EFFECTIVE ANTIALIASING OF COMPUTER GENERATED IMAGES

Dr. Geoffrey Y. Gardner, Staff Scientist
Edwin P. Berlin Jr., Computer Design Engineer
Grumman Aerospace Corporation
Bethpage, New York 11714

ABSTRACT

Aliasing in computer generated images produces artifacts which degrade training effectiveness. A rigorous implementation of low-pass filtering required to prevent aliasing requires computation too extensive to be incorporated into real-time computer image generation (CIG) systems. As a result, current CIG systems employ poor approximations to proper filtering, and aliasing still occurs.

This paper discusses the theory of image filtering and demonstrate a new real-time antialiazing technique developed from the theory. The new technique represents a much closer approximation to the rigorous solution and therefore produces images of much higher quality than current real-time techniques. At the same time it requires less computation.

INTRODUCTION

Computer Image Generation systems represent a continuously varying image by limited sampling in both time and space. The dynamic image is represented as a sequence of static images or frames sampled in time at a rate between 30 and 60 frames per second. Each frame is represented by a finite number of spatially sampled picture elements or pixels generally arranged in a raster format with equal spacing. The Nyquist Sampling Theorem states that for any sampled signal any frequency greater than one half the sampling frequency will not be properly represented. The result is called aliasing because the undersampled frequency components appear as different lower frequencies. Aliasing is a critical problem in CIG displays for training because it produces artifacts which are distracting and confusing to the trainee. Temporal aliasing occurs when the dynamic image changes significantly between sampled frames. Spatial aliasing occurs when any frame contains spatial frequencies greater than one half the spatial sampling frequency. In general, temporal aliasing occurs only during extreme vehicle maneuvers such as high speed rotation. Spatial aliasing, however, occurs in every frame containing high frequency information, such as that inherent in well-defined boundaries between surfaces. In addition, spatial aliasing anomalies are aggravated by dynamic presentation of the image sequence because each successive image is sampled slightly differently. As a result crawling, jumping and scintillation effects are produced at surface boundaries. Because the eye is so sensitive to dynamic changes, the temporal effects of spatial aliasing are extremely distracting and therefore degrade the training effectiveness of the CIG system. At the present time, therefore, spatial aliasing is by far the more critical problem to be solved to improve image quality. Consequently, we have concentrated our study on developing an effective technique for spatial antialiasing.

In seeking a solution to the spatial antialiasing problem, we must consider the frequency content of the images we will be generating and be careful to sample at a frequency at least twice that of the highest frequency in a given image. However, since sharp surface boundaries and corners, which contain infinite frequency components, are common, any finite sampling of the unfiltered image will produce aliasing. Constrained by a system with finite sampling, the ideal solution is to prefilter the image before sampling to remove all frequencies higher than half the sampling frequency. This can be done in two ways. First, the Fourier transform of the image can be computed and multiplied by the frequency domain transfer function of the appropriate low-pass filter. The inverse transform of the product will produce the filtered image. Alternatively, we could perform the equivalent operation in the spatial domain by convolving the image with the impulse response of the low-pass filter. A rigorous implementation of this convolution requires the integration of weighted intensity over the entire image for each pixel. Either of these techniques requires extremely involved mathematics to process complex surface boundaries. In addition, the immense computation load required is prohibitive.

Because of these problems, current CIG systems have retreated from proper image prefiltering to a variety of compromises which fail to reduce aliasing to an acceptable level. At the same time these compromises have added an undesirable amount of computation to achieve even a limited amount of antialiasing. In general, these compromise solutions attempt to adjust the intensity of any pixel which contains portions of more than one surface. The intensity of the pixel is represented by an average of the intensities of all surfaces within the pixel, each weighted by the fraction of the pixel area that it covers. Although seemingly straightforward, this approach is very costly in computation since the determination of subpixel areas is time-consuming even for straight edge boundaries. To simplify the computation some systems use an approximation, subdividing the pixel into an nxn array of subpixels for which the quantized image is generated and averaged for the total pixel. The problems of this technique are readily apparent. First, the computation load for an antialiased pixel is increased by a factor of n^2 (with n as high as 8). Second, a quantization problem still exists with the potential for producing distracting jumps in intensity in a dynamic image. Finally, even an exact, continuous area weighting does not satisfy the sampling

theorem filtering requirement, so that aliasing can not be properly suppressed. In brief, current techniques provide too little antialiasing at too great a cost.

