

Sampled-grating and crossed-grating models of moire patterns from digital imaging

John Krumm, MEMBER SPIE
Steven A. Shafer, MEMBER SPIE
The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Abstract. Traditional "crossed-grating" moire as well as the newer "sampled-grating" (scanning) moire have proved to be effective methods of shape measurement. There is speculation that the moire patterns of a sampled grating, which are due to aliasing, can be modeled with crossed gratings. We compare the two by writing notationally consistent models of each and show that while crossed gratings can correctly predict the frequencies of a sampled grating, they cannot correctly predict the amplitudes. Our sampled-grating model is a new formulation that accounts for multiple stages of sampling and transmission. We show how neglecting multiple stages can lead to mistakes in moire analysis. We demonstrate our models with an experiment using a digital imaging system.

Subject terms: moire patterns; aliasing; scanning moire; digital imaging; crossed gratings; sampled gratings; CCD image sensors; Nyquist sampling theorem; image processing.

Optical Engineering 30(2), 195-206 (February 1991).

CONTENTS

1. Introduction
2. Moire patterns from crossed gratings
3. Moire patterns from sampled gratings
 - 3.1. Model of multiple stage reconstruction and sampling
 - 3.2. Moire patterns
4. Experimental verification of sampled-grating moire
5. Crossed-gratings models of sampled gratings
 - 5.1. Approximation with two Ronchi rulings
 - 5.2. Approximation with two arbitrary gratings
 - 5.3. Visual comparison
6. Conclusion
7. Acknowledgments
8. References

1. INTRODUCTION

Moire patterns can be caused by crossed gratings or by the discrete sampling of a single grating. The crossed-gratings case can be seen in everyday surroundings in layers of fences or screens. Figure 1 shows a superposition of two crossed gratings that give rise to a moire pattern. If one of the crossed gratings is distorted, the moire pattern will be distorted too. This phenomenon has been extensively exploited in metrology applications. One of the constituent gratings is projected onto a distorted object and observed through a second, flat grating. Since the resulting moire pattern is sensitive to small distortions in the object grating, accurate measurements of the object's shape can be derived. The object's shape is a function of both the geometric layout of the original, undistorted gratings and the resulting moire pattern. Post¹ has shown that an analysis of the profiles of the crossed gratings can be used to sharpen the moire fringe patterns, making it easier to track the fringes and thus easier to determine the shape of the object. Therefore, both the geometry and brightness profiles of the moire method are important to consider.

Moire patterns can also be caused by aliasing in the discrete sampling of a single grating. The sampled-grating case can often be seen on television when a relatively high-spatial-frequency grating, say a striped shirt, is shown. Idesawa, Yatagai, and Soma² have shown that the flat grating in traditional, crossed-grating moire metrology can be replaced by a digital camera, still maintaining observable moire patterns using only a single, object grating. We will distinguish these two methods by calling the former "crossed-grating" moire and the later "sampled-grating" moire. The moire patterns from both methods are similar, and it has been shown that sampled-grating moire can automate many time-consuming, subjective procedures of the traditional, crossed-grating paradigm.³

Idesawa et al. assert that the camera can be modeled by a second grating, thus allowing the application of a considerable amount of past study on crossed gratings to sampled gratings. However, this is not completely accurate. The two phenomena are fundamentally different, and it has never been shown explicitly how well crossed gratings can model the effects of a digital camera. The moire patterns that result from crossed gratings are due to the heterodyning of their respective spatial frequencies, and such patterns require at least two gratings. Moire patterns from digital cameras (or any periodic, discrete, sampling device) come from aliasing of the spatial frequencies of the imaged grating and may be produced with a single grating. There have been moire pattern models that explicitly account for this discrete sampling, but they do not account for lens blur or multiple stages of filtering and sampling as we do.

Our goals in this paper are to develop a useful model of sampled-grating moire and then to compare it to crossed-grating moire. Toward these ends, we first develop a model of crossed-grating moire that will be notationally consistent with the sampled-grating model. Both models are formulated in the Fourier domain for convenience. Our sampled-grating model is a new formulation that accounts for the video transmission (including lens blur) and multiple sampling stages in most digital imaging systems. We use our multiple stage sampling model to dem-

Paper 2835 received Nov. 16, 1989; revised manuscript received Sept. 24, 1990; accepted for publication Oct. 7, 1990.
© 1991 Society of Photo-Optical Instrumentation Engineers.

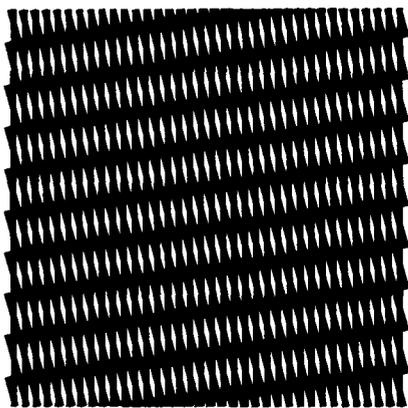


Fig. 1. Crossed gratings making a moire pattern.

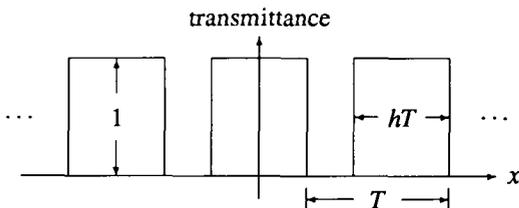


Fig. 2. Square wave transmittance profile.

onstrate how ignoring multiple sampling stages can lead to mistakes in moire analysis. We present results of an experiment that partially verifies the model used in the sampled-grating formulation. In our comparison of the two types of moire patterns, we examine the ability of crossed gratings to model sampled gratings by developing a model of sampled-grating moire patterns using crossed gratings, and we verify the model qualitatively with a simple experiment.

The next section contains our crossed-grating model, followed by our sampled-grating model in Sec. 3. We verify the sampled-grating model in Sec. 4, and Sec. 5 presents a comparison of the two models.

2. MOIRE PATTERNS FROM CROSSED GRATINGS

In this section we derive the Fourier transform of the superposition of an arbitrary number of flat, periodic, crossed gratings with arbitrary profiles. Gratings are usually considered as transparencies with regions of transmittance varying between zero and one. We can express the transmittance of any infinite, 2-D, periodic grating whose transmittance profile satisfies the Dirichlet conditions and that is oriented so that its lines run perpendicular to the x axis, as the following Fourier series:

$$g(x,y) = \sum_{n=-\infty}^{\infty} c_n \exp\left(\frac{jn2\pi x}{T}\right), \tag{1}$$

where T is the period of the grating and the c_n are the complex coefficients of the Fourier series of the 1-D transmission profile along the x axis. We note that any physically realizable transmittance grating will satisfy the Dirichlet conditions.

If the profile of the grating $g(x,y)$ is a square wave, as shown in Fig. 2, the c_n will be given by $c_n = h \text{sinc}(nh)$, where $\text{sinc}(x) = \sin(\pi x)/\pi x$ and h is the fraction of a period that is transparent. If $h = 1/2$, then the grating is called a Ronchi ruling, shown in Fig. 3.

The Fourier transform of $g(x,y)$ is

$$G(f_x, f_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x,y) \exp[-j2\pi(f_x x + f_y y)] dx dy. \tag{2}$$

The frequencies f_x and f_y are in units of cycles/unit length. We use uppercase letters to indicate the Fourier transforms of corresponding lowercase, spatial functions. Substituting the Fourier series for the general grating $g(x,y)$ into the Fourier transform equation gives

$$G(f_x, f_y) = \sum_{n=-\infty}^{\infty} c_n \delta\left(f_x - \frac{n}{T}, f_y\right). \tag{3}$$

This is a set of evenly spaced, Dirac δ -functions spread along the f_x axis. The volumes of the δ 's are given by the Fourier series coefficients. Figure 4 shows a plot of the Fourier transform of the Ronchi ruling in Fig. 3. The dots show the position of the δ -functions, with the corresponding value of n in parentheses above each. The area of a dot is proportional to the magnitude of its coefficient. For a Ronchi ruling, however, every δ with an even, nonzero n will have a coefficient of zero. We show these with small dots.

