Stochastic Sampling of Two-dimensional Images

László Szirmay-Kalos

Department of Process Control, Technical University of Budapest,

Budapest, Műegyetem rkp. 11, H-1111, HUNGARY

Internet: szirmay@fsz.bme.hu

ABSTRACT

Stochastic sampling is a special anti-aliasing technique that originated in signal theory examining functions of one variable (time). However, in computer graphics two-dimensional images are sampled instead of one dimensional signals, thus the original results must be extended to two-dimensional space. This extension can provide better understanding of the application of stochastic sampling in computer graphics and can lead to new methods which have no one-dimensional counterpart. In this paper a two-dimensional adaptive filtering technique is discussed, which can eliminate aliasing artifacts without introducing two much noise in the resulting image.

Keywords: Stochastic sampling, jitter, filtering, anti-aliasing.

INTRODUCTION

From the information or signal processing point of view, modeling can be regarded as the definition of the intended world by digital and discrete data which are processed later by image synthesis. Since the intended world model, like the real world, is continuous, modeling always involves an analog-digital conversion to the internal representation of the digital computer. Later in image synthesis, the digital model is resampled and requantized to meet the requirements of the display hardware, which is much more drastic than the sampling of modeling, making this step responsible for the generation of artifacts due to the approximation error in the sampling process. Sampling methods applying regular grids produce regularly spaced artifacts that are easily detected by the human eye, since the eye is especially sensitive to regular and periodic signals. Random placement of sample locations can break up the periodicity of the aliasing artifacts, converting the aliasing effects to random noise which is more tolerable for human observers. Two types of random sampling patterns have been proposed [Cook1986], namely the **Poisson disk distribution** and the **jittered** sampling.

Jittered sampling is based on a regular sampling grid which is perturbed slightly by random noise. Unlike the application of dithering algorithms, the perturbations are now assumed to be independent random variables. Compared to Poisson disk sampling its result is admittedly not quite as good, but it is less expensive computationally and is well suited to image generation algorithms designed for regular sampling grids.

Jittered sampling was first examined in 1962 by Balakrishnan who analyzed it as a negative phenomenon in sampling of continuous time functions [Balakrishnan1962]. More than twenty years later Cook [Cook1986] realized that the effects of stochastic sampling can be advantageous in computer graphics to reduce aliasing artifacts. He relied on the original paper of Balakrishnan and took advantage of notational simplicity of single variate functions. His followers used the same approach, and have achieved important practical results [Cook *et al.*1987]. However, the image space is two-dimensional, thus discussing its sampling using one-dimensional analogy may miss important properties. This paper tries to fill that gap by analyzing the stochastic sampling in the two-dimensional space and by proposing sampling strategies that are based on this more general formulation.

JITTERED SAMPLING IN THE TWO-DIMENSIONAL SPACE

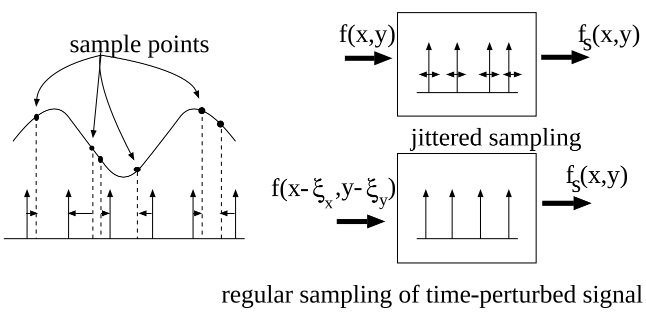

Figure 1: Signal processing model of jittered sampling

First the theory of jittered sampling will be discussed in two dimensions. Suppose function $f(x,y)$ is sampled and then reconstructed by an ideal low-pass filter. The perturbations of the various sample locations are assumed to be uncorrelated random vector variables defined by the probability density function $p(\alpha,\beta)$. The effect of jittering can be simulated by replacing $f(x,y)$ by $g(x,y) = f(x-\xi_x(x,y), y-\xi_y(x,y))$, and sampling it by a regular grid, where function $\vec{\xi}(x,y) = [\xi_x(x,y), \xi_y(x,y)]$ is an independent stochastic vector process whose probability density function is $p(x,y)$ (Fig. 1).

