

A Weighted Error Metric and Optimization Method for Antialiasing Patterns

Samuli Laine and Timo Aila

Helsinki University of Technology, Telecommunications Software and Multimedia Laboratory
Hybrid Graphics Ltd.

Abstract

Displaying a synthetic image on a computer display requires determining the colors of individual pixels. To avoid aliasing, multiple samples of the image can be taken per pixel, after which the color of a pixel may be computed as a weighted sum of the samples. The positions and weights of the samples play a major role in the resulting image quality, especially in real-time applications where usually only a handful of samples can be afforded per pixel. This paper presents a new error metric and an optimization method for antialiasing patterns used in image reconstruction. The metric is based on comparing the pattern against a given reference reconstruction filter in spatial domain and it takes into account psychovisually measured angle-specific acuities for sharp features.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation – Antialiasing

1. Introduction

In supersampling or multisampling algorithms the process of image synthesis is divided into two stages: sampling and reconstruction. The choice of sampling positions plays a critical role in the resulting antialiasing quality. This is especially true when only few samples can be afforded for each pixel, which is often the case in real-time rendering.

In this paper, we consider only *linear* reconstruction filters. A linear reconstruction filter computes the final color of a pixel as a weighted sum of the samples. In high-quality systems the support of the filter typically extends beyond the boundaries of a single pixel, and thus the same samples are used for reconstruction of several pixels. This is commonly referred to as *sample sharing*.

Usually sampling is the dominating cost compared to reconstruction, thus making it critical to maximize the quality resulting from a given number of samples. Optimally the sampling positions should be selected so that the perceived aliasing artifacts are minimized in the resulting image. This is a somewhat of a fuzzy goal, and numerous sampling patterns have been introduced, along with a few techniques for analyzing the patterns [CA87, Shi91, DEM96].

It is known that visually more pleasing results are obtained by employing a different pattern for adjacent pixels [KH01, DN02], as it further converts aliasing artifacts to less disturbing noise.

Sampling and filtering In supersampling or multisampling algorithms, the image is first sampled at the sampling points of the sampling pattern. The final color of a pixel is then computed as a weighted sum of the sampled colors. However, this is only an approximation of the mathematically correct procedure, which would first filter the image analytically with the desired reconstruction filter and sample the result at the centers of pixels.

Overview of contributions Our error metric evaluates the perceptual difference between sampling and filtering, comparing a given sampling pattern to an arbitrary *reference reconstruction filter*. We take into account the fact that the observed amount of aliasing depends on the screen-space slope of a sharp feature [Nai98]. We also present an efficient optimization method for generating patterns with good characteristics. The analysis is equally applicable to modern hardware-friendly sample sharing patterns that take as few as 1.25 samples per pixels [AM03], and to high-quality offline

patterns that consist of dozens of samples per pixel [Coo86]. As the number of samples per pixel increases, the importance of carefully choosing the positions and weights of the sampling points diminishes. Therefore, the presented metric is primarily useful for evaluating and optimizing low-cost antialiasing patterns.

The rest of this paper is organized as follows. The related techniques are briefly reviewed in Section 2. Section 3 introduces a classification scheme for antialiasing patterns in order to simplify the discussion. The new error metric is derived in Section 4, and the related optimization method in Section 5. Test results and discussion are provided in Section 6. Finally, conclusions and future work are offered in Section 7.

2. Related Work

The four major sources that contribute to aliasing in rendered images are discontinuity edges, shader undersampling, intersecting surfaces and objects falling between sampling points [Cro77]. Supersampling is a general technique that reduces all aliasing artifacts by increasing the sampling rate. Multisampling differs from supersampling in that the shader is not evaluated separately in all sampling points, and thus multisampling cannot reduce aliasing due to shader undersampling. Programmable shaders can band-limit themselves (texture fetches already do), and in most cases this should be a more efficient way for solving the shader aliasing problem than supersampling.

