of system-level experiments that attempted to explain the systematic behaviour of this impairment when the drive was operated as shown in figure 1.

System-Level Experiments

In the standard operating test, described in the introduction, the tape is re-written on each pass and the error rate is taken from read-after-write measurements. In the first variation of this standard program new tape was written only on the first pass, and in all subsequent passes the read error rate was recorded. This produced essentially the same error rate drift, indicating that the behaviour is dominated by an impairment to the read channel rather than a writing effect.

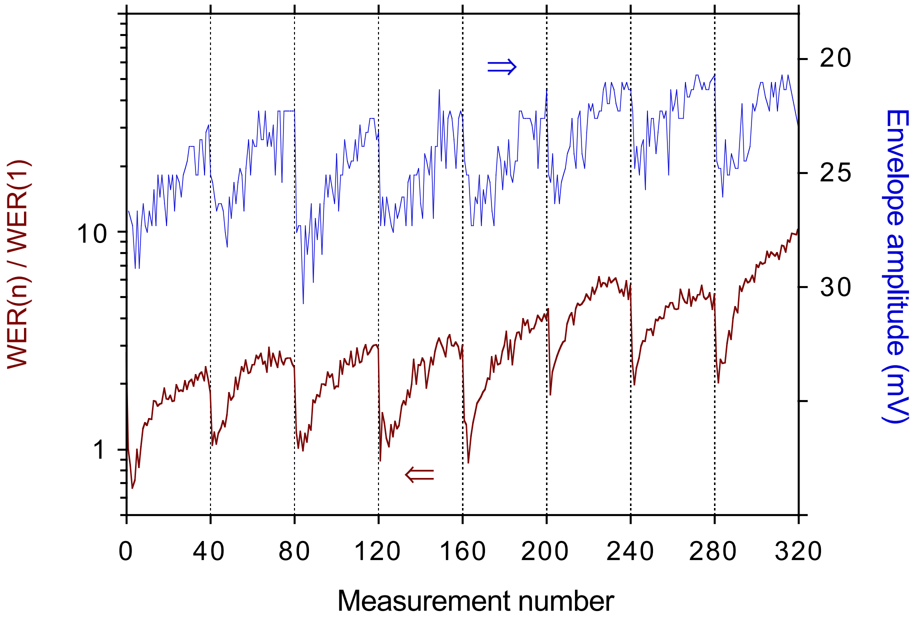

Fig 3 Simultaneous measurement of error rate and the read head's analog signal amplitude. Note that the amplitude scale is inverted. Error rate data from the first four drifts of this experiment were shown in fig 1.

A second experiment was performed to differentiate between media and read-head effects on error rate drift. Data was written on new tape, the cassette was ejected, read by a second drive (2nd pass over the media), and then returned to the first drive for 39 more read-only passes. The cassette was then transferred back to the second drive and a short final read error rate measurement was made on the repeatedly used section of tape. The first drive, which had experienced 40 passes of media, showed the usual ramping error rate. The second drive, which had experienced only two passes of media, showed only a slight increase. This suggested that drift is caused mainly by changes to a drive's read head rather than by changes in the media.

During a read-only pass a HP DDS drive can record the position on tape of any C1 blocks[*] that are in error; this is called "tape mapping". If two or more consecutive blocks are erroneous the event is called a "burst error", and may indicate that a dropout has occurred. To determine whether drift is caused by an increase in dropouts or by an increase in random errors, tapes were mapped on each pass through an error rate drift experiment. The change that occurred in the burst length distribution during the drift sequence was statistically consistent with the drift being caused entirely by random errors, with no

---

[*] A C1 block contains both user data and error correction parity bytes, see reference 1

increase in dropouts. When combined with an observation that the read head's analog signal amplitude decreases during drift (figure 3), this suggested that the increase in error rate is caused principally by reduced signal-to-noise ratio (SNR).

