

Experimental evaluation of user capacity in holographic data-storage systems

Geoffrey W. Burr, Wu-chun Chou, Mark A. Neifeld, Hans Coufal, John A. Hoffnagle, and C. Michael Jefferson

An experimental procedure for determining the relation between the number of stored holograms and the raw bit-error rate (BER) (the BER before error correction) of a holographic storage system is described. Compared with conventional recording schedules that equalize the diffraction efficiency, scheduling of recording exposures to achieve a uniform raw BER is shown to improve capacity. The experimentally obtained capacity versus the raw-BER scaling is used to study the effects of modulation and error-correction coding in holographic storage. The use of coding is shown to increase the number of holograms that can be stored; however, the redundancy associated with coding incurs a capacity cost per hologram. This trade-off is quantified, and an optimal working point for the overall system is identified. This procedure makes it possible to compare, under realistic conditions, system choices whose impact cannot be fully analyzed or simulated. Using LiNbO_3 in the 90° geometry, we implement this capacity-estimation procedure and compare several block-based modulation codes and thresholding techniques on the basis of total user capacity. © 1998 Optical Society of America

OCIS codes: 090.0090, 040.1520, 050.7330, 070.2580, 090.2900, 210.2860.

1. Introduction

Volume holography is a page-oriented data storage paradigm offering high density and fast parallel access.^{1,2} High-density storage arises from the superposition of multiple holograms within the same physical volume of a storage material (a stack of multiplexed holograms). The Bragg selectivity of volume holography makes possible independent readout of each of the stored data pages. Each reconstructed hologram creates a two-dimensional (2-D) intensity pattern of bright and dark pixels, which is detected in parallel by a CCD camera and converted to binary data. The fidelity of the retrieved binary data is an important measure of the overall storage-system performance. For holographic storage to become a competitive technology, acceptable goals for the bit-error

rate (BER) at the system output (user BER) must be met.

There are two contributions to the fidelity of retrieved data in holographic storage systems. The first contribution concerns the signal-to-noise ratio (SNR) performance of the optoelectronic hardware comprising the storage and the retrieval subsystems. The SNR's of these subsystems should be made as high as possible. This task is complicated, however, by the complexity of the noise and the interference environment associated with holographic storage. The second contribution to data fidelity concerns the interface coding and signal processing used to achieve an acceptable user BER from a given input SNR. Like conventional storage technologies, such as magnetic disk drives or optical disks, a digital holographic storage system incorporates redundancy in the form of modulation codes and error-correction coding³⁻⁵ (ECC). A portion $(1 - r)$ of each data page is given over for coding to achieve the extremely low user-BER specification expected of current data-storage devices (usually $<10^{-12}$). Although this redundancy is a sacrifice in the number of user bits per hologram, coding allows the system to operate at a lower SNR and thus store more holograms. As a result, the overall user capacity tends to increase through the use of coding, despite the redundancy.³ The portion of the data page left for user data is

G. W. Burr, H. Coufal, J. A. Hoffnagle, and C. Michael Jefferson are with the IBM Almaden Research Center, IBM Corporation, 650 Harry Road, San Jose, California 95120. W.-C. Chou and M. A. Neifeld are with the Optical Sciences Center, Department of Electrical Engineering, University of Arizona, Tucson, Arizona 85721.

Received 24 February 1998; revised manuscript received 18 May 1998.

0003-6935/98/235431-13\$15.00/0

© 1998 Optical Society of America

termed the code rate r . At the optimal working point of the system, the user capacity of each hologram stack (the product of the code rate, the number of multiplexed holograms, and the number of pixels per hologram) is maximized. This capacity per stack can be extended either to total system capacity (by multiplication by the number of stacks in the system) or to volumetric or areal density (by division by the volume or area per stack).

This paper describes a procedure for experimentally measuring the stack capacity in a holographic storage system. An experimental evaluation of the raw-BER behavior associated with holographic storage is combined with modulation codes and ECC and to determine the optimal amount of redundancy required and thus the maximal stack capacity. The experimental procedure we describe is useful as a tool for optimizing the performance of any functioning holographic storage system. Although the paper describes and demonstrates this procedure as it applies to systems that use photorefractive crystals, extensions to other types of storage material are described.

Previous research on system optimization by means of noise modeling has been limited to a few simple noise sources and has not included extensive experimental verification.⁵⁻⁷ In contrast, experimental demonstrations have listed results (capacity and BER) but have not attempted to optimize performance or to describe the scaling of the BER with capacity.^{2,4} In the present study a realistic relation between the raw BER and the number of stored holograms M is obtained experimentally in the presence of multiple noise sources. As a result, system choices that affect signal levels or the raw BER—such as phase masks, beam ratios, and coding and signal-processing techniques—can now be compared in terms of user capacity per stack. In addition, system choices that change the stack volume—such as lens focal length, aperture stop, pixel size, and material thickness—can be compared in terms of either total system capacity, volumetric density, or areal density.

In Section 2 we describe the various noise sources that are present in holographic storage systems and motivate the need for an experimental evaluation of the trade-off between the raw BER and the number of holograms. In Section 3 a modified recording schedule is described that strives to equalize the raw BER of the first and the last stored holograms. The capacity of such a procedure is compared with the conventional approach of equalizing the diffraction efficiency. The experimental measurements and data-extraction procedures are described in Section 4 and typical results are shown. The use of modulation codes and ECC and the resulting experimentally obtained user capacities per hologram stack are discussed in Section 5, resulting in identification of the optimal working point of the system. Sections 6 and 7 offer generalizations of the measurement procedure and conclusions, respectively.

2. Relating the Number of Stored Holograms and the Raw Bit-Error Rate

To optimize the capacity of a volume holographic storage system, it is important to understand the various noise contributions and their scaling with the number of stored holograms. It is generally true that the retrieved holographic data signals decrease and the various noise contributions increase as more holographic pages are added. The maximum stack capacity is reached when the next stored hologram would push the user BER over specification. The decoding algorithms that produce a low user BER from a poor SNR are then being used to their full capabilities: storing another hologram would cause a further decrease in the SNR, for which the codes would not be able to produce the desired user BER.

Consider a simple model in which only the finite dynamic range (a decrease in the diffraction efficiency per hologram in the presence of a constant noise floor) tends to decrease the SNR as additional holograms are superimposed. For instance, in the 90° geometry,¹ division of the dynamic range is the dominant SNR contribution, and the noise floor is mainly due to the detection electronics. In this simple noise model the SNR scales with only the signal strength (the number of detected photons in the holographic reconstruction). From the operation of the recording schedule in photorefractive materials, it is well known that signal strength falls as one over the square of the number of holograms M .⁸ It is then possible to combine this scaling of the SNR as a function of M with the SNR gain associated with the modulation codes and the ECC to determine the optimal code rate and capacity.³

Unfortunately, this simple picture is incomplete because there are other noise sources present besides the constant background noise. Some of these noise contributions scale linearly with the signal strength. If these signal-dependent noise sources were the only noise terms present, the SNR would remain constant no matter how many holograms were stored. A constant noise floor is still what causes the SNR to fall as the number of stored holograms rises. However, the presence of signal-dependent noise sources, even at modest levels, affects the point at which the falling signal levels come too close to the noise floor. If we assume that all noise sources are statistically independent and add in quadrature, we can write the SNR as

$$\text{SNR} = \frac{SM^{-2}}{N_1M^{-2} + N_{\text{floor}}}, \quad (1)$$

where S is the signal strength, N_1 is the signal-dependent noise, and N_{floor} represents the constant background-noise floor. For example, S could be the mean number of detected photons per pixel, N_1 the standard deviation of S owing to interpage cross talk,⁶ and N_{floor} the number of detector noise electrons, referenced through the quantum efficiency of the detector. When M is small, the fixed N_{floor} is unimportant, and the SNR remains at its initial

value $\text{SNR}_0 = S/N_1$. At large M , the noise floor term is dominant. The effect of the signal-dependent noise term N_1 is to bring the system to a given SNR target (say, SNR_t) at a lower number of holograms M_{max}' than would be obtained in the case of $N_1 = 0$ (M_{max}). Given the above assumptions, the number of pages that can be stored is reduced from M_{max} to

$$M_{\text{max}}' = M_{\text{max}} \left[1 - \left(\frac{\text{SNR}_t}{\text{SNR}_0} \right) \right]. \quad (2)$$

Because experimental systems typically operate with a rather low initial SNR (i.e., $\text{SNR}_0 \sim 4\text{SNR}_t$), significant capacity costs can be associated with these signal-dependent noise sources.

