To appear in *International Trends in Optics*, 2002

Optical data storage

Hans Coufal and Geoffrey W. Burr

IBM Almaden Research Center
650 Harry Road
San Jose, California 95120

Tel: (408) 927–2441, 927–1512
Fax: (408) 927–2100
E-mail: *coufal, burr@almaden.ibm.com*

1 Introduction

The distribution of audio information in disk format—sound stored as modulations in a flat surface—has had a long and successful track record. At first, analog signals were imprinted on shellac and then later vinyl platters. But recently, sound has been almost exclusively distributed as and replayed from digitally encoded data, molded into the surface of compact polymer disks. This migration was driven by consumer demand for higher fidelity, i.e., higher bandwidth in recording, replication and playback, which in turn required dramatic increases in the amount of stored information. At first the diameter of the disk was increased, but this soon reached its practical limit. With a mechanical stylus, increasing the areal storage density increased the susceptibility to wear and tear, which forced a transition from mechanical detection to a non–contact optical scheme.

Optical detection retrieves the stored data by sensing changes in the intensity or polarization of a reflected laser beam. In the form of the read-only compact disk, this data storage medium became the dominant vehicle for music distribution (CD) and later for computer software (CD-ROM) [1, 2]. With the ferocious appetite of consumers for ever more information at ever higher data rates, the CD–ROM is currently undergoing a metamorphosis into the Digital Versatile Disk (DVD) [3]. The DVD standard offers higher areal density per layer, and as many as four layers of prerecorded information, offering sufficient readout bandwidth and capacity for distribution of several hours worth of high quality compressed video. Table I compares the pertinent characteristics of the CD and DVD formats, highlighting the significant progress that has been made. Since its introduction into the market, DVD devices have reportedly experienced the fastest growth of quarterly sales and market penetration of any consumer electronics technology, as of the last quarter of 2001.

In addition to this distribution of prerecorded content on removable and interchangeable media, the desire of consumers to record information by and for themselves has been an important economic and technical driver. The wax coated drum of Edison, with analog information engraved with a vibrating needle, has given way to the CD 'burner,' with carefully modulated and interleaved digital data imprinted with a flashing laser beam. Recordable optical disks rely on using high laser power to write marks that low laser power can still read. The high laser power must locally modify the optical properties of the recording medium to provide a significant change of signal; conversely, the low readout power must *not* affect the properties of the media. Processes such as ablation of a dye layer are irreversible and thus are well suited for "write-once, read many" times (WORM) applications. The prevalent implementation of this technique is the CD–R: WORM media with a form factor and functionality nearly identical to a CD or CD-ROM once it has been recorded. Other recording processes—foremost the change between amorphous and crystalline structure in a thin metal, semiconductor or alloy film—can be reversible and allow erasable or ReWritable (R/W) media. For these types of media, the thermal characteristics are as important as the optical parameters. One has to ensure that after short transient heating with the tightly focussed laser, each very small spot will cool fast enough to rapidly quench the film into an amorphous phase. On the other hand, the same R/W medium must cool slowly after a large area is uniformly illuminated, allowing recrystallization of the film and erasure of previously recorded data. The issues of optical, thermal and device engineering for WORM and R/W recording are extremely complex and go beyond the scope of this brief summary. Details can be found, for example, in Reference [2].

In addition to the pervasive 120mm diameter CD/DVD disk, many other disk diameters have been used. One early niche application was the laser disk, used for the distribution of high quality prerecorded movies before DVD–level areal densities and data compression were available. For high–end data storage applications, diameters of 5.25", 12" and up to 14" are commonly used to increase the capacity per platter up to 200GB. Smaller disks for small form factor, portable consumer electronics devices are becoming increasingly popular. Examples are MP3 players, digital still and digital video cameras using optical disk media with diameters as small as 1".

The fact that the optical recording media are removable, have high storage capacity and a relatively low price also make them a prime candidate for massive data warehouses. Jukeboxes for home audio and video applications allow a home user to have all the multimedia content that they own at their fingertips. In a similar fashion, commercial data libraries and warehouses use data silos with robotics and industrial strength CD-R recorders and CD-ROM players to provide massive amounts of near-line storage at low cost.

Optical readout of cheaply replicated, injection molded CDs, CD-ROMs, and the rapidly emerging DVD disks has clearly come to dominate multi–media content and software distribution. On the other hand, general–purpose read–write information storage is still the exclusive domain of magnetic recording. Magnetic recording in the form of hard disk drives delivers higher areal densities, higher data rates and faster access times than optical storage, while supporting millions of write/erase cycles and many years of retention time at very competitive cost. To obtain these high performance characteristics, the removability offered by the once ubiquitous floppy disk had to be sacrificed. Thus the key advantages of optical recording are clearly the removability and interchangeability of the media, as well as the parallel replication of pre–recorded disks in seconds via injection molding. Optical recording on WORM media, typically in the form of "burning" a CD-R, plays an important role as inexpensive backup for data stored on magnetic disks, and is starting to compete with magnetic tapes for archival data storage in large libraries. All attempts to bring a commercially viable optical tape technology to market have failed up until now [4].

In 1997, the National Storage Industry Consortium (NSIC) and the Optoelectronics Industry Development Association (OIDA) convened a meeting of 50 leading industry experts to develop a comprehensive Optical Disk Storage Roadmap [5]. This roadmap summarized the anticipated component, system and market development. While the overall trends were predicted correctly, technical and economic challenges considerably slowed down the actual pace of development.

	CD	DVD
Capacity	0.65 GB	4.7 GB (single layer)
Laser wavelength	780 nm	650 nm
Numerical Aperture	0.45	0.6
Track Density	16,000/inch	34,000/inch
Minimum length pit ($\sim$2 bits)	0.833–0.972 μm	0.4–0.44 μm
Areal density (user)	0.39 Gb/sq.in.	2.77 Gb/sq.in.
Reference velocity	1.2 m/s	3.49 - 3.84 m/s
Data rate	1.47 Mb/s	11.08 Mb/s

Table 1: Comparison between CD and DVD [5–7]

Key drivers for future developments were expected to be the rapid expansion of the Internet and the emergence of HDTV with their requirements for higher bandwidth and storage capacity. The Internet has expanded, but high–speed broadband service to the home is not yet pervasive. Standards issues and broadcaster resistance have pushed back the advent of HDTV. At the same time, the pervasive use of computers and storage in consumer electronics devices (and eventually common household appliances) is requiring more storage in smaller form factors. Thus all of these requirements are still driving optical data storage to reach for ever higher storage densities [5].

