Toward Permanence in Digital Data Storage

Barry M. Lunt, Matthew R. Linford, Robert C. Davis, Sarah Jamieson, Anthony Pearson, Hao Wang; Brigham Young University; Provo; Utah/United States of America (APA)

Introduction

Since the invention of the computer many decades ago, computer storage has always been ephemeral, with the notable exception of magnetic core storage. Hard disc drives, invented in the 1960s and improved in many ways over the decades, remain ephemeral, with data lifetimes estimated at 1-7 years. Flash memory, the dominant non-volatile solid-state option available today, has a data lifetime of 1-12 years, depending on the density of the memory. Recordable optical discs, the only other recordable data storage format available today, has a data lifetime of 1-25 years, depending on many factors. Most people in the computing culture have grown up with an understanding of these limitations and have learned to accept them, since it has always been this way, and since losing computer data has always been part of the computing experience. Our position is that it should not be this way, and does not need to stay this way.

The Need for Permanence

Since mankind began to keep records, it has always been their intent to make records that would endure. In many cases, great pains were taken to produce records that would persist for centuries, using materials such as pottery, gold, parchment, papyrus, and inks that would last as long as possible. The fact that many of these documents are readable today is evidence of the permanence of these records and the effectiveness and durability of these early methods and materials. If these earlier records had been stored using today's digital storage technology, only their materials would remain, such as papyrus, parchment, gold, or pottery – all markings would have disappeared, and we would know nothing of their language or their writing, creating a dark age regarding their history. Historians would have very little to work with to piece together their history.

From this we believe it can be argued that persistence of markings is the *sine qua non* of records. If the markings fade with time, no one can ever learn to decipher them. If the Rosetta Stone had been blank, it would have been completely worthless. While persistence of markings is clearly a necessary condition for permanence of records, it is not necessarily a sufficient condition. For example, Linear A, a writing method used by the residents of the island of Crete somewhere around 2000 BCE, is still unreadable today because we have lost the ability to decipher it. The same was true of Egyptian hieroglyphs until they were deciphered with the help of the Rosetta Stone in the 1820s. But whatever storage method we adopt for archiving computer data, it must also provide persistence of markings.

The second condition necessary for long-term archival of records is persistence of a method for reading (deciphering) the marks. History gives several examples, the most widespread of which is the use of Latin today. Although no country or ethnic group speaks Latin today, Latin is still taught, and many people still learn it – because there are thousands of records that exist today in Latin. But if the markings on all those Latin records were to suddenly fade to nothing, what would be the reason to study or teach Latin? And the reason there are so many Latin records available today is because it was such a widely adopted language – most of the Roman Empire employed or learned Latin and it continued as the lingua franca of Europe for centuries. So, if we adopt a storage format that is widely adopted and easily deciphered today, chances are very good that many centuries in the future, those same records, if the marks persist, will be decipherable.

Thus, we argue that persistence of marks is indeed the *sine qua non* of archival storage. Solving that problem must come first, or there is no data to decipher in the future. For deep archival storage (meaning we store and forget), we cannot depend on active management as we do today, for it is inevitable that some records will not be properly managed and will therefore be lost. And at the level of individual records, active management simply does not happen. The data should not require special conditions for its storage. Thus our efforts have focused on storage formats that are widely adopted, materials that are extremely durable, and recording methods that produce permanent marks. This approach produces digital data storage technologies that will persist without requiring active management and that need no special storage conditions. Such records may be stored like a journal – write it, put it on the shelf, and centuries later it is available to be pulled off that shelf and read.

One other point should be made. There are businesses today whose primary function is to provide computer users a way to transfer their computer files from an older storage technology to a newer technology. Such businesses depend on some level of persistence of marks – a few years or a few decades. Thus, if data persists, it is logical that there will be companies far into the future that will provide ways to transfer data from older formats to newer formats, much as we can transfer songs from LPs to CDs today. But this can only happen if the data persists; if it does not persist, we face a future "digital dark age"[1,2].

