

Noise reduction of page-oriented data storage by inverse filtering during recording

Geoffrey W. Burr, Hans Coufal, Robert K. Grygier, John A. Hoffnagle, and C. Michael Jefferson

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

Received July 30, 1997

Two novel techniques for eliminating deterministic noise from a page-oriented memory are presented. The first technique equalizes the output response of ON pixels by adjustment of the exposure of each pixel during the recording of each data page. A test image transmitted through the system measures the spatial nonuniformities, and the appropriate inverse filter is imposed upon the data page and recorded in the storage material. On readout, the output signal values are then spatially uniform, perturbed only by random noise sources. Experimental results of using this predistortion technique in a pixel-matched holographic storage system are shown. Under conditions of high volumetric density, raw bit-error-rate (BER) improvements of 6–8 orders of magnitude are obtained (from 10^{-4} to $<10^{-10}$). The second technique uses a phase shift during holographic storage to subtract from bright OFF pixels. Under conditions of low spatial light modulator contrast, BER improvements of 6 orders of magnitude (from 10^{-2} to 10^{-8}) are demonstrated. © 1998 Optical Society of America

OCIS codes: 210.2860, 070.6020, 120.2440.

Page-oriented data storage is a feature common to several three-dimensional optical storage technologies^{1–3} intended to replace conventional bit-serial storage devices. By distributing data throughout a volume and accessing it in parallel, these technologies can offer both high density and fast data rates. Output data leave the storage material as a complex two-dimensional image, which is then brought into focus and aligned onto a pixelated detector array. Binary decisions (was this pixel ON or OFF?) are made either by use of thresholding or in conjunction with modulation decoding.⁴ Errors occur when random noise or pixel-to-pixel variations cause an ON pixel to appear dark or an OFF pixel to appear bright. Random or effectively random noise sources include scattered light, interpage cross talk, shot noise, and noise in the detector electronics. Pixel-to-pixel variations arise from nonuniformities in input or detector pixel response, variations in intensity across the original input image, spatial profiles induced by the recording process, optical imperfections, and interpixel cross talk from nearby pixels.

These spatial variations in the brightness of output pixels are nonrandom; two stored copies or holograms made in succession will be nearly identical. Most of the variations are present in the transmission of an image through the storage system. In this Letter a novel predistortion technique is introduced that dramatically reduces nonrandom pixel-to-pixel variations during the recording of page-oriented data. On readout, the ON signal levels are spatially uniform across the data page and only random noise remains to affect the bit error rate (BER). Although the technique is described as it applies to holographic storage in photorefractives, it is applicable to page-oriented optical storage in general.

The first step in the predistortion technique is the measurement of spatial nonuniformities in the form of a template, which can either be a practice hologram or a transmitted image of the desired data page,

with averaging of multiple frames to reduce detector noise and appropriate spacing of holograms to avoid interpage cross talk.

The second step is to create an inverse filter from this template, choosing the exposure of each pixel with the goal of equalization. ON pixels that were weak in the template will receive more exposure when the corrected hologram is recorded. Pixels that were too strong will be brought back down to the average by reduction of their exposure. If the spatial light modulation (SLM) is capable of multiple brightness levels, the exposure of a pixel can be reduced by setting of the appropriate gray level on the SLM. If the SLM is binary, one can reduce the exposure of a pixel by turning it off before the total exposure time is completed.

The diffraction efficiency of each signal component of a hologram can be written as $\eta_i \propto I_S I_R t_i^2$, where I_R is the reference beam intensity and I_S and t_i are the intensity and the exposure time of the i th pixel, respectively. If the SLM is gray-scale, the predistortion technique has linear control over η_i through the choice of I_S . If the SLM is binary, then the term under control is t_i^2 .

Experiments with the predistortion technique were performed on the IBM DEMONstration platform.⁴ This system is a pixel-matched volume holographic system using angle multiplexing in LiNbO₃:Fe. The SLM is effectively binary because of the driver electronics. Figure 1 shows before-and-after histograms when a practice hologram is used as a template to measure the spatial nonuniformities. In this experiment a large aperture was placed at the Fourier transform plane, so interpixel interference is not a large effect. The inversion of the template was used to create a sequence of images, which were displayed during recording of the corrected hologram. The first image in the sequence was identical to the data page used to record the uncorrected hologram, whereas the

