

Some Effects of Fourier-domain Phase Quantization*

Abstract: If the Fourier transform of a function $g(x)$ is quantized, the function recovered by inverse transformation differs from $g(x)$. By means of a biased limiter model, the effects of Fourier-domain phase quantization are studied. Amplitude information is assumed fully retained, while phase is quantized to N equally spaced levels. The recovered function is shown to consist of several different contributions, the relative strengths of which depend on the number of phase quantization levels. Several specific examples are given. Motivation and interpretation are presented in terms of digitally constructed holograms.

1. Introduction

• The problem

When constructing holograms with a digital computer, it is necessary to deal with discrete arrays of points and to limit the physical quantities computed to a discrete set of possible values. The problem of sampling the hologram in a sufficiently fine array to avoid aliasing errors has been studied by others¹ and is not of concern here. Rather, we limit attention to the problem of quantization and its effects on the images formed by Fourier-transform holograms.

Stated in the most general possible way, the problem can be described as follows. Let the function $g(x)$ have a Fourier transform $G(\nu)$, i.e.,

$$G(\nu) = \int_{-\infty}^{\infty} g(x) \exp[-i2\pi x\nu] dx. \quad (1)$$

Let $G(\nu)$ be subjected to quantization, yielding a new spectrum $\hat{G}(\nu)$. How does the inverse Fourier transform $\hat{g}(x)$ of $\hat{G}(\nu)$ relate to the original function $g(x)$?

As we shall see, the problem stated above has direct relevance to digital holography. However, in its most general form it also has important implications in many other fields, including Fourier spectroscopy, pattern recognition and communication bandwidth compression. Our point-of-view will emphasize holography, but many of the results can be directly transcribed for other applications.

• Types of quantization

The Fourier transform $G(\nu)$ is in general complex-valued, and as a consequence, two distinctly different types of quantization can be envisioned. First, it is possible to quantize separately the real and imaginary parts of $G(\nu)$, each to a finite set of N allowable values. For example, we might choose the quantization levels to be equally spaced, in which case the quantization operation can be mathematically described (for odd N) by

$$\begin{aligned} \operatorname{Re} \{\hat{G}(\nu)\} &= -\left(\frac{N-1}{2}\right)a, \text{ when} \\ &\quad -\infty \leq \operatorname{Re} \{G(\nu)\} < -\left(\frac{N}{2} - 1\right)a \\ &\quad \vdots \\ &= -a, \text{ when } -\frac{3a}{2} \leq \operatorname{Re} \{G(\nu)\} < -\frac{a}{2} \\ &\quad \vdots \\ &= 0, \text{ when } -\frac{a}{2} \leq \operatorname{Re} \{G(\nu)\} < \frac{a}{2} \\ &\quad \vdots \\ &= +a, \text{ when } \frac{a}{2} \leq \operatorname{Re} \{G(\nu)\} < \frac{3a}{2} \\ &\quad \vdots \\ &= \left(\frac{N-1}{2}\right)a, \text{ when} \\ &\quad \left(\frac{N}{2} - 1\right)a \leq \operatorname{Re} \{G(\nu)\} < +\infty. \end{aligned} \quad (2)$$

* Work sponsored by the Office of Naval Research.
J. W. Goodman is with the Dept. of Electrical Engineering, Stanford University and A. M. Silvestri was with Technical Operations, Inc. at the time the manuscript was submitted; he is now with the Information Technology Corporation, Burlington, Massachusetts 01803.

with a similar equation for $\text{Im} \{\hat{G}(\nu)\}$. This quantization scheme is shown diagrammatically in Fig. 1(a), where the dashed lines represent quantization boundaries and the dots represent the possible quantized values.

A second quantization scheme is found by representing $G(\nu)$ in terms of its amplitude $|G(\nu)|$ and its phase $\phi(\nu)$,

$$G(\nu) = |G(\nu)| \exp [i\phi(\nu)]. \quad (3)$$

In this case we separately quantize $|G(\nu)|$ and $\phi(\nu)$ to discrete values, perhaps with different numbers of quantization levels allotted to amplitude and phase. For the special case of equally spaced quantization values, we can represent the quantization operation mathematically by

