Towards compliant data retention with probe storage on patterned media

Pieter Hartel Leon Abelmann [*]

Abstract

We describe how the compliance requirements for data retention from recent laws such as the US Sarbanes Oxley Act may be supported by a tamper-evident secure storage system based on probe storage with a patterned magnetic medium. This medium supports normal read/write operations by out-of-plane magnetisation of individual dots. We report on an experiment to show that in principle the medium also supports a separate class of write-once operation that destroys the out-of-plane magnetisation property of the dots irreversibly by precise local heating. The write-once operation can be used to support flexible data retention by tamper-evident writing and physical data deletion.

1 Introduction

Laws such as the US Sarbanes-Oxley Act of 2002 (SOX) [15], or the EU data retention directive [21] stipulate that business critical information must be retained for a precise amount of time, after which it must be destroyed. The act of destroying the information must be recorded, and all cached copies, backups and any derived information that could be used to reconstruct the business critical information, must be destroyed as well. Information that was destroyed too early ultimately lead to the collapse of Enron and its accountant Arthur Anderson. Information that is destroyed too late is a minefield of liability [7]. In a 2005 study [3], 216 respondents from large and small companies (65% from the US, and 35% from the rest of the world) stated that SOX and data retention are the two major issues their organisations are struggling with.

Not only do modern organisations have to deal with compliance, their workforce is also increasingly mobile. Consider the following scenario of a loyal employee who processes and stores customer data that is subject to compliance regulations on her laptop. She takes her laptop with her when she visits customers. The main threat is someone tampering with the information on her laptop. The risk of tampering is particularly acute when she is out of the office because it is difficult always to keep an eye on your belongings in airports, hotels, at customer sites etc.

Assume that to mitigate the risk, the organisation has set the following inefficient (but effective) data retention policy: (1) a full backup of all business critical information must be made before leaving the office, (2) a secure hash of the information must be carried as a separate item (for instance printed on a piece of paper). Each time the employee suspects foul play (to be confirmed by checking whether the hash of the information still matches the hash on the piece of paper), she restores the laptop to the state it was in when she left the office, and (3) the information and the backups are deleted when the retention period expires.

Problem The data retention policy sketched above is inefficient for at least three reasons. Firstly, making and restoring full backups can take a long time. Secondly, a backup represents yet another copy of the information that is subject to compliance regulations. Thirdly, the employee risks loosing all the changes that she made to the information after leaving the office.

We propose a partial solution to these problems by combining an idea of Molnar et al. [12] for *tamper-evident storage* with our own work on *probe storage* on *patterned media*. We introduce each of these three elements below.

Molnar et al. [12] describe how standard Programmable Read-Only Memory (PROM) can be used to build *tamper-evident storage*. The basic idea is to store each bit of information in a cell occupying two bits using Manchester encoding: the bit 1 is encoded as the cell 10 and 0 is encoded as 01. The value 11 indicates a cell that has not yet been used (all bits in a PROM are initialized to 1). The value 00 indicates a cell that has been tampered with for the following reason. The physical properties of a PROM make it *impossible* to change a 0 back into a 1 (except by exposing the entire memory module to ultra violet light, which would reset the entire module). Therefore, the only way to tamper with information (which is encoded as 01 or 10) is to clear a bit. This immediately results in an invalid code 00, which provides the evidence of tampering.

A *Probe storage* device with a patterned medium does not exist at present, although an advanced prototype of the Millipede, a probe storage device with uniform media capable of storing over 641 Gbit per square inch has been demonstrated by IBM research [17]. The Millipede basically consists of a voice-coil based actuator, a polymer medium and an MEMS array of 64×64 read/write probes with tips as sharp as 10 nm.

A *patterned medium* [23] consists of a regular arrangement of magnetic dots separated by sub-micron distances that can be magnetised in two out-of-plane directions. The magnetic dots can be read and written any number of times. However, the out-of-plane magnetisation property of a dot can also be irreversibly destroyed by precise local heating. The idea is to use this feature to create a storage device that operates as a RAM, until part of the space is frozen. From then on the frozen part of the device operates as a tamper-evident store, while the rest continues to operate as a RAM. The ability to freeze parts of the store incrementally provides flexibility that cannot be matched by CD, DVD, (EE)PROM or any other current technology.

