Permanent Digital Data Storage: A Materials Approach

Barry M Lunt
Brigham Young University
265 Crabtree Building, BYU
Provo, UT 84602
+1 (801) 422-2264
luntb@byu.edu

Robert Davis
Brigham Young University
N-215 ESC, BYU
Provo, UT 84602
+1 801-422-3238
davis@byu.edu

Douglas Hansen
Millenniata, Inc.
915 S. 500 E. Suite 112
American Fork, UT 84003
+1 801-358-0935
doug@millenniata.com

Matthew R. Linford
Brigham Young University
C-306 BNSN, BYU
Provo, UT 84602
+1 801-422-1699
mrlinford@chem.byu.edu

Hao Wang
Brigham Young University
N-215 ESC, BYU
Provo, UT 84602
+1 801-422-1699
shuaiwanghao1@gmail.com

John Dredge
Millenniata, Inc.
915 S. 500 E. Suite 112
American Fork, UT 84003
+1 801-358-0935
john.dredge@millenniata.com

ABSTRACT

Permanent marks, interpreted as bits, are the *sine qua non* of deep archival storage. Until about 2006, this area of research apparently did not exist, but advances in the past several years at Brigham Young University and at Millenniata, Inc., have produced one product (the M-Disc - a permanent recordable optical disc of DVD capacity), a follow-on recordable optical disc of Blu-ray capacity, work on multi-layer optical discs of Blu-ray capacity on each layer, work on a permanent solid-state storage medium, and work on a permanent optical tape storage medium. This paper explains the approach used to develop these products and advance the research for permanent digital data storage.

INTRODUCTION

For deep archival storage, the ideal is to be able to create the desired artifact, then store it and forget about it, knowing that whenever we wish to access it again, it will still be there, and we will still be able to read or observe it. While this has been the case with non-digital artifacts throughout history, as exemplified by the many historical artifacts remaining from ancient civilizations, this has never been the case with digital artifacts. Starting with ½-inch tape and hard drives, and moving on through SRAM, DRAM, ROM, floppy discs, optical discs, and flash memory, the situation has always been that digital data has been quite ephemeral, especially as compared to historical artifacts. Even the printed book has a lifetime much greater than the best digital storage option, even though books are a technology many centuries old.

In this day of going fully digital, we could greatly benefit from digital storage media which are permanent. While such a development does not solve all the problems of accessibility for centuries to come, it is the *sine qua non* of deep archival storage.

HISTORICAL EXAMPLES

There are many historical examples of artifacts which have survived for centuries and even for millennia. Such examples include coins (metal and ceramic), written documents, vessels, buildings, weapons, clothing, and many types of works of art. Some of these have survived primarily due to the optimal storage conditions in which they were left; a classic example of this would be the Dead Sea Scrolls, found between 1946 and 1956 in caves near the Dead Sea. Lying in the dry climate, with low humidity and little light, and remaining undisturbed for centuries, these documents were remarkably well preserved, especially considering that most of them were made from parchment and papyrus.

Undoubtedly there have been many other documents produced over the centuries which were not stored in ideal conditions, and which deteriorated quickly, leaving us no knowledge of their ever having been created. Others documents were stored in less than ideal conditions, and suffered severe deterioration but were not utterly destroyed.

A materials approach to studying these surviving artifacts reveals much, as would be expected. Artifacts made of gold essentially do not deteriorate; artifacts made of brass, bronze, and silver suffer some deterioration, but have often lasted for millennia. Artifacts made of pottery, ceramic, and other vitreous materials, while brittle, have still not deteriorated much, and remain today as widespread tokens of bygone civilizations.