$$\begin{aligned} |\hat{G}(\nu)| &= 0, \quad \text{when } 0 \leq |G(\nu)| < \frac{a}{2} \\ &= a, \quad \text{when } \frac{a}{2} \leq |G(\nu)| < \frac{3a}{2} \\ &= 2a, \quad \text{when } \frac{3a}{2} \leq |G(\nu)| < \frac{5a}{2} \\ &\vdots \\ &= (M-1)a, \quad \text{when } \left(M - \frac{3}{2}\right)a \leq |G(\nu)| < \infty \end{aligned} \quad (4)$$

and

$$\begin{aligned} \phi(\nu) &= 0, \quad \text{when } -\frac{\pi}{N} \leq \phi(\nu) < +\frac{\pi}{N} \\ &= \frac{2\pi}{N}, \quad \text{when } \frac{\pi}{N} \leq \phi(\nu) < \frac{3\pi}{N} \\ &\vdots \\ &= \frac{N-1}{N} 2\pi, \quad \text{when } \frac{2N-3}{N} \pi \leq \phi(\nu) < \frac{2N-1}{N} \pi. \end{aligned} \quad (5)$$

This particular quantization scheme is illustrated in Fig. 1(b), where again the dashed lines represent quantization boundaries and the dots represent quantized values.

• Holographic motivation

Our emphasis in this paper will be on the second of the two quantization schemes. The reason for concentrating on this specific quantization method is its particular relevance to the binary Fraunhofer hologram technique described by Lohmann and his co-workers.^{2,3} The reader can find a detailed description of the technique used




Figure 1 Types of Fourier-domain quantization. (a) real and imaginary part quantization; (b) amplitude and phase quantization.

in computing and plotting binary Fraunhofer holograms in the references cited above. For our purposes here it suffices to note that when the computed hologram is plotted prior to photoreduction, the effective phase transmission through an elementary cell of the hologram is controlled by causing the plotter to stroke at one of a certain quantized set of positions within the cell. In addition, the effective amplitude transmittance is controlled by causing the stroke area to be one of a certain set of quantized values. The result is amplitude and phase quantization of the Fourier spectrum of the holographic image.

By way of further motivation, a new wavefront reconstruction device called the kinoform has recently been described.⁴ Again the reader can find a detailed discussion of the kinoform in the reference cited; for our purposes we simply note that a kinoform produces a wavefront which differs from the ideally desired wavefront in two respects. First, all amplitude information is quantized to a single value, and second, the phase information is quantized to a finite set of values. Thus a kinoform produces a wavefront which is quantized in both amplitude and phase.

• Restrictions of our analysis

The analysis to be presented in this paper cannot be regarded as a complete solution to the amplitude and phase quantization problem. Specifically, we shall focus our attention exclusively on phase quantization, with the effects of amplitude quantization ignored. The results remain meaningful, however, for two reasons. First, it is of some academic interest to understand in detail the effects of phase quantization without any accompanying amplitude quantization. Second, and more important, it has been found experimentally that in the Fourier plane, amplitude information is of far less importance than phase information, and that amplitude information can be fully retained or fully discarded without causing




Figure 2 Model for analyzing phase quantization.

any significant change in the image. This fact is particularly true when the object is diffuse, and has been amply demonstrated in experiments with kinoforms⁴ and with "phase only" acoustic holograms.⁵ More recently an analysis by Kermisch⁶ has verified that for diffuse objects the error that results from discarding amplitude information is small. In view of the relative insignificance of amplitude information in these important situations, we feel justified in ignoring the effects of amplitude quantization by assuming, for mathematical convenience, that amplitude information is fully retained. Equivalently, we assume an infinite number of amplitude quantization levels, but only a finite number of phase quantization levels.

2. Analysis

• The model

The Fourier spectrum $G(\nu)$, which may be written

$$G(\nu) = |G(\nu)| \exp [i\phi(\nu)], \quad (6)$$

is to be quantized in phase alone, yielding a new spectrum

$$\hat{G}(\nu) = |G(\nu)| \exp [i\hat{\phi}(\nu)]. \quad (7)$$

The type of phase quantization assumed is precisely that described by Eq. (5); i.e., the interval $(0, 2\pi)$ is divided into N equal quantization cells, each of width $2\pi/N$. Whenever $\phi(\nu)$ falls into a particular quantization cell, the quantized phase is assigned the value at the midpoint of the quantization interval.

