Near-Field Recording
TeraStor's Mass Storage Benchmark
Technology for the Next Decade

Gordon Knight, Ph.D.

TeraStor Corporation
3510 North First Street
San Jose, CA 95131-1011
Gknight@terastor.com
+1-408-324-2110

Abstract: TeraStor is developing storage products based on near-field recording that deliver a significant areal density improvement over conventional technologies. The solid immersion lens (SIL) has its origins in microscopy and was refined for use in storage applications at Stanford University. Near-field recording, using the SIL, has a substantial history and the products being developed are the result of many years of research and refinement. The basic properties and advantages of near-field recording are discussed in detail to help the reader understand how density increases are achieved.

Introduction

Near field recording (NFR) is based on several basic principles. NFR uses a combination of first surface recording and a flying optical head to achieve it goals. Digital Equipment Corporation originally developed these concepts in the late 1980s. It is believed that over $60M was invested by Digital in their development. Quantum Corporation acquired the patents on this technology as part of their acquisition of Digital's storage business in 1994. Co-exclusive patent rights have been granted by Quantum to TeraStor for both the flying optical head and first surface recording. The basic technology behind the Solid Immersion Lens (SIL) was developed and patented by Dr. Gordon Kino at Stanford University in 1992. Stanford has granted exclusive patent rights for the SIL technology to TeraStor.

The primary advantage of NFR technology is achieving the highest areal density, or number of magnetically charged bits, that can fit into a given area. Products developed by TeraStor will target markets that value high capacity and random access performance, but have a strong desire to lower their storage cost for a given capacity. These markets include backup, digital asset management, and nearline storage management, among others.

NFR technology is form factor independent and permits both removable and fixed drive applications. TeraStor will commercialize products incorporating NFR in mid-1998. TeraStor has also licensed Quantum to deliver compatible NFR-based products into the market.

The Growing Need for Storage

Data storage needs are rapidly changing, yet the data storage solutions currently being produced with conventional technologies are not growing in capacity as fast as the users' needs are growing. This problem is being created in large part by the rapid convergence of different forms of information or "content" into digital form. This convergence is driven by a desire either on the part of the creators of the data or users of the data to be able to share, transport, replicate and distribute data for mass-market usage, and to ensure the data integrity, longevity, and future restoration capability.

This digital content is becoming vital for applications in areas such as multimedia, graphics design, desktop video, intranets, digital broadcast, and digital photography, in addition to others. This digital content is being driven by a wide range of users including students, business workers, artists, photographers, librarians, journalists, and many others. The convenience of digital data, the ability to easily replicate and protect it all lead to the same conclusions; data is being generated at all levels of the storage market like never before and something extraordinary will need to happen to deal with the growth.

Capacity requirements for mass storage devices range from several megabytes in a handheld computer or digital camera, to gigabytes used in today's personal computer systems, to terabytes and beyond used in massive near-line archival systems or on-line database systems. While all of these storage capacities are achievable today, many different technologies must be used, each having a different set of product, performance and cost issues. These issues make the use of currently available technology in large capacity storage applications either extremely expensive, as in the case of hard disk systems, or uncomfortably slow, as in the case of large off-line and near-line libraries using sequential access media. What is needed is a new mass storage technology providing much higher storage capacities than current technologies, with high-performance random access at a lower cost of storage than other on-line, near-line and offline alternatives. This new storage technology should provide the assurance of an aggressive, continuous path to even greater capacities and performance at lower costs for years to come.

Existing Technology Limitations

Storage technologies being used today present end-users with a set of tradeoffs in determining which product will provide the best solution for their requirements. While one may meet performance requirements, its capacity is wrong for the problem. Another may meet capacity requirements, while not meeting cost requirements. Highlighted below are some of the issues with each of today’s available storage product classes.

Hard Disk Drives

Hard disk drives (HDDs) can store from 1 to more than 20 GB. Because of their simplicity of design and relatively high performance, HDDs are ideal for on-line storage of operating systems, frequently used application software and data files. This is because they feature fast random access to data stored anywhere on the disk and allow the data to be retrieved with high transfer rates. In desktop PCs, single HDDs with capacities of 2GB to 4GB are commonly used for this purpose. In larger systems, several HDDs can be used either separately or configured into an array system (RAID) of secure, reliable disks providing tens or even hundreds of gigabytes of storage capacity, but at significant increase in hardware and/or software costs above that of a single drive.

