

Research on Another Permanent Data Storage Solution

Barry M. Lunt, Matthew R. Linford, Robert Davis, Brigham Young University, Provo, Utah

Introduction

There are three main classes of data storage in use today: magnetic (including hard-disk drives – HDDs – and tape); optical discs (including CDs, DVDs, and Blu-ray); and solid state (including SRAM, DRAM and Flash). With one notable exception among DVDs, none of these options provides a way to permanently store data, with data lifetimes ranging from milliseconds (DRAM) to about 30 years (magnetic tape). This ephemerality is particularly concerning to those dealing with archival issues.

Additionally, there are several characteristics which are particularly important to archivists when it comes to preserving digital data. These include: 1) no active maintenance or migration is required to preserve the actual data; 2) no special storage conditions are required to preserve the media or the data; 3) a minimum lifetime of at least 100 years, preferably 500 years (this means the average lifetime is MUCH greater); 4) no power is required to maintain the data; 5) the media is easily transported; 6) the format is widely adopted; and 7) the medium has a large storage capacity. The right archival solution should include all these characteristics.

The options available today for storing digital data include only four: 1) magnetic hard-disk drives; 2) magnetic tape; 3) optical discs (including CDs, DVDs, and BDs); and 4) Flash memory. These have relatively short lifetime expectancies (LEs) for the data, as shown in Table 1. Clearly, for archival purposes, there is only one solution on the market today, and it is relatively new.

Table 1: Life Expectancy of data for today's digital data storage options.

Media	Life Expectancy of Data
Magnetic tape	10-50 years
Magnetic hard-disk drives	1-7 years
Flash drives and Solid-state drives	10-12 years
Recordable optical discs	1-25 years
Millenniata recordable optical disc	1,000 years (advertised)

Previous Research

In a paper presented in 2009, Lunt & Linford¹ reported on a revolutionary new type of DVD that was designed to last 1,000

years, and thus be the first true archival-quality data storage solution. This DVD was rigorously tested against the very best archival-quality and standard-quality DVDs in a test performed by the Naval Air Warfare Center Weapons Division (China Lake, California, USA) in 2009². Figure 1 shows the results of this testing. The discs from each brand of disc included in the study were written to, and the PI8(max) errors were counted. The all discs were subjected to 24 hours of elevated temperature (80°C), humidity (85% RH), and full-spectrum light (1200 W/m²). Based on their summary, it appears that this media is greatly superior to all those tested (pages i and 45, Svrcek).




Figure 1: The results of the testing by the Naval Air Warfare Center Weapons Division, showing that all discs except the Millenniata discs failed the accelerated testing rather dramatically.

Many historians have labeled our current period of time as the “Digital Dark Age(s)”^{3,4}, pointing to the fact that digital records do not persist. At the level of organizations (businesses, governments, agencies), people actively manage the problem. However, at the individual level, people store their precious personal digital photos, stories, and other personal digital artifacts, on various forms of digital media, none of which will persist unless actively managed – which is very unlikely to happen at the individual level. Additionally, the present method of active management (which usually means storing data on large hard-disk drive installations) does not meet any of the other desirable characteristics for archival storage mentioned in the previous section.

The study by the Naval Air Warfare Center Weapons Division involved accelerated aging. But what of natural aging? In their 2011 paper, Lunt, Hansen & Linford⁵ reported on a study done on top-quality, archival CD-Rs and DVD-Rs used in two academic libraries. These discs were not circulated – they never left the temperature-humidity-light-controlled conditions of their storage, and were periodically evaluated for data permanence. It is a classic case of a natural aging test. This study included about 26,500 discs (both CD and DVD), all of which were archival-grade discs. The alarming results were that both collections of discs were experiencing an annual loss of 2% of their discs, meaning that every year, about 2% of their discs permanently lost at least one file or became completely unreadable. This is an alarming data

loss, yet it occurred under optimal conditions, with the best possible quality of discs. It also serves to underline the problem of digital data archiving.

