Lifetime of digital media: is optics the solution?

Erich Spitz [1], Jean-Charles Hourcade [2], Franck Laloë [3]

[1] Academie des Sciences, Academie des Technologies ; Thales, RD 128, 91767 Palaiseau Cedex, France (erich.spitz at thalesgroup.com)
[2] Academie des Technologies ; Red Cat Technologies, 3 Rue Cauchois, 75018 Paris France (jean-charles.hourcade at numericable.com)
[3] Laboratoire Kastler Brossel, Ecole Normale Supérieure ; CNRS and UPMC, 24 rue Lhomond, 75005 Paris, France (laloe at ens.fr)

ABSTRACT

While the short term and mid-term archiving of digital data and information can be handled reasonably well with modern techniques, the long term aspects of the problem (several decades or even centuries) are much more difficult to manage. The heart of the problem is the longevity of storage media, which presently does not go beyond a few years, maybe one or two decades in the best cases. In this article, we review the various strategies for long term archiving, with two main categories: active and passive. We evaluate the various recording media in terms of their longevity. We then discuss the recordable optical digital disks (RODDs) and the state of the art in this domain; the present situation is that, with the techniques that are implemented commercially, good prospects for long term archiving are not available. Nevertheless, the conceptual simplicity of RODDs could be exploited to create new recordable digital media; the improvements that are needed seem to be reachable with reasonable development effort. Since RODDs are now in strong competition with other systems (hard disks or flash memory for instance) that constantly make enormous progress, there seems to be little hope to see RODDs win the race of capacity; nevertheless, longevity could provide them with a new market, since the need for long term archiving is so pressing everywhere in the world.

INTRODUCTION

An enormous amount of numerical information is created every day in the world, a fraction of which has more than short term interest and should really be preserved in the long run. For instance, all of us would like to preserve family memorabilia, documents, photos and souvenirs passed on by our forebears - as well as those we possess today and wish to hand down to our own children and grandchildren. There are also the medical files (in particular images) that must be kept for 30 years, legal documents (property deeds, for example) the value of which must last for centuries, administration files (for example those used to calculate pension rights), etc. Another important category is given by public data with high intrinsic value that we must clearly preserve: scientific data such as data from satellites and space probes, data from large scale particle accelerators, data bases with biological or sociological measurements. Moreover, the private sector has to guarantee future access to information related to complex products and infrastructures, both material and immaterial (building plans, power stations circuits, *etc*.); oil companies must preserve

Quantum Sensing and Nanophotonic Devices VII, edited by Manijeh Razeghi, Rengarajan Sudharsanan, Gail J. Brown,
Proc. of SPIE Vol. 7608, 760802 · © 2010 SPIE · CCC code: 0277-786X/10/$18 · doi: 10.1117/12.848948

geological data related to exploration sites, where exploitation can be stopped for several decades before starting up again; the aerospace sector must by law keep its aircraft drawings and associate data for 70 years. In this article, we shall only discuss the conservation of the category of information that has long-lasting value, information that it is highly desirable to archive in the long term covering several decades, centuries, or maybe even more - who knows ?

Fortunately, digital information can be copied without errors an almost unlimited number of times. Indeed, this is a marvellous and new feature: before, the very act of copying always implied a progressive accumulation of errors, which were not necessarily visible at the first copy; they actually increased with each new copy until, in the end, the information had completely disappeared. This wonderful property of digital information was made possible by the extensive use of error correction codes, which are in turn a long term consequence of fundamental work of renowned mathematicians. This started with Evariste Galois in the 19^{th} century and lead, much later, to the famous article of Irving Reed and Gustave Solomon "Polynomial codes over certain finite fields", published in 1960 in the "Journal of the Society for Industrial and Applied Mathematics". Without error corrections codes, no digital device or computer would be usable at all: for instance, each time an author saved his file while writing it, there would be a significant chance of introducing a new error (the basic error rate on a hard disk, without any error correction code, is about 10^{-5}). In 1960, there was no practical way to implement the Reed Solomon codes, so that their contribution remained dormant for some time; but now these codes are implemented everywhere and used extensively. This is one more example of the exceptional creative power of research driven by pure scientific curiosity, which can lead to unexpected and very important applications, sometimes in completely unexpected domains.

In theory, thanks to this absence of copying errors, provided someone takes care of copying data in time, there should be no reason why digital information could not be transmitted from generation to generation, over centuries or millenia. Of course, just copying is not sufficient; for instance, it is also clearly necessary to provide information on what information is archived (metadata); a careful choice of formats has to be made to avoid the rapid obsolescence of file formats; several copies have to be made to secure the information against sudden catastrophes (fire, flood, theft, .. as well as hard disk crashes!), *etc*. But this is just applying common sense and can be mastered with sufficient care. So, if we assume that this has been done, is the problem solved?