The transformation from $\phi(\nu)$ to $\hat{\phi}(\nu)$ may be described in an analytically tractable fashion if the model of Fig. 2 is utilized. The function $\cos \phi(\nu)$ is applied at the input, which subsequently divides into N separate branches. In

the k th branch we place a phase delay of amount $k(2\pi/N)$, followed by an ideal biased limiter which has the input-output characteristic

$$g(y_k) = \begin{cases} 1, & y_k > \cos(\pi/N) \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

The output of the limiter is then multiplied by the complex number $\exp [ik(2\pi/N)]$. Finally, all branches are summed to yield an output $P(\nu)$. For $\cos \phi(\nu)$ at the input, we find $P(\nu)$ is given exactly by

$$P(\nu) = \exp [i\hat{\phi}(\nu)], \quad (9)$$

where $\hat{\phi}(\nu)$ is the value of the quantized phase described by Eq. (5).

Our procedure for finding the function $\hat{g}(x)$ produced after phase quantization will be the following: 1) find a convenient expression for $P(\nu)$ in terms of $\phi(\nu)$, 2) multiply $P(\nu)$ by $|G(\nu)|$, and 3) take the inverse Fourier transform of the product to obtain $\hat{g}(x)$.

• Nonlinear analysis applied to a single branch

We first analyze the nonlinear transformation introduced by the k th branch of the model of Fig. 2, i.e., the mapping of $\cos [\phi(\nu) - k(2\pi/N)]$ into $P_k(\nu)$, the k th component of $P(\nu)$. Our analysis rests on the "transform method" for analyzing nonlinear devices.⁷ Let

$$y_k(\nu) = \cos [\phi(\nu) - k(2\pi/N)] \quad (10)$$

represent the input to the k th biased limiter, and let $g[y_k(\nu)]$ represent its output. Then $g(y_k)$ may be written as

$$g(y_k) = \int_C f(i\xi) \exp [i\xi y_k] \frac{d\xi}{2\pi}, \quad (11)$$

where ξ is a complex variable, $f(i\xi)$ is defined by

$$f(i\xi) = \int_{-\infty}^{\infty} g(y_k) \exp [-i\xi y_k] dy_k, \quad (12)$$

and C is a contour chosen in the region of convergence of $f(i\xi)$. For the biased limiter of interest here, $f(i\xi)$ is given by⁸

$$f(i\xi) = \frac{\exp [-i\xi \cos(\pi/N)]}{i\xi} \quad (13)$$

and the region of convergence is the lower half of the ξ plane.

Next we insert Eq. (10) and Eq. (13) into Eq. (11), obtaining

$$g(y_k) = g_k = \int_C \frac{\exp [-i\xi \cos(\pi/N)]}{i\xi} \times \exp \{i\xi \cos [\phi - k(2\pi/N)]\} \frac{d\xi}{2\pi}. \quad (14)$$

The identity

$$\exp [i\xi \cos \theta] = \sum_{l=-\infty}^{\infty} (i)^l J_l(\xi) \exp [il\theta] \quad (15)$$

may then be used to reduce Eq. (14) to

$$g_k = \sum_{l=-\infty}^{\infty} C(k, l) \exp \{il[\phi - k(2\pi/N)]\}, \quad (16)$$

where the constants $C(k, l)$ are given by

$$C(k, l) = (i)^l \int_C \frac{\exp [-i\xi \cos (\pi/N)]}{i\xi} J_l(\xi) \frac{d\xi}{2\pi}. \quad (17)$$

The constants $C(k, l)$ could be evaluated by means of an appropriate complex integration, if desired. However, it is easier to use the following simpler method. Equation (16) must hold for any $\phi(\nu)$, and therefore in particular it must hold for the special case

$$\phi(\nu) = 2\pi x_0 \nu. \quad (18)$$

Since the $C(k, l)$ are independent of $\phi(\nu)$, if they can be found for the special case of Eq. (18), they will be known for all $\phi(\nu)$. For this special case we have as input to the k th biased limiter

$$y_k(\nu) = \cos [2\pi x_0 \nu - k(2\pi/N)]. \quad (19)$$

The output of the limiter is a periodic series of pulses, with period x_0^{-1} , pulse width $(Nx_0)^{-1}$, and with the center of the first pulse (in the positive ν direction) occurring at displacement $k(Nx_0)^{-1}$. The waveform is illustrated in Fig. 3.