Hard disks usually meet all performance requirements. Where they most often fail to meet requirements is in cost per unit of capacity. This causes end users to architect near-online storage systems, which decrease the cost of storage while increasing access times to some of the data. In environments where data use is infrequent, this can work quite well. In environments where data access is random and constant, near-online storage does not work well. For instance database tables are seldom appropriate for migration to near-line storage. In addition to cost, another attribute missing from hard disks is removability. Today, hard disks are sometimes used in the backup process to hold a mirror image of a production drive, which is then copied to removable media to be stored off site and on low cost media. One thing is for certain; cost makes it impossible to store all of our data in on-line disk drives.

Tape Drives

Tape drives can store from less than 1GB to more than 100GB on removable cartridges or reels of magnetic tape. Their form factors vary widely, as do their performance characteristics. Because of its high storage capacity and low media cost, tape has been the storage medium of choice for backup and archive applications. However, since data must be written and read from the tape sequentially, the amount of time to locate any one piece of data is considerably longer than when using a HDD since the HDD can randomly seek to and access any given data file. With exception at only the very highest end of the product spectrum, data transfer rates are also slower than with HDDs. In addition, when compared to HDDs, tape drives are also perceived as unreliable since tape mechanisms are complex and the tape medium itself is susceptible to a range of environmental threats. For these reasons, tape technology has been limited to data backup and archiving applications (singly or in tape arrays), for temporary file storage and off-line storage in an environmentally controlled vault for disaster recovery purposes.

While its low cost per unit stored is attractive, sequential access tape mechanisms have actually limited the way backup software products can be designed. Backup is looked at as a sequential operation, and being that the backup operation is a streaming function, this is probably OK. Where the sequential access method breaks down is in data restore. Access times to backup copies of data are severely limited by the time needed for a tape drive to locate the required data. In some tape drives this can average several minutes, while the worst case can sometimes be measured in hours. In a better case, it would be possible to handle backups as on-line. But this would require a technology that allows for the low cost of storage, like tape, while presenting the data as on-line like hard disk.

Magneto-Optical Drives

Magneto-optical (MO) drives can store from several hundred MB to several GB of data on a removable disk and come in a variety of form factors. As the need for larger storage capacities on a removable media continues to grow, technologies such as optical storage have emerged to compete with tape and removable HDDs for some applications. All of these optical products employ some sort of laser-based reading and, in the case of writable or rewritable products, laser-based writing technology.

Read-only optical technologies such as CD-ROM and DVD-ROM are very successful as a publishing and distribution medium. Since these products are designed to utilize media that is very low cost but highly reproducible (via a stamping or molding process with a precision master) design tradeoffs are made that are favorable to its use as a tool to access large amounts of published data. These same design tradeoffs, however, become detrimental from a performance standpoint, when these devices are modified to become rewritable mass storage devices and compare unfavorably with higher performance rewritable mass storage products such as HDDs. (see below)

Rewritable optical technologies based on the CD-ROM and DVD-ROM platforms described above, such as CD-R (CD-Recordable), CD-RW (CD-rewritable), DVD-R and DVD-RAM have tried to overcome the performance limitations of a storage technology designed for publishing and distribution applications. These technologies fall short of user needs in the area of performance for use with application software and on-line data files. Thus rewritable products based on CD or DVD technology will continue to fill a market need for short run distribution of high-capacity data files on CD-ROM or DVD-ROM compatible media. The initial appeal of the DVD-RAM disks are simply because they offer a potentially higher capacity random access removable media device than available from any other technology. This appeal is in spite of the fact that performance of these drives is not comparable to hard disk drives. As an example of this, DVD-RAM, when introduced, is expected to offer a storage capacity of just 2.6 GB and transfer rates of just 20 megabits/second (2.5 Mbytes/second) -- much slower than the performance delivered by today's HDDs and unsuitable for serious data storage applications.