[*] Faculty of Electrical Engineering, Mathematics and Computer Science, Univ. of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands, email {pieter.hartel, leon.abelmann}@utwente.nl

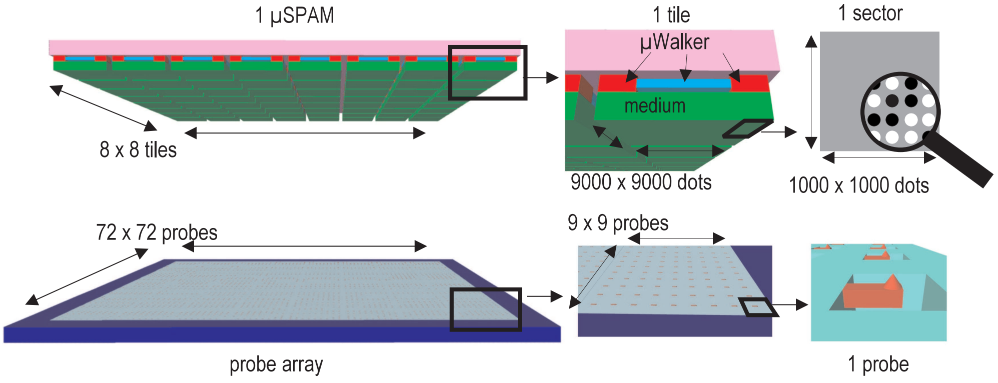

Figure 1: µSPAM approach to Probe storage on patterned media.

Contribution The contributions of the paper are (1) to study the feasibility of the freeze operation itself and then (2) to show how probe storage on patterned medium with a freeze operation might be used to support flexible data retention policies.

The paper is organised as follows. Related work is discussed in the next section. A brief introduction into our approach towards realising probe storage on patterned media is given in Section 3. Section 4 describes an experiment to measure the temperature that must be achieved to destroy the out-of-plane magnetic property of a dot. A proposal for a tamper-evident storage and deletion is presented in Section 5. The last section concludes.

2 Related work

During the last ten years, several recording systems based on probe microscopy technology have been proposed. The IBM lead already presented above is followed by other companies such as HP [13], Samsung [10], Seagate [6], LG [5] and a number of universities such as Exeter, Carnegie Mellon, DSI Singapore, Tohoku, Yonsei, and Twente. Probe storage is also being combined with disk storage; Hong et al. [2] propose a hybrid disk / probe storage architecture where probe storage bricks are used to extend the MTBF of the system, estimated 23 years.

Our proposal supports data integrity but not confidentiality and authenticity, for which cryptographic tools would be needed. Not using cryptography in our proposal avoids all troubling issues of key management. Kher and Kim [4] provide a good survey of crypto based distributed secure storage.

3 Probe storage on patterned media

The design for the Twente Micro Scanning Probe Array Memory (µSPAM) (Figure 1) is made out of two or more silicon wafers bonded to each other. One half contains the (magnetic) medium. An electrostatic stepper actuator, such as the µWalker [20] or Harmonica drive [19] is used to move the medium. The other half consists of one large array of probes.

The patterned medium The medium for the µSPAM is a regular matrix of magnetic single domain dots. Such a discrete medium is expected to be able to support higher bit densities compared to the continuous polycrystalline media used in the hard disk today [22].

We have already shown that a matrix with a period of 200 nm can be achieved [24]. An scanning electron microscope and magnetic image is shown in Figure 2. An improved setup with periodicities down to 150 nm has recently been realised [8]. We are currently aiming at a period of 100 nm, being 50 nm dot size and 50 nm spacing. This will give a capacity of 10 Gbit/cm^2 (=65 Gbit/inch^2).

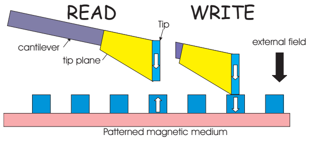

Figure 3: The principle of an MFM-measurement

The probes For reading, the MFM (Magnetic Force Microscopy)-principle has been chosen [16]. An MFM-probe is made by placing a small magnetic element, the tip, on a cantilever spring. Typical dimensions are a cantilever length of 200 µm, element length of 4 µm and diameter of 50 nm and a distance from the surface of 30 nm.

