Drive-Independent Data Recovery: The Current State-of-the-Art

Charles H. Sobey, *Senior Member, IEEE*, Laslo Orto, and Glenn Sakaguchi

***Abstract*—The term "data recovery" herein refers to accessing logically and/or physically damaged storage media, for which no functioning backup exists. The state-of-the-art physical techniques for recovering data from failed hardware can all be described as "part replacement." To achieve high data density and high manufacturing yields, modern drives are "hyper-tuned" in the factory so that their data layout, zone frequencies, and various channel settings are optimized for each head, surface, and zone. This greatly complicates part replacement because a transplanted headstack, for example, no longer matches the servo, preamp, and read channel parameters that were optimized for the original headstack. Methods and challenges are discussed for replacing, or "refreshing," firmware and system area information and for replacing all of the drive's electronics. The data recovery industry, limitations of current techniques, and some probable future directions in data recovery are also presented. It is predicted that data recovery will be more important in the future as drives are exposed to more extreme mobile environments. Drive manufacturers may be able to differentiate themselves from their competition by designing for recoverability.**

***Index Terms*—Calibration, Computer crime, Data recovery, Defect management, Digital magnetic recording, Disk drives, ECC, Logical block address (LBA), Maintenance tracks, Optimization, PRML, Servomechanism, System area.**

I. Introduction

THE TERM "DATA RECOVERY" often refers to restoring or retrieving data (*i.e.*, files, blocks) from backup media. Depending on the field, data recovery also refers to the results of data mining; detecting data in waveforms (often involving phase-locked loops); decryption; or decompression. In this paper, data recovery refers to accessing logically damaged and/or physically damaged media, specifically from hard disk drives (HDDs), to obtain files or blocks that have no functioning backups – or are *themselves* backups.

Although the techniques for recovery from logical damage are interesting and challenging, they are more closely related to the operating system and the software programs used to create the data, not the HDD itself. The interested reader will find more information on logical recoveries in [1-3]. The techniques for recovering data from modern disk drives experiencing hardware failures have been shrouded in secrecy and are largely unpublished, except for [4]. In fact, the authors could not find any reference to peer-reviewed papers on the topic of hardware data recovery, as defined herein, except [5].

The closest other references are papers that describe magnetically imaging a disk by scanning an MR or GMR head over its surface [6-11]. This is presented as a way to analyze the written magnetization patterns, with the possibility of data recovery by generating a readback waveform from the image of the magnetization. In [4], the excessive requirements in time, storage, processing, and complexity of such a data recovery method are briefly discussed.

The reasons for the secrecy surrounding hardware data recovery techniques and the lack of public information include the desire to protect intellectual property (trade secrets) of data recovery companies, general lack-of-knowledge in the data recovery community about the inner workings of HDDs, misdirection (so that the true scope of recovery capabilities is not known to data saboteurs), and obscuring the often crude and simple nature of techniques that, up until recently, have exemplified the state-of-the-art.

These techniques for recovering data from physically damaged HDDs can all be described as *part replacement*. Printed circuit boards (PCBs) are swapped; heads are transplanted; motors and base castings are "replaced" by re-mounting the disks onto the spindle of a donor drive; and firmware or system information is replaced or "refreshed" by rewriting it. Placing the disks in a donor drive swaps everything – except for the on-disk system information, which is described in more detail later. Data stored on portions of the magnetic layer of the disk that have been physically removed, such as due to a slider (head) scraping away the surface as in Fig. 1, cannot be recovered – unless the future holds a way to reassemble the saved magnetic debris.

Data recovery is difficult now, and is getting more difficult. In order to simultaneously achieve higher data density and higher manufacturing yields, drives are "hyper-tuned" in the factory. This precisely matches the head/disk/preamp to other system components. These parameters are stored in the *system area* on the disk. Hyper-tuning is the reason some drive models almost always fail attempts at traditional part replacement. This trend is expected to continue, resulting in a requirement for drive-independent methods of data recovery for practically all drives built in the coming years.