Alas, the reality is that there is a long way between these exciting theoretical possibilities and practice. In the short term, no special problem occurs: the creation, transmission and storage of digital data is relatively easy both from practical and technical viewpoints. Nevertheless, in the long run (more than a few years) the conservation of this information gives rise to many difficulties, which have several different components. This article focuses on the major of them, the lifetime of the physical media themselves, the problem that is at the heart of any archiving project. Clearly, the key to the conservation of digital information over centuries is the use of very stable recording media.

.

The physical and chemical stability of the presently available digital recordable media is indeed the core of the archiving problem. If all the supports are degraded so that we are no longer able to recover the digital data, in other words if the information itself has disappeared because the media has become totally unreadable, even the best metadata will not help much! In reverse, we can reasonably think that, after a century, if a given data is still present and "usable" on a given support, and if it is really interesting, someone will do what is necessary to extract it. Experience shows that, if people know that some really essential information is still there, they will always find a way to have access to it (devices to read the data, software, etc.).

In this article, we specifically discuss recordable digital supports, those that can be used to keep personal data; we shall not cover the case of industrially "pressed" optical disks, sold over the counter, with data (music, movie films, etc.) already contained. Moreover, we shall examine solely those products that are either currently available, or not too far from being marketed - not those that are not yet in an advanced stage of laboratory development, despite the interest and attractiveness in some purported and advanced solutions. This article is a shortened version of a report[1] written by the authors for two French Academies (Sciences and Technologies); the full version (in French) of this text will presumably be publicly available shortly.

1. Archiving strategies

In general terms, what strategies can be envisioned to ensure long term data conservation?

1.1 Passive method: "Archive and forget"

This is the time-old strategy, the most natural process: we tend to keep our precious souvenirs deep down inside a cupboard, out of sight, out of daily reach. Instinct drives us to using this strategy for the simple reason that this is the process whereby family souvenirs and treasures have been handed on to us, from generation to generation. It is also the process adopted by libraries who keep precious books and manuscripts from the past on their shelves.

Manufacturers of CD-Rs indeed play on this instinct and orient their marketing campaigns accordingly. Thus, in 2003, we could see large print advertisements in a prime French daily newspaper that ran "You can now save your perfect and inalterable recordings, for all eternity!" Sadly enough, there is a wide gap between these claims and reality: recent precise measurements have shown how much of a dangerous delusion there is in these statements. Yet this need to preserve objects was so deeply ingrained in everybody that many continue to believe in it, not least the public at large who sometimes confuses the differences between recordable disks and industrially pressed disks (which do have much longer life expectancies).

1.2 Active method : "Archive and migrate"

The active method calls for a constant follow-up of the data ensuring that, when the previous support shows signs of ageing, they are migrated properly to new supports, and so on without any end. If, for example, the support is a magnetic tape, the time lapse between copies is typically 4 to 5 years. Technical surveillance is needed here to ensure also that the migration process is launched before the initial support has degraded too far, *i.e.,* as long as the error corrections codes operate

[1] E. Spitz, J-C. Hourcade and F. Laloë, "Préservation à long terme de l'information numérique", report to the French Academies of Science and of Technologies, november 2009.

with perfection. It is also recommended to duplicate onto various support materials (and if possible with varied brands) and that the results be stored in different places, thereby avoiding the risk of loss due to natural catastrophe (fire, flooding, etc). This active method is currently used by major institutional actors whose missions explicitly include long term data conservation: national libraries, large international scientific projects, etc. Banks also use this strategy, as do certain key industries and other public and/or private establishments. The process involves teams of specialists, who are required to apply appropriate operational protocol standards. They often use robots that can handle large ensembles of magnetic cassette tapes (containing hundreds of Tera-bytes of information), capable of retrieving and reading segments on demand. These robots automatically roll and unroll tapes, to prevent their surfaces from mechanically adhering to each other or undergoing magnetic print-through from spire to spire. A constant check of the error rates determines appropriate scheduling for the next data migration.