Interpage cross talk resulting from residual Bragg mismatch, interpixel cross talk owing to diffraction-misalignment-aberrations, and nonuniformity across the stored data page are all noise sources that scale linearly with the signal strength. Shot noise is a noise source that scales as the square root of the signal strength and can become important as the signal levels approach the detector-based noise floor. There are also noise sources that increase not with M but with the total optical exposure of the material. These include noise gratings, erasure effects, and the photovoltaic noise buildup empirically observed in $\text{LiNbO}_3:\text{Fe}$.⁹ Consideration of these contributions is particularly important for photorefractive materials, for which modifying the recording schedule to obtain 5% more diffraction efficiency can double the total exposure time.⁹ The term erasure effects refers to loss of fidelity attributed to the effects of absorption.^{10,11} The absorption creates variations in local intensity; thus holograms erase nonuniformly across the crystal volume. This can cause spatial variation in the hologram strength across the image (image-plane geometry) or a broadened point-spread function (Fourier plane geometry).

Up to this point, we have separated the holographic storage system into two stages: a hardware portion, which stores M holograms at a particular SNR, and a coding portion, which takes this SNR and outputs digital data at a particular user BER. However, the SNR is not a particularly convenient variable to measure, particularly in the presence of multiple noise sources. Formulas for the SNR tend to introduce assumptions about the statistics of the noise processes. In addition, algorithms for making binary decisions on the received data (such as local thresholding or modulation decoding) work on small areas of the data page, making it difficult to gather enough data to measure the local mean and variance accurately. In contrast, the raw BER at the output of the binary-decision process is easy to measure, involves no assumptions about noise models, and can be gathered over the whole data page. For obtaining the final system output, the performance of the ECC (from raw BER to user BER) can be analyzed mathematically in a straightforward manner (to avoid trying to measure user BER's of 10^{-12}). Although

dividing the system in the middle of the coding stage is not as pleasing intuitively as separation into hardware and coding, doing so minimizes the effect of the measurement process on the results.

It is relatively straightforward to model the effects of one or two of the noise sources we have described. However, the raw BER for M holograms stored in a real holographic storage device is due to all the factors listed above, combined in a manner not easily predicted or modeled. The experimental difficulties are exactly opposite: Suppressing all but a few noise sources is difficult, but measuring the sum effect of all the noise sources is straightforward. After the holographic storage system is functioning, one can just store M holograms and measure the average raw BER. Although this is straightforward, it will not indicate what raw BER would have been obtained had one tried 100 holograms more or 100 holograms fewer. So, although user capacity can be evaluated (at whatever raw BER comes out of the measurement), optimization becomes a matter of trial and error.

The above problem is similar to that faced in the past when trying to measure the $M/\#$ of a photorefractive crystal.⁹ The $M/\#$ is a system-material figure of merit for dynamic range and is defined as the coefficient of proportionality between the square root of the hologram diffraction efficiency η and the number of holograms⁹:

$$\eta = \frac{M/\#^2}{M^2}. \quad (3)$$

The $M/\#$, pronounced M number, is a useful metric for comparing material candidates both in saturable materials (such as photopolymers) and in read-write materials (such as photorefractives). In a photorefractive crystal, a consequence of the read-write ability is that recorded holograms are slowly erasing during subsequent recording exposures. To store multiple equal-strength holograms, one follows a schedule of decreasing exposure times.⁸ The first hologram is recorded to a strong diffraction efficiency, erases exponentially during the remaining exposures, and finishes at the same diffraction efficiency that the last hologram reaches with its short exposure time. (After recording is finished, the holograms would be fixed thermally or electrically, or the gating light that permitted recording would be removed.) The derivation of the recording schedule leads both to Eq. (3) and to a definition for the $M/\#$ (Ref. 9):

$$M/\# \equiv \left(\frac{A_0}{\tau_r} \right) \tau_e. \quad (4)$$

The $M/\#$ is the product of the recording slope A_0/τ_r and the erasure-time constant τ_e . A high dynamic range results when holograms record quickly and erase slowly. (Although A_0 , the square root of the saturated diffraction efficiency, and τ_r , the recording-




Fig. 1. Evolution of the diffraction efficiency (top) and the raw BER (bottom) for the first hologram in a recording schedule. The raw BER is the BER presented to the ECC, which must then correct down to the user-BER specification. The arrow on the raw-BER axis points in the direction of a lower (better) raw BER.

time constant, could be considered as separate variables, one needs only to measure their ratio.)

The terms A_0/τ_r in and τ_e in Eq. (4) can be measured from the recording and the erasure dynamics of a single hologram.⁹ Essentially, one follows the evolution of the diffraction efficiency η of the first hologram of the recording schedule. The measurement of the capacity that we describe in this paper is directly analogous to this $M/\#$ measurement: an experiment that follows the evolution of the raw BER of the first hologram. In addition, because the first hologram experiences different noise sources than does the last, we also measure the raw-BER properties of the last hologram.

Figure 1 represents the evolution in both diffraction efficiency (top) and raw BER (bottom) for the first hologram. As the hologram grows stronger during recording, the signal rises out of the constant noise floor and the raw BER improves rapidly (that is, gets smaller). The signal strength continues to increase during the recording exposure, but the raw BER saturates as the signal-dependent noise sources become dominant. During the recording of the remaining holograms, the first hologram loses diffraction efficiency. During this erasure, the raw BER also gets worse (larger). Initially, this increase is from loss of fidelity during erasure and is not tied directly to the signal strength. As the optical exposure of the material increases, photovoltaic noise and noise gratings can increase and affect both stored holograms and the fidelity of subsequently recorded holograms. This effect is particularly noticeable if the material is close

to the Fourier transform plane where the data-bearing object beam achieves a tight focus.¹² Finally, as the first hologram continues to erase, the signal strength plummets toward the constant noise floor and the raw BER climbs rapidly.

Knowing when this first hologram rises to the target raw BER (the maximum tolerable raw BER allowed by the error-correction algorithms) dictates the maximum duration of an acceptable exposure schedule. As we show below, this duration can be converted to a maximum number of stored holograms, allowing the raw BER and the capacity per stack to be related. However, this measurement of a single first hologram seems to suffer from the same type of limitation as the previously dismissed experiment (i.e., recording M holograms and measuring the raw BER). After a particular exposure time t_1 is used to record the first hologram, it would seem that we have information about only schedules with this first exposure time, and optimization would again be a matter of trial and error. In this case, however, it is possible to use a strong first hologram to test what *would* have happened with a shorter first exposure time. The key is in the way this test hologram is measured. In the remainder of this section we describe the appropriate measurement procedure. Then in Section 3 we describe how to use the measured data to evaluate the number of holograms that can be stored.