We will call the orientation of $g(x,y)$ zero because the δ 's of its Fourier transform occur along the f_x axis. If $g(x,y)$ is rotated around the origin by an angle θ , its Fourier transform rotates in exactly the same way. This is shown in Figs. 5 and 6, where we show a Ronchi ruling rotated 15° and its Fourier transform. The dots in Fig. 6 are labeled like those in Fig. 4. If $g_i(x,y)$ is a rotated version of $g(x,y)$, its Fourier transform is

$$G_i(f_x, f_y) = \sum_{n_i=-\infty}^{\infty} c_{n_i} \delta\left(f_x - \frac{n_i}{T_i} \cos\theta_i, f_y - \frac{n_i}{T_i} \sin\theta_i\right), \tag{4}$$

where we have begun subscripting with i to account for multiple gratings. Only the positions of the δ 's change, not the coefficients.

If two of these transmittance gratings, $g_0(x,y)$ and $g_1(x,y)$, are superimposed, the resulting net transmittance will be given by the product of the constituent transmittances, $g_0(x,y)g_1(x,y)$. The Fourier transform of the net transmittance will be the 2-D convolution of the Fourier transforms, $G_0(f_x, f_y) * G_1(f_x, f_y)$. We show the two Ronchi rulings of Figs. 3 and 5 crossed in Fig. 7. The Fourier transform of the crossed gratings is shown in Fig. 8. The parameters for the Ronchi rulings in the figures are $T_0 = T_1 = 1$, $\theta_0 = 0$, and $\theta_1 = 15^\circ$.

We can think of the convolution operation graphically by considering the Fourier transforms of the two constituent gratings in Figs. 4 and 6. The first step in a 2-D convolution is to flip either one of the Fourier transforms around both the f_x and f_y axes. Since the Fourier transforms of the gratings as we have defined them are symmetric lines of δ 's, this operation leaves the function unchanged. The next step is to record the value of the product of the two functions as a function of the offset of one of them. If we slide, say, the Fourier transform of the rotated grating around in (f_x, f_y) space over the Fourier transform of the unrotated grating, we see that the product will only be nonzero when two δ 's overlap, at which point the product will be the product of the coefficients of the overlapping δ 's. As the Fourier transform of the rotated grating is moved, the position of its center δ will indicate the value of the offset. Since two different lines cannot intersect at more than one point, it is clear that there

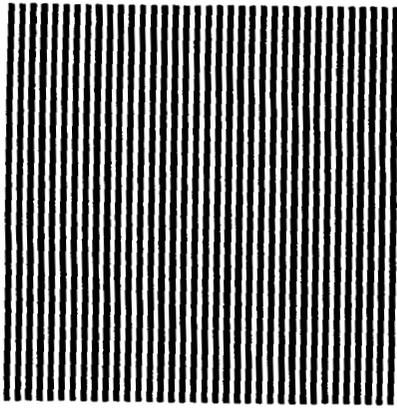


Fig. 3. Ronchi ruling at 0°.

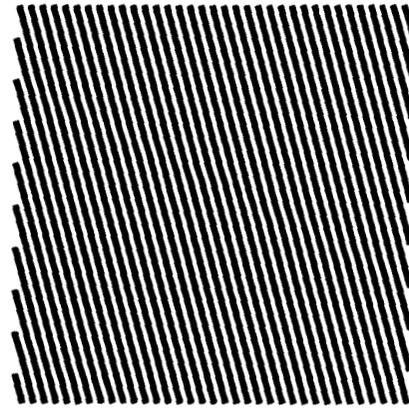


Fig. 5. Ronchi ruling at 15°.

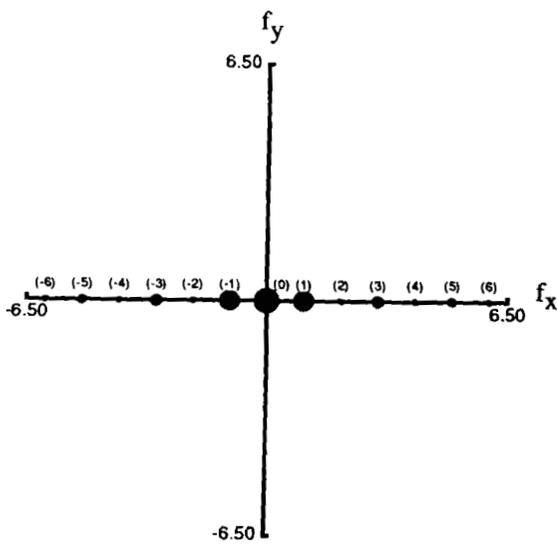


Fig. 4. Fourier transform of Fig. 3.

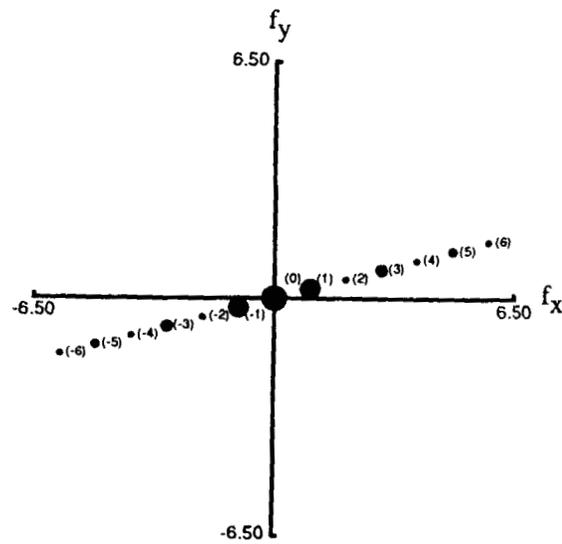


Fig. 6. Fourier transform of Fig. 5.

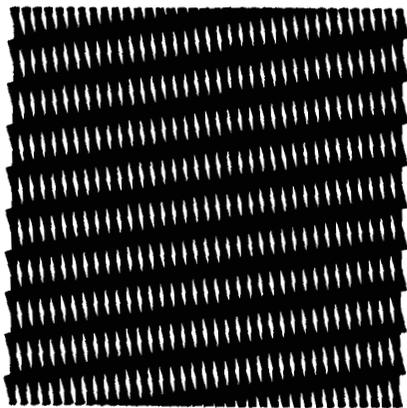


Fig. 7. Crossed Ronchi rulings at 0° and 15°.

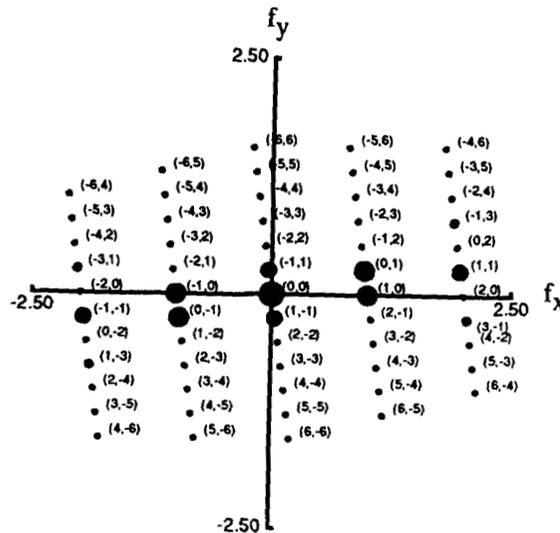


Fig. 8. Fourier transform of Fig. 7.

will be, at most, only one pair of overlapping δ 's for any given offset. Thus, we overlap every possible pair of δ 's, recording the product of their coefficients at the position of the center peak of the sliding Fourier transform.

If $g_i(x,y)$ has period T_i , angle θ_i , and Fourier series coefficients c_{n_i} , then the Fourier transform of two superimposed gratings will be

$$G_0(f_x, f_y) * G_1(f_x, f_y) = \sum_{n_0=-\infty}^{\infty} \sum_{n_1=-\infty}^{\infty} c_{n_0} c_{n_1} \delta\left(f_x - \frac{n_0}{T_0} \cos\theta_0 - \frac{n_1}{T_1} \cos\theta_1, f_y - \frac{n_0}{T_0} \sin\theta_0 - \frac{n_1}{T_1} \sin\theta_1\right). \quad (5)$$

In general, the locations of the δ 's in frequency space are given by vectors

$$(f_x, f_y) = \left(\frac{n_0}{T_0} \cos\theta_0 + \frac{n_1}{T_1} \cos\theta_1, \frac{n_0}{T_0} \sin\theta_0 + \frac{n_1}{T_1} \sin\theta_1\right) \quad (6)$$

for $n_0, n_1 = -\infty, \dots, -1, 0, 1, \dots, \infty$. This vector interpretation was suggested by Bryngdahl.⁴ The ordered pairs in the Fourier transform in Fig. 8 show (n_0, n_1) . The amplitudes of the corresponding complex exponentials in the spatial domain are given by $c_{n_0} c_{n_1}$. A moire pattern is obvious when a δ with sufficient amplitude falls closer to the origin than the fundamental frequencies of either of the constituent gratings. In the case of Fig. 8, the fundamental frequency of both gratings is 1 cycle/unit length. The lowest frequency shown in the Fourier transform of the superimposed gratings is for $(n_0, n_1) = \pm(-1, 1)$, whose frequency is about 0.26 cycles/unit length in the direction of 97.5° . These are the frequency and direction of the dark, nearly horizontal bands seen in the moire pattern in Fig. 7.