Figure 3 shows the principle of an MFM-measurement. The magnetic tip is attracted or repelled, depending on the stray field of the medium. The tip is affected by the magnetic orientation of a dot. The displacement of the cantilever might be measured by means of the change in capacity between the cantilever and the medium. Bits can be written by the combined field of the magnetic tip and an externally applied field generated by a coil placed on the backside of the medium [24].

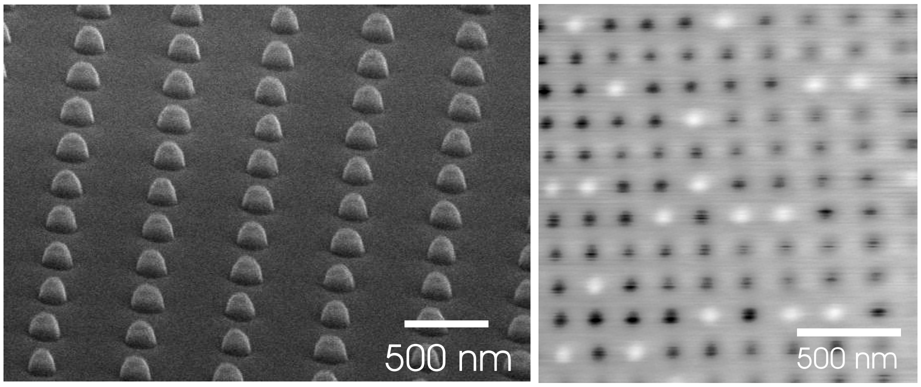

Figure 2: Scanning Electron Micrscope (left, 500 nm pitch) and Magnetic Force Microscopy image (right, 200 nm pitch) of patterned media.

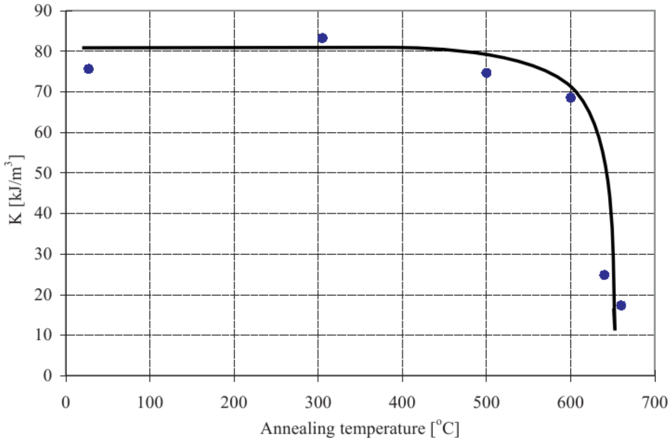

Figure 4: Anisotropy as a function of the annealing temperature

4 Change in anisotropy by annealing

To support the freeze operation, it should be possible to destroy the out-of-plane magnetic property of the dots. The magnetic material used in the dots consists of a stack of extremely thin Co and Pt layers, each no more than 1 nm thick [24]. The interfaces between the Co and Pt film cause an easy axis of magnetisation perpendicular to the interfaces. Due to the delicate structure of these films, they do not support high temperatures over long periods of time. Above a certain temperature, the interface between the Co and Pt mixes, and the perpendicular anisotropy of the film is destroyed. As a result the easy axis of magnetisation turns in the film plane. This is an irreversible process. After heat treatment, the interfaces cannot be restored.

To determine at which temperature interface mixing occurs, we have measured the anisotropy constant K of samples subjected to six different temperatures. The anisotropy constants were calculated by a Fourier transformation of the torque curve obtained with an applied field of 1350 kA/m. Figure 4 shows the dependence of the anisotropy value as a function of the annealing temperature. The anisotropy of the unannealed film is 80 kJ/m^3. This value is maintained up to an annealing temperature of 500 °C. Above 600 °C the value of K drops dramatically.

The unannealed sample (a) and (b) a sample annealed at 700 °C were investigated by low angle X-ray diffraction (XRD) to investigate the interfaces (Figure 5). A peak around 8 degrees on the 2 θ axis is visible on the sample without annealing. This peak is due to the periodicity of the Co-Ni and Pt multilayers. From the XRD measurement we can determine that each layer has a thickness of 0.6 nm. In the annealed sample, this peak has disappeared. From these measurements we can conclude that after an annealing treatment at a temperature higher than 600 °C, the interfaces have mixed, the perpendicular anisotropy is lost and the out-of-plane magnetic properties of the film are destroyed.