To start the experiment, we record the test hologram with a long exposure by using the same object- and reference-beam intensities that the system will be using for recording during normal operation. Before the test hologram is measured it must experience the same erasure exposure that would occur during the recording schedule. These conditions are created by the recording of other holograms for a total exposure time t_{meas} . The test hologram is then re-Bragg matched and measured multiple times. Between each measurement, the power in the Bragg matched readout beam is reduced. As long as there is no scatter noise, there is no difference between this strong hologram read by a weak readout beam and a weak hologram reconstructed by a high-power readout beam. In this experiment we measure the one strong test hologram repeatedly and infer what would happen with weak holograms.

As the amount of power used to read out the test hologram is reduced, the constant noise floor causes the raw BER to rise. Given a chosen raw-BER target (say, 10^{-3}), we can identify the ratio between the lowest acceptable readout power (that still achieves the raw-BER target) and the full readout power with which the system will normally operate. This is identical to the ratio between the diffraction efficiency of the weakest acceptable hologram and the diffraction efficiency of our test hologram. With this ratio, we can determine what first recording exposure t_1 *would* have been necessary to reach the target raw BER at the end of a recording schedule of length t_{meas} by use of all the available readout power. A similar measurement of a just-written hologram gives the

required exposure of the last hologram t_M . If effects such as coherent saturation change the raw-BER dynamics during the recording schedule,¹² this last hologram will be competing with more noise sources than were present when the first hologram was recorded; if not, one measurement suffices to describe both t_1 and t_M . Along with the erasure-time constant τ_e and the total exposure time since recording the first hologram t_{meas} , these two exposure times (for the first and the last holograms) can be used to determine a unique recording schedule.

Previously, recording schedules were used to equalize the diffraction efficiencies of multiple holograms.⁸ Here we propose a paradigm shift: a modified recording schedule that equalizes the raw BER of the holographic pages instead of the diffraction efficiency. Our initial purpose is to complete this experimental procedure for capacity estimation. The derivation of such a “flat-BER” recording schedule results in a unique value for M , the number of holograms that can be stored (given the constraints on total exposure time t_{meas} and the raw BER). However, at the end of Section 4 we also show that this new flat-BER schedule can also provide a significant capacity improvement over the conventional flat- η schedule.

To complete the capacity-evaluation experiment, separate recording schedules (and thus estimates of M) are computed for a range of target raw BER's and for a range of total recording times t_{meas} . Conveniently, all the required data are provided by the measured erasure dynamics of a single test hologram: the t_{meas} values come from measuring periodically during erasure and the target raw BER's from making each measurement with a range of readout powers. Each ordered pair (t_{meas} , raw-BER target) has its own pair of measured data (t_1 , t_M) and thus a unique value for the number of holograms that can be stored M . In Section 3 we describe the specific derivation of the flat-BER recording schedule, starting from measured values for t_1 , t_M , τ_e , and t_{meas} and ending by producing a value for M .

3. Recording Schedule for the Flat Bit-Error Rate

The measurement technique described in Section 2 produces three time variables: t_{meas} , t_1 , and t_M . The term t_{meas} is the exposure time between the recording of the first and the last holograms. The term t_1 is the exposure time for a first hologram such that, after t_{meas} seconds of erasure, this hologram has exactly the target raw BER. The term t_M is the exposure time needed to make the last hologram reach the target raw BER immediately after recording. The experimental procedure described above produces a reference-to-readout-power ratio: this is the factor (always ≥ 1.0) by which the reference power has to be reduced for the hologram to reach the target raw BER at readout. Because the diffraction efficiency increases quadratically during exposure, the corresponding change in the recording exposure is the square root of this factor. For example, suppose we have begun the experimental capacity measure-

ment by recording the first hologram for 30 s and then erased for some time t_{meas} . In making our measurements, we find that the reference beam must be reduced to one ninth of its original power for the hologram to rise to a raw BER of 10^{-3} . Then we know that, if we begin a schedule with a hologram of $t_1 = 30 \times \sqrt{1/9} = 10$ s and finish after t_{meas} seconds, the first hologram will be at the target raw BER of 10^{-3} with all the available readout power.

In a conventional recording schedule for equalized diffraction efficiency, one starts with a measured erasure-time constant τ_e and the desired number of stored holograms M . Starting from a chosen final exposure time t_M , the hologram exposures are calculated recursively back toward the first exposure t_1 . This final exposure involves a choice, expressed here as the factor X , where

$$t_M = \frac{\tau_e}{M + X}, \quad (5)$$

and X is always greater than zero. As X approaches zero, the recording schedule calls for an initial exposure t_1 that becomes infinitely large. Smaller values of X yield a better diffraction efficiency but also a longer total exposure time. In effect, one is drawing up an infinite recording schedule for $M + X$ holograms and then starting with the $(X + 1)$ th exposure (skipping the inconveniently long first exposures). This first exposure time can be written as⁸

$$t_1 = \tau_e \ln \frac{X + 1}{X}. \quad (6)$$

In our experimental procedure the measurements of the diffraction efficiency of the first hologram at different values of t_{meas} produce the exponential time constant τ_e . (This is the exponential decay constant of the amplitude $\sqrt{\eta}$ of the first hologram.) Knowing the appropriate value of X as constrained by our knowledge of t_1 [Eq. (6)], we can solve for the number of holograms M as

$$M = \left(\frac{\tau_e}{t_M} - X \right), \quad (7)$$

where M is rounded down to an integer. So the usual flat- η recording schedule begins with a known τ_e , a desired M , and a chosen value of X , and the result is t_1 . In a flat-BER schedule, we know τ_e and have a desired t_1 and t_M , and the result of the recording schedule is M , the number of holograms one can store.

This recording schedule will start and end with the two desired exposure times, but there is no guarantee that the total exposure time t_{total} produced by this schedule is the same as the t_{meas} used in the raw-BER experiment. Because the measurement indicates only the raw BER after exactly t_{meas} of the total post-recording exposure, it is fairly important to force the schedule to match t_{meas} . The total exposure time of the schedule can be found from summing the individ-

ual exposure times by backward recursion.⁹ The m th exposure time is calculated from the $(m + 1)$ th exposure as

$$t_m = t_{m+1} \exp(t_{m+1}/\tau_e), \quad (8)$$

where one starts from the final (the M th) exposure t_M , as given by Eq. (5), and works backward to t_1 .

If the summed $t_{\text{total}} = t_{\text{meas}}$, this recording schedule will satisfy both criteria: the first and the last holograms will be of equal diffraction efficiency and both at the target raw BER. If $t_{\text{total}} > t_{\text{meas}}$, however, then after we use this recording schedule the first hologram will be more than (worse than) the target raw BER (because it would have been erased more than in the t_{meas} experiment). So even though the first and the last holograms would have equal diffraction efficiency, the first hologram would have a higher raw BER. As mentioned in Section 3, one cause of this higher raw BER can be nonuniform erasure from absorption effects that reduce the hologram fidelity.^{10,11} To make sure that the first hologram meets the BER specification, we have to back off on the number of holograms M and construct a schedule with a different time constant, τ_{chosen} , than the one that governs the evolution of the diffraction efficiency τ_e . Because we want the number of holograms in Eq. (7) to decrease, we pick a value of $\tau_{\text{chosen}} < \tau_e$. With this artificially chosen time constant, the first hologram will end up with a stronger diffraction efficiency than will the last. One keeps reducing τ_{chosen} and recalculating X , M , and t_{total} until $t_{\text{total}} \sim t_{\text{meas}}$. The ratio between the diffraction efficiencies of the first and the last holograms will then be

$$\frac{\eta_1}{\eta_M} = \left(\frac{t_1}{t_M} \right)^{2(1-\tau_{\text{chosen}}/\tau_e)}, \quad (9)$$

but the first and the last holograms will both be at the target raw BER.