For every relatively prime pair (n_0, n_1) in the Fourier transform of the superimposed gratings (Fig. 8), there will be a set of δ 's that fall on a line through the origin of frequency space and whose n 's are given by (mn_0, mn_1) , $m = -\infty, \dots, -1, 0, 1, \dots, \infty$. These lines of δ 's in frequency space form periodic patterns in the spatial domain because their harmonics are spaced at equal intervals. The line of δ 's whose relatively prime pair is $(n_0, n_1) = (-1, 1)$ in Fig. 8 is responsible for the obvious moire pattern in Fig. 7 because it happens to have a low fundamental frequency with a relatively large amplitude. There are, however, an infinite number of relatively prime pairs, all of which serve as the fundamental frequency of a line of δ 's and all of which could be considered moire patterns. The well-known method of partial sum extraction⁵ is to pick a relatively prime pair and sum up all of the frequency components along that line. Normally, higher absolute values of n have associated with them lower values of c_n , so these higher frequency moire patterns have small amplitudes and are much less noticeable. Lines of δ 's with higher fundamental frequencies are sometimes ignored as noise.⁵

If k transmittance gratings are superimposed, Eq. (5) for the Fourier transform generalizes to

$$G_0(f_x, f_y) * G_1(f_x, f_y) * \dots * G_{k-1}(f_x, f_y) = \sum_{n_0=-\infty}^{\infty} \sum_{n_1=-\infty}^{\infty} \dots \sum_{n_{k-1}=-\infty}^{\infty} \left[\left(\prod_{i=0}^{k-1} c_{n_i} \right) \delta\left(f_x - \sum_{i=0}^{k-1} \frac{n_i}{T_i} \cos\theta_i, f_y - \sum_{i=0}^{k-1} \frac{n_i}{T_i} \sin\theta_i\right) \right]. \quad (7)$$

This is a distribution of δ -functions at

$$(f_x, f_y) = \left(\sum_{i=0}^{k-1} \frac{n_i}{T_i} \cos\theta_i, \sum_{i=0}^{k-1} \frac{n_i}{T_i} \sin\theta_i\right), \quad (8)$$

with coefficients of c_{n_i} . Equation (5) serves as our model of two crossed gratings, and Eq. (7) is our model of k crossed gratings.

3. MOIRE PATTERNS FROM SAMPLED GRATINGS

The previous section shows how moire patterns develop from crossed gratings. Digital cameras cause moire patterns by discrete sampling. The development of moire patterns in sampled images is fundamentally different from crossed gratings. Sampled-grating moire patterns occur when a periodic pattern is sampled at a rate less than twice the highest spatial frequency of the periodic pattern. This is the familiar phenomenon of aliasing. Spatial frequencies that are greater than half the Nyquist frequency are aliased into lower frequencies. Thus, much like the crossed-gratings case, high frequency patterns beget patterns of lower frequencies.

At least two sampled-grating moire models have appeared in the literature. Bell and Koliopoulos⁶ consider discrete, 2-D sampling of a projected, cosinusoidal grating and develop a rule for producing unambiguous moire fringes. Cetica, Francini, and Bertani⁷ specialize Bell and Koliopoulos's result to one dimension and report an experiment with a linear, photodiode array. Neither model takes into account the effects of camera blurring or the multiple stages of sampling that are common in digital imaging systems. We develop a model that accounts for these factors in the next subsection and then apply it to a flat, periodic grating. It is seen that blurring and multiple-stage sampling have a significant impact on the brightness and geometry of moire fringes.

3.1. Model of multiple stage reconstruction and sampling

We give a brief formulation of a single-stage, image sampling system to demonstrate aliasing and to establish our notation; similar formulations can be found in most texts on image processing. The model can be easily explained in terms of a digital camera recording an image, although we show how the model can be applied, recursively, to subsequent stages of image transfer and display. The input to the sampling system is an arbitrary image of intensities $g(x,y)$. Before sampling, it is blurred by a convolution with $p(x,y)$ —the point spread function (PSF). For the case of a camera, $g(x,y)$ is the ideal input image and $p(x,y)$ represents the effects of misfocus, diffraction, and pixel response, all of which can be described by convolution. (We are neglecting geometric and photometric scale factors here.) The image sampled by the photosensitive camera cells is then $p(x,y) * g(x,y)$. The Fourier transform of this image is $P(f_x, f_y)G(f_x, f_y)$, where $P(f_x, f_y)$ is the Fourier transform of the PSF and $G(f_x, f_y)$ is the Fourier transform of the ideal image. The camera's sampling grid can be represented by a grid of δ -functions, each centered on a photosensitive cell on the image plane. If the rectangular spacing of the cells in the x and y directions is T_x and T_y , the sampling grid is given by

$$\sum_{n_x=-\infty}^{\infty} \sum_{n_y=-\infty}^{\infty} \delta(x - n_x T_x, y - n_y T_y).$$

Here we have neglected the finite size of the sampling array. This can be accounted for by multiplying the array by a zero-one function in the shape of the sampling grid (usually a rectangle). Bell and Koliopoulos do this in their analysis of moire

patterns, but for our purposes it only leads to unnecessary complications. This is because although a truncated sampling grid can have a significant impact on the frequency-domain representation of a sampled-grating, its spatial-domain representation (where moire pattern analysis is ultimately carried out) is simply a truncated version of the infinite moire pattern. The sampled image is given by the product of the blurred image and the sampling array as

$$[p(x,y) * g(x,y)] \sum_{n_x=-\infty}^{\infty} \sum_{n_y=-\infty}^{\infty} \delta(x-n_x T_x, y-n_y T_y) \quad (9)$$

The Fourier transform of this is the convolution of the Fourier transform of the blurred image and the Fourier transform of the sampling grid:

$$[P(f_x, f_y)G(f_x, f_y)] * \frac{1}{T_x T_y} \sum_{n_x=-\infty}^{\infty} \sum_{n_y=-\infty}^{\infty} \delta\left(f_x - \frac{n_x}{T_x}, f_y - \frac{n_y}{T_y}\right) \quad (10)$$

This equation indicates that the Fourier transform of the blurred image is repeated at intervals of $(1/T_x, 1/T_y)$ in frequency space. These copies of the Fourier transform are called the *spectral orders* of the function $g(x,y)$, with the 0th spectral order centered at the origin of frequency space.

Figure 9 shows the Fourier transform of a function $g(x,y)$ that is to be sampled. In this case, the function happens to be a grating. The δ -functions are shown as \times 's. After sampling, the Fourier transform looks like Fig. 10. The original Fourier transform has been repeated into an infinite number of spectral orders. The 0th spectral order is shown as \times 's, while the others are shown as \bullet 's. The centers of the spectral orders are shown with larger symbols.

Interpolation is applied to the samples to extract a continuous estimate of the original image. This can be formulated in the spatial domain as a convolution of an interpolation function, usually some kind of narrow pulse, with the samples. The convolution operation means that the interpolation function is repeated at every sample point and modulated by the values of the discrete samples. In frequency space, interpolation amounts to multiplying the distribution of spectral orders by the Fourier transform of the interpolation function. This function may be chosen to reduce the effects of aliasing or noise, but it is always bandlimited by the Nyquist frequencies, $[\pm 1/(2T_x), \pm 1/(2T_y)]$. These limits are shown as a rectangle in Fig. 10. If frequencies from spectral orders other than the 0th fall inside the nonzero region of the interpolation function, aliasing will occur. This has happened in Fig. 10.

