Predicting Archival Life of Removable Hard Disk Drives

Paul Williams; ProStor Systems, Inc.; Boulder, CO., USA / David S. H. Rosenthal; Stanford University Libraries; Stanford, CA., USA / Mema Roussopoulos; Institute of Computer Science, FORTH, Greece / Steve Georgis; ProStor Systems, Inc.; Boulder, CO., USA

Abstract

A collaboration between the developer of a Removable Hard Disk Drive (RHDD) storage technology product and academics presents an accelerated life test of non-spinning powered down hard disks. The results are used to predict the reliability of removable disk cartridges stored in shirtsleeve environments as archival media. The contributions of this paper include an initial study of the reliability of data stored off-line on RHDDs, and a description of the industry standard accelerated life test process together with the methodology used to predict storage media reliability based on the results. Our study shows that removable hard disk drive technology has considerable potential as a reliable archival medium.

1. Introduction

Attention has recently been drawn to the lack of suitable data on which designs for long-term storage can be based. At the 2007 FAST conference the keynote [1] and papers from Google [2] and Carnegie-Mellon [3] discussed the mismatch between manufacturer's specifications and actual performance in practice of disk drives. At the 2008 conference, data from NetApp showed that other components of storage systems also contribute to a significant rate of storage corruption [4]. This paper is the result of collaboration between a storage vendor and academics to see what data can be made available, to encourage system designers and vendors to talk the same language, and to investigate some of the impediments to better storage media performance data

Current data collections involve large arrays of spinning drives; this work looks at data retention on non-spinning, idle drives. These are becoming important. MAID (Massive Array of Idle Disks) technology, in which densely packed SATA disks are spun down whenever possible to reduce power, wear and vibration, is becoming available. Removable hard disk cartridges (RHDD) are becoming an alternative to tape cartridges for off-line storage.

Idle drives are a more tractable case for performance data, because there are many fewer parameters to consider. They are much less subject to the continued access, thermal and power variations, vibration and other environmental factors which spinning drives inevitably encounter.

Nevertheless, some of the formidable difficulties in the way of providing storage system designers with high-quality performance data of disk in general still apply to idle disks. Disks are remarkably reliable. It was only because Google and Carnegie-Mellon studied vast numbers of drives for long periods of time that they were able to reveal clearly the disparities between specifications and performance. A particular disk product may only be on the market for 18 months where the service life of a spinning drive might be 60 months. Reliability data based on experience in service of a particular product will typically be available after the product is obsolete. Useful reliability data about particular products will necessarily be a prediction based on testing a sample of early production units.

To provide meaningful predictions of storage reliability using feasible sample sizes and test durations requires accelerated life test techniques. The storage devices are subjected to conditions extreme enough to cause many failures. The failure data is adjusted to predict failures in normal conditions using a model of the effect of environmental conditions on failure rates.

We report on an accelerated life test on RHDDs and a small sample of desktop SATA drives, using industry-standard techniques [e.g.,5], performed by an independent testing lab under contract to the developer of the disk cartridge technology. We provide an overview of the accelerated aging test process, report the data collected, describe the model used to generate predicted reliability from the data and show the results. This analysis concludes that, stored in realistic conditions, the RHDDs are reliable enough for archival use. The predicted failure rate is less than 1% over 30 years. The factors leading to the gap between prediction and experience for spinning disks should be less significant in the idle case.

The contribution of this paper is an initial study of the reliability of data stored off-line on RHDDs using (and describing) the industry-standard test process, and the industry-standard methodology for predicting service reliability. We hope this will help storage system designers better understand the disk medium, and the limitations of data available about it.

2. Removable Hard Disk Drives

RHDDs encapsulate industry-standard 2.5” SATA HDDs (notebook computers drives) in a ruggedized cartridge that has the removability, portability, durability and off-line storage characteristics of legacy magnetic tape, but with the random-access performance advantages of disk. RHDDs are designed to protect the embedded HDD against shock, vibration and Electrostatic Discharge (ESD) RHDDs are rapidly gaining acceptance as tape replacement for data protection and archive application in servers and workstations. One example of RHDD is the multi-vendor RDX format developed by ProStor Systems and manufactured by several storage industry suppliers, such as Tandberg Data, Imation and Dell Computer.