The limiter output $g_k(\nu)$, being periodic, may be expressed as a Fourier series. In particular, we may write

$$g_k[\nu - (k/Nx_0)] = \sum_{l=-\infty}^{\infty} b_l \exp [i2\pi lx_0 \nu], \quad (20)$$

where, for the particular pulse train of Fig. 3, the complex Fourier coefficients b_l are given by

$$b_l = (1/N) \operatorname{sinc} (l/N), \quad (21)$$

where $\operatorname{sinc} x = \sin \pi x / \pi x$. It follows directly that

$$g_k(\nu) = \sum_{l=-\infty}^{\infty} (1/N) \operatorname{sinc} (l/N) \exp [i2\pi lx_0 \nu] \times \exp [-i2\pi(lk/N)]. \quad (22)$$

Comparison of Eq. (16) and Eq. (22) shows that

$$C(k, l) = (1/N) \operatorname{sinc} (l/N). \quad (23)$$

Finally, to find the output $P_k(\nu)$ of the k th branch we must multiply $g_k(\nu)$ by the complex constant $\exp [i2\pi k/N]$, which yields

$$P_k(\nu) = g_k(\nu) \exp [i2\pi k/N] = \sum_{l=-\infty}^{\infty} (1/N) \operatorname{sinc} (l/N) \times \exp [-i2\pi(k/N)(l-1)] \exp [il\phi]. \quad (24)$$




Figure 3 Periodic pulse train produced by the k th biased limiter.

• Total output of the phase quantizer

The total output of the phase quantizer may now be found by summing $P_k(\nu)$ over all N branches. Thus,

$$P(\nu) = \sum_{k=0}^{N-1} \sum_{l=-\infty}^{\infty} (1/N) \operatorname{sinc} (l/N) \times \exp [-i2\pi(k/N)(l-1)] \exp [il\phi]. \quad (25)$$

Interchanging orders of summation we find

$$P(\nu) = \sum_{l=-\infty}^{\infty} \operatorname{sinc} (l/N) \exp [il\phi] \left\{ \frac{1}{N} \sum_{k=0}^{N-1} \exp [-i2\pi(k/N)(l-1)] \right\}. \quad (26)$$

The summation in brackets is readily seen to be nonzero only for certain values of l ; specifically,

$$\frac{1}{N} \sum_{k=0}^{N-1} \exp [-i2\pi(k/N)(l-1)] = \begin{cases} 1, & l-1 = 0, \pm N, \pm 2N, \dots \\ 0, & \text{otherwise.} \end{cases} \quad (27)$$

It is therefore advantageous to change the index of summation to

$$m = (l-1)/N, \quad (28)$$

in which case we obtain

$$P(\nu) = \sum_{m=-\infty}^{\infty} \operatorname{sinc} [m + (1/N)] \times \exp [i(Nm + 1)\phi(\nu)]. \quad (29)$$

The total spectrum after quantization may now be found, noting $\hat{G}(\nu)$ is simply the product of $P(\nu)$ and $|G(\nu)|$. Thus,

$$\hat{G}(\nu) = \sum_{m=-\infty}^{\infty} \operatorname{sinc} [m + (1/N)] |G(\nu)| \times \exp [i(Nm + 1)\phi(\nu)]. \quad (30)$$

To find the x -domain function $\hat{g}(x)$ after quantization, we take the inverse Fourier transform of Eq. (30) to yield our final analytical result,

$$\hat{g}(x) = \sum_{m=-\infty}^{\infty} \text{sinc}[m + (1/N)]g_m(x), \quad (31)$$

where

$$g_m(x) = \int_{-\infty}^{\infty} |G(\nu)| \exp[i(Nm + 1)\phi(\nu)] \times \exp[i2\pi\nu x] d\nu. \quad (32)$$

The section to follow is devoted to interpretation and discussion of the result represented by Eqs. (31) and (32).

3. Interpretation, discussion and examples

• General comments

Our main analytical result, Eq. (31), demonstrates that the recovered function $\hat{g}(x)$ consists of a summation of several different contributions. The portion contributed by the $m = 0$ term we call the "primary image" term, since, as is readily verified, it contributes a function which is simply an attenuated version of the original object. Thus for $m = 0$ we find the image contribution to be

$$\text{sinc}(1/N) \cdot g(x). \quad (33)$$

As the number of quantization levels increases, the argument of the sinc function approaches zero, and the strength of the primary image increases, approaching that of the original function $g(x)$.