For a typical digital image capture and storage system, this model is incomplete. Once the image is sampled into charge levels in the camera, the samples are converted into an analog, video signal for transmission, then amplified, filtered, and re-sampled by an analog-to-digital converter. Thus, one set of samples is turned into another, and the sampling rates may not be equal. We show the frequency domain effects of a second stage of sampling in Figs. 11 and 12. The 0th spectral order in these two figures (shown with \times 's) contains all of the frequencies from inside the interpolation rectangle of Fig. 10. If the second stage has smaller sampling periods than the first (i.e., higher sampling frequencies), the spectral orders after the second stage will be spread in frequency space such that no overlap occurs, as in Fig. 11. The box indicates the Nyquist bounds of the second sampling stage. If, on the other hand, the sampling rates of the

second stage are lower than those of the first stage, more aliasing could occur. This is the case in Fig. 12. We see that the Nyquist bounding box includes some frequencies from neighboring spectral orders and excludes some from the 0th order.

We can think of such a system as a multiple stage reconstruction and sampling process. We assume the system has a series of spatial sampling stages separated by linear components. As shown in Fig. 13, each pass produces a new set of samples that are fed back through the system. The index i is incremented for each pass. The initial input to the system is the original, ideal image from the scene. It is first blurred and then sampled. The samples are then sent back through the system, where they are reconstructed into a continuous image and blurred (in the same step), and then sampled again. The process can be repeated over and over.

As mentioned, the reconstructed image is bandlimited by the Nyquist frequencies. Since a pass through the loop does not impose any higher frequencies on the reconstructed image, the final reconstructed image, no matter how many times the loop has been traversed, will be bandlimited by the minimum Nyquist frequency encountered during any pass. Thus, no aliasing can occur after the minimum sampling frequency has been encountered. Alternatively, if the sampling rate decreases between any two passes, there may be a complex interaction of frequencies as aliased frequencies become aliased themselves. These arguments apply independently to the x and y directions.

The process can be characterized mathematically by putting subscripts on Eq. (10), which describes one pass in the frequency domain:

$$S_i(f_x, f_y) = [P_i(f_x, f_y)S_{i-1}(f_x, f_y)] * \frac{1}{T_{x_i} T_{y_i}} \sum_{n_{x_i}=-\infty}^{\infty} \sum_{n_{y_i}=-\infty}^{\infty} \delta\left(f_x - \frac{n_{x_i}}{T_{x_i}}, f_y - \frac{n_{y_i}}{T_{y_i}}\right) \quad (11)$$

We have replaced G with S to stand for "samples." One pass may be interpreted as follows: The samples $S_{i-1}(f_x, f_y)$ are converted to a continuous signal by filtering with $P_i(f_x, f_y)$. P_i represents not only the interpolation function but the amplification and filtering that occur before the i th sampling stage. The analog signal is then sampled with x and y spacings of T_{x_i} and T_{y_i} , giving a new set of samples in $S_i(f_x, f_y)$.

The process of creating a stored, digital image of a scene may be cast in this multiple stage model. We start with the ideal image $G(f_x, f_y)$, which serves as $S_0(f_x, f_y)$, the initial sampled image. While this is not really a sampled image, it is the appropriate place to start the looping. The first sampling loop ($i=1$) is the digital camera. $P_1(f_x, f_y)$ accounts for lens defocus, diffraction, and the area integration of the pixels spaced at intervals of T_{x_1} and T_{y_1} . The resulting samples are manifest as charges present on the image plane. The second sampling loop ($i=2$) is the analog-to-digital conversion. This stage starts by converting the image plane charges to a 1-D analog signal using $P_2(f_x, f_y)$. P_2 takes into account the cascaded effects of interpolating the discrete charges into an analog signal as well as any filtering of the analog signal that occurs before the conversion to digital. Since the image plane samples are usually scanned out row by row, there is no blurring in the y direction in this step. For the same reason, $T_{y_2} = T_{y_1}$. T_{x_2} and T_{x_1} are not necessarily equal because the sampling rate on the image plane may be different from the sampling rate of the digitizer. We could go around the loop once more to model the conversion to

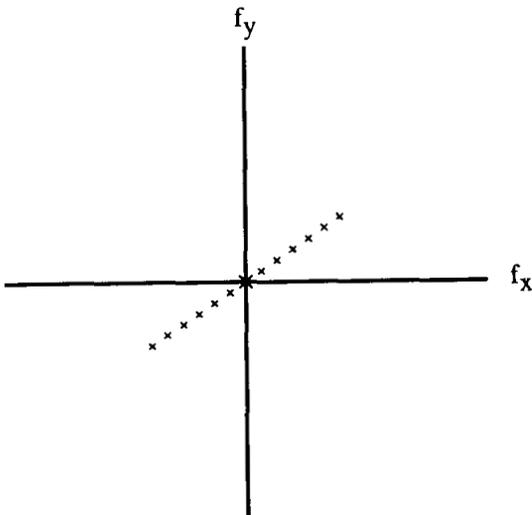


Fig. 9. Fourier transform of grating.

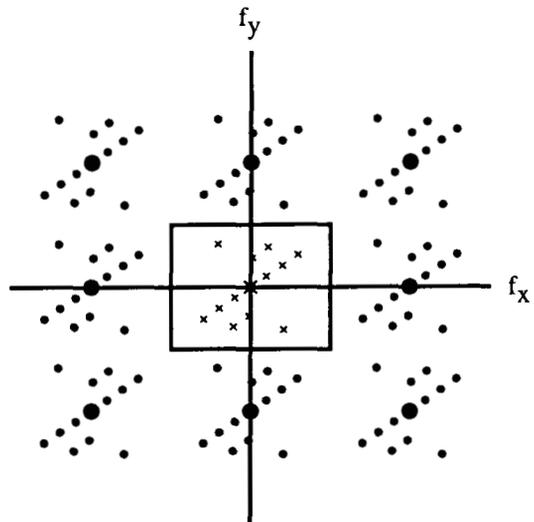


Fig. 11. Spectral orders after second stage sampling with no additional aliasing because of high frequency sampling.

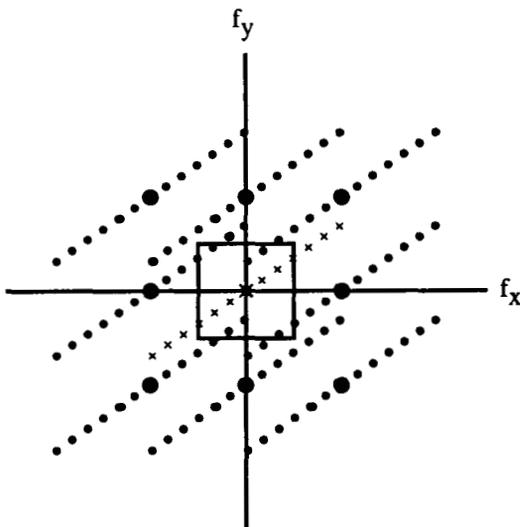


Fig. 10. Spectral orders after first stage sampling.

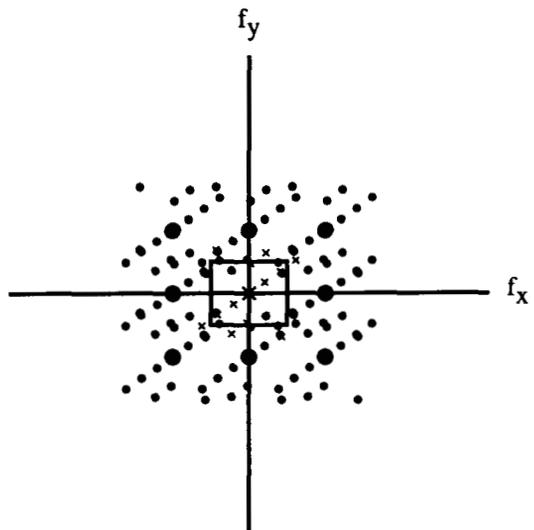


Fig. 12. Spectral orders after second stage sampling showing additional aliasing because of low frequency sampling.

discrete charges on the phosphors of a display device. Yet another pass could model certain facets of human vision since the discrete sampling of the human fovea can produce moire patterns, suggesting that it, too, has a fairly regular sampling grid (although not rectangular⁸). This would require that the model be modified to accommodate a nonrectangular sampling grid.