One difference between 2.5” laptop drives and standard desktop or server 3.5” drives that is important for use in idle drive systems is the technique used to park the heads when the drive is idle. 3.5” drives use contact start-stop (CSS) technology, parking the heads on a special landing zone on the disk medium when the drive isn't spinning. 2.5” drives, designed for laptop use, use ramp loading technology to physically remove the heads and lock them away from the medium when the drive spins down. This provides much greater non-operating shock tolerance, but it may also be important for long-term data retention.

MAID systems using 3.5” drives typically spin the drives up periodically to exercise the mechanics and reduce the risk of failures caused by CSS, among other reasons. RHDDs are designed to be removed from the system for storage, making periodic exercise of this kind impractical.

If idle disk drives are to be used as an archival medium, storage system designers need to know how long the data written to them will remain readable. Existing research on spinning disks has suggested a number of factors which affect their reliability, including power-on hours, duty cycle, operating environment, rotational vibration and access patterns (e.g., [6]). Idle disks operate at very low duty cycles and are mostly powered-down. Consequently operating environment, vibration and access patterns are not significant concerns.

3. Background

The conditions, under which the powered-down RHDDs are stored, in particular temperature and relative humidity, can affect their ability to retain data. ProStor Systems designed and contracted with an independent testing service to conduct an accelerated life test so as to understand these reliability factors and the limits of long-term data preservation using RDX cartridges.

Through an understanding of the design of HDDs, known failure mechanisms and aging effects, it is possible to postulate a set of archival life reliability factors for RHDDs and design experiments to test for these factors.

These potential factors identified are:

1. Magnetic thermal decay of recorded bits and control signals
2. Media corrosion
3. Media lubricant evaporation
4. Fluid dynamic bearing oil evaporation
5. Electronics corrosion and degradation

These archival life factors are all functions of temperature, humidity and time making the factors excellent candidates for accelerated testing stresses.

4. QALT

Quantitative accelerated life testing (QALT) [7] consists of a series of related tests designed to quantify the life characteristics of a component or system under normal use conditions by testing the units at higher stress levels in order to accelerate the occurrence of failures. These tests provide valuable information about a product's performance under normal use conditions that can allow a manufacturer to make predictive statements about its products field performance. The obvious benefit of quantitative accelerated life testing is the time savings, which is based on the decrease in test duration due to increased stress levels.

With stress related acceleration, one or more environmental factors that are known to cause the product to fail under normal conditions such as temperature, voltage, humidity, vibration etc. are increased in order to cause the product to fail more quickly in the test. The stress and levels of stress used in accelerated tests must be chosen so that they accelerate the failure modes of the product but do not introduce failure modes that would not normally occur under normal conditions. These stress levels may fall outside the product specification limits for the product but usually well inside the real design limits.

The life data obtained from these tests require accelerated life data analysis techniques, which include a mathematical model to translate from accelerated conditions to the product under normal use conditions. This model can be used to calculate important reliability statistics. These include: Reliability or the probability of success, or the converse (the probability of failure), the mean life of a population, the failure rate per unit time (hours), and the B(X). The distribution parameter B(X) where X is the given proportion of the population which will fail by the time being evaluated. For example if it is desirable to know at what time 1% a population will fail, then B(1) would be evaluated to produce a number of hours corresponding to this cumulative failure proportion.

5. Experiment Design

For this particular accelerated life test we have already discussed the expected failure mechanisms and the associated stresses which precipitate them. These factors are temperature and relative humidity. The experimental design utilized three levels of stress for each in order to enable the use of non-linear life models in analyses of the resultant data. This is important because real world practical models are exponential in their nature [8]. To accomplish a three level stress experiment with two different stresses, six test cells are required. To save on samples and test resources one of the cells from each stress can be common. This improves the efficiency of the test by reducing the total number of stress cells to five. The levels of stress were chosen to stress the drives to the maximum limits of their capabilities without creating unrelated failure mechanisms. Also, larger sample sizes were used in the lower stress cells in an attempt to observe enough failures for statistical requirements. Longer test intervals were also used in these low stress cells to further improve the efficiency of the test and to reduce hazards from handling the drives. This exercise resulted in a 5 stress cell experiment detailed in Table 1, where the Test Cell is a letter designation to simplify tracking. The rest of table 1 includes, Test Stress consists of Temperature in degrees Celsius and Hum RH is Relative Humidity. Samples refer to the number of RDX removable cartridges used in each test cell and the test interval is the number of hours between each test point.