Signal Processing-Based Experiments

An attempt was made to directly measure the change in SNR as drift occurred. A drive was cleaned using the high humidity recovery process then, whilst writing pseudorandom data, seventeen repeat passes were made over a length of tape 12,288 frames long. At the end of passes 1, 5, 9, 13 and 17 the RAW error rate was recorded. At three equally spaced points during these same passes a digital storage oscilloscope was used to record (with 4× oversampling) the analog signal output from the read head. Following the seventeenth pass the drive moved onto unused media and, with no further cleaning, performed a second 17-pass drift sequence. The resulting 30 data files (or "waveforms"), each one track in length (160 kB), were processed using a software simulation of the phase-locked loop data recovery algorithm. The simulation extracts the amplitude and mean noise of the coherent signal, the channel frequency response and the incoherent noise power spectrum. From these both the coherent and incoherent signal-to-noise ratios can be calculated.

The results showed that the coherent signal amplitude decreased as drift progressed (figure 4), as did the mean noise, but the behaviour was less regular than that of the error rate. Overall the coherent SNR did decrease, but the change was even more erratic and the data did not provide strong, direct evidence for an SNR origin to the error rate increase (that was suggested indirectly by other experiments).

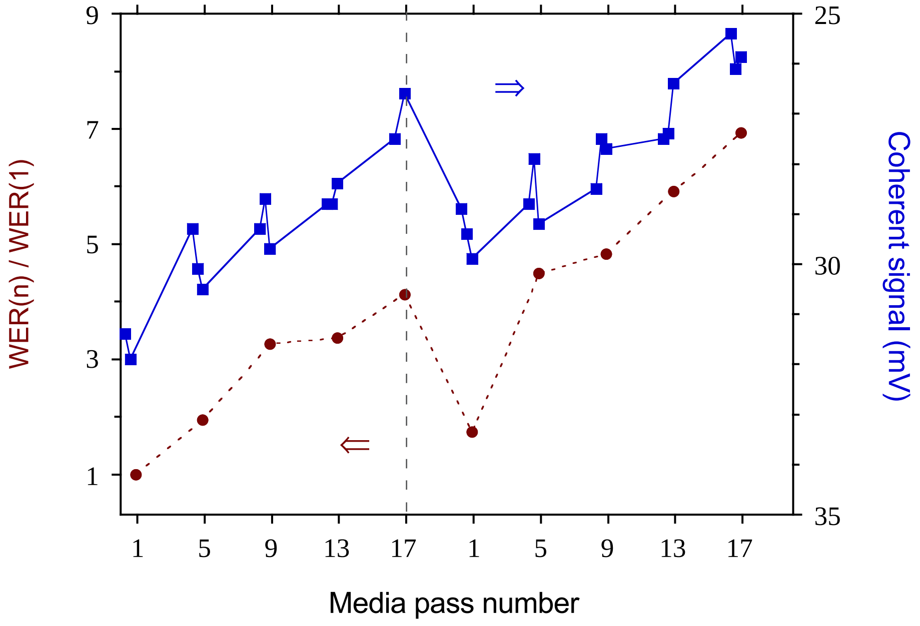

Fig 4   Simultaneous measurements of uncorrected error rate and coherent signal amplitude. Note that the amplitude scale is inverted. After pass 17 a second drift sequence is started on unused media.

The difference in smoothness between the error rate data and the low-level data can probably be attributed to the large difference in the duration of the measurements; each error rate point resulted from the average behaviour over many hundreds of tracks (12,288 in the case of figure 4), but each waveform is only 1 track long. Future work will attempt to improve the signal processing-based experiments, and provide direct evidence for a link between SNR and the drifting error rate.