This strategy is not easy to implement by a private person, because mainly of the constant follow-up protocols involved. It is all the less easy since the existence of error correction codes may give the false impression that CDs are in perfect condition, as we mentioned earlier; but they will suddenly cease to be readable, when the level of errors accumulated exceeds a given threshold. It is therefore necessary to apply deeper levels of analysis, beyond a simple “readability” test. Data that is left unattended will necessarily “die” after only a few years. The option of allowing everyone to ensure his/her own information migration protocols therefore seems somewhat hazardous

Another problem, of course, is the cost. Costs incurred here relate essentially to personnel, technical premises and possibly include cost of the robot systems employed, but they are far less dependent on the cost of the support materials and devices; assessments that are based on the latter only are totally unrealistic, since this part is negligible. A reasonable evaluation is a figure between 500 $ and 1 000 $ per To and per year (see § 1.4).

1.3 Remote archiving: “Not seeing what is happening”

Remote archiving is a third possibility; it consists in entrusting data to a service contractor. For a negotiated price, the contractor will accept responsibility for data conservation. Thanks to the Internet, the contractor can be physically located at a very great distance from the data source. It is even possible to have backup copies of the same data in several different countries or continents. Modern techniques enable a single file to be distributed to different support machines, even if they are separated by thousand of kilometres. Multiple remote archiving provides excellent guarantees against sudden loss of data, due to natural events or accidents. Some see this as the panacea, all the more since companies such as Google offer free storage to its clients (the cost being borne by advertisement revenues). This notion of having one’s data spread throughout the world is fashionable, as are concepts such as “cloud computing”, “grid computing”, etc.. But, although this is a satisfactory approach to data backup, it does not guarantee long-term archiving possibilities with reasonable safety.

Companies that propose remote backup often operate through a process akin to that adopted by insurance companies: a small-scale digital data storage company will insure itself, subscribing a policy contract with a larger company, hiring storage space to duplicate its own data, and the process repeats with even larger companies. The end result is that the risk of sudden accidents or impairment to one device or process in the line is considerably decreased.

The method of remote data delegation does not preclude problems from occurring. Of course, the Internet itself does not store anything; it just provides “pipes” to connect users to distant conservation storage facilities, or storage tanks. Data stored in Asia or Africa are, of course, also physically located on hard drives or tapes, which have the same lifetime limitations, the same sudden breakdowns, the same ageing processes than everywhere. It is therefore self-evident that the contractor must also solve the same technical problems than here, and opt between one of the above two strategies. One then hopes that the strategy will be selected with utmost care, in order to reach a reasonable degree of certainty to ensure perfect conservation and service continuity!

Moreover, with this strategy, the conservation system – which was not exactly transparent – becomes even more opaque. The question arises: is it really safer to safeguard one’s personal treasures spread out or duplicated on a global set of sites, where there is no possible way to ascertain how the conservation process is implemented? If we analyse the situation, we realise that the delegation strategy is probably more suited to conserving short-term or mid-term data, say over a span of a few years, than much longer times. However, although it is common practice to approach distant service companies for relatively short-term missions, it is not common at all to do this for operations spanning several decades. What, for example, will happen if the company in question goes bankrupt and disappears? Who will be responsible for returning the data to the initial owners? Will the latter have enough information to know how and to whom they should address a claim for return of data? Then we are confronted with the issue of data confidentiality, which we know can be crucial inasmuch as encryption codes are never 100% sure. Lastly, it is also self-evident that this strategy is risky in terms of geopolitical implications.

In the same vein, recently, we read quite a lot about “virtualising files” (the term ‘virtualising’ covers different protocols from one service company to another). The method probably can provide large-scale services, but in terms of long-term data conservation, there is in fact no particular gain to be made.

1.4 Going back to analogue recording

The world-famous North American film industries use digital techniques extensively to make films, but seems today on the verge of an amazing step backwards. A well-documented 70 pp. document, “The digital dilemma”, issued in 2007 by the “Science and Technology council of the Academy of Motion Pictures, Arts and Sciences” sets out their reasons for this. In particular, in Ch.6, entitled *“Digital motion picture archiving economics”,* this document contains a detailed argumentation related to the costs associated with migration strategies, with the perspective of conservation times exceeding the century. The bottom line, in essence, is that the annual cost as a whole lies probably somewhere between 500 $ and 1 000 $ per year and per Terabyte.

The conclusion drawn by this report is that in order to save motion picture films on the long-term, the best way is to make large dimension analogue copies (even if the original is filmed using digital techniques) and separate monochromatic copies (to avoid colour chemical degradation), and, final step, to store these copies in deep underground repositories where ambient temperature and humidity can be closely monitored and controlled. Digital archiving is not eliminated as an option, but its use would be limited to short and medium term purposes. Here we have a beautiful illustration of the difficult problems raised by digital long-term data conservation.