By measuring high-angle-XRD, the crystal structure of the films can be investigated (Figure 6). In the annealed sample, we can find a strong reflection peak around 41.7 degrees in 2 θ axis. This peak can be characterized to a fct Co-Pt (111) plane [25]. It suggests that a new crystalline structure of fct Co-Pt was formed in the film. This crystal exhibits magnetic anisotropy in the [001] direction, i.e. there are tilted magnetic easy axes in the film (not perpendicular, not in plane). So there is no risk that after excessive heating the perpendicular anisotropy can be restored by crystallisation.

Further research is needed to determine the time required, the amount of energy dissipated, the wear on the tip, and the effect of heating one dot on the neighbouring dots. The energy must be conducted off the medium for it not to bend or warp, and to avoid increasing the error rate of the perpendicular read processes due to thermal noise. It is not unlikely that by tailoring the materials and layer structures, the interface mixing temperature can be reduced. In any case it will be necessary to use the write-once operation sparingly.

5 Tamper-evident writing and deletion

Having described the most important aspects of the physical realisation of probe storage on patterned media, we turn our attention to the system aspects. We argue that the compliance inspired data re-

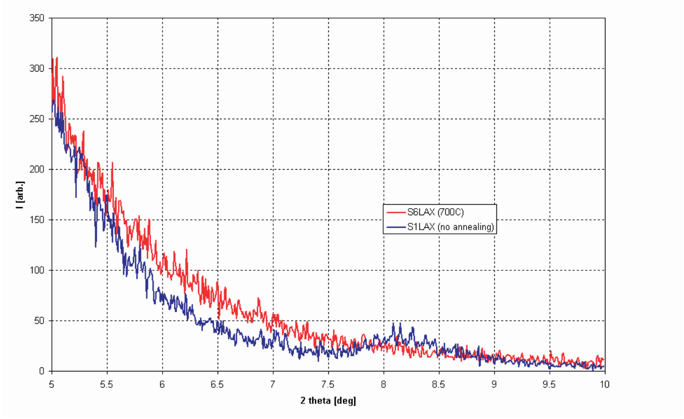

Figure 5: X-ray diffraction under a low angle of two samples, one with and one without annealing.

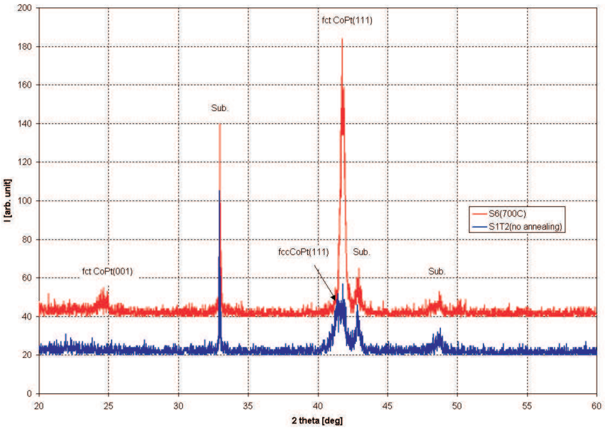

Figure 6: X-ray diffraction under a high angle of the same two samples as for low angle XRD.

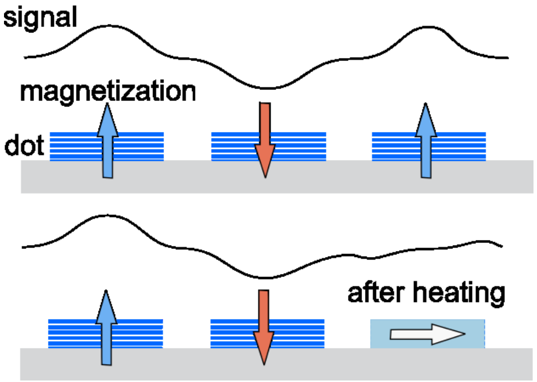

Figure 7: Above: dots are magnetised upwards or downwards; Below: dots have a perpendicular or in-plane easy axis.

tention requirements can be served flexibly by a small set of operations. The first operation writes the information to the medium in such a way that physical and logical tampering can be detected by a second verify operation. The third operation physically deletes information. We describe a set of low level bit operations first, followed by the description of the high level sector operations to freeze, verify, and delete.