Rewritable magneto-optical technology (MO), while overcoming the poor performance of the CD and DVD based platforms, still comes up far short of the performance of HDDs. As an example, the planned extension of MO technology to "ASMO" technology in 1998, while providing capacities as high as 6 GB, will feature access times of 25 to 30 msec and transfer rates of 4 to 40 Mbit/sec, far short of HDDs, yet the product will be relatively expensive when compared with HDDs.

Magneto-optical products continue to face a number of challenges in the market. While they combine two admirable features, random access and removability, the cost of data storage is so high compared to other options that is it usually only used in low capacity applications. Today, the cost of storing data on magneto-optical media is nearly the same cost as keeping the data stored on on-line disks. Furthermore, the use of magneto-optical drives in on-line applications is precluded by their long seek and latency times compared to HDDs.

There are other practical limits to what can be achieved with these conventional mass storage technologies as we go forward in time. In particular, magnetic recording technology used in HDDs faces a potential electromagnetic threshold known as the superparamagnetic limit of the recording material -- the point at which the magnetic domains are so small that their ability to retain a magnetic charge at room temperature is lost. Similarly, the smallest achievable bit cell size in MO technology is limited by the shortest wavelength of light that can be used in lasers, as well as by the highest numerical aperture of a lens that it is practical to achieve in mass volume production. These thresholds present some very real problems that threaten to slow down the rate of increase in achievable areal density of these conventional technologies just at the time when they need to be increasing.

Areal Density - The Key Metric

In the mass storage industry, the metric used most commonly to define achievable storage capacity of any technology is ***areal density.*** The areal density is calculated by multiplying the number of achievable data bits per inch of longitudinal track times the number of data tracks per inch. The resultant areal density is defined in millions or billions of bits per square inch of recording area:

$$\text{Areal Density (bits/in}^2\text{)} = \text{(bits-per-inch) x (tracks-per-inch)}$$

To grow the areal density over time, both the number of bits-per-inch (BPI) and the number of tracks-per-inch (TPI) need to be increased by making the bits smaller and stored closer together. In 1996, the highest achievable average areal density in a shippable product was between 900 and 1,000 Mbits/square inch. In 1997, it will be between 1.3 and 1.5 Gbits/square inch. In 1998, TeraStor's first product is targeted at 15 Gbit/sq in., up to 10 times the areal density of previous technologies.

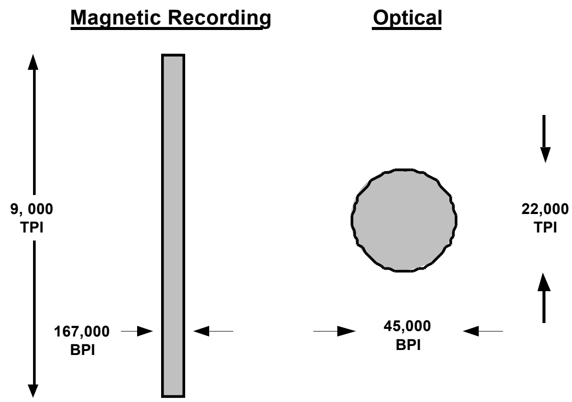

Figure 1. Recorded area comparison between magnetic recording and optical

All storage technologies compete with one another to continually raise the achievable limit of areal density. Each storage technology has particular advantages and disadvantages when it comes to increasing the areal density. In HDDs, BPI is typically much, much higher than TPI because of the properties of the magnetic media and the read/write element of the head, and the tracking methods used. As a result HDDs with 9,000 TPI and 167,000 BPI are expected to be available in 1998. Increases in the TPI are usually more difficult than increases in BPI for magnetic HDDs.

In optical recording, the achievable TPI is usually much higher than magnetic HDDs while the BPI is usually lower. The higher TPI is achievable because high track densities can be stamped or molded with servo tracks in plastic media rather than discrete servo wedges that are recorded by servo-track writers for the HDDs. The laser beam produces a round spot the width of the stamped track. Unlike HDD embedded servo tracks, the stamped tracks allow for continuous track following. The bit density of optical recording is determined and limited by the diameter of the spot size that a laser can be focused to on the surface of the media and is a function of the wavelength of light from the laser. Optical drives with 22,000 TPI and 45,000 BPI are expected to be available in 1998. Figure 1 illustrates the relative recorded area between a magnetic recorded bit and an optical recorded spot.