The remainder of this article highlights different approaches towards achieving this goal of higher density optical storage, either by increasing the density at the surface of an optical recording material, or by utilizing the unique capability of optics to access the volume of suitable media.

2 Approaches to increased areal density

For many years, the demands for better data storage have been met by evolutionary advances: steady increases in the areal density and other performance specifications of magnetic and optical recording devices. The areal density of hard disks has been growing continuously through more than seven orders of magnitude over the last fifty years, but during the last five years the growth accelerated to a compounded annual growth rate of slightly more than 100%! Optical storage, with a much shorter track record, has increased in storage density by a factor of 5 from the original CD standard to the recent DVD format (Table I). The reason for this discrepancy is removability. Since high performance magnetic recording does not support removability, hard drives need not work with the media of previous generations and the only standards to satisfy are those on data input/output and form factor, thus encouraging strong technical competition that has resulted in tremendous progress. Optical storage, however, is dominated by interchangeable media and backward compatibility. This compatibility facilitates the introduction of each new generation of technology in the market, but forces a time–consuming standards process for each higher density generation.

Still, optical storage has substantial potential as a storage technology. While the limits of magnetic recording are still being debated [8], the limits of conventional optical recording are well understood [5]. Current optical storage technology is already working close to the optical diffraction limit. However, significant future increases in density are possible by taking advantage of the wavelength and/or numerical aperture dependence of the diffraction limit, or by going beyond it with near field techniques. In addition, if the signal/noise ratio is sufficient then

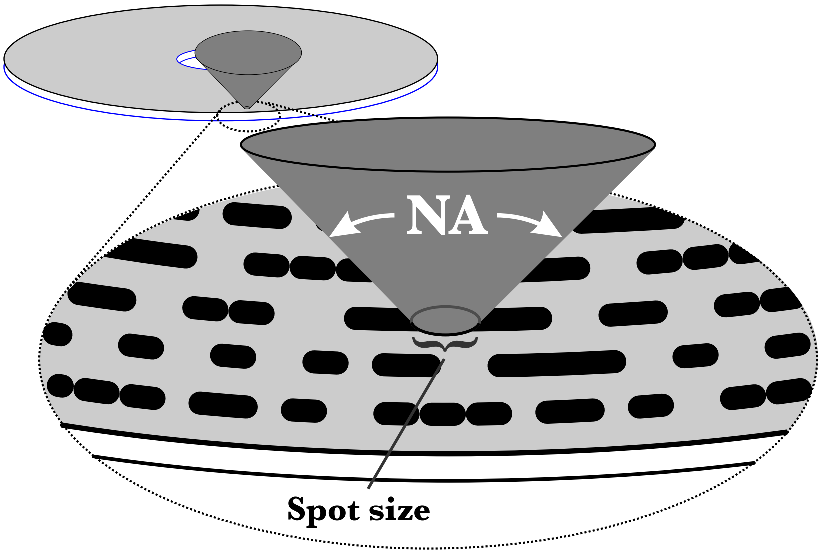

Figure 1: Conventional optical storage uses a tightly focussed laser beam to access individual bits in a single layer.

---

gray–scale techniques allow the storage of more than one bit per location [9].

The diameter of the diffraction limited spot is directly proportional to the laser wavelength λ and inversely proportional to the numerical aperture (NA) of the imaging lens (Figure 1). The area of the spot is then proportional to the square of these parameters [10]

$$A \sim \left(\frac{\lambda}{\mathrm{NA}}\right)^2, \tag{1}$$

and the resulting maximum areal density is simply the inverse of this area times the number of bits per spot b,

$$\mathcal{D} \sim b \left(\frac{\mathrm{NA}}{\lambda}\right)^2. \tag{2}$$

As described in Table I, the differences in capacity and data rate between the CD and DVD formats is a clear consequence of reducing the diffraction–limited spot size of the focus at the medium. Other factors were also involved in the density improvement found in DVDs—such as stronger modulation coding, signal processing, error correction, and more aggressive tolerancing—but it is unclear how much more improvement could be extracted from these areas in future optical disk standards.

2.1 Short wavelength lasers

Early optical storage products used infrared wavelengths of 830 nm down to 780nm, simply because these were the only diode lasers available with the reliability, optical power, quantities and cost that the industry required for a consumer product. As the development from CD to DVD highlights, the obvious path to further increases in storage density has been the development of a suitable short wavelength laser. While blue lasers have been used in laboratories around the world for high density recording studies, they have yet to show up in consumers products. Instead of leapfrogging from 800 nm to 400 nm lasers (using GaN) and thus immediately quadrupling

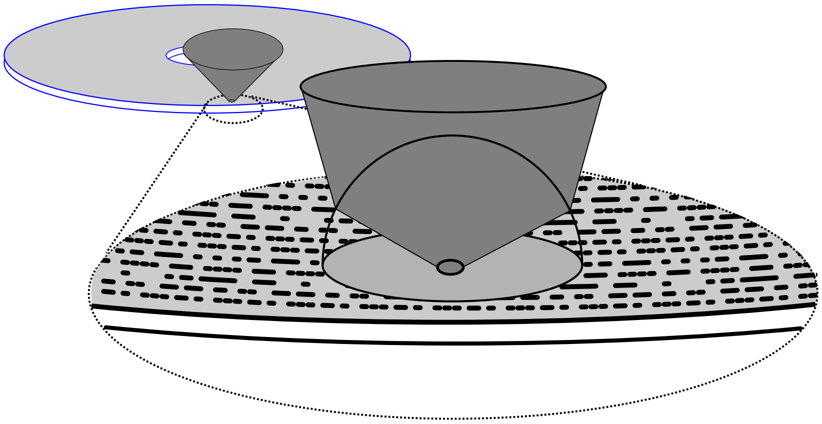

Figure 2: A solid–immersion lens (SIL) can increase the effective NA beyond 1.0, further increasing density but requiring evanescent coupling between the SIL and disk.

