SPIE Conference on *Three– and Four–Dimensional Optical Data Storage*
July 31, 2001 Paper **4459–52**

Progress in read–write fast–access volume holographic data storage

Geoffrey W. Burr[1], Erwin Mecher[2], Thorsten Juchem[3], Hans Coufal[1], C. Michael Jefferson[1], Mark Jurich[1], Francisco Gallego[2], Klaus Meerholz[2], Norbert Hampp[3], John A. Hoffnagle[1], Roger M. Macfarlane[1], and Robert M. Shelby[1]

[1] *IBM Almaden Research Center,*
650 Harry Road, San Jose, California 95120
[2] *Institut für Physikalische Chemie,*
Ludwig Maximilians Universität München, 81377 München, Germany
[3] *Institut für Physikalische Chemie,*
University of Marburg, D–35032 Marburg, Germany

ABSTRACT

We review recent progress made towards commercializable read–write, fast–access holographic data storage. This includes a recent demonstration of high areal density holographic storage [1], systems architectures for extending this high density to high capacity using phase–conjugate readout [2], and recent experimental progress along these lines. Other topics include using signal processing to relieve alignment and distortion constraints [3], optical elements for improving beam uniformity [4], and most importantly, requirements and prospects for improved photorefractive materials for two–color, gated nonvolatile holographic storage [5].

Keywords: volume holographic data storage, read-write holographic data storage, phase–conjugate readout, nonlinear signal processing.

1. INTRODUCTION TO HOLOGRAPHIC DATA STORAGE

An intriguing approach for next generation data storage uses optical holography to store information throughout the three–dimensional volume of a material [6–11]. By superimposing many holograms within the same volume of the recording medium, holograms can potentially store data at a volumetric density of one bit per cubic wavelength [12]. Given a typical laser wavelength of 500 nm or so, this density corresponds to more than 10^{12} bits (1 Terabit) per cubic centimeter.

In holographic storage, data are transferred to and from the storage material as 2–D images composed of thousands of pixels, each of which represents a single bit of information. Since an entire "page of data" can be retrieved by a photodetector at the same time, rather than bit–by–bit, the holographic scheme promises fast readout rates as well as high density [6–11]. If a thousand holograms, each containing a million pixels, could be retrieved every second, for instance, then the output data rate would reach 1 Gigabit per second. Despite this attractive potential, however, research into holographic data storage all but died out in the mid–1970s due to the lack of suitable devices for the input and output of pixelated 2–D data pages. In the early 1990s, interest in volume–holographic data storage was rekindled by the availability of devices that could display and detect 2–D pages, including charge coupled devices (CCD), complementary metal–oxide semiconductor (CMOS) detector chips and small liquid–crystal panels.

In this paper, we describe our recent progress towards fast access, read–write holographic data storage systems. We describe a research platform for high–areal density storage, on which we have experimentally demonstrated areal densities equivalent to 390 bits per square micron. This density exceeds the storage capabilities of DVD disks by nearly 80×, and of commercial magnetic disks by a factor of 10×. We then describe a nonlinear signal–processing method which was one of the key enablers of this demonstration. Then we describe a system architecture that can allow us to extend this high density to high capacity, which combines phase–conjugate readout and multiplexed volume holography by using an intermediate "buffer" hologram. Finally, we discuss non–volatile, read–write storage in photorefractive crystals using two–color, gated lithium niobate.

For further information, contact G. W. Burr at *burr@almaden.ibm.com*

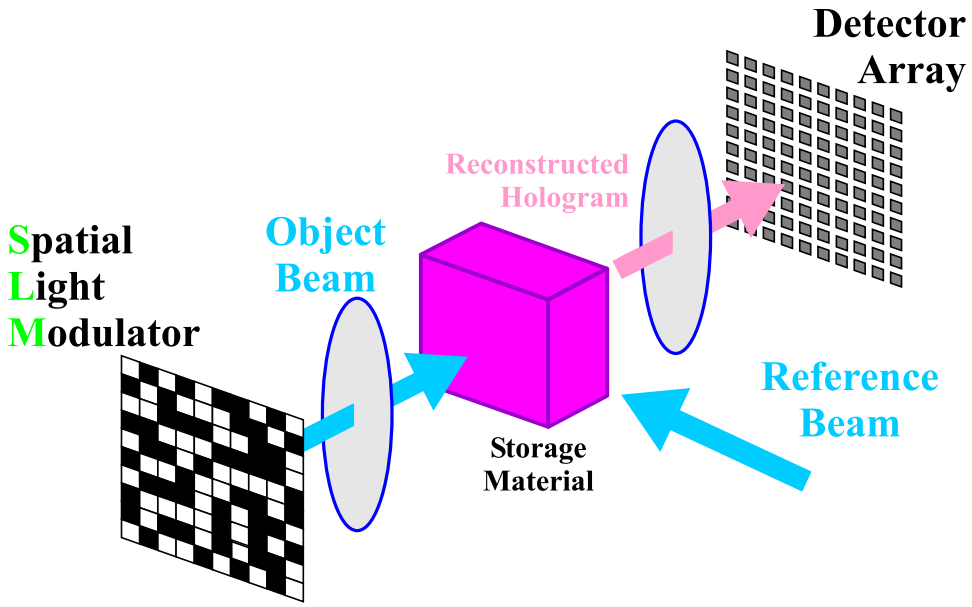

Figure 1. Data are imprinted onto an object beam by shining laser light through a spatial light modulator. A pair of lenses image the data through the storage material onto a pixelated detector array, such as a charge–coupled device (CCD). A reference beam intersects the object beam in the storage material, allowing the holograms to be stored and retrieved later.

1.1. Holography

To make a hologram, two coherent laser beams (the reference and data–bearing object beams) are overlapped in a photosensitive medium such as a photopolymer or inorganic crystal. The resulting optical interference pattern is preserved by the storage media as a change in its optical properties. When the recorded hologram is illuminated by the readout beam, some of the light is diffracted to "reconstruct" a weak copy of the object beam. If the object beam originally came from a 3–D object, then the reconstructed hologram makes the 3–D object reappear.

If the hologram is recorded in a thick material, the reconstructed object beam will only appear when the readout beam is nearly identical to the original reference beam. Since the diffracted wavefront accumulates energy from throughout the thickness of the storage material, a small change in either the wavelength or angle of the readout beam generates enough destructive interference to make the hologram effectively disappear through Bragg selectivity. This destructive interference allows multiple stored holograms to be superimposed within the same common volume, with independent access to each hologram through illumination with its original reference beam. Several different techniques have been developed to define a set of suitable reference beams by, for example, slightly changing the angle, wavelength or phase of the original light beam [11]. Using so–called "angle multiplexing," as many as 10,000 holograms have been stored in the same 1 cm^3 volume [13].

1.2. Storing and retrieving digital data

To use volume holography as a storage technology, digital data must be imprinted onto the object beam for recording and then retrieved from the reconstructed object beam during readout (Figure 1). The device for putting data into the system is called a spatial light modulator (SLM)—a planar array of thousands of pixels. The data are read using a similar array of detector pixels, such as a CCD camera or CMOS sensor array. The object beam often passes through a set of lenses that image the SLM pixel pattern onto the output pixel array, as shown in Figure 1. To maximize the storage density, the hologram is usually recorded where the object beam is tightly focused. To access holographically–stored data, the correct reference beam must first be directed to the appropriate spot within the storage media. The hologram is then reconstructed by the reference beam, and a weak copy of the original object beam continues along the imaging path to the camera, where the optical output is detected and converted to digital data.