Manuscript received July 12, 2005.

C. H. Sobey is with ChannelScience, Plano, TX 75024 USA (phone: 972-814-3441; fax: 972-208-9095; e-mail: csobey@channelscience.com).

L. Orto is with ActionFront Data Recovery Labs, Inc., Markham, Ontario L3R 2Z5 Canada (e-mail: laslo@actionfront.com).

G. Sakaguchi is with ActionFront Data Recovery Labs, Inc., Markham, Ontario L3R 2Z5 Canada (e-mail: sagu@actionfront.com).

Fig. 1. The dull ring near the middle diameter of this spinning disk is the result of a head crash. The headstack shown is a new replacement from a donor drive. The flex circuit at the top of the picture is connected to the old damaged headstack. The new headstack's flex circuit (lower right) is connected to circuitry that replaces the drive's electronics, see section V.

Furthermore, the need for data recovery is expected to grow. This is not because drives are being built more poorly. Instead, the networked lifestyle and the expectation that all of one's data are always available, even when mobile, will put massive amounts of information in more vulnerable places. For example, as disk drives continue to enter new non-traditional markets such as automobiles, cell phones, navigation, and personal mobile entertainment devices, HDDs will experience more and larger extremes of temperature, humidity, shock, vibration, and neglect. "Neglect" in the sense that drives employed in such uses are less likely to be backed-up as often as drives connected to home PCs or corporate data centers.

In addition, the superparamagnetic effect [12-17], which causes bits to decay with time, temperature, and external magnetic field (*e.g.*, from adjacent-track writing), also appears to be causing some recent data losses. The longer-term unknown effects of the current change from longitudinal to perpendicular recording may also result in additional, unexpected failure modes [18-21].

Due to the uniqueness of this topic in the HDD literature, the next section provides a brief overview of the data recovery industry. Section III defines and discusses the meaning of "drive-independent." This is followed by descriptions of the part replacement methods of data recovery: System area refreshing in section IV and drive electronics substitution in section V. The conclusion, section VI, discusses trends in the development of drive-independent data recovery techniques.

II. Overview of the Data Recovery Industry

It is no longer typical for HDD manufacturers to provide a list of "approved" data recovery companies. Most simply suggest performing a Google® search on the term "data recovery," which as of this writing generates over 3.1 million hits. This number has doubled in the past year. The latest trend, however, is to offer a branded data recovery service that is actually provided by a traditional, large data recovery company under the HDD or tape drive manufacturer's name.

Most people never need data recovery services. Some may need them only once in a lifetime. Large companies and government agencies might need such services once or twice a year – four or more times could be considered excessive. In this business environment, long-term relationships are difficult to build. Furthermore, without standards or industry certification the first-time customer of data recovery services has a difficult time selecting a reputable company or one that has successful experience with their particular issue.

Industry best-practices include never writing anything to the damaged drive. This is probably the most common mistake of do-it-yourself recovery attempts. Instead, the malfunctioning drive is "cloned"[1] as soon as possible. The cloned image on a new drive is then used for all data recovery procedures. Also, most reputable companies do not charge if they cannot recover the customer's data. Some do charge "attempt fees" or evaluation fees; some provide priority service (including on-site support) at additional cost.

The cost of a typical recovery is from US$500 to US$2 500, depending on the file system (servers, RAID arrays, and Unix are more expensive), the type of repair, the time required, the type of data required (text is easier to recover than corrupted databases), and the probability of success (and hence of payment). Some extreme cases, including those requiring on-site support, can cost tens of thousands of dollars. Worldwide, it is estimated that the market for data recovery in 2005 will be just over US$100 million [22]. The top five international recovery companies make up about three quarters of this market. The remaining market is divided up among thousands of small local data recovery service providers, most of whom provide only logical recoveries.

III. Drive-Independence

In the strictest sense, we define drive-independent data recovery as recovering the user's data without any hardware from the original drive (except the media, of course) and without any firmware or other information from the original drive, with the possible exception of its model number. In this case, disk drive parts from donor HDDs may be used and similar drives may be examined to ascertain critical information about the drive family. A portion of ActionFront's inventory of donor drives is shown in Fig. 2.