---

the storage density, the development has been gradual. 4X and 8X MO drives employ visible wavelengths of 650–685 nm, while DVDs rely on 635–650 nm lasers [5]. Continued advances to even shorter wavelength have been made with novel diode laser materials and concepts. Progress with prototypes of green ZnSe lasers at 490 nm has been substantial and blue GaN lasers with wavelength as short as 375 nm have been explored [5]. The NSIC/OIDA roadmap predicted the general availability of optical storage products with these short wavelength lasers around 2002 [5]. That this has not occurred reflects both technical difficulties in fabricating cheap, long–lifetime blue lasers as well as economic reality: HDTV and broadband Internet, the applications that were to drive the need for this technology, are penetrating the market much more slowly than expected. However, other promising evolutionary approaches might provide significant improvements in areal storage density while the performance, lifetime, and cost of blue laser diodes slowly improve.

2.2 Increased numerical aperture

Instead of reducing the laser wavelength, a corresponding increase in numerical aperture can achieve the same increase in areal density. The increase in density when making the transition from CD to DVD was partially due to an increase in NA. There are, however, practical limits to the increase in NA, such as manufacturing tolerances for diffraction limited optics. The key issue here is depth of focus, which is directly proportional to the wavelength, but inversely proportional to the square of the numerical aperture [10]:

$$\delta \sim \frac{\lambda}{\mathrm{NA}^2}. \tag{3}$$

To maximize density, the data layer of the media should be as close as possible to the focus of a diffraction–limited readout beam. Inevitably, manufacturing tolerances on the drive and media make a fast focusing servo essential. For ultimate performance, even dual stage servos with multiple lenses are being considered. Obviously, these issues become more difficult as the optical depth of focus decreases. One aspect that does improve with smaller depth of focus is the thickness of the cover layer over the recorded data (which takes up some of the working distance

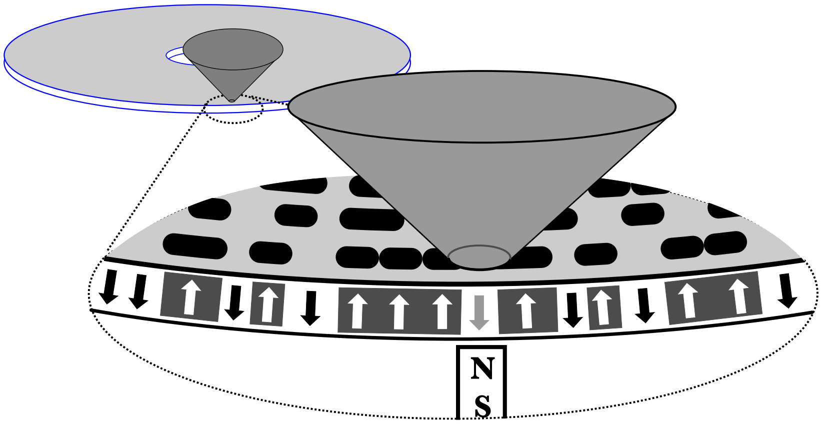

Figure 3: Magneto–optical storage uses a tightly focussed laser beam to record bits, using local heating of a magnetic film in the presence of an external magnetic field.

of the optics and forces correction for spherical aberration). Smaller depth of focus would allow this layer to be thinner without affecting the tolerance of the system to scratches and other defects on the out–of–focus disk surface.

Although the maximum NA in air is 1.0, microscopists have long improved their resolution beyond this by using oil–immersion objective lenses. Rather than coating optical disks with oil, however, numerical apertures greater than 1 can similarly be attained by solid–immersion lenses brought into close proximity to the disk media [11], as shown in Figure 2. If the gap between the bottom of the lens and the recording media is smaller than the light wavelength, then evanescent coupling of light across this gap can result in a subwavelength diameter recording spot on the media. Similar to the read–write head in magnetic recording, it is possible to fly a solid immersion lens on an air–bearing over the media [11]. This approach to high storage density, while pursued agressively for several years, has yet to deliver a successful commercial storage product. In addition, because of the susceptibility of this tiny air gap to external contaminants on the disk, such near–field approaches would seemingly have to sacrifice media removability: one of the very advantages that optical storage has historically had to offer.

2.3 Magnetic super–resolution

In magneto-optic recording, another interesting effect can be utilized to increase the resolution. In this technology, bits are stored by heating a magnetic film locally with a laser in the presence of an externally applied magnetic field (Figure 3). In the portions of the film that exceed the Curie temperature, the local magnetization will orient with the applied field. This magnetization change is then sensed optically by detecting the slight rotation of the polarization due to the Faraday effect [10]. Here, thermal diffusion processes rather than optics can dominate the resolution. A short exposure with a laser beam of non–uniform spatial profile, such as a simple Gaussian beam, can produce a temperature distribution that exceeds the Curie temperature only for a brief time, and only in the very center of the beam. Subsequent heat diffusion away from the spot quickly brings the temperature below the Curie point, resulting in magnetic reversal in an area of sub-optical-wavelength dimensions, including marks as small as 100 nm [12].

However, such marks must be still be optically detected. With dozens of such marks within the

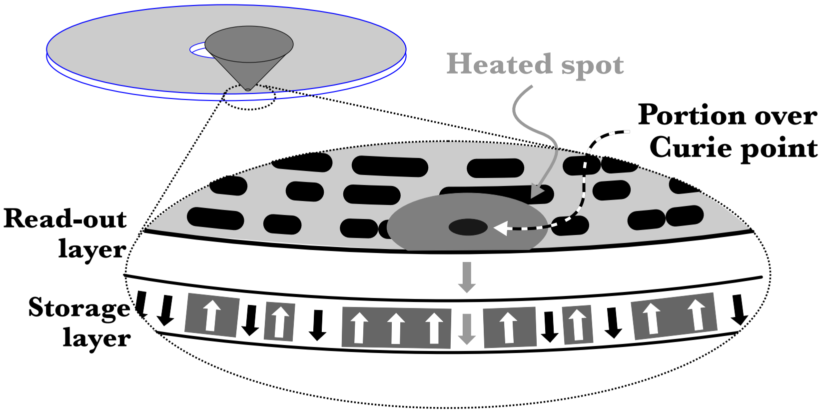

Figure 4: Magnetic Super–resolution uses two layers, with the heat from the laser spot transferring just the bits of interest from the bottom storage layer into the top read–out layer.

