Compensation for pixel misregistration in volume holographic data storage

Geoffrey W. Burr and Timo Weiss

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

Received November 13, 2000

We describe what we believe to be a novel postprocessing algorithm for compensating for misregistrations between a detector array and the coherent image of a pixelated two-dimensional data page. A lookup table of baseline local offsets is combined with the dynamically measured global offset of the received data page, producing an estimate of the total lateral shift of each small block of pixels. A serial algorithm then reallocates the appropriate portion of the signal detected by each pixel to its neighbors, accounting for both the linear and the quadratic contributions introduced by coherent illumination of square-law detectors. This procedure can relax the tight constraints on page registration, optical distortion, and material shrinkage that currently hamper page-oriented holographic storage systems. Experimental results from a pixel-matched 1-Mpixel volume holographic system are presented, showing an increase in position tolerance (for a raw bit-error rate $<10^{-3}$) from $\pm 16\%$ to $\pm 40\%$ of the pixel pitch.

OCIS codes: 210.0210, 100.0100.

Volume holographic data storage is a topic of recent interest[1] because it can offer both high density and fast parallel readout. Digital data, optically input as large pages of bright and dark pixels, are superimposed within a common volume by use of appropriate multiplexing techniques. Upon readout, as many as 1×10^6 parallel output data channels can be present at the pixelated detector array.[2]

These data channels are independent, however, only if the pixel array imposed by the input spatial light modulator (SLM) can be accurately imaged to the array of detector pixels. To do this, one must ensure that the image of each pixel arrives at the center of, and is confined to a small region around, the associated detector pixel. Any change of scale (magnification error), lateral position (misalignment), or deviation from the original square grid (optical distortion) can lead to interpixel cross talk and loss of signal. Misfocus, diffraction from spatial band limiting,[3] and the remaining optical aberrations (e.g., coma, spherical aberration) also lead to interpixel cross talk through pixel blurring.

Several authors have proposed signal-processing schemes for page-oriented data storage,[4–6] extending existing one-dimensional algorithms to two dimensions while retaining a simple linear channel model. However, the optical detection process in holographic storage is inherently nonlinear. Although the intermingling of signals from neighboring pixels, which the signal processing looks to undo, takes place in the amplitude domain with coherent light,[5] each detector pixel measures the total intensity integrated over its spatial extent.[3] As a result, one cannot take the square root to return to incident amplitude. Also, misalignment of the data page throws off precomputed equalization kernals (unlike temporal channels, in which jitter is mostly a zero-mean random noise source). This creates an undesirable compromise: Signal processing to compensate for pixel blur demands even more stringent attention to pixel alignment.

Pixel alignment has usually been addressed through careful design and engineering: Aberrations are suppressed through either optical design or phase-conjugate readout,[7] and position feedback and fixturing achieve and maintain alignment. In general, though, once data pages reach 1 Mpixel (1024×1024 pixels) in size, a significant portion of the signal-to-noise ratio budget is consumed simply by residual optical distortion: One portion of the page must always be partly misaligned to bring another portion into perfect alignment.[2,8] In this Letter we propose and demonstrate a low-complexity nonlinear equalization technique that can compensate for both optical distortion and misalignment, correcting a moderate pixel blur in the presence of a significant pixel offset.

We derive this algorithm by directly considering the physical readout process in holographic data storage. Figure 1 shows a one-dimensional slice through three detector pixels, which measure signals r_0, r_1, and r_2. The incident optical field, imaged from the three associated SLM pixels, is slightly shifted relative to the camera pixels and blurred by the point spread function (PSF) of the optical system. For SLM pixels of linear fill factor ff_S and a Fourier-transform aperture of width D, the field from each pixel is the convolution

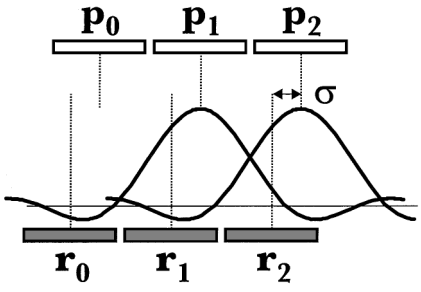

Fig. 1. Spatially blurred images of three SLM pixels (p_0, p_1, p_2) are slightly shifted relative to three CCD pixels (r_0, r_1, r_2), reducing the signal in the desired targets pixels and creating cross talk in their neighbors.

