

Reconstruction Filters in Computer Graphics

Don P. Mitchell

Arun N. Netravali

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Problems of signal processing arise in image synthesis because of transformations between continuous and discrete representations of 2D images. Aliasing introduced by sampling has received much attention in graphics, but reconstruction of samples into a continuous representation can also cause aliasing as well as other defects in image quality. The problem of designing a filter for use on images is discussed, and a new family of piecewise cubic filters are investigated as a practical demonstration. Two interesting cubic filters are found, one having good antialiasing properties and the other having good image-quality properties. It is also shown that reconstruction using derivative as well as amplitude values can greatly reduce aliasing. This paper has been accepted for presentation in SIGGRAPH 88.

CR Categories and Subject Descriptions: I.3.3 [**Computer Graphics**]: Picture/Image Generation; I.4.1 [**Image Processing**]: Digitization

General Terms: Algorithms

Additional Keywords and Phrases: Antialiasing, Cubic Filters, Filters, Derivative Reconstruction, Reconstruction, Sampling

1. Introduction

The issues of signal processing arise in image synthesis because of transformation between continuous and discrete representations of images. A continuous signal is converted to a discrete one by *sampling*, and according to the sampling theorem [SHA49], all the information in the continuous signal is preserved in the samples if they are evenly spaced and the frequency of sampling is twice that of the highest frequency contained in the signal. A discrete signal can be converted to a continuous one by interpolating between samples, a process referred to in the signal-processing literature as *reconstruction*.

Many conversions between continuous and discrete representations may occur in the course of generating an image. For example when ray tracing a texture-mapped surface, a photograph may be sampled by a digitizer to define the texture, then the texture samples are interpolated and resampled when a ray strikes the textured surface, the ray samples are interpolated and resampled to generate pixel values, and the pixels are interpolated by a display and finally resampled by retinal cells when the image is viewed. Resampling may be more explicit, as in enlarging or reducing a digital image or warping an image (e.g., with Catmull and Smith's algorithm [CAT80]). Each of these conversions can introduce conspicuous errors into an image.

Errors introduced by sampling (e.g., aliasing) have received considerable attention in the graphics community since Crow identified this as the cause of certain unwanted artifacts in synthetic images [CRO77]. Aliasing in images was discussed in the classic 1934 paper by Mertz and Gray [MER34]. Their discussion contains a description of artifacts well-known to graphics researchers today and shows that the condition for preventing aliasing was known, as a rule of thumb, long before Shannon's proof of the sampling

theorem:

The interference usually manifests itself in the form of serrations on diagonal lines and occasional moiré effects in the received picture. Confusion in the signal may be practically eliminated by using an aperture of such a nature that it cuts off all [Fourier] components with n numbers greater than $N/2$ [half the scanning rate]

By comparison, the problems introduced by reconstruction have been somewhat neglected in the graphics literature. Reconstruction can be responsible for aliasing and other types of distortion that mar the subjective quality of an image. This paper will focus on the effects of reconstruction and how to design filters for graphics applications.

2. Aliasing Caused by Reconstruction

Aliasing in synthetic images is a serious problem and still not completely solved. In other digital-signal-processing applications, aliasing is eliminated by prefiltering signals before sampling, as illustrated in Figure 1. Note that it is the prefiltered signal that is reconstructed in this case.




Figure 1. Sampling and Reconstruction

While prefiltering is the classic solution to aliasing problems, there is a special problem encountered in computer graphics. Many synthetic images originate from what we will call *procedural signals*, in which the signal is only implicitly defined by an algorithm for computing point samples. Operations that require an explicit representation of the signal cannot be performed, and in particular, prefiltering is impractical. This difficulty is unique to computer graphics, and ray tracing is the clearest example of it [WHI80].