---

diameter of each focussed spot, deconvolving each mark individually from the aggregate optical signal would be quite challenging. An elegant technique called Magnetic Super-Resolution (MSR) can overcome this readout problem using the highly non–linear behavior of thermal diffusion [12]. The MSR technique adds a magneto-optic read-out layer above the magnetic storage layer. For readout of a particular sub–wavelength spot, this top layer is briefly heated with a laser just as in the recording process described above: the power density in the center of the beam barely exceeds the Curie temperature. As above, the magnetization of the hot spot orients with the applied field, which in this case is the field due to the recorded bit in the lower layer. The surrounding area of the upper layer does not reach the Curie point and is not affected. After the magnetization at that location has been "revealed," the read-out layer contains only the sub–resolution spot with reversed magnetization surrounded by a magnetically homogeneous background (Figure 4). This sub–resolution spot can then be optically detected with sufficient signal/noise ratio, despite the much larger diameter of the focussed laser beam.

Taking advantage of similar magnetic multilayer techniques that have been recently introduced to increase the areal density of conventional magnetic recording, further increases in optical storage density are possible by introducing additional layers for magnetic amplification (Magnetically Amplified Magneto-Optic System or MAMMOS [13].

3 Volumetric optical recording

Both magnetic and conventional optical data storage technologies, where individual bits are stored as distinct magnetic or optical changes on the surface of a recording medium, are approaching physical limits beyond which individual bits may be too small or too difficult to store and retrieve. Storing information throughout the volume of a medium—not just on its surface—offers an intriguing high-capacity alternative

Three dimensional optical storage literally opens up another dimension to increase the capacity of storage media. In principle, a volume element with the dimensions of the wavelength should suffice to store one bit. The volumetric density that can be achieved in the diffraction limit would then scale with $1/\lambda^3$. Tradeoffs between density and data rate make it possible, at least in principle, to forgo some of this density to obtain blistering data rates. Some of the techniques that have been proposed do not require mechanical motion, enabling access times in the range of tens of microseconds, albeit to a small storage volume with moderate capacity.

3.1 Volumetric addressing techniques

Several approaches for volumetric optical data storage have been explored. These can be distinguished by the method used to address the stored data. Some of the techniques simply extend the multilayer nature of conventional optical storage already begun with the 2–layer DVD standard, while others take advantage of the 3–D character of optics. The techniques described here are

- Adjusting focus to access data on a particular layer;
- Using an interferometer (sensing differences in path length) to address a layer;
- Addressing a particular point, line, or plane in a medium by intersecting two laser beams;
- Using material with extremely narrow spectral sensitivities to address data; and
- Addressing data by the spacing and direction of interference fringes, known as “K–vector” or holographic addressing;

3.2 Addressing by depth of focus

Depth of focus, a phenomenon familiar to anyone that has used a microscope, is an already integral part of the current DVD standard. To increase the capacity per disk over the single–layer limit of 4.7GB/disk [14], the standard can double this capacity with two data layers, or quadruple it with a two–sided, two layer–per–side option (the disk must be flipped to access data with a single objective).

Using depth of focus has the clear advantage of being similar to existing CD and DVD technology, and it could possibly be compatible with the large installed base of optical storage devices. As shown in Figure 5, the key disadvantage is that the individual layers must be widely separated to avoid crosstalk errors on readout, or tracking problems with the focus servo. However, the required diffraction–limited performance is difficult to accomplish over a sizable depth range, limiting the maximum number of layers.

One way to improve this is to use fluorescence for readout: since the readout light is incoherent, the effects of interference and scattered light can be reduced leading to higher signal–to–noise ratio (SNR). These reduced noise effects trade off with the omnidirectional output of the fluoresence from the illuminated voxel (volumetric pixel). However, two–photon fluorescence leads to significant signal only from focussed volume and permits higher resolution, leading to better SNR and increased volumetric density.

Multilayer Fluorescent Disks (MFD) employ media with many (10 to possibly as many as 100) layers [15]. Each layer contains pits or grooves filled with fluorescent material, which absorbs the incident laser light and re-emits incoherent red–shifted fluorescent light. Portions of the disk manufacturing process are similar to CD/DVD fabrication: a mastering process (here for each individual layer) followed by replication via hot embossing or stamping. The pits are filled with fluorescent dyes—which require appropriate absorption/emission spectra, high conversion efficiency, short time response, and high saturation levels—and then the layer

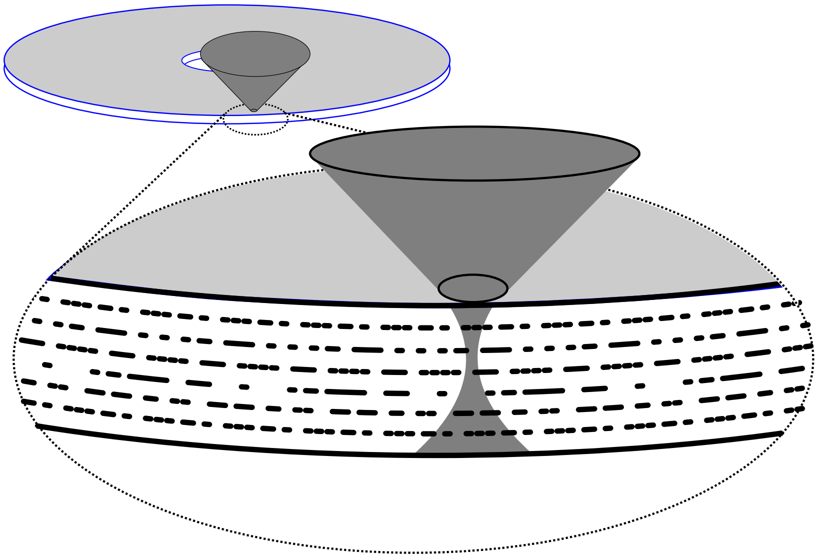

Figure 5: The effective optical areal density can be increased by using depth of focus to access multiple bit layers within a disk.

substrates are bonded together. Such ROM media fluorescent dyes are under development, as well as WORM media based on thermal bleaching and multiphoton processes for red and green light [15]. Reversible photochromic material have been demonstrated in the lab.