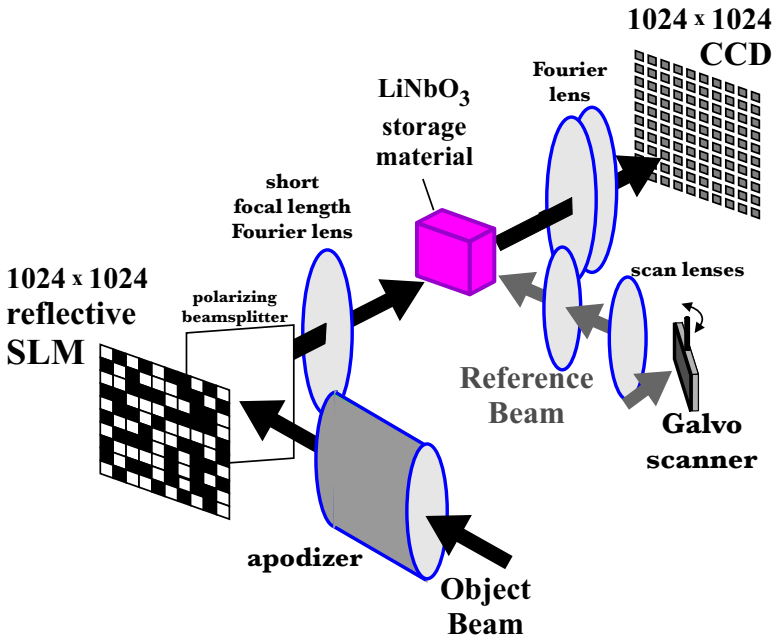

Figure 2. Salient features of the DEMON II holographic digital data storage engine. Utilizing 30mm focal length Fourier transform lenses in the 90 degree geometry with a 1 million pixel SLM, this system has demonstrated an areal storage density of nearly 400 bits per square micron.

The output speed of a storage device can be jointly described by two parameters: the readout rate (in bits per second) and the latency, or time delay between asking for and receiving a particular bit of data. The latency tends to be dominated by mechanical movement, especially if the storage media has to be moved. The readout rate is often dictated by the camera integration time: the reference beam reconstructs a hologram until a sufficient number of photons accumulate to differentiate bright and dark pixels. A frequently mentioned goal is an integration time of about 1 millisecond, which implies that 1000 pages of data can be retrieved per second. If there are 1 million pixels per data page and each pixel stores one bit, then the readout rate is 1 Gigabit per second. This goal requires high laser power (at least 1 W), a storage material capable of high diffraction efficiencies, and a 'megapel' detector (one with a million pixels) that can be read out at high frame rates. Despite these requirements, even faster readout and lower latency could be reached by steering the reference beam angle non–mechanically, by using a pulsed laser, and by electronically reading only the desired portion of the detector array.

Both capacity and readout rate are maximized when each detector pixel is matched to a single pixel on the SLM, but for large pixel arrays this requires careful optical design and alignment. In order to study the recording physics, materials, and systems issues of holographic digital data storage in depth, we have built several precision holographic recording testers on which this pixel–to–pixel matching has been achieved. In Reference [10], we described our earlier work on testing of holographic storage materials using our 'PRISM' tester [14–16] and on development of signal processing and coding techniques using our 'DEMON I' platform [17–20]. Here, we describe our subsequent efforts towards high areal density and high read–write system capacity using the 'DEMON II' and 'DEMON III' platforms.

2. HIGH–DENSITY HOLOGRAPHIC STORAGE

The DEMON II holographic storage platform, shown in Figure 2, was designed to demonstrate high density holographic data storage, while including aspects of our two previous test platforms [14, 17]. DEMON II combines the large data pages of the PRISM tester [14], the dynamic SLM and the 90 degree geometry configuration of the DEMON I platform [17], and the short focal length optics needed for high density [21].

The laser light is provided by a diode pumped solid state laser (532 nm, doubled Nd–YAG); waveplates and polarizing beamsplitters provide control over the power in reference and object beam. The object beam was apodized

with a pair of aspheric optical elements, achieving illumination uniformity within 5% rms with 80% of the input optical power [4]. The SLM is a reflective device fabricated by IBM Yorktown [22], containing 1024×1024 pixels and illuminated via a polarizing beamsplitter cube. The magnification from the 12.8 micron pitch of the SLM pixels to the 12 micron pitch of the 40Hz, 1024×1024 pixel CCD camera is built into the Fourier optics (effective focal length 30mm). The use of two separate elements in the back Fourier lens (between the storage material and the detector array) allows the magnification of the optical system to be varied over a range of ±0.5%. Linear stages provide two axes of motion for the storage material, and three axes of motion for the detector array. A pair of scan lenses provide an improved relay of the reference beam from the galvo mirror to the $LiNbO_3$ crystal, providing diffraction limited performance over an angular scan range of ±15 degrees. The entire system, including the laser, occupies 0.8×0.8 square meters.

The short focal length of the DEMON II optics allows the system to demonstrate high areal storage densities (the storage capacity of each stack of holograms, divided by the area of the limiting aperture in the object beam). Since the lenses in the object beam implement a two–dimensional spatial Fourier transform, an aperture placed in the central focal plane of the 4–f system (just in front of the storage material) can be described as a spatial low–pass filter. The smaller the volume allocated to each stack of holograms, the larger the capacity of a given large block of storage material. However, if the aperture is decreased too far, some of the information from the SLM fails to pass through the aperture. The size of the smallest tolerable aperture corresponds to the spatial equivalent of the Nyquist sampling condition. That is, the spatial frequency sampling on the SLM (one over its pixel pitch) is twice the maximum spatial frequency allowed to pass the limiting aperture. Apertures smaller than this Nyquist aperture do not pass all the information imposed by the SLM through to the detector array. The Nyquist aperture is equal to the inverse of the pixel pitch of the SLM, scaled by the wavelength and the focal length of the lenses. The design of the imaging optics is then complicated by this need for short focal length, since the maximum ray angle (and thus the potential for optical aberrations) is greatly increased. The optical distortion (displacement of pixel centers from a rectangular grid) in the DEMON II platform is consequently much larger than in the earlier two testers, reaching approximately 0.03% (0.3 pixels) in the corners of the received data page.

Once the optical system is configured to successfully pass the SLM data to the detector array despite the small aperture at the storage material, multiple holograms must still be recorded and retrieved independently. This adds two additional terms to the overall signal–to–noise budget: inter–page crosstalk caused by incomplete Bragg mismatch, and the decrease in diffraction efficiency caused by sharing the finite dynamic range of the storage material among the superimposed holograms. In the 90° geometry in $LiNbO_3$, interpage crosstalk is less important than the limited dynamic range. If M is the number of superimposed holograms, then the diffraction efficiency falls as $(M\#/M)^2$ [23]. Here, $M\#$ is a scaling coefficient which describes both system and material properties [24]. Achieving high areal density is then a balancing act between the inter–pixel crosstalk introduced by the small aperture, and the loss of signal associated with recording multiple holograms.

2.1. High areal density holographic data storage

In our first high–density experiment, we angle–multiplexed one thousand volume holographic data pages, each containing one million pixels, in a common volume of lithium niobate [1]. After each hologram exposure, the crystal was displaced both horizontally and vertically to average out the BER degradation caused by the photovoltaic effect in the absence of a random phase mask [1,24,25]. This increases the maximum tolerable exposure time and modulation depth, improving the achievable $M\#$ and thus the number of holograms that can be superimposed [1].

However, this crystal motion implies that each stack of multiplexed holograms is partly exposed while neighboring stacks are recorded. To account for this experimentally, nine stacks of one thousand holograms each were spatially multiplexed in the crystal, arranged over the object beam entrance face in a 3×3 grid on 1.6mm centers. This guarantees that every hologram in the center stack received optical exposure corresponding to 1000 overlapping holograms, regardless of its relative position. Each hologram had an average diffraction efficiency of 1.1×10^{-6} or a realized $M\#$ of 1.07. With each hologram containing 1.01 million raw pixels and the effective area defined by the stack spacing of 1.6mm, the resulting areal density in channel bits was 394 pixels/μm^2 (254 Gpixels/in^2) [1]. For comparison, a CD has a density of ∼0.7 bit/μm^2, a DVD disk ∼4.5 bit/μm^2, and the magnetic disk in the 1GB IBM MicroDrive, ∼23 bit/μm^2.