EFFECTIVE ANTIALIASING

The Ideal Solution

The theoretically correct solution to the aliasing problem is to filter the image before sampling. The frequency domain transfer function of the ideal presampling filter, $H(\omega_x,\omega_y)$ is shown in Fig. 1 (the sample spacing is assumed to be 1 in the x and y directions). The impulse response of this filter is:

$$h(x,y) = \frac{\sin(\pi x)}{\pi x} \cdot \frac{\sin(\pi y)}{\pi x} \quad -\infty \leq x,y \leq \infty \qquad (1)$$

and the filtered image $F_F(x,y)$ is this convolved with the original image $F(x,y)$:

$$F_F(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} h(x-R,y-S)F(R,S)dRdS \qquad (2)$$

Direct evaluation of this integral for each sample is computationally very costly. It is desirable to find some approximations to this integral which give suitable results without a heavy computational burden.

Approximating the Ideal Solution

Our goal is to develop a satisfactory approximation to the convolution in Eq (2) while minimizing the computation required. There are two important problems that we must solve. First, we must find a way to perform the convolution before sampling. Second, we must drastically reduce the number and complexity of the calculations.

For an image represented by a set of surfaces and their boundaries high spatial frequencies will occur only at abrupt changes in intensity. Such changes occur at the surface boundaries; regions within the boundaries will have low spatial frequencies and, therefore, will require no antialiasing. This implies that antialiasing can be limited to image regions near boundaries. Thus we can significantly reduce the computation of Eq (2) by restricting it to only a fraction of the image points.

We must further limit the computation by placing finite limits on the integral in Eq (2). We will do this by approximating the impulse response, Eq (1), by a weighting function with a finite domain. This domain, which we will call a sampling region, must cover a distance of at least two sampling periods in x and y for acceptable filtering. Restricting the calculation of the convolution not only reduces the computation load dramatically, but also allows us to make two more simplifications within each sampling region. First, within each sampling region, we will approximate all boundaries as straight lines. Second, we will assume surface intensity to be constant within the sampling region. With these approximations Eq (2) reduces to

$$F_F(x,y) = \sum_{i=1}^{N} I_i \iint_{R_i} h'(x,y)dxdy \qquad (3)$$

for N surfaces visible within the sampling region R, with the i^{th} surface visible over a region R_i so that $\sum_{i=1}^{N} R_i = R$. The function $h'(x,y)$ is the weighting function approximation to $h(x,y)$.

For two surfaces separated by a single boundary, Eq (3) becomes

$$F_F(x,y) = I_1 \iint_{R_1} h'(x,y)dxdy + I_2 \iint_{R-R_1} h'(x,y)dxdy \qquad (4)$$

We have now reduced the convolution to a form in which the integration is dependent only on the region of integration. In Eq (4) the regions of integration are determined by the location of the boundary. We have developed a technique to evaluate the integral of our weighting function explicitly as a function of a single continuous measure, ℓ, of boundary position within the sampling region. This enables us to express the contribution of surface 1 to the filtered intensity of a sampling region as

$$F_{F_1}(x,y) = I_1 W(\ell) \qquad (5)$$

where $W(\ell)$ is a new weighting function applied to the surface intensity and defined by

$$W(\ell) = \iint_{R_\ell} h'(x,y)dxdy \qquad (6)$$

$W(\ell)$ is thus the volume defined under $h'(x,y)$ over R_ℓ and can be normalized to a maximum of one. $W(\ell)$ will equal zero when the surface lies outside the sampling region, will equal one when the surface covers the region and will vary continuously and nonlinearly as the boundary moves across the region. Since the total volume under $h'(x,y)$ over R will be normalized to one, Eq (4) can be written as

$$F_F(x,y) = I_1 W(\ell) + I_2(1-W(\ell)) \qquad (7)$$

Equation (7) represents a manageable approximation of the ideal filtering defined by Eq (2). We have achieved a tremendous reduction in computation by restricting filtering to image points near boundaries, approximating the impulse response by a bounded weighting function, linearizing the boundary and assuming constant intensity within a sampling region, and devising a means of representing surface intensity weighting as a function of boundary position in the sampling region. In addition to its computational efficiency, this technique has the important feature of approximating the ideal filtering on the unsampled image. Because the filtering operation depends on a single continuous measure

of boundary position, it will not introduce any quantization effect that could produce distracting jumps in intensity in a dynamic image.