The Near-Field Solution

TeraStor Corporation has developed a new storage technology that allows an order of magnitude higher areal density when compared to the conventional magnetic and optical technologies that exist today at a significantly lower cost of storage than any current technology. Because of these attributes, users will redefine the boundaries of the storage hierarchy and change user expectations of the capacity, cost and performance features available for any given storage application.

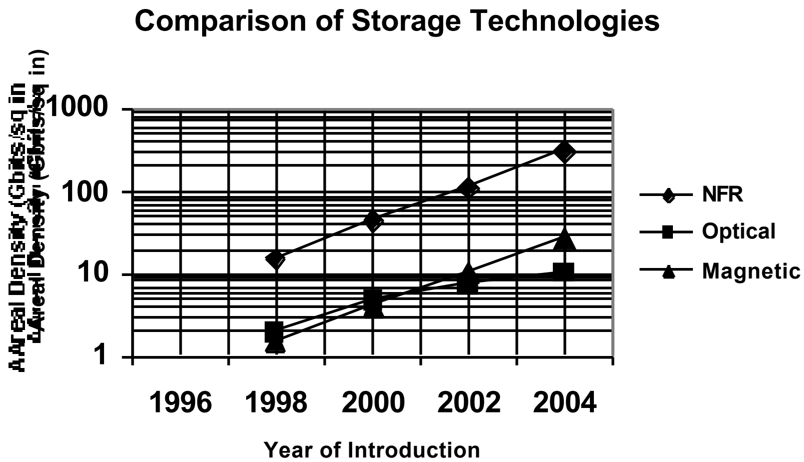

Figure 2. Areal Density Growth Comparison

This new storage technology, called Near Filed Recording (NFR), provides the highest areal density achievable in the storage industry today. The technology provides the capacity and cost of storage comparable to high-end tape drives, but with the performance (seek time and transfer rate) and costs comparable to mid-range HDDs. In addition, TeraStor's well-defined NFR technology roadmap ensures that this areal density advantage will be sustainable for at least the next decade by continuous increases to both the BPI and TPI.

NFR technology is a combination of design elements adapted from several related technology fields, including magnetic recording as in hard disk drives, optical recording, consumer electronics and microscopy. The key elements of TeraStor's *NFR* technology include the flying optical head, the solid immersion lens (SIL), first surface recording, and crescent recording.

Each of these key elements contributes to achieving the areal density, capacity, performance and cost advantages of *NFR* technology over current technologies. The flying optical head, which is similar to a HDD flying head, provides the simplicity and low cost associated with hard disk drives. This allows the recording element to be placed close enough to the recording media so the distance separating them is less than the wavelength of the laser light, placing it into the near field. One of the key optical components inside the flying head is the SIL, which is used to tightly focus the laser beam to produce an ultra-small spot. The energy from this ultra-small spot is then transferred or coupled onto the first surface (top surface) of the disk in an effect called evanescent coupling. In the near field recording process, a tiny magnetic coil in the flying head writes data to the heated spot. The ultra-small bit domains are written in overlapping sequences, creating a series of crescent-shaped bit domains. This crescent recording effectively doubles the linear bit density, allowing *NFR* to achieve an even higher areal density.

This combination of technology and techniques comprises TeraStor's ***NFR*** technology, which results in storage densities an order of magnitude higher than current conventional technologies. This areal density advantage is sustainable to capacities of hundreds of gigabytes per disk as the laser wavelength, SIL shape and material, read channels, as well as other elements of the system, improve over time.

TeraStor is the only company that can combine all of these elements into one technology, which leapfrogs the areal density of current mass storage technologies, while maintaining competitive costs and high performance levels. TeraStor products based on NFR will draw heavily on existing MO and HDD materials and processes, allowing low manufacturing costs and rapid production ramps, comparable to those currently seen in high volume HDD products.