After recording, one hundred pages were sampled from the central stack with a readout power of 345mW. One of the retrieved pages contained text instead of binary data, and is shown in Figure 4 to show the qualitative

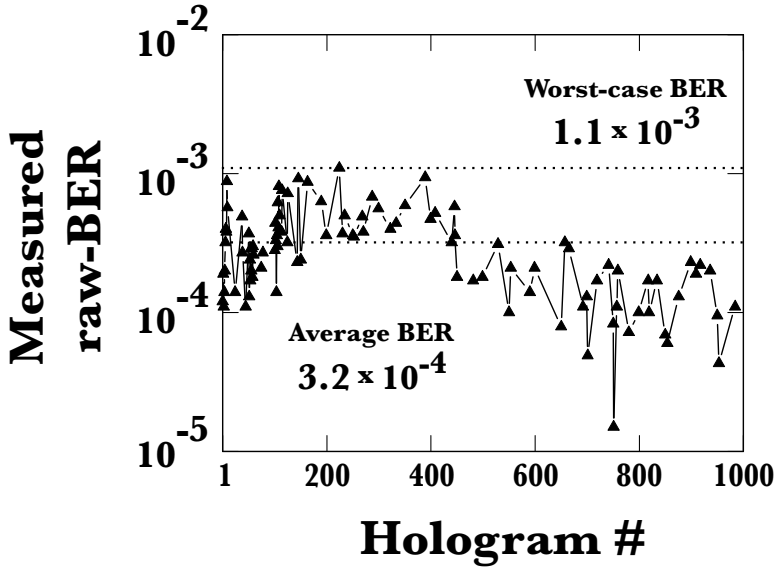

Figure 3. Measured raw–BER results from the DEMON II tester at high areal density [1].

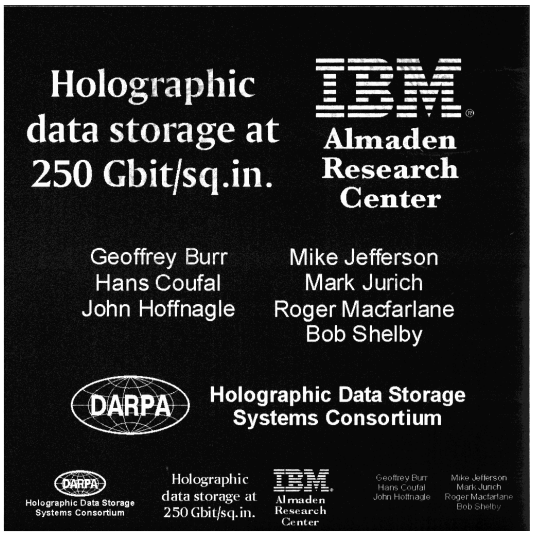

Figure 4. Retrieved data page #1000 from 250 Gbit/sq. in. demonstration using DEMON II.

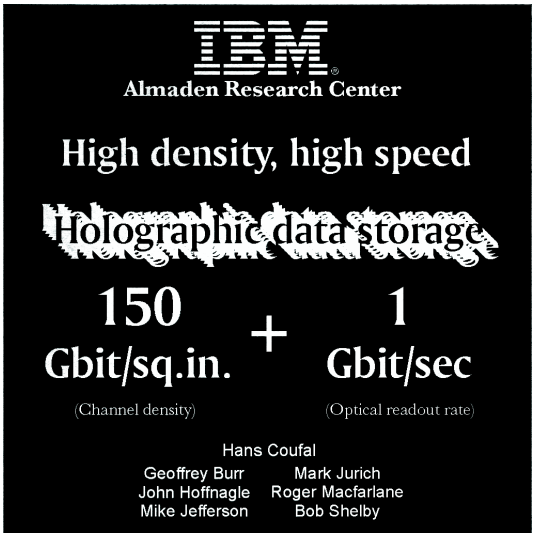

Figure 5. Retrieved data page #600 from 1 Gbit/second at 150 Gbit/sq. in. demonstration using DEMON II.

reconstruction fidelity. The other retrieved data pages were post–processed with a novel algorithm to compensate both global page misregistration and the local pixel offsets of optical distortion [3], as described in the next section. These processed data pages were then decoded with a strong 8–bits–from–12–pixels modulation code [17]. As shown in Figure 3, the worst–case raw–BER before error correction was $\sim 10^{-3}$, sufficient to deliver a user–BER of 10^{-12} with about the same level of error–correction as a CD player [1]. The overall system code–rate of 0.61 leads to a storage density in user bits of 241.5 bits/μm^2 (155.8 Gbits/in^2). Given the 5.5mm hologram thickness, the demonstrated channel density of 0.072 pixels per cubic micron corresponds to 1.08% of the well–known theoretical volumetric density limit of $1/\lambda^3$ [12].

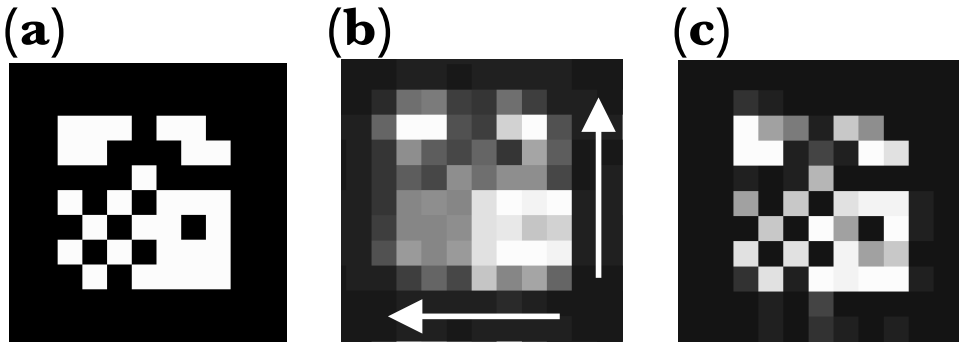

Figure 6. A 9×9 pixel pattern is imaged from SLM to CCD (a) under perfect conditions; (b) with half–a–pixel offset in both x and y; (c) after post–processing with the shift–compensation algorithm [3].

2.2. High speed plus high density

By adding a Pockels cell immediately after the laser in DEMON II, we have recently demonstrated high areal density storage at a readout rate of 1 Gbit/second. The readout shutter time of 1 millisecond per hologram supplied the optical flux corresponding to a 1 Gbit per second data stream. (The additional laser timing, camera readout, and electronic decoding required for this same data rate were demonstrated at low areal density (<10 Gb/sq.in.) by the DARPA-sponsored HDSS consortium [26]).

Although we have previously demonstrated high readout rate (1 Gbit/sec) at low density (0.23Gb/sq.in.) [15], as well as the high density (254 Gb/sq.in.) at moderate readout rate (40Mb/sec) described above and in Reference [1], the strong dependence of output signal strength on both density and data-rate makes it difficult to achieve both experimentally. According to this tradeoff, the 25-fold increase in readout rate achieved here should have required a 5-fold loss in density; however, here we have significantly beaten this by retaining 60% of the density achieved at low readout rate.

This high speed demonstration is primarily enabled by increased laser power (10–fold). However, this added laser power cannot be used to increase the density in the 254 Gb/sq. in. experiment. The reason for this is that the scatter noise does not follow the M^2 vs. t_{int} tradeoff alluded to in the previous paragraph. Extra laser power applied in the presence of thermal noise can be used either toward higher density (more holograms) or faster speed (shorter integration times). The laser power increases signal levels, making it possible to then decrease signal with more holograms or less integration time yet retain the original signal–to–noise ratio (SNR). In the presence of significant scatter noise, however, more laser power increases both signal and scatter, so that when more holograms are stored and signal levels decrease again, the SNR is lower than it was before. In contrast, decreasing the integration time decreases *both* signal and scatter, returning the SNR to the same acceptable level.

3. POST–PROCESSING TO COMPENSATE FOR PIXEL SHIFTS

The high density and fast readout offered by volume holographic data storage are due in large part to the arrangement of data into large pixelated pages [10]. But to retrieve any stored data, the pixel array imposed by the input spatial light modulator (SLM) must be accurately delivered to the array of detector pixels. Displacement of individual pixel images away from their target detector pixels (because of magnification error, misalignment, optical distortion, or material shrinkage [16]) quickly leads to uncorrectable levels of error.