0146-9592/01/080542-03$15.00/0

of the rectangular SLM pixel and the sinc function PSF:

$$h(x) \equiv c \int_{-ffs/2}^{ffs/2} \mathrm{sinc}\left[\frac{D}{D_N}(x - x')\right] \mathrm{d}x', \tag{1}$$

where $\mathrm{sinc}(x) \equiv \sin(\pi x)/(\pi x)$. Here, c is chosen such that $\int_{-\infty}^{\infty} h^2(x)\mathrm{d}x = 1$, x and x' are in units of pixels, and $D_N \equiv \lambda f/\delta$ is the Nyquist aperture defined in terms of the wavelength, focal length, and pixel pitch. We use p_0, p_1, and p_2 both to identify SLM pixels and to represent the transmitted signal intensity. If the PSF is no broader than the pixel pitch and the pixel offset is significant, we can assume that all of the signal detected by pixel r_2 comes from pixels p_1 and p_2 and write

$$r_2 = \int_{-ff_d/2}^{ff_d/2} [\sqrt{p_2}\, h(x - \sigma) + \sqrt{p_1}\, h(x - \sigma + 1)]^2 \mathrm{d}x\,. \tag{2}$$

Here, ff_d is the fill factor of the detector, and σ is the total local offset between the SLM image and detector that is due to misregistration, optical distortion, magnification error, and material shrinkage[9] combined. Note that when σ flips sign (SLM pixels are shifted left), the $+1$ offset changes sign and p_3 replaces p_1 in Eq. (2); i.e., the right-hand neighbor pixel produces cross talk. When Eq. (2) is rewritten as

$$r_2 = p_2 H_{00}(\sigma) + 2\sqrt{p_1 p_2}\, H_{01}(\sigma) + p_1 H_{11}(\sigma)\,, \tag{3}$$

the detected signal is decomposed into linear contributions from the two pixel intensities and a nonlinear factor through their constructive interference. 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, as a function of the known local shift σ. The corresponding destructive interference demanded by conservation of energy is distributed across the next ring of pixels and is ignored as a second-order effect. The presence of a pixel-matched phase mask would spatially modulate the sign of the $\sqrt{p_1 p_2}$ term, causing problems only if the phase mask pattern is unknown to the receiver. The $H_{xx}(\sigma)$ weights can be computed from Eqs. (1)–(3) or measured empirically with known data during system calibration. In either case, large SLM and CCD pixels and apertures close to the Nyquist aperture lead to weighting functions that resemble broadened and shifted sinc functions for $|\sigma| \le 0.5$.

If p_1 is known, then Eq. (3) can be used to solve for p_2 and vice versa. For instance, if pixel p_0 in Fig. 1 has been intentionally left OFF, then p_1 is simply $r_1/H_{00}(\sigma)$. We can then substitute this value of p_1 into Eq. (3) to solve for p_2, using the binomial equation for $\sqrt{p_2}$. The shift-compensation algorithm that we propose here simply proceeds in this way along the entire row, solving for each pixel's true value by use of the just-processed neighbor pixel. At each pixel, we take the measured signal r, subtract the portion that belonged to the previous pixel, subtract a further portion that is due to interference, and then factor in the missing signal that should have been there but that actually fell into the next pixel. Although the algorithm requires *a priori* information about the local shift σ and the associated weighting factors $H_{xx}(\sigma)$, no information about the binary pixel states or special encoding arrangement is required.