Technology Components

The TeraStor flying optical head acts as an aerodynamic, mechanical platform for mounting the objective lens, the SIL lens and the tiny write coil and positions this assembly at a close distance above the surface of the spinning disk. If the distance between the SIL lens and the recording disk is much less than the wavelength of the laser, the resolution of the spot within the SIL is maintained across the air gap through evanescent coupling. Using a red laser as a reference point, the distance between the bottom of the SIL lens and the recording surface would have to be a fraction of the wavelength of a red laser, or less than 0.685 microns. By using a flying head, achieving this close proximity is actually much easier than one might expect. Drawing from the flying heads used in HDD technology, TeraStor is able to produce a laser-based flying head that flies at a distance less than 6 micro-inches or 0.15 microns, well within the 0.685 micron requirement (see Figure 3).

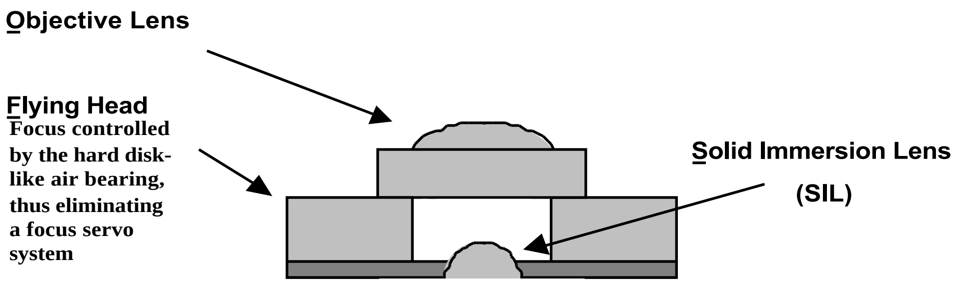

Flying Head Cross-section

Figure 3. Flying Head With Objective Lens and SIL

The NFR optical head fly height is less than 6 micro-inches -- significantly higher than heads used in current generation HDDs. In fact, because the air-bearing slider accurately controls the fly height, there is no need for a servo control system to maintain the focus between the lens and the media. The areal bit density resulting from NFR is not directly related to fly height, as it is in HDD systems. In future generations, significant increases in areal density can be achieved without reducing fly height to a distance anywhere close to that of present HDD technologies.

The evanescent coupling or transfer of laser energy heats the spot on the recording surface to the Curie point of about 200 degrees Celsius in roughly one nanosecond. At this temperature, the irradiation heats the molecules in that spot to a finite depth, enabling magnetization when placed within a magnetic field. This magnetic field (positive or negative) is pulsed into the heated spot by a planar coil embedded within the head substrate.

The planar magnetic coil is, as the name suggests, a flat coil that rests on the same plane as the flying head surface. Significantly, this extremely light, small coil resides inside the

flying head assembly, rather than underneath the substrate of the recording media, as is the case with traditional Magneto Optical technology (see Figure 4).
This orientation has two important advantages:
1. Direct overwrite. Through high-speed switching of magnetic pulses, this small head-based coil is able to directly overwrite without going through a complete rotation of the disk. This enables a significant increase in write performance over what is available using MO technology.

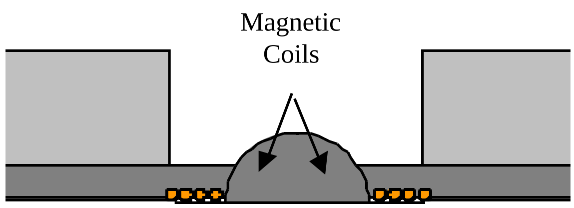

Figure 4. Magnetic Coil in Flying Optical Head

2. Two-sided recording. By embedding the magnetic coil in the head rather than the backside of the disk, TeraStor's solution supports two-sided disks with two heads on line per disk.

In current diffraction-limited optical systems, spot size can be reduced in one of two ways: by decreasing the wavelength of the laser or increasing the numerical aperture (NA) of the objective lens. This approach has clear theoretical limits. For example, in a standard diffraction system, the numerical aperture of a lens in air can never be greater than one.

In the early 1990s, scientists at Stanford University overcame the previous limits on spot size through the use of a new optical system that provides a totally different approach to reducing spot size. This approach increases the effective NA above the theoretical limit of one by adapting an old, but relatively obscure, technique known as liquid immersion microscopy.