A more fortunate situation occurs when, with the original τ_e , one finds that $t_{\text{total}} < t_{\text{meas}}$. This situation implies that the final holograms are more prone to error, possibly because photovoltaic noise is affecting the image fidelity (and thus new holograms) more quickly than it is affecting already stored holograms. Forcing all the holograms to have the same diffraction efficiency would leave the first hologram at a raw BER that is less than (better than) the target. We can repeat the above procedure of using an incorrect time constant to *increase* the number of holograms, in this case by setting $\tau_{\text{chosen}} > \tau_e$.

Note that an alternative approach might be to select a different ECC code for holograms at different points within the recording schedule—stronger codes for weak holograms and weaker codes for the strong holograms. However, the performance of ECC tends to improve in discrete steps as redundancy is added, making it more difficult to tune the performance of all the holograms. In addition, there is significant complexity involved in implementing multiple ECC codes in hardware and then rapidly switching between

them during readout. In comparison, the approach we outlined above involves only a slight reorganization of the recording exposures.

4. Experimental Implementation

The experimental demonstration of this procedure for estimating capacity was performed on the IBM DEMON holographic demonstration platform.⁴ Holograms were stored in LiNbO₃:Fe in the 90° geometry. An input field of 320 × 240 pixels on a spatial light modulator (SLM) was pixel matched onto the CCD, permitting the implementation of various modulation codes and thresholding techniques. The 8-mm-thick LiNbO₃ crystal was located approximately 20 mm in front of the Fourier transform plane, at which was placed a 9 mm × 9 mm aperture. The crystal was 0.015% Fe doped, cut for the 90° geometry (c axis in the reference-object interaction plane, at 45° to the entrance faces). For recording the object beam contained approximately 500 μW. Approximately 80% of this was in the central dc order of size 1.3 mm × 1 mm at the 15 mm × 15 mm crystal entrance. The 125-mW reference beam was 5 mm wide and 11 mm tall and covered the central three (vertical) diffraction orders of the SLM pattern that pass through the Fourier transform aperture. The same reference-beam power was used for readout.

The attenuated readout beams needed for the capacity-estimation technique were implemented by an acousto-optic modulator (AOM) controlled by a Hewlett-Packard Model 5359 time synthesizer as a fast shutter. The AOM deflected the constant laser power into the system for the desired portion of the 16.66-ms camera-integration time, thus dictating the readout energy. The resulting change in the diffracted signal detected by the camera (in units of camera counts, the 8-bit values returned by the EPIX frame grabber) was measured to be linear with the AOM deflection time. Before storing each hologram, the SLM pixel grid was x - y registered with respect to the CCD camera. A typical first exposure was ~30 s, resulting in a diffraction efficiency of $\sim 1.7 \times 10^{-4}$. Because the detectable signal level in the presence of detector noise was $\sim 2 \times 10^{-8}$, this first hologram could survive at least $4\tau_e$ seconds of erasure. With the intensity values listed above, the erasure-time constant was ~550 s.

Erasure of the first hologram was done by the recording of multiple holograms (each at a 10-s exposure) within a small angle range that was spaced several degrees away from the hologram of interest. Before the first hologram was measured, the time synthesizer was used to reduce the readout energy until the reconstructed hologram no longer saturated the camera. The reference-beam angle was scanned (by use of a galvomounted mirror) to locate the Bragg matching angle, and the CCD camera was x - y registered with respect to the reconstructed hologram. Then 18 measurements of the hologram were taken over a range of readout energies (as controlled by the AOM deflection time). These values covered an order of magnitude in detected signal strength (from a

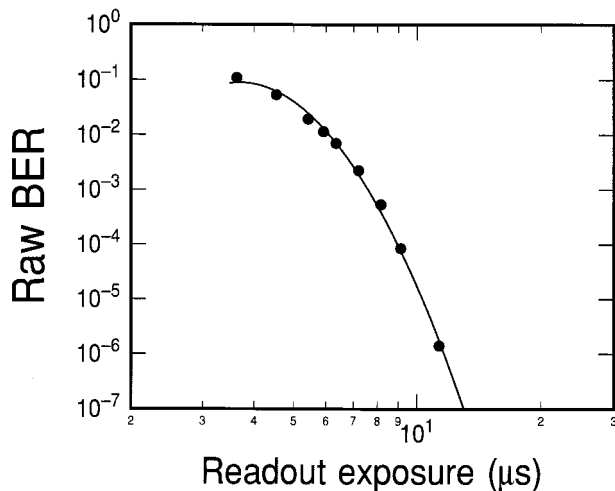


Fig. 2. Measurement of the raw BER as a function of the readout exposure time and a quadratic fit. The camera-integration time is 16.66 ms.

mean camera-count level of 100 in the ON pixels down to a mean of <10). For each measurement the readout deflection time was adjusted for any slight variations in the total reference-beam power, measured with a calibrated pick-off powermeter. The captured data pages were decoded off-line with the appropriate modulation decoder or thresholding algorithm. This erasure, re-Bragg matching, and measurement cycle was then repeated several times, resulting in a set of data corresponding to a range of t_{meas} values. A typical measurement of the raw BER versus the readout exposure is shown in Fig. 2. For raw-BER values of less than 1 error/page, we extrapolated from the measured data by using a quadratic fit on the log-log plot, as shown. Before fitting, the weakest zero-error data point is placed on the graph at a BER level of 0.1 error/page. Measurements with multiple pages and extrapolated raw BER's showed that a quadratic relation holds for raw BER's down to $<10^{-10}$. (It might be tempting to ascribe this quadratic behavior to the operation of a Gaussian-like noise distribution. However, such a dependence would be exactly quadratic on only a semilog plot that is linear in readout exposure.)

As was described above, the data from Fig. 2 can be used to identify the minimum real exposure time that can be used to meet a given raw-BER target. All that is needed is the exposure time that was used at the beginning of the experiment. To improve the repeatability of the experiment, we averaged the absolute hologram strength over several trial exposures (taken at different reference-beam angles). In these experiments, we observed no change in recording behavior as a function of optical exposure, so the dependence of the raw BER on the recording time was the same for both the first and the last holograms (when measured immediately after recording). The particular t_1 values were obtained in two steps: First, we compared the relative change between two plots like those shown in Fig. 2 that were taken of the same