If $|\sigma| > 0.5$, then the missing signal ascribed to the next pixel is large and the algorithm becomes prone to error propagation. For this reason, we prefer to process each line twice: once from left to right to compensate for rightward shifts of as much as half a pixel (as in Fig. 1), and once from right to left for leftward shifts. Each pixel is processed on one pass and skipped on the other. Note that a pixel shift of exactly one pixel requires only bookkeeping (r_1 is really p_0, r_2 is p_1, etc.) and blank columns at the page boundaries. By extension, the shift-compensation algorithm needs to work only for all shifts in the range $-0.5 < \sigma \le 0.5$ to achieve complete alignment insensitivity. Unfortunately, the transition at $|\sigma| \sim 0.5$ is problematic, especially if the estimation of the local offset σ is flawed. To reduce this effect, we soften the transition for any pixel with $0.4 < |\sigma| < 0.6$ by processing these pixels on both passes, producing a weighted average of the two estimates.

To extend the shift-compensation algorithm to two dimensions, one could expand the analysis of Eq. (3) to three neighboring pixels (horizontal, vertical, and diagonal), producing an equation with four linear and six nonlinear terms. However, when we measured the ten associated H_{xx} parameters, using the DEMON2 platform,[8] we found that all the parameters were separable into x- and y-dependent functions. (This implies that diffraction, not aberrations, dominated the PSF.) So we simply process each two-dimensional page with the one-dimensional algorithm four times: twice horizontally as described above on all rows to produce an interim page and then twice vertically (top to bottom and bottom to top) on all columns of this interim page. On a 600-MHz Pentium III computer, this processing takes 2.5 s for a 1024×1024 data page. In Fig. 2(a), we show a small 9×9 pixel block as it should ideally be received. Figure 2(b) shows the same pattern imaged through DEMON2 when the SLM is shifted by a half-pixel in both x and y. The DEMON2 platform pixel matches 1-Mpixel pages

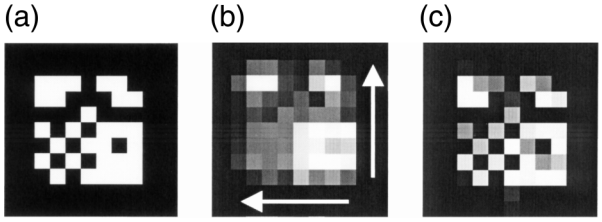

Fig. 2. A 9×9 pixel pattern is imaged from SLM to CCD (a) under perfect conditions, (b) with a half-pixel offset in both x and y, and (c) after postprocessing with the shift-compensation algorithm.

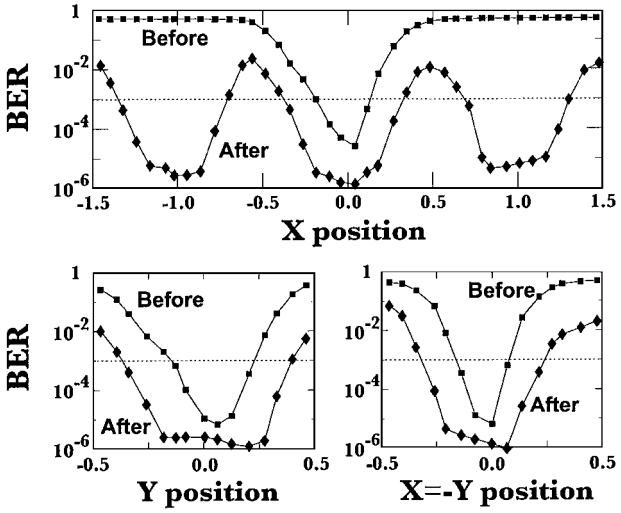

Fig. 3. Raw BER (with an 8:12 modulation code[10]) before and after shift-compensation postprocessing as a function of (a) x shift, (b) y shift, and (c) shift along the $x = -y$ diagonal.

through an aperture of $1.36D_N$ (1.7 mm $\times$ 1.7 mm aperture, $f = 30$ mm, $\lambda = 532$ nm, SLM $\delta = 12.8$ μm). In Fig. 2(c), we show these data after postprocessing with the shift-compensation algorithm: The original pixel pattern can be recovered by use of thresholding. Our initial implementation of the full two-dimensional algorithm, which would have been needed if aberrations had dominated the PSF, gave similar (but slightly inferior) results.