Using these materials as our examples, it seems obvious that for deep archival storage, artifacts should be made from materials which either do not oxidize (such as gold), or else are by nature fully oxidized or chemically-reacted materials (such as vitreous

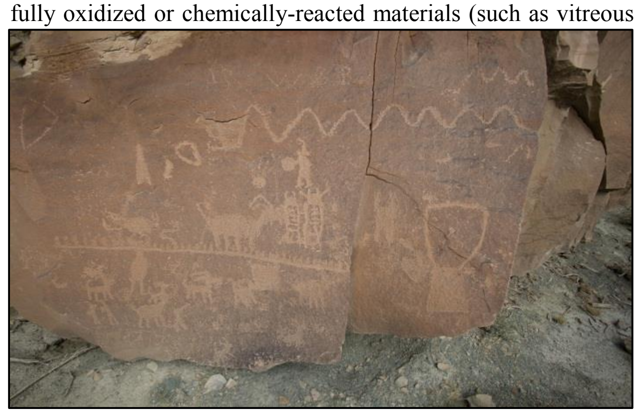

Figure 1: A petroglyph left by the Fremont Indians in Utah, in a place known as 9-Mile Canyon.

Figure 2: A close-up of the petroglyph of Figure 1, showing the etchings in the rock.

materials). Likewise, if we want to preserve writings, the surface on which we write must be have these same characteristics (does not oxidize, or is already fully oxidized), and the ink we use must be treated in such a way that it becomes permanent (such as painting clay then firing it in an oven).

A classic example of an artifact which endures for millennia is petroglyphs (see Figure 1). These writings were made by chipping away a thin layer of dark rock, exposing a lighter layer underneath (see Figure 2). Because these drawings involve physical changes in very durable materials, petroglyphs have endured for centuries in the worst of storage conditions – in the open weather.

NEW TECHNOLOGY, OLD MATERIALS

If we take a materials approach to the problem of permanent digital data storage, it can be seen that if we use extremely durable materials, and if we make significant changes to these materials, these changes will persist, giving us a permanent storage medium.

Accordingly, when this research began in 2006, the first group of materials we studied were metals that either do not oxidize (such as gold), or that form a thin self-limiting layer of oxide (such as aluminum). We knew we could make permanent marks in these metals using relatively cheap solid-state lasers, such as those in CD, DVD, and Blu-ray drives, and we experienced a great deal of success with these materials.

APPLICATION TO CURRENT STORAGE METHODS

Current digital data storage options include magnetic hard-disk drives, magnetic ½-inch tape, flash memory (including USB sticks or "thumb drives", and solid-state drives or SSDs), and optical discs (including CDs, DVDs, and Blu-ray discs). While these storage options have greatly increased in density while dropping in price, they have not addressed the issue of data permanence.

A new technology always has a better chance of success in meeting customer needs if it is not too different from existing technologies. This is particularly true in consumer products for digital data storage. Ac cordingly, we chose to use the existing standard of DVDs, changing only the materials of the recording layer so as to make the data perm anent. This allowed our new DVD-compatible optical disc (the M-Disc) to be widely accessible, as it was readable in all DVD drives. Additional research is currently underway in permanent solid-state storage and in permanent optical tape storage, as described below.

STATE OF THE ART IN PERMANENT OPTICAL STORAGE MEDIA

As anyone experienced in optical data storage will know, merely making a mark in a layer of recording material is not sufficient to produce a robust data storage solution. In our optical disc research, a recording stack was developed consisting of a small number of layers of inorganic, stable materials. Using standard optical disc recording and playback lasers, we were able to reduce the jitter to below the specification for DVDs, and the marks produced (see Figure 3) provided excellent optical contrast and extreme durability.

Optical data recording has a significant advantage when it comes to data longevity, and that is that the recording and playback process does not involve any contact between the media and the recording and playback mechanism. This separation between the recording and playback device and the media allows an infinite number of playback iterations, which is unique when compared to tape. And the relative simplicity of the playback mechanism means that, if the data persists on the media, future optical playback systems will readily be capable of being adapted as necessary to read data stored permanently on optical discs.

A WORKING DEFINITION OF PERMANENT DATA STORAGE

The title of this paper includes the word "permanent". But if permanent means that something endures forever, nothing that science acknowledges is permanent. If we accept a more practical definition, it means that something lasts an extraordinarily long time. But that is a relative term, and lacks a specific meaning.