Efforts to Date

The first fruit of our research was a DVD-format recordable optical disc, now available as the M-Disc. As has been reported elsewhere[3], the data on this new type of DVD is extremely persistent, far better than even the best archival-quality DVDs available today. A Blu-ray version of this optical disc is also now available. Figure 1 shows how this technology differs from other recordable optical discs available today; the recording layer of the M-Disc is made from materials that are much different from materials in other recordable optical discs. The recording marks made on this disc, as seen in Figure 2, are physical, and not just based on optical contrast; they are the result of permanently opening up holes in the recording layer with the recording laser.

The M-Disc has established that it is possible to create a digital data storage medium that will allow permanent marks to be made, and that will thus allow data to persist for centuries. However, we recognize the importance of providing more than a single permanent storage medium. Accordingly, our most recent research has been focused on two other formats, which are discussed in the following sections.

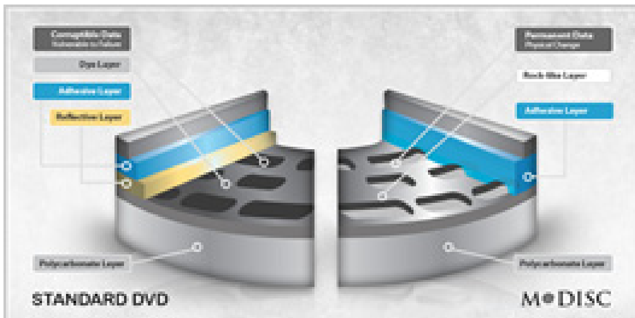

Figure 1: *The recording layer of the M-Disc is dramatically different from that of other recordable optical discs. (from www.millenniata.com)*

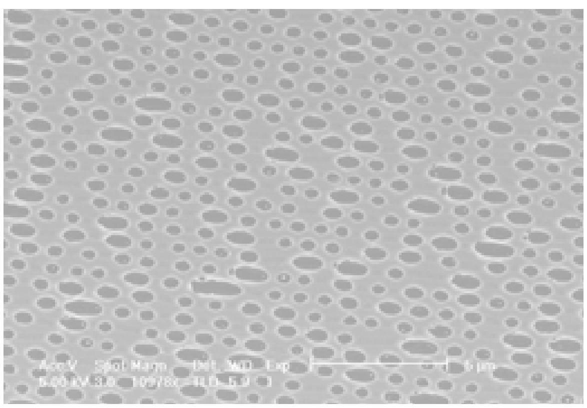

Figure 2: *Marks on the M-Disc, as imaged by an SEM (scanning electron microscope).*

Solid-State Permanent Storage

The early days of solid-state storage (the 1960s) saw the development of two permanent storage media: read-only memory (ROM) and programmable ROM (PROM). The main problem with ROM was that it was programmable only in the factory, with the bits actually being manufactured on each chip. While fully permanent, it was not programmable by the user, and thus filled only a small niche in storage needs.

PROM was promising for permanent storage, and was quite popular for a number of years. It did, however, suffer from one rather serious lifetime problem, which was the formation of dendrites (see Figure 3). The growth of these tree-like structures is driven by the presence of moisture, applied voltage, and available material. As can be seen in Figure 3, dendrites can grow until they short out a blown fuse, which then compromises the data stored there.

Today, technologies that allow the user to program them only once are known as WORM (write once, read mostly), and have again grown in popularity. However, the only such technology that is truly permanent is the M-Disc. Solid-state options for WORM are dominated by flash memory that has a controller that does not allow the data to be erased. But since data stored in flash memory slowly leaks away, this option does not solve the permanence issue.