Linear signal processing techniques have had success deblurring conventional 1–D channels found in modern communication and data storage devices. However, these techniques are not directly applicable to holographic storage, because the optical detection in holographic storage is inherently nonlinear. With coherent light, the intermingling of signal from neighboring pixels takes place in the amplitude domain [27], yet the signals accessible to a post–processing algorithm are the spatial integrals of the square of this optical amplitude (i.e., intensity) over the area of individual pixels [21]. Because of the spatial integration, one cannot simply take the square root to return to incident amplitude. In addition to pixel blur, once data pages reach a megapel (1024×1024 pixel) in size, a significant portion of the SNR budget is consumed simply by residual optical distortion: one portion of the page must always be partly misaligned in order to bring another portion into optimal alignment [1, 15]. To jointly address these two

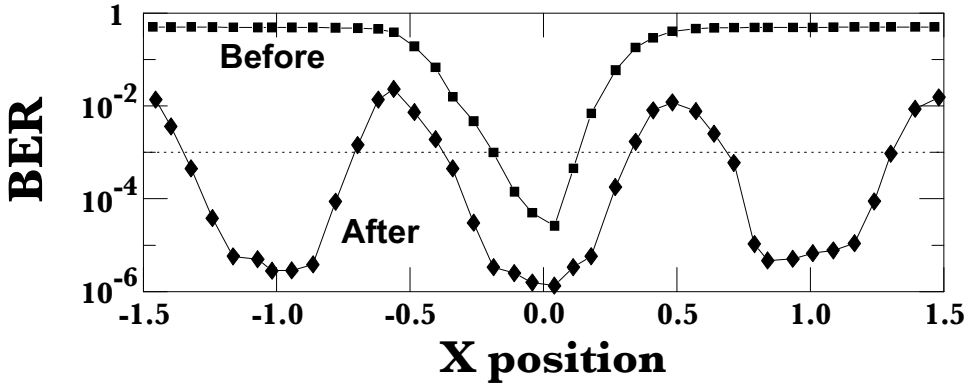

Figure 7. Dependence of raw–BER (with an 8:12 modulation code [17]) before and after shift–compensation post–processing, as a function of x shift [3].

problems, we have recently derived an algorithm that can compensate for both optical distortion and misalignment, correcting a moderate pixel blur in the presence of a significant pixel offset [3].

The algorithm comes directly from the detection physics in holographic data storage. Consider the readout signal received at a detector pixel when the incoming data page is shifted such that two SLM pixel images are contributing (the correct SLM pixel, and one neighbor). In the simplest 1–D case, we can decompose this detected signal r_2 into linear contributions from the two SLM pixel intensities p_1, p_2, and a nonlinear factor through their constructive interference, as

$$r_2 = p_2 H_{00}(\sigma) + 2\sqrt{p_1 p_2} H_{01}(\sigma) + p_1 H_{11}(\sigma). \quad (1)$$

The weights $H_{00}(\sigma), H_{11}(\sigma)$ and $H_{01}(\sigma)$ represent the normalized signal integrated by the detector pixel from the correct SLM pixel alone, the signal from the neighboring SLM pixel alone, and the additional contribution when both SLM pixels are present, respectively. This equation can be inverted to iteratively to solve for each pixel's correct value (p_2) from the just–processed neighbor pixel (p_1) and the received data signal (r_2). At each pixel, we take the measured signal, subtract the portion that belonged to the previous pixel, subtract a further portion due to interference, and then add in the missing signal that should have been here but which actually fell into the next pixel. Because the point–spread function in the DEMON II system [1] is dominated by diffraction effects, the 2–D data page can be processed by using this simple 1–D algorithm repeatedly, first on all the rows, and then on the columns [3].

In Figure 6(a), we show a small 9×9 pixel block as it should ideally be received. Figure 6(b) shows the same pattern imaged through DEMON II when the SLM is shifted a half–pixel in both x and y. The DEMON II platform pixel–matches megapel pages through an aperture of 1.36 D_N (1.7 x 1.7mm^2 aperture, f=30mm, λ=532nm, SLM δ= 12.8 microns). In Figure 6(c), we show this data after post–processing with the shift–compensation algorithm: the original pixel pattern is effectively recovered [3]. To process each block of pixels, the algorithm combines the dynamic global shifts, as measured by dedicated fiducial marks, with the static local baseline offsets taken from a lookup table [3]. Assuming that 10^{-3} is the maximum acceptable raw bit–error–rate (BER) that can be corrected by error–correction codes of moderate overhead [1], the shift–compensation algorithm increases the position tolerance of the DEMON II tester from ±16% to ±40% of the pixel pitch [3]. This is demonstrated experimentally (with images) in Figure 7.

4. PHASE–CONJUGATE READOUT

The success of the DEMON II platform is due jointly to excellent imaging fidelity (pixelated data arrives at the right detectors) and tight focusing of the object beam (holograms can be stored using very little volume). These two features were made possible by optical design: minimizing the optical aberrations, particularly distortion, of the short focal length optics. To scale fast–access holographic storage to high capacity, however, this same high density must be achieved at many storage locations, and without moving the storage media. The correspondingly greater demands on optical imaging performance soon limit the capacity achievable along this path to commercially uninteresting

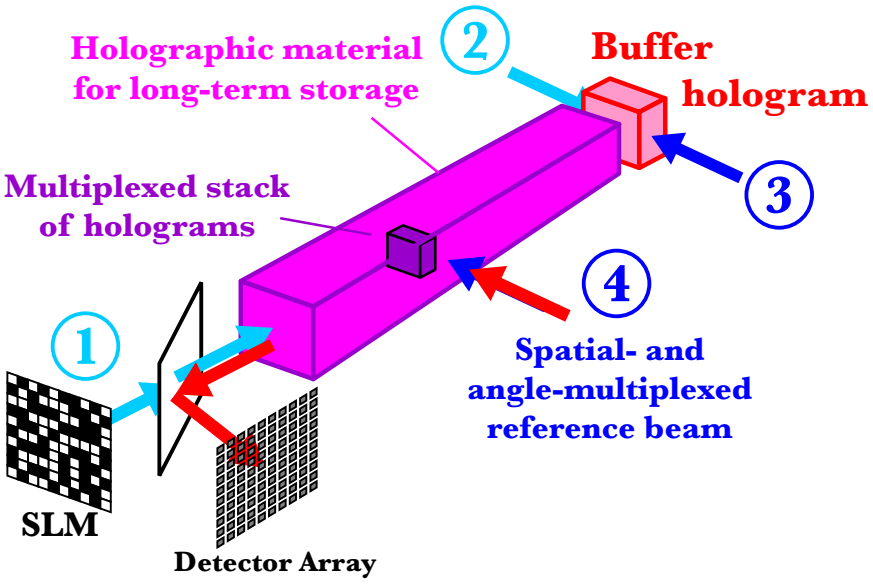

Figure 8. Phase–conjugate holographic storage system using a buffer hologram A temporary buffer hologram is recorded by an object beam containing the data from the SLM (1) and a reference beam (2). This hologram is illuminated with a phase–conjugate beam (3), reconstructing the phase–conjugate of the original object beam, which is then stored permanently with a spatial– and angle–multiplexed reference beam (4) [2].

limits. However, several researchers have long proposed bypassing these imaging constraints with "phase–conjugate" readout.

Once the light from the spatial light modulator has been recorded using a reference beam, the resulting hologram can be reconstructed with a phase–conjugate or "time-reversed" copy of the original reference beam. The wavefront diffracted by the phase–conjugate readout beam then retraces the path of the incoming object beam in reverse, canceling out any accumulated phase errors from lens aberrations or material imperfections. This allows data pages to be retrieved with high fidelity using an inexpensive lens [28, 29], or even without imaging lenses for an extremely compact system [30, 31].

However, two uncertainties prevented earlier work from proceeding. First, researchers were worried that imperfections in producing the phase–conjugate reference beam would introduce errors in each retrieved data page. In the simplest case, the phase–conjugate reference beam is a separate light beam that is carefully aligned to propagate in exactly the opposite direction to the original reference beam. However, even minor differences between the two beams will distort the reconstructed data pages. Alternatively, an extremely accurate phase–conjugate beam can be produced by a self–pumped phase–conjugate mirror [32]. Recently, we showed that a beam reflected from a phase–conjugate mirror could retrieve pages containing one million pixels onto a camera with very few bit errors, providing at least one potential solution to this first uncertainty [33].

The second concern was that many pairs of phase–conjugate reference beams would be needed to read the many different holograms recorded within the same volume—-and maintaining these beams over long periods of time would be impossible from a practical point of view. This problem also kept researchers from using the phase–conjugate mirror or PCM, since the barium titanate crystal takes some time to respond when the input beam changes.