In liquid immersion microscopy, the lens and the object to be studied are both placed in a medium such as oil, thereby increasing the magnification of the object beyond what is possible with the microscope alone. The Stanford scientists -- most notably, Dr. Gordon Kino -- simply inverted this effect through the use of a solid lens shaped like a truncated sphere (thus the term "solid immersion lens"). When placed between the standard objective lens and the writing surface, the SIL focuses the incident rays of the laser from the objective lens to a single spot at the base of the partial sphere, as shown in Figure 5. The resulting spot is approximately half the diameter of the spot achieved by conventional means that use only an objective lens.

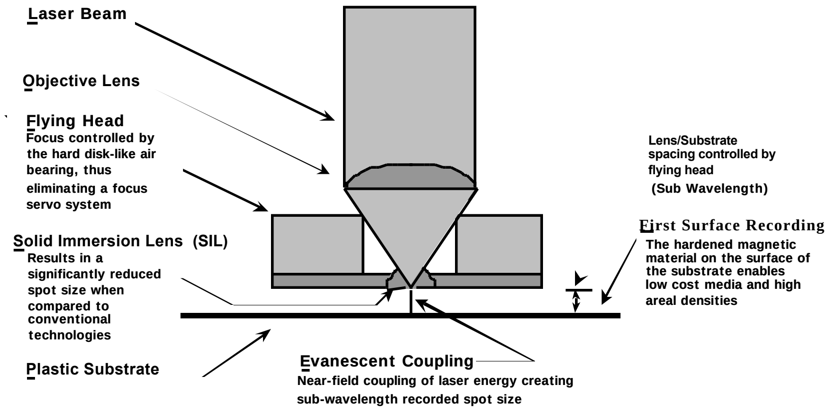

Figure 5. Architecture of TeraStor NFR technology

In technical terms, the SIL lens slows down the laser light beam to a fraction of its normal speed in air, thus shortening the beam's wavelength and enabling a very fine spot size. Notably, the composition of this SIL lens and its refractive index plays an important role in this process. In particular, materials with a high index of refraction will produce a higher effective NA than those with a lower index of refraction. Air, for example, has an index of refraction of 1.0, while common glass has an index of refraction of 1.5, and diamond has an index of refraction of 2.4. The focused spot size can be calculated using the following equation:

$$\text{Diameter of Spot} = \frac{\text{Wavelength of Laser Light}}{\text{2NA(Objective Lens) x n(Index of Refraction of SIL)}}$$

For example, with a common red laser of 0.685 microns wavelength, and a numerical aperture of 0.65, the focused spot diameter would equal 0.53 microns. If this same focused beam is put into a SIL with an index of refraction of 2.0, the spot diameter is reduced by a factor of two to 0.26 microns.

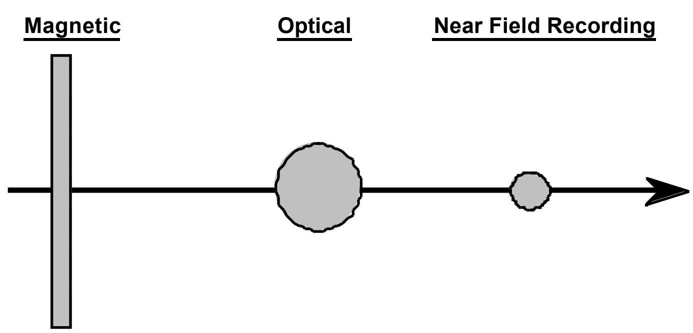

Figure 6. Recorded Area Comparison

The objective lens focuses the laser light in a manner that results in a spot that exists only within the SIL lens. Thus, if the lens were placed near a recording surface, only a small fraction of the light from the laser would appear on that surface (the on access rays). To make the small spot at the bottom of the SIL lens appear on the surface, the Stanford researchers took advantage of another technique known as evanescent coupling.

The concept of evanescent coupling comes from quantum physics and is best introduced by analogy. Let's say, for example, that a doctor puts on a stethoscope and slowly brings the amplifier near a patient's chest. The doctor would begin to hear a heartbeat just before actually coming in contact with the skin. At this point, the stethoscope is within the near field of the heart's sound waves, causing the waves to couple with the stethoscope and produce a startlingly high-quality sound in the instrument even though no physical contact has been made. A similar effect occurs when an energized microwave waveguide (a tube-like device) is placed closely parallel to a second, non-energized waveguide: a portion of the microwave energy begins traveling down the second waveguide. The microwave energy in the first tube couples with the second tube and produces a corresponding energy field in the second tube.