Despite the large measurement error, analysis of the calculated channel frequency response was able to give some indication about the cause of reduced signal amplitude. Each channel frequency response curve was normalised by the response from the first measured waveform, captured on the first pass over the media. In the frequency range 5 MHz to 15 MHz the normalised curves were of the form $\exp(-cf)$, where f is frequency and c is a constant that increased as the drift progressed. This behaviour is consistent with the drift being caused by spacing loss (or “Wallace loss”) due to increased separation between the head’s gap and the magnetic recording media[2]. The factor c is then equal to $2\pi d/v$, where v is the velocity of the head with respect to the tape and d is the increase in head-tape spacing since the first pass. The calculated change in spacing is shown in figure 5. When combined with the system-level experiments this suggests that the impairment underlying error rate drift is a physical separation between the head and the media, and that this separation is located on the read head.

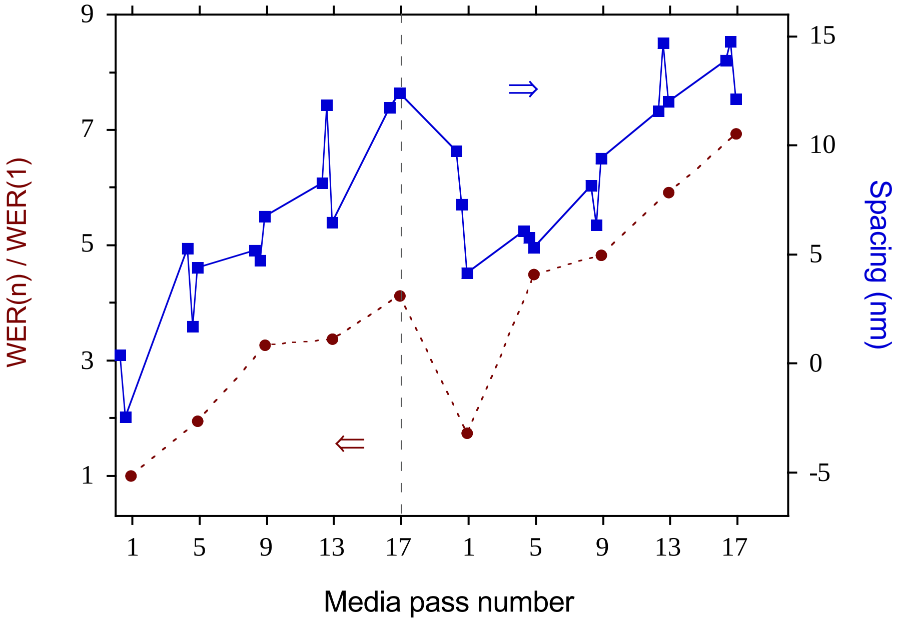

Fig 5 The change in head-tape separation during drift, as estimated from the slope of the normalised channel frequency response. After pass 17 a second drift sequence is started on unused media.

Optical and Atomic Force Microscopy

Using an optical microscope a patchy white "stain" can be seen on the glass parts of HP DDS-2 and DDS-3 write and read heads. The stain is tenaciously attached to the head and cannot be removed by any known form of light cleaning (e.g., alcohol and swab). Running the drive at high humidity with unused tape does appear to reduce stain on the glass but does not eliminate it, whilst operating at low humidity appears to give more stain. This form of contamination was a naturally favoured explanation for error rate drift, although the theory was untested. Other possible sources of spacing loss are gap erosion, pole or MIG material recession, and a macroscopic physical change in the head profile.

To identify the source of spacing that causes error rate drift an atomic force microscope was used to examine the surface of a DDS-2 drive's read head at several stages during a drift experiment. The first AFM height image was measured just after the drive had been through an 80% RH recovery process. The ferrite was completely clean, but on the glass were patches of material several nanometres thick that coincided with the position of white stain seen with an optical microscope. The AFM phase image showed strong contrast between the raised patches and the underlying glass, indicating that they were a different material. Although we cannot be certain, it appears that following a high humidity recovery process remaining patches of white stain do not contribute to spacing loss because the glass to which they adhere is recessed from the ferrite.

