Holographic Memory (Holomem)
Introduction & Tutorial

by
Dr. Richard G. Zech[1]

My CeBIT 2001[2] Report contains a substantial analysis of holographic memories. The motivation for this effort was a claim by a German startup Optostor AG[3] that it had essentially invented the "perpetual motion machine" of data storage. The company provided a read-only demonstration of its technology at CeBIT 2001, which, of course, came no where near matching its claims for storage density and capacity.

Although I had not been expecting an actual operating holomem at CeBIT 2001, I was pleased to see it. I spoke at some length with Optostor's technology director to learn as much as possible. My intent was to give them a full and positive write up. However, I quickly realized he knew little about the technology, and, worse, the numbers he was quoting were impossible to obtain. I shared my concerns, and showed him how to calculate the correct numbers. I must have made an impression, because, Optostor reduced its claim for storage density (although the proposed new capacity still could not be achieved for other than a "write-only" holomem).

After researching the matter carefully, I concluded that as research scientist and engineer who helped pioneer this field, I had duty to speak up. Rather than simply condemning what I thought might be a fraud, I decided to simply present the facts carefully analyzed, proceeded by an introduction and tutorial to holomems. The reader can then decide. The bottom line is a detailed exposition of the basics of holomems, supported by specific data provided by the excellent Holomem research of IBM.

---

[1] Dr. Zech is a 40-year optical storage veteran. His dissertation was on data storage in volume holograms.

[2] CeBIT (Hannover, Germany) is the world's largest computer and communications trade show and conference.

[3] Optostor did not survive the start up phase.

Focus on Computer Storage Products and Technology

The CD/DVD Perspective

The CeBIT 2001 Executive Report PLUS

(June 2001)

The Hannover Messe Fairgrounds

Home to CeBIT

The enormous size of the CeBIT Fairgrounds is shown by this areal photograph. The distance from Hall 1 to Hall 10 (north gate at lower left to south gate at upper right) is more than one kilometer. More than three US COMDEX/Fall Shows (largest US IT exposition) would easily fit within the 28 Halls of the Fairgrounds.

CeBIT 2001
Table of Contents

Part 4
MO, WORM, & Holographic Optical Storage

4.1 Introduction & Summary

Part 4 of the CeBIT 2001 Report is a synopsis of the latest developments in 2”, 3½”, and 5¼” MO, 12” (phase change WORM), and 3-D holographic optical storage (crystal and photopolymer WORM).

Very few companies exhibited “high-performance” optical storage products of any kind. Fujitsu overwhelmingly dominates the 3½” MO drive market, Maxoptix and Sony are the only surviving 5¼” MO drive manufacturers, and Plasmon is the sole source for 12” WORM (ATG appears to have exited the business). Sony did not exhibit its 5¼” MO drive, but its major OEM customer, Hewlett Packard, did at its main stand in Hall 1. Neither Sony nor Maxoptix exhibited their advanced 5¼” MO UDO (Ultra Density Optical) or OSD (Optical Super Density) drive.

Fewer than 20 system integrators and media suppliers exhibited solutions based on these types of products. All surviving European MO optical disk library (ODL) manufacturers (primarily, DSM and Plasmon) were at CeBIT 2001. Grundig (K+S Systeme, the best MO product integrator ever to exhibit at CeBIT) has apparently left the business.

The failure of TeraStor’s near-field recording (NFR) 5¼” MO product development effort came as no surprise to most OEMs and product integrators, whose CeBIT 2000 consensus was that TeraStor’s products were vaporware. The future of 5¼” MO, in this group’s view, is now in the hands of Sony and Maxoptix. The markets and applications for 5¼" MO are relatively small, but the margins for application-specific solutions are very attractive. Hence, these companies believe that either a successful UDO or OSD product development will help maintain and grow the archiving, tape-alternative, and near-line storage markets. A comparison of UDO and OSD drive and media specifications is given in Table 4.1-1.