In a broader sense, drive-independence also refers to being able to effectively utilize part replacement independent of the drive model, regardless of the part(s) that needs replacing. Drive-independent data recovery, then, refers to a collection of techniques that works across drive families, as opposed to one-off patches and tricks that can only be applied to specific problems with specific drive models.

The ultimate part replacement operations are re-mounting disks onto new drives and transplanting headstacks. In these

[1] Cloning, also called "mirroring," refers to using special software to copy an exact image of every logical block from one (failed) drive to the same logical blocks on another (good) drive, which is usually a different model.

Fig. 2. A very small portion of the inventory of donor drives required for robust part replacement and experimentation.

two extreme cases there are six difficult challenges to overcome for successful data recovery:

1. Re-optimize preamp read settings;
2. Recalibrate repeatable run-out (RRO) and head offsets;
3. Control spindle rotation and head positioning, typically using the magnetic servo patterns on the disk surfaces;
4. Determine the layout and format of each surface, defects, and defect mapping strategies [23, 24];
5. Detect the binary data in the analog head signal; and
6. Decode the precoding, scrambling, RLL, parity-assist, ECC, and any other codes to reveal *user* data [25, 26].

Of course, the sectors or blocks created from the detected and decoded user bits must still be assembled into useful files. It is at this latter task where logical recoveries typically start. Interestingly, data forensic examinations can only begin after the physical and then the logical recoveries have been completed [27]. In this paper, we examine two drive-independent data recovery techniques that span both the strict and the broad interpretations of the term: 1) Replacing or "refreshing" the system area information (broad sense) and 2) Replacing the drive's electronics (strict sense).

IV. Refreshing the System Information

"System information" includes the drive-specific hyper-tuned parameters mentioned in section I. The system information, or the area where it is located, goes by many names, including system area, maintenance tracks, negative cylinders, reserved cylinders, calibration area, initialization area, and diskware. These are always areas that are not directly accessible over the ATA/IDE or SCSI interface. Often they are located at the extreme outer diameter (OD); however, they have been found at other radii as well.

A typical drive's boot-up procedure begins at power-on with electronic sub-system self-checks; then spindle motor spin-up; headstack unlatching or unloading; initial acquisition of servo wedge timing; seeking to the system area; and reading in its information for additional boot procedures and/or needed parameters and firmware. Clearly if the system area information is corrupted, the drive is not likely to function. For this reason, many manufacturers store multiple copies of the system area information, often one (or more) per surface. Furthermore, if a module of this information is corrupt in one copy, a good copy of the module (*i.e.*, one with a valid checksum) from another area can be used.

The system area may become corrupted due to malfunctioning circuits, firmware bugs, exceeding the operational shock specifications of the drive, or position system errors. Another, more common, reason for system area corruption is a loss of power during an update of the system area itself. This might occur when system logs are being updated or when the G-list is being changed. The G-list, or grown defect list, holds information about the location of defects that have been found in the field during drive operation. The G-list is typically used for sector swapping, or sector reallocation [23, 24]. Related to this is the P-list, or primary defect list, that stores the location of media defects that were found during manufacturing. This is typically used for sector slipping and is not updated in the field.

Corrupted system area information can be rewritten as a form of part replacement. For some drive models, the system area contains only a small amount of information, such as a unique drive serial number, the P-list and G-list, often a "translator" that converts between logical and physical block addresses including the effects of head and track skewing, S.M.A.R.T. data [28], and a (possibly encrypted) drive password. This small amount of drive-specific information usually indicates that the drive is more amenable to part replacement, including head transplantation. An example of the identification portion of the system area of a 2.5″ Hitachi drive is shown in Fig. 3.