We focus on reconstruction of images from samples, and thus the sampling strategies for higher-dimensional integrals, global light transport, soft shadows, and other such phenomena are outside the scope of this article.

Psychovisual results Mitchell and Netravali [MN88] observe that the theoretically optimal *sinc* is not a good reconstruction filter in practice, and define a family of bi-cubic filters that are parameterized by two parameters. Optimal values for these parameters are then determined by conducting psychovisual tests. The $(\frac{1}{3}, \frac{1}{3})$ filter is found to produce good overall image quality, whereas the $(\frac{3}{2}, -\frac{1}{4})$ notch filter strongly suppresses postaliasing patterns but results in blurry images. The resulting filters have been recently implemented in high-end graphics hardware [DN02]. We use the $(\frac{1}{3}, \frac{1}{3})$ filter as the reference reconstruction filter in our test measurements in Section 6.

Naiman [Nai98] observes that jaggedness is more disturbing in some orientations of edges or features. He provides a table of psychovisual acuities for each direction. For example, almost vertical or horizontal features require the most antialiasing efforts.

Bouville et al. [BTB91] utilize psychovisual results that resemble Naiman's acuity factors. They argue that sampling

density can be halved by replacing a regular sampling grid with a Quincunx pattern (Figure 1d).

Spectral analysis of sampling patterns Several authors [DW85, Coo86, Mit87, HDK01] have studied the spectral characteristics of sampling patterns. Human eye rarely aliases even though the density of cones in extrafoveal regions is relatively low. Yellott [Yel83] observes that the cones (sampling points) are organized into a Poisson disk distribution, i.e., the points are spread randomly so that no two points are closer to each other than a specified distance. He further argues that such sampling patterns are very good in converting aliasing to less disturbing noise.

The power spectrum of a Poisson disk pattern roughly corresponds to blue noise, which implies two criteria. First, the spectrum is noisy and lacks high spikes of energy that could result in coherent aliasing artifacts. Second, the energy is concentrated in high frequencies, corresponding to irregular noise. While the spectral analysis clearly correlates with the quality of a sampling pattern, quality estimation of a pattern is somewhat ambiguous, and the technique does not suit that well to repetitive patterns that have only a few samples per pixel [AM03]. Also, psychovisual acuities are difficult to take into account.

Analysis using discrepancy Discrepancy is a quantity that measures the ability of a sampling pattern to estimate areas of subregions inside a pixel. Shirley [Shi91] computes discrepancy by sampling the area with rectangles of different sizes and aspect ratios, and Dobkin et al. [DM93] measure discrepancy based on arbitrary edges through a pixel. Matousek [Mat99] offers a more complete discussion on different types of discrepancy.

Quasi Monte Carlo methods [Kel01] typically use deterministic low-discrepancy sequences (e.g. [Hal70]) for generating the sampling patterns. Dobkin et al. [DEM96] use simulated annealing for creating low-discrepancy sampling patterns.

Glassner [Gla95] points out that commonly used simple discrepancy metrics do not directly correlate with the quality of images generated with the patterns. Also, discrepancy does not take reconstruction filters into account. In a comment on Cook's work [Coo86], Pavlidis [Pav90] emphasizes that errors arising from reconstruction can be more significant than the ones resulting from sampling. Our method takes reconstruction filters explicitly into account.

Other analysis methods Chen and Allebach [CA87] analyze sampling patterns for estimation of band-limited signals, without employing reconstruction filters. Their method is based on the mean square error between the reference signal and a sampled version, so that the effect of an input signal is eliminated by computing a maximum error over a class of signals. The metric favors regular grids, and is therefore not well suited for comparing the quality of antialiasing patterns.




Figure 1: Common sampling patterns. From left to right: 2×2 uniform supersampling, 2×2 jittered, rotated grid, Quincunx, Flipquad and Fliptri.