To process each block of pixels, the algorithm needs only the total lateral shift: the sum of global misregistration and the local offset that is due to magnification error, material shrinkage, and optical distortion. The dynamic global shifts are measured by use of dedicated fiducial marks; the static local baseline offsets can be mapped into a lookup table. For this work, we used a unique pair of x and y offsets, with a resolution of 0.025 pixel, for each block of 10×10 pixels across the 1-Mpixel data page. One set of $H_{xx}(\sigma)$ values was used across the entire page: two dimensions (symmetry broken by CCD pixel geometry), three tables per dimension (H_{00}, H_{01}, H_{11}), and 21 unique values per table to cover $0 \leq |\sigma| \leq 0.5$) Figure 3 shows the resulting pagewide performance for images transmitted through the DEMON2 system as a function of x misalignment, y misalignment, and equal x and y misalignment. Similar results have been demonstrated with holograms stored at areal densities as great as 250 Gbits/in.^{2} (Ref. 8). The raw bit-error rate (BER) was measured, before and after application of the shift-compensation algorithm, at the output of a subsequent 8-bits-from-12-pixels modulation decoder.[10] As expected, the postprocessed BER increases when either $|\sigma_x|$ or $|\sigma_y|$ is near 0.5 and decreases again for pixel shifts close to ± 1. Based on the assumption that 10^{-3} is the maximum acceptable raw-BER that can be corrected by error-correction codes,[8] Fig. 3 represents an increase in position tolerance from $\pm 16\%$ to $\pm 40\%$ of the pixel pitch.

One of the factors leading to a high BER at $|\sigma| \sim 0.5$ is error propagation, since noise-induced mistakes in signal estimation have a strong influence on succeeding pixel values. One way to reduce this factor is to insert rows and columns of OFF pixels at which the algorithm can reseed itself, trading code rate for performance gain. Such an approach also permits parallel execution, reducing the buffer size and processing latency. Alternatively, one could intentionally introduce a small ($\sim 1\%$) magnification error so that the changing local offsets pass quickly through $|\sigma| \sim 0.5$ in small isolated patches. Finally, Eq. (3) could simply be used as an improved channel model within existing parallel detection[6] and sequence-estimation schemes.

In conclusion, we have presented a postprocessing algorithm that compensates for local pixel misregistrations within coherent two-dimensional data pages. Experimental imaging results from a pixel-matched 1-Mpixel volume holographic system were presented, showing an increase in position tolerance from $\pm 16\%$ to $\pm 40\%$ of the pixel pitch while maintaining $<10^{-3}$ raw BER. We expect that this procedure can help relax—and with further improvement, completely remove—the tight constraints on page registration, optical distortion, and material shrinkage that currently hamper page-oriented holographic data storage systems.

G. Burr's e-mail address is burr@almaden.ibm.com.

References

1. H. J. Coufal, D. Psaltis, and G. Sincerbox, eds., *Holographic Data Storage* (Springer-Verlag, Berlin, 2000).
2. 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, Opt. Lett. **22,** 1509 (1997).
3. M.-P. Bernal, G. W. Burr, H. Coufal, and M. Quintanilla, Appl. Opt. **37,** 5377 (1998).
4. J. Heanue, K. Gurkan, and L. Hesselink, Appl. Opt. **35,** 2431 (1996).
5. V. Vadde and B. V. K. Vijaya Kumar, Appl. Opt. **38,** 4374 (1999).
6. B. King and M. A. Neifeld, Appl. Opt. **37,** 6275 (1998).
7. F. Zhao and K. Sayano, Opt. Mem. Neural Netw. **6,** 261 (1997).
8. G. W. Burr, C. M. Jefferson, H. Coufal, M. Jurich, J. A. Hoffnagle, R. M. Macfarlane, and R. M. Shelby, Opt. Lett. **26,** 444 (2001).
9. R. M. Shelby, D. A. Waldman, and R. T. Ingwall, Opt. Lett. **25,** 713 (2000).
10. G. W. Burr, J. Ashley, H. Coufal, R. K. Grygier, J. A. Hoffnagle, C. M. Jefferson, and B. Marcus, Opt. Lett. **22,** 639 (1997).