Demonstration of Results

We demonstrated the effectiveness of our antialiasing technique on extreme cases where aliasing is severe -- thin bright surfaces on a dark background. We defined surfaces with width equal to one pixel spacing using both linear and curved boundaries. The surface intensity was assigned the maximum value (1) and the background was assigned the minimum value (zero). Because each surface was so thin, in general two boundaries passed through a given sampling region. Accordingly, we modified Eq (7) to represent the contribution of the surface between the boundaries

$$F_F(x,y) = I_1(W(\ell_1) + W(\ell_2)-1) + I_2(2-W(\ell_1)-W(\ell_2))$$

$$= W(\ell_1) + W(\ell_2)-1$$

since $I_1 = 1$ and $I_2 = 0$.

To provide a benchmark, we used a current technique of area-weighted intensity as approximated by the n x n subpixel array. Since no improvement in image quality was perceivable for n greater than 8, we used this value as representative of exact area-weighting.

To demonstrate aliasing clearly we generated a 128-pixel-wide image of three thin surfaces with linear bounds at small angles to the scan lines. Figure 2 shows the unfiltered image with and without zoom-central enlargement by our scan converter. Antialiasing using the n x n subpixel array technique for values of n ranging from 2 to 8 is shown in Fig. 3a, b and c. Although considerable smoothing is apparent, some aliasing is still evident in the variation in intensity in the two lower lines. This residual aliasing holds the potential for distracting edge crawling in a dynamic display.

Figure 3d shows the same image antialiased by our technique. Since no variation in intensity is perceivable along any of the lines, and since filtered intensity at any pixel is a continuous function of boundary position, no dynamic discontinuities will be produced as the boundary moves in the image.

To demonstrate the behavior of our technique on surfaces with curved boundaries, we generated a test image of seven elliptical surfaces of 1 pixel width. Figure 4 shows the unfiltered image with and without enlargement. Figure 5a, b and c shows the image smoothed by the nxn subpixel technique for values of n ranging from 2 to 8. Evidence of aliasing is still present. Figure 5d shows the image after filtering with our technique. All perceivable aliasing has been removed.

CONCLUSION

Despite the assumptions and approximations we used to make our technique computationally efficient, it performs antialiasing in a manner far superior to the area-weighted intensity technique because our technique represents a much better approximation to ideal filtering. First, it uses nonuniform weighting of intensity, and second, it applies the weighting over a larger region. To demonstrate the importance of each of these factors, we generated a test image of 21 concentric circles each 1 pixel wide. Figure 6a shows the unfiltered image. When uniform weighting is applied over two sampling periods (Fig. 6b), the region is large enough but the weighting is improper. Figure 6c shows the result of applying nonuniform weighting over a region covering only one sampling period. Figure 6d shows the result of applying uniform weighting over one sampling period, equivalent to area-weighted intensity in which both weighting and region size are improper. Effective antialiasing using nonuniform weighting over two sampling periods is shown in Fig. 6e.

The success of our technique is due to the fact that it is based on proper filtering before sampling. Because it uses a continuous measure of boundary position it eliminates the potential for distracting jumps in intensity in a dynamic image. The technique can be extended to antialias surfaces with multiple boundaries, such as small objects, by combining the position information of each boundary. Since each position is measured continuously, the combined weighting of surface intensity will vary without discontinuity.

We have demonstrated our technique on images including boundaries composed of lines or second-order curves. However, curves of any order can be handled as long as the equation of the boundary is known analytically.

Finally, the effectiveness of our technique is achieved at a low computational cost. The intensity weighting function depends on a single measurement of boundary position and involves a simple computation or table lookup. In the test images generated our technique took less computation time than the n x n subpixel array technique with n = 2.

ACKNOWLEDGEMENT

The research described in this paper was sponsored by the Air Force Human Resources Laboratory (AFHRL) under contract F33615-79-C0029. The authors are grateful to Mr. William Foley and Capt. Michael Ingalls for the insight they provided into the problem of aliasing.