To explain the role that reconstruction plays in aliasing, it will be helpful to review briefly the theory of sampling and define the operations of sampling and reconstruction more precisely. In one dimension, a signal can be represented by a continuous function $f(x)$. Producing a discrete signal by sampling is equivalent to multiplying by an infinite train of impulses known as a comb function:

$$f_s(x) = f(x) \cdot comb(x) \tag{1}$$

where

$$comb(x) = \sum_{n=-\infty}^{\infty} \delta(x-n) \tag{1b}$$

Unit spacing between samples is assumed in equation (1b), and $\delta(x)$ is the Dirac delta function. In this case, the sampling theorem states that $f(x)$ can be reconstructed exactly from its samples if it contains no frequencies greater than 0.5 cycles per sample. This critical frequency is called the *Nyquist frequency*.

Reconstruction is accomplished by convolving (indicated by $*$) the discrete signal with a reconstruction filter kernel, $k(x)$:

$$f_r(x) = f_s(x) * k(x) \tag{2}$$

$$= \int_{-\infty}^{\infty} f_s(u) \cdot k(x-u) du \tag{2b}$$

$$= \sum_{n=-\infty}^{\infty} f_s(n) \cdot k(x-n) \tag{2c}$$

Except in the mathematically ideal case, some error is introduced in the process of sampling and reconstruction, and $f(x)$ will be somewhat different from $f_r(x)$. To analyze this error, it is useful to view the problem in the frequency domain. The Fourier transform of the signal is its spectrum $F(\nu)$, and the Fourier transform of the filter is its frequency response $K(\nu)$. Since multiplication in the spatial domain is equivalent to convolution in the frequency domain (and vice versa), sampling can be described by:

$$F_s(\nu) = F(\nu) * Comb(\nu) \tag{3}$$

and reconstruction by:

$$F_r(\nu) = F_s(\nu) \cdot K(\nu) \tag{4}$$

The Fourier transform of a comb function is also a comb function (with reciprocal spacing between impulses).




Figure 2. Sampling and Reconstruction in Frequency Domain

Figure 2 illustrates the consequences of the convolution in equation (3). The spectrum of a sampled signal $F_s(\nu)$ is the sum of an infinite sequence of shifted replicas of the original signal's spectrum, each centered at the location of an impulse in the comb. Equation (4) states that, in the frequency domain, reconstruction can be interpreted as the multiplication by $K(\nu)$ which is intended to eliminate all the *extraneous replicas* of the signal's spectrum and keep the original *base-band* centered at the origin. $K(\nu)$ is indicated by the dashed curve in Figure 2.

However, Figure 2 also demonstrates a problem. The replicas of the signal spectrum overlap, and the reconstruction filter can not isolate a pure version of the base-band signal. When part of the energy in a replica of the spectrum leaks into the reconstructed signal, aliasing results. If the bandwidth of the signal were narrower or the sampling rate higher, the copies would not overlap, and exact reconstruction would be possible.

Even if the replicated spectra do not overlap, aliasing can result from poor reconstruction, as illustrated in Figure 3. When aliasing is a consequence of undersampling (or lack of prefiltering), it is referred to as *pre-aliasing*, and when it results from poor reconstruction, it is called *postaliasing*.




Figure 3. Postaliasing Resulting from Poor Reconstruction Filter

Figure 4 shows an extreme example of aliasing in an image. In this figure, the two-dimensional signal, $f(x, y) = \sin(x^2 + y^2)$, was sampled on a 128×128 pixel grid. Then, these samples were reconstructed with a cubic filter (to be described later in the paper) and resampled to 512×512 pixels.

The rings on the left side of the image are part of the actual signal, but the rings on the right side are Moiré patterns due to prealiasing. In the center of the image is a fainter set of concentric rings resulting from postaliasing. Postaliasing occurred when the discrete image of 128×128 pixels was enlarged to 512×512 pixels by resampling. Note that this conspicuous postaliasing pattern results from "beating" between the signal and its alias. This can also be understood from Figure 3, where it can be seen that at the Nyquist frequency (indicated by ν_N) the signal's spectrum and its nearest replica come close together. Power in the spectrum very near the Nyquist frequency is thus the cause of the most difficult type of aliasing to remove from an image. This problem has been noted by other graphics researchers [COO87] and by Mertz et al. [MER34].