Table 4.1-1

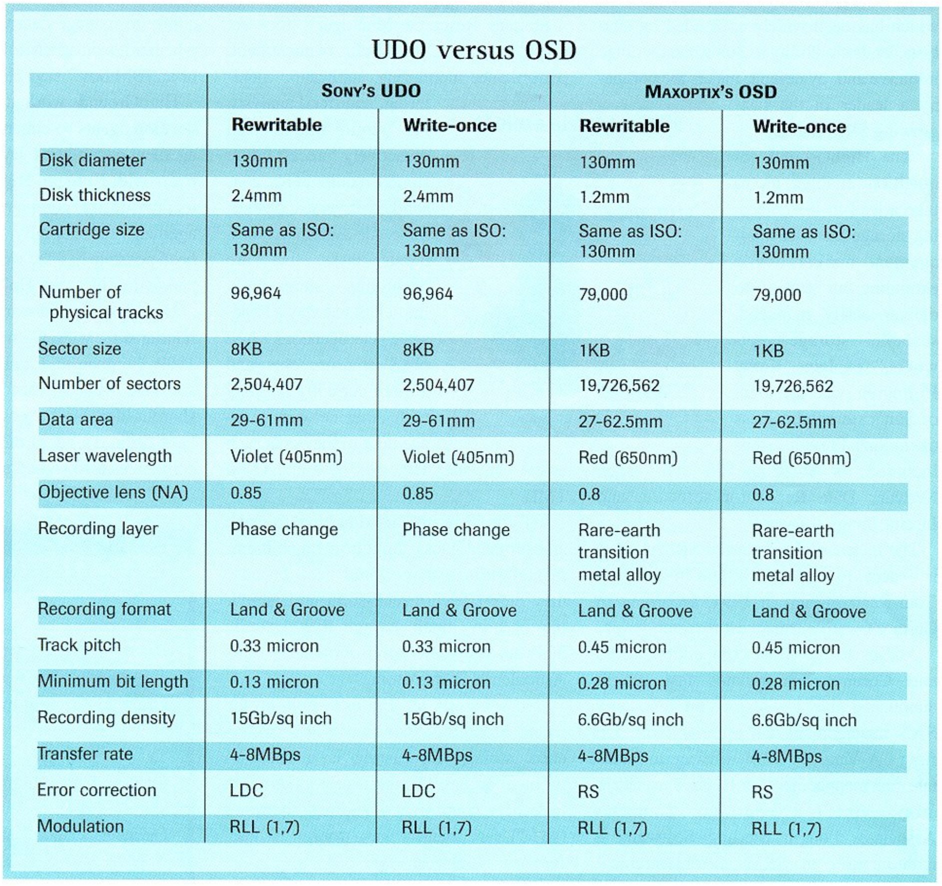

UDO versus OSD

	Sony's UDO		**Maxoptix's OSD**	
	Rewritable	**Write-once**	**Rewritable**	**Write-once**
Disk diameter	130mm	130mm	130mm	130mm
Disk thickness	2.4mm	2.4mm	1.2mm	1.2mm
Cartridge size	Same as ISO: 130mm	Same as ISO: 130mm	Same as ISO: 130mm	Same as ISO: 130mm
Number of physical tracks	96,964	96,964	79,000	79,000
Sector size	8KB	8KB	1KB	1KB
Number of sectors	2,504,407	2,504,407	19,726,562	19,726,562
Data area	29-61mm	29-61mm	27-62.5mm	27-62.5mm
Laser wavelength	Violet (405nm)	Violet (405nm)	Red (650nm)	Red (650nm)
Objective lens (NA)	0.85	0.85	0.8	0.8
Recording layer	Phase change	Phase change	Rare-earth transition metal alloy	Rare-earth transition metal alloy
Recording format	Land & Groove	Land & Groove	Land & Groove	Land & Groove
Track pitch	0.33 micron	0.33 micron	0.45 micron	0.45 micron
Minimum bit length	0.13 micron	0.13 micron	0.28 micron	0.28 micron
Recording density	15Gb/sq inch	15Gb/sq inch	6.6Gb/sq inch	6.6Gb/sq inch
Transfer rate	4-8MBps	4-8MBps	4-8MBps	4-8MBps
Error correction	LDC	LDC	RS	RS
Modulation	RLL (1,7)	RLL (1,7)	RLL (1,7)	RLL (1,7)

Source: *Infostor*, January 2001, p.15

New for this year is a section dedicated to 3-D holomems. Despite its long and mainly unsuccessful development history, 3-D holomems still generate excitement. A small company called Optostor AG actually demonstrated a read-only system. It also made some claims that may be too good to be true. In the interest of truth in advertising, we provide a detailed overview of 3-D holomems, and a reality check for Optostor's planned storage system. We report, you decide.

4.2 3-D Holographic Memories

Almost every year for the past 5 years, the promise of some type of 3-D holographic memory (holomem) system has hung ripe over CeBIT. The greater the demand for high-performance storage, the greater the level of interest in 3-D holomems. Unfortunately, 3-D holomems have generally turned out to be the *hype du jour*. Still, hope springs eternal, and those seeking something better than the old standby magnetic disk array or tape farm persist in reacting to every new prospect. Those with a need to do long-term archiving are especially interested.

It was therefore with great expectations that we visited the stand of Eutelis Consulting GmbH, where we met with the technical director of **Optostor AG** and observed a demonstration of a prototype 3-D holomem. Optostor claims to have a real product under development with some very interesting specifications. In this section of the CeBIT 2001 report, we will analyze the company's claim, and give our opinions about the viability of the company's design.

First, however, some basic concepts and insights about this esoteric storage technology are required, if our analysis of Optostor is to be understood. Some may find this material difficult, but it has been simplified as much as possible. As compensation, some of the real secrets of 3-D holomems will be revealed for the first time.

4.2.1 3-D Holographic Storage Background

The hologram was invented by 1971 physics Nobel laureate Dr. Dennis Gabor in 1948 at Imperial College (England) to solve an aberration problem in electron microscope images. The technology was made practical with the availability of the He-Ne gas laser and the invention of the off-axis (Fresnel) hologram by Emmett N. Leith and Uris Upatnieks at the University of Michigan in 1962. Finally, the concept of the 3-D holomem was proposed by P. J. van Heerden, a research scientist at Polaroid, in 1963.

Van Heerden showed that a theoretical maximum volume storage density of $\rho_V \sim 1/\lambda^3$ was possible for binary data (for $\lambda = 500$nm, this calculates to 8 bits/μm^3, or 131 Tb/in^3). The corresponding *theoretical* areal density ρ_A is ~ 8N bits/μm^2 (~ 5.2N Gb/in^2). N is (or can be interpreted as) the number of holograms that can be independently *stacked* in a common volume of the storage medium. For N = 1,000, the areal storage density is ~ 5.2 Tb/in^2. This storage density was so astounding a possibility in 1963 that industrial and government laboratories throughout the world launched a multi-billion US$ R&D effort to implement storage systems that exploited van Heerden's theory. Thus began the 38-year quest for a commercial 3-D holomem.