What would permanent mean, when it comes to data storage? We should probably look to non-digital data storage media for answers. Books are generally considered rather permanent, and if properly manufa ctured and stored, can easily endure for centuries. It would be great if a digital data storage medium could have a similar life expectancy (LE) – a term which should be defined. We accept the definition given by ANSI/AIIM:

"Life Expectancy: Length of time that information is predicted

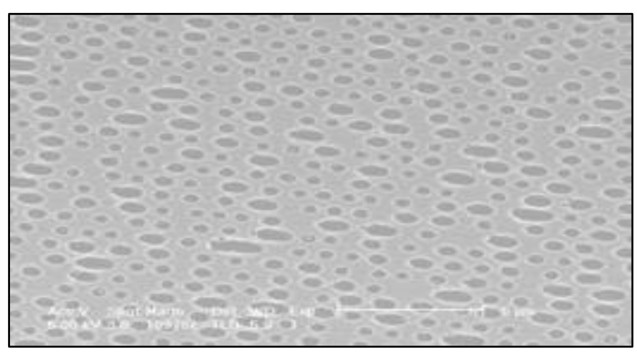

Figure 3: Marks made on the M-Disc, as seen with an SEM.

to be retrievable in a system under e xtended-term storage conditions. Rating for the ‘life expectancy’ of recording materials and associated retrieval systems. NOTE: The number following the LE symbol is a prediction of the minimum life expectancy in years for which information can be re trieved without significant loss when prope rly stored under extended-term storage conditions, e.g., LE-100 indicates that information can be retrieved for at least 100 years of storage”[5]

One definition for permanence has been proposed for paper:

> “Permanence: The ability of paper to last at least several hundred years without significant deterioration under normal use and storage conditions in libraries and archives.”[6]

From a practical perspective, permanence for digital data storage must be defined similarly. We would propose the following:

Permanence: The ability of a digital data storage medium to last at least two hundred years without significant deterioration under normal use and storage conditions in libraries and archives. This means there is a 99.99% confidence of complete data recovery using the intended read mechanism or hardware.

Using the above definition, there is only one digital data storage medium which even comes close to meeting this standard of permanence – the M-Disc, from Millenniata, Inc. All other digital data storage media have very serious limitations when it comes to permanence.[7,8,9,10,11]

The M-Disc is presently available in both DVD-R and BD-R formats, with capacities of 4.7 GB and 25 GB, respectively. Future enhancements include the development of a dual-layer BD-R format and a 2-sided, 2-layer BD-R format, for capacities of 50 GB and 100 GB, respectively.

VERIFICATION OF LIFE EXPECTANCY

One of the first questions raised by the definition we propose is how the LE of a medium can be verified. The science of accelerated testing has long been accepted as a reliable way to determine LE, using the Arrhenius and Eyring equations as the scientific foundation.[12]

For optical discs, the standard method for determining the LE for a given medium has been accepted to be the international standard, ISO/IEC 10995[13]. This standard outlines four test conditions (85°C, 85% RH; 85°C, 70% RH; 65°C, 85% RH; 70°C, 75% RH), as well as the hours required for each set of conditions and the criteria for failure. These tests, even when run in parallel, require a minimum of 2,500 hours of accelerated aging, which is over 104 days at 24 hours/day. Adding in the required testing at appro-priate intervals means that such a test requires many months, much test equipment, and careful testing and analysis. In short, such a test is a major undertaking, and is rarely done in its entirety.

Nevertheless, a full ISO/IEC 10995 test has been completed on the M-Disc, and the preliminary results have been analyzed. One parameter that has shown a clear difference is jitter, which is an electronic measure of the variation in the shape and position of the physical bits on the disc. If this parameter increases over time, it is a clear indication that something physical is happening to the bits on the disc – a sign of degradation. Figure 4 combines the jitter measurements in all four test conditions specified in the ISO/IEC 10995 test, where the numbers indicate the temperature (in degrees C) and the relative humidity (in percent). The comparison was between the best archive media (JVC Archival Grade DVD-R and Verbatim UltraLife Gold Archival DVD-R) and the M-Disc. In all four test conditions, it is readily evident that the jitter in the archive media is increasing significantly over time, while the jitter in the M-Disc increases only slightly. For example, in the harshest of the four test conditions (85°C/85% RH), over the course of the 600 hours of the test, the jitter in the archive media increased by nearly 5% (far exceeding the specification) and was still climbing, while the jitter in the M-Disc increased by less than 1%, and had apparently stopped increasing. This trend is evident in all four test