The impossibility of aliasing after encountering the minimum sampling rate can be seen in Eq. (11). Suppose that in pass $i = 1$ (e.g., the sampling on the camera's image plane), the sampling periods T_{x1} and T_{y1} are larger than the sampling periods of the digitizer, T_{x2} and T_{y2} in pass $i = 2$. That is, the sampling frequency of the first stage is lower than that of the second stage. The Fourier transform of the first set of samples will be a distribution of spectral orders at intervals of $(1/T_{x1}, 1/T_{y1})$. When this signal is reconstructed in preparation for resampling in the second stage, it will be bandlimited by the Nyquist frequencies of the first sampling stage, $[\pm 1/(2T_{x1}), \pm 1/(2T_{y1})]$. The Fourier transform of the second set of samples will again be a distribution of spectral orders, centered at intervals of $(1/T_{x2}, 1/T_{y2})$. Since these centers are farther apart than the maximum dimensions of

the Fourier transform of the reconstructed image, there will be no overlap in frequencies, and thus no aliasing.

If any stage in this system causes aliasing, the system is not linear, since aliasing introduces new frequencies. However, once the minimum sampling rate stage has been passed, the system may be treated as linear. The practical effect of this is that the transfer functions P_i after the minimum sampling rate stage may be multiplied into one effective transfer function to describe all subsequent stages.

3.2. Moire patterns

We can use the multiple stage sampling and reconstruction model above to predict the development of moire patterns in an image sampling system. Equation (11) gives the expression for one pass through the sampling system with an arbitrary input function $S_{i-1}(f_x, f_y)$. Equation (4) is a grating with an arbitrary profile at an angle θ . If we substitute this grating equation (with $i = 0$) into the sampling equation for $S_{i-1}(f_x, f_y)$, we see that the result of one stage of sampling on a rotated grating is

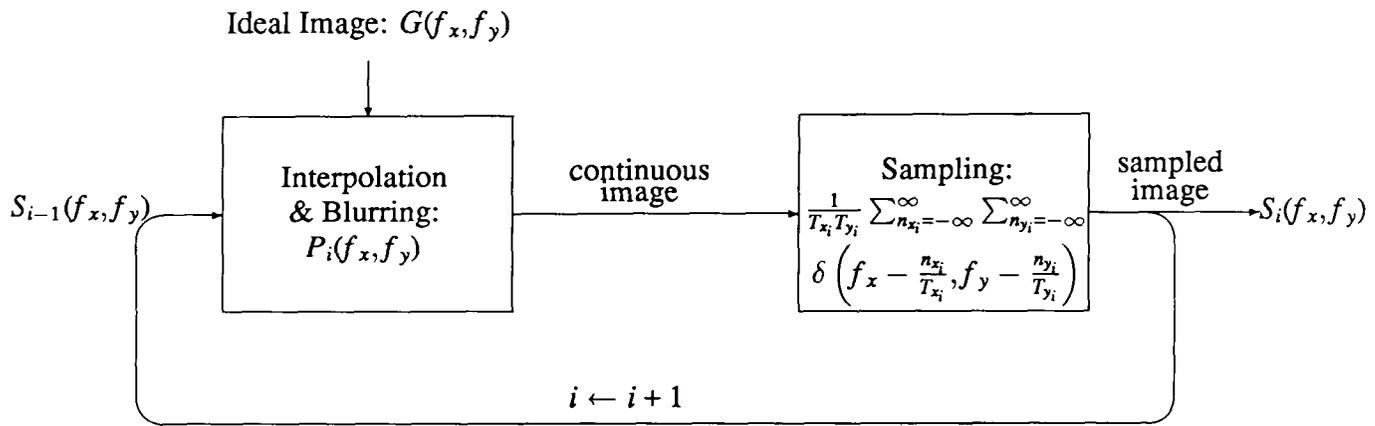


Fig. 13. Block diagram illustrating multiple stage reconstruction and sampling.

$$S_i(f_x, f_y) = \frac{1}{T_{x_i} T_{y_i}} \sum_{n_{x_i}=-\infty}^{\infty} \sum_{n_{y_i}=-\infty}^{\infty} \sum_{n_0=-\infty}^{\infty} c_{n_0} P_1 \left(\frac{n_0}{T_0} \cos \theta_0, \frac{n_0}{T_0} \sin \theta_0 \right) \times \delta \left(f_x - \frac{n_0}{T_0} \cos \theta_0 - \frac{n_{x_i}}{T_{x_i}}, f_y - \frac{n_0}{T_0} \sin \theta_0 - \frac{n_{y_i}}{T_{y_i}} \right). \quad (12)$$

This is a distribution of δ 's whose locations are given by

$$(f_x, f_y) = \left(\frac{n_0}{T_0} \cos \theta_0 + \frac{n_{x_i}}{T_{x_i}}, \frac{n_0}{T_0} \sin \theta_0 + \frac{n_{y_i}}{T_{y_i}} \right). \quad (13)$$

The grating variables are given by those variables with a 0 subscript. The δ 's of the grating's Fourier transform are indexed by n_0 , while the spectral orders are indexed by (n_{x_i}, n_{y_i}) . The coefficients of the δ -functions are independent of the particular spectral order but depend on the relative displacement from the center of the spectral order. As with crossed gratings, a moire pattern will be especially obvious if the first fundamental frequency from another spectral order falls closer to the origin than the fundamental frequency of the 0th spectral order, because the first fundamental usually has significant amplitude. See Rosenfeld and Kak⁹ for a simple example.

In most digital imaging systems, this image will be reconstructed and sampled again. This process can have a significant effect on the moire patterns observed in the final image. Given a suitable set of geometric parameters (i.e., the rotation of the grating and the sampling intervals), there could actually be moire patterns developed in the first stage that are aliased into other moire patterns in subsequent stages. This was shown in Figs. 10 and 12, where frequencies that were aliased in the first sampling stage were again aliased in the second stage. These extra, aliased frequencies cannot be accounted for in a single-stage sampling model, and they would likely confound current moire fringe analysis techniques.

In addition to affecting the geometry of moire fringes, multiple sampling stages also affect their brightness. Each sampling stage has associated with it its own PSF that describes the combined effects of interpolated reconstruction and transmission to the next sampling stage. Mathematically, the applications of the sampling grids and PSF must be interleaved, which means that in general, no single PSF can be used to characterize multiple sampling stages. Previous models of sampled-grating moire have used either a simple, single-stage model with only one PSF or else no PSF at all. As was shown in Post's work on fringe sharpening,¹ a careful analysis of the brightness profiles of moire

fringes can lead to more robust detection of the fringes. Thus, it is important to have an accurate model of the sampling system's effect on the fringe profiles.

4. EXPERIMENTAL VERIFICATION OF SAMPLED-GRATING MOIRE

We performed an experiment to verify our model of sampled-grating moire patterns by taking a digital image of a square-wave grating using a standard CCD camera and digitizer. A 64×64 subsection of the image is shown in Fig. 20. We compared the discrete Fourier transform (DFT) of the actual image with the predicted Fourier transform. The experiment is described in more detail in another paper¹⁰ and is only summarized here. The system could be modeled in two stages corresponding to CCD sampling and video digitizer sampling. We found it convenient to express linear dimensions in terms of "y-pixels"—the vertical distance between CCD pixel centers. Using this unit, along with the fact that the image is scanned from the CCD in rows, both vertical sampling periods were equal to one, i.e., $T_{y_1} = T_{y_2} = 1$ y-pixel. The imaged grating had a period of $T_0 = 1.21$ y-pixels, an angle of $\theta_0 = 88.38^\circ$, and a duty ratio $h = 0.48$. The CCD sampling periods were $(T_{x_1}, T_{y_1}) = (1.70, 1.00)$ y-pixels, and the digitizer sampling periods were $(T_{x_2}, T_{y_2}) = (1.28, 1.00)$ y-pixels. Since our setup happened to have its lowest x sampling rate on the CCD, we did not have to concern ourselves with aliasing in the digitization stage.

Because we were purposefully inducing aliasing in the system, the PSFs of the system components could not be cascaded into one. We had to separately account for the PSFs of the lens/CCD combination (P_1) and the interpolation/video signal filtering (P_2) that occurred after the charges were read off the CCD. While it is difficult enough to determine the PSF of a bare CCD, it is virtually impossible to measure both PSFs in a digital camera system without detailed optical and electronic analysis. Our simple models based on end-to-end input/output analysis captured the general low-pass nature of the components, but little else. We suspect also that the electronics of our camera do not admit to a linear model. Thus, our predictions for the intensity of the sampled-grating moire pattern were inaccurate. And although our model is a good one for describing linear components interleaved with discrete sampling, it can be difficult to characterize the linear components.