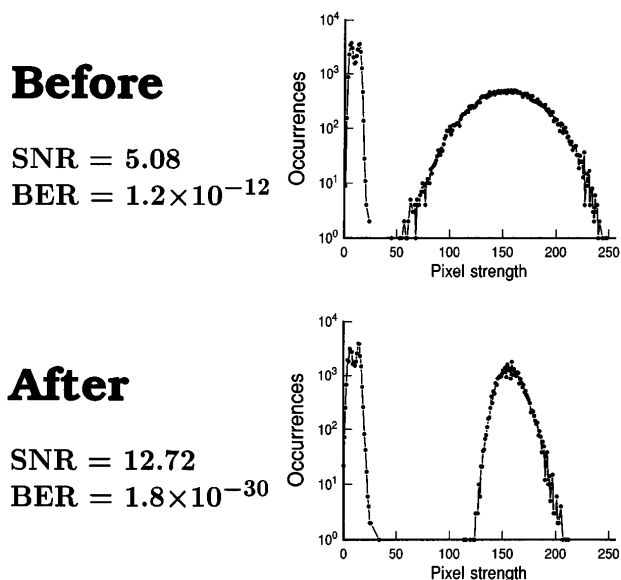


Fig. 1. Histograms of detected pixel values for holograms before and after use of the predistortion technique.

last image was almost completely dark, with only the 5–10 worst ON pixels still transmitting light. There were 20 images in the sequence; the first was recorded for $\sim 50\%$ of the total exposure, and the remaining time was divided equally among the other images. We define SNR as $(\mu_{\text{ON}} - \mu_{\text{OFF}})/(\sigma_{\text{ON}}^2 + \sigma_{\text{OFF}}^2)^{1/2}$, where μ and σ are the mean and the standard deviation of the two brightness distributions, respectively. The BER is estimated with the assumption of a 6-bit–8-pixel modulation code.⁴ The estimation procedure involves calculation of separate Gaussian probability density functions and pixel-swap probability for each of the 140,000 possible ON–OFF swaps on the data page.

Instead of storing a practice hologram before each corrected hologram, one can measure the spatial nonuniformities with a transmitted image. However, a template created from an image does not include the spatial profile added by the recording process. This profile can come from the reference beam itself or from absorption of the reference beam as it propagates across the signal path. For a data page stored anywhere outside the Fourier transform plane, variations across the hologram aperture can affect the uniformity of the detected data page.

Suppose that the predistortion procedure was carried out on a page composed of all ON pixels but without these recording effects' being taken into account. Since the resulting corrected hologram has all other spatial variations removed, any nonuniformity left is due only to recording effects. Then the product of the transmitted image of each data-bearing page and this one envelope template hologram reflects all the spatial nonuniformities. The same envelope template can be used for an entire multiplexed stack of holograms with no noticeable change in BER performance, reducing the capacity overhead to a single non-data-bearing hologram per stack.

The improvement in spatial uniformity with the predistortion technique comes at the cost of some dynamic range. Since the technique reduces the expo-

sure of the average pixel, a longer total exposure time is required for a corrected hologram to reach the same average diffraction efficiency. For Fig. 1, the ratio between the total exposure times of the after and before holograms was approximately 2.25. The additional erasure caused by the reference beam during these longer exposures must be weighed against any improvements in initial BER.

To increase storage capacity, predistortion must allow more holograms to be written, given a final BER target. Since the average diffraction efficiency of M superimposed holograms decreases as $1/M^2$, and several of the random noise sources (such as scatter noise and detector noise) do not decrease with signal level, the final BER for large M is influenced strongly by hologram strength.

There are situations, however, in which it is easy to show that the predistortion technique increases storage capacity: e.g., when the initial BER of a single hologram is already over the target level. One such situation occurs when one tries to maximize volumetric storage density by decreasing the size of the aperture at the Fourier transform plane. Since the image at the CCD camera is low-pass filtered, interpixel cross talk is a significant noise source. Figure 2 shows the BER's of transmitted images in the DEMON system as a function of the aperture size. The local minimum near 2.8 mm is where the first null of the isolated-ON-pixel response coincides with the center of the neighboring OFF pixel. This dip in BER is quite pronounced in the DEMON system because the effective CCD areal fill factor is small ($\sim 15\%$).⁴

The predistortion technique is more difficult to apply in the presence of interpixel cross talk, since any correction applied to a pixel affects its neighbors as well. For the initial demonstration of the technique, several practice holograms stored in succession were used to refine the template information and implement the required lateral inhibition between pixels. After each practice hologram, the relative exposures that had been applied to each ON pixel were updated for the next hologram by the pixel brightness in the just-measured hologram. The lower curve in Fig. 2 corresponds to the BER of corrected holograms after several of these perturbation cycles. After four cycles the BER did not improve further, nor did it change if the fourth template was repeatedly used without




Fig. 2. Raw BER versus aperture size for transmitted images and holograms corrected with the predistortion technique.