MFD drives are similar to CD/DVD drives with differences in the optical head: dichroic mirrors and filters to separate the fluorescent signal light, optics for spherical aberration correction, and a more sensitive data detector. Tracking and focusing techniques are similar to other optical disks, with additional electronics and possibly data encoding needed to find and maintain focus on the different layers. A ROM system with more than twenty layers was demonstrated at CD density (650 nm wavelength semiconductor laser, 1 mW in the reading spot, 3 spot tracking, astigmatic focusing) [15]. The SNR for the experimental disk was better than 30 dB (1.5 MHz bandwidth) [15].

In addition, optical cards with 10 layer media with 1 micron sized marks have been demonstrated. CCD readout was used, image processing was applied to read data. Plans are in place to introduce MFD–ROM systems with 10 to 20 layers at DVD density (approximately 50 to 100 GByte per disk), and optical ROM cards with 1 GByte capacity, followed later by WORM versions of these devices [15]. As with CDs and DVDs, the advent of blue lasers would be expected to support an increase in capacity by a factor of four, although this would require the dye chemistry to be re–engineered to these wavelengths.

3.3 2–photon absorption for addressing of a bit cell

Another means of limiting the interaction volume of the addressing beam in the medium is by using 2–photon absorption, which scales with the square of the beam intensity. This effectively shortens and narrows the focus volume to where the intensity is sufficient for writing. These approaches often require femtosecond lasers with high peak intensities, but could enable media with hundreds of layers allowing capacities approaching 1 Terabyte per disk. Here the media is initially homogeneous (not layered), with two–photon effects used to either initiate or quench

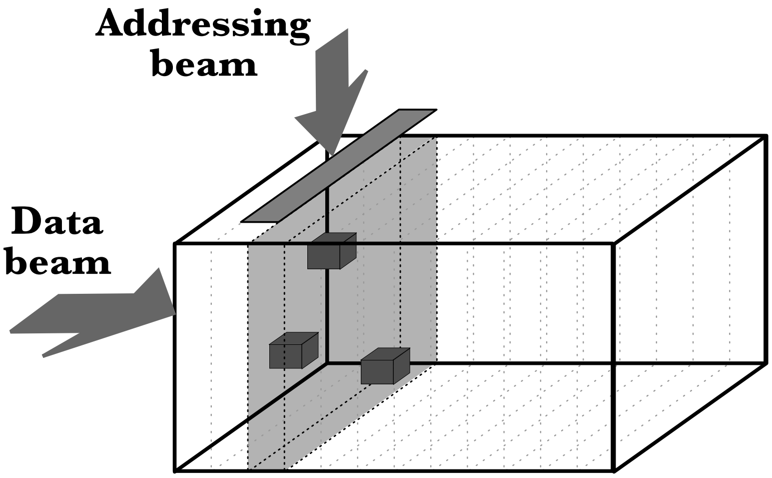

Figure 6: Orthogonal beams can write and read data in parallel into a three–dimensional volume, using 2–photon fluorescence at the intersection of the beams (After Reference [18]).

---

fluorescence at each voxel [16, 17]. Confocal microscopy has been used to detect extremely small marks in initial studies, but it is unclear whether such an approach would work on a rapidly spinning disk.

Further gains in depth resolution can be realized by crossing two laser beams of the same or different wavelength in a volume element using 2–photon absorption to initiate a change of its optical properties only in the commonly illuminated volume [18, 19], as shown in Figure 6. Here just one volume element within a larger piece of recording media can be selected for recording or readout without affecting adjacent volume elements, thus avoiding inter–track and inter–layer interference. Extensions of this technique use columns or sheets of light to address a large number of volume elements at the same time, although these extensions make diffraction–limited performance for sub–micron sized marks more difficult to obtain.

Even so, effective areal densities greater than 100 Gb/in^2 would appear achievable by using many layers with relatively large marks (i.e., greater than 1 micrometer). Once written, the marks (or columns or pages) may be nondestructively read by the 2–photon fluorescence. Demonstrated capabilities include the recording and reading of media with more than 100 data layers, recording tracks of 2x2 μm^2 data marks, and the construction of several proof-of-principle portable readout systems [18–20].

This 2–photon 3D optical storage technology seeks to provide disk drive systems with high capacity (100–500 GB/disk) and high data rates (1–10 Gb/s) using inexpensive, easily manufactured, and long-lived plastic media. The drive technologies are potentially backward compatible with conventional DVD media. The media will be removable and offers the possibility of wide wavelength and mechanical tolerances in drives for near-line/on-line servers.

However, efficient 2–photon absorption recording requires moderate average laser power (50 mW) but high peak intensity at the focused spot (on the order of GW/cm^2), necessitating the use of a pulsed laser source for recording. Suitable diode pumped solid-state pulsed lasers are available, albeit at a significant price and size disadvantage as compared to semiconductor laser diodes. These laser requirements remain the largest risk to this approach.

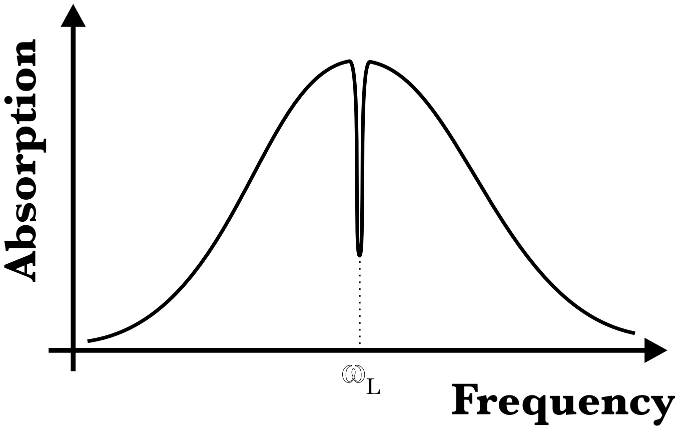

Figure 7: After illumination with a narrow-band laser at ω_L, a persistent and very narrow spectral hole has been "burned" into the absorption spectrum.