Jittered sampling can be analyzed by comparing the spectral power distributions of $g(x,y)$ and $f(x,y)$. Since $g(x,y)$ is a random process, if it were stationary and ergodic [Lamperti1972], then its frequency distribution would be best described by the power density spectrum which is the Fourier transform of its autocorrelation function.

The autocorrelation function of $g(x,y)$ is derived as an expectation value:

$$R(x,y,u,v) = E[g(x,y)\cdot g(x+u,y+v)]. \quad (1)$$

If $u \neq 0 \vee v \neq 0$, then $\vec{\xi}(x,y)$ and $\vec{\xi}(x+u,y+v)$ are stochastically independent random variables, thus we get :

$$R(x,y,u,v) = E[g(x,y)\cdot g(x+u,y+v)] =$$
$$E[g(x,y)]\cdot E[g(x+u,y+v)]. \quad (2)$$

The expectation value of $g(x,y)$ is

$$E[g(x,y)] = E[f(x-\xi_x(x,y), y-\xi_y(x,y))] =$$
$$\int_{\alpha=-\infty}^{\infty}\int_{\beta=-\infty}^{\infty} f(x-\alpha,y-\beta)\cdot p(\alpha,\beta)\; d\alpha\; d\beta =$$
$$(f*p)|_{x,y}, \quad (3)$$

where $f*p$ is the convolution of the two functions. Thus the autocorrelation function for any $u \neq 0 \vee v \neq 0$ is:

$$R(x,y,u,v) = (f*p)|_{x,y}\cdot(f*p)|_{x+u,y+v}. \quad (4)$$

If $u,v = 0$, then:

$$R(x,y,0,0) = E[g^2(x,y)] =$$
$$\int_{\alpha=-\infty}^{\infty}\int_{\beta=-\infty}^{\infty} f^2(x-\alpha,y-\beta)\cdot p(\alpha,\beta)\; d\alpha\; d\beta, \quad (5)$$

that is the second moment of $g(x,y)$. The autocorrelation function, for any u,v, is:

$$R(x,y,u,v) = (f*p)|_{x,y}\cdot(f*p)|_{x+u,y+v} +$$
$$[E[g^2(x,y)] - E^2[g(x,y)]]\cdot\delta(u,v), \quad (6)$$

where $\delta(u,v)$ is the delta function which is 1 for $u,v=0$ and 0 for $u \neq 0 \vee v \neq 0$. This delta function introduces an "impulse" in the autocorrelation function at $u,v=0$. The size of the impulse in the autocorrelation function is the square variance of the random variable $g(x,y)$, that is

$$\sigma^2_{g(x,y)} = E[g^2(x,y)] - E^2[g(x,y)].$$

Unfortunately $g(x,y)$ is usually not a stationary process, thus in order to analyze its spectral properties, the power density spectrum is calculated from the "average" autocorrelation function which is defined as:

$$\hat{R}(u,v) = \lim_{X,Y\to\infty} \frac{1}{4XY} \int_{x=-X}^{X} \int_{y=-Y}^{Y} R(x,y,u,v)\, dx\, dy. \tag{7}$$

The power density of $g(x,y)$ is the Fourier transform of the autocorrelation function as defined by the following formula ($\jmath = \sqrt{-1}$):

$$S_g(\zeta,\eta) = \mathcal{F}^{\zeta}\mathcal{F}^{\eta}\ \hat{R}(u,v) = \int_{u=-\infty}^{\infty} \int_{v=-\infty}^{\infty} \hat{R}(u,v)\cdot e^{-2\pi\jmath\zeta u} e^{-2\pi\jmath\eta v}\ du\ dv. \tag{8}$$