Some drive models have larger system areas, which may span tens of tracks. This typically indicates that a drive employs hyper-tuning and hence is much less amenable to traditional part replacement, especially head transplantation and system area refresh. Their system areas contain all of the information listed above, plus some or all of the following: program overlays (executable code) for seldom-used functions or functions subject to revision; drive-specific tables such as RRO compensation, writer/reader offsets, data rates, zone table, many read channel parameter settings, gains, bias currents, and servo parameters; test routines; calibration routines; factory defaults; system logs; and extensive details about drive components.

It is possible that a drive submitted for data recovery because it will not initialize has hardware that is completely intact, therefore requiring no part replacement. Approximately 30% of the time this "hardware" failure is due to system area corruption. In these cases, the system area information is the part that needs to be "replaced." Rewriting it is the fastest, most-reliable, lowest cost data recovery method (other than restoring a recent working backup to a different drive).

The format of the system area is not published and must be determined for each drive family. This information is commercially available only for a few drive models [29]. Some data recovery companies augment such information with their own investigations of other drives. Fig. 4 shows a screen capture of such an in-house tool, displaying information from the system shown in Fig. 3.

If the system information is rewritten from archival copies of the data from other similar drives, the hyper-tuned parameters will not match those needed for the drive's original components. These head-specific parameters must be re-optimized [30] and rewritten to the system area, just as if a head transplant had been performed. Parameter re-optimization is not supported by any commercial utility and is an area of ongoing research for data recovery companies.

V. Replacing the Drive Electronics

If it were available, system area refresh that re-optimizes key parameters as needed for headstack transplants or disk removal would be the preferred method for drive-independent data recovery. However, there are data recovery situations for which even this powerful method would be inadequate. For example, many consecutive servo sectors could be corrupted causing the drive to shut down or to constantly recalibrate; precise control of headstack placement may be necessary to bypass badly damaged portions of the disk; the commands to control the read channel parameters for re-optimization may not be known; it may be necessary to capture system area information from a drive that will not spin-up or will not initialize; it may be necessary to read data from normally inaccessible locations (*e.g.*, passwords in the system area, defective sectors, spare sectors); decaying bits or poorly written transitions may require additional off-line signal processing to be read adequately; and it may be necessary to bypass the normal drive boot-up sequence, especially if the sequence is written to initiate automatic encryption or destruction of data as a security measure. Extreme cases in which the disk is badly damaged and not flyable, such as those with which the defense, intelligence, and law enforcement communities may be involved, will not yield to recovery methods that involve drive hardware with specific servo signal, timing, zone layout, or frequency requirements.

In these cases and others, it is desirable and often necessary to replace all of the control, signal processing, and decoding functions of the drive with systems that are completely under the control of the data recovery specialist. For flyable media, the most cost-effective way to spin the disk is with its original motor and base casting or with from of a donor drive. All that is required is a standard HDD motor controller and related programming capability.

Once a compatible headstack is in place and the disks are spinning, the signal from the preamp needs to be acquired and used: first for servo positioning and then for data detection. To acquire a good signal, the read bias currents must be approximated for each head. In general, it can never be assumed that proprietary specifications for any chip or drive are available to the data recovery company. Therefore, even the programming of the preamp must be determined by experimentation. Often this involves probing the serial programming interface to the preamp while issuing various read and write commands. After the preamp is configured, the readback waveform can be digitized.

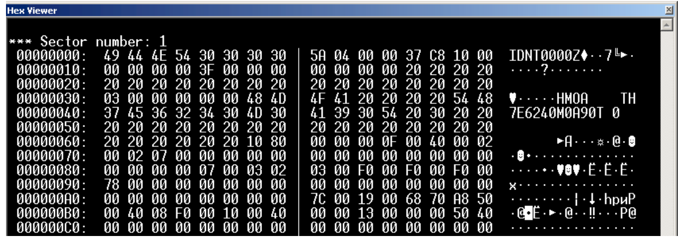

Fig. 3. The beginning of the identification sector ("IDNT") of the system area for a 2.5″ Hitachi drive is shown. The left column is the offset index from the beginning of the sector; the next two wide columns contain the hexadecimal interpretation of the data stored there. The rightmost column shows the ASCII equivalent of the hex values, when an equivalent exists. The model number is interpreted as "HTE726040M9AT00."