3.4 Interferometry

Taking advantage of a relatively short coherence length and variable path difference between two laser beams, one can obtain constructive interference somewhere along the path of a collimated laser beam to address volume elements throughout the depth of the media [21]. While similar to the conventional depth of focus approach, it is different enough to be not easily compatible with CD or DVD. It has the advantage of diffraction limited performance throughout thick media, and that it is similar to optical coherence domain reflectometry, widely explored as a method for imaging through thick biological samples [22]. While simple in concept, it does require interferometric stability and has not been explored in detail for data storage.

3.5 Persistent spectral hole burning

Spectral hole burning uses media that can support dopants with extremely selective spectral response (Figure 7). To access this narrow spectral response requires a tunable narrowband laser, which tends to be a complex and expensive light source. Thus high storage densities are essential to amortize the cost of this light source over a large storage capacity. Unfortunately, these very high densities are currently only possible at cryogenic temperatures adding additional complexity and cost. However, the latency and the data transfer rates that could be achieved by this technique are superior, allowing spectral hole burning to play in the memory as well as in the storage arena.

Several groups have been pursuing variations of spectral hole burning for data storage and memory applications [23, 24]. PSHB can alternatively be addressed using time-domain interference techniques, in which holograms laid down in frequency spectrum allow time sequences to be replayed as "photon echoes" [25, 26]. Both time– and spectral–access alternatives support spatial multiplexing across numerous locations, and would support the prospect of combining this technique with the spatial holographic techniques to be described below. Crosstalk and distortions between spectral channels, and photon flux are typically the important limiting factors. Time-domain access relieves the requirement for laser tunability, although the laser source would still need to be narrow band and extremely stable, and also enables content addressability (accessing data by the correlation or similarity with search data). However, efficient address schemes are yet to be developed and random access to data is difficult to provide, with ”First-in-first-out” (FIFO) being the standard.

Tunable solid state narrow band lasers are currently being developed both for these PSHB applications as well as for telecommunications applications. Depending on the choice of PSHB medium, the required tuning range spans either 10 GHz (0.25 Angstrom at about 750 nm), 200 GHz (15 Angstrom at 1500 nm) or 2 THz (30 Angstrom at 650 nm). None of these tuning ranges is extremely difficult— although mode hopping can be problematic—but opto-mechanical approaches would sacrifice one of the key features of spectrally sensitive storage: rapid random access.

3.6 Holographic storage

Holographic data storage is a volumetric approach which, although conceived decades ago, has made recent progress toward practicality with the appearance of lower-cost enabling technologies, significant results from long-standing research efforts, and progress in holographic recording materials [27,28]. In holographic data storage, an entire page of information is stored as an optical interference pattern by intersecting two coherent laser beams within a thick, photosensitive optical material, as shown in Figure 8. The object beam contains the information to be stored; the reference beam is designed to be simple to reproduce for readout. The resulting optical interference pattern is stored as a modulation of the absorption, refractive index or thickness of the photosensitive medium. When this hologram is illuminated with one of the two optical waves that was used during recording, some of the incident light is diffracted by the stored grating to reconstruct a weak copy of the other wavefront. Illuminating the stored grating with the reference wave reconstructs the data–bearing object wave, and vice versa. Interestingly, a backward–propagating or phase–conjugate reference wave, illuminating the stored grating from the “back” side, reconstructs an object wave that also propagates backward toward its original source.

A large number of these interference gratings or patterns can be superimposed in the same thick piece of media and be accessed independently, as long as they are distinguishable by the direction or the spacing of the gratings. Such separation can be accomplished by changing the angle between the object and reference wave or by changing the laser wavelength. Any particular data page can then be read out independently by illuminating the stored gratings with the reference wave that was used to store that page. The theoretical limits for the storage density of this technique are around tens of terabits per cubic centimeter [27, 28].

In addition to high storage density, holographic data storage promises fast access times, because laser beams can potentially be moved rapidly without inertia, unlike the actuators in disk drives. With the inherent parallelism of its page–wise storage and retrieval, a very large compound data rate could be reached by having a large number of relatively slow (and therefore cheap) parallel channels.

A rather unique feature of holographic data storage is associative retrieval, or content–addressable data storage [29]. If a partial or search data pattern is imprinted on the object beam, illuminating the stored holograms reconstructs all of the reference beams, each weighted by the similarity between the search pattern and the content of their particular data page. By determining, for example, which reference beam has the highest intensity, the closest match to the search pattern can be found without initially knowing its address.

Despite of all of these advantages and capabilities, holographic storage has been—for nearly four decades—an intriguing but elusive alternative to conventional data storage techniques. With

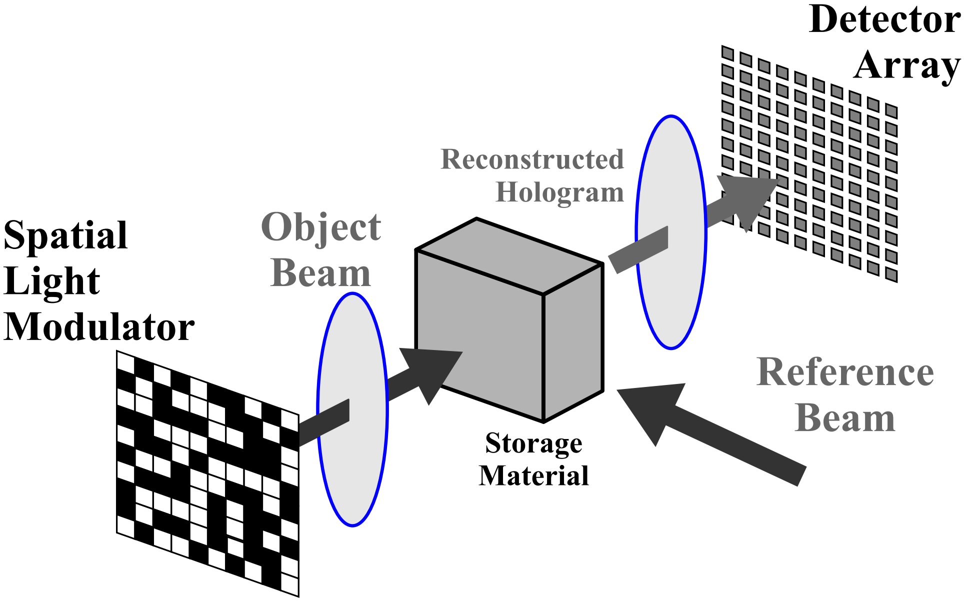

Figure 8: In holographic storage, a data page imprinted onto an optical beam is stored by recording optical interference patterns in the storage material. Illuminating this hologram with the same reference beam reconstructs the object beam, which is then detected by a pixellated detector array.