Using the same set of samples as in Figure 4, a much better reconstruction filter can be applied (a 30-unit-wide windowed sinc filter). Figure 5 demonstrates a dramatic reduction of the postaliasing pattern, but the prealiasing is unaffected. The spectrum of this reconstruction filter is very close to the ideal step shape shown in Figures 2 and 3.

3. Other Image Defects Caused by Reconstruction

Notice in Figure 2, that a reconstruction filter $K(\nu)$ has two tasks. First it must remove the extraneous replicas of the signal spectrum (to prevent aliasing). Second, it should pass the original signal base band, but the signal can be distorted if this is not done perfectly. This second type of reconstruction error will be referred to as *base-band attenuation*.

From the previous section, one might assume that the literature of signal processing provides a complete solution to the reconstruction problem in graphics; however, there is a serious difficulty with the ideal sinc filter that is not obvious from studying its frequency response. Figure 6 shows a simple figure reconstructed with the same filter used in Figure 5. The rippling pattern radiating from the edges is called *ringing*. Ringing is strongly suggested by the form of the impulse response of the sinc filter, as shown in Figure 7:




Figure 7. Impulse Response of Ideal Sinc Filter

Classical digital filter design places a heavy emphasis on the frequency response of a filter. That works well in the audio domain, but when considering the appearance of images, it is important to also pay attention to the shape of the impulse response.

The response of human viewers to various spatial effects of filters is not yet a well-understood science and is largely subjective in nature. Filters that have some aliasing problems or certain types of base-band attenuation may turn out to give visually-pleasing results. Schreiber and Troxel have discussed the spatial effects of reconstruction filters [SCH85], and they mention some of the important defects that can occur when judging the quality of an image subjectively: *sample-frequency ripple*, *anisotropic effects*, *ringing*, *blurring*, and *aliasing*. Each of these effects will be considered in detail in the following section.

Unfortunately, it is often necessary to trade off one type of distortion for another, and the design of a single filter perfect for all applications is almost certainly impossible. As Figure 6 illustrated, perfect antialiasing resulted in the serious defect of ringing. However, Brown realized that a moderate amount of ringing can improve the subjective quality of an image by enhancing the appearance of sharpness [BRO69]. He found that a single transient lobe of ringing was effective at sharpening, but multiple transients (as in Figure 6) always degrade image quality.

Many of the concepts presented so far have been illustrated in one dimension for simplicity. However, image reconstruction takes place in two dimensions and involves the convolution of a 2D lattice of samples with a filter $k(x, y)$. In this paper, we will consider only *separable* filters, where the samples are convolved with the product $k(x)k(y)$. Separable filters are computationally more efficient than nonseparable because the filtering operation can be performed in separate passes vertically and horizontally. If the filter kernel is N samples wide, the reconstruction can be performed with $O(N^2)$ multiplications for the general filter $k(x, y)$ but with $O(N)$ if the filter is separable.

4. Piecewise Cubic Reconstruction Filters

Rather than discuss the issues of filter design abstractly, this paper will apply them to the study of a family of filters defined by piecewise cubic polynomials. Cubic filters are sufficiently complex to have a broad range of behaviors, but they are simple enough to be computationally attractive. Hou and Andrews have studied the filtering properties of the cubic B-spline [HOU78], and two studies have been made of the one-parameter family of cardinal cubic splines [KEY81,PAR83].

The general form for a symmetric cubic filter is:

$$k(x) = \begin{cases} P|x|^3 + Q|x|^2 + R|x| + S & \text{if } |x| < 1 \\ T|x|^3 + U|x|^2 + V|x| + W & \text{if } 1 \leq |x| < 2 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Several obvious constraints can be placed on this function to reduce the number of free parameters. First, the filter should be smooth in the sense that its value and first derivative are continuous everywhere. Discontinuities in $k(x)$ will lead to high-frequency leakage in the frequency response of the filter which can allow aliasing. In addition, the problem of *sample-frequency ripple* can be designed out of the filter by requiring (for all x):

$$\sum_{n=-\infty}^{\infty} k(x-n) = 1 \quad (7)$$

This means that if all the samples are a constant value, the reconstruction will be a flat constant signal. Figure 8 demonstrates this defect by using an unnormalized Gaussian filter to reconstruct a 512 x 512 image from 64 x 64 samples. In the frequency domain, sample ripple can be viewed as an alias of the image's DC component. It can be shown that the condition given by equation (7) means that the frequency response of these cubic filters will be zero at all integer multiples of the sampling frequency except zero, eliminating all extraneous replicas of the DC component.