In addition to the primary image, there are contributions from the $m \neq 0$ terms which we refer to here as "false images" because, under certain circumstances they can produce separate and distinct images of the original object. As can be directly seen from our analytical results, the strength of the m th false image is attenuated by the factor $\text{sinc}[m + (1/N)]$. Since m is always an integer, and since the value of the sinc function for integer argument is zero, we see that as the number N of quantization levels increases, the strengths of the false image terms decrease.

In general, the shape of a false image is different from that of the primary image, since Eq. (32) demonstrates that the phase distribution across the spectrum of the m th false image is $(Nm + 1)\phi(\nu)$. The effect of phase multiplication in the Fourier domain is in general complicated, but several important conclusions can be drawn. First, there do exist functions with transforms having the property

$$(Nm + 1)\phi(\nu) = \phi(\nu). \quad (34)$$

When this is the case, the false image $g_m(x)$ will have exactly the same shape as $g(x)$. An example of a transform pair having this property is

$$g(x) = \begin{cases} 1 - |x|, & |x| \leq 1 \\ 0, & \text{otherwise,} \end{cases} \quad (35)$$

$$G(\nu) = \text{sinc}^2 \nu,$$

which has a zero-phase spectrum for all ν .

A second general effect of phase multiplication becomes evident when we consider a function $g(x)$ which may consist of a small nonzero object $g_1(x)$ in a larger, empty object field. If $g_1(x)$ is centered at the middle of the object field, we have

$$G(\nu) = G_1(\nu) = |G_1(\nu)| \exp[i\phi_1(\nu)], \quad (36)$$

but if $g_1(x)$ is shifted to be centered at a point x_0 away from the origin, we have

$$G(\nu) = |G_1(\nu)| \exp\{-i[2\pi x_0 \nu - \phi_1(\nu)]\}. \quad (37)$$

The effect of phase multiplication will then be to shift the centers of the false images to new points $(Nm + 1)x_0$ in the image field, possibly separating them completely from the primary image and from each other. Hence the origin of the term "false image." Examples of this phenomenon are shown in the next section. When effects of sampling in the x - and ν -domains are taken into account, these displaced false images become potential sources of aliasing errors.

Finally we mention the effects of phase multiplication when the object $g(x)$ is diffuse (i.e., has a random phase at each point x). For diffuse objects, Eq. (34) would never be satisfied for all ν , since the phase function $\phi(\nu)$ in the spectrum is in general very complicated [it can in fact be shown to be statistically distributed with uniform probability on the interval $(0, 2\pi)$]. In view of our comments in Section 1, however, the amplitude information $|G(\nu)|$ is relatively unimportant, and we could, with reasonable accuracy, approximate the spectrum of the m th false image in the following way:

$$\begin{aligned} G_m(\nu) &= |G(\nu)| \exp[i(Nm + 1)\phi(\nu)] \\ &\approx |G(\nu)|^{|Nm+1|} \exp[i(Nm + 1)\phi(\nu)] \\ &= \begin{cases} G^{Nm+1}(\nu), & Nm + 1 > 0 \\ [G^{|Nm+1|}(\nu)]^*, & Nm + 1 < 0. \end{cases} \end{aligned} \quad (38)$$

We conclude that, to a first approximation, the m th false image may be found as an $|Nm + 1|$ -order self convolution of the primary image. The effect of the conjugate sign in Eq. (38) is simply to mirror-reflect the convolution about the x -origin, and to conjugate the corresponding false image.

• Some examples

We now illustrate our analytical conclusions with some specific examples. Our first example concerns the effect




Figure 4 Effects of shifting $g(x)$, two quantization levels: (a) the original $g(x)$; (b) $\hat{g}(x)$ recovered with no shift; (c) $\hat{g}(x)$ recovered with shift $x_0 = 1/4$; (d) $\hat{g}(x)$ recovered with $x_0 = 1/2$; (e) $\hat{g}(x)$ recovered with $x_0 = 1$.

of shifting a small object in a larger object field. The object chosen is a simple one [see Fig. 4(a)],

$$g(x) = \text{rect}(x - x_0), \quad (39)$$

where the function $\text{rect}(x - x_0)$ is defined by

$$\text{rect}(x - x_0) = \begin{cases} 1, & |x - x_0| \leq \frac{1}{2} \\ 0, & \text{otherwise.} \end{cases}$$