Geometrically, we found good agreement between our model and experiment. This was partly because we could easily mea-

sure the geometric parameters of the system and partly because we could ignore the second sampling stage because no more aliasing would be induced after the CCD. The assessment of the model was based on visually matching positions of the actual peaks in the DFT to the peaks predicted from Eq. (13). The primary moire pattern was caused by the δ with $(n_0, n_x, n_y) = (-1, 0, 1)$. With these n 's and the geometric parameters above, Eq. (13) predicts the fundamental frequency of the obvious moire pattern to be 0.175 cycles/y-pixel at 97.65° . The corresponding peak in the actual DFT was 0.181 cycles/y-pixel at 96.79° . This error in frequency magnitude is about 1.2% of the maximum possible frequency of 0.5 cycles/y-pixel. Our predictions for the positions were worse for those frequency vectors with higher values of n_0 and n_y because these values tend to amplify slight errors in our measurements of the geometric parameters of the system, as can be seen from Eq. (13). Practically, this shows that moire patterns taken from higher order harmonics of the object grating (high value of n_0) or higher spectral orders (high values of n_x or n_y) should be suspect because they magnify errors in the measured geometric parameters of the system.

Interestingly, all of the measurable frequency peaks had $n_x = 0$. Since T_{x1} appears only as n_x/T_{x1} in Eq. (13) for the position of the δ 's, the CCD sampling period in the x-direction had no effect on the moire patterns. This is because the grating was oriented so that its lines were almost parallel to the rows of the CCD. If the grating were rotated through another 90° , we would see the dependence on T_{y1} lessen until all the $n_y = 0$. We suspect that Idesawa et al.² had a similar situation since they actually consider the camera to be modeled by one grating corresponding to the horizontal scanning lines of the camera.

5. CROSSED-GRATINGS MODELS OF SAMPLED GRATINGS

Equation (7), which describes crossed-grating moire, and Eq. (12), which describes sampled-grating moire, are surprisingly similar, given the fundamentally different processes that they model. They are both regularly spaced distributions of δ -functions. It has been asserted that the moire patterns that result from digitally sampling gratings and those from crossed gratings can be considered equivalent.² The intuitive feeling behind this assertion is that the rows and columns of the sampling grid can be considered as gratings that when combined with an object grating, make moire patterns like crossed gratings. This is important because it means that the considerable amount of past work on metrology using crossed-grating moire can be directly applied to the relatively new idea of "scanning moire"—moire with one grating and a digital camera.

To assess the degree of similarity between crossed gratings and sampled gratings, we try to model a sampled grating with crossed gratings. We imagine a single grating $g_0(x,y)$ (the object grating) imaged by a digital camera. The grating has period T_0 , angle θ_0 , and Fourier series coefficients c_{n_0} . If aliasing occurs, moire patterns will result in the reconstructed image. We would like to know what set of crossed gratings of the form of Eq. (4) would give the same moire pattern for the same object grating. These "effective" gratings will be chosen such that they give approximately the same Fourier transform as the Fourier transform of the sampled, reconstructed image of grating $g_0(x,y)$. The effective gratings will serve as a sort of camera model. It is important to realize that the camera model is made of *gratings* (patterns of some transmission profile which is periodic along some direction), not arbitrary transmission functions. If we al-

lowed arbitrary functions, then the camera model would be a simple inverse Fourier transform. We are restricting ourselves to gratings in order to test the assertion that sampled-grating moire patterns can be considered the same as crossed gratings. In addition, we will only attempt to model 1/2 passes through our multiple stage model, meaning that the input grating is first blurred, sampled, and then reconstructed. This is necessary because crossed gratings result in a continuous image, not a sampled image which results from one pass through the loop. The Fourier transform of a reconstructed, sampled image of grating $g_0(x,y)$ is given by

$$\frac{1}{T_{x1}T_{y1}}P_2(f_x, f_y) \left[\sum_{n_{y1}=-\infty}^{\infty} \sum_{n_{x1}=-\infty}^{\infty} \sum_{n_0=-\infty}^{\infty} c_{n_0}P_1\left(\frac{n_0}{T_0}\cos\theta_0, \frac{n_0}{T_0}\sin\theta_0\right) \times \delta\left(f_x - \frac{n_0}{T_0}\cos\theta_0 - \frac{n_{x1}}{T_{x1}}, f_y - \frac{n_0}{T_0}\sin\theta_0 - \frac{n_{y1}}{T_{y1}}\right) \right]. \quad (14)$$

Here P_1 is the Fourier transform of the PSF of the camera (including lens and pixel shape) and P_2 represents the interpolation function.

The set of effective gratings is not unique. In fact, even the number of effective gratings is not determined. (For instance, any grating can be composed of two or more parallel gratings, each of which makes up part of the first grating's profile.) It is reasonable, however, to consider a set of two effective gratings, representing the rows and columns of the sampling grid, respectively. We first derive an approximate pair of effective gratings by assuming that they are each Ronchi rulings. We then relax the conditions on the two effective gratings by allowing arbitrary profiles.

Our strategy in both cases is to compare an equation that characterizes the assumed form of the effective gratings to Eq. (14), which gives the moire patterns from a digital sampling system. We then instantiate the free parameters of the effective gratings and evaluate how well the two equations match.

5.1. Approximation with two Ronchi rulings

Equation (7) gives the expression for k superimposed transmission gratings. We will let $k = 3$ for two effective gratings and one arbitrary, object grating. The c_n for a Ronchi ruling are $\text{sinc}(n/2)/2$. The expression for the superposition of the arbitrary grating $G_0(f_x, f_y)$ and the two Ronchi rulings representing the camera is then

$$\frac{1}{4} \sum_{n_0=-\infty}^{\infty} \sum_{n_1=-\infty}^{\infty} \sum_{n_2=-\infty}^{\infty} c_{n_0}\text{sinc}\left(\frac{n_1}{2}\right)\text{sinc}\left(\frac{n_2}{2}\right) \times \delta\left(f_x - \frac{n_0}{T_0}\cos\theta_0 - \frac{n_1}{T_1}\cos\theta_1 - \frac{n_2}{T_2}\cos\theta_2, f_y - \frac{n_0}{T_0}\sin\theta_0 - \frac{n_1}{T_1}\sin\theta_1 - \frac{n_2}{T_2}\sin\theta_2\right), \quad (15)$$

where T_0 , θ_0 , and c_{n_0} characterize the arbitrary grating. The free parameters are the periods and angles of the Ronchi rulings, i.e., T_1 , T_2 , θ_1 , and θ_2 . We see from Eq. (14), which gives the expression for the sampled grating G_0 , that a reasonable choice is to have $T_1 = T_x$, $T_2 = T_y$, $\theta_1 = 0$, and $\theta_2 = \pi/2$. The periods and angles of the effective gratings then match the periods and angles of the rows and columns of the sampling grid.

The *positions* of the δ -functions in both the sampled-grating case and the crossed-gratings case are thus equivalent. The amplitudes, however, are not. The most significant difference is that the sampled grating is bandlimited by the Nyquist frequencies, while the crossed gratings have no such limits. Likewise, the crossed gratings are not subject to the interpolation function or camera blur. Thus, two Ronchi rulings can accurately generate a superset of the frequencies from an aliased grating, but they will not give a good indication of the amplitudes.

We can demonstrate the Ronchi ruling approximation in pictures. Figures 3 and 14 show the two constituent, orthogonal Ronchi rulings that serve to simulate the sampling grid of a camera. Figures 4 and 15 show the respective Fourier transforms. When these two gratings are crossed, they look like Fig. 16, and the Fourier transform looks like Fig. 17. The Fourier transform is a grid of δ 's whose spacing is the reciprocal of the periods of the two Ronchi rulings. If the grating periods are equal to the sampling periods of the camera, as we have shown they should be, then the δ 's of the two orthogonal Ronchi rulings will coincide with the centers of the spectral orders of a sampled grating. In Fig. 18 we show the two orthogonal Ronchi rulings crossed with a third Ronchi ruling at 15° . This is the same 15° Ronchi ruling shown with its Fourier transform in Figs. 5 and 6. Finally, Fig. 19 shows the Fourier transform of the three crossed Ronchi rulings. Although it is difficult to tell from the figure, the effect of convolving the Fourier transform of the tilted ruling with that of the two orthogonal rulings has been to repeat the Fourier transform of the tilted ruling on a grid of δ 's, just like the repeated spectral orders of a sampled grating.