4.1. Buffer hologram

To solve this problem, we have proposed [2]—and are currently completing the evaluation of—a novel architecture that allows phase–conjugation and multiplexed holographic storage to co–exist. The technique involves separating the phase–conjugation and hologram storage processes into two successive steps by using a 'buffer' hologram [2], as shown in Figure 8. Data to be recorded are modulated onto the object beam (1) with the SLM and focused into a long storage crystal. The object beam travels down the crystal, confined by total internal reflection, and passes into a buffer crystal, where it interferes with beam (2) and records a hologram. This hologram is then immediately read

with beam (3), the phase–conjugate of beam (2), reconstructing a phase–conjugate object beam which travels back into the storage crystal. This new object beam can now be recorded, and then later reconstructed, with beam (4) at one of the storage locations. This permits all the same angle–, phase–code–, wavelength–, and spatial–multiplexing approaches used for conventional volume holograms [6, 8].

Holograms can then easily be multiplexed at a large number of separate storage locations using only one SLM and one detector array. The buffer hologram technique still requires a pair of phase–conjugate reference beams, although only a single pair which do not need to ever change angle. Thus the self–pumped phase–conjugate mirror can be used, providing accurate phase–conjugation and adaptation to system misalignments without requiring the storage system to repeatedly wait for the PCM reflectivity to build up. Alternatively, the single phase–conjugate reference beam needed for the buffer hologram could be generated by careful alignment of a counter–propagating beam. This might have advantages over the self–pumped PCM, for example when implementing the buffer hologram in a wavelength–multiplexed system or when high modulation depth is required while recording the buffer hologram. Other advantages of the buffer hologram system include the ability to align the detector array to the SLM through the buffer hologram without sacrificing storage capacity. Finally, the strong buffer hologram can be monitored during recording by diverting, at the beamsplitter, a small portion of the returning phase–conjugate object beam to the detector while continuing to pass most of the object beam power from the SLM to the media.

4.2. Materials requirements

The long–term storage material in a phase–conjugate/buffer hologram system must provide highly sensitive, non–volatile storage. One candidate is two–color, gated volume holography in $LiNbO_3$ (see next section). The object and reference beams use long wavelength light (say, red or IR) which the crystal only absorbs in the presence of short wavelength gating light (green). The storage crystal can then be made extremely long, the gating light used to activate storage locations, and stored holograms read out without erasure.

Since this technique for multiplexed phase–conjugate holograms records two holograms for each stored data page, it would seem to inherently slow down the recording process. However, once the diffraction efficiency of the buffer hologram exceeds the power efficiency of the original object beam (typically ∼1–10%), then the recording of the storage hologram is actually accelerated. The buffer hologram does require a dynamic holographic material with high sensitivity, so that each new data page completely and rapidly overwrites the previous one. However, dynamic range for multiple holograms, dark storage lifetime, hologram thickness, and optical quality are less important, and could be traded off during material optimization for more sensitivity. Low scattering and uniform spatial frequency response are still needed, however, and the material must be able to tolerate a large number of read–write–erase cycle without degradation. There are several read–write materials whose strengths and weaknesses fit well here, including photorefractive polymers [34,35], bacteriorhodopsin [36], and fast photorefractive materials such as strontium barium niobate (SBN).

Erasure of the buffer hologram is preferably induced externally, either with electric field or an incoherent erase beam. This is particularly true when using the self–pumped PCM, because while the weak return beam is being used at the buffer material, the strong beam must also be present. If both beams are present during recording, then the modulation depth at the buffer material is greatly reduced. When both are present during readout and the material is not gated, then the erasure of the buffer hologram is greatly accelerated. If the buffer material is gated, then the PCM can be used upon readout and the strong pump beam safely ignored. For an ungated buffer material, however, adding buffer readout power—to get a stronger phase–conjugate object beam for transfer into the long–term hologram—is actually disadvantageous. The benefits of an initially stronger object beam are outweighed by the faster buffer erasure, so that increases in buffer readout power end up reducing the amount of object beam *energy* reaching the long–term hologram. The best option, in the absence of a gated buffer material, is to write a strong hologram and read it with as weak a readout beam as is feasible. It is also important to consider the impact of the erasure induced in the long–term storage material during transfer from the buffer, and the need to transfer holograms of varying strength within a recording schedule.

4.3. DEMON III test platform

We have built yet another test platform, called DEMON III, to evaluate the phase–conjugate/buffer hologram technique, using conventional, one–color lithium niobate. The optical layout is shown in Figure 9. A beam from a Coherent DPSS 532nm laser is expanded and split with various beam expansion optics to create three beams: an object beam for illuminating the SLM, a reference beam for conventional spatial– and angle–multiplexing, and a

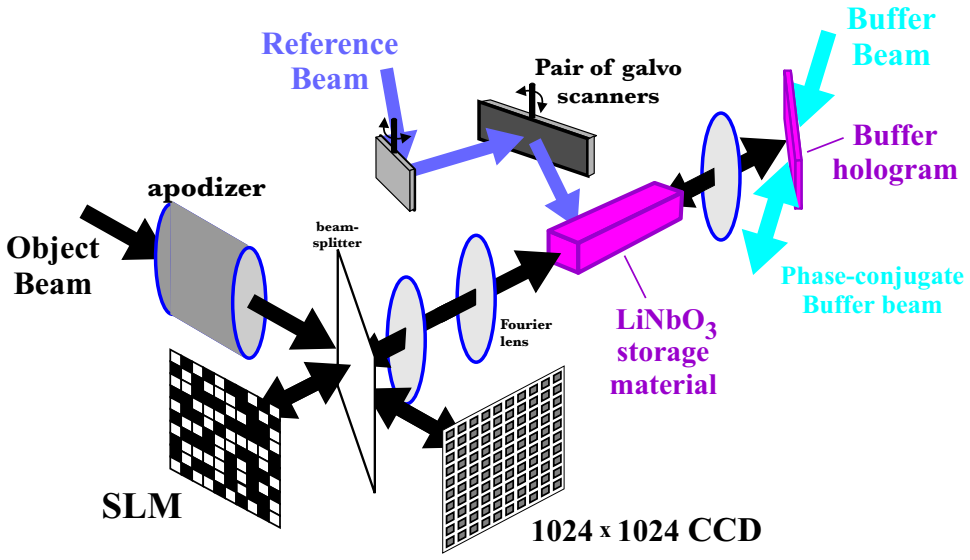

Figure 9. DEMON III holographic storage platform, for testing the use of a buffer hologram with multiplexed phase–conjugate readout.

third beam for writing the buffer hologram. The SLM and CCD are identical to those in the DEMON II platform. In DEMON III, however, a 1:1 relay lens delivers the SLM image to the input plane of a short–focal length Fourier transform lens, which then focuses the object beam into the $2\times2\text{mm}^2$ aperture of a long $LiNbO_3$:Fe bar (cut for the 90° geometry). The object beam expands out of the back end of the $LiNbO_3$ onto the buffer hologram material. The phase–conjugate reference beam for the buffer hologram can either be provided by a self–pumped phase–conjugate mirror in $BaTiO_3$ [2, 32], or by a second beam aligned to be counter–propagating. After the object beam is phase–conjugated, it can be stored in the long $LiNbO_3$ bar with the reference beam. A pair of galvo mirrors provides control over beam position along the bar as well as horizontal incidence angle.

To date, we have demonstrated pixel–matching of megapel pages in this platform through a simple biconvex lens of low F/#, despite significant aberrations in the image before phase–conjugation. The buffer hologram is written then immediately read by a pair of beams carefully aligned to be phase–conjugate, simply by maximizing the SNR of the resulting holograms with a random phase mask in the object beam. These experiments, performed without the lithium niobate crystal in place, have used photorefractive polymer material [37], bacteriorhodopsin [38], and SBN as the buffer hologram material. The remaining experimental difficulty is the low diffraction efficiencies attained with all three of these materials, approximately 0.1% at best, which is not quite sufficient for transfer of these phase–conjugate object beams into the long–term $LiNbO_3$:Fe crystals.