With these constraints, the number of free parameters are reduced from eight to two, resulting in the following family of cubic filters:

$$k(x) = \frac{1}{6} \begin{cases} (12-9B-6C)|x|^3 + & \text{if } |x| < 1 \\ (-18+12B+6C)|x|^2 + (6-2B) & \\ (-B-6C)|x|^3 + (6B+30C)|x|^2 + & \text{if } 1 \leq |x| < 2 \\ (-12B-48C)|x| + (8B+24C) & \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Some values of (B, C) correspond to well-known cubic splines. (1,0) is the cubic B-spline, (0, C) is the one-parameter family of cardinal cubics with (0, 0.5) being the Catmull-Rom spline, and (B, 0) are Duff's tensioned B-splines [DUF86].

In two or more dimensions, visible artifacts can be caused by angle-dependent behavior or *anisotropic effects*. Figure 9 illustrates this problem by reconstructing with the separable filter $k(x)k(y)$ using parameter values of (0,0). Even though sample-frequency ripple has been designed out of $k(x)$, in two dimensions the pixel structure is highly conspicuous because the impulse response and the sampling lattice are not radially symmetric.

The phenomenon of *ringing* has already been seen in Figure 6. Filters in the cubic filter family can also exhibit this problem as seen in Figure 10, where parameter values of (0,1) were used. Ringing results when $k(x)$ has negative side lobes, and although some ringing can enhance sharpness, a filter that becomes negative is problematic. In Figure 10, a typical problem is seen where portions of the image near an edge have become negative and have been clamped to zero. This results in pronounced black spots (e.g., at the top of the statue's head). Similar clamping occurs to white, but is less noticeable because of the eye's nonlinear response to contrast. Schreiber and Troxel have suggested that subjectively even sharpening can only be produced by introducing ringing transients in a suitably nonlinear fashion [SCH85]. These conspicuous clamping effects could also be eliminated by reducing the dynamic range of the image or raising the DC level of the image.

Parameter values of $(3/2, -1/4)$ result in an image that is unnecessarily *blurry*, as seen in Figure 11. The cubic B-spline also suffers from this problem. In viewing many reconstructions with filters in this family, ringing, anisotropy, and blurring are the dominant behaviors, and in a small region of the parameter space, a satisfactory compromise seems to exist which is seen in Figure 12, using parameter values of $(1/3, 1/3)$. This is quite good, considering that the image is being magnified from 64×64 pixels. There is some degree of sharpening, and almost no visible evidence of the sampling lattice.

To get a better idea of which regions of the parameter space yield which type of behavior, a simple subjective test was designed. On a neutral background, four images were displayed typifying the effects of ringing, blurring, anisotropy, and an example of the most satisfactory behavior. In the center of the display, images reconstructed from filters with random values of (B, C) were displayed, and the test subject was asked to choose which of the four behaviors it exemplified. Nine expert observers (researchers working in graphics or image processing) took part and over 500 samples were taken. It would not be credible to suggest that a single ideal parameter pair can be deduced from subjective testing. The motivation for this experiment was simply to draw approximate boundaries between regions of differing behavior as shown in Figure 13. The test subjects were quite consistent with one another in their judgements.