Sampling patterns The quality of uniform sampling patterns (Figure 1a) can be improved by jittering (Figure 1b) [CPC84]. N-rooks [Shi91] patterns divide a pixel into $N \times N$ grid and place one sample to each row and column. A special case with four samples per pixel is the rotated grid supersampling (Figure 1c). Kollig and Keller [KK02] construct low-discrepancy patterns that combine the strengths of jittered and N-rooks sampling.

Poisson disk patterns have good spectral characteristics, and thus also desirable properties for image synthesis [DW85]. Molnar [Mol91] places unweighted sampling points so that the minimum distance of a Poisson disk process is locally modified according to the impulse response of a reconstruction filter. Hierarchical construction of blue noise sampling patterns have been studied by McCool and Fiume [MF92], Hiller et al. [HDK01], and Ostromoukhov et al. [ODJ04]. Poisson disk patterns are often advocated as theoretically optimal in the sense of signal processing. Nevertheless, Goss and Wu [GW00] observe that the spectra of Poisson disk patterns exhibit clearly visible banding.

Sample sharing techniques use the same samples in the reconstruction of multiple output pixels. Quincunx [BTB91, NVI01] generates two samples per pixel and uses five samples for reconstruction (Figure 1d). Flipquad [AMS03] results in higher quality images, while also creating two samples per pixel and using only four samples for reconstruction (Figure 1e). Due to mirroring along the x- and y-axes there are four distinct patterns. Fliptri [AM03] is similar to Flipquad, but takes on average only 1.25 samples per pixel (Figure 1f). These techniques typically introduce a slight low-pass effect on the final image.

3. Classification of Antialiasing Patterns

In our analysis, we consider only *repeating* sampling patterns with the additional constraint that the same number of samples is taken for each pixel. Most existing sampling patterns fall into this category. Using a repeating sampling pattern greatly diminishes its storage requirements, which is crucial if the sampling pattern is to be implemented in hardware. However, it is not necessary to repeat the pattern after one pixel; e.g. in SAGE graphics architecture [DN02] the pattern repeats in 128×128 pixel blocks.

Pattern	s	r	n
Rotated grid [Shi91]	1	1	4
Quincunx [NVI01]	1	1^+	2
SAGE [DN02]	128^2	25	16
Flipquad [AMS03]	4	1^+	2
Fliptri [AM03]	4	1^+	1.25

Table 1: Classification of some common patterns in our $P(s, r, n)$ notation, where s is the number of distinct pixel-sized sample sets, r is the number of pixels accessed by the reconstruction filter and n is the average number of samples per pixel.

We define a pattern family $P(s, r, n)$, where s is the number of distinct pixel-sized sample sets, r is the number of pixels accessed by the reconstruction filter and n is the average number of generated samples per pixel. In general, n does not have to be an integer, e.g., the Fliptri pattern has $n = 1.25$. We write r^+ in place of r when the reconstruction filter also takes into account the samples that lie exactly on the border of multiple pixels (see Figure 2b). Table 1 summarizes the values of s , r and n for some common patterns, and Figure 2 shows three examples. In the remainder of this paper, we assume that every sample has an associated weight. For example, in the sampling patterns of Figure 1, Quincunx and Fliptri have samples with distinct weights, whereas in uniform, jittered, and rotated grids, and in Flipquad every sample has an equal weight.

When $s > 1$, the pattern defines multiple distinct *reconstruction kernels* of size r . Figure 3 shows the structure of a $P(4, 4, n)$ pattern and its four reconstruction kernels. A separate set of sample weights is required for each reconstruction kernel. Our error metric handles every reconstruction kernel separately.

The simplest pattern family is $P(1, 1, n)$, with the special case of $n = 1$ corresponding to no filtering. In this family, the samples are located in the same places in all pixels, and the reconstruction filter uses only the samples inside the current pixel (Figure 2a). This leads to simple software and hardware implementations, but no sample sharing occurs and many sampling points are typically needed for acceptable image quality.