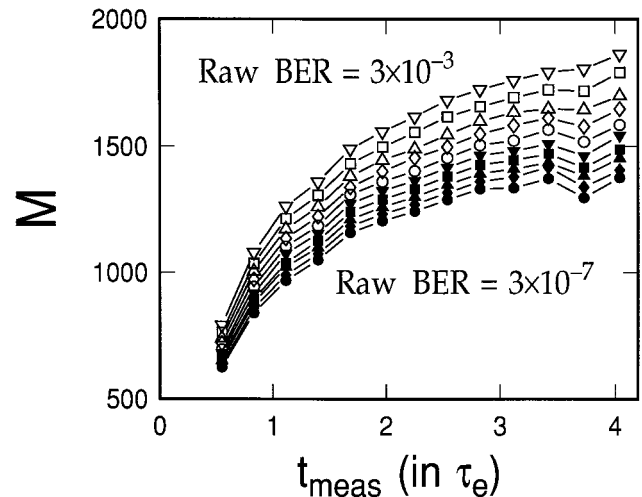


Fig. 3. Number of holograms that can be stored (M) as a function of the total recording exposure time t_{meas} (expressed in units of the erasure-time constant τ_e). The different curves correspond to raw-BER values ranging from 3×10^{-7} to 3×10^{-3} . All the curves represent the 6:8 modulation code.

hologram just after recording and at time t_{meas} . The data of interest here are the readout powers needed to make the raw BER reach its target value. The relative ratio between these two readout powers (after recording and after time t_{meas}) was turned into an absolute exposure time t_1 by use of data that had been averaged over several hologram exposures. This reflects the experimental realities in our particular system: the same hologram exposure time does not always produce the same absolute diffraction efficiency (because of fluctuations in the laser or the delivery fiber), yet a hologram (once it is written) evolves independently of its initial absolute diffraction efficiency.

So values of t_1 and t_M were obtained at each t_{meas} point. From these three variables and the erasure time constant τ_e , the procedure outlined above was used to obtain M , an estimate of the number of holograms that can be stored. A typical result is shown in Fig. 3, with M plotted as a function of t_{meas} for various target raw BER's. Here t_{meas} is shown in units of τ_e . The number of stored holograms saturates as the total exposure time approaches $3\tau_e$, as is expected from the operation of the recording schedule.⁸ Further increases in the total exposure time provide little increase in capacity.

There is, however, a significant capacity gain associated with using this new recording schedule because it equalizes the performance variable of interest (the raw BER) instead of the diffraction efficiency η . Under the conventional flat- η schedule, many holograms end up below (better than) the raw-BER target. By using the flat-BER schedule, we bring these holograms back up to the target raw BER by recording them for a commensurately shorter time, thus reducing the total exposure time. One can then store additional holograms and yet still finish within the original total exposure time.

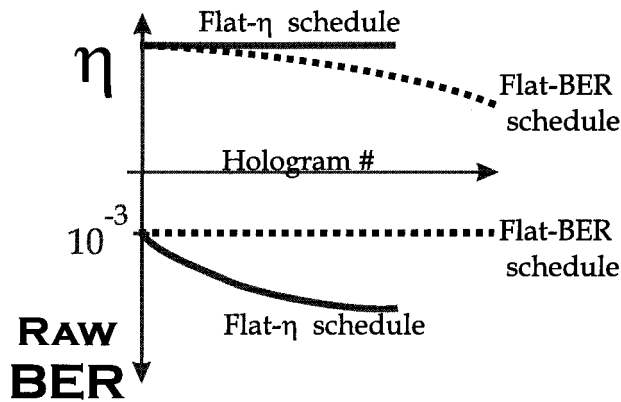


Fig. 4. Expected diffraction efficiency and the raw BER as a function of hologram number for a flat-BER and a flat- η schedule. Because the last hologram does not need as much diffraction efficiency as the first, the dynamic range can be taken away from the last hologram and used to store additional holograms. The arrow on the raw-BER axis points in the direction of a lower (better) BER. (This graph is an illustration, not measured data.)

Figure 4 demonstrates how the flat-BER schedule provides an increase in capacity over the flat- η schedule. The most common situation is shown, in which the first hologram of the conventional recording schedule (solid curve) has a higher (worse) raw BER, even though it has the same signal strength. For making the flat- η schedule meet the raw-BER specification for the first hologram, the other end of the schedule is a much lower raw BER than is really needed. The flat- η schedule (dotted curve) essentially takes away this excess dynamic range and puts it toward more holograms. The last hologram then has a lower diffraction efficiency than the first but has the same raw BER.

Figure 5 shows the results of the capacity-estimation experiment for both the conventional flat- η and the novel flat-BER schedules. The ex-

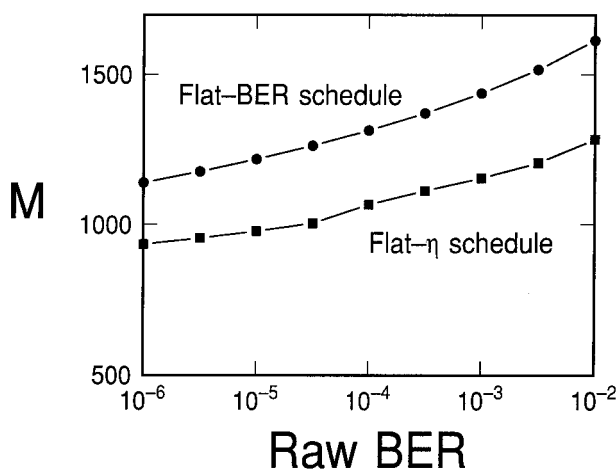


Fig. 5. Measured data for the number of holograms that can be stored (M) as a function of the raw BER with a flat-BER and a flat- η schedule. The flat-BER schedule takes the dynamic range from holograms that are under (better than) the raw-BER specification to maximize the number of stored holograms.

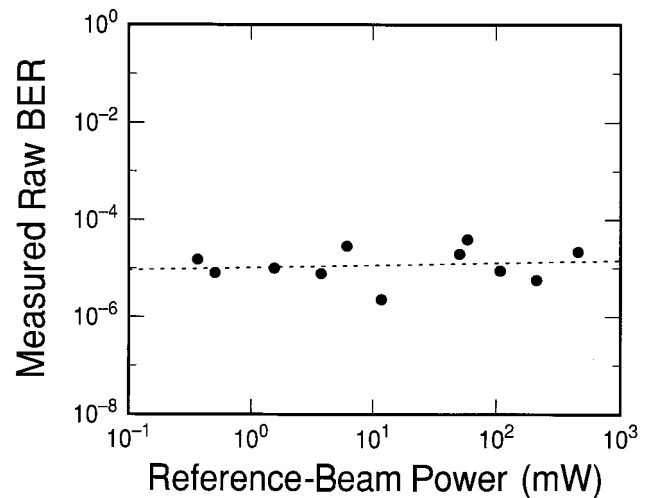


Fig. 6. Measured raw BER at a constant diffracted energy as a function of the readout power. Because scatter noise is not significant, the raw BER is independent of the readout power.

pected number of holograms is shown as a function of the raw BER, demonstrating an $\sim 20\%$ gain between the flat- η and the flat-BER schedules. As long as the flat- η schedule fails to equalize the BER, switching to a flat-BER schedule will permit a larger number of stored holograms.

The final experimental detail is the assumption concerning the role of scatter noise. In Fig. 6 we show an experiment that proves that scatter noise was indeed not significant in these experiments. The vertical axis shows the raw BER of a set of just-written holograms, each measured at the same signal strength at the camera (sufficient for a mean camera-count level of 20 in the ON pixel distribution). The horizontal axis indicates the wide range of readout powers used to achieve this reconstructed signal strength (high readout powers imply weak holograms). Because the scatter noise was not significant, this curve is flat and the raw BER is independent of the readout power. This is the condition assumed by our capacity-measurement procedure, because it draws conclusions about weak holograms (right-hand side of Fig. 6) from measurements of weakly reconstructed strong holograms (left-hand side of Fig. 6). If scatter noise had been significant, then this raw-BER curve would have increased rapidly with the readout power.