There are also technical processes that ensure highly resistant engraving of analogue images on optical substrates, which may contain several thousand “pages” of documents, images, *etc*. But, of course, there is neither sound track nor motion picture files. Moreover, these techniques are far too expensive for systematic use by private persons.

2. Digital recording media

We now briefly review the digital supports that are available today, and we shall return to the question of optical digital disks (ODDs) in the following section.

2.1 Magnetic tapes

The advent of magnetic tapes largely precedes the massive diffusion of digital techniques; they were omnipresent in analogue recording processes (magnetic tapes were invented in 1928). They began to be used for digital data storage as of 1951. Historically therefore, they represent the first digital storage devices that were in widespread use, well before the marketing of magnetic disks; even today they still represent the base-line for “active strategy” (generally in the format of tape cassettes). In terms of costs and storage capacity, tapes prove to be particularly economic. In contradistinction, however, they are relatively sensitive to time factors: the oxide layer tends to become unstuck, peels off and ends up as oxide dust, or it sticks to the reading head. The plastic support itself becomes brittle with time. To be certain that a tape remains operational, it must be unrolled and rerolled at regular intervals.

Tapes also require relatively long access times for a given piece of data. They are thus primarily used in a backup safeguard mode for data conservation, but not as the communication support *per se* (over Internet, for example).

2.2 **Magnetic hard disks**

Hard disk drives (HDDs) were introduced in 1956, initially as devices reserved for big computing centres. The underlying physical principle is the same, but the access time to the data is far faster. In recent years, hard disks drives have not stopped making spectacular improvements, both in terms of capacity and price reduction. The very widespread impact of basic research carried out by Albert Fert and Peter Grunberg is well-known; for their work, they were jointed awarded the 2007 Nobel Prize for Physics. Today, we can purchase over the counter a mobile 1 Tbyte HDD device for less than 100 Euro, this being almost inconceivable a few years ago. HDDs, however, remain complex mechanical devices where the mobile reading head detects and reads data from a distance of around 10 nanometres from the surface of the magnetic disk spinning at about 10 000 rev/min. The active layer of the disk is about 10 nanometres thick, under a protection layer only 5 nanometres thick. Given the required accuracy of the mechanical components, we readily understand that they constitute fragile devices that cannot easily withstand shock. Moreover, even in the absence of any mechanical disturbance, they occasionally and unexpectedly “crash”, leading to complete loss of data.

There is a non-zero probability that a HDD that has been stopped will never restart, and here again most of the time the result is total loss of data. It is for this reason that in large computer centres, the HDDs are continuously powered up and remain in this mode for their total operational life. This is an important point to bear in mind: if we wish to evaluate the viability of a given conservation strategy for digital data over the networks, connecting to large-scale computer centres, then we must include the associate power consumption and environmental constraints.

Recent studies[2] show that it is difficult to predict failures and HDD behavioural events over several years[3]. If we use HDDs for the purpose of long term conservation, we should allow for a generous redundancy of the data recording processes (for instance the so-called technique of mirror sites); this is a “doubly active” conservation process, not only in terms of the migration management and technical teams but also in terms of machines that never stop (contrary to magnetic tapes that are stored in robots or cabinets).

The HDDs in continuous motion throughout the world provide permanent platforms for the success of research engines such as Google. This corporation possesses a large number of hangar-like buildings - their exact number is, apparently, a secret - spread over various continents and housing myriads of HDD assemblies. This enables a high degree of data redundancy, spread over the world, and avoids the risks of catastrophic failure or loss (fire and flooding). When failures do occur, they are repaired rapidly, but the down-time variations do have an observable effect on the response of the search engine on a day-to-day basis; everyone can see *de visu* that Google is sensitive to HDD breakdowns.

2.3 Flash memory

Over the past few years, flash memories have become very popular, essentially because they are so convenient. They are found in photographic cameras, video recorders, USB keys, walkmen, and they now even sometimes replace the HDD in laptop computers. Flash memories rely on the physical move of electronic charges in small transistors with two grids, a control grid and a floating grid in an oxide suspension. Flash memories are, however, subject to wear, since the number of transition cycles is limited (somewhere between 10 000 and 100 000 cycles). In terms of rapidity of access to the data, they prove very efficient, even better in some cases, than HDDs; as they have no moving parts, there is no need for a high precision, rapidly rotating, device. The main reasons why flash memory has not superseded HDDs in the PC market are the cost per bit storage capacity, and the somewhat limited number of possible writing cycles.