From about 1963 through the mid-1970s, 2-D and 3-D holomems were viewed as universal storage solutions. Interest in holomems declined after this period because of the extremely difficult engineering challenges to systems implementation. The hologram storage medium was (and largely remains) the major item on the critical path. A few startups tried to develop either disk (notably, Tamarack Systems) or tape 3-D holomems in the late 1980s, but failed. Finally, interest was rekindled in the early 1990s by significant improvements in key components. From about 1994 to the present, companies such as Holoplex, IBM, Lucent Bell Labs and SIROS (Optitek) began serious attempts to commercialize 3-D holomems (so far, only Holoplex has succeeded).

A major force behind the effort to implement real 3-D holomems over the last 5 years of the 20th century was the DARPA-sponsored HDSS (**H**olographic **D**ata **S**torage **S**ystem)/PRISM (**P**hoto**R**efractive **I**nformation **S**torage **M**edia) initiative. The goals of this project included demonstrating a 1 Tb (128 GB) capacity and a 1 Gbps data rate. The program was largely successful (in fact, a 10 Gbps data rate demonstration was claimed), and the project was completed in November 2000. However, no commercial or military systems resulted. SIROS has exited the 3-D holomem part of its business.

IBM **(www.research.ibm.com/thinkresearch)** is viewed by many as the 3-D holomem technology leader, but it does not appear to be ready to announce a product any time soon. Lucent Bell Labs has also done some very good R&D. Lucent, Imation, and private investors recently (January 2001) formed InPhase Technologies **(www.inphase-technologies.com)**, located in Longmont, CO to develop an holographic disk storage system. A shipping product is probably at least 2-3 years away. Rockwell is addressing real-time correlators. Holoplex (funded by Hamamatsu) has actually shipped a few systems; it is focused on fingerprint identification applications (a type of correlator). Some research continues in major

Japanese companies, for example, Hitachi, NHK, Pioneer, and Sony. None of these companies has shown anything at a trade show like CeBIT.

4.2.2 Basic 3-D Holomem Theory

A hologram is a true amplitude and phase representation of some object or data pattern. It is created with coherent laser light by interfering the light transmitted by or reflected from a data pattern (object or signal beam) with an unmodulated carrier (reference beam). The interference pattern created by the interaction of the two beams is captured by a storage medium, such as a photopolymer film or an electro-optical crystal (these are referred to as volume phase storage media, and are the focus of our interest here). For certain recording setups and storage medium thicknesses, a 3-D (volume) hologram is recorded, which has useful properties for data storage. The hologram is read out (reconstructed) by illuminating the hologram with only the reference beam. In principal, a perfect 3-D image of the original object is reproduced. Holograms have been used for things as practical as credit card and CD authenticators or interferometric profilometry to imposing full-color works of art.

A 3-D holomem is a collection of 3-D holograms, which share a common storage medium volume (the hologram array), and the means to write and read these holograms. 3-D holograms are stored throughout a part of the volume of the storage medium defined by an area large enough to ensure adequate resolution of the input device pixels (see SLM below). This permits multiple holograms to be recorded in a common volume and individually retrieved by, for example, (Bragg) angle or wavelength multiplexing. This collection of holograms is called a stack, even though all of the spatially varying information that represents the holograms is randomly intermingled and cannot normally be individually distinguished or erased. The areal density is defined as the number of stacks x data per page divided by the stack area. Two, or more, stacks form an hologram array. Block (cube), disk, tape and card formats have been evaluated, and each has its own set of advantages and trade offs. Add required hardware and software, and one has a 3-D holomem.

The concept of a 3-D holomem can be best understood by reference to and a review of the schematic diagram of Figure 4.2-1. Think of the data page input (spatial light modulator) as being imaged one-to-one onto the data page output (photodetector array). Writing of the hologram in the storage medium delays the completion of the imaging and detection processes, until hologram reconstruction takes place. Important details will now be discussed with this frame of reference.

The 3-D holograms each store a page of data. The page is created by an input device called a spatial light modulator (SLM). The SLM typically is an array of N_p pixels (for example, 1024 x 1024), which can thought of as on-off switches that convert electrons to photons. Generally, the Fourier transform of the SLM is projected by a lens into the volume of the storage medium. Fourier transform holograms maximize storage density. A reference beam is added, and the resulting 3-D interference pattern is captured (frozen) in the storage medium. This process is repeated for different reference beam angles until a stack of N holograms is formed in a common storage medium volume (a process called incoherent adding or superposition).

Using only the reference beam (now the called reconstruction beam, which must be a very accurate and precise replica of the original reference beam), the individual holograms can be read out and imaged (inverse Fourier transformed) onto a photodetector array (PDA). The amount of power in the reconstructed image of the n^{th} of N holograms divided by the power in the reconstruction beam incident on the hologram (P_H) is called the “diffraction efficiency” (η_n). Generally, “write scheduling” is used to make η_n about the same for all N holograms in a stack. Write scheduling compensates for partial erasure of prior-written holograms in real-time media such as lithium niobate and for nonlinear characteristics in write-once media, such as photopolymers. The PDA is the output device that converts photons back to electrons. Note that $\eta_n P_H$ divided by the average number of pixels in the holographic image ($½N_p$) is the average power per reconstructed pixel (P_D) incident on a photodiode of the PDA. The corresponding energy per reconstructed pixel (E_D) is $P_D\tau_i$, where τ_i is the PDA integration time. E_D ,together with signal-dependent and random noise sources, ultimately determines raw (also, soft or uncorrected) bit error rate (BER), read data rate, and achievable user capacity.