4.4. Extended shift–compensation algorithm

The shift–compensation procedure described above can relax the tight constraints on page registration, optical distortion, and material shrinkage that currently hamper page–oriented holographic storage systems. Since a shift of an integral number of pixels requires only careful bookkeeping, improvement of this algorithm to tolerance of a ±50% pixel shift would imply that the data could be retrieved as long as it falls somewhere on the detector array. This would provide many exciting opportunities to push system design (disk rotation rate), to relax mechanical constraints (disk wobble or readout pulse timing tolerances), and to tolerate material imperfections (shrinkage and expansion, whether induced thermally or optically).

Recently we have modified the shift–compensation algorithm so that it can successfully correct for arbitrary page misalignments. This modification requires that the SLM and CCD pattern have different magnifications, which we obtain for DEMON III platform by the simple expedient of removing the magnification optics. The SLM pattern, with 12.8μm pixel pitch, is then reconstructed by the phase–conjugate readout onto the CCD camera of 12μm pitch (6.25% magnification error). Lines that are 15 pixels apart on the SLM can be perfectly captured by CCD pixel columns that are 16 pixels apart. Everything in–between must be interpolated by the shift–compensation algorithm.

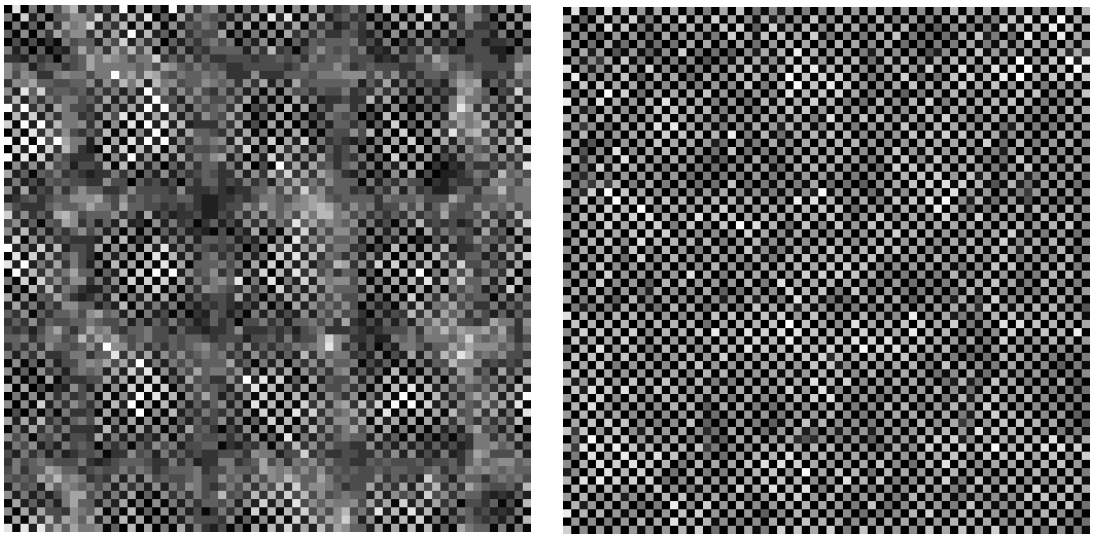

Figure 10. 1 × 1 pixel chessboard pattern (a) after phase–conjugate reconstruction onto CCD detector with 6.25% magnification error (slips half a pixel every 8 pixels), and (b) after post–processing with modified shift–compensation algorithm.

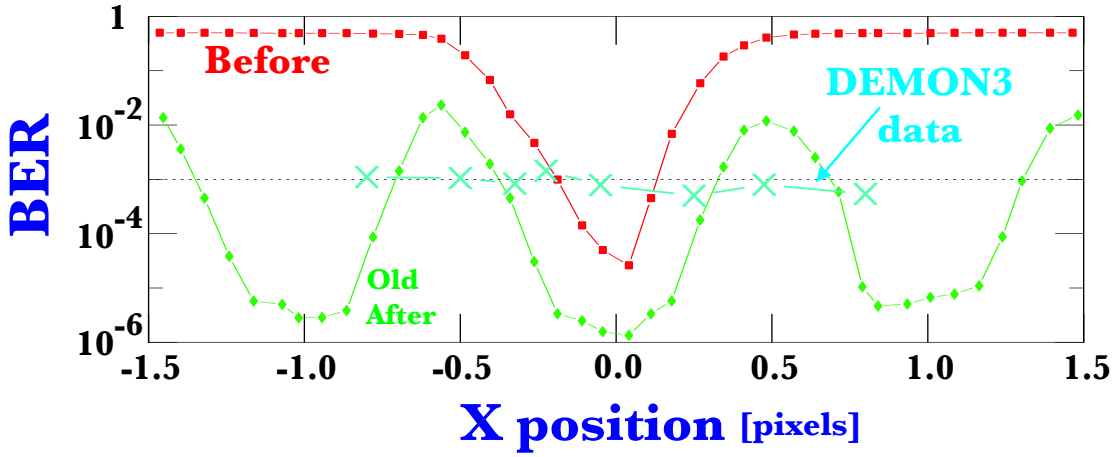

Figure 11. BER (with 8:12 modulation code) as a function of page misalignment without shift–compensation, with the original shift–compensation algorithm, and with the modified shift–compensation algorithm. With the original algorithm, the poor performance at a half–pixel shift means that the page-wide BER becomes unacceptable at these shifts. With the magnification error that the modified algorithm corrects, some portions of the page are at half–pixel shifts while others are aligned directly on pixels and the BER does not change with global page shift.

The bookkeeping is somewhat intricated, since the signal received at pixel 200,200 is now no longer associated with the SLM pattern that was transmitted at pixel 200,200, and all the data must be correctly moved to the appropriate pixel location in order to allow decoding and BER analysis.

Figure 10 shows a portion of a 1×1 pixel chessboard pattern, first as it is received by the CCD camera (part (a)), and then after correction with the modified shift–compensation algorithm. In Figure 11 shows BER as a function of shift for pages imaged on DEMON II without shift compensation, pages imaged on DEMON II with the original shift compensation algorithm, and holograms reconstructed with the modified shift compensation algorithm on DEMON III with the 6.25% magnification error described above. The presence of the magnification error means that all possible local shifts are represented on each data page, independent of global shift. Thus the modified algorithm provides true invariance to page misalignment, magnification, and optical distortion. Although isolated

pages with BERs as low as 2×10^{-5} have been retrieved in the presence of this magnification error, we continue to work to improve the robustness and decrease the BER that this modified shift–compensation algorithm can produce.

We anticipate that the successful use of phase–conjugation and shift–compensation in holographic storage will enable compact and affordable high–capacity systems, with only a moderate increase in the overall system complexity. However, such systems will require a recording material that supports both read–write access and non–volatile storage.

5. NON–VOLATILE, READ–WRITE HOLOGRAPHIC STORAGE MEDIA

Photorefractive media—such as lithium niobate, strontium barium niobate, and barium titanate—have long been developed for holographic storage [11]. These materials react to the light and dark regions of an interference pattern by transporting and trapping electrons, which subsequently leads to a local change in index of refraction. The trapped charge can be rearranged by later illumination, so it is possible to erase recorded holograms and replace them with new ones. The ability of a photorefractive material to erase through charge re-excitation also results in the undesired erasure of stored holograms during normal readout, and even gradual erasure in the dark through thermal excitation. Recorded holograms can be "fixed"—i.e. made semi–permanent and resistant to erasure during readout—by separate thermal or electronic processes. However, this fixing process affects all the stored holograms within a volume simultaneously, and tends to be slow and cumbersome.

An alternate method for achieving non–volatile storage in photorefractive materials is by recording at a wavelength of light which is only absorbed by the crystal in the presence of a third "gating" beam of different wavelength. This third beam is present only during recording and is switched off while the information is read out, allowing the data to be retrieved without erasure. This low absorption also allows holograms to be read out *through* other storage locations, greatly extending the potential capacity of the phase–conjugate buffer system described above. Conventional photorefractive materials can be optimized for this gated, two–color recording process by changing the ratio of lithium to niobium in the compound [5, 39] or by doping crystals with two dopants, such as manganese and iron [40, 41].