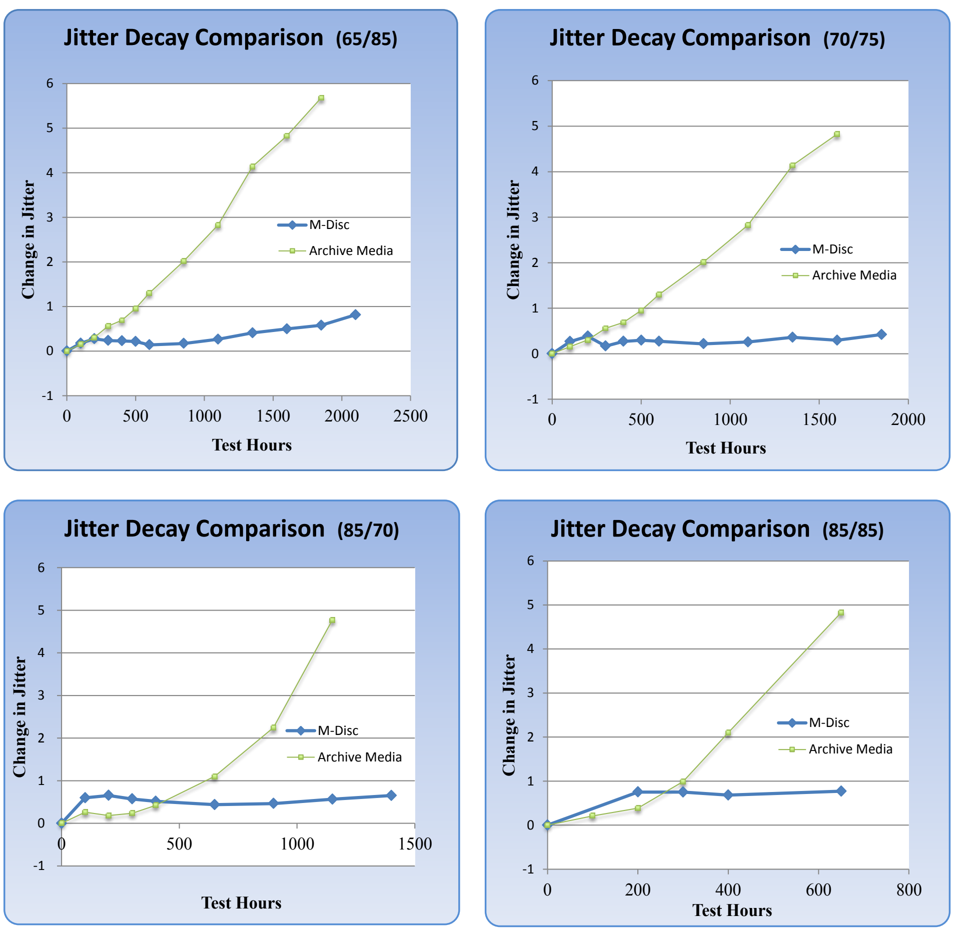

Figure 4: Jitter as a function of test hours, for all 4 test conditions of the ISO/IEC 10995 test.

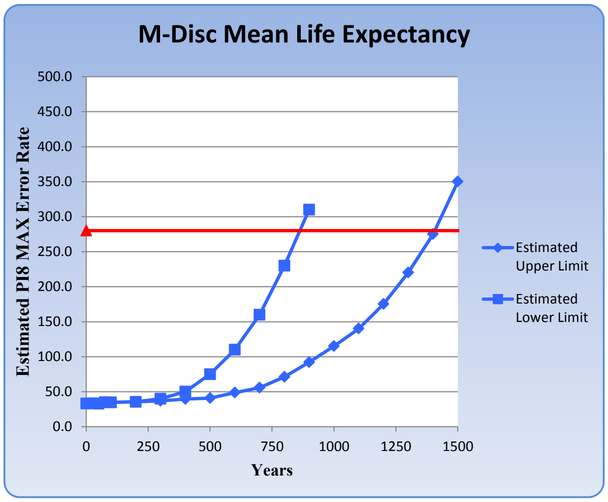

Figure 5: Mean LE for M-Disc, based on ISO/IEC 10995 test results.

conditions, and clearly shows that the best archive media available is not stable with time, while the M-Disc is extremely stable in all four test conditions.