We quantize phase to only two levels, 0 and π . Thus Eq. (31) becomes

$$\hat{g}(x) = \sum_{m=-\infty}^{\infty} \text{sinc}(m + \frac{1}{2}) g_m(x), \quad (40)$$




Figure 5 Effects of changing the number of quantization levels: (a) the original $g(x)$; (b) $\hat{g}(x)$ recovered when $N = 2$; (c) $\hat{g}(x)$ recovered when $N = 3$; (d) $\hat{g}(x)$ recovered when $N = 4$.

where the spectrum $G_m(\nu)$ of $g_m(x)$ is

$$G_m(\nu) = |\text{sinc } \nu| \exp[i(2m+1)\phi_1(\nu)] \times \exp[-i2\pi(2m+1)x_0\nu]. \quad (41)$$

The phase ϕ_1 is the phase of $\text{sinc } \nu$, and jumps periodically between 0 and π . It follows that

$$\exp[i(2m+1)\phi_1(\nu)] = \exp[i\phi_1(\nu)] \quad (42)$$

and, therefore, that

$$\hat{g}(x) = \sum_{m=-\infty}^{\infty} \text{sinc}(m + \frac{1}{2}) \times \text{rect}[x - (2m+1)x_0]. \quad (43)$$

Figures 4(b), (c), (d) and (e) show the recovered function $\hat{g}(x)$ when $x_0 = 0, \frac{1}{4}, \frac{1}{2}$ and 1, respectively. For any $x_0 > \frac{1}{2}$, the false images are distinct and separate from the primary image.

The effect of changing the number N of quantization levels is next considered for the particular function

$$g(x) = \text{rect}(x - 1) \quad (44)$$

shown in Fig. 5(a). The results are shown in Figs. 5(b), (c) and (d) for $N = 2, 3$ and 4, respectively. The conclusions can be summarized as follows: 1) as N increases,




Figure 6 Intensity of the false images as a function of N .

the strength of the primary image grows and the strength of the false images decreases; 2) as N increases, the false images move farther and farther away from the primary image, assuming that the object is not centered at the origin. Note that for the particular case $N = 3$, the $m = \pm 1$ false images, centered at $x = -2$ and $x = 4$,

do not have the same shape as the primary image, but rather are considerably spread out along the x -axis, a consequence of the fact that in this case the phase does not satisfy Eq. (34).

To develop a general feeling for the manner in which the false image strengths fall off with increasing N , we have plotted in Fig. 6 the value of $\text{sinc}^2 [m + (1/N)]$. The intensity of the m th order false image falls asymptotically as

$$I_m \sim 1/m^2 N^2 \quad (m \neq 0) \quad (45)$$

for large N .

4. Concluding remarks

The analysis presented represents an exact solution to the phase-quantization problem when the quantization levels are equally spaced. A slight modification of the phase delays and limiter thresholds in Fig. 2 would allow an exact solution for unequal quantization intervals.

The usefulness of the solution in predicting image degradations caused by phase quantization in holography is probably limited to the case of reasonably small N . As the number of phase quantization levels increases, eventually a point will be reached where amplitude quantization effects predominate.

References

1. W. T. Cathey, Jr., "The Effect of Finite Sampling in Holography," *Optik* **27**, 317 (1968).
2. A. W. Lohmann and D. P. Paris, "Binary Fraunhofer Holograms Generated by Computer," *Appl. Opt.* **6**, 1739 (1967).
3. B. R. Brown and A. W. Lohmann, "Computer Generated Binary Holograms," *IBM J. Res. Develop.* **13**, 160 (1969).
4. L. B. Lesem, P. M. Hirsch and J. A. Jordan, Jr., "The Kinoform, A New Wavefront Reconstruction Device," *IBM J. Res. Develop.* **13**, 150 (1969).
5. A. F. Metherell, "The Relative Importance of Phase and Amplitude in Acoustical Holography," in *Acoustical Holography, Vol. I* (A. F. Metherell, H. M. A. El-Sum and L. Larmore, editors), Plenum Press, New York 1969.
6. D. Kermisch, "Image Reconstruction from Phase Information Only," *J. Opt. Soc. Am.* **60**, 15 (1970).
7. D. Middleton, *Introduction to Statistical Communication Theory*, McGraw-Hill Book Co., Inc., New York 1960, pp. 105-119.
8. *Ibid.*, p. 116.

Received December 12, 1969