5. Error-Correction and Modulation Coding

ECC techniques are based on the use of structured redundancy that is added to user data so that errors occurring during storage or retrieval can be detected and corrected. The powerful class of Reed-Solomon (RS) block codes are examples of linear codes that utilize b -bit symbols to form code-word blocks containing c symbols, where $c = 2^b - 1$.¹³ These codes have been applied widely in both communication and storage channels because of their strong ECC power and their convenient decoding algorithms. For example, RS codes of length $c = 255$ symbols can

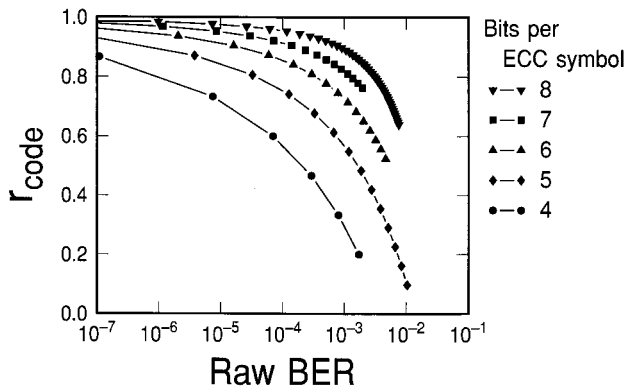


Fig. 7. Code rate of the RS ECC as a function of the lowest (worst) correctable raw BER for different numbers of bits per ECC symbol. The codes correct to a 10^{-12} user BER and are of maximal length for the symbol size.

achieve decoded user-BER values of 10^{-12} from raw-BER values of 10^{-3} by use of only modest amounts of redundancy (code rate of $r = 0.898$). High-speed decoding hardware for several members of this class of codes are available.

The raw-BER requirement for holographic storage systems in the absence of ECC becomes a raw BER equal to 10^{-12} to achieve acceptable user data fidelity. This raw-BER constraint is unachievable in most systems, and ECC provides a mechanism through which larger raw-BER values can be made acceptable at the cost of code rate.

Figure 7 shows a plot of the RS code rate required to achieve a user BER of 10^{-12} as a function of the raw BER. Several RS code-word lengths are considered. We see from this figure that longer code words offer a more efficient error-correcting ability and that the ability to tolerate very large raw-BER ($>10^{-2}$) is achieved only at the expense of significant code rate. In the context of holographic storage higher raw BER's engender only moderate increases in page capacity (number of stored pages), as seen from Fig. 5. This combination of rapidly decreasing code rate and slowly increasing page capacity suggests that an optimal raw BER (one that maximizes the amount of user data per hologram stack) can be identified. Additional capacity gained by means of an increase in the number of stored pages (i.e., through an increase in the tolerable raw BER) is then counterproductive because the cost of ECC overhead becomes greater than the apparent capacity gain. This phenomenon is demonstrated experimentally below in this section.

Whereas ECC operates on binary-valued data, correcting from the raw BER to the final user BER, modulation codes operate on analog data to deliver an improved raw BER from the same SNR. Modulation codes utilize redundancy to improve the binary decision-making process (the step in which bright and dark pixels are converted to binary ones and zeroes). Often we can treat the detection process as a simple thresholding operation in which the modulation code assists in the determination of the best threshold. One simple coding method is to include a

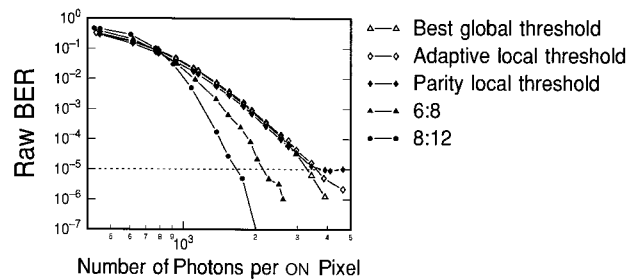


Fig. 8. Raw BER versus the signal strength in the presence of a constant noise floor. Two modulation codes (6:8 and 8:12) are compared with global, local, and adaptive thresholding.

few known pixels on each data page from which a threshold can be derived. These known pixels can be included within each block of pixels on the page or seeded just once on the page and the threshold adaptively updated as the page is decoded.^{14,15} With this latter method, error propagation is a possible result.

In contrast to these simple thresholding methods, block-based modulation codes utilize distributed redundancy, with n bits represented on the data page as m pixels ($m > n$).⁴ The encoding process ensures that each block contains a known number of ON and OFF pixels (that is, bright and dark pixels, respectively). At readout a correlation detector uses the brightness levels of the m received pixels to return the n -bit label of the best matching code word. For codes in which the number of possible code words is not much larger than 2^n , the same performance can be achieved if the m pixels are sorted and thresholded with an implicit threshold. When the number of possible code words is significantly larger than 2^n , the correlation decoder can return the correct symbol label even when an applied threshold would cause an error.⁴

In Fig. 8 the performance of two block-based modulation codes are compared with three thresholding techniques. This graph shows the experimentally measured raw BER as a function of the signal level in the presence of a constant noise floor (detector noise is dominant).⁴ The signal level is measured in photons per ON pixel. The global thresholding technique uses *a priori* information to obtain the best possible threshold; however, even in such an unrealistic scenario the global threshold performs poorly. The parity thresholding scheme implements local thresholding over regions of approximately 250 pixels. A small, spatially distinct portion of the page is used to store the number of ON pixels for each region. (These few parity bits are protected by differential encoding, in which two pixels are used to represent each bit. For instance, ON-OFF might represent binary 1, and OFF-ON binary 0.) For decoding, the local threshold for a block is computed from the total sum within the block (computed from the detected page), the number of ON pixels (from the transmitted parity information), and the contrast ratio between the ON and the OFF levels (two floating-point constants within the decoder).

The adaptive thresholding algorithm is taken directly from Ref. 15, in which the threshold is calcu-

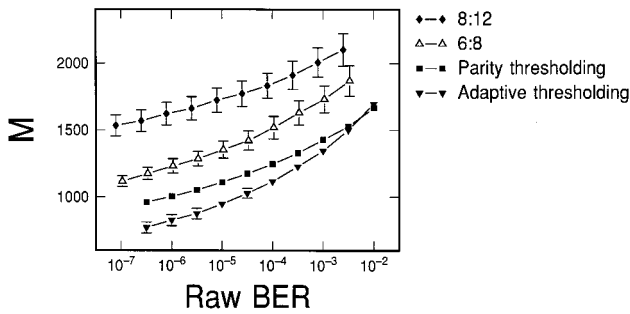


Fig. 9. Number of holograms that can be stored (M) as a function of the raw BER. Long recording schedules are assumed (saturation region of Fig. 4). Several modulation codes and thresholding schemes are compared, combining results from two separate experiments.

lated from a formula by use of the last eight decoded ON values, the last four OFF values, and two floating-point constants (optimized in our case to $p = 0.6$ and $q = 0.3$). Because our pages were uniformly illuminated, the data were input to the algorithm in a back-and-forth raster fashion rather than spiraling out from the center of the page, as was reported in Ref. 15. Similar performances are observed for all three thresholding schemes, suggesting that the improvement one gains by thresholding over localized regions is roughly balanced by the occasional suboptimal threshold. Superior raw-BER performance, however, is observed for the two modulation codes identified in the figure by their $m:n$ ratios. The 8:12 modulation code is a stronger code than the 6:8 code, and it exhibits some degree of soft error correction and obtains an improved raw BER, as shown in Fig. 8.