Figure 3: *A dendrite has grown from the remains of a programmed cell (blown fuse) in a PROM storage device.3*

Our initial research focused on two classes of materials to develop fuses (the basic storage elements of solid-state WORM storage). The requirements for such materials include the following:

1. Electrically resistive but not insulative.
2. Programmable with voltages from 3-25V.
3. Extremely stable in the unprogrammed state, at temperatures up to 150°C.
4. Extremely stable in the programmed state, at temperatures up to 150°C.
5. Compatible with existing integrated-circuit manufacturing processes and equipment.
6. Relatively inexpensive.
7. Entirely resistant to dendrite growth.

The first class of materials studied was that used in developing the M-Disc, as has been reported elsewhere[4]. These materials have been shown to satisfy requirements #1, 3, 4, and 6 above. We investigated their acceptability for requirements #2, 5, and 7.

"Bow-tie" test structures were fabricated out of these soft metals (see Figure 4), and current was passed through them, creating a rapid rise in the temperature of the neck region due to I^2R heating, followed by melting (see Figure 5).

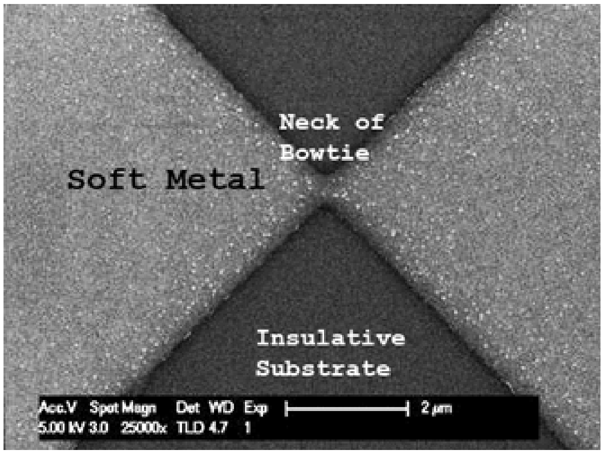

Figure 4: *A horizontal planar fuse of soft metal on an insulative substrate. Moderate resistance and high current density in the neck region causes rapid localized heating and melting.*

We found that these materials met requirements #2 (it required about 15 Volts to blow the fuses) and #5, but after programming the fuses and subjecting them to prolonged exposure to voltage, humidity, and air, we were disappointed to see that dendrites had grown in them (see detail in Figure 5). Thus, we ended further research into these materials for permanent solid-state storage.

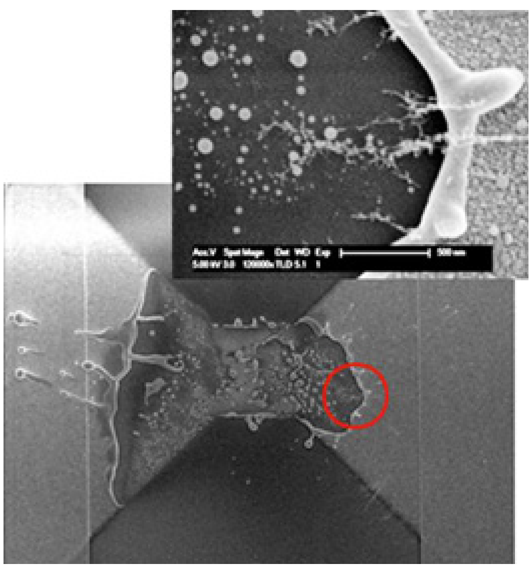

Figure 5: *The above fuse has been blown, but with time, dendrites have started to form.*

The second class of materials studied for this application were hard metals, such as rhodium (Rh), iridium (Ir), tungsten (W), molybdenum (Mo), and tantalum (Ta). These metals have very high melting points, and are resistive enough that sufficient I^2R heating was deemed possible. Accordingly, test structures as in Figure 4 were made from these metals and the fuses were blown. After further time for reliability testing, it was found that dendrites also grew in these materials (see Figure 6). The voltage required to blow the fuses was about 20 Volts. Thus, these materials were shown to satisfy all the requirements except #7.