One last detail remains to be considered. What type of 3-D hologram page addressing method is to be employed: (a) a “moving parts” design, which is less expensive, but limits access speed significantly and may be less reliable, or (b) a “no moving parts” (all-photonic) design, which is much more expensive, but is very fast and much more reliable? The analogy to modern hybrid (E/O/E) versus all-optical network switching would not be too far afield. This is a very complex issue with many factors to be analyzed. The choice is generally application driven. It is but one illustration of the trade offs that make it difficult to obtain all the attractive features of a 3-D holomem simultaneously. Fortunately, we can assume here either is appropriate, and not worry about the details.

From an end user/applications perspective, good reasons for a high level of interest in 3-D holomems existed in the 1960s, and still do today. Some of the most important (in a theoretical and idealized sense) are:

- Very high storage densities (from 50 GB/in^2 up to many Tb/in^2)
- Very high capacities (from 200 GB disks up to many TB crystals).
- Very high data rates (from 50 MBps to many GBps; this results from the high degree of write/read parallelism inherent to 3-D holomems).
- Very fast page access times (from many msecs down to < 10 μsec).
- Very low cost per TB (media cost/TB, for example, would be much less than magnetic tape for most designs).
- Very high reliability (assuming no moving parts and a completely photonic implementation).
- Very stable long-term storage (probably true only for certain write-once storage media).

With all of these exciting (potential) characteristics, one might reasonably ask why no commercial 3-D holomem systems for data storage (particularly, archiving) have ever come to market. Some of the basic reasons are listed below (although their real complexity cannot be examined or appreciated here):

- Absence of a viable storage medium. No photonic analog to magnetic media exists. The characteristics, requirements and specifications for 3-D holomem storage media are so complex that this component deserves a white paper in its own right. Lithium niobate, an electro-optical or “photorefractive” crystal is the best known, but has serious practical limitations. The best hope of the technology currently resides with certain types of “write-once/read many” (WORM) photopolymer

(PP) storage media. InPhase and Polaroid-licensee Aprilis have done excellent research in this area. Complutense University (Spain) and the National Research Council of Canada recently announced a new photopolymer medium that has favorable properties for both 3-D holomems and fiber optic gratings. The key requirements that PP storage media must have are high diffraction efficiency, relatively high write speed, no inherent or write/read-dependent noise, and no dimensional variations due to writing processes or temperature fluctuations. PP media approximately meet these basic criteria.

- Key component availability is limited. These include the laser (writing and reading device), spatial light modulator (input device), photodetector array (output device), and page addressing mechanism or subsystem. These components do exist, but are (a) often very expensive, (b) designed and engineered mainly for laboratory use, and (c) lacking reliability characteristics suitable for commercial data storage devices (especially, long-term archiving).
- A 3-D holomem is a complex, analog, interferometric device. It is far more difficult to engineer and employ commercially than, for example, a near-field recording (NFR) optical disk drive, which (thanks to Terastor) is well known to be extremely challenging. Very sophisticated servo systems will be required to compensate and control thermal variations (athermalization), vibration, and shock that affect everything from laser power to storage medium alignment. The latter is especially problematic, because the positioning accuracy and precision at the reconstruction beam-stack area interface are often measured in microradians and nanometers for ultra-high density 3-D holomems.
- Manufacturing cost is very high for the 3-D holomem drive. Creating a *pro forma* bill of materials (BOM) from a detailed system block diagram and applying reasonable component cost estimates will quickly convince the reader that a 1 TB holomem (block or disk format) will cost more than ten times a 1 TB magnetic disk drive. Of course, the cost/TB of the storage medium, which is removable, is likely to be very reasonable in the tradition of optical data storage.
- No photonic infrastructure or second sourcing exists for either the drive or storage medium. One must accept the fact that support for a large-volume, low-cost holomem storage product line does not exist today, despite the explosive growth of optical networking. Nor are there standards. In the early days of the storage industry, this did not matter

very much. Today, it will take a very significant alteration of the price/performance/reliability metric to induce market acceptance over a time interval short enough to ensure the cash flow-viability of a start up in this arena.