Table 1: Test Cell Definitions

	Test Stress			
Test Cell	Temp C	Hum RH	Samples	Test interval
A	80	85	10	336
B	80	55	10	500
C	80	10	15	500
D	70	85	15	750
E	60	85	30	1000

6. Test Samples

A sample of 80 RDX 160GB (80GB/Platter) 2.5" form factor hard drives was used for this experiment. The particular drives were selected because they were the first commercially available drives using perpendicular recording technology at this areal density. These samples were distributed as detailed in Table 1.

Additionally, a small sample of five 500GB 3.5" hard drives (using the same areal density and perpendicular recording technology as the 160GB RDX drives) were placed into the chamber with test cell A. This group is later referred to as group F in the results section. The relative performance of these drives to the RDX removable drives in this cell are of interest to benchmark desktop backup solutions performance over time.

7. Test Procedure

The testing was carefully monitored and rigorously performed and our methodology follows industry-standard testing practices [5]. First, each drive was functionally tested and random data written to the entire drive. Then each drive was read comparing the data to what was originally written to verify the data was written properly and to record any errors. This process established that the removable disk was error-free before any stresses were applied. The samples were separated into the 5 test groups and permanently marked to identify them.

The following steps were executed for each test interval:

1. The cartridges were placed in the environmental chamber according to sample group and environmental settings.
2. The chamber was ramped to the desired temperature and humidity levels over a period long enough to insure temperature and humidity ramp times specified by the hard drive manufacturer were not violated. Once the temperature and humidity levels were reached they were maintained at the desired levels for the specified interval.
3. At the completion of the interval, the humidity was reduced to 10% RH and the temperature was maintained at the elevated level before being ramped to 20º C over a period long enough to insure temperature and humidity ramp times specified by the hard drive manufactures were not violated. This was done in order to dry out the chamber and hard drives so that condensation did not occur in either.
4. The cartridges were then removed from the chamber and all data sectors were read while checking for errors.
5. All cartridges that accurately read the data were returned to the chamber and the interval cycle was repeated until the desired total duration was achieved. See Table 2 in section 8 for details.

Failure was defined as an unrecoverable read error while reading any data sector from the disk or a mis-comparison of the data to what was originally written. Once a cartridge experienced an unrecoverable read error or mis-compare, it was removed from the test and logged as a failure. Each cartridge was monitored until either failure or test completion. During the course of testing, the environmental conditions in the chambers were continuously monitored to ensure operation at specified levels.

8. Test Results

The results of the accelerated life tests are shown in the table below. "(Courtesy of: Percept Technology Labs, RDX Removable Disk Archivability Study, Test Report [9])."

Table 2: Results by Test Cell

	Test Stress					
Cell Number	Temp C	Temp F	Hum RH	Samples	Failures	Total time
A	80	176	85	10	4	4000
B	80	176	55	10	4	2000
C	80	176	10	15	2	3500
D	70	158	85	15	5	3000
E	60	140	85	30	1	4000
F	80	176	85	5	5	4000

9. Reliability Analysis

The results of the accelerated testing for achievability were analyzed using ReliaSoft's ALTA reliability software, which is well suited for this type of quantitative data. ALTA is a high-quality commercially available Accelerated Life Test Analysis (ALTA) statistical tool. It is widely accepted as a standard tool for this type of analysis within the Reliability Engineering community [10].

First the time to failure data was coded into ALTA for each of the environmental condition cells of this experiment. This coding includes the time to each of the failures, the time at which drives that did not fail were suspended as well as the environmental conditions of temperature and humidity for each data point.

ALTA provides a tool called distribution wizard which helps determine the best lifetime distribution for the data set. The wizard estimates the parameters for each of several distributions, compares the log-likelihood values and then recommends the one that is the best statistical fit for the data. The distribution wizard was employed to determine the best distribution for the experimental data and Lognormal was chosen as the best fit for the data.

ALTA provides a choice of Life-Stress relationships including models for Arrhenius, Eyring, Inverse Power Law (ILP), Temperature-Humidity and Temperature-nonthermal. These models allow the extrapolation of test data to enable Reliability prediction for other environmental conditions as well as the desired time. The Temperature-Humidity life-stress model is a variation of the Eyring relationship [8,11] which is used for the analysis of data from temperature and humidity tests. This is the model chosen because of its applicability to the stresses used in the archive test.