Progress to Date

The development of the Millenniata DVD (M-DISC) has been reported elsewhere^{6,7,8}, and it appears that it is a viable solution. However, it is a single solution at the present, and the urgency of this problem prompts that more options be available. Table 1 shows that the M-DISC meets the majority of the requirements for true archival storage, with its major weakness being its relatively low storage capacity of 4.7 Gbytes (standard DVD capacity).

Table 2: Desirable storage attributes for archiving, and how the M-DISC performs in each attribute.

Attribute	Score for M-Disc (10=perfect)
No active maintenance required	10
No special storage conditions required	9
Last ≥100 years	10
No power required	10
Easily Transported	10
Widely-adopted format	10
Large capacity	4

It is often pointed out that format obsolescence is another very important consideration, and this is true. But without data which persists, format obsolescence is a moot point – persistence of data is the *sine qua non* of archiving. Another major point in response to this concern is to look at history. In all instances when a given format (or language) has been widely adopted, future generations have learned to read that format or language. For instance, Latin is not the official language of any people or country, yet it is still taught in universities around the world. The reason? Because there are many documents originally written in Latin, and scholars wish to be able to read them. Optical discs are the most widely-adopted digital data storage format in the history of the world – there are hundreds of billions of them, and tens of billions of drives. And finally, if it is easy to find a way to read these discs today, the future progress of technology will only make it easier in the future.

Solid-State Storage

One area that has consistently been able to show a great deal of increase in storage capacity is that of solid-state storage, including DRAM, SRAM, and Flash memory. None of these are workable solutions for archival storage, since the data does not last for more than about 12 years (for Flash memory), or requires power to maintain the data (for DRAM and SRAM). However, if a type of solid-state memory were developed that was inherently permanent, the density could be increased to allow the storage

capacity to meet archival needs, much as we see solid-state drives (SSDs) being used today.

Accordingly, we have begun research into a WORM (write-once, read mostly) type of solid-state storage that overcomes all presently-known failure mechanisms. These failure mechanisms include metal migration, latent defects, and dendritic growths.

The first two failure modes are common to nearly all types of integrated circuits (ICs), and have thus been addressed by the IC industry for several decades. Though the physics of these failure modes are still at work, IC design has long established design rules that, when followed, dramatically minimize these failures. The expected lifetime of most ICs today is over 99 years, though the exact number is highly subject to debate. However, there are many examples of solid-state devices which have operated for decades, so it appears that these failures are not a significant factor today.

The last failure mode – dendritic growths – has been a factor from the beginning for WORM solid-state storage. Figure 2 is an example of such a growth in a WORM storage cell. Because the dendritic growth is composed of conductive material, it has shorted out the previously-blown fuse, thus compromising the programmed data. Dendritic growths are still not fully understood, but it is known that if there are no remaining materials when the fuse is blown (when the data is programmed), there will not be dendritic growth. Thus, if the fuses are made from highly reliable materials, and if the materials are completely removed in the programming process, both the manufactured 1's (intact fuses) and the programmed 0s (blown fuses) will be permanent (cannot be changed) and will have a very long life expectancy (LE).




Figure 2: A dendrite has grown from the remains of a programmed cell (blown fuse) in a WORM storage device.⁹

Accordingly, research has progressed on three different classes of materials from which these fuses could be made, all of which have shown significant promise. Figure 3 is an example of a programmed fuse made from class one materials; note the absence of any remaining fuse materials from which dendrites could grow, except for that material which has folded back onto the main fuse material. Also note the very significant gap created in the programming process, making it very difficult for any type of growth or defect to compromise the programmed data.




Figure 3: A programmed fuse, made from class one materials.