This integral can be expressed using some identity relations:

$$S_g(\zeta,\eta) = \mathcal{F}^{\zeta}\mathcal{F}^{\eta}\hat{R}(u,v) = \mathcal{F}^{\zeta}\mathcal{F}^{\eta} \lim_{X,Y\to\infty} \frac{1}{4XY} \int_{-X}^{X} \int_{-Y}^{Y} R(x,y,u,v)\ dxdy =$$

$$\lim \frac{1}{4XY} \int_{-X}^{X} \int_{-Y}^{Y} \left[(f*p)|_{x,y} \cdot \mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p)|_{x+u,y+v} + \sigma_g^2\right] dxdy. \tag{9}$$

Since

$$\mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p)|_{x+u,y+v} = e^{2\pi\jmath(\zeta x+\eta y)} \cdot \mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p),$$

we have

$$S_g(\zeta,\eta) = \left[\lim_{X,Y\to\infty} \frac{1}{4XY} \int_{-X}^{X} \int_{-Y}^{Y} (f*p)|_{x,y} e^{2\pi\jmath(\zeta x+\eta y)} dxdy\right] \cdot$$

$$\cdot \mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p) +$$

$$+ \lim_{X,Y\to\infty} \frac{1}{4XY} \int_{-X}^{X} \int_{-Y}^{Y} \sigma_g^2\ dx\ dy =$$

$$\lim \frac{1}{4XY} [\mathcal{F}_X^{\zeta}\mathcal{F}_Y^{\eta}\ (f*p)]^* \cdot [\mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p)] + \overline{\sigma^2_{g(x,y)}}, \tag{10}$$

where superscript $*$ means the conjugate complex pair of a number, $\overline{\sigma^2_{g(x,y)}}$ is the average variance of the random variable $g(x,y)$ for different x,y values, and $\mathcal{F}_{X,Y}$ stands for the truncated Fourier transform defined by the following equation:

$$\mathcal{F}_X^{\zeta}\mathcal{F}_Y^{\eta}\ w(x,y) = \int_{x=-X}^{X} \int_{y=-Y}^{Y} w(x,y) \cdot e^{-2\pi\jmath\zeta x} e^{-2\pi\jmath\eta y}\ dx\ dy. \tag{11}$$

Let us compare this power density ($S_g(\zeta,\eta)$) of the perturbed signal with the power density of the original function $f(x,y)$, which can be defined as follows:

$$S_f(\zeta,\eta) = \lim_{X,Y\to\infty} \frac{1}{4XY} |\mathcal{F}_X^{\zeta}\mathcal{F}_Y^{\eta} f|^2 \tag{12}$$

This can be substituted into Eq. 10 yielding:

$$S_g(\zeta,\eta) = \frac{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p)|^2}{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta} f|^2} \cdot S_f(\zeta,\eta) + \overline{\sigma_g^2} \tag{13}$$

The spectrum consists of a part proportional to the spectrum of the unperturbed $f(x,y)$ signal and an additive noise carrying $\overline{\sigma_g^2}$ power in a unit frequency range. Thus the perturbation of the sample positions can, in fact, be modeled by a linear network or filter and some additive noise (Fig. 2).

The gain of the filter perturbing the sample positions by an independent random process can be calculated as the ratio of the power density distributions of $f(x,y)$ and $g(x,y)$ ignoring the additive noise:

$$Gain(\zeta,\eta) = \frac{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta}(f*p)|^2}{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta} f|^2} = \frac{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta} f|^2 \cdot |\mathcal{F}^{\zeta}\mathcal{F}^{\eta} p|^2}{|\mathcal{F}^{\zeta}\mathcal{F}^{\eta} f|^2} = |\mathcal{F}^{\zeta}\mathcal{F}^{\eta} p|^2 \tag{14}$$

Thus, the gain is the Fourier transform of the probability density used for jittering the sample positions.

The most often used jitter is the **white noise jitter** which distributes the jitter values uniformly between $[-T/2,-T/2]$ and $[T/2,T/2]$, where T is the periodicity of the regular sampling grid and jittering in x and y directions are stochastically independent. The gain of the white noise jitter is:

$$Gain_{wn}(\zeta,\eta) = \mathrm{sinc}^2(\pi\zeta T) \cdot \mathrm{sinc}^2(\pi\eta T) \tag{15}$$

where $\mathrm{sinc}(x) = \sin x/x$.