To emphasise changes in the head's topography the first image was subtracted from all subsequent height images, after aligning them in both tilt, angle and position. Following a lengthy period of continuous movement over media at the normal recording speed ("streaming") no change was seen to occur in either the drive's error rate or read head topography. But when the drive performed a repeated cycle of media reuse the error rate increased and contamination built-up on the surface of the read head (figure 6).

The increase in thickness of the surface contamination (~20 nm) was similar to the spacing estimated from the signal processing-based experiments, and was significantly larger than any change in the overall head profile, in height difference between the two poles, or in recession of the glass below the ferrite. In fact no change in the latter three items could be seen above noise and distortion introduced by post-processing of the image.

Contamination built-up on both the glass and on the ferrite. The former was clearly visible under an optical microscope as "white stain", but the latter could only be seen when a Nomarski or differential interference contrast (DIC) filter was used. The increase in thickness was similar on both the ferrite and the glass, and it is not clear which of the two is more culpable in creating the spacing loss. If there is a single cause it is more likely to be contamination on the ferrite, since it is more continuous and occurs right at the gap.

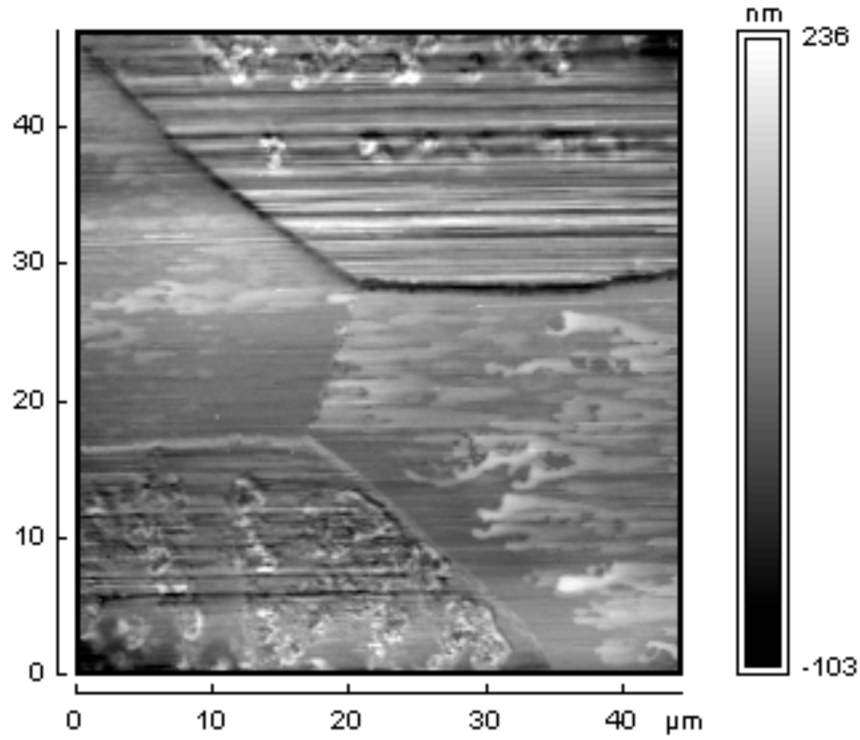

Fig 6 Change in topography of a DDS-2 read head from a clean state just after an 80% RH recovery process to a stained state following repeated media reuse. Contamination is thicker on the leading pole (right of image).

The AFM experiment provided strong evidence that the origin of error rate drift is surface contamination on the drive's read heads.

Chemical Analysis by ToF-SIMS

Composition of the stain

The chemical composition of the surface contamination was identified using Time-of-Flight Secondary Ion Mass Spectroscopy[*]. Measurements were made on glass and ferrite regions of the stained head shown in figure 6, and on the same head after the surface had been sputter etched to a depth of 170 nm. The latter gave very similar spectra to a second head that had been cleaned using the 80% RH recovery process with abrasive tape.