---

2 E . Pinheiro, W.D. Weber and L.A. Barroso (Google Corp.) “Failure trends in large disk drive population”, 5[e] USEMIX conference FAST (2007). B. Schroeder and G.A. Gibson (Carnegie Mellon Univ.), “Disk failures in the real world”, 5[e] USEMIX conference FAST (2007).

3 RAID techniques (Redundant Array of Independent Disks), now a relatively common feature, prove efficient to counter sudden hard disk crashes, but do not guarantee very long-term conservation of data.

For the moment, nobody really considers flash memories as long-term storage devices. We can easily imagine that avoiding electronic charge migration is more difficult than avoiding a change of the orientation of magnetic domains; a concern is for instance ionising radiations (cosmic rays, or solar wind, for example). But there are yet to come detailed studies on these issues; it would prove very interesting, perhaps, to see if accessible modifications (such as voltage change) could help here. Moreover we can envision "periodic refresh modes" of the data stored on the flash memories. If we are considering long-term data conservation, the question of the number of useful cycles in the device's life becomes relatively secondary: if we could assemble a flash memory with only, say, 4 to 5 cycles, but with a guaranteed conservation factor of several decades, this would constitute a major leap forward!

2.4 New devices

Many other possible support alternatives are being developed. A lot has been heard about IBM's Millipede, based on MEM type techniques, or about atomic microscope memories. It is difficult here to give an exhaustive list of all possibilities, given that not all the projects underway have given rise to publications. A totally new and unexpected idea could arise at any moment, putting to use some new discovery in a wide range of physical science fields. Among the better known projects, let us just mention:

- Phase change random access memories (PRAMs), the elements of which are made of chalcogens;
- Resistive random access memories (RRAMs); HP Labs have announced a prototype baptised "memristor";
- The spin-torque transfer RAM (STT-RAMs), improved magnetic RAMS (MRAMs); a dozen or so major electronics companies in the world are working on market-ready development of these and similar products;
- Carbon nanotube memories (NRAMs), which are seen as highly promising in terms of data density, access speed, inventiveness to change and even in terms of durability (ruggedness); this project, that came out of a Harvard University lab., is being developed by a private company;
- Another university project, this one from Berkeley[4], where the carbon nanotube is used as a magnetic conduit in which a minute mass, in fact a magnetic monocrystal, moves;
- The so-called "Digital Rosetta Stone" which designates a joint university project (Keio and Kyoto Universities) and the Sharp Corporation[5] who in essence constructed a pile of magnetic slices, subsequently buried in a silicon mould.

This list is not exhaustive; what is however interesting to note is that the very idea of creating a real long-life memory is "in the air", and the need for such a device is making itself more and more demanding.

---

4 G.E. Begtrup *et al.* "Nanoscale reversible mass transport for archival memory", *Nanoletters*, Vol. 9, 1935-38 (2009).

5 http://techon.nikkeibp.co.jp/english/NEWS_EN/20090618/171883/)

3. The recordable optical digital disks (RODD)

Optical digital disks (ODDs) *a priori* seem adapted to very long-term data conservation. In essence, they comprise a simple support that is easy to manipulate, totally separate from the reading/recording devices, they have no moving parts, which ensures zero surface wear. We now discuss the present situation of ODDs.

3.1 General situation

Annual world production of optical digital disk is huge; over 10 billion ODDs are produced each year. Their main attractiveness lies in the fact that they can store data, can be sent elsewhere by post, can provide a useful support for participants after a conference, etc. For the time being, RODDs come in three main formats : CD-Rs or CD-RWs, DVD-Rs or DVD-RWs, and the latest is BD-R (Blu Ray). These supports are conceptually very simple: they are each made in one piece with no moving parts; they are completely passive once recorded – they are read in a purely optical mode and therefore there is no mechanical contact nor wear; their sole function is to send back light to a light detector. It is probably for this reason that many initially thought that they are eternal, but no wear during the reading function is not sufficient to ensure a perfect longevity! Unfortunately, today, we now know that in reality such supports are continuously degrading, even when they are not in use.

For several years now, we have been hearing about super-resolution properties and holographic disks. In the former, the hope is to be able to place close to 500 Gbytes on a single disk; TDK has just announced a 320 Gb disk made of 10 layers of recorded data. Alas, even this impressive capacity already looks limited when compared to the capacity of hard disks already available commercially: in fact, one gets the feeling that the RODD's have already lost the race for large capacities. The second process, holography, would enable techniques to access information in parallel, but it is still limited today exclusively to laboratory programmes that seem somewhat far removed from industrial markets (cf. however, § 5.1.3).