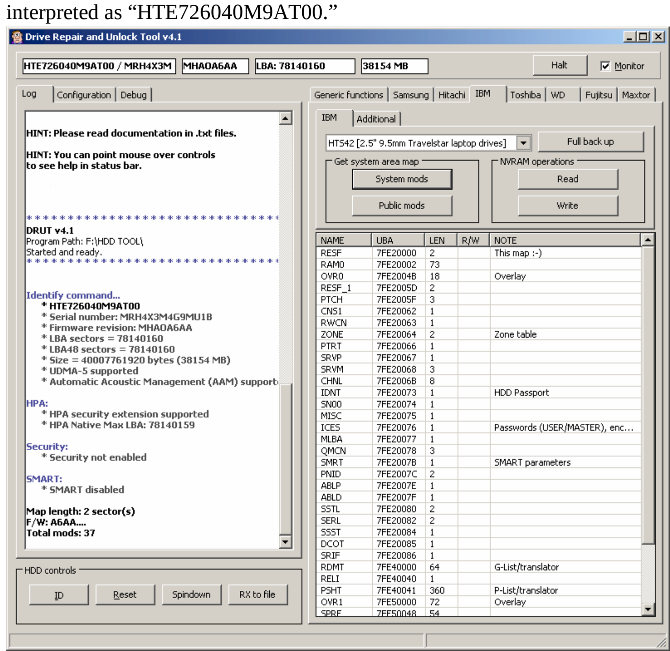

Fig. 4. Example of an in-house system area reader, *Drive Repair and Unlock Tool*. The pane on the left displays basic identifying information from the drive's system area, such as that shown in Fig. 3. The right pane lists the locations of various pieces of important system are information. Notable entries include the "Zone table," "Passwords," "SMART parameters," and the G-list and P-list. The "Overlay" entries typically indicate executable firmware.

Complete sectors of *channel* bits from a wide variety of disk drives, and even from tape and optical drives, are routinely recovered from digitized waveforms by commercial channel models, such as [31]. However, for greater speed, these functions must be moved from the general purpose PC to specialized hardware. Recovery of *user* bits (converted to ASCII data) in this manner, using ActionFront's SignalTrace™ system, was first reported publicly in [4].

Fig. 5 shows the current version of the SignalTrace™ PCB. The analog waveform from the headstack's preamp is passed through a differential probe to a low-noise preamplifier and variable gain amplifier (VGA) before being fed into a 10-bit 2Gsample/s analog-to-digital converter (ADC). The ADC typically oversamples the readback waveform at 6 to 8 times above the drive's data rate for a particular zone. This oversampling ratio is controlled for the highest data rate drives by slowing down the spindle speed appropriately.

The ADC's output is synchronized, via a master clock, with a 1:8 demultiplexer. This brings the 2Gsample/s data stream down to eight parallel 8-bit bytes of data at 250 MHz. Only the

Fig. 5. Latest generation of SignalTrace™ hardware, which replaces most drive electronics functions. Note: Heat sinks obscure the FPGA and DMUX.

most significant 8 bits of the ADC's 10 bits are used because 8-bit resolution is usually sufficient (HDD read channels typically have no more than 6-bits of resolution).

The 8 bytes of data are fed into an FPGA that has been programmed to perform filtering, timing extraction, and data detection. The read channel's continuous-time filter is approximated by a 32-tap FIR, whose output is resampled at the data rate under the control of a phase-locked loop algorithm similar to that in [25]. The filter coefficients are initially set based on a minimum mean squared error (MMSE) criterion, assuming a PR4 or EPR4 target response as appropriate for the zone. A 5-tap adaptive FIR equalizes the samples that are then sent to an implementation of a configurable-target Viterbi detector.