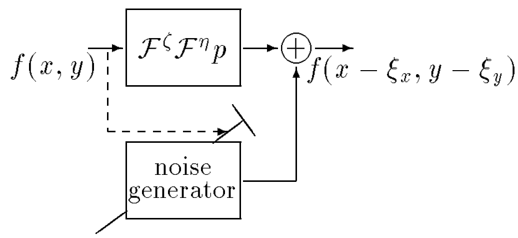

Figure 2: System model of sample position perturbation

The white noise jitter is a fairly good low-pass filter suppressing the spectrum of the sampled signal above the Nyquist limit, and thus greatly reducing aliasing artifacts.

Jittering trades off aliasing for noise. In order to intuitively explain this result, let us consider the sample position perturbation for a sine wave. If the extent of the possible perturbations is less than the length of half a period of the sine wave, the perturbation does not change the basic shape of the signal, just distorts it a little bit. The level of distortion depends on the extent of the perturbation and the "average derivative" of the perturbed function as suggested by the formula of the noise intensity defining it as the variance $\overline{\sigma_g^2}$. If the extent of the perturbations exceeds the length of period, the result is an almost random value in place of the amplitude. The sine wave has disappeared from the signal, only the noise remains.

FILTER CONSTRUCTION

Filters based on stochastic sampling can thus be constructed by selecting the probability density of the jitter appropriately. To do so, two different criteria must be taken into consideration. This filter must suppress frequencies above the Nyquist limit without destroying the low-frequency ranges. On the other hand, the noise introduced by the method should be minimized. However, the two different criteria are contradicting. The optimal low-pass filter would have a sinc-like inverse Fourier transform, which would require a sinc like probability density function. The sinc function is known to have quite wide interval where it significantly differs from zero, which makes the average square variance $\overline{\sigma_g^2}$ big, and thus the resulting image rather noisy. On the other hand, the noise could be minimized by a concentrated, delta function-like density function which has no high frequency suppression at all.

Obviously, a compromise must be chosen between aliasing and noise. Examining the images generated by white-noise jitters, we can conclude that the compromise offered by this jitter is not at all optimal, but sacrifices the minimal noise criterion for the elimination of the aliasing artifacts (Fig. 5). The compromise should depend on the properties of the image. If it contains high frequency components, then the primary objective is the elimination of those frequency components that are above the Nyquist limit. If the high frequency components are not significant, the noise must be minimized. Moreover, this compromise is not necessarily global. The priorities of the filter may be required to change over the image. This idea leads to an adaptive filtering method that evaluates the frequency characteristics of a small portion of the image and adapts itself to the measured properties and prefers high-frequency suppression to noise or vice versa. In order to make this idea practically useful, the calculations needed to measure the frequency characteristics should not be much more complex than the simplest stochastic sampling algorithm. An algorithm meeting this requirement is discussed in the next section.

ADAPTIVE FILTER CONSTRUCTION

In order to introduce the proposed adaptive filter, we first describe the basic idea informally, then a rigorous mathematical analysis is provided, finally simulation results are shown.

Informal description of the algorithm

Let us examine the compromise of the noise and aliasing by controlling the noise factor. As calculated, the intensity of the noise resulting from the stochastic sampling is $\overline{\sigma_g^2}$. Looking at Fig. 3, we can see that the intensity of the noise is the square variance of a random variable controlled by the probability

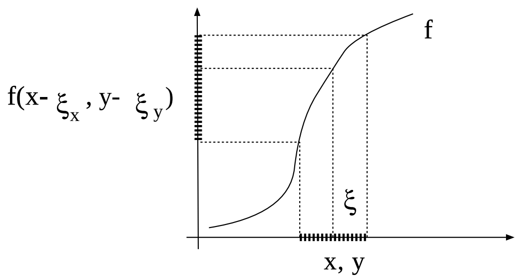

Figure 3: Noise intensity

density function of the filter and amplified by the gradient of the sampled function f. This means that the noise can be kept under control if the direction of the jitter is selected to minimize the amplification of the gradient of the signal. This requires the calculation of the gradient of $f(x,y)$ in the center of the pixel, and the perturbation of the sample position to be selected from a line perpendicular to the gradient. This jitter is called the **line-jitter**. Comparing this to the white-noise jitter, it is quite obvious that restricting jitter to a single line instead of a rectangle decreases the high-frequency suppression of the filter.