A few subtle details about 3-D holomem performance and operation need also be shared:

- The capacity (C) of a 3-D holomem is upper bounded by the well-known equation:

  $$C = N_H N_p = N N_L N_p$$

  where N_H is the total number of hologram, N_p is the number of pixels per page, and N_L is the number of hologram storage sites (each with a stack height = N). The value of C will always be less than the number given by the above equation, and often much less. User capacity will be determined by the acceptable level of corrected bit error rate at the system level. As in the case of optical disc drives, this implies the need for a relatively low raw BER (for example, 10^{-4}). In turn, a minimum level of signal-to-noise ratio (SNR) is required. This helps explain why we stated earlier that E_D is such an important variable.
- A 3-D holomem is operationally a write-once (WORM) system, even if the storage medium itself is erasable. This is because individual holograms cannot be selectively erased (some exceptions exist, but are still in the research stage or are impractical). In fact, for certain applications (for example, medical records or movie film archiving), a true and stable write-once storage medium is preferable. Unless a very stable erasable storage medium with a fast erase/reset mechanism and very high speed writing sensitivity can be developed, write-once media is a compelling choice.
- A significant asymmetry exists between write and read data rates. This is mainly an issue of storage media write sensitivity. Electro-optical crystals require high power densities and relatively long exposure times to record holograms. Photopolymers, which have an inherent chemical gain mechanism, are orders of magnitude faster than electro-optical crystals. Recent advances in photopolymer storage media design have shown the potential to reduced the read/write data rate asymmetry significantly.

- Optimizing performance requires a careful write scheduling of holograms (this means that the write time of a hologram is a function of its position or order in the stack). Current practice is make the diffraction efficiency of every hologram the same. Recently, the concept of making the raw bit error rates similar for each hologram evolved. Regardless, the implication is that “stack-at-once” (to borrow a term from CD-R) writing is required, or some method must be devised to control all the variables in real time during reconstruction of the holograms (extremely complex). Hence, as a practical matter, one should consider a 3-D holomem as having a CD-ROM operational model without the replication (3-D hologram stacks cannot be copied). Dedicated writers located in well-controlled environment would fill hologram stacks in serial fashion (mastering), and dedicated (lower-cost) readers would be used to read the data. This is a valid operational model for many applications, particularly archiving. Obviously, a 3-D holomem using a disc format must anticipate a better operational model; the details of how this might be done are unknown.
- Key performance parameters are coupled in non-obvious ways. For example, true page access time (t_{acc}) is the sum of the page address time (a function of the page addressing mechanism) plus page read time. Page read time is inversely proportional to the (burst) read data rate. The read data rate (R) is proportional to diffraction efficiency (through the PDA integration time τ_i), which is inversely proportional to the square of the stack height, N. Finally, capacity is directly proportional to N. For a block-oriented 3-D holomem, $R = \text{constant}/C^2$, a non-obvious result.

For those left more confused than enlightened by the preceding discussion, be advised that we will be publishing a detailed report on 3-D holomem technology (optical, x-ray, and electron beam) and products later this year. For the present, please consult Appendix C/Part 1 (*Special Appendix 4)* for general reference materials and the CeBIT Glossary following Part 5.

4.2.3 Analysis of the Optostor 3-D Holomem

We now turn our attention to Optostor AG. The company provided the first real trade show demonstration of a 3-D holographic memory system (HMS) prototype. See Figure 4.2-2 for an example of its hologram array.

Although the system was operating in the read-only mode, the reading of bit pages was clearly demonstrated. The following background information will help in understanding Optostor's concept:

- The storage medium (doped lithium niobate) was developed at the University of Cologne by Dr. Theo Woik, et. al. The system is being developed by startup Optostor AG, a "32%-subsidiary" of Eutelis Consulting GmbH. Optostor has been developing the system for about two years. This has produced some 20 patents applications.
- The prototype uses Bragg-angle multiplexing to achieve hologram stacking (incoherent superposition of holograms, each separated by a minimum of one Bragg angle).
- Fourier transform holograms are recorded, which have a defined minimum area requirement and maximize areal storage density.
- The laser is a frequency-doubled YAG system emitting at 532nm. Power output, beam quality, and coherence length are unknown. However, commercial devices are relatively mature, although expensive. Beam quality suitable for holographic recording with CW power in excess of 1 W can be obtained
- The input device (SLM) is a Sony liquid crystal display (LCD) with 1024 x 1024 (1K x 1K) pixels (picture elements). The LCD uses switchable input cells with two reflective or transmissive states to create binary 0s and 1s. Each pixel can be thought of as a mark or space; coding determines the number of bits each represents (similar to pulse position modulation used in early optical disk drives).
- The photodetector array is a CCD (charge-coupled device); type and frame rate, or fps, were unspecified. Both CCD and CMOS PDAs are under consideration.
- The storage medium is doped-lithium niobate, a well-known type of real-time, volume phase hologram recording material. It is either iron (Fe) doped or iron-silicon (Fe-Si) alloy doped (our German may be failing us here). More importantly, this storage material can be coated in high-quality sheets at least 3mm thick. This is an unusual format for data storage in lithium niobate, which is generally used in a block format. Lithium tantalate has also been mentioned as a potential storage medium. Since lithium niobate self erases during hologram readout, some type of fixing is needed to ensure long-term stability.

The means for fixing the storage medium is unknown (heating has been mentioned as a means of stabilizing the holograms). Optostor says that dimensional and environmental stability of its lithium niobate are very good.