Magnetic read and write a bit We require two read and two write operations of bits. In the normal mode of operation we have a medium with a regular matrix of magnetic single domain dots with a preferred axis of magnetisation perpendicular to the medium. This is illustrated in the top half of Figure 7, which shows the substrate and three layered dots. The first and last are magnetised in the upwards direction, the middle in the downwards direction. The magnetic write bit operation `mwb` sets the direction of the magnetisation (up is 1, down is 0) and the corresponding read `mrb` senses the direction of the magnetisation.

Electrical read and write of a bit By heating an individual dot by means of an electric current flowing from the probe tip via the dot to the medium, the multilayer structure of a dot is destroyed and as a result the easy axis of magnetisation rotates into the medium. The data stored by `mwb` is lost and there is no correlation between the original out-of-plane direction of the magnetisation and the in-plane direction. However, we have a new way of representing data: the initial state of a dot with out-of-plane magnetisation is 1, `ewb` writes 0 to a dot by setting the magnetisation in-plane. The bottom half of Figure 7 shows that the layered structure of the last dot is permanently destroyed. The electrical write bit `ewb` is thus an irreversible process, which can only change a 1 into a 0. We are confident that even a skilled focused ion beam (FIB) operator would find it difficult to reconstruct a perfect out-of-plane dot because she would have to remove the debris of an in-plane dot first, and then deposit several thin Co and Pt layers in a sub-micron area just to reconstruct one dot. Using magnetic imaging techniques [9], a forensics team would probably not have any difficulty identifying a reconstructed out-of-plane dot from an original out-of-plane dot.

The read operation `erb` detects the presence or absence of an out-of-plane dot by performing an atomic sequence of `mrb` and `mwb` operations as follows: (1) `mrb` original, (2) `mwb` inverse of the original, (3) `mrb` and check that the inverse has indeed been read back, (4) `mwb` original, (5) `mrb` and check that the original has indeed been read back. If any of the two checks fail we assume that the dot has lost its out-of-plane property and return 0, else `erb` returns 1.

Since the medium and the head are stationary, `erb` is at least 5 times slower than `mrb` but there are no positioning delays incurred as would be the case with a spinning disk. `ewb` is also slower than `mwb` because of the local heating process. However, as stated before, the idea is to use the `erb` and `ewb` operations sparingly.

Magnetic read and write a sector Following Pozidis et al. [17] we assume a standard sector size of 512 bytes and about 15% sector overhead for the sector header, ECC, and CRC. All data can be read and written using the efficient magnetic sector write `mws` and read `mrs`, subject to error correction appropriate to the medium, the tips etc.

Freeze a sequence of sectors Our approach is inspired by the Venti archival system [18] where the integrity of a sequence of blocks is protected by protecting the integrity of the hash. While Venti stores the hash separately from the block, we are able to bind the hash physically to the blocks. The operation `freeze`, which, when given a suitable n and a sequence of 2^n sectors aligned on a 2^n boundary, performs the following atomic sequence of steps: (1) read $2^n - 1$ sectors representing the information to be protected using $2^n - 1$ calls to `mrs`, (2) calculate a secure hash (e.g. SHA-256) of the sectors just read, (3) write the 512 byte Manchester encoding of the 256 byte hash in the last (2^n-th) sector using the tamper-evident write operation `ewb`, (4) check that the hash can be read back using `erb`, or else fail.

All properly aligned sequences of 2^n sectors can be frozen individually, thus providing significant flexibility. Frozen data can still be read efficiently, and as often as needed.

The `freeze` operation, when applied to a sequence of 2^n sectors that has already been frozen will fail unless the Manchester encoding of the new and the old hash are exactly the same. This would require the hash of the old and the new information to be the same, which we assume to be unlikely due to the collision resistance of secure hash functions such as SHA-256.

Verify a frozen sequence Firstly, we analyse how the `verify` operation detects all possible attempts to tamper with the frozen information. An attack perpetrated using one or more `mwb`, `ewb`, or `mws` operations on the information is either harmless or it is detected. There are three cases to consider. Depending on the ECC, (a) one class of tampering will be corrected by the ECC thus not interfering with the integrity of the information; (b) a second class of tampering will cause the ECC to fail, signalling the user that her data is corrupted, and (c) the third and final class of tampering yields a valid code word that is different from the original information. This third case is not detected by the ECC but it *is* detected when the `verify` operation checks whether the hash of the information corresponds to the hash written in the last sector.