By comparing spectra from the clean and stained heads the contamination was shown to consist primarily of Fe, Al and Si. There are several sources of these elements in the system so their origin is ambiguous, but the stain also contained small quantities of Co, Ni and other high molecular weight additives (e.g., a Lanthanide) that could only come from

---

[*] Auger emission spectroscopy failed due to charging of glass parts of the head.

the media. The stain contained very little organic material so is not binder or lubricant transferred from the media, and is not a “friction polymer”[3]. A small amount of cross-contamination did take place between the glass and the ferrite, but the dominant process appears to be transfer of Fe, Al and Si from the media. The composition of the stain on the glass and ferrite was identical.

Changes in the media

ToF-SIMS was used to compare the surface of tape from a multiple-pass drift experiment with the surface of an adjacent region of unused tape. This was done to try to identify the mechanism of material transfer from the media to the head. The largest absolute change was an increase in iron (14% to 19%) and a fall in positively charged fragments with mass 29, 41 and 43 Daltons. There was a large relative decrease (greater than 2×) in many peaks between 37 and 100 Daltons, and an increase in C, CH and peaks at 143 and 145 Daltons. There was little change in spectra from negatively charged particles.

The tape contains various organic components (e.g., binders, lubricant), so the observed change in the ToF-SIMS spectrum could be due to a change in the relative concentration of these at the surface, and/or a chemical change in one or more of these components. Clearly a subtle change occurs in the organic composition of the media, but the actual change could not be identified.

Systematic Aspects of Error Rate Drift

There are two parts to error rate drift: an impairment (i.e., stain), and a mechanism that causes the impairment to increase. Physical spacing does not in itself cause error rate *drift*; a fixed spacing loss will make the error rate worse, but the error rate will remain constant. For drift to occur there must be a systematic increase in the physical spacing. System-level experiments were carried out to investigate this aspect of the error rate drift phenomenon.

To identify the crucial step that results in error rate drift the standard operating test was modified by removing or increasing component parts, and the affect of the change on the size of drift was observed. This showed that error rate drift is *not caused* by simply streaming over media, stopping the tape, allowing the drum to spin down, or changing the direction of tape pulling. The error rate increased only when the media was reused. This was found to be true even when the media and drive were exposed to a high humidity (80% RH) and were left unused for a considerable time (24 hours) between each low humidity pass. Similar results were obtained by Kuroe *et al*[4].

These experiments also showed that the error rate of a stained drive will gradually recover if the heads are streamed over unused media. This occurs even when the drive is operated at low humidity, and does so with a recovery ‘time constant’ of about a few thousand tracks (although the actual value is very variable and depends on the relative humidity).

This form of recovery is clearly seen in figures 4 and 5, where the signal amplitude and spacing change after moving onto unused tape between media passes 17 and 1.

Discussion

The results from this investigation suggest a model in which the overall thickness of spacing, and hence the level of error rate degradation, is determined by a dynamic equilibrium between stain formation and stain removal. The stain is formed by transfer of material from the media onto the head, and is removed by abrasive wear of the head by the media. Whilst the drive is running with a fixed set of conditions (e.g., humidity, number of previous uses of media) the equilibrium level of spacing will remain constant and the error rate will not drift. Changing an important condition will upset the balance and the system will move towards a new equilibrium thickness of spacing.

Error rate drift occurs when the media is reused. This suggests there is some change in the media each time it interacts with the drive's heads and that this change upsets the equilibrium - either by increasing the rate of material transfer or by reducing the rate of wear. The second of these is supported by experiments that show a reduction in the media's abrasivity[5] (a qualitative measure of its ability to wear) with multiple passes (figure 7). Our observation that the error rate of a stained drive slowly recovers when it subsequently streams on unused media also supports the proposed model. The smaller error rate drift at high humidity (figure 2) is explained by the increase in abrasivity that occurs when the humidity of the operating environment is increased[5].