ABOUT THE AUTHORS

Dr. Geoffrey Y. Gardner is Staff Scientist and principal investigator of the Advanced Computer Image Generation project at Grumman Aerospace Corporation. He holds a BA from Harvard and a PhD (EE) from the Polytechnic Institute of New York.

Edwin P. Berlin, Jr., BSEE, Massachusetts Institute of Technology, designs high-throughput computers for Grumman Aerospace Corporation. He holds a patent on a three-dimensional display device and has another pending.

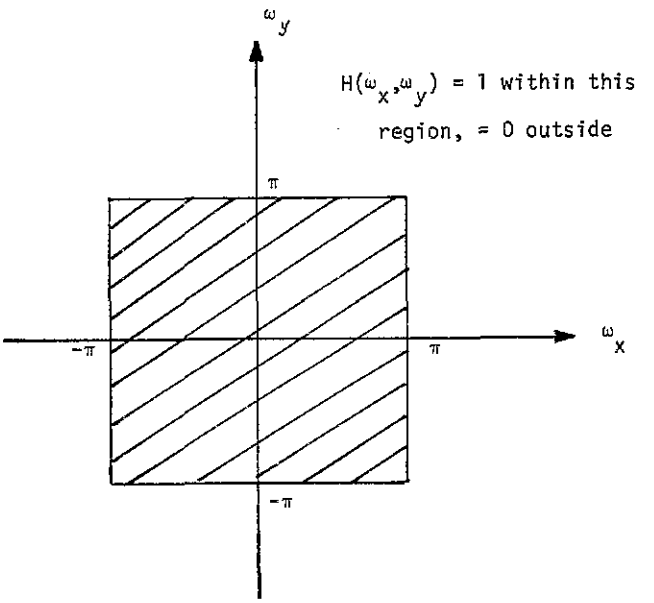

Fig. 1 Frequency Domain Transfer Function of Ideal Presampling Filter