On the basis of the signal-level data presented in Fig. 8, one could derive a capacity comparison by using the $1/M^2$ assumptions, as described in Section 2. However, this would reflect the contribution of only the one noise source that was present (detector noise). With the experimental procedure described in Section 4, the expected user capacity with these different decoding strategies can be compared under more realistic system conditions. Following the discussion of Section 4, we measured the expected capacity of the DEMON holographic storage system with the codes described above. Figure 9 shows the maximum number of holograms that can be stored (M) for a long recording schedule (i.e., the saturation value from Fig. 3) as a function of the raw BER. This graph presents the results of two separate experiments averaged together to compare several modulation codes and thresholding schemes. The error bars indicate the variation between the two experiments. For each code the required conversion from symbol-error rate to raw BER is included. In support of the predictions of Fig. 9, we note that 1200 holograms were multiplexed in the DEMON system under similar experimental conditions and retrieved with no error (raw BER $< 2 \times 10^{-8}$)⁴ by use of the 8:12 code.

For the parity coding scheme, errors can occur from

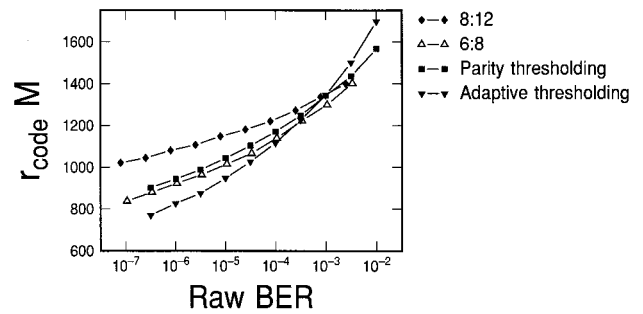


Fig. 10. Number of user and ECC bits that can be stored ($r_{\text{code}} \times M$) as a function of the raw BER. Here a low code rate is traded off against high performance (permitting more stored holograms). The error bars from Fig. 9 still apply but have been omitted for readability.

pixels that fall on the wrong side of the computed threshold or in rare cases when the differential-encoded parity information is decoded incorrectly and a wildly incorrect local threshold is obtained. To compare the importance of these two possibilities, we decoded the same parity-encoded data pages again but used *a priori* knowledge instead of the decoded parity information. The resulting capacity was only slightly increased over the curve shown, so we conclude that errors made in decoding the parity information are a second-order effect.

As expected, the stronger codes permit a larger number of holograms to be stored. However, any increase in the number of stored holograms permitted by these codes must be balanced against the code rate: the loss of bits *per* hologram. This trade-off is shown in Fig. 10, for which we multiplied the number of holograms M by the code rate r_{code} to obtain an estimate of raw capacity (user bits as well as redundant ECC bits). Here we dropped the error bars for readability, although they still apply. We see from this figure that the 8:12 code still outperforms the other methods; however, it does not win by a large margin. When this small capacity improvement is weighed against the required decoding complexity, it is not clear that the 8:12 code offers an overall superior solution under these experimental conditions. (We note that the aperture at the Fourier transform plane was fairly large and that the holograms were isolated in angle so that neither interpixel nor interpage cross talk was a significant source of noise.)

The discussion above together with the data of Fig. 10 suggest that large capacity can be obtained by operation of the holographic storage system at a high raw BER. Although the results shown in Fig. 10 do include the effect of the modulation-code rate, the vertical axis does not yet measure user capacity because it does not include the code rate of the ECC. This can be included simply by multiplication of each point in the raw capacity curve of Fig. 10 by the corresponding point in the 8-bit-per-symbol ECC code-rate curve of Fig. 7 (the top curve), because the horizontal axis of both curves is the raw BER. The resulting graph of the expected user capacity ($r_{\text{ECC}} r_{\text{code}} M$) is shown in Fig. 11 as a function of the




Fig. 11. Total user capacity in bits ($r_{\text{ECC}} \times r_{\text{code}} \times M$) as a function of the raw BER for an output user-BER specification of 10^{-12} . An optimal working point for the system can be identified. Error bars have been omitted for readability.

raw BER. This can be converted directly to user bits per stack by multiplication by the number of pixels per hologram. As mentioned in Section 1, the user capacity per stack can be extended to total system capacity (in user bits) or volumetric density (in user bits per unit volume) by incorporation of the total number of stacks or the volume per stack, respectively.

Under these experimental conditions (no interpage or interpixel cross talk), the strong raw-BER performance of the modulation codes is almost exactly balanced by their lower code rate. Given the complexity of the 8:12 decoder and keeping in mind the desire for a high readout rate as well as high capacity, we see that the thresholding schemes would appear to be advantageous here. As the areal density is increased by aperturing of the data-bearing image beam, however, the resulting interpixel cross talk will increase the N_1 term from Eq. (1) and possibly push the balance back toward the block-based codes. Future experiments with the same capacity technique detailed in this paper are the best (and possibly only) way to settle this question.

It is also possible to convert the horizontal axis of Fig. 10 from the raw BER to the number of pages and show the total user capacity as a function of M (Fig. 12). Figures 11 and 12 can be used to identify the optimal working point for the holographic storage

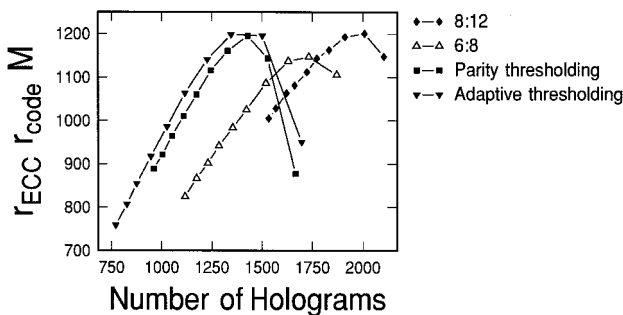


Fig. 12. Total user capacity in bits ($r_{\text{ECC}} \times r_{\text{code}} \times M$) expressed as a function of the number of holograms that can be stored (M). Recording too many holograms reduces capacity because the code-rate expense of the required ECC code outpaces the increase in the number of holograms. Error bars have been omitted for readability.

system. Figure 11 represents this optimum in terms of the best raw BER, whereas Fig. 12 depicts the same optimum as the number of holograms that should be recorded to maximize capacity. Note that the optimal operating point appears to be at a raw BER of 10^{-3} , independent of the decoding scheme used. In contrast, the optimal number of holograms depends on the decoding strategy. For example, the same user capacity is achieved either by means of storing 1350 holograms with adaptive thresholding and a strong RS code ($t = 14$ bytes of correction) or by means of storing 2000 holograms with the 8:12 code and nearly the same RS code ($t = 13$).

6. Generalizations of the Procedure

The capacity experiment described in this paper can be generalized to evaluate the capacity of write-once or saturable materials. In the simplest case of these materials, the index change with exposure is completely linear and all recording exposures are equal. For exposure-scheduling purposes, however, this is identical to the ordinary flat- η recording schedule used for photorefractives constructed with a time constant to infinity. To convert to the flat-BER schedule described in this paper, one would simply switch to a large but finite time constant and derive a recording schedule as if using a photorefractive material. In the case in which the first holograms are more error prone than the last, the time constant would be large and positive. If the last holograms needed a larger dynamic range than the first, then the time constant would be large and negative. (Here the recording schedule is operating under the pretense that the diffraction efficiency of the first hologram *grows* during the recording of subsequent holograms!)