Conventionally, lithium niobate is grown in the congruent melting composition, expressed by the quantity $c_{\mathrm{Li}} \equiv$[Li]/([Li]+[Nb]) = 48.5%, because the identical compositions of the melt and the crystal promote high optical quality and large boules. Reduced stoichiometric lithium niobate shows both one–color sensitivity in the blue–green spectral region as well as two–color sensitivity for writing in the near IR and gating with blue–green light [5, 42–44]. From this it can be seen that the gating light also produces erasure. This is a consequence of the broad spectral features of reduced or Fe–doped lithium niobate. Considerable progress is envisaged if a better separation of gating and erasing functions can be achieved by storing information in deeper traps and/or using wider bandgap materials. Figure 12 compares one–color and two–color writing in a sample of reduced, near–stoichiometric lithium niobate to illustrate the nondestructive readout that can be achieved. The gating ratio, the increase in sensitivity of the material in the presence and absence of gating light, was in excess of 5000 [5]. In general, the higher the gating ratio, the better the persistence of the stored holograms during readout.

The most important photorefractive properties for two–color holographic data storage are gating ratio (which measures the degree of non–volatility), sensitivity, $M\#$ (which describes the dynamic range available for multiple holograms), dark decay, and optical quality. Photorefractive sensitivity for two–color recording in lithium niobate is linear in the gating light intensity, I_g, only at low values of I_g because of competition between gating and erasing. Hence the sensitivity in terms of incident intensities S_{η_2} can be defined in a similar way to that for one–color processes as

$$S_{\eta_2} = \frac{\sqrt{\eta}}{I\ l\ t} \tag{2}$$

but for a fixed and reasonably low value of $I_g = 1$ W/cm^2. Here I is the total intensity, l is the medium thickness, and t is the exposure time; this form of sensitivity is usually given in units of cm/J. The sensitivity in terms of absorbed power, $S_{\eta_1} = S_{\eta_2}/\alpha$ where α is the absorption coefficient at the writing wavelength, reveals an interesting way to compare two–color and conventional one–color materials. In terms of this sensitivity, all the samples we have studied [5], including the single photon Fe–doped material written at 488nm, are almost equally sensitive. This suggests that the sensitivity is determined by the amount of light that can be absorbed at the writing wavelength. So far, the maximum absorption of writing light that we have found in reduced SLN is 6% for $I_g = 1$ W/cm^2.

There is an interesting difference in the behavior of one- and two-color materials with regard to dynamic range. In a one-color material, the M# is proportional to the modulation index or fringe visibility of the optical interference

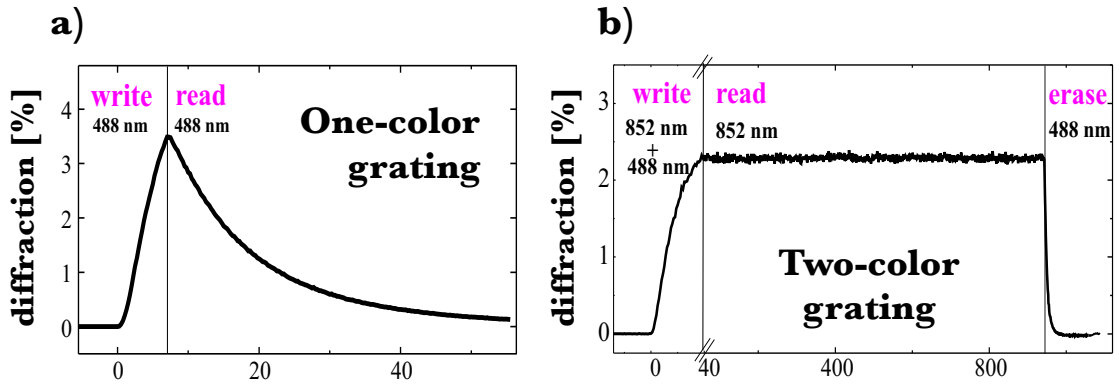

Figure 12. Typical write–read–erase curve for holographic gratings in LiNbO3 crystals. (a) One–color scheme, in which an argon ion laser at 488nm, 1W/cm^2 is used for both writing (two beams) and reading(one beam). (b) Two–color scheme, in which a laser diode at 852nm (4W/cm^2 total intensity) is used for writing and an argon ion laser at 488nm, 1W/cm^2 is used for the gating step. Nondestructive reading was done with one of the unattenuated writing beams (2W/cm^2) and erasing with the gating light [5].

pattern, $m = 2\sqrt{I_1 I_2}/(I_1 + I_2)$. However, in a two-color material, the writing light $(I_1 + I_2)$ does not erase the hologram and the M# is proportional to $\sqrt{I_1 I_2}$. As a result, for object and reference beams of equal intensity, the M# is proportional to the writing intensity. While this provides a general way of increasing the dynamic range in a two-color material, the writing power requirements in the present material system become rather high in order to achieve a substantial increase in M#.

Instead of amplifying the role of the intrinsic shallow levels with stoichiometry, an alternative scheme for implementing two–color holography in lithium niobate is the introduction of two impurity dopants [40, 41]. One trap, such as Mn, serves as the deep trap from which gating occurs, while a more shallow trap such as Fe, provides the more shallow intermediate level for gated recording. While this scheme provides more opportunities for tuning through choice of dopants, in general it is difficult in $LiNbO_3$ to separate the two absorption bands enough to provide high gating ratios and thus truly non–volatile storage. In addition, while $M\#$ improves monotonically with writing intensity for stoichiometric lithium niobate, with the two–trap method, $M\#$ is maximized at a particular writing intensity, thus creating the potential for an undesirable tradeoff between recording rate and dynamic range.

Gated, two–color photorefractive materials have received much attention over the last three years. These studies have led to improvements in both the sensitivity and dynamic range of the materials, increasing both the speed with which data can be written and the capacity. Further improvements, however, are still needed before prototype read–write holographic storage systems can be built.

6. CONCLUSIONS AND OUTLOOK

Holographic data storage has several characteristics that are unlike those of any other existing storage technologies. Most exciting, of course, is the potential for data densities and data transfer rates substantially exceeding those of magnetic data storage. In addition, as in all other optical data storage methods, the density increases rapidly with decreasing laser wavelength. In contrast to surface storage techniques such as CD–ROM, where the density is inversely proportional to the square of the wavelength, holography is a volumetric technique, making its density proportional to one over the third power of the wavelength. In principle, laser beams can be moved with no mechanical components, allowing access times of the order of 10 microseconds, faster than any conventional disk drive will ever be able to randomly access data.

The research efforts of the last few years have demonstrated that holographic storage systems with desirable properties can be engineered and built in the laboratory. However, existing and other developing storage technologies also continue to evolve and improve at a tremendous pace, making the next few years crucial for holographic storage. The next steps are to optimize the storage media, to demonstrate these systems outside the laboratory environment, and to design and build systems that are cost–competitive with existing technologies.