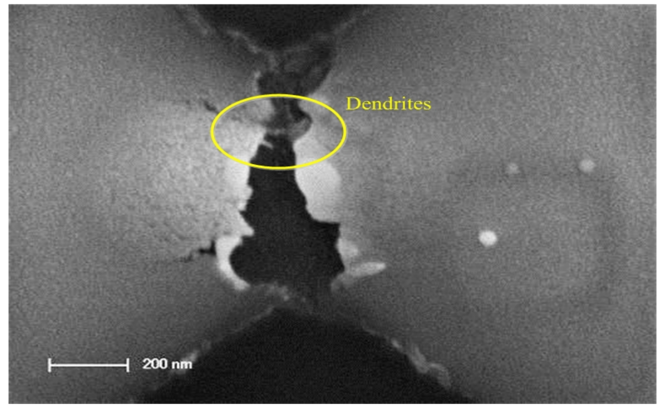

Figure 6: *A close-up SEM of the neck region of a blown hard-metal fuse. Dendrites can be seen growing across the blown portion of the fuse.*

After the failures of these first two material classes, a third class of materials was researched. These are the carbonaceous materials, including several allotropes of carbon such as graphite, graphene, and glassy carbon. We knew that these materials would probably not melt with I^2R heating, but they could be oxidized. We also knew that these materials would be extremely stable in either their programmed or unprogrammed state, and that present integrated circuit (IC) manufacturing equipment could be used to fabricate them. So, as before, test structures were made of these materials, and efforts were made to program (blow) the resulting fuses, and prolonged reliability tests were performed with an applied voltage bias. The results are shown in Figure 7. The dark part of Figure 7 is the carbonaceous fuse; the small vertical gap in the middle (neck) of the bowtie region is where the fuse was blown. The good news was that even after prolonged reliability testing, no evidence of dendrites was seen. The voltage required to blow these fuses was approximately 20 Volts. Thus, it appears that this class of materials satisfies all seven of the requirements for a fuse material.

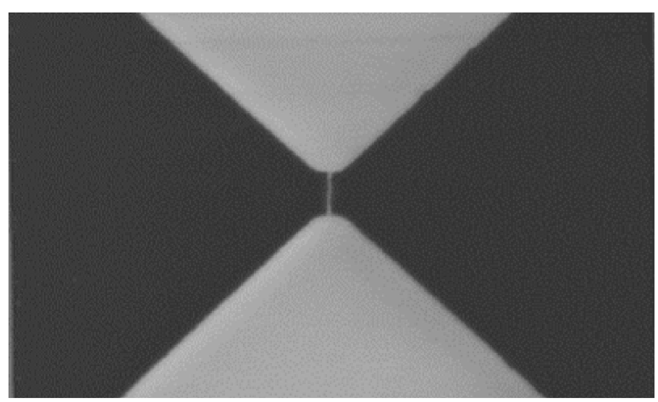

Figure 7: *A blown fuse made of carbonaceous material, after prolonged exposure to reliability testing.*

The cell size for a planar fuse would be approximately the same as the cell size for flash memory. By extension, this means that the potential density for this structure is equivalent to flash, which also means that the costs would be approximately the same. However, we have also been researching vertical fuse structures in a crossbar array, as shown in Figure 8. Simulations indicate that such a structure can be easily manufactured, and that the potential density is four to twelve times greater than flash memory. We also know that such a structure could be made in multiple layers, allowing the density of this type of permanent solid-state storage to greatly exceed that of flash memory. With such increases in density, the cost of this memory would also be much lower than that of flash memory.

Further research and development in this area is presently focusing on producing structures which are fully compatible with IC manufacturing processes, including the packaging of the ICs. We expect to complete this next phase of research within the calendar year, and that the design of an IC fab prototype would begin shortly thereafter.