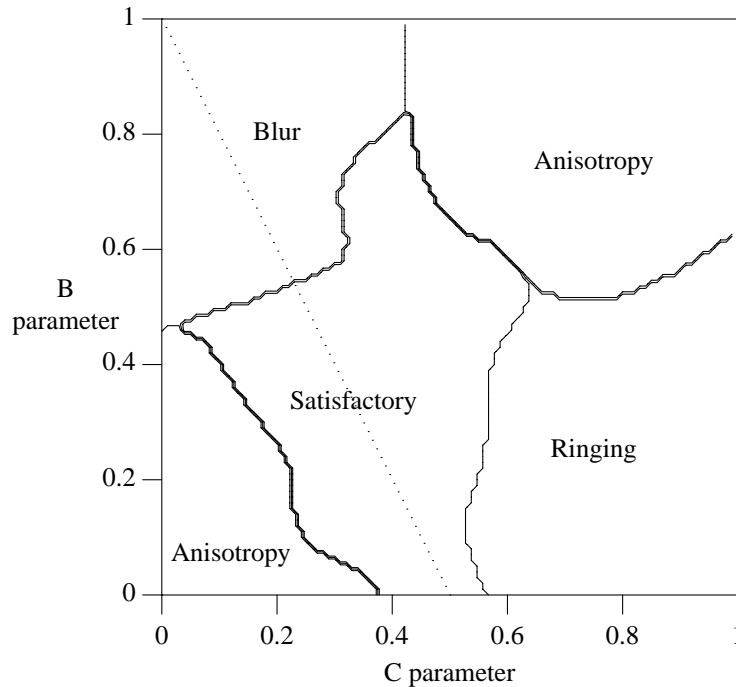


Figure 13. Regions of Dominant Subjective Behavior

To help choose a good filter from the two-parameter space, some quantitative analysis can be done to remove one more degree of freedom. Keys and Park et al. studied the cardinal cubic splines because these cubics exactly interpolate at the sample positions [KEY81,PAR83]. Using standard numerical analysis, Keys concluded that the Catmull-Rom spline was best. Park et al. reached the same conclusion using an equivalent analysis in the frequency domain. Figure 14 illustrates this technique:




Figure 14. $f(x)$ and $f_r(x)$

As the sample spacing h diminishes, the function and its reconstruction become closer. The difference $f(x) - f_r(x)$ can be expanded into a power series in h to study how parameters affect various orders of behavior. Details of this type of analysis can be found in Keys' paper, and when applied to the two-parameter family, the following is obtained:

$$f(x) - f_r(x) = (2C + B - 1) h f' r(x) + O(h^2) \quad (9)$$

$r(x)$ is a polynomial factor. When $2C + B = 1$ (indicated by dotted line in Figure 13), quadratic convergence of fit is achieved. This line contains the cubic B-spline and the Catmull-Rom spline (which actually has cubic convergence). Within the interval of $B = 5/3$ to $B = 0$, good subjective behavior is found with a simple trade-off between blurring and ringing. Outside this interval, $k(x)$ becomes bimodal or exhibits extreme ringing. The filter $(1/3, 1/3)$ used to generate Figure 12 is recommended by the authors, but other observers may prefer more or less ringing.

5. Postaliasing Revisited

A systematic consideration of subjective appearance along with quantitative analysis has yielded an excellent piecewise cubic filter. However, the issue of postaliasing, defined in section 2, has been ignored. In fact the $(1/3, 1/3)$ filter has only fair antialiasing properties and was used to generate Figure 4. Postaliasing is usually not strong enough to cause visible "jaggies" on edges unless a very poor filter is used (e.g., a box filter); however, an image with periodic patterns can have conspicuous postalias Moiré effects unless careful precautions are taken. Synthetic images that contain brick walls, ocean waves, or the ubiquitous checker-board pattern are examples of images that might have this difficulty. There are several approaches to fixing this problem.