REFERENCES

1. 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), pp. 444–446, 2001.
2. G. W. Burr and I. Leyva, "Multiplexed phase-conjugate holographic data storage using a buffer hologram," *Optics Letters* **25**(7), pp. 499–501, 2000.
3. G. W. Burr and T. Weiss, "Compensation of pixel misregistration in volume holographic data storage," *Optics Letters* **26**(8), pp. 542–544, 2001.
4. J. A. Hoffnagle and C. M. Jefferson, "Design and performance of a refractive optical system that converts a Gaussian to a flattop beam," *Applied Optics* **39**(30), pp. 5488–5499, 2000.
5. H. Guenther, R. M. Macfarlane, Y. Furukawa, K. Kitamura, and R. R. Neurgaonkar, "Two–color holography in reduced near–stoichiometric lithium niobate," *Applied Optics* **37**, p. 7611, 1998.
6. D. Psaltis and F. Mok, "Holographic memories," *Scientific American* **273**(5), p. 70, 1995.
7. J. F. Heanue, M. C. Bashaw, and L. Hesselink, "Volume holographic storage and retrieval of digital data," *Science* **265**, p. 749, 1994.
8. J. H. Hong, I. McMichael, T. Y. Chang, W. Christian, and E. G. Paek, "Volume holographic memory systems: techniques and architectures," *Optical Engineering* **34**, pp. 2193–2203, 1995.
9. D. Psaltis and G. W. Burr, "Holographic data storage," *IEEE Computer* **31**(2), pp. 52–60, 1998.
10. J. Ashley, M.-P. Bernal, G. W. Burr, H. Coufal, H. Guenther, J. A. Hoffnagle, C. M. Jefferson, B. Marcus, R. M. Macfarlane, R. M. Shelby, and G. T. Sincerbox, "Holographic data storage," *IBM J. Research and Development* **44**, pp. 341–368, May 2000.
11. H. J. Coufal, D. Psaltis, and G. Sincerbox, eds., *Holographic Data Storage*, Springer–Verlag, 2000.
12. P. J. van Heerden, "Theory of optical information storage in solids," *Applied Optics* **2**(4), pp. 393–401, 1963.
13. G. W. Burr, *Volume holographic storage using the 90° geometry*. PhD thesis, California Institute of Technology, Pasadena, Calif., 1996.
14. M.-P. Bernal, H. Coufal, R. K. Grygier, J. A. Hoffnagle, C. M. Jefferson, R. M. Macfarlane, R. M. Shelby, G. T. Sincerbox, P. Wimmer, and G. Wittmann, "A precision tester for studies of holographic optical storage materials and recording physics," *Applied Optics* **35**(14), pp. 2360–2374, 1996.
15. R. M. Shelby, J. A. Hoffnagle, G. W. Burr, C. M. Jefferson, M.-P. Bernal, H. Coufal, R. K. Grygier, H. G. . Günther, R. M. Macfarlane, and G. T. Sincerbox, "Pixel–matched holographic data storage with megabit pages," *Optics Letters* **22**(19), pp. 1509–1511, 1997.
16. R. M. Shelby, D. A. Waldman, and R. T. Ingwall, "Distortions in pixel-matched holographic data storage due to lateral dimensional change of photopolymer storage media," *Optics Letters* **25**(10), pp. 713–715, 2000.
17. G. W. Burr, J. Ashley, H. Coufal, R. K. Grygier, J. A. Hoffnagle, C. M. Jefferson, and B. Marcus, "Modulation coding for pixel–matched holographic data storage," *Optics Letters* **22**(9), pp. 639–641, 1997.
18. G. W. Burr, H. Coufal, R. K. Grygier, J. A. Hoffnagle, and C. M. Jefferson, "Noise reduction of page–oriented data storage by inverse filtering during recording," *Optics Letters* **23**(4), pp. 289–291, 1998.
19. G. W. Burr, G. Barking, H. Coufal, J. A. Hoffnagle, C. M. Jefferson, and M. A. Neifeld, "Gray–scale data pages for digital holographic data storage," *Optics Letters* **23**(15), pp. 1218–1220, 1998.
20. G. W. Burr, W.-C. Chou, M. A. Neifeld, H. Coufal, J. A. Hoffnagle, and C. M. Jefferson, "Experimental evaluation of user capacity in holographic data storage systems," *Applied Optics* **37**(23), pp. 5431–5443, 1998.
21. M.-P. Bernal, G. W. Burr, H. Coufal, and M. Quintanilla, "Balancing inter–pixel crosstalk and thermal noise to optimize areal density in holographic storage systems," *Applied Optics* **37**(23), pp. 5377—5385, 1998.
22. J. L. Sanford, P. F. Greier, K. H. Yang, M. Lu, R. S. Olyha, Jr., C. Narayan, J. A. Hoffnagle, P. M. Alt, and R. L. Melcher, "A one–megapixel reflective spatial light modulator system for holographic storage," *IBM Journal of Research and Development* **42**, pp. 411–426, May/July 1998.
23. D. Psaltis, D. Brady, and K. Wagner, "Adaptive optical networks using photorefractive crystals," *Applied Optics* **27**(9), pp. 1752–1759, 1988.
24. F. H. Mok, G. W. Burr, and D. Psaltis, "System metric for holographic memory systems," *Optics Letters* **21**(12), pp. 896–898, 1996.
25. A. M. Glass, D. von der Linde, and T. J. Negram, "High–voltage bulk photovoltaic effect and the photorefractive process in $LiNbO_3$," *Applied Physics Letters* **25**(4), pp. 233–235, 1974.

26. S. S. Orlov, E. Bjornson, W. Philips, Y. Takashima, X. Li, L. Hesselink, R. Okas, and R. Snyder, "High transfer rate (1 Gbit/sec) high–capacity holographic disk digital data storage system," in *Conference on Lasers and Electro–optics*, OSA Technical Digest, pp. 190–191, Optical Society of America, (Washington, DC), 2000.
27. V. Vadde and B. V. K. Vijaya Kumar, "Channel modeling and estimation for intra–page equalization in pixel–matched holographic data storage," *Applied Optics* **38**(20), pp. 4374–4386, 1999.
28. F. Ito, K.-I. Kitayama, and H. Oguri, "Compensation of fiber holographic image distortion caused by intrasignal photorefractive coupling by using a phase–conjugate mirror," *Optics Letters* **17**(3), pp. 215–217, 1992.
29. M. C. Bashaw, A. Aharoni, and L. Hesselink, "Phase–conjugate replay for a–axis strontium barium niobate single–crystal fibers," *Optics Letters* **18**(23), pp. 2059–2061, 1993.
30. F. Zhao and K. Sayano, "Compact read–only memory with lensless phase–conjugate holograms," *Optics Letters* **21**(16), pp. 1295–1297, 1996.
31. J. J. P. Drolet, E. Chuang, G. Barbastathis, and D. Psaltis, "Compact, integrated dynamic holographic memory with refreshed holograms," *Optics Letters* **22**(8), pp. 552–554, 1997.
32. J. Feinberg, "Self–pumped, continuous–wave phase conjugator using internal reflection," *Optics Letters* **7**(10), pp. 486–488, 1982.
33. G. W. Burr and R. M. Shelby, "Pixel–matched phase–conjugate holographic data storage," *SPIE Optical Processing & Computing newsletter* , p. 8, Jan. 1999.
34. S. Ducharme, J. C. Scott, R. J. Twieg, and W. E. Moerner, "Photorefractive behavior in a polymer?," *Physics Review Letters* **66**, p. 1846, 1991.
35. E. Mecher, R. Bittner, C. Brauchle, and K. Meerholz, "Optimization of the recording scheme for fast holographic response in photorefractive polymers," *Synthetic Metals* **102**(1–3), pp. 993–996, 1999.
36. N. Hampp, C. Brauchle, and D. Oesterhelt, "Bacteriorhodopsin wildtype and variant aspartate–96 → asparagine as reversible holographic media," *Biophysics Journal* **58**, pp. 83–93, 1990.
37. 7DCST:PCBM:DPP:PVK in the ratio 25%:0.5%:29%:45.5%, 53μm thick, 53V/μm, supplied by Erwin Mecher, Francisco Gallego, and Klaus Meerholz, Ludwig Maximilians Universitat Munchen, Institut fur Physikalische Chemie, 81377 Munchen, Germany.
38. M–type holograms in $\sim$80μm thick material optimized for low light optical interferometry, supplied by Thorsten Juchem and Norbert Hampp, Institut fur Physikalische Chemie, Universitat Marburg D–35032 Marburg, Germany.
39. H. Guenther, G. Wittmann, R. M. Macfarlane, and R. R. Neurgaonkar, "Intensity dependence and white–light gating of two–color photorefractive gratings in $LiNbO_3$," *Optics Letters* **22**, pp. 1305–1307, 1997.
40. K. Buse, A. Adibi, and D. Psaltis, "Non-volatile holographic storage in doubly doped lithium niobate crystals," *Nature* **393**, pp. 665–668, 1998.
41. L. Hesselink, S. S. Orlov, A. Liu, A. Akella, D. Lande, , and R. R. Neurgaonkar, "Photorefractive materials for nonvolatile volume holographic data storage," *Science* **282**, pp. 1089–1094, 1998.
42. F. Jermann, M. Simon, and E. Kratzig, "Photorefractive properties of congruent and stoichiometric lithium niobate at high light intensities," *Journal of the Optical Society of America* **A 12**, pp. 2066–2070, 1995.
43. Y. S. Bai, R. R. Neurgaonkar, and R. Kachru, "High-efficiency nonvolatile holographic storage with two-step recording in praseodymium-doped lithium niobate by use of continuous-wave lasers," *Optics Letters* **22**, pp. 334–336, 1997.
44. D. Lande, S. S. Orlov, A. Akella, L. Hesselink, and R. R. Neurgaonkar, "Digital holographic storage system incorporating optical fixing," *Optics Letters* **22**, pp. 1722–1724, 1997.