To calculate the gradient we must use discrete samples of a possibly non band-limited signal, thus aliasing might also occur in this calculation. To reduce aliasing the tricks of stochastic sampling can be used again. If the sample points used to calculate the gradient is perturbed by a white-noise jitter, then the high-frequency components will be suppressed and noise will be added to the calculated gradient. This means that in image portions having low frequency characteristics the calculated gradient dominates the noise, but for high frequencies the results calculated in this way are random variables. The direction calculated in this way approximates the real gradient for low-frequency signals, but becomes a random variable for signals above the Nyquist limit. This property makes this strategy adaptive, which behaves as a line jitter for low-frequency signals but tends to behave as a white-noise jitter for high frequencies.

Formal analysis

Calculation of the signal gradient

The gradient of a signal can be approximated by the following differences:

$$\mathbf{grad} f \approx \Delta(x,y) =$$
$$[f(x+\tfrac{1}{2}, y-\tfrac{1}{2}) - f(x-\tfrac{1}{2}, y-\tfrac{1}{2}),$$
$$f(x-\tfrac{1}{2}, y+\tfrac{1}{2}) - f(x-\tfrac{1}{2}, y-\tfrac{1}{2})]. \quad (16)$$

However, this approximation also uses samples of the signal which can cause aliasing artifacts in the gradient. This can be recognized having evaluated the Fourier transform of the gradient signal:

$$\mathcal{F}^{\zeta}\mathcal{F}^{\eta}\Delta(x,y) =$$
$$[e^{-\jmath\pi\eta}\jmath \cdot \sin\pi\zeta, e^{-\jmath\pi\zeta}\jmath \cdot \sin\pi\eta] \cdot \mathcal{F}^{\zeta}\mathcal{F}^{\eta} f. \quad (17)$$

For small frequencies, we can use the Taylor's approximation:

$$\mathcal{F}^{\zeta}\mathcal{F}^{\eta}\Delta(x,y) \approx [\jmath\pi\zeta, \jmath\pi\eta] \cdot \mathcal{F}^{\zeta}\mathcal{F}^{\eta} f =$$
$$\mathcal{F}^{\zeta}\mathcal{F}^{\eta}[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}]. \quad (18)$$

Thus, for small frequencies, this really approximates the gradient of f. For high frequencies, however, the results are very far from the gradient of the function.

To avoid aliasing, the trick of stochastic sampling is used. If the sample point is perturbed by random vector variable ξ, we get:

$$\mathbf{grad} f(x,y) \approx$$
$$[f(x+\tfrac{1}{2}+\xi_x, y-\tfrac{1}{2}+\xi_y) - f(x-\tfrac{1}{2}+\xi_x, y-\tfrac{1}{2}+\xi_y),$$
$$f(x-\tfrac{1}{2}+\xi_x, y+\tfrac{1}{2}+\xi_y) - f(x-\tfrac{1}{2}+\xi_x, y-\tfrac{1}{2}+\xi_y)]. \quad (19)$$

The effect of perturbing the sample point by a white-noise jitter distributed between $[-1/4, -1/4]$ and $[1/4, 1/4]$ is equivalent to the application of a low-pass filter of gain $\mathrm{sinc}^2(\pi\zeta/2) \cdot \mathrm{sinc}^2(\pi\eta/2)$ and addition of some random noise. The effective gain of the gradient filter with random sampling is:

$$Gain(\zeta, \eta) =$$
$$\mathrm{sinc}^2(\pi\zeta/2) \cdot \mathrm{sinc}^2(\pi\eta/2)[\sin^2\pi\zeta, \sin^2\pi\eta], \quad (20)$$

which forces the signal gradient to vanish for frequencies above the Nyquist limit, but gives back the real gradient for lower frequencies. For higher frequencies the noise added to the signal will determine the values sampled.