However, before any data sectors can be detected, the FPGA must acquire the timing of the servo sectors by locating occurrences of the servo preamble frequency, and sometimes also of a specific gap or byte pattern. Once found, a timing window is set up for repetitive detection of the servo sectors.

The servo timing, sync marks, gray encoded track ID, wedge ID, and burst configuration all must be determined experimentally by examining a working drive of a similar model and date code. Based on this, timing windows are then set up to detect the servo track ID and wedge ID information, which are usually stored as di-bits. As in read channel chips, as much of the signal processing for the data detection as possible is re-used for detecting the servo information. Once the track ID has been detected reliably, additional timing windows are set up to control the integration of the bursts.

The burst values and track ID are sent to a Motorola Coldfire processor that applies a servo control algorithm to determine the appropriate position error signal (PES). The resulting PES control signal is sent to the voice coil motor driver. Seek time is not an important consideration for data recovery applications, so most emphasis is placed on accurate track following. For example, tracks with too many servo sectors that contain errors might be still be readable using non-traditional algorithms, such as applying servo corrections that were calculated for other tracks.

Once stable tracking has been achieved, the data areas are captured and buffered and processed as described above. The detected channel bits are passed through FPGA implementations of the sync detector, descrambler, RLL (run-length limited) decoder, and ECC (error correction code) decoder. If the ECC indicates that uncorrectable errors are present, the sector is retried. If necessary the filter and detector parameters are adjusted on-the-fly for additional retries.

The decoding steps listed above are very difficult to determine. They are proprietary and largely unpublished. Often the HDD manufacturers themselves do not know the details of these read channel blocks. Therefore, first principles must be used on each drive model to determine the decoding implementations. For example, the sync mark (end-of-preamble mark), and possibly a redundant sync mark, must be located in the Viterbi output decisions by analyzing the expected sector size based on coding overhead and the expected position of the sync mark(s).

The descrambler and RLL decoding are especially difficult to determine when the corresponding read channel specification is not available. The techniques used to solve this problem provide a key competitive advantage and are not described herein. The location of any split sectors (data sectors that are interrupted by a servo sector) must be determined for each zone. These sector portions are then reassembled. When parity-assist post-processing is used, the parity bits are typically stripped off with the expectation that Reed-Solomon error correction, in conjunction with sector re-reads, will correct any data errors.

As a specific example, ActionFront's SignalTrace™ system completes all of the functions listed in this section for a 2002 Western Digital drive, with a capacity of 40GB/platter, at a rate of approximately 2GB/hour. Mapping the physical sectors into logical sectors, including the effects of skewing and defect management, is completed in a separate step.

VI. Conclusion

Data recovery is difficult now, and is getting more difficult. The hyper-tuning that simultaneously enables higher data density and higher yields causes the data recovery industry's traditional hardware repair method of part replacement to fail in more drive models. While some drive models currently have recovery success rates above 90%, and others are above 60%, an increasing number have practically no chance of recovery for most part replacements. For this reason drive-independent data recovery is needed and its capabilities must be enhanced.

The current state-of-the-art research for system area refreshing focuses on developing algorithms that can quickly and adequately re-optimize all important channel, preamp, and servo system parameters *without writing over data*. This capability is needed both when the system area information is corrupted and when a headstack transplant is necessary.

The current state-of-the-art research for drive electronics replacement focuses on developing faster and more robust methods for determining the servo sector track ID and wedge ID encoding and the data sector encodings. Additionally, timing, equalization, and detection methods are being advanced to recover data from the drives that are being built today and in the future. These are likely to employ iterative equalization and decoding, LDPC (low-density parity-check) codes, and new timing recovery schemes [32-34].

Similar drive-independent data recovery techniques can be applied to magnetic tape and optical media, as well. When these media are discarded, sensitive information is often vulnerable to malicious recovery [35]. This is true even when files have been "deleted" [36, 37].