- The hologram array has dimensions of 30 x 30 x 3 mm^3. The 30 x 30 mm^2 area is subdivided into 100 3 x 3 mm^2 sub-areas (essentially contiguous cubes). This provides 100 hologram storage locations (stack sites). The total number of holograms per hologram array must therefore be (8 x 1 TB)/$(1024)^2$ = 7.63 million. This means the stack height must be 76,294 holograms. The computed storage density is ~ 5.73 Tb/in^2. Our baseline knowledge of lithium niobate suggested that this would likely result in a **"write only"** memory. We shared this insight with Optostor during our meeting, but did not received a satisfactory explanation.
- Basic technical specifications given by Optostor were: 1 TB system capacity; 10 msec page access time; and ~ 100 Mbps read data rate. A 100-year archival life (> 500 years calculated) is claimed. The write data rate was not given, nor was any reliability data provided. Note that page size in bits divided by page access time, that is, $(1024)^2$/0.01 sec is ~ 100 Mbps. This means that the 10 msec number is really the page read time, and 100 Mbps is the *burst read data rate*. Since no page addressing mechanism was identified, the time to access successive pages is unknown (5-10 msec is possible with a servo-controlled galvo mirror scanner), nor is average page positioning time.
- Optostor's CEO said on a BBC TV segment covering CeBIT 2001 that system capacity would grow to 5 TB in 2-3 years and to 100 TB in 5-10 years. He also said that initial pricing for a 1 TB system would be DM 200.000,- (< US$ 100,000). Storage media modules are expected to cost DM 1.000,- (< US$ 500). Optostor plans to deliver 5 beta units some time in 2001 and to start production in late 2002. Initial customers include banks, government health agencies, insurance companies, and video/film producers. The "killer application" is apparently archiving (European laws are much stricter about archiving medical, financial and government records than are those of the US or most Asian countries. A German law that came into effect in 2001 requires that medical records be preserved for 30 years).
- After CeBIT 2001, we learned that Optostor moved up to a 2 TB capacity 3-D holomem with a data rate of 12 MBps (about 100 Mbps,

as before). This time, the numbers work out much better, but still dramatically push or exceed engineering limits. The hologram array size is now 50 x 50 x 3 mm^3 and the total number of holograms is 15.3 million. With the same 1K x 1K SLM as before, this does indeed calculate to 2 TB capacity. Assuming 2mm x 2mm holograms (minimum size for double-Rayleigh resolved Fourier transform holograms with a F/2 lens), 625 hologram locations (stack sites) are required. To obtain 15.3 million holograms requires a stack height of 24,414 holograms, which is an improvement over 76,294, but is still a very large value for N. The storage density is now *only* 4.1 Tb/in^2.

We now provide some "reality test" technical analysis. First, refer to Table 4.2-1. This table provides a summary of data reported by IBM for its DEMON II 3-D holomem test system (very sophisticated and proven; the end result of the company's participation in the HDSS/PRISM project). The test bed uses a 90-degree write/read geometry, which is the best case for Bragg-angle multiplexing. The specific reference is **J. Ashley, et al, Holographic Data Storage, *IBM J. Res. Develop.*, Vol 44, No 3, May 2000, pp. 341-368**. This excellent paper is probably the most complete reference available today on the realities of 3-D holomem design. Second, refer to Table 4.2.-2 (in two parts), which functions as a *pro forma* Optostor 3-D holomem spec sheet (with best case assumptions). Note that we have chosen a raw BER of 10^{-4} as the key analytical metric. This is a reasonable choice, given the anticipated applications.

There are so many unknowns that some assumptions and design choices must be made on Optostor's behalf. The first has to do with write/read geometry. The 90-degree model cannot be used, because the hologram stacks are distributed over a 50 x 50 mm^2 hologram array. We chose reference and signal beam angles which ensure a minimum of 50% beam overlap (it is 100% for the 90-degree geometry) in the stack (a consideration forced by the 2 x 3 stack aspect ratio). The "effective" hologram thickness is less than 3mm, but we will ignore this. A 2-D angular addressing scheme is assumed, because the Bragg angle for a 3mm thick hologram is too large to permit 1-D angular addressing. These, and many other technical details are summarized in Table 4.2-2.

Some of the numbers in the Table 4.2-2 will certainly raise red flags for experts in the field. The bottom line is this: *Optostor has over specified its*

3-D holomem. Required system margins to ensure both short- and long-term reliability at customer locations are absent. This means that the 2 TB capacity cannot be realized for a raw bit error rate of 10^{-4}. The energy per pixel E_D received at the photodiodes of the PDA is equivalent to ~ 43 photons (below the noise floor of most PDAs), when 2,000 are needed. The signal-to-noise ratio (SNR) is less than 1. For archiving applications, the raw bit error rate 10^{-4} should be an end of life specification, so the situation is actually worse than first appears. Scaling back N to be consistent with IBM's experimental data yields a capacity of ~ 0.1 TB, not 2 TB. If Optostor's lithium niobate is remarkably better than IBM's' and if better design parameters are possible, capacity might be increased to 0.2-0.3 TB (all other things equal). If, however, the end of life raw bit error rate must be < 10^{-4}, then achievable user capacity will again be decreased. Finally, a rough bill of materials analysis, using current component costs, suggests that a DM 200.000,- (< US$ 100,000) HMS selling price is not realistic for a lithium niobate-based 3-D holomem. Storage media pricing seems on target, if not even a little too high.

Should Optostor be taken seriously? This is a difficult question to answer. Its 3-D holomem concept appears flawed on first principles. On the other hand, the company does claim to have filed for 20 patents, some of which could teach a new and better way to implement 3-D holomems (we suspect, however, that most of these filings are for the lithium niobate storage medium). Optostor would not be the first company to believe that the storage medium makes the storage system. We doubt that its lithium niobate is radically better than other variants (no data on write speed, maximum index change, scattering noise, fixing time, etc. are available to make a better judgment). Going from laboratory (research) prototype to data center (commercial) system will not be easy, even if the application is restricted to archiving.