The PI8 maximum error rate is the main specification by which the lifetime expectancy (LE) is estimated. The specification limit is 280. Figure 5 shows the mean LE for the M-Disc, based on the results of the testing that has been done. From these results, it appears that the LE of the M-Disc is between 750 and 1400 years.

Research has been conducted in applying these same recording-layer materials (from the M-Disc) to a Blu-ray type disc. These results have been highly successful, resulting in the announcement at the 2013 Consumer Electronics Show (Las Vegas, NV, Jan 9, 2013) of the availabil ity of a Blu-ray density M-Disc, to be manufactured by Ritek, with availability in July 2013. Lifetime testing has not yet been performed on this disc, but because it uses the same materials approach of the M-Disc and is thus a unique type of archi val disc, it is highly probable that the lifetime will also be dramatically better than anything presently on the market. We also believe that these same materials can be applied to produce a permanent optical tape; this is discussed in more detail later in this paper.

STATE OF THE ART IN PERMANENT SOLID-STATE STORAGE MEDIA

Solid-state storage consists today of DRAM, SRAM, and Flash memory; the characteristics of these memory types are

Table 1: Characteristics of main memory types available today

Characteristic	DRAM	SRAM	Flash
Write speed	50 ns	10 ns	500 ns
Read speed	50 ns	10 ns	100 ns
Data retention	Volatile	Volatile	8-10 years
Relative cost	1.0	4.0	1.0
Typical density	16 Gb	4 Gb	16 Gb

summarized in Table 1. Two of these memory types are volatile, which means that the data disappears if power is lost; these memory types are thus useless for archival storage. Flash memory is the only main type of memory available today that is nonvolatile, which is what has made "flash drives" (also known as "thumb drives") not only possible but also very popular. However, for archival storage they are not practical, since the data is stored as a charge on a floating gate (see Figure 6), and this charge will eventually leak away. Most estimates are that this data will persist for only 8-10 years.

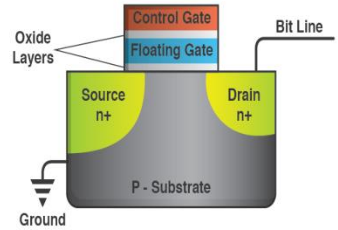

Figure 6: Diagram of a flash memory cell, which stores data as a charge on a "floating" gate.

The materials approach to understanding these solid-state storage media explains why they are not permanent. For DRAM and SRAM, their structure is such that they cannot store data at all, unless they have power present; the "1" state is only distinguishable from the "0" state if power is continuously applied. Flash memory is quite different, requiring power to write and read the data, but not requiring power to maintain the data, at least for several years. However, because it is erasable, there is very little electrical difference between the "1" state and the "0" state, and there is no material difference between these states.

If a storage medium is erasable, by nature it is not as permanent as a medium that cannot be erased. If data is recorded by a physical change in the medium (as in scribing on gold plates; see Figure 7), such data storage is inherently much more permanent that any non-physical change. Another way of saying this is that the greater the difference between the "1" state and the "0" state, the more persistent the data will be.

Figure 7: An Etruscan gold plate with lengthy inscription, dated to about the 5th century BCE.