Fig. 3. Histogram for a hologram with four gray levels, created with predistortion.

Before

$$\text{SNR} = 2.28$$

$$\text{BER} = 1.1 \times 10^{-2}$$



After

$$\text{SNR} = 6.13$$

$$\text{BER} = 2.2 \times 10^{-8}$$




Fig. 4. Histograms for holograms before and after predistortion with subtraction from too-bright OFF pixels.

subsequent refinement. A practical implementation would use a gray-scale SLM to iterate with practice images instead of holograms.

Figure 3 shows a histogram that corresponds to a hologram with four gray levels, created by use of several perturbation cycles of the predistortion technique and a large aperture at the Fourier plane. Instead of the ON pixel exposures' being optimized to reach an average μ_{ON} , some exposures are optimized to $0.33 \mu_{\text{ON}}$ and some to $0.66 \mu_{\text{ON}}$.

A second predistortion technique, applicable to page-oriented holographic memories, uses a phase shift during recording to subtract from bright OFF pixels. Essentially, two holograms are stored in sequence, with a 180° phase shift between them. Figure 4 shows before-and-after histograms for holograms stored by predistortion with subtraction. A wave plate was inserted in the DEMON system after the SLM to generate a low-contrast image. The two exposures were of identical length; one exposure was with predistortion (equalizing the ON pixels by adjustment of their exposures), and the other was with all pixels off. The phase shift was introduced when the crystal was moved between exposures. Several perturbations were used to narrow the ON distribution. The total exposure time can be reduced by use of the inverse of the data page to subtract from the OFF pixels, with a phase shift of 180° minus the phase difference between the two transmission states of the SLM.

This subtraction technique is not so useful for correcting OFF pixels when the undesired signal comes

from interpixel cross talk, because the field amplitude of the cross-talk signal tends to switch sign across the area of a CCD pixel. Since the subtraction technique can only subtract uniformly, there are many cases in which any additional recording only makes the OFF signal increase. Judging from Fig. 2, the best approach is to rely on choice of aperture rather than subtraction to affect the OFF pixels and then to use predistortion to equalize the ON pixels.

An alternative to predistortion might be to use the same hologram template information to perform electronic postprocessing.⁶ In known dark regions of the data page, signal values are increased before thresholding. However, since this also broadens the additive noise on both distributions, the signal-noise ratio of each subregion in the data page remains unchanged. The postprocessing essentially allows the detector to follow a spatially varying threshold. Since this is what a block-based modulation code does, the BER performance with postprocessing is the same as in the before cases that we have shown, albeit at a higher code rate. So postprocessing does not improve the BER over modulation coding, nor can it correct for fluctuations that are data dependent or that change during the recording schedule.

In conclusion, two novel techniques for removing deterministic noise sources from a page-oriented memory have been presented and demonstrated in a pixel-matched holographic storage system. The detected intensities of ON pixels can be equalized by appropriate choice of relative exposure during recording, resulting in spatial uniformity on readout at the cost of some dynamic range. Experimentally, the initial BER of single holograms can be reduced to less than 10^{-25} when interpixel cross talk is negligible. Alternatively, the technique can be used to store gray-scale holograms even with binary SLM. Under conditions of high volumetric density, the raw BER can be improved from 10^{-4} to $<10^{-10}$. A second technique, applicable only to holographic storage, uses a phase shift during recording to subtract from bright OFF pixels caused by low-contrast SLM.

This study was supported in part by the U.S. Defense Advanced Research Projects Agency under agreement MDA972-95-3-0004. We thank the members of the IBM Holographic Optical Storage Team and the Holographic Data Storage Systems consortium for helpful discussions.

References

1. D. Psaltis and F. Mok, *Sci. Am.* **273**(5), 70 (1995).
2. S. Hunter, C. Solomon, S. Esener, J. E. Ford, A. S. Dvornikov, and P. M. Rentzepis, *Opt. Memory Neural Networks* **3**, 151 (1994).
3. A. Renn and U. P. Wild, *Appl. Opt.* **26**, 4040 (1987).
4. G. W. Burr, J. Ashley, H. Coufal, R. K. Grygier, J. A. Hoffnagle, C. M. Jefferson, and B. Marcus, *Opt. Lett.* **22**, 639 (1997).
5. D. Psaltis, D. Brady, and K. Wagner, *Appl. Opt.* **27**, 1752 (1988).
6. X. An and D. Psaltis, *Opt. Lett.* **20**, 1913 (1995).