10. Temperature-Humidity Relationship

The temperature-humidity (T-H) relationship, a variation of the Eyring relationship,[8,11] has been proposed for predicting the life at use conditions when temperature and humidity are the accelerated stresses in a test. This combination model is given by:

$$L(V,U) = Ae^{\frac{\Phi}{V} + \frac{b}{U}} \quad [8,11]$$

where:
ϕ is the thermal activation energy.
b is the activation energy for humidity.
A is a constant parameter.
U is the relative humidity (decimal or percentage).
V is temperature (in absolute units, Kelvin in our work).

T-H Acceleration Factor

The acceleration factor for the T-H relationship is given by:

$$A_F = \frac{L_{USE}}{A_{Accelerated}} = \frac{Ae^{\frac{\Phi}{V_u} + \frac{b}{U_u}}}{Ae^{\frac{\Phi}{V_A} + \frac{b}{U_A}}} = e^{\Phi\left(\frac{1}{V_u} - \frac{1}{V_A}\right) + b\left(\frac{1}{U_u} - \frac{1}{U_A}\right)}$$

where:
LUSE is the life at use stress level.
LAccelerated is the life at the accelerated stress level.
Vu is the use temperature level.
VA is the accelerated temperature level.
U_A is the accelerated humidity level
U_u is the use humidity level.

The objective of this analysis was to extract a prediction of the lifespan for RDX removable disk drives at various environmental conditions as well as set environmental limits for various desired archival periods. Initial environmental usage conditions are were set as 20°C / 68°F at 30% RH to correspond to a standard commercial Archival condition. Graph 1 is the result of this analysis. In the legend each test cell is identified along with the number of Failures (F) and Suspensions (S) for each.

Graph 1: ALTA Life Distribution Plot

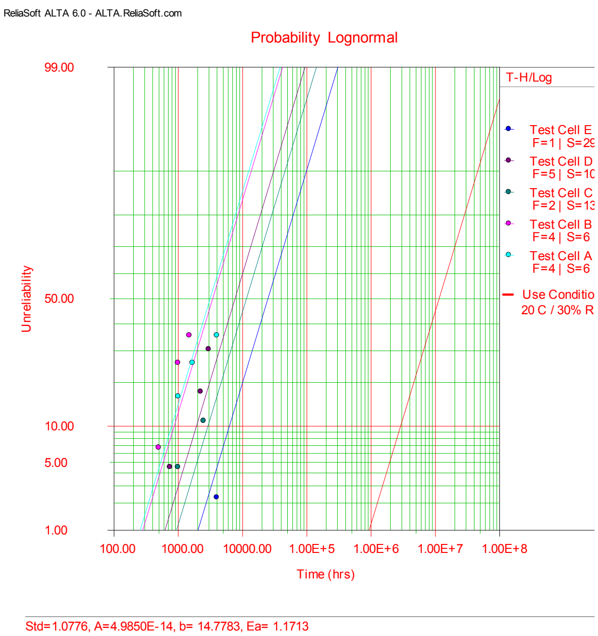

In Graph 1 the X axis is time in hours and the Y axis is Unreliability or the cumulative failure proportion of the test samples. For example, take the upper leftmost point in Graph 1, it represents a failure that occurred at 1500 hours and represented ~35% cumulative failure for the population for test cell B. It should also be stated that Unreliability (probability of failure) is equal to 1 – Reliability (probability of success). As an example unreliability on the graph of 1% corresponds to a 99% Reliability. Each set of colored points represents the failures for a particular test cell and the corresponding colored line represents the best statistical fit for those points. The legend identifies which test group to which the data belongs. The rightmost (red) line with no points surrounding it is a prediction of the life characteristics for the environmental usage condition input. In this case it was set as 20°C / 68°F at 30% RH.

Using the tools within ALTA, the temperature / humidity limits under which the RDX removable disk drive can be safely stored for 20 years, 25 years and 30 years can be predicted. As expected, the higher the temperature in the storage environment, the lower the humidity has to be in order to get the maximum length of storage life. All of this analysis was performed with a data Reliability requirement for the data of 99%. Graph 2 displays these results:

Graph 2: Archival life vs. Storage conditions:

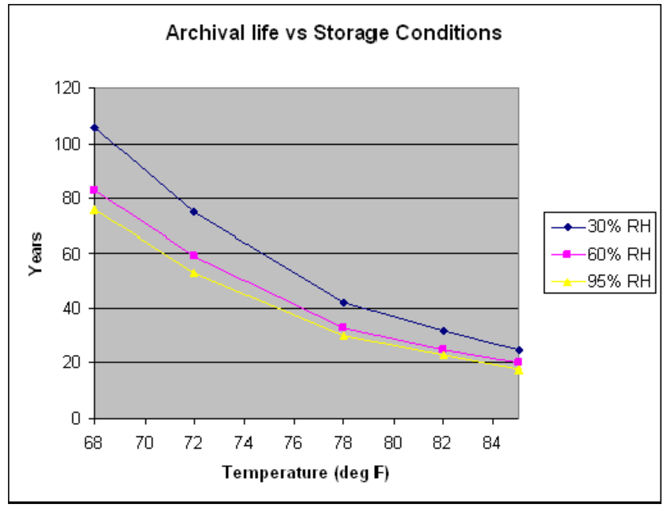

11. Mobile vs. Desktop Drives

It was also desirable to understand the difference between the group of RDX removable drives and the standard 3.5" CSS drives commonly used in MAID systems and desktop external USB devices. In group A, 10 samples of RDX 160GB 2.5" form factor hard drives were subjected to 176ºF / 80ºC 85% for 4000hrs with 4 out of 10 eventually failing as a result of the experiment. In group F, 5 samples of 500GB 3.5" form factor drives were placed into the chamber with test group A. During the course of this experiment all 5 of the samples failed. The 500GB 3.5" equates to the same areal density as the 2.5" 160GB RDX drives. Both groups utilize the same perpendicular recoding technologies sharing common heads and media technologies.

Graph 3: Life expectancy comparison RDX vs. 3.5" Hard Drives

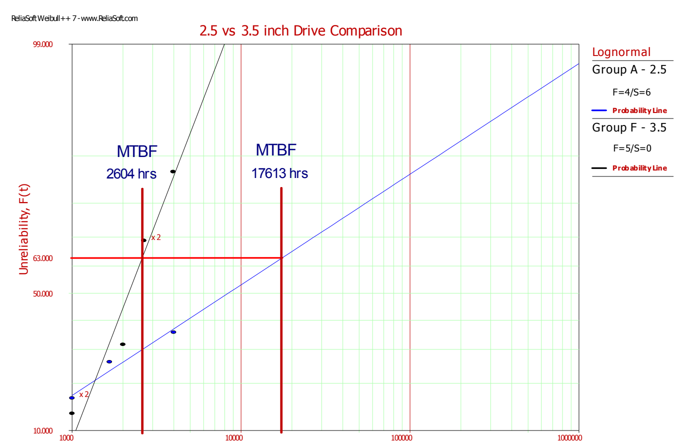

To compare these two groups which were stored at the same environmental conditions Reliasoft's Weibull++ was used to create a life distribution plot of time vs. cumulative failure percentage for the two populations. In this case a reasonable comparison point between the two groups is the characteristic or "average" life or MTBF. This point is chosen because it is the key parameter for the distribution of times to failure. For the 3.5" drives the MTBF is 2604 hours and for the RDX the MTBF is 17,613 hours. This ratio of >6.7 times is a very strong indicator that 3.5" CSS drives are far more vulnerable to data loss than is the RDX removable disk when used as archival media. This comparison is clear in Graph 3.

12. Related Work

Many storage systems have been designed for specific failure modes (e.g., [12, 13]) and there is great need for real data from the field to support the use of these failure models. Published work presenting and analyzing failure data in real storage systems is only starting to appear.

Recent FAST conferences drew attention to the lack of suitable data on which designs for long-term storage can be based. Steve Kleiman's keynote [1] and papers by Pinheiro et al. [2] and Schroeder et al. [3] discuss the mismatch between manufacturers' specifications of disk drives and actual performance in practice. Specifically, Schroeder et al. study and analyze data from 100,000 drives with SCSI, FC, and SATA interfaces and show many disparities between manufacturer specifications and actual disk performance. Pinheiro et al. study disk replacement data from more than 100,000 hard drives in operation at Google, including serial and ATA drives and similarly find disparity in the annual replacement rates measured in actual use and the rates predicted by the vendors. Jiang et al [4] show that other components of the storage system contribute to these failure rates. Bairavasundaram et al [14] show that in addition to these visible failures, silent data corruption occurs at significant rates.

Elerath and Shah discuss in detail the factors that can lead to the disparity between specifications and actual measured drive reliability [6, 15]. These include thermal variation, duty cycle, architecture and logic of the system in which the hard drive is used, and the data collection and analysis process of the failures.