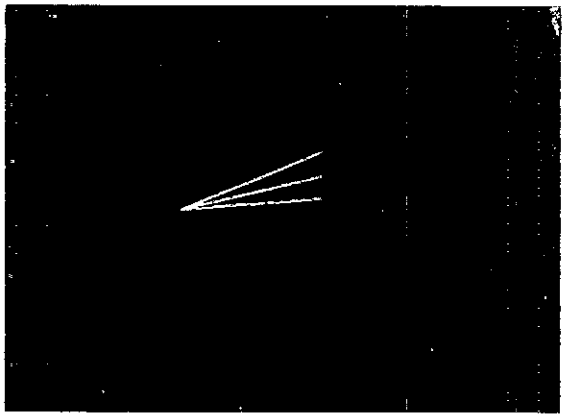

(a) Normal Size Image on 512 Pixel Wide Display

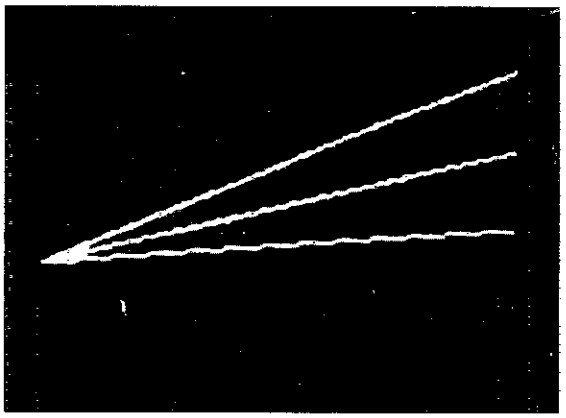

(b) Same Image Enlarged to Show Aliasing More Clearly

Fig. 2 Unfiltered Image of Thin Surfaces with Linear Boundaries

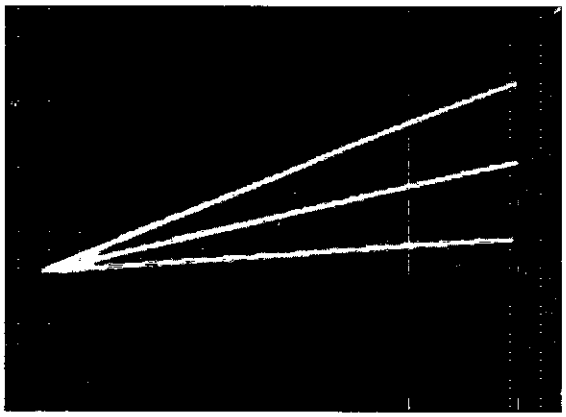

(a) 2 x 2 Subpixel Array

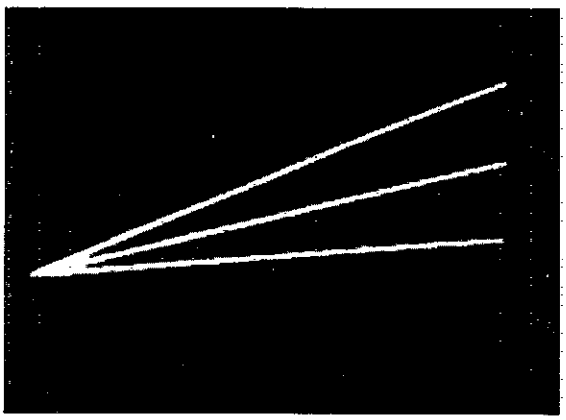

(b) 4 x 4 Subpixel Array

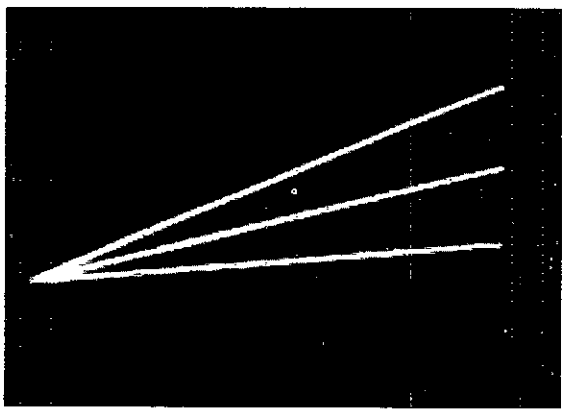

(c) 8 x 8 Subpixel Array

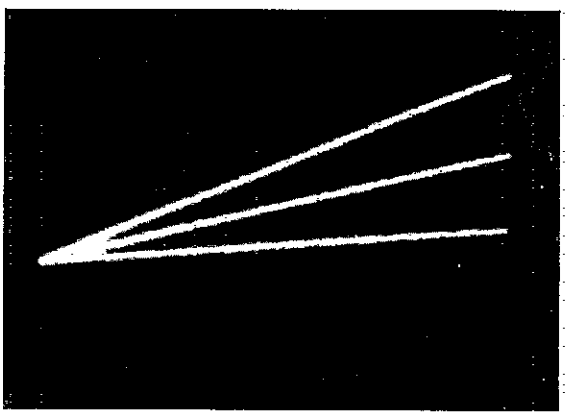

(d) Nonuniform Weighting Over Two Sampling Periods