Evanescent coupling was once thought to be impossible, since a sudden discontinuity, such as a wave jumping across barriers, was believed to actually violate the laws of physics. But equations developed in the last century by Scottish physicist James Maxwell indicate that any form of electromagnetic radiation enclosed within a boundary will produce waves outside that boundary that decay at an exponential rate. If an object is close enough to that boundary (within a fraction of a wavelength), it will receive some of the transmitted radiation. In some cases, evanescent coupling can produce a very high-quality effect, inducing more than 50 percent of the energy in the source radiation to couple with the material in the near field. This is precisely what happens in NFR.

In this case, a recording surface is placed within the near field of a SIL lens and a laser spot is focused on the bottom of the lens. The close proximity evokes an evanescent coupling of laser energy, resulting in a "transfer" of the spot from the inside of the SIL to the surface of the disk, creating a small but well-defined spot on the recording disk.

Flying close to the disk surface would mean little if the recording material or storage layer was buried under layers of protective plastic substrate. Traditional MO media uses multiple layers of sputtered material with an intervening substrate layer that places the recording material too far away from the flying head/SIL lens to be practical for NFR. As shown in Figure 7, TeraStor solves this problem by sputtering the magnetic storage layer onto the surface of the disk (and not putting layers of plastic on top). This hardened magnetic material on the top surface of the disk supports NFR and enables high-volume, low-cost production by manufacturers familiar with present MO recording media. The disk itself can be a low-cost plastic injection molded substrate, providing significant cost-savings relative to metal or glass substrate media.

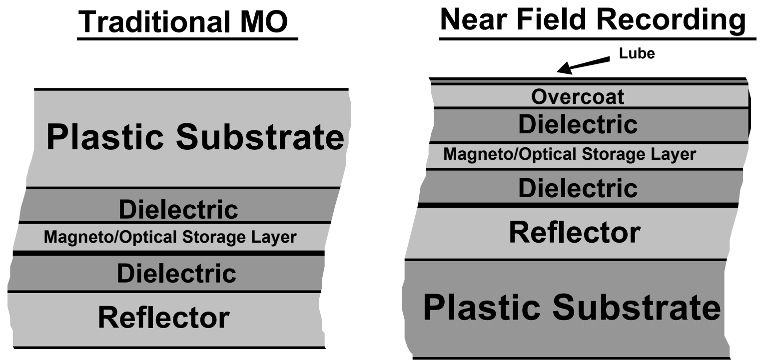

Figure 7. Cutaway view of Magneto/Optical and NFR recording media

Up to this point, the TeraStor solution described has combined a SIL lens, evanescent coupling, a flying head and a planar magnetic coil enabling the production, on first surface recording media, of a bit cell that is smaller in diameter than is attainable using any other recording technology. As the final step in the process of NFR, TeraStor has implemented a crescent recording technique where domains partially overwrite each other, producing distinctive crescent-shaped domains increasing the areal density of the recording surface.

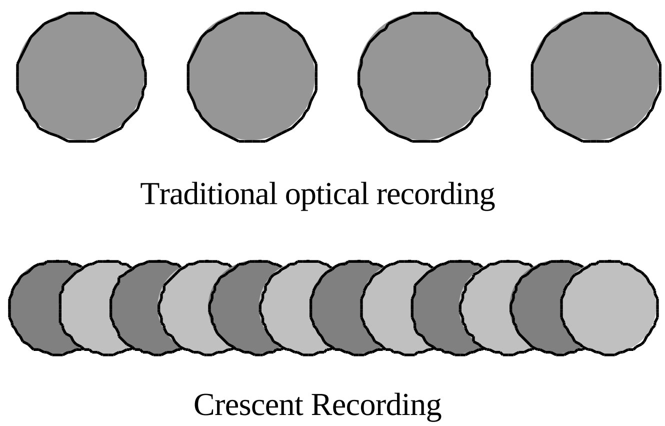

Figure 8. Crescent recording compared to traditional optical recording

The areal density is achieved by the combination of the small spot size produced by the SIL and evanescent coupling along with the overlapping pattern of crescent recording. This results in nearly equivalent linear bit density or BPI as compared to traditional HDDs, but with significantly higher TPI, as depicted in Figure 9.