In our opinion, based on the known, public domain facts, Optostor appears to have rediscovered a road well-traveled by other companies before it, all of which were turned back by very challenging engineering problems. Moreover, it seems unlikely that a small research group at the University of Cologne and a 2-year old startup have accomplished what the DARPA HDSS/PRISM project (comprised of some of the best US companies and universities in the field) could not in 5 years for a cost of over US$ 32 million. But, given Germany's reputation for excellence in chemistry, we could be wrong. Also, it is unsettling that a company spawns a media frenzy before it appears to understand fully the basics of a complex data storage technology. Only three years ago, another European company (Opticom; Oslo, Norway) claimed it could store 70 TB on a credit card-size storage medium comprising a polymer-protein matrix. After a flurry of publicity, Opticom delivered nothing.

We recommend reserving judgment on Optostor's product development for another year, but also withholding orders or investments until a lot more capability is actually demonstrated by Optostor. Most importantly, it must be determined whether or not Optostor's lithium niobate storage medium is exceptional and really capable of the writing and long-term preservation of archival records.

We wish Optostor well. After 38 years of research and investment, a real 3-D holomem commercial success would be greatly welcomed.

Table 4.2-1
IBM DEMON II 3-D Holomem Tester Data

specification	Units	size
max. volume density	bits/μm^3	400
max. areal density	bits/mm^2	5.5 x 10^8
max. areal density	Tb/in^2	0.36
number of holograms (data pages)/stack, N		1,350 (for 10^{-4} raw BER)
hologram stacking method		Bragg-angle multiplexing, 90 deg setup (best case)
page access mechanism		mechanical (galvo scanner, +/- 15 deg range)
Coding		8 bit/12 pixel modulation
raw bit error rate (BER)		10^{-4} (measured)
corrected bit error rate		10^{-12} (Reed-Solomon EDAC)
storage material		Fe-doped lithium niobate
fixing method		not specified
insertion loss		not specified
hologram type		Fourier transform
FT Lens focal length	mm	30
hologram area	mm^2	1.6 x 1.6 (~ 28% larger than minimum)
storage medium thickness	mm	5.5
laser type		frequency-doubled YAG with λ = 532nm
spatial light modulator		1024 x 1024 pixel, reflective (IBM Yorktown; 12.8 μm pixel pitch)
photodetector array		1024 x 1024 element CCD (operating at 41 fps; 12 μm photodiode pitch)
photodetector threshold energy E_0 **+**	J	~ 7.5 x 10^{-16} (~ 2,000 photons for 10^{-4} raw BER; 1 photon yields 1 electron assumption)

The symbol “**+**“ means we calculated or assumed reasonable values for this variable or parameter.

Table 4.2-2 (Part 1)
Optostor AG 3-D Holographic Memory System

specification	units	size
Capacity	TB	2 (total, not user)
write data rate	MBps	not specified
page read time	msec	10
page address time	msec	unknown
page access time	msec	unknown (defined as the sum of page read and page address times).
read data rate	MBps	~ 12.5 (burst, not sustained)
areal density **+**	Tb/in^2	~ 4.1 (versus theoretical max. = 12.6)
areal density **+**	bits/μm^2	6,640 (versus theoretical max. = 20,000; IBM has achieved only ~ 400)
volume density **+**	bits/μm^3	~ 2.13 (versus theoretical max. = 6.64)
system reliability (MTBF)	power on hours	not specified
hologram storage life	years	100
system cost	DM	200.000,- (< US$ 100,000)
media cost	DM	1.000,- (< US$ 500)
availability		beta units – late 2001 (?) first production – late 2002 (?)
storage medium		doped lithium niobate (n ~2.3 assumed)
fixing method		not specified
insertion loss		not specified
storage medium area	mm^2	50 x 50
storage medium thickness	mm	3
storage medium packaging		unknown
laser type		frequency-doubled (SHG)YAG with λ = 532nm (power, beam quality and coherence not specified)
spatial light modulator		1024 x 1024 pixel LCD (Sony)
photodetector array		1024 x 1024 element CCD
hologram stacking method		Bragg-angle multiplexing
hologram type		Fourier transform
minimum hologram size **+** (surface area)	mm^2	~ 2 x 2 (double-Rayleigh resolution criterion for FT lens with F# = 2)
hologram size used **+**	mm^2	~ 2 x 2 (based on a stated laser beam diameter of 2mm)
hologram thickness	mm	3

The symbol “**+**“ means we calculated or assumed reasonable values for this variable or parameter.