3.2 State of the art

RODDs in their present form are far from simple objects with evolutions that can be predicted fairly easily, as initially hoped for. Under market-place pressure, the solutions that have been chosen by the manufacturers privilege capacity and speed of data retrieval as their main targets, and also of course a low price; longevity of the product is scarcely taken into account.

Several manufacturers have proposed what they call "archiving disks", sometimes with the attribute "gold", this being justified by the fact that the normal reflecting metallic surface (aluminium) is replaced by a gold layer. If oxidation were the main cause for ageing, then gold might indeed prove to be an attractive alternative! But, alas, several studies, in particular those of the French Laboratoire National d'Essais (LNE), have shown that it is not the case; in reality, these archival disks are not necessarily better than standard disks, actually sometimes even worse.

Construction of optical digital disks using photosensitive organic dyes is a complex process that is difficult to control. The manufacturers, for example, cannot control the quality of ingredients as essential as the dye solvent, which can vary from one delivery to another; the exact chemical formula, moreover, may be a trade secret that the dye manufacturers protect. The consequence is that the manufacture of optical digital disks is inhomogeneous and not totally reproducible[6], even within a single given model from a given single manufacturer. There can be important variations in quality as a function of the batch, and date. In a sampled small fraction of CD-Rs where the life expectancy was discovered to be catastrophic (one year or less), it turned out that many of the disks incriminated had been manufactured in 2003, but no precise explanation for this problem could be identified. A recent study (2008) carried out by the [French] Directorate for National Archives (DAF), showed that the CD-R models that had appeared to be best quality in the previous analysis campaign (2006) turned out, in fact, to be least rugged[7]. Another study by the LNE[8] shows that local authority archives (French territorial departments), recorded on CD-Rs are now partly illegible, so that data have been lost. The LNE, however, did observe that some of the CD-Rs are of relatively good quality and seem ready to be able to conserve their data for 10 or 20 years, maybe even more. This is a favorable case where the RODD has a better longevity than a magnetic support; but it does seem very difficult to pinpoint the rules that would enable in advance that given manufactured batches would be of good, *i.e.,* up to standard quality.

The continuous race towards higher and higher storage capacity does not help matters. The underlying physical reason is easy to describe: the degradation process for a physical object, for example, by diffusion of water molecules from one layer to another, is all the more rapid that the physical scale is smaller (the specific diffusion factor varies as the square of the distance covered). An object containing very dense, minute engravings will therefore and necessarily be more sensitive to similar processes than objects with larger physical dimensions. Still in the perspective of proposing higher and higher storage capacities, the "Blu Ray" disks store their data very close to the surface and are protected by an extremely thin layer; it is therefore likely that they will have a shorter longevity, despite the quality of their protecting layer.

---

6 J.M. Lambert and Y. Saunders (LNE) « La conservation des données sur CD-R »[CD-R data conservation], an LNE study dated July 2004 that in a batch of 70 off the shelf CD-R engraved in yr. 2000 some 15% were partly or totally illegible.
Joint study INA-LNE-INA « Etude portant sur la fiabilité du stockage de données sur CD-WORM »["Study as to reliability factors for data storage on CD-WORMs" in yr. 2006 ; some 8% of CD-Rs are substandard after 4 years; those purchaed in yr.2003 prove the least stable.
J.M. Lambert (LNE), « Etude du vieillissement en conditions climatiques sévères de deux références de CD-R » ["Study of storage ageing factors under severe climatic conditions"] , R&D Fact sheet by LNE, April 2007.

7 J.M. Lambert and J. Perdereau (LNE), « Qualité des CD-R disponibles sur le marché pour l'archivage des données numériques », LNE report dated July 2008 ; « Qualité des DVD+R et DVD-R disponibles sur le marché pour l'archivage des données numériques », LNE report, dated October 2008 soulignant la faible qualité des premiers DVD-R, qui s'est heureusement améliorée depuis.

8 Y.Saunders (LNE) « Qualité des disques optiques numériques conservés dans différents services d'archives », final report commissioned by the Direction des Archives de France (DAF) [National Archives] (December 2006) indiquant que « pour un disque sur trois, le risque de perdre des données est bien réel, et très élévé pour un sur cinq ».
This study and several others can be consulted on the DAF site
http://www.archivesdefrance.culture.gouv.fr/gerer/archives-electroniques/stockage/

3.3 Physicochemical processes taking place in the spontaneous decay of RODDs

We may legitimately be astonished to find such a large gap between the hopes that we placed initially in optical digital disks (ODDs), supposed to last forever, and what we learn from the measurements from artificial ageing processes, such as those conducted by the LNE. The reason is that there is a long and arduous path between the concept of an almost perfect recording/writing process, carried out using a laser beam with no mechanical contact (engraving makes us think of the hieroglyphic inscriptions on timeless Egyptian monuments), and available RODDs. The process is complex and involves different and, as yet, not fully understood mechanisms.