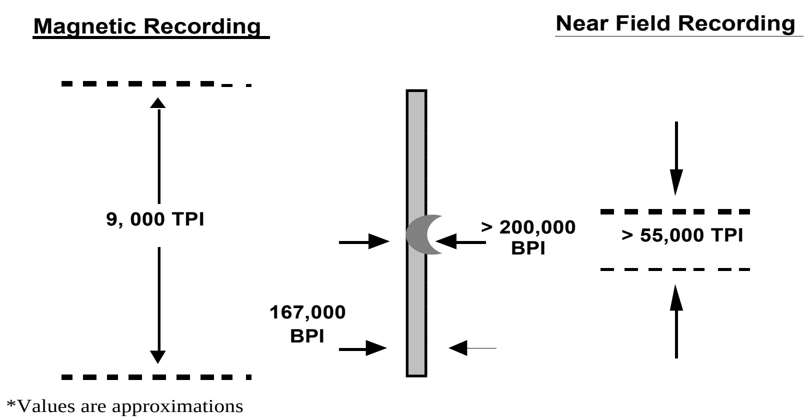

Figure 9. Areal Density Comparison

Conclusion

NFR represents an entirely new category of mass storage technology with distinct advantages over all previous technologies. This breakthrough technology delivers more storage capacity per square inch at a lower cost per gigabyte than any of today's other mass storage technologies, including HDD, tape and MO systems. At the same time, the technology is also easily implemented in both removable media and fixed media devices.

NFR enables direct overwrite and very small bit cells. Further, the technology provides direct access capabilities. All of these features enable the high capacity and performance levels required to serve the needs of high-end data center and workstation environments, as well as the cost savings necessary to support desktop systems. The random access performance of NFR makes it equally suitable for storing system files, applications and actively used data, including video, 3-D images, Internet downloads, CAD/CAM files and more.

With its easy support for standard interfaces such as SCSI, FC-AL, IEEE 1394 and IDE, NFR technology is bound to begin displacing tape, MO and removable HDD technologies in backup and archive applications as soon as products become available.

As the following chart indicates, TeraStor's NFR technology provides a variety of technical benefits compared to HDDs, magneto-optical, and tape drives. Notably, NFR provides a less expensive substrate, a fast direct overwrite capability, and greater bit density per square inch than competing technologies. NFR also allows HDD data transfer rates and does not require an expensive, complex focus servo control to enable accurate recording. In addition, because of its vertical bit cell writing technique, NFR has no known areal density limits, such as the superparamagnetic limit identified for HDD technology.

Compared to tape technology, TeraStor's NFR technology provides equal or greater storage capacities, higher transfer rates, random access, and low-cost, long-life recording media. This high-performance, low-cost solution provides reliable backup and archiving capabilities with a media life of at least 30 years -- well in excess of anything tape drives can offer. NFR provides a clear capacity advantage over MO technologies. In addition, NFR allows much faster transfer rates and access times and much lower drive costs. NFR enables a much higher capacity or lower cost/GB than is possible with the removable HDDs currently available. DVD-RAM, like its older counterpart, the compact disk, is mainly a device for publishing and distributing mass-produced software content, thus performance has never been near that required for a system data storage device. Eventually, NFR technology will be used as the primary, nonvolatile data storage devices designed into systems. Since NFR provides up to a 10x areal density improvement over current HDD technology, it is expected to be an attractive choice for users as an add-on storage technology for desktop systems and workstations.

Table 1 TeraStor NFR vs. Alternative Technologies

	HDD	Optical	Tape	NFR
Random access	x	x		x
Sector addressing	x	x		x
High performance	x			x
Removability	x	x	x	x
Low Cost/GB			x	x

We can only imagine how this technology will change the world of storage management. The value of small tape libraries will diminish with the first product introductions. The changes allowed to backup strategies and architectures are hard to predict, but one can easily imagine making backup, and restore for that matter, an on-line phenomenon. Finally, what contributes to the choice of tape for near-line storage systems will become less clear as high capacity removable and fixed disk products based on near field recording are introduced in the future.