Fig. 3 Antialiasing of Linear Boundary Image

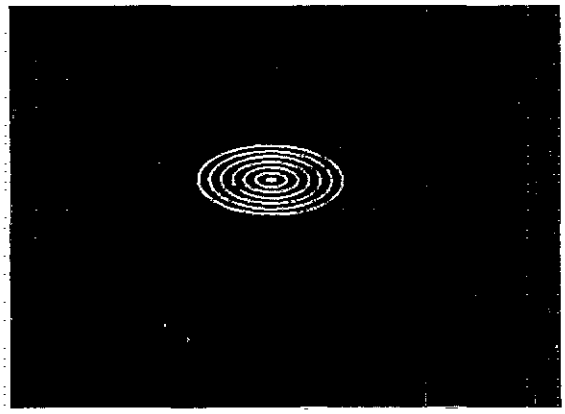

(a) Normal Size Image on a 512 Pixel Wide Display

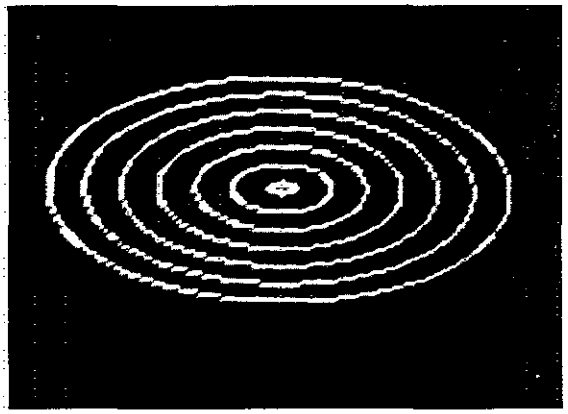

(b) Same Image Enlarged to Show Aliasing More Clearly

Fig. 4 Unfiltered Image of Thin Surfaces with Curved Boundaries

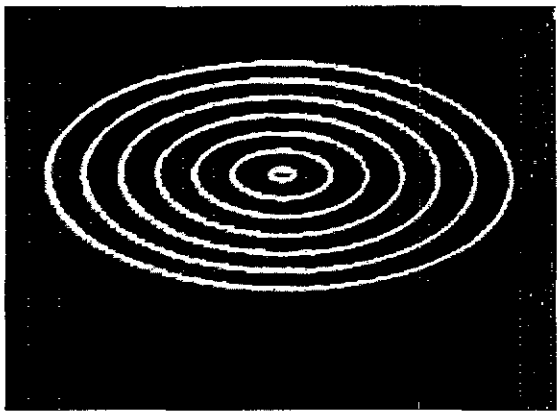

(a) 2 x 2 Subpixel Array

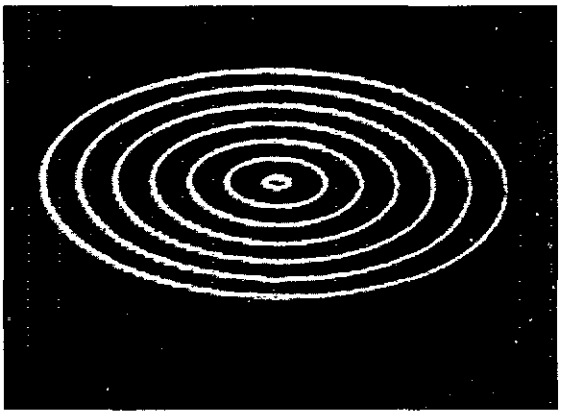

(b) 4 x 4 Subpixel Array

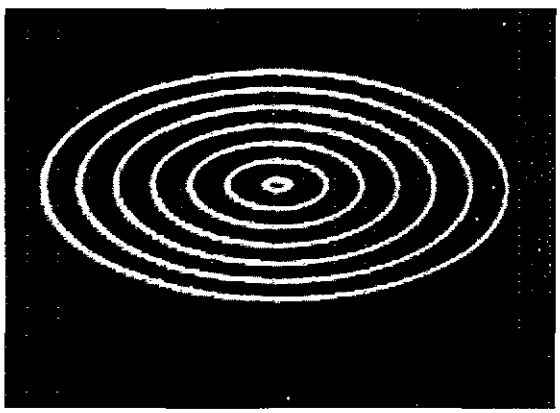

(c) 8 x 8 Subpixel Array

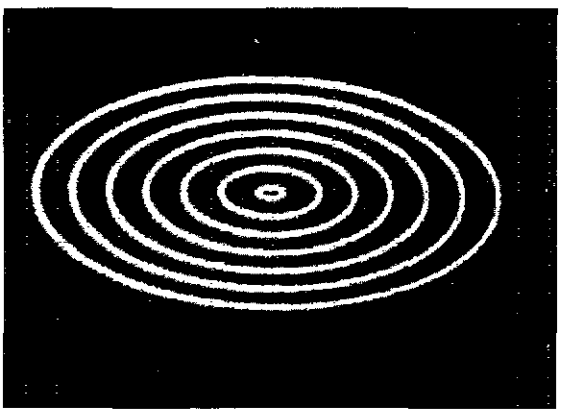

(d) Nonuniform Weighting Over Two Sampling Periods