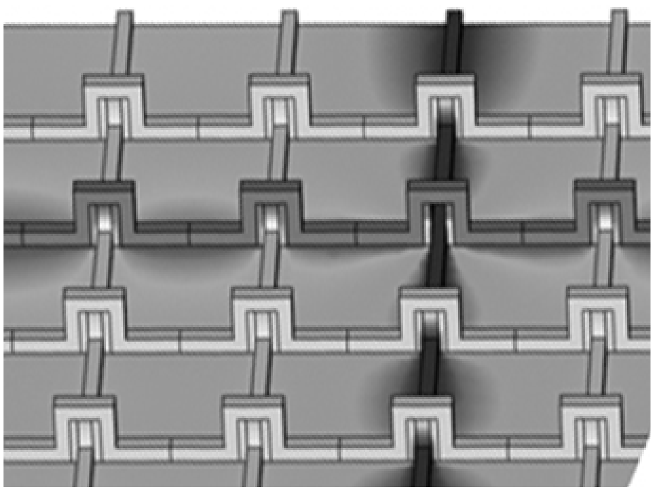

Figure 8: *A simulation model for vertical fuses in a crossbar array.*

Optical Tape Permanent Storage

Magnetic ½-inch tape has been widely used for decades to store digital data, and its main application at present is backing up data from hard drives. The lifetime of data on magnetic tape is estimated to range from 10 to 50 years, depending on several factors[5]. Clearly, there is a need for something more permanent in tape storage.

As was mentioned previously, the main difference between M-Discs and other "archival-quality" optical discs is the recording layer (see Figure 1). The materials used for this unique and permanent recording layer were investigated for use on tape. The requirements for materials to be used as a permanent tape storage solution would include the following:

1. Optically absorptive at the desired wavelength.
2. Writable with less than 50 mW of optical power.
3. Extremely stable in its unrecorded state.
4. Extremely stable in its recorded state.
5. Compatible with existing web-printing processes.
6. Relatively inexpensive.
7. Flexible and resistant to tension.
8. Adheres well to the Mylar substrate.

We already knew that the recording layer used in the M-Disc satisfies the optical response of requirements #1 and #2. We also knew it met requirements #3, 4, 5 and 6. To test for requirements #7 and #8, we prepared samples via sputtering using a PVD-75 (from the Kurt J. Lesker Co., Clairton, PA) onto blank Mylar ½" tape samples. Two kinds of adhesion tests were performed on these samples: a Scotch tape adhesion test, and a weight stand abrasion test.

The Scotch tape adhesion test was performed by adhering a sample of Scotch tape of about 4 cm length onto the recording layer side of the ½" tape. The Scotch tape was then removed. This cycle was repeated 50 times, then both the Scotch tape and the ½" tape were visually studied for any evidence of material removal. If no material appeared to have been removed, a fresh piece of Scotch tape was used and the procedure was repeated another 50 times. This was repeated two more times, so that in the end, four pieces of Scotch tape about 4 cm in length were adhered to and then removed from the same sample of ½" tape, each Scotch tape piece being adhered and removed 50 times, for a total of 200 iterations. There was still no evidence that any of the recording layer material had been removed from the sample of ½" tape.

The weight stand abrasion test is depicted in Figure 9. The ½" tape sample was attached to a weight of 5, 10, 15 and 20 Newtons, which covers the range of tension to which ½" tape is subjected during normal use. The 90° bend exceeds the normal bend of about 60° in most tape drives. Also, the 1.2 cm diameter roller shown in Figure 9 over which the tape is draped at 90° is a stationary roller, which is much worse than in actual tape drives, where all rollers are required to roll. The tape sample was moved back and forth over this roller, with the weight attached, a total of 100 times for each weight sample. Again, the sample of ½" tape was examined visually for any evidence of material removal. Since no material removal was observed, we felt that requirements #7 and #8 were fully satisfied on the ½" tape samples.