5.2. Approximation with two arbitrary gratings

When we relax the conditions on the profiles of the two effective gratings, we still find that the geometric parameters should match those of the sampling grid. Thus, the equation for the crossed-gratings moire pattern becomes

$$\sum_{n_0=-\infty}^{\infty} \sum_{n_1=-\infty}^{\infty} \sum_{n_2=-\infty}^{\infty} c_{n_0} c_{n_1} c_{n_2} \delta\left(f_x - \frac{n_0}{T_0} \cos\theta_0 - \frac{n_1}{T_1}, f_y - \frac{n_0}{T_0} \sin\theta_0 - \frac{n_2}{T_2}\right). \quad (16)$$

Here again the positions of the δ -functions in both the crossed-gratings case and the aliased-grating case are equivalent. The remaining free parameters are the Fourier series coefficients, c_{n_1} and c_{n_2} , of the two effective gratings. We would like these coefficients to account for the PSF of the camera and the interpolation function of the reconstruction. We see, however, that this is in general impossible because these coefficients are not sensitive to the necessary variables. From Eq. (14), the Fourier transform of the PSF appears as $P_1(n_0 \cos(\theta_0)/T_0, n_0 \sin(\theta_0)/T_0)$. P_1 is a function of n_0 , while c_{n_1} and c_{n_2} have no such dependence. Since the n 's vary independently of each other in the three sums, $c_{n_1} c_{n_2}$ cannot, in general, be equal to $P_1(n_0 \cos(\theta_0)/T_0, n_0 \sin(\theta_0)/T_0)$. A similar argument applies to the interpolation function $P_2(f_x, f_y)$, which must be zero outside the Nyquist frequencies. For a given n_1 and n_2 , the $c_{n_1} c_{n_2}$ product will multiply δ -functions at frequencies both inside and outside the Nyquist bounds because n_0 is still free to vary and take the δ -functions inside and outside the Nyquist bounds. Thus, by relaxing the conditions on the profiles, we cannot generally get better results than from the two Ronchi rulings.

We can reach a more concrete conclusion if we disregard the PSF and the interpolation function for the sampled grating case

(by setting them to one). By comparing Eqs. (14) and (16) we then have, for the effective gratings, $c_{n_1} c_{n_2} = 1/(T_x T_y)$. We arbitrarily set $c_{n_1} = 1/T_x$ and $c_{n_2} = 1/T_y$. The Fourier series for $g_1(x, y)$, the effective grating with the vertical lines, becomes

$$\begin{aligned} g_1(x, y) &= \frac{1}{T_x} \sum_{n_1=-\infty}^{\infty} \exp\left(\frac{jn_2 \pi x}{T_0}\right) \\ &= \sum_{n_1=-\infty}^{\infty} \delta(x - n_1 T_x). \end{aligned} \quad (17)$$

Similarly,

$$g_2(x, y) = \sum_{n_2=-\infty}^{\infty} \delta(y - n_2 T_y). \quad (18)$$

These are each a series of evenly spaced lines on otherwise opaque transparencies. Their product is a grid of pinholes, coincident with the centers of the pixels of the camera's image plane. These effective gratings model a camera with no blur and no Nyquist limitations.

The practical conclusion from these two-grating approximations is that the PSF and interpolation function cannot be accurately modeled using crossed gratings while the constituent frequencies can be. Thus, the body of research developed for crossed-grating moire can be applied to sampled-grating moire if the profiles of the moire patterns are not important. It should be noted that the crossed-gratings model predicts frequencies beyond the Nyquist bounds, which will never appear in sampled-grating moire.

If we want to simulate more passes through the sampling and reconstruction loop, we can add two more gratings for each pass. The periods of the gratings should be equal to the sampling periods. Their angles should match the angles of the sampling grid. In this way, we can generate a superset of the constituent sinusoids of the sampled moire pattern, although the amplitudes will still not match.

5.3. Visual comparison

We present here a brief experiment to show how a sampled-grating moire pattern can be approximated with two crossed gratings. The sampled-grating moire pattern to be approximated is the one used for the experiment described in Sec. 4 and shown in Fig. 20. The figure was halftoned and reduced for publication, so the actual amplitudes are not shown exactly; we are concerned here with a qualitative comparison.

According to the development in Sec. 5.1, the two gratings that simulate the camera should have periods equal to the CCD pixel spacing given in Sec. 4. Thus, (using arbitrary units) we set $T_1 = T_{x1} = 1.70$ and $T_2 = T_{y1} = 1.00$. We also have $\theta_1 = 0$ and $\theta_2 = \pi/2$. Matching the actual object grating, we set $T_0 = 1.21$, $\theta_0 = 88.38^\circ$, and the duty ratio $h = 0.48$.

Figure 21 shows the superposition of three gratings, one of which has the geometric parameters of the imaged square wave and two of which have the geometric parameters of the CCD as described above. We arbitrarily used Ronchi rulings for the two orthogonal gratings that simulate the camera.

According to our analysis in Sec. 5, we can generate a superset of the constituent frequencies of an aliased moire pattern by matching the geometry of the sampling system with crossed gratings. A visual comparison of the crossed gratings in Fig. 21 with the aliased moire pattern in Fig. 20 shows this to be qual-

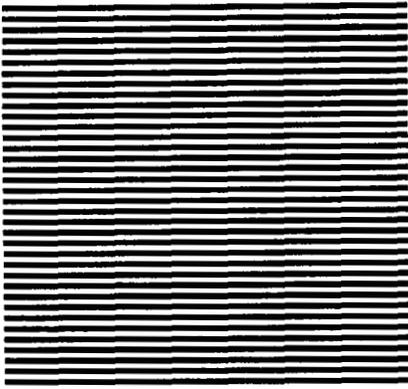


Fig. 14. Ronchi ruling at 90°.

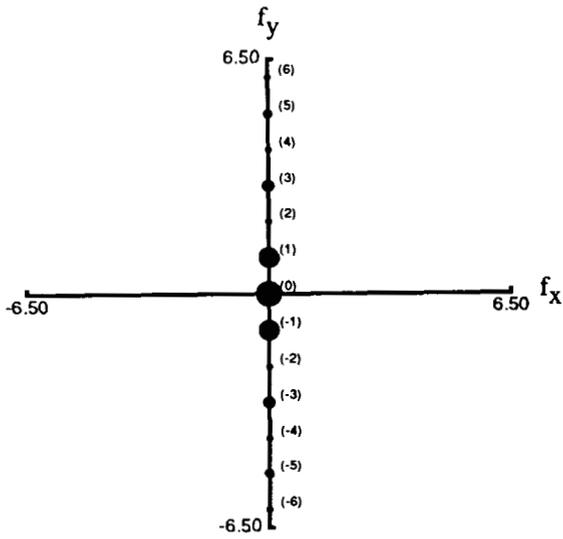


Fig. 15. Fourier transform of Fig. 14.

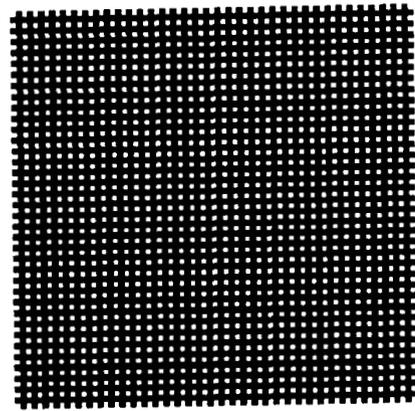


Fig. 16. Ronchi rulings at 0° and 90°.

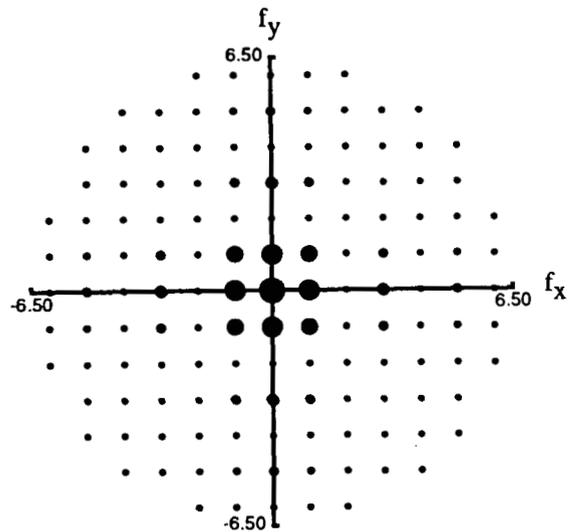


Fig. 17. Fourier transform of Fig. 16.