Secondly, the only type of attack possible on the hash is to ”burn” another dot using the `ewb` operation. This would have the effect of either changing a Manchester encoding 01 into 00 or an encoding 10 into 00, both of which are detected as an invalid code.

Finally, since the hash immediately follows the data (i.e. there are no pointers linking the data to be protected to the hash that can be tampered with as in Venti [18]), there is no way of separating the data from the hash, so that the integrity of the hash truly protects the integrity of the data.

Delete a frozen sequence Our proposal for secure deletion is to physically change the medium where the data was stored, but more subtly than just by shredding the entire device [14]. The `mwb` operation might suffer from memory remanence problems. Therefore, the `ewb` operation must be used to burn each and every bit of the data, leaving the hash intact to witness the fact that the data was once written with the integrity check in place and that it has now been destroyed. Without hardware support, secure deletion is a difficult problem for at least two reasons. Firstly, Bauer and Priyantha argue that due to memory remanence properties of the media, recovering data deleted normally from a disk drive is remarkably easy [1]. Secondly, Mitra and Winslett add that even when the data is deleted, correlation of derived data sets may make it possible to recover the original, deleted data [11].

Our approach is to avoid making copies of the data. Probe storage on patterned media makes it possible to store a working copy of the data, until the data is finalised, at which point it should be kept for as long as the data retention policy prescribes. There is no need to copy the data from a disk to another system such as an optical write once medium which has the correct archival properties; our system provides working storage and archival storage all on one medium.

6 Conclusions

Probe storage on patterned media is a promising technology for developing secure storage. The capacity of such devices will be huge, and the tamper resistance can be excellent. The experiment reported in this paper demonstrates that in principle it is possible to use a patterned magnetic medium in two essentially different ways: for normal read-write purposes and for write-once/read purposes. It is physically impossible to alter the data after the write once operation has been used. Having described an approach to storing data such that (1) any logical or physical tampering can be detected, and (2) the data can be deleted such that even a skilled forensics lab will be unable to recover the data, we believe that we have provided at least a partial answer to the data retention requirements induced by modern compliance legislation. The main advantage of the approach is its combination of integrity and flexibility: normal data processing can be interleaved with tamper-evident freeze, verify and delete operations. Much work remains to be done, including finding an effective way of making backups.

Acknowledgements

We thank Rogelio Murillo, Sebastiaan Konings, and Takahiro Onoue for their help with the measurements.

References

[1] S. Bauer and N. B. Priyantha. Secure data deletion for linux file systems. In *10th USENIX Security Symp.*, pages 153–164, Washington D. C., Aug 2001. USENIX Association.

[2] B. Hong, F. Wang, S. A. Brandt, D. D. E. Long, and T. J. E. Schwarz. Using MEMS-based storage in computer systems–MEMS storage architectures. *Trans. Storage,* 2(1):1–21, Feb 2006.

[3] J. Hurley. The CSO’s security compliance agenda: Benchmark research report. *Computer Security Journal,* 22(1):37–44, Dec 2006.

[4] V. Kher and Y. Kim. Securing distributed storage: challenges, techniques, and systems. In *StorageSS '05: Proceedings of the 2005 ACM workshop on Storage security and survivability*, pages 9–25, Fairfax, VA, USA, 2005. ACM Press.

[5] Y.-S. Kim, S. Jang, C. Sunyong Lee, W.-H. Jin, I.-J. Cho, M.-H. Ha, H.-J. Nam, J.-U. Bu, S.-I. Chang, and E. Yoon. Thermo-piezoelectric Si3N4 cantilever array on CMOS circuit for high density probe-based data storage. *Sensors and Actuators, A: Physical*, 135(1):67–72, Mar 2007.

[6] E. M. Kurtas, M. F. Erden, and X. Yang. Future read channel technologies and challenges for high density data storage applications. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages V737–V740, Philadephia, Pennsylvania, Mar 2005. IEEE.

[7] M. C. S. Lange. Sarbanes-Oxley has major impact on electronic evidence - several provisions of act govern document-retention policies. *The National Law Journal*, Jan 2003.