As a result of the high photosensitivity of these saturable materials, the multiple measurements required for our procedure, if made immediately after recording, would probably affect the remainder of the experiment. To match the conditions that resemble the operation of the storage system, one should measure the holograms only after the material is fully exposed and fixed. This fixing process would include any UV postexposure that might be required either to use up remaining monomers or unreacted absorbers or to destroy the sensitizer. As a result the various choices of total recording exposure that were implemented for photorefractives by measurement at certain time delays t_{meas} would now have to be implemented at different spatial locations on a given sample of write-once material. Holograms should be written just strongly enough to attain a low raw BER—then, as before, the amount that the readout beam needs to be weakened indicates exactly what exposure time is required.

A method for scheduling the exposures to equalize the diffraction efficiencies of holograms in saturable materials has been published by Pu and Psaltis.¹⁶ This method takes into account any reciprocity failure at the beginning or the end of recording by deviation from equal-time exposures at these points.

One can combine this method with the capacity measurement by merely switching from η equalization to raw-BER equalization. This should result in a flat-BER schedule that retains the corrections for reciprocity failure.

A second consideration relevant to the implementation of this measurement technique in saturable materials is the inherent assumption that background scatter noise is not a significant factor. This is generally true for the 90° geometry because of the large angle between the strong reference beam and the object-beam path leading to the light-sensitive detector array. However, scattering is important in the transmission-geometry configurations mandated by thin (<2-mm) saturable materials. If background scatter is significant, the contribution to noise and thus the raw BER will increase as the readout energy increases. During the course of the measurement procedure, scatter noise would be accurately included for only the final measurement (the last value of t_{meas}), for which the initial hologram has decayed so much that the readout beam hardly needs to be weakened at all. However, it is difficult to know which value of t_{meas} will optimize capacity before performing the experiment. One option for including the effects of scatter noise in the capacity-estimation procedure would be to record several test holograms of varying exposures at the outset and then to measure each with nearly the full readout power, thus introducing the real amount of scattered light. (Because a good measure of τ_e is also needed, one should also measure the diffraction efficiency of the strongest test hologram throughout the experiment.) An alternative method for gauging scatter noise might be to split a constant amount of power between a Bragg-matched readout beam and a non-Bragg-matched beam. The added complexity of a third beam and the need to keep the two readout beams from writing undesired gratings during the measurements probably make the first solution the more practical of the two.

As a final note, the procedure could be expanded further to incorporate the effects of interpage cross talk. After performing the experiment with isolated holograms as described above, the experiment can be repeated with several test holograms angularly spaced at the separation that would be used to implement the M holograms (first null, second null, etc.). This should be especially interesting because the erasure effects already observed to affect fidelity should also tend to broaden the angular selectivity through nonuniform erasure of the hologram. It is conceivable that the angular spacing between holograms should also be scheduled to compensate for subsequent exposure.

7. Conclusions

In conclusion, we have described and implemented an experimental procedure to estimate the total user capacity of holographic data-storage systems in the presence of multiple noise sources. In this procedure, the raw-BER evolution of a hologram was

measured with weakened readout beams while subsequent holograms were multiplexed into the storage location. Only one recording schedule was able to meet the resulting requirements for first exposure, last exposure, and total recording time, which led to a value for M , the number of holograms that can be stored at a given raw BER. By the inclusion of the relation between the ECC code rate and the raw BER, the optimal working point for the system could be identified. By the inclusion of the modulation-code rate, the important raw BER-code rate trade-off could finally be quantified in terms of total user capacity or volumetric user density (previously an intractable problem in the presence of several noise sources). Under experimental conditions in which the noise buildup with total optical exposure and the background detector noise dominate, strong modulation codes showed little advantage in total capacity over local thresholding techniques.

System choices that can be evaluated with this procedure include the choice of phase mask, the choice of beam ratio during recording, component choices affecting hologram stack size (focal length, pixel size, aperture stop, and material thickness), and signal-processing techniques such as predistortion.¹⁷ Modifications for implementing this procedure in the presence of scatter noise or interpage cross talk and in saturable materials have been discussed.

The authors acknowledge helpful discussions with members of the Holographic Data Storage Systems Consortium. This work was partially supported by the U.S. Defense Advanced Research Projects Agency through the Holographic Data Storage Systems Consortium under agreement MDA972-95-0004.

References

1. D. Psaltis and F. Mok, "Holographic memories," *Sci. Am.* **273** (5), 70–76 (1995).
2. J. F. Heanue, M. C. Bashaw, and L. Hesselink, "Volume holographic storage and retrieval of digital data," *Science* **265**, 749–752 (1994).
3. M. A. Neifeld and M. McDonald, "Error correction for increasing the usable capacity of photorefractive memories," *Opt. Lett.* **19**, 1483–1485 (1994).
4. 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," *Opt. Lett.* **22**, 639–641 (1997).
5. J. F. Heanue, M. C. Bashaw, and L. Hesselink, "Channel codes for digital holographic data storage," *J. Opt. Soc. Am. A* **12**, 2432–2439 (1995).
6. C. Gu, J. Hong, I. McMichael, R. Saxena, and F. Mok, "Cross-talk-limited storage capacity of volume holographic memory," *J. Opt. Soc. Am. A* **9**, 1–6 (1993).
7. C. Gu, F. Dai, and J. Hong, "Statistics of both optical and electrical noise in digital volume holographic data storage," *Electron. Lett.* **32**, 1400–1402 (1996).
8. D. Psaltis, D. Brady, and K. Wagner, "Adaptive optical networks using photorefractive crystals," *Appl. Opt.* **27**, 1752–1759 (1988).
9. F. H. Mok, G. W. Burr, and D. Psaltis, "System metric for holographic memory systems," *Opt. Lett.* **21**, 896–898 (1996).
10. S. Campbell, S.-H. Lin, X. Yi, and P. Yeh, "Absorption effects

- in photorefractive volume-holographic memory systems. I. Beam depletion," *J. Opt. Soc. Am. B* **13**, 2209–2217 (1996).
11. G. W. Burr, "Volume holographic storage using the 90° geometry," Ph.D. thesis (California Institute of Technology, Pasadena, California, 1996).
 12. M.-P. Bernal, G. W. Burr, H. Coufal, J. A. Hoffnagle, C. M. Jefferson, R. M. Shelby, and M. Quintanilla, "Experimental study of the effects of a six-level phase mask on a digital holographic storage system," *Appl. Opt.* **37**, 2094–2101 (1998).
 13. I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *J. Soc. Indust. Appl. Math.* **8**, 300–304 (1960).
 14. M. M. Wang, S. C. Esener, F. B. McCormick, I. Cokgor, A. S. Dvornikov, and P. M. Rentzepis, "Experimental characterization of a two-photon memory," *Opt. Lett.* **22**, 558–560 (1997).
 15. X. A. Shen, A.-D. Nguyen, J. W. Perry, D. L. Huestis, and R. Kachru, "Time-domain holographic digital memory," *Science* **278**, 96–100 (1997).
 16. A. Pu and D. Psaltis, "High-density recording in photopolymer-based holographic three-dimensional disks," *Appl. Opt.* **35**, 2389–2398 (1996).
 17. G. W. Burr, H. Coufal, J. A. Hoffnagle, and C. M. Jefferson, "Noise reduction for page-oriented data storage by inverse filtering during recording," *Opt. Lett.* **23**, 289–291 (1998).