If the signal is bandlimited and samples carry information about the derivative as well as about signal amplitude, a better job of reconstruction can be done [PET64]. Given samples (at unit spacing) of a signal and of its derivative, a reconstruction can be done in the following form:

$$f(x) = \sum_{n=-\infty}^{\infty} \left[f_n g(x-n) + f'_n h(x-n) \right] \quad (10)$$

In an extension of the sampling theorem, if the signal contains no energy above the sampling frequency (twice the allowed bandwidth of sampling without derivatives), then it can be perfectly reconstructed by the filter kernels:

$$g(x) = \frac{\sin^2 \pi x}{\pi^2 x^2} \quad (11)$$

$$h(x) = \frac{\sin^2 \pi x}{\pi^2 x} \quad (11b)$$

This is analogous to the ideal sinc reconstruction formula in the standard case where no derivative information is present. A common approximation to these ideal reconstruction formulae is Hermite cubic interpolation:

$$g(x) = \begin{cases} 2|x|^3 - 3|x|^2 + 1 & \text{if } |x| \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

$$h(x) = \begin{cases} x^3 - 2x|x| + x & \text{if } |x| \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (12b)$$

Figure 15 shows the aliasing test pattern (still starting with 128 x 128 samples) reconstructed with the Hermite cubic postfilter. The effect is dramatic when compared to Figure 4. The postaliasing artifact in the middle of the image is nearly gone, and the prealiasing pattern on the right is less intense.

The theory of *derivative reconstruction* may have some practical value in computer graphics. For example, it may be possible to extend Whitted's ray-tracing shading model [WHI80] to generate derivatives with respect to the screen coordinates. This is not an easy problem, but we have demonstrated the feasibility of this extension by deriving the formulae for Lambert and Phong shading on quadric surfaces. It is possible that the density of rays used to reconstruct an image could be reduced in this manner by gathering more useful information from each visible surface calculation.

A second approach to improving postaliasing properties is suggested by the success of stochastic sampling on the prealiasing problem [COO86,DIP85,MIT87]. However, preliminary experiments conducted by the authors with *stochastic-phase reconstruction* have yielded very poor results. The amount of noise needed to obscure postaliasing seriously degraded image quality.

Finally, it was observed in section 2 that signal energy very near the Nyquist frequency is most responsible for conspicuous Moiré patterns. It is possible to cut out this component by *notch-filter reconstruction*. The frequency response of the two-parameter cubic filter in equation (8) is:

$$\begin{aligned} K(\nu) = & \frac{3-3B}{(\pi\nu)^2} \left[\text{sinc}^2(\nu) - \text{sinc}(2\nu) \right] \\ & + \frac{2C}{(\pi\nu)^2} \left[-3\text{sinc}^2(2\nu) + 2\text{sinc}(2\nu) + \text{sinc}(4\nu) \right] \\ & + B \text{sinc}^4(\nu) \end{aligned} \quad (13)$$

This function goes to zero at $\nu = 1/2$ when $B = 3/2$. In fact, the frequency response is zero at all integer and half-integer multiples of the sampling rate except zero. The filter $(3/2, -1/4)$ is quadratically convergent, and the result of reconstruction with it can be seen in Figure 16, in which the postaliasing artifact is almost completely eliminated. Unfortunately, this filter is quite blurry as was seen in Figure 11. The behavior of this notch filter can be seen in its frequency response in Figure 17 compared with the cubic B-spline filter

(1, 0) in Figure 18. The log magnitudes of the frequency responses are plotted below:




Figure 17. Frequency Response of Cubic Notch Filter




Figure 18. Frequency Response of Cubic B-Spline Filter

6. Conclusions

Designing reconstruction filters for computer graphics applications requires a balanced analysis of formal quantitative properties and subjective image quality. There are many trade-offs, and it may be impossible to find a filter that yields good image quality and has good antialiasing properties.

A new family of cubic filters has been analyzed, and two interesting filters have been found. The $(1/3, 1/3)$ filter yields excellent image quality, and the notch filter $(3/2, -1/4)$ strongly suppresses postaliasing patterns.

If derivative values can be generated by a procedural signal, an image with less aliasing is possible by reconstruction with Hermite interpolation or some other suitable filter.

More work remains to be done. While the authors do not believe simple filters will be found that improve much on the cubic filters derived here, there are other avenues for progress. Adaptive filters might allow good image quality with strong antialiasing only where it is needed in problem areas. The effects of the

reconstruction in the display and eye might be allowed for given models of the visual system [NET88].