Analysis of the line-jitter

Let $[v_x, v_y]$ be a unit vector that is perpendicular to the gradient, that is:

$$[v_x, v_y] = [-\frac{\partial f}{\partial y}, \frac{\partial f}{\partial x}] \Big/ \sqrt{\left(\frac{\partial f}{\partial y}\right)^2 + \left(\frac{\partial f}{\partial x}\right)^2} \tag{21}$$

Then the sampled signal is $f(x + v_x \cdot l, y + v_y \cdot l)$ where l is a random variable. Let $[-L/2, L/2]$ be the domain of this random variable. This leads to a stochastic sampling scheme where the sampling position is selected from a line segment of length L centered at the pixel point, unlike white noise jitter which selected the position form a square. Assume that the point is selected from the line segment by a uniform distribution. Thus the two dimensional probability density function of the sampling location can be expressed as:

$$p(x, y) = \frac{1}{L} \int_{l=-L/2}^{L/2} \delta(x - v_x \cdot l, y - v_y \cdot l)\, dl. \tag{22}$$

The gain of the filter realized by this stochastic sampling can be evaluated using Equ. 14 as follows:

$$Gain(\zeta, \eta) = |\mathcal{F}^\zeta \mathcal{F}^\eta p|^2 =$$

$$|\frac{1}{L} \int_{l=-L/2}^{L/2} \int_x \int_y \delta(x - v_x \cdot l, y - v_y \cdot l) \cdot e^{-2\pi \jmath (\zeta x + \eta y)}\ dx\ \ dy\ \ dl|^2 =$$

$$\text{sinc}^2(\pi L(\zeta v_x + \eta v_y)). \tag{23}$$

Comparing this frequency function to that of the white-noise jitter (Eq. 15), it has worse high frequency suppression, limiting the use of this jittering type when the signal sampled is under the Nyquist limit. The adaptive property of the gradient calculation, however, ensures that v_x and v_y are approximately constant only if the signal has no high frequency components. For signals of higher frequencies v_x and v_y are random variables, which makes the line-jitter similar to the white-noise jitter that is good in eliminating aliasing artifacts.

SIMULATION RESULTS AND CONCLUSIONS

In order to demonstrate the merits of the new adaptive algorithm, the widely accepted test figure of a chess-table is used and is rendered by mathematical sampling, white-noise jitter and finally by the new adaptive jitter. The images demonstrate that the adaptive jitter does not add too much noise to the image where it is not necessary unlike white-noise jitter, but it is not worse than that in eliminating aliasing. In low-frequency ranges of the picture demonstrating the adaptive jitter, that is in the front section of the picture, the noise does not destroy the image. The result is similar to that of the mathematical sampling. In the back section, however, where frequency is higher than allowed by the sampling theorem, the Moire pattern generated from the chess-board is fully replaced by noise.

Concerning the effectiveness of the algorithm, it should be mentioned that it requires just three more samples and the calculations required are simple.

ACKNOWLEDGEMENT

This work has been supported by the Hungarian National Scientific Research Fund (OTKA), under the reference Number F-015884.

References

A.V. Balakrishnan. On the problem of time jitter in sampling. *IRE Trans. Inf. Theory*, Apr:226–236, 1962.

Robert L. Cook, Loren Carpenter, and Edwin Catmull. The reyes image rendering architecture. In *Proceedings of SIGGRAPH '87*, pages 95–102, 1987.

Robert L. Cook. Stochastic sampling in computer graphics. *ACM Transactions on Graphics*, 5(1):51–72, 1986.

John F. Lamperti. *Stochastic Processes*. Springer-Verlag, 1972.