This inherent complexity has an immediate repercussion on the difficulty of studying ageing factors and phenomena. If we were analysing a single well-identified chemical process, we could attribute an activation energy level E_a, and predict that the associate speed of evolution would vary as the exponential of the ratio ($-E_a / k_B T$), where T is the absolute temperature and k_B Boltzmann's constant. We could then measure this rate of evolution at different temperatures and deduce a value for E_a, thereby enabling prediction of evolution at any given temperature. This standardised method would make it possible - using accelerated ageing measurements at high temperatures - to deduce by extrapolation the ageing rate at ambient room temperature, which is much lower. Roughly speaking, this is the principle of an ISO standard, but with two parameters (temperature, humidity) that are taken into account, and therefore two exponential extrapolations instead of just one. The method ignores numerous other ageing mechanisms, such as those due to pollution through chemical contaminants, illumination factors*, etc.*.

When we realise that the extrapolations are from measurement times of typically one month to several years, we immediately understand the dangers of the methodology: very small measurement errors lead to large bars of errors in the final result. What is more serious is that, if two basic ageing processes act in parallel, one with a large weight but a long time scale, the other with a small weight and a short time constant, the accelerated ageing measurements only detect the first process (give or take some very small discrepancies). When extrapolating to longer time periods, the values found are then totally false; the end result is that the extrapolated value is far too optimistic.

Actually, evoking this possibility is not unreasonable when we examine the experimental results. If indeed the RODD ageing process was unique, the plots measured under artificial ageing models would appear as simple, monotonous, time functions, all more or less similar to the next (within a possible change of scale). But one only need look briefly at the LNE measurements to see that this is not the case: a great variety of behaviours becomes apparent, with sometimes sudden changes in the slope, etc. This gives the impression that for certain RODDs there is a first process that is followed by another different process – see for instance the break observed in some plots, which really looks like a signature for a multiprocess phenomenon. There are numerous possibilities to explain this observation, such as bleaching of the dye (for example, due to ambient light), diffusion of water molecules between the structural layers, oxidation or peeling of the metallic layers, volume changes in the polycarbonate material, *etc*. Certain processes can lead to a latent behaviour, a phenomenon that is well-known to chemists, who speak of "initiation processes", or "induction phase", followed by a "rapid phase". For example, we can imagine that the faults or cracks that occur in the polycarbonate matrix diffuse thermally, and group together to form empty voids, which in turn diffuse and lead to the observable ageing phenomena.

At the end, the only certainty that emerges from artificial ageing measurement protocols is in fact the raw result: the resistance or ruggedness of certain RODDs to extreme temperatures and humidity levels. Of course, it remains logical to use these data as markers to longevity, since for the moment nobody is in a position to propose a better method. We can entertain the hope that those RODDs that resist best will also prove to be the supports that retain*, i.e.,* conserve their data legible under real storage conditions, but even this is not absolutely certain. The net result is that these measurements are undoubtedly precious, but we must be cautious with respect to lifetime predictions measured in years, since they are probably false by a large factor.

3.3 Other variants

Some hope might come from holographic disks. The motivation for their development seems, mostly, to be increased storage capacity and a possibility to implement parallel reading with several information beams. Moreover, if we can succeed in writing the data in the physical heart of the support and not in a thin surface layer, we can reasonably expect to minimise external physico-chemical attacks.

However, the recording medium itself must be stable and not be subject to spontaneous degradation (for example, when exposed to light). An American company (In Phase Technology) markets a product that is announced to have long term stability. The disks have a photopolymer layer that is light sensitive, and the disks need to be inserted into a special cartridge holder. They can only be read by equipment from the same company. The order of magnitude of the price for such readers and the disks is about one hundred times that of the sort of equipment currently purchased by the general public.

Let us also mention the UDO (for Ultra Density Optical disks), marketed by a British company, Plasmon UK. The process - which implements a support medium that is their propriety brand - uses surface phase changes on the disk surface; the company guarantees a 50 year life expectancy for the product. The customer thereby becomes "prisoner" of a single supplier and, should the company go bankrupt, nothing really guarantees the availability of the data stored over the long term.