Figure 4 is an example of another class of materials from which a fuse has been constructed, and which fuse has then been electrically blown, as was the fuse in Figure 3. While the gap in this fuse is not as large as that in Figure 3, it is still quite large (approximately 250 nm), which is easily large enough to prevent dendritic growth. Additionally, the material from which the fuse is made is known to be extremely resistant to dendritic growths. And finally, note the complete absence of any remains of the blown fuse. Such a structure, made from extremely long-lasting materials, would have a very long LE.




Figure 4: A programmed fuse, made from class two materials.

Being careful to not put all our eggs into one or two baskets, we have also begun research on a third class of materials which are extremely resistant to dendritic growth, extremely long-lasting, and which can be programmed without leaving any fuse remains. This research is still preliminary, and no results have yet been demonstrated.

Estimated Density

Present cell sizes for DRAM and Flash memory are approximately $5F^2$, where F = the state-of-the art design rules

(approximately 13 nm). The cell size for this technology would be even smaller – probably about $2F^2$, providing a density improvement of $(5/2)^2$, or about 6.25 x. Additionally, these devices could readily be made in layers approximately 10 μ m thick, providing 2 to 10 layers of devices in the same real estate, providing an overall density improvement of up to 62x. Future work could develop the ability to make even more layers. If this technology were available today (10 layers at a cell size of $2F^2$), this means permanent storage in the form of solid-state drives would be capable of storing more than 16 Tbytes, at a price roughly equivalent to today's storage for 256 Gbytes. This is a dramatic improvement in density, and the data would be stored permanently.

Future Work

Future work will be focused on making one or more of these materials classes into a viable commercial product, capable of being manufactured in a standard IC manufacturing facility. It is estimated that this work would be completed in 3-5 years. In its final format as permanent USB drives or solid-state drives, we estimate that it will be cost competitive when initial manufacturing challenges are solved and densities reach the limits inherent to this new technology.

Summary

This paper has provided details on the LE of the data on magnetic tape, hard-disc drives, flash memory, and optical discs. For archival purposes, only the M-DISC presently meets the requirement of long-term data retention (more than 100 years). This disc has been rigorously tested, and the results are very positive.

The main weakness of the M-DISC is its relatively small capacity. This weakness is being addressed by our current research on a permanent solid-state solution, one whose density will enable a capacity of several terabytes in the format of a solid-state drive. Though it will be much more expensive than an M-DISC, the convenience of having several terabytes in a single, removable form of storage should be very attractive to the archival community.

References

- [1] Lunt, Barry M., Matthew R. Linford, "Towards a True Archival-Quality Optical Disc", *Proceedings of 2009 International Symposium on Optical Memory-Optical Data Storage Conference*, Nagasaki, Japan, Oct 2009.
- [2] Svrcek, Ivan, "Accelerated Life Cycle Comparison of Millenniata Archival DVD", *Life Cycle and Environmental Engineering Branch, Naval Air Warfare Center Weapons Division*, China Lake, CA, USA, Nov 10, 2009.
- [3] 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.
- [4] 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.

- [5] Lunt, Barry M., Douglas Hansen, Matthew R. Linford, "Optical Disc Life Expectancy: A Field Report", *Proceedings of ISOM-ODS 2011*, Lihue, Kauai, Hawaii, USA, July 17-21, 2011.
- [6] Lunt, Barry M., "How Long is Long-Term Data Storage?", *Proceedings of Archiving 2011*, Salt Lake City, UT, USA, May 16-19, 2011.
- [7] Lunt, Barry M., Matthew R. Linford, "Family History Archives: Research on New Media", *Proceedings of FHT 2010*, Salt Lake City, UT, USA, Feb 2010.
- [8] Lunt, Barry M., Douglas P. Hansen, Matthew R. Linford, "Optical Disc Life Expectancy: A Field Report", *Proceedings of ISOM/ODS 2011*, July 17-21, 2011, Lihue, Kauai, Hawaii, USA.
- [9] Pecht, Michael, "Handbook of Electronic Package Design", Marcel Dekker, Inc., New York, NY, 1991, p. 597.