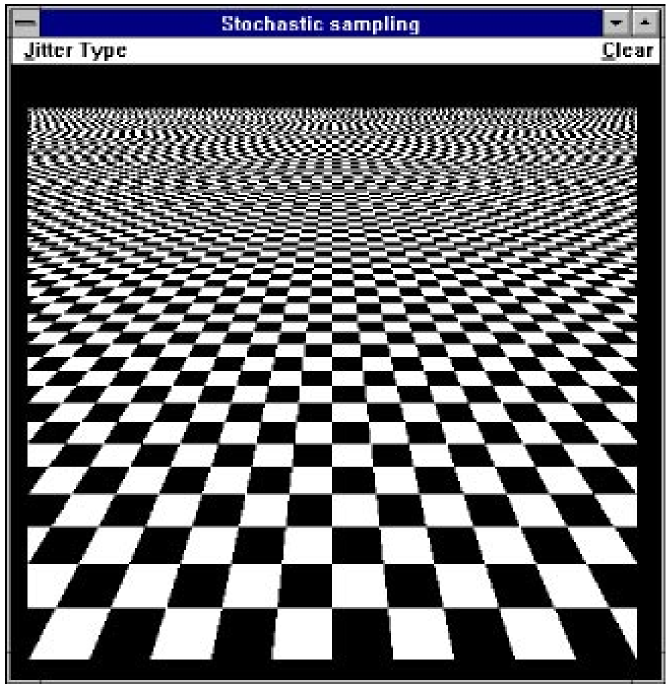

Figure 4: Mathematically sampled image

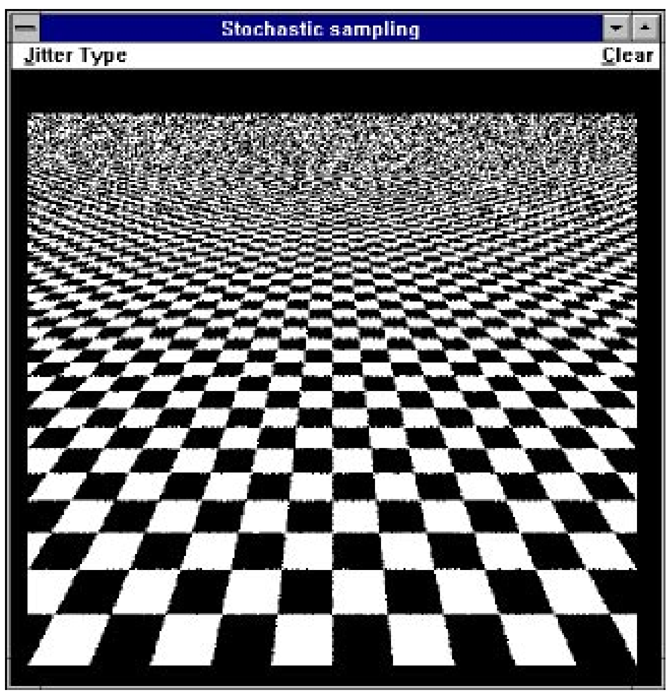

Figure 5: Image sampled using white noise jitter

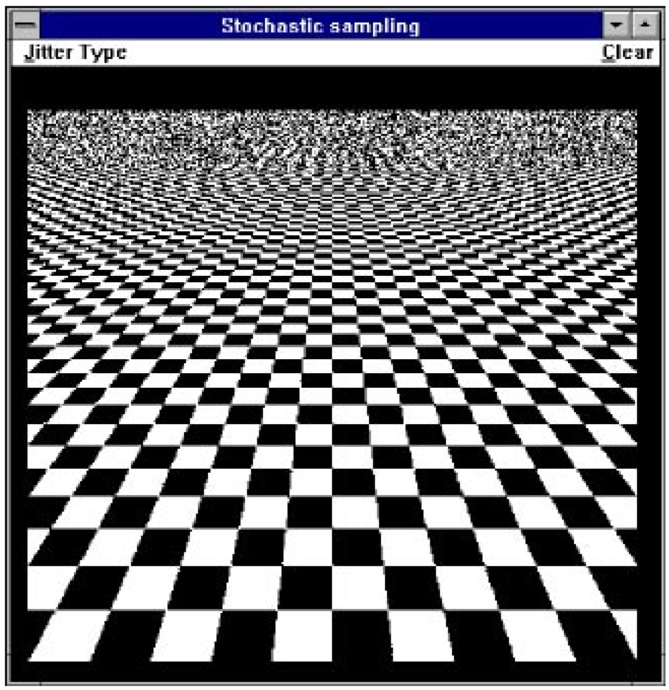

Figure 6: Image sampled using adaptive jitter