[8] R. Luttge, H. A. G. M. van Wolferen, and L. Abelmann. Nanolithography for patterned magnetic data storage media. In *Innovative Mass Storage Technologies (IMST)*, Enschede, The Netherlands, Jun 2007. to be published in Proceedings EIPBN.

[9] I. D. Mayergoyza, C. Serpico, C. Krafft, and C. Tse. Magnetic imaging on a spin-stand. *Journal of Applied Physics*, 87(9):6824–6826, May 2000.

[10] D.-K. Min and S. Hong. Design and analysis of the position detection algorithm for a probe storage. *IEEE Sensors Journal*, 6(4):1010–1015, Aug 2006.

[11] S. Mitra and M. Winslett. Secure deletion from inverted indexes on compliance storage. In *StorageSS '06: Proceedings of the second ACM workshop on Storage security and survivability*, pages 67–72, Alexandria, Virginia, USA, Oct 2006. ACM Press.

[12] D. Molnar, T. Kohno, N. Sastry, and D. Wagner. Tamper-evident, history-independent, subliminal-free data structures on PROM storage -or- how to store ballots on a voting machine (extended abstract). In *IEEE Symp. on Security and Privacy (S&P)*, pages 365–370, Berkeley, California, May 2006. IEEE Computer Society, Los Alamitos, California.

[13] S. Naberhuis. Probe-based recording technology. *J. of Magnetism and Magnetic Materials*, 249(3):447–451, Sep 2002.

[14] Office of DASD. *Design Criteria Standard for Electronic Records Management Software Applications*. Department of Defence, Jun 2002.

[15] J. Patzakis. New accounting reform laws push for Technology-Based document retention practices. *Int. Journal of Digital Evidence*, 2(1):paper 2, Spring 2003.

[16] S. Porthun, L. Abelmann, and C. Lodder. Magnetic force microscopy of thin film media for high density magnetic recording. *J. of Magnetism and Magnetic Materials*, 182(1-2):238–273, Feb 1998.

[17] H. Pozidis, P. Bächtold, J. Bonan, G. Cherubini, E. Eleftheriou, M. Despont, U. Drechsler, U. Dürig, B. Gotsmann, W. Häberle, C. Hagleitner, D. Jubin, A. Knoll, M. A. Lantz, A. Pantazi, H. E. Rothuizen, A. Sebastian, R. Stutz, and D. W. Wiesmann. Scanning probes entering data storage: From promise to reality. In *IEEE Conf. on Emerging Technologies - Nanoelectronics*, pages 39–44, Singapore, Jan 2006. IEEE.

[18] S. Quinlan and S. Dorward. Venti: A new approach to archival data storage. In *1st USENIX Conf. on File and Storage Technologies (FAST)*, pages 89–101, Monterey, California, Jan 2002. USENIX Association.

[19] E. Sarajlic, E. Berenschot, N. R. Tas, H. Fujita, G. Krijnen, and M. C. Elwenspoek. Fabrication and characterization of an electrostatic contraction beams micromotor. *IEEE Int. Conf. on Micro Electro Mechanical Systems (MEMS)*, pages 814–817, Jan 2006.

[20] N. R. Tas, J. Wissink, A. F. M. Sander, T. S. J. Lammerink, and M. C. Elwenspoek. Modeling, design and testing of the electrostatic shuffle motor. *Sensors and actuators A (Physical)*, 70:171–178, 1998.

[21] M. Taylor. The EU data retention directive. *Computer Law & Security Report*, 22(4):309–312, 2006.

[22] B. D. Terris and T. Thomson. Nanofabricated and self-assembled magnetic structures as data storage media. *J. of Physics D: Applied Physics*, 38(12):R199–R222, Jun 2005.

[23] B. D. Terris, T. Thomson, and G. Hu. Patterend media for future magnetic data storage. *Microsystem technologies*, 13(2):189–196, Jan 2007.

[24] R. Murillo Vallejo, M. H. Siekman, T. Bolhuis, L. Abelmann, and J. C. Lodder. Thermal stability and switching field distribution of CoNi/Pt patterned media. *Microsystem technologies*, 13(2):177–180, Jan 2007.

[25] H. Zeng, M. L. Yang, N. Powers, and D. J. Sellmyer. Orientation-controlled nonepitaxial $l1_o$ CoPt and FePt films. *Applied Physics Letters*, 80:2350–2352, 2002.