The next test was to observe the power required to make marks in this permanent recording layer. This was done and marks of micron dimensions were produced at reasonable powers. A detailed account of this research will be forthcoming. Assuming the marks can eventually be made as small as the marks on Blu-ray discs, the number of tracks could be equal to that on today's ½" tape (2176 tracks). With the minimum pit length of 0.15 µm (as on today's Blu-ray discs), and assuming 846 meters of ½" tape in a cartridge (typical for LTO 6 tape cartridges), this would allow 1.534 TB on a single cartridge, which is close to the native capacity of LTO 6 cartridges (2.5 TB).

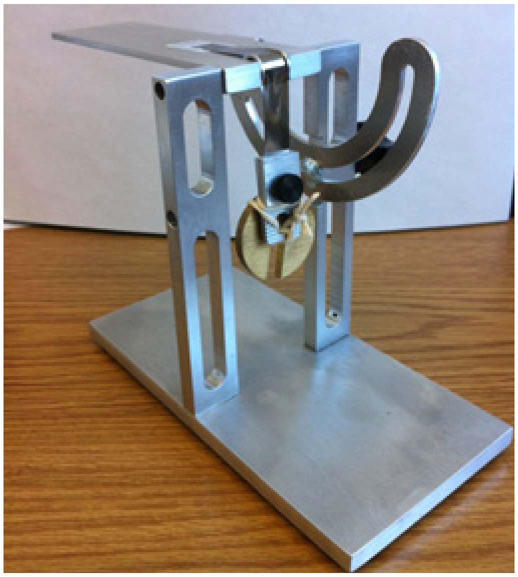

Figure 9: *The weight stand test fixture, over which the tape is draped and tensioned, then moved back and forth.*

Summary

The above results have demonstrated viability of these two additional technologies to meet the need for permanent deep archival of digital data. The solid state option has the potential to equal flash memory in density and performance, and for the data to persist at least 1,000 years. The optical tape option has the potential to equal magnetic tape in density and performance, and for the data to persist much longer than today's magnetic media. These two permanent storage technologies could be available in as little as three years.

References

[1] Smit, Eefke, Jeffrey Van der Hoeven, David Giaretta, "Avoiding a Digital Dark Age for data: why publishers should care about digital preservation", *Learned Publishing*, 24: 35-49, Jan 2011.

[2] Kuny, Terry, "A Digital Dark Ages? Challenges in the Preservation of Electronic Information", *63rd IFLA (*International Federation of Library Associations and Institutions*)*, http://www.ifla.org.sg/IV/ifla63/63kuny1.pdf, September 1997.

[3] Pecht, Michael, "Handbook of Electronic Package Design", Marcel Dekker, Inc., New York, NY, 1991, p. 597.

[4] Abbott, Jonathan, Travis L. Niederhauser, Barry M. Lunt, Douglas P. Hansen, Ray T. Perkins, David A. Bell, Erik C. Bard, Mark O. Worthington, C. Michael Miller, Daniel F. Hyatt, Matthew C. Asplund, Gui Lin Jiang, Matthew R. Linford, Richard R. Vanfleet, Robert C. Davis, "Carbon Coated Tellurium for Optical Data Storage", *ACS Applied Materials and Interfaces*, vol 2, #8, (Jan 2010), pp 2373-2376.

[5] Navale, Vivek, "Predicting the Life Expectancy of Modern Tape and Optical Media", *RLG DigiNews*, Aug 15, 2005, 9:4; also at www.rlg.org/en/page.php?Page_ID=20744#article3

Primary Author

Barry M. Lunt is a full professor of Information Technology at BYU in Provo, UT. He has taught college for 24 years and has worked full-time with IBM, Bell Labs and Micron. Recent research is in permanent storage for digital data. He has published over 40 papers and two books.

Contact Information for Primary Author

Barry M. Lunt
265 CTB
Brigham Young University
Provo, UT 84602
801-422-2264 (office)
801-422-0490 (fax)
luntb@byu.edu

Co-Authors

Matthew R. Linford
Robert C. Davis
Sarah Jamieson
Anthony Pearson
Hao Wang