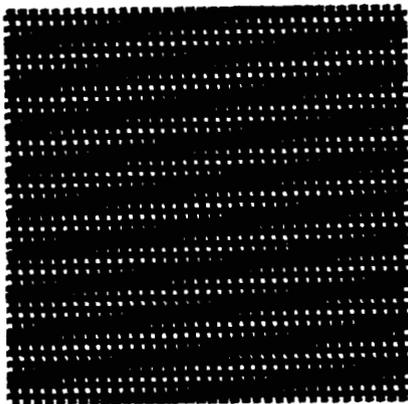


Fig. 18. Ronchi rulings at 0°, 90°, and 15°.

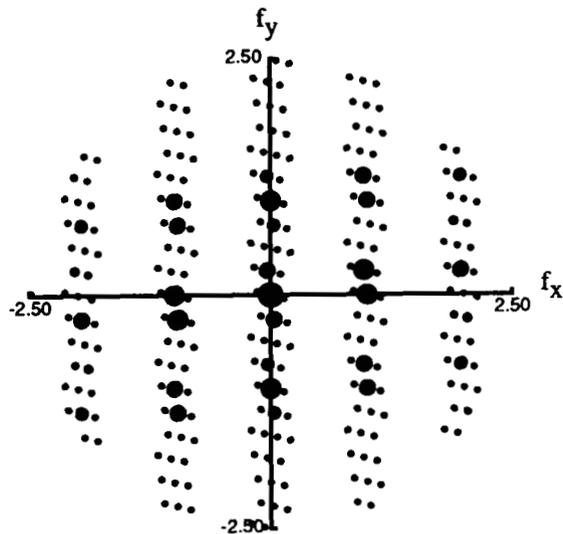


Fig. 19. Fourier transform of Fig. 18.



Fig. 20. Sampled-grating moire pattern.

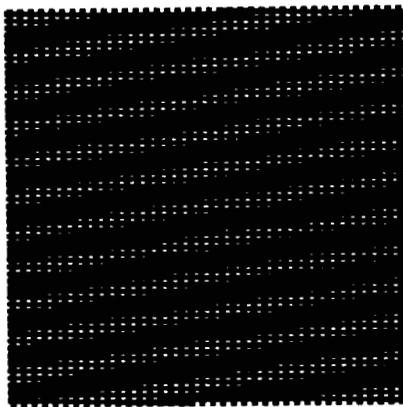


Fig. 21. Sampled-grating moire pattern simulated with three crossed gratings.

itatively true. The periods and angles of both moire patterns are about equal. In Fig. 22 we show the same crossed-gratings model without the vertical-line-grating that models the columns of the camera. Essentially the same moire pattern is visible in both. This is to be expected. As we have demonstrated, for our particular set of parameters, the sampling rate in the x direction had virtually no effect on the moire pattern because the lines of the object grating were almost parallel to the horizontal scanning lines.

6. CONCLUSION

Past attempts at characterizing sampled-grating moire have used either a crossed-gratings model or a single stage sampling model. We were able to evaluate these models by developing notationally consistent, frequency space models of crossed-grating moire and multiple stage, sampled-grating moire, of which single stage sampling is a special case. In evaluating the crossed-gratings model, we found that if we chose the periods and angles of the crossed gratings to match those of the sampling grid, the resulting crossed-grating moire pattern would contain a superset of the frequencies of the sampled grating. We demonstrated this fact with equations and a simple, qualitative experiment. We also showed that the crossed-gratings model cannot correctly predict the amplitudes of sampled-grating moire.

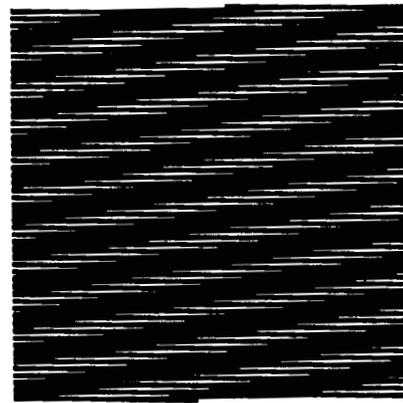


Fig. 22. Sampled-grating moire pattern simulated with two crossed gratings.

The single stage, sampled-grating models proposed in the literature are better than the crossed-gratings model for characterizing moire patterns due to aliasing. But we have shown how our model of multiple stage sampling accounts for multiple stages of aliasing, which has a significant impact on both the geometry and brightness of sampled-grating moire patterns. The geometry can become complicated if later sampling stages have lower sampling rates than earlier stages. Thus, in practically applying sampled-grating techniques, it is important to be aware of the sampling rates of all stages of the imaging system. Otherwise, there could be unanticipated aliasing effects that would likely lead to mistakes in shape measurement. In terms of brightness, we have shown that a single PSF cannot, in general, characterize a multiple-stage sampling system. We demonstrated with an experiment how our model does a good job of predicting the frequencies present in an actual sampled grating.

This research demonstrates clearly the justification for using a crossed-gratings model to characterize sampled gratings, but also shows its limitations. We have demonstrated the sorts of errors that can result from a crossed-gratings model or single stage, sampled-grating model and how these errors can be avoided using our multiple stage, sampled-grating model. The multiple stage sampling model should be used to avoid errors whenever sampling considerations become important in a digital imaging system.

7. ACKNOWLEDGMENTS

Thank you to Carol Novak of CMU for supplying camera data for our experiments. Thank you to Karel Vander Lugt of Augustana College for early inspiration on the Fourier transform.

This research was supported by the Defense Advanced Research Projects Agency, DoD, through ARPA Order No. 4976, monitored by the Air Force Avionics Laboratory under Contract F33615-87-C-1499, and by the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the National Aeronautics and Space Administration under Contract 957989. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the U.S. Government or the Jet Propulsion Laboratory.

8. REFERENCES

1. D. Post, "Sharpening and multiplication of moire fringes," *Exper. Mechan.* 7(4), 154-159 (1967).
2. M. Idesawa, T. Yatagai, and T. Soma, "Scanning moire method and automatic measurement of 3-D shapes," *Appl. Opt.* 16(8), 2152-2162 (1977).

3. Y. Morimoto, Y. Seguchi, and T. Higashi, "Application of moire analysis of strain using Fourier transform," *Opt. Eng.* 27(8), 650-656 (1988).
4. O. Bryngdahl, "Moire: formation and interpretation," *JOSA* 64(10), 1287-1294 (1974).
5. K. Paturski, S. Yokozeki, and T. Suzuki, "Moire profile prediction by using Fourier series formalism," *Jpn. J. Appl. Phys.* 15(3), 443-456 (1976).
6. B. W. Bell and C. L. Koliopoulos, "Moire topography, sampling theory, and charge-coupled devices," *Opt. Lett.* 9(5), 171-173 (1984).
7. M. Cetica, F. Francini, and D. Bertani, "Moire with one grating and a photodiode array," *Appl. Opt.* 24(11), 1565-1566 (1985).
8. D. R. Williams, "Aliasing in human foveal vision," *Vision Res.* 25(2), 195-205 (1985).
9. A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, Second Edition, Vol. 1, pp. 83-87, Academic, New York (1982).
10. J. Krumm and S. A. Shafer, "A sampled grating model of moire patterns from digital imaging," *Tech. Rep. CMU-RI-TR-89-19*, Carnegie Mellon Univ. (1989).



Steven A. Shafer received his Ph.D. degree in computer science from Carnegie Mellon University in 1983 and is now an associate professor there. He works in physics-based vision and the description of appearance: what it is that makes things look the way they do. This work includes the analysis of color, highlights, shadows, material types, and imaging systems. To study these, Dr. Shafer established the Calibrated Imaging Laboratory in 1985; it is a facility for controlled experiments in machine vision. Dr. Shafer has also studied 3-D shape inference from images and worked on the software architecture for the NAVLAB mobile robot van. In addition to his research, he is associate director of the pioneering Ph.D. program in robotics at Carnegie Mellon.



John Krumm grew up in Sioux Falls, South Dakota, graduating from Augustana College there in 1983 with majors in mathematics, physics, and computer science. Although South Dakota is a farming state, he never milked a cow, but did get to ride on a tractor occasionally. He did his master's work at the University of Minnesota, working on object recognition for robot assembly. He is now in the Robotics Ph.D. program at Carnegie Mellon University, working on computer vision in the Calibrated Imaging Laboratory. His primary interest is local spatial frequency analysis of images for computer vision, exploiting a common, simultaneous representation for spatial and frequency phenomena such as shape, texture, aliasing, and optical effects. His other research interests include moire patterns, shape from focus, and stereo vision.