Fig. 5 Antialiasing of Curved Boundary Image

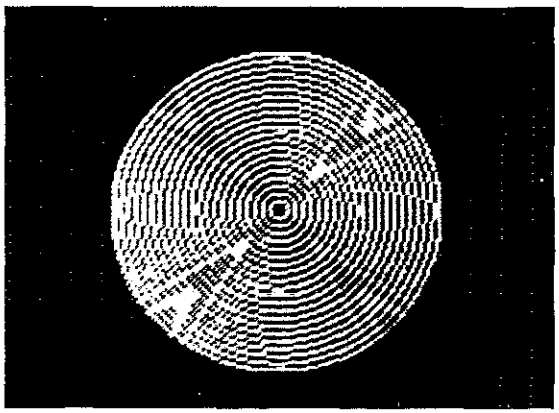

(a) Unfiltered Image

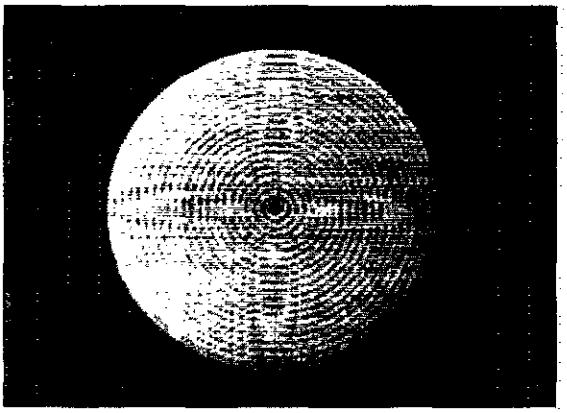

(b) Uniform Weighting Over Two Sampling Periods

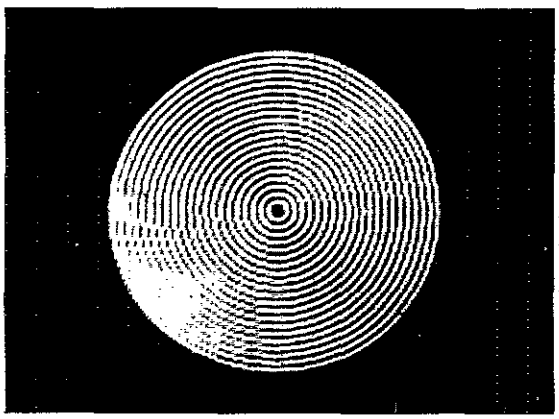

(c) Nonuniform Weighting Over One Sampling Period

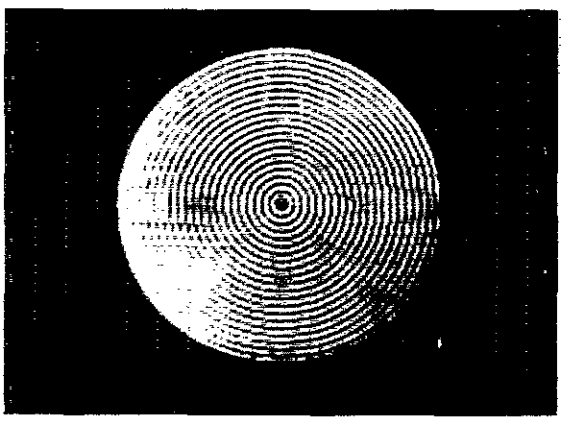

(d) Uniform Weighting Over One Sampling Period

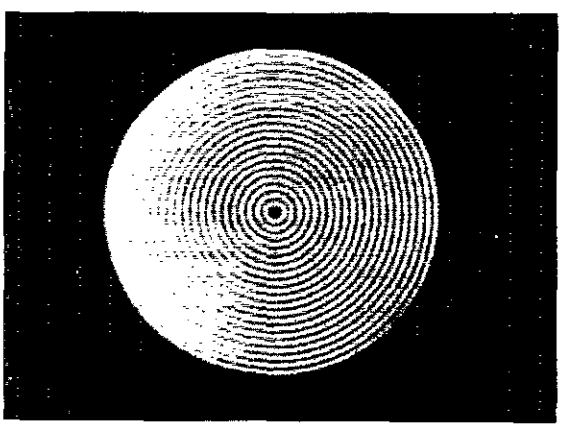

(e) Nonuniform Weighting Over Two Sampling Periods

Fig. 6 Effect of Weighting Distribution and Region on Antialiasing of 21 Concentric Circles