Table 4.2.-2 (Part 2)

Optostor AG 3-D Holomem System

specification	units	size
total number of holograms (data pages) **+**		~ 15.3 million
number of stacks (hologram storage locations) **+**		625
number of holograms (data pages)/stack, N **+**		24,410 (a maximum of ~ 1,350 for 10^{-4} raw BER per IBM)
page access mechanism		not specified (appears to be mechanical)
data encoding		not specified
error detection and correction coding		not specified (Reed-Solomon typical)
external signal/reference beam angles **+**	deg	-46.7/+46.7 (internal angles of –18.4/+18.4; provides good beam overlap in the crystal)
minimum Bragg angle **+**	radians/degrees	1.27×10^{-4}/0.0072
hologram angular separation **+**	radians/degrees	2.54×10^{-4}/0.0144 (2x min. Bragg angle to prevent crosstalk)
min. angular access range (1-D, θ only) **+**	radians/degrees	> $2\pi/360$ (not physically possible)
min. angular access range (symmetrical 2-D, θ–ϕ) **+**	radians/degrees	> 0.04/2.24 (157 θ x 157 ϕ holograms)
read power, P_H **+**	W	1 (at hologram stack)
max. diffraction efficiency (N = 1, optimum write) **+**	%	20 (post fixing, including insertion loss; 100 for ideal volume phase change media)
$1/N^2$ bias buildup loss (due to hologram stacking) **+**		$1/(24{,}414)^2 = 1.68 \times 10^{-9}$ (signal strength loss compared to N = 1)
PDA frame read rate **+**	fps	100 (required for 100 Mbps read rate)
CCD integration time, τ_i **+**	msec	10
imputed IBM photodiode energy, E_0 **+**	J	~ 7.5×10^{-16} (2,000 photons for 10^{-4} raw BER per IBM; 1 photon yields 1 electron assumption)
power (per pixel) at photodetector, P_D **+**	W	~ 1.6×10^{-15}
energy (per pixel) at photodetector, E_D **+**	J	~ 1.6×10^{-17}
E_D/E_0 ratio **+**		~ 0.02 (must be =/> 1)

The symbol "**+**" means we calculated or assumed reasonable values for this variable or parameter.

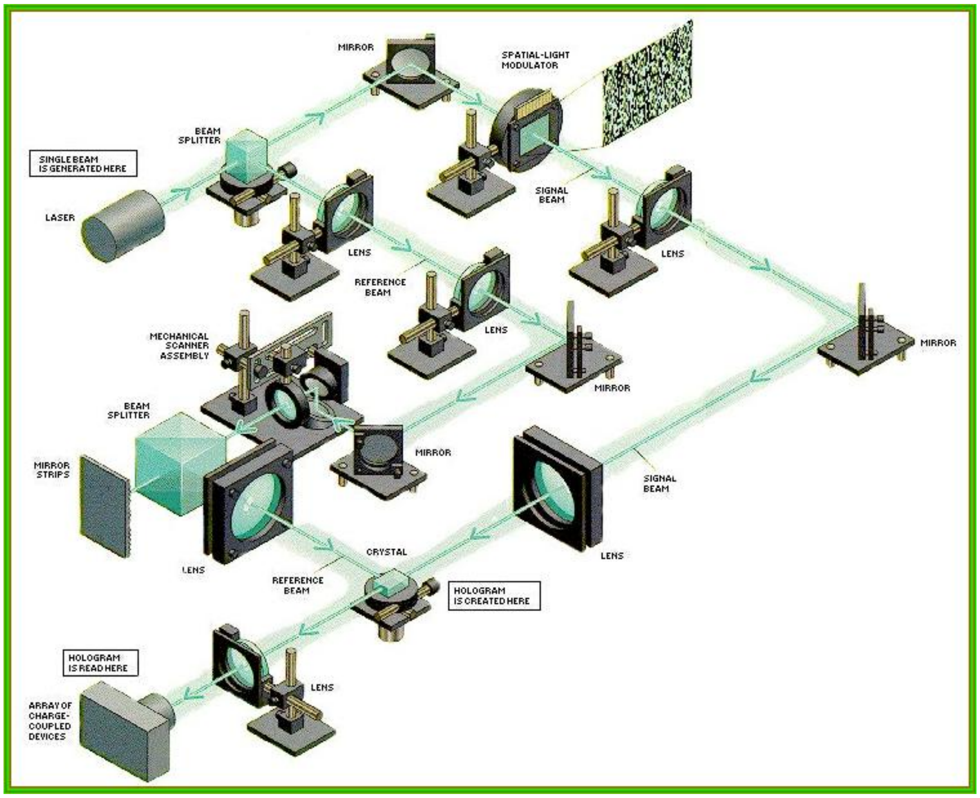

Figure 4.2-1 – Block diagram for a typical 3-D holomem. This architecture is called BORAM (block-oriented random access memory). Bragg-angle multiplexing is used to stack the holograms. A mechanical scanner (moving parts) page addressing method is shown. This excellent illustration comes from D. Psaltis and F. Mok, Holographic Memories, *Scientific American*, November 1995, Vol. 273, No. 5, pp. 70-76.

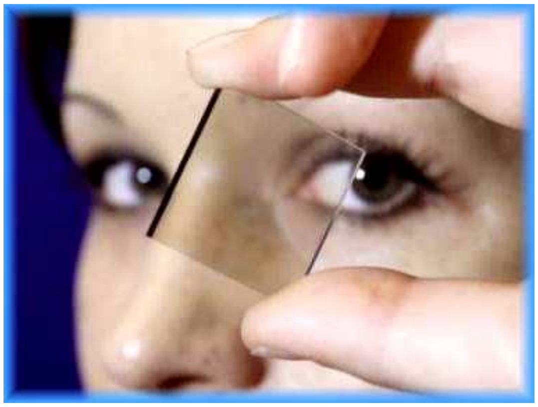

Figure 4.2-2 – Optostor lithium niobate 3-D holomem storage medium example. The dimensions are 30 x 30 x 3 mm^3, which is an unusual format for an electro-optical crystal (more like a disk, than a block). We cannot tell if any data is actually stored.