The early da ys of solid-state storage (the 1960s and 70s) included the development of s everal types of non-volatile memory: ROM (read-only memory), PROM (programmable read-only memory), EPROM (erasable PROM), and EEPROM (electrically-erasable PROM). ROM cannot be programmed by the user, and so is not useful for end users – it must be programmed when it is manufactured. PROM was very useful for niche applications, but it suffered from a defect known as dendrites (see Figure 8) – small tree-like growths that tended to grow and short out the

Figure 8: A dendrite has grown from the remains of a programmed cell (blown fuse) in a PROM storage device.[14]

blown fuse, thus compromising the data. EPROM and EEPROM are the immediate predecessors of today's flash memory, which has been discussed already. Thus, historically, none of the solid-state storage options has been viable for permanent digital data storage.

For solid-state storage to be permanent, both the programmable cells (where the bits are actually stored) and the read/write circuitry must be permanent. Fortunately, half of this problem is already solved, since integrated circuits (ICs) have lifetimes typically rated in FITs – Failures In Time, or the number of devices that fail in 10^9 hours of operation (approximately 1,000,000 years). Typical FIT numbers for ICs are in the range of 50, which means that the typical IC can be expected to operate for over 22,000 years.

One way in which the data in a permanent solid-state memory device could be stored would be to represent a 1 with an intact fuse, and a 0 as a blown fuse. The user could then program their own data into the memory. But the material for the fuse must be as long-lasting as the IC itself. To solve this other half of the problem of permanence for solid-state digital data storage, we need a materials approach – we must find a material which is extremely stable, somewhat resistive but not an insulator, and which does not grow dendrites when the fuse is blown. This would produce fuses which are extremely long-lasting either intact or as programmed. Fortunately, such materials exist, and we have been successful in producing and programming fuses of these materials.

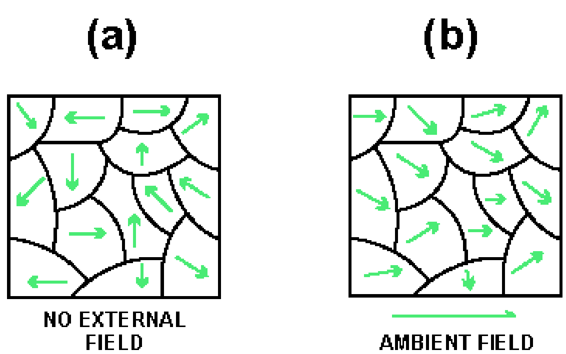

Figure 9: Schematic drawing of magnetic domains, subject to no external magnetic field (a), and subject to an ambient magnetic field (b).

There is still much development work to be done before ICs using these design concepts can be commercially available, but the fundamental concepts have been proven in the lab and in lifetime testing.

STATE OF THE ART IN PERMANENT ½-INCH TAPE STORAGE MEDIA

Current ½-inch tape products are all magnetic, and suffer two main degradation mechanisms. The first is the slow relaxation of the magnetic domains; these domains are how the data is stored in all magnetic storage products. These domains are depicted in Figure 9. When not under the influence of any external magnetic field (Figure 9a), these domains are randomly oriented. After being subjected to an ambient magnetic field, these magnetic domains change to produce a net magnetization in one of two directions. This difference in the direction of the remaining magnetic field is the difference between 1s and 0s in digital data. With time (and with temperature), these magnetic domains begin to relax, slowly reverting to their original random orientation, and slowly degrading the difference between the encoded 1s and 0s. Eventually, so many of these bits will have degraded that reading a file back will have become impossible.

The other main degradation mechanism of magnetic tape is the delamination of the recording layer from the plastic substrate of the tape. In modern ½-inch magnetic tape, the recording layer is applied to the polyester substrate in a printing process. The binder in the recording material is what keeps the recording material bound to the substrate. But binder materials are generally organic, and they degrade with time, which leads to small pieces of the recording layer peeling off ("delaminating") from the polyester substrate. Wherever these pieces peel off, the data is irretrievably lost.

The two main advantages of magnetic storage are that it is non-volatile, and that it can be re-recorded an infinite number of times. However, as mentioned before, as soon as we make a medium in which the data can be erased (changed), we have lowered the barrier to losing that data with time. We also need to solve the problem of delamination. A materials and processes approach to this problem would be to choose a material which can be readily and permanently altered with a laser in a way that can be detected with a laser, and then to deposit this layer of materials in such a way that they are permanently bonded to the polyester substrate without the use of any binders. Because we have already done this with the recording layer of the M-Disc, our research has focused on applying these materials and processes to a new optical tape, one which is as permanent as the M-Disc.