If a recovery is to be attempted from media that are not flyable, due to excessive damage including bending, abrasion, corrosion, fragmentation, penetration, etc., the layout information, signal processing, and decoding methods discussed herein must still be used. However, other means for recovering a readback signal from the recorded pattern are needed. These methods are always time-consuming, extremely difficult, expensive, and have a low probability of success. This is another area of current research, primarily for intelligence investigations of national security importance, not criminal investigations nor corporate data recovery.

Drive manufacturers can take steps to improve the recoverability of their products. These include publishing the specifications for their system area information and their encoding schemes, and providing a way for their in-house optimization routines to be run in the field. It is interesting to speculate that a drive's reputation for recoverability might be a differentiator of considerable value in some markets.

Acknowledgment

The authors thank Nick Majors, Dmitry Postrigan, Christian Sorensen, and Lucian Vanta, all with ActionFront Data Recovery Labs, Inc., for many helpful discussions.

References

[1] M. Dowler, "Beginners Guides: Hard Drive Data Recovery," *PC Stats*, August 16, 2003, http://www.pcstats.com/articleview.cfm?articleID=1139.

[2] Various Authors, "File Allocation Table", *Wikipedia*, http://en.wikipedia.org/wiki/File_Allocation_Table.

[3] Computer Forensics Tool Testing (CFTT), National Institute of Standards and Technology, http://www.cftt.nist.gov/.

[4] C.H. Sobey, *Recovering Unrecoverable Data: The Need for Drive-Independent Data Recovery*, April 14, 2004, available for free download at http://www.ActionFront.com.

[5] S.K. Kasanavesi, T.D. Milster, D. Felix, and T. Choi, "Data Recovery from a Compact Disc Fragment," *Proc. SPIE*, September, 2004, pp. 116-127.

[6] R.S. Indeck and J.H. Judy, "Measurements of Surface Magnetic Fields of Perpendicular and Longitudinal Magnetic Recorded Transitions Using a MR Transducer," *IEEE Trans. Magn.*, Sepetmber 1984, pp. 730-732.

[7] R. Gomez, A. Adly, I. Mayergoyz, and E. Burke, "Magnetic Force Scanning Tunneling Microscope Imaging of Overwritten Data", *IEEE Trans. Magn.*, September 1992, pp. 3141-3143.

[8] S.Y. Yamamoto, D.C. Vier, and S. Schultz, "High Resolution Contact Recording and Diagnostics with a Raster-scanned MR Head," *IEEE Trans. Magn.*, Septermber 1996, pp. 3410-3412.

[9] S.Y. Yamamoto, S. Schultz, Y. Zhang, and H.N. Bertram, "Scanning Magnetoresistance Microscopy (SMRM) as a Diagnostic for High Density Recording," *IEEE Trans. Magn.*, January 1997, pp. 891-896.

[10] I.D. Mayergoyz, C. Serpico, C. Krafft, and C. Tse, "Magnetic Imaging on a Spin-Stand", *J. Appl. Phys.*, Vol. 87, No.9, May 1, 2000, pp.6824-6826.

[11] I.D. Mayergoyz, C. Tse, C. Krafft, and R.D. Gomez, "Spin-stand Imaging of Overwritten Data and Its Comparison with Magnetic Force Microscopy," *J. Appl. Phys.*, 1 June 2001, pp. 6772-6774.

[12] S.H. Charap, P.L. Lu, and Y. He, "Thermal Stability of Recorded Information at High Densities," *IEEE Trans. Magn.*, January 1997, pp. 978-983.

[13] Y. Zhang and H.N. Bertram, "Thermal Decay of Signal and Noise in High-density Thin Film Media," *IEEE Trans. Magn.*, September 1999, pp. 4326-4338.

[14] N. Li, F. Meyer, and B. Zak, "Drive Level Media Thermal Stability Test," *IEEE Trans. Magn.*, July 2001, pp. 1547-1549.

[15] R.M. Brockie and H.J. Richter, "Thermal Decay Study of Various Magnetization Patterns," *IEEE Trans. Magn.*, September 2002, pp. 1946-1948.