A novel project, called "Millenniata" has recently gained attention; it is a product that called for collaboration by staff of the University of Utah[9]. It announces that it will be market ready before the end of 2009 and that the product will be a RODD with a thousand year life expectancy; nevertheless, it is a RODD with a fairly classic layer on a polycarbonate substrate. We can then call to question the longevity of a polycarbonate disk over one thousand years. But, whether or not there is a real and viable technical solution here, it certainly serves to demonstrate that ideas are bubbling around the concept.

There is one exception on this somewhat sombre and dismal horizon: the "Century Disk", an invention that dates back to the mid-eighties. The Century disk was launched initially by a French

[9] *Cf.* http://www.millenniata.com/index.html

company called Digipress, and was later sold out to several successive buyers. The NASA has used these disks to put them onboard its space probes, to provide information about our civilisation on Earth to other distant civilisations! The disks are made of tempered glass, and are recorded via a lithographic process that makes the disks very stable through time. They come out with flying colours at each test phase by LNE. They are therefore of exceptionally high quality and they can effectively guarantee data conservation for a century at least. However, for the moment, their main setback is that the manufacturing process is not yet automated; each disk is made manually. The price of a Century Disk CD or DVD is about 100 Euro, which is too much for a large market.

3.5 **A few ideas towards the realization of a RODD with excellent longevity**

Ideally, stable future RODDs should be both readable and recordable on the same machines as ordinary RODDs. A very widespread diffusion of a standard is a guarantee for long-term availability of the device. If digital disks can be read on fairly common equipment, sold in very high numbers and available everywhere, then the long-term stability of the system as a whole is reasonably assured.

However, requesting total compatibility without any modification to the today's equipment may be asking too much, given the current context. After all, if the reading of future RODDS requires a slight change of the reading equipment, and if these disks have a really long lifetime, a reasonable modification of present day readers remains perfectly acceptable. Our digital readers today can already handle several different types of discs, and adapt to the support characteristics (automatic power setting of the laser beam, for example). They automatically adapt to various formats, as indeed necessary if we wish to play disks that run music and various codex's for video. Increased reader capacity towards longer lifespan disks cannot therefore be excluded. Nevertheless, one prime condition - if one wishes to guarantee a long-standing market for the system - is that upgrading this must not entail a radical hike in the price for the equipment.

We can likewise look to holographic disks as an alternative, where the information is stored in a volume format. This would be a way to avoid surface diffusion and pollution problems. Moreover, to the extent that the holographic content uses a much larger volume of the disk and thereby procures a much higher theoretical storage (3 to 4 Tbytes), we can readily appreciate the interest that such a device raises. But, in this alternative, the compatibility with today's readers would be completely lost.

The long term future for ODDs is in fact uncertain. Even though they are manufactured in billions, the spectacular progress of competing technologies, notably flash memories, is a real threat. ODDs have largely lost the race for high capacity faced with HDDs; the price per Gbyte is comparable, but the transfer rates are not as high. Their main hope for survival seems to lie in an exploration of the ODD's specific features: simplicity of the storage process, separation of the functions. We can raise the hope that marketing such a product with guaranteed longevity (say, a century or more), even if their capacity remains limited for the moment, will create a new market, since the need for archiving is so universal and pressing. We hope that this opportunity will be exploited by manufacturers.

CONCLUSION

For all digital data supports, the market forces have put the emphasis only onto capacity and speed of reading/writing. From this standpoint, ODDs seem already to have lost the battle, given the spectacular progress registered in other storage solutions; they no longer present a specific clear advantage in terms of price per unit. Their best chance of survival could therefore lie in a market niche reorientation: to become **the** long term archiving solution. With archival planning framed in decades and centuries, this would generate a new market associated to growing needs in the world.

Technically, in the long term, we have no particular reason to be pessimistic: whether the solution comes from one side or the other, there should emerge a stable device with excellent long-term properties for digital data storage. Indeed, no insurmountable technical hurdle to be crossed appears; the condition is that we must simply deploy sufficient efforts to attain the 3 or 4 century target (as for ordinary paper).

The best candidate apparently for this purpose remains the optical digital disk, for several reasons. In principle, these disks can be very simple physical objects, where all the physical processes involved can reasonably be well understood and controlled. Compared to hard disks, which include both the storing media as well as the recording and reading devices (magnetic heads) in a single device, they are much less complicated and allow a complete separation of functions. There is no special reason why a recordable digital disk with a lifetime of several centuries should not be implemented and introduced to the market-place within a few years, provided the manufacturers realise the needs of Society and can identify the existence of the corresponding market.