---

the recent availability of relatively low-cost components, such as liquid crystal displays as input devices and camera chips from electronic cameras and camcorders for detector arrays, interest in practical holographic storage devices was rekindled. Recently, largely as part of an effort that was sponsored by the United States Defense Advanced Research Projects Agency (DARPA), members of the Holographic Data Storage Systems (HDSS) Consortium were able to demonstrate impressive performance with experimental platforms: at the IBM Almaden Research a data density of 250 Gb/in^2 was achieved [30], at Stanford University a data rate of 10 Gb/sec was reached for the read-out of holographic data [31], and a team at the Rockwell Science Center demonstrated access times on the order of 10 microseconds [32]. Each of these demonstrations were the product of extensive studies of recording physics, systems tradeoffs, signal processing, and coding techniques. These experiments also showed that it will be difficult, but not impossible, to obtain a combination of all three desirable performance characteristics in one hardware platform.

Holographic storage research has branched into three related disciplines: the content– addressable storage described above, read–write fast–access holographic memories designed to exploit the non–mechanical access, and write–once holographic disk storage driving at high effective areal density (to attain high–capacity with the same form factor disks as CDs and DVDs, for instance).

Several small commercial efforts have been launched in the last 18 months to aggressively pursue write–once holographic data storage. Buoyed by years of improvements to write–once media such as photopolymers and photochromic direct–write media, these efforts are now moving to solve the systems and drive issues involved in creating a holographic "3–D disk." These issues include tracking and servoing, signal processing and coding, and obtaining sufficient laser and media performance to be able to write and read holograms with a single laser pulse. Additional media issues would include improving characteristics such as dynamic range, scatter, sensitivity, shelf–life before and after recording, and thermal expansion properties. The real challenge is that the obvious application areas (low–cost data archiving, possible next–generation distribution format for data and multimedia) call for cheap media (which may well be possible) and inexpensive and robust disk readers (which, given the engineering issues, may not).

Less active is the area of read–write holographic data storage, mainly because of the much slower pace of materials improvement. In order to have sufficient capacity accessed without mechanical motion, it is necessary to have thick read–write material which can support numerous holograms per spot. While photorefractive single crystal materials have been known for years, their liabilities have changed little. Since the interference fringes are stored by spatially redistributing electronic charge between traps, holograms stored in these materials will erase if these traps are re–excited, either by thermal processes or by the simple act of illuminating the hologram for readout. And if the absorption coefficient is reduced to slow this process then the associated recording speed suffers.

Non-destructive read-out in such photorefractive crystals would require a process that renders the stored information semi–permanent. Such fixing processes have been demonstrated by applying an electric field (domain reversal) or by annealing with heat (ionic charge compensation). But the required high voltages, ovens, or high-power lasers are unfeasible from a system–design point of view, where "stacks" of superimposed holograms would need to be fixed in–situ without any affect on neighboring stacks [28]. Such fixing techniques would also make it much more difficult to verify information immediately after it has been written, as is typically done with magnetic and optical recording today. Other proposed methods of avoiding hologram erasure are materials solutions that unfortunately create more systems problems, such as readout at a different wavelength (the resulting diffracted object beam is a highly distorted copy of the original information) or recording with nonlinear effects (which would also distort the reconstructed optical wavefront). [27] Alternatively, these perpetually–erasing holograms could be refreshed by nearly continuous re–recording, at a cost in the input–output performance and complexity of the system.

The most exciting approach under current development is the use of gated photorefractives, in which two different laser wavelengths are used to excite charge out of two trap levels. The first light source excites the recording medium from its ground state to an intermediate energy state. Longer wavelength light then has sufficient energy to write the data hologram in the material by exciting these electrons from the intermediate state up into the conduction band. The spatially distributed charge eventually ends up in the lower trap, creating an interference pattern that is Bragg–matched to the long wavelength readout light but stored in trap levels too deep to be photo–excited [28]. Since the gating light source need not be coherent or even very narrow in wavelength, incandescent lamps with suitable color filters or light-emitting diodes can be used successfully [28].

In any holographic storage device, the retrieval of the stored information relies on the fact that the interference pattern that is used to read out the information is identical to the one that is formed by the superposition of object wave and reference beam during recording. If the interference pattern changes during this process the reconstructed image will be different from the input data page. Temperature differences between recording and read-out obviously effect this grating via thermal expansion. Even during the recording process itself the shrinkage that is typically associated with a photopolymerisation can cause a non-negligible difference between the interference pattern of the recording laser beams and the recorded pattern in the medium! These concerns are the focus of research in several groups that develop holographic materials, but no simple solution has been found yet that avoids these issues completely.

A truly useful holographic recording medium would have all the right optical and mechanical properties, would be almost completely temperature insensitive, and could be economically manufactured. While recent progress in holographic data storage has been impressive, several technical issues have to be resolved before a storage device can be built at a price that is competitive with established data storage technologies. As with all of the optical storage technologies described here, it is also important to not overlook the fact that well-entrenched competitor technologies continue to improve every year.

4 Conclusion

The evolutionary approaches that are based on current technologies and have the potential to be forward compatible—such as the use of depth of focus—have the lowest risks in terms of media, storage device and cost. However, the benefits for potential density gains are also limited, especially if the key features of optical data storage—removeability and interchangeability—are to be retained. And the window of opportunity for becoming a viable storage technology is closing.

On the other hand, the revolutionary approaches promise much higher densities at substantially higher risk. Some of this risk arises because laboratory development tends to demonstrate features individually, avoiding effects which only show up when performance is pushed across–the–board. For instance, as density and data rates increase, maintaining an acceptable signal–to–noise ratio and accessing each desired record without crosstalk become severe concerns that cannot be considered independently. However, some of the more exotic techniques have additional attributes that may make them advantageous for niche applications.