[16] Z. Li, G. Sandler, E. Champion, and S. Li, "Adjacent Track Erasure Analysis and Modeling at High Track Density," *IEEE Trans. Magn.*, September 2003, pp. 2627-2629.

[17] J. Feng, "Long-Range Erasure of Nearby Recorded Tracks," *IEEE Trans. Magn.*, July 2004, pp. 2588-1590.

[18] R. Wood, "The Feasibility of Magnetic Recording at 1Terabit per square inch," *IEEE Trans. Magn.*, January 2000, pp.36-42.

[19] R.D. Cideciyan, E. Eleftheriou, and T. Mittelholzer, "Perpendicular and Longitudinal Recording: A Signal-Processing and Coding Perspective," *IEEE Trans. Magn.*, July 2002, pp. 1698-1704.

[20] W.R. Eppler *et al.*, "Perpendicular Recording Near 100Gb/in^2," *IEEE Trans. Magn.*, March 2003, pp. 663-667.

[21] B.R. Acharya, J.N. Zhou, M. Zheng, G. Choe, E.N. Abarra, and K.E. Johnson, "Anti-Parallel Coupled Soft Under Layers for High-Density Perpendicular Recording," *IEEE Trans. Magn.*, July 2004, pp. 2383-2385.

[22] Internal data from ActionFront Data Recovery Labs, Inc., March 2005.

[23] US Patent 5,812,755, F.A. Kool and J.S. Packer, "Logical and Physical Zones for Management of Defects in a Headerless Disk Drive Architecture," Sept. 1, 1995.

[24] US Patent 6,212,647B1, J.R. Sims, III and K. Way, "Systems and Methods to Perform Defect Management to Block Addressable Storage Media," April 3, 2001.

[25] R.D. Cideciyan, F. Dolivo, R. Hermann, W. Hirt, and W. Schott, "A PRML System for Digital Magnetic Recording," *IEEE Journal on Selected Areas in Communications*, January 1992, pp. 38 - 56.

[26] R.D. Cideciyan, J.D. Coker, E. Eleftheriou, and R.L. Gailbraith, "Noise Predictive Maximum Likelihood Detection Combined with Parity-Based Post-Processing," *IEEE Trans. Magn.*, March 2001, pp. 714 – 720.

[27] D. Farmer and W. Venema, *Forensic Discovery*, Addison-Wesley, 2005.

[28] Additional information on S.M.A.R.T (Self-Monitoring, Analysis, and Reporting Technology) is available at http://www.t13.org.

[29] *PC-3000* and *Data Extractor* from ACE Laboratories, Ltd., Rostov-on-Don, Russia; http://www.acelab.ru/products/pc-en.

[30] US Patent 6,252,731 B1, B.A. Sloan, K.L. Enarson, and S.A. Wakefield, "Parametric Optimization Using Disc Drive Read Channel Quality Measurements," June 26, 2001.

[31] PRML*pro*™ channel model, trial version available for free download at http://www.ChannelScience.com.

[32] R. Koetter, A.C. Singer, and M. Tuchler, "Turbo Equalization," *IEEE Signal Processing Magazine*, January 2004, pp. 67-80.

[33] H. Song and B.V.K.V. Kumar, "Low-Density Parity-Check Codes for Partial Response Channels," *IEEE Signal Processing Magazine*, January 2004, pp. 56-66.

[34] J.R. Barry, A. Kavcic, S.W. McLaughlin, A. Nayak, and W. Zeng, "Iterative Timing Recovery," *IEEE Signal Processing Magazine*, January 2004, pp. 89-102.

[35] S. Garfinkel and A. Shelat, "A Study of Disk Sanitization Practices," *IEEE Security and Privacy*, January-February, 2003.

[36] P. Gutmann, "Secure Deletion of Data from Magnetic and Solid-State Memory," *Sixth USNIX Security Symposium Proceedings*, July 22-25, 1996.

[37] S. Bauer and N.B. Priyantha, "Secure Data Deletion for Linux File Systems," *Proceedings of the 10th USENIX Security Symposium*, August 13-17, 2001, pp. 153-164.