Finally, the problem of reconstruction from nonuniform sampling is not entirely solved. Reasonable filters have been proposed [MIT87], but more analysis could be done. "Ideal" nonuniform reconstruction filters are known which are analogous to the sinc filter used with uniform samples. A greater challenge will be to understand the subjective issues involved in designing filters that are well suited to computer graphics.

7. Acknowledgements

We would like to thank our colleagues who volunteered to help with subjective testing. We would also like to thank Jim Bergen from the David Sarnoff Laboratory, Jim Johnston from Bell Labs Signal Processing Research Department., and William Schreiber from MIT's Advanced Television Research Program for their views on filter design. We would also like to thank Larry O'Gorman, Bruce Naylor, Rob Pike, David Thomas and Pamela Zave and the SIGGRAPH reviewers for their helpful comments.

8. References

- [BRO69] Brown, Earl F., "Television: The Subjective Effects of Filter Ringing Transients", *Journal of the SMPTE*, Vol. 78, No. 4, April 1969, pp. 249-255.
- [CAT80] Catmull, Edwin, Alvy Ray Smith, "3-D Transformations of Images in Scanline Order", *Computer Graphics*, Vol. 14, No. 3, pp. 279-285.
- [COO86] Cook, Robert L., "Stochastic Sampling in Computer Graphics", *ACM Trans. Graphics*, Vol. 5, No. 1, January 1986.
- [COO87] Cook, Robert L., *personal communication*, August, 1987.
- [CRO77] Crow, Franklin C., "The Aliasing Problem in Computer-Generated Shaded Images", *Comm. ACM*, Vol. 20, No. 11, November 1977, pp. 799-805.
- [DIP85] Dippe, Mark A. Z. and Erling Henry Wold, "Antialiasing Through Stochastic Sampling", *Computer Graphics*, Vol. 19, No. 3, July 1985, pp. 69-78.
- [DUF86] Duff, Tom, "Splines in Animation and Modeling", *State of the Art in Image Synthesis*, SIGGRAPH 86 Course Notes.
- [HOU78] Hou, Hsieh S., Harry C. Andrews, "Cubic Splines for Image Interpolation and Digital Filtering", *IEEE Trans. Acoustics, Speech, and Signal Processing*, Vol. ASSP-26, No. 6, December 1978, pp. 508-517.
- [KEY81] Keys, Robert, G., "Cubic Convolution Interpolation for Digital Image Processing", *IEEE Trans. Acoustics, Speech, and Signal Processing*, Vol. ASSP-29, No. 6, December 1981, pp. 1153-1160.
- [MER84] Mertz, Pierre, and Frank Grey, "A Theory of Scanning and its Relation to the Characteristics of the Transmitted Signal in Telephotography and Television," *Bell System Tech. J.*, Vol. 13, pp. 464-515, July 1934.
- [MIT87] Mitchell, Don P., "Generating Antialiased Images at Low Sampling Densities", *Computer Graphics*, Vol. 21, No. 4, July 1987, pp. 65-72.

- [NET88] Netravali, Arun N., Barry G. Haskell, *Digital Pictures: Representation and Compression*, New York, Plenum, 1988.
- [PAR83] Park, Stephen K., Robert A. Schowengerdt, "Image Reconstruction by Parametric Cubic Convolution", *Computer Vision, Graphics, and Image Processing*, Vol. 23, No. 3, September 1983, pp. 258-272.
- [PET64] Petersen, Daniel P., David Middleton, "Reconstruction of Multidimensional Stochastic Fields from Discrete Measurements of Amplitude and Gradient", *Information and Control*, Vol. 7, pp. 445-476.
- [SCH85] Schreiber, William F., Donald E. Troxel, "Transformation Between Continuous and Discrete Representations of Images: A Perceptual Approach", *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 2, March 1985, pp. 178-186.
- [SHA49] Shannon, Claude E., "Communication in the Presence of Noise.", *Proc. IRE* Vol. 37, 1949, pp. 10-21.
- [WHI80] Whitted, Turner, "An Improved Illumination Model for Shaded Display", *Comm. ACM*, Vol. 23, No. 6, June 1980, pp. 343-349.