These revolutionary approaches are several years away from functioning prototypes. Issues here are mainly the need for better media, and the relatively high complexity and cost of the "drive." Applications that take advantage of some of the unique properties will in all likelihood be targeted by first product designs. Whether revolutionary optical storage technologies will be able to become mainstream, or whether conventional optical disk storage will continue to evolve, will depend on both technical feasibility and commercial viability. The demand is there, but the jury is still out on the best approach to meet the ferocious appetite of the information age.

References

[1] A. B. Marchant. *Optical Recording.* Addison–Wesley, Massachusetts, 1990.

[2] M. Mansuripur and G. Sincerbox. Principles and techniques of optical data storage. *Proceedings of the IEEE*, 85(11):1780–1796, 1997.

[3] ECMA–267 standard: 120mm DVD – Read–only disk. http://www.ecma.ch, December 1999.

[4] W. S. Oakley. A novel digital optical tape recorder. volume 2604 of *Proceedings of the SPIE*, page 265, 1995.

[5] *NSIC–OIDA Optical Disk Storage Roadmap.* National Storage Industry Consortium and Optoelectronics Industry Development Association, San Diego, 1997.

[6] http://www.mpeg.org/MPEG/DVD/.

[7] http://www.ee.washington.edu/conselec/CE/kuhn/doit96/dhome.htm.

[8] D. A. Thompson and J. S. Best. The future of magnetic data storage technology. *IBM Journal of Research and Development*, 44(3):311–322, 2000.

[9] T. L. Wong and M. P. O'Neill. Multilevel optical recording. *Journal of the Magnetics Society of Japan*, 25(3):433–436, 2001.

[10] M. Mansuripur. *The physical principles of Magneto–optical recording.* Cambridge University Press, 1995.

[11] B. D. Terris, H. J. Mamin, and D. Rugar. Near–field optical data storage. *Applied Physics Letters*, 68(2):141–143, 1996.

[12] M. Kaneko, K. Aratani, and M. Ohta. Multilayered magnetooptical disks for magnetically induced superresolution. *Japanese Journal of Applied Physics Part 1*, 31(2B):568, 1992.

[13] H. Awano, S. Ohnuki, H. Shirai, and N. Ohta. Magnetic amplifying magneto–optical system. In H. Birecki and J. Kwiecien, editors, *Optical Data Storage*, volume 3109 of *Proceedings of the SPIE*, page 83, 1997.

[14] Y. V. Martynov and H. A. Wierenga. Migration path of optical storage drives and media. *Journal of Information Storage and Processing Systems*, 2(1):93–100, 2000.

[15] C3D. White paper. Technical report, Constellation 3D, June 2000. http://www.c-3d.net/whitepaper.html.

[16] S. Kawata and Y. Kawata. Three–dimensional optical data storage using photochromic materials. *Chemical Reviews*, 100(5):1777–1788, 2000.

[17] K. Yamasaki, S. Juodkazis, M. Watanabe, H. B. Sun, S. Matsuo, and H. Misawa. Recording by microexplosion and two–photon reading of three–dimensional optical memory in polymethylmethacrylate films. *Applied Physics Letters*, 76(8):1000–1002, 2000.

[18] S. Hunter, F. Kiamilev, S. Esener, D. A. Parthenopoulos, and P. M. Rentzepis. Potentials of two–photon based 3–D optical memories for high performance computing. *Applied Optics*, 29(14):2058–2066, 1990.

[19] F. B. McCormick, H. Zhang, A. Dvomikov, E. Walker, C. Chapman, N. Kim, J. Costa, S. Esener, and P. Rentzepis. Parallel access 3–D multilayer optical storage using 2–photon recording. In *Advanced Optical Data Storage: Materials, Systems, and Interfaces to Computers*, volume 3802 of *Proceedings of the SPIE*, pages 173–182, 1999.

[20] I. Cokgor, F. B. McCormick, A. S. Dvornikov, M. M. Wang, N. Kim, K. Coblentz, S. C. Esener, and P. M. Rentzepis. Multilayer disk recording using 2–photon absorption and the numerical simulation of the recording process. In *Optical Data Storage 1997*, pages 54–55, Apr 1997.

[21] S. R. Chinn and E. A. Swanson. Multilayer optical storage by low–coherence reflectometry. *Optics Letters*, 21(12):899–901, 1996.

[22] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, and J. G. Fujimoto. Optical coherence tomography. *Science*, 254(5035):1178–1181, 1991.

[23] W. E. Moerner. Molecular electronics for frequency domain optical storage: persistent spectral hole–burning–a review. *Journal of Molecular Electronics*, 1(1):55–71, 1985.

[24] W. E. Moerner, editor. *Persistent Spectral Hole Burning: Science and Applications.* Springer Verlag, New York, 1988.

[25] T. W. Mossberg. Time–domain frequency–selective optical data storage. *Optics Letters*, 7(2):77–79, 1982.

[26] E. S. Maniloff, A. E. Johnson, and T. W. Mossberg. Spectral data storage using rare–earth–doped crystals. *MRS Bulletin*, pages 46–50, Sep 1999.

[27] G. Sincerbox, editor. *Selected papers on holographic data storage*, volume MS95 of *SPIE Milestone Series.* SPIE, 1994.

[28] H. J. Coufal, D. Psaltis, and G. Sincerbox, editors. *Holographic Data Storage.* Springer–Verlag, 2000.

[29] G. W. Burr, S. Kobras, H. Hanssen, and H. Coufal. Content–addressable data storage by use of volume holograms. *Applied Optics*, 38(32):6779–6784, 1999.

[30] G. W. Burr, C. M. Jefferson, H. Coufal, M. Jurich, J. A. Hoffnagle, R. M. Macfarlane, and R. M. Shelby. Volume holographic data storage at an areal density of 250 Gigapixels/in^2. *Optics Letters*, 26(7):444–446, 2001.

[31] S. S. Orlov, W. Phillips, E. Bjornson, L. Hesselink, and R. Okas. Ultra–high transfer rate high capacity holographic disk digital data storage system. In *Proceedings of 29th Applied Imagery Pattern Recognition Workshop*, pages 71–77, New York, 2000. Institute of Electrical and Electronics Engineers.

[32] J. A. Ma, T. Chang, S. Choi, and J. Hong. Ruggedized digital holographic data storage with fast access. *Optical and Quantum Electronics*, 32(3):383–392, 2000.