Figure 10 shows some marks which have been made on this recording layer in ½-inch tape, using less than 10 mW of optical power, which is readily available

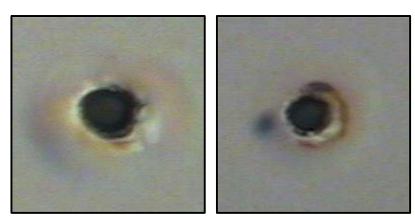

Figure 10: Marks made in our recording layer on ½-inch tape, using less than 10 mW of optical power.

from today's solid-state lasers as used in optical disc drives. These marks are expected to be as permanent as those on the M-Disc – they should endure for centuries.

We have also used a very different process for depositing the recording layer on the polyester substrate, a process which uses no binder, and which is widely used in industry today. We have tested the adhesion of the recording layer to the substrate, using both the tape test method and using a tape tension & bend method. Both of these tests were rigorously applied, and no material was removed. We have shown that the recording layer is permanently bonded to the substrate, so delamination will not occur.

Using a mark density equal to that in Blu-ray discs, and a track density equal to that of today's ½-inch tap e, we would expect that an optical tape cartridge should be able to hold multiple Terabytes, which is roughly equivalent to the density of today's ½-inch magnetic tape.

CONCLUSION

We believe there is an urgent need for some way to store digital data permanently[3,4], and that the M-Disc is one viable solution to this need. We also believe that we are well on our way to providing two additional media for permanent digital data storage.

We readily acknowledge that these permanent media are not a complete solution – there must also exist some way to read the data stored on them far into the future. But data on a permanent medium is the *sine qua non* of deep archival data storage. Some future generation may struggle to learn how our data was stored, but if all the marks on the Rosetta stone had faded away, deciphering them would never have been possible. It is relatively easy and inexpensive to read the marks made on optical discs - the hardware is widely available and optical discs are the most widely adopted digital storage medium in history. And wide adoption is a powerful predictor of relative permanence of readability, as witnessed by the many people who still learn to read and write Latin; though a "dead" language, there are hundreds of thousands of documents in that language. If the marks on these Latin documents were to suddenly disappear, Latin would become irrelevant.

REFERENCES

1. 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.
2. 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.
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. ANSI/AIIM TR2-1998, "Glossary of Document Technologies", Association for Information and Image Management International, Silver Spring, MD.
6. ANSI/NISO Z39.48-1992 (R1997), "Permanence of Paper for Publications and Documents in Libraries and Archives".
7. Byers, Fred, "Optical Discs for Archiving", Information Technology Laboratory, Information Access Division, NIST, OSTA, Dec 6, 2004.
8. Iraci, Joe, "The Relative Stabilities of Optical Disc Formats", *Restaurator: International Journal for the Preservation of Library and Archival Material*, 2005, 26:2, 134-150.
9. 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
10. Shahani, Chandru J., Basil Manns, Michele Youket, "Longevity of CD Media: Research at the Library of Congress", Preservation Research and Testing Division, Washington, DC.
11. Slattery, O., Lu, R., Zheng, J., Byers, F., Tang, X, "Stability Comparison of Recordable Optical Discs- A study of error rates in harsh conditions," *Journal of Research of the National Institute of Standards and Technology*, 109, 517-524, 2004.
12. Gillen KT, Clough RL, and Wise J, "Extrapolating Accelerated Thermal-Aging Results: A Critical Look at the Arrhenius Method", Albuquerque, NM, Sandia National Laboratories.
13. ISO/IEC 10995: Information technology – Digitally recorded media for information interchange and storage – Test method for the estimation of the archival lifetime of optical media", ISO/IEC 2008, Geneva, Switzerland.
14. Pecht, Pecht, Michael, "Handbook of Electronic Package Design", Marcel Dekker, Inc., New York, NY, 1991, p. 59