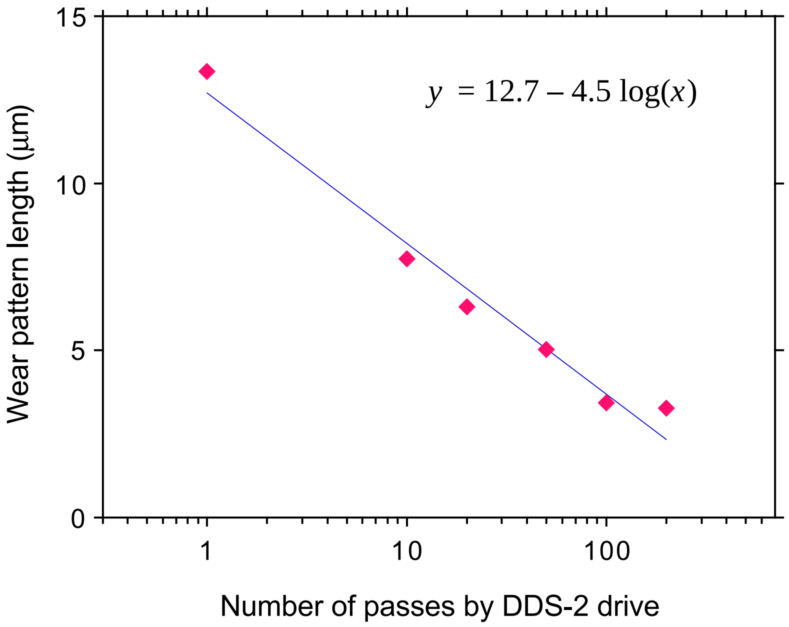

Fig 7 Change in relative abrasivity of HP DDS-2 media with use. The abrasivity was measured using an AlFeSi wear bar test with 50 m tape at 10-20 g tension (Data from experiments by K.Kusumoto[6]).

One factor that probably increases the rate of wear of the stain is the thickness of the stain itself. Without such feedback the two proposed mechanisms would result in an equilibrium rate of change in spacing, and hence a uniform rate of drift, rather than an equilibrium value of the overall thickness.

Conclusions

The root-cause of error rate drift is a layer of contamination on the surface of read heads that results in spacing loss and a consequent reduction in signal amplitude. The contaminant is inorganic Fe, Al and Si transferred onto the heads from the media. The thickness of the contamination layer, and hence the extent of signal attenuation, is determined by a dynamic equilibrium between competing processes: transfer of stain material from the tape to the head, and removal of the stain by abrasion from the tape. Each time a piece of tape is rubbed by a head its abrasivity is reduced, shifting the system’s equilibrium and slightly increasing the spacing loss. The contamination only builds up with media reuse, and repeatedly reusing the same section of tape causes the error rate to progressively increase.

As an outcome of this work the usage of Hewlett-Packard DDS format tape drives was carefully scrutinised and the firmware was modified to reduce the number of avoidable repeat passes over media. This substantially reduced the magnitude of error rate drift in normal operation, particularly at low humidity.

References

1. **Tan E. and Vermeulen B.** Digital audio tape for data storage, *IEEE Spectrum 1989*, **26**, 34-38
2. **Mallinson J. C.** The foundations of magnetic recording, *Academic Press, San Diego, 1993*. p.151
3. **Bhushan B. and Hahn Jr F. W.** Stains on magnetic tape heads, *Wear 1995*, **184**, 93-202
4. **Kuroe A., Shinoda F. and Mikoda M.** An experimental analysis of the recorded signal decrease phenomena in the repeated tape running, *Electrical Information Society 1991*, **J74-C-II**, 194-200 (in Japanese)

5 **Jorgensen F.** The complete handbook of magnetic recording, 3rd edition, *Tab Books Inc., Blue Ridge Summit, PA, 1988.* p424. Note: some material about wear was omitted from the 4th edition.

6. **Kusumoto K.** Private communication, *Magnetic Media R&D Dept., Hitachi Maxell, Oyamazaki, Otokuni, Kyoto, Japan, 1994*