While vendor-published failure-prediction metrics such as MTTF have been criticized by the research community [16, 17], it is important to understand how these metrics are derived. One intent of our work in detailing the accelerated life tests typically used by industry and the methodology used to predict storage media reliability using the test results is to help storage system designers better understand the limitations of data available about the disk medium. Because disks are already very reliable, meaningful data about actual performance is available only by studying very large populations of drives over very long periods. Because disk technology changes so rapidly, data about current products will inevitably be predictions based on accelerated life tests, not experience. The predictions will inevitably be based on models of how environmental factors affect performance.

13. Conclusions

Analysis of data from an accelerated life test predicts that, if data is written to 160GB RDX removable hard drive cartridges based on 2.5" laptop disk technology, and the cartridges are then stored in realistic conditions for 30 years, more than 99% of the drives will then read their entire contents with no errors. A small sample of 3.5" CSS drives included in the test demonstrated much lower data retention.

The related work above focuses on measuring, analyzing, and predicting failure rates of spinning hard drives, i.e., hard drives that were powered on, spinning, and in operation. To the best of our knowledge, this is the first published work on non-spinning media and thus complements these studies. Although its industry-standard methodology shares limitations with the tests on which manufacturers base their published specifications, the environmental effects in the idle case are much simpler. This should lead to more realistic performance predictions.

We hope this paper is a step towards getting storage system researchers and storage media reliability engineers talking the same language. Clearly, storage system designers need higher-quality, more timely data about the reliability of the components from which they must construct their systems. Equally, the difficulties in the way of storage media vendors providing such data are formidable.

References

1. Steve Kleiman, "Trends in Managing Data at the Petabyte Scale", 5th USENIX Conference on File and Storage Technologies (FAST), February 2007.
2. Eduardo Pinheiro, Wolf-Dietrich Weber, Luiz Andre Barroso, "Failure trends in large disk drive population", 5th USENIX Conference on File and Storage Technologies (FAST), February 2007.
3. Bianca Schroeder, Garth A. Gibson, "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? 5th USENIX Conference on File and Storage Technologies (FAST), February 2007.
4. W. Jiang, C. Hu, and Y. Zhou, A. Kanevsky, "Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage Subsystem Failure Characteristics", FAST 2008.
5. European Computer Manufacturers Association EMCA-379 Standard: Test Method for the Estimation of the Archival Lifetime of Optical Media, First Edition, June 2007.
6. J.G. Elerath, S. Shah, "Server class disk drives: how reliable are they?", Proceedings of the Annual Reliability and Maintainability Symposium, January 2004.
7. Reliasoft, Quantitative Accelerated Life Testing, Data Analysis Software, http://alta.reliasoft.com/
8. Reliasoft ALTA, Accelerated Life Testing Reference, © 1996-2001, pp179.
9. Percept Technology Labs, "RDX Removable Disk Archivability Study, Test Report", July 2007.
10. G. Cole, "Estimating Drive Reliability in Desktop Computers and Consumer Electronics Systems", Seagate Technology Paper TP-338.1, November 2000.
11. W. Nelson, Accelerated Testing: Statistical Models, Test Plans and Data Analyses, (Wiley Series in Probability and Mathematical Statistics-Applied Probability),1990, pp100.
12. P.F. Corbett, R. English, A. Goel, T. Grcanac, S. Kleiman, J. Leong, Sunitha Sankar, "Row-diagonal parity for double disk failure correction", Proceedings of Conference on File and Storage Technologies (FAST), 2004.
13. S. Ghemawat, H. Gobioff, S.T. Leung, "The Google file system", Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP), October 2003.
14. L. Bairavasundaram, G. Goodson, B. Schroder, A. Arpaci-Dusseau, R. Arpaci-Dusseau, "An Analysis of Data Corruption in the Storage Stack", FAST 2008.
15. J.G. Elerath, "Specifying reliability in the disk drive industry: No more MTBFs", Proceedings of the Annual Reliability and Maintainability Symposium, 2000.
16. J.G. Elerath, "AFR: problems of definition, calculation and measurement in a commercial environment", Proceedings of the Annual Reliability and Maintainability Symposium, 2000.
17. J. Yang and F.B. Sun, "A comprehensive review of hard-disk drive reliability", Proceedings of the Annual Reliability and Maintainability Symposium, 1999.