The simplest pattern families that allow sample sharing are $P(1, 1^+, n)$ containing Quincunx, and $P(4, 1^+, n)$ containing Flipquad and Fliptri (Figure 2b). In the family $P(4, 1^+, n)$ a good set of samples is typically determined for only one pixel, and the rest of the sets are then obtained by mirroring along the axis of translation. This implicitly results in interleaved sampling [KH01]. This is a special case because in general s distinct sets of samples must be considered.

The simplest pattern family with a larger reconstruction




Figure 2: a) In $P(1, 1, n)$ only the samples inside a pixel are used for reconstruction. b) In $P(4, 1^+, n)$ samples on the border of the pixel are used as well. c) In $P(1, 4, n)$ samples from four pixels are used for reconstruction.




Figure 3: Patterns in family $P(s, r, n)$ have exactly s distinct reconstruction kernels for different pixels. a) For example, in pattern family $P(4, 4, n)$ the pattern loops in 2×2 tiles. b) This results in four reconstruction kernels. The capital letters stand for distinct pixel-sized sets of n samples. The output color is always reconstructed to the center of the reconstruction kernel, marked with large dots. If the size of the reconstruction kernel is even ($2 \times 2, 4 \times 4, \dots$ pixels), the output pixel grid needs to be shifted 0.5 pixels vertically and horizontally. This implementation trick removes the need of resorting to a 3×3 pixel reconstruction kernel with a reconstruction filter that has a radius of 1.0.

filter is $P(1, 4, n)$ (Figure 2c). As in $P(1, 1, n)$, the samples are located in the same places in every pixel, but the reconstruction filter now shares samples between adjacent pixels. Four sets of sample weights are needed because each sample is used in reconstruction of four output pixels.

The most complex pattern families discussed in this paper are $P(4, 4, n)$ and $P(16, 16, n)$. They combine the properties of $P(4, 1^+, n)$ and $P(1, 4, n)$ families: multi-pixel reconstruction filter and multiple sets of sample positions and weights. We observed that the patterns in these families produce superior images compared to the previously discussed families, with respect to the number of samples taken per pixel.

4. Error Metric for Antialiasing Patterns

We now turn into defining an error metric that estimates the *approximation error* caused by sampling, when compared against the mathematically correct analytical filtering using a given *reference reconstruction filter*. Our analysis is based on evaluating the approximation error in presence of discontinuities in the image, i.e., regions where abrupt changes in intensity are present. We emphasize that the exact sources of these discontinuities are not relevant for the analysis.

Discontinuities in a two-dimensional image can have any shape, but the vast majority are at least locally shaped according to a straight line. Thus we concentrate on the analysis of approximation errors in case of linear discontinuities.

It must be emphasized that in order to evaluate the approximation error in presence of discontinuities, some kind of restrictions on the shape of the discontinuities must be placed. Otherwise the analysis would need to consider an infinite-dimensional space of all possible shapes that discontinuities may have.

Restricting the analysis to discontinuities with specific shapes does not mean that the error metric is unable to estimate the approximation error in presence of irregularly shaped discontinuities. On the contrary, it is reasonable to assume that if a sampling pattern is able to produce good results, i.e. small approximation errors, for all linear discontinuities, is also gives good results for discontinuities with more complicated shapes.

The following analysis is not limited to any particular reference reconstruction filter. We derive the error metric from a geometrical point of view, first for one-dimensional and then for two-dimensional sampling patterns.

4.1. Error Metric in One Dimension

We focus on filtering a family of monochromatic one-dimensional *image functions* $f(x)$ using both the reference filter and a given sampling pattern and measuring the average approximation error. We choose the image functions so that each function represents a single one-dimensional edge, i.e., two constant parts are separated by a sharp discontinuity. The position of the edge is parameterized by a parameter a , giving a family of image functions $f_e(x, a)$. Without loss of generality, we choose values 1 and 0 as the intensities on different sides of the edge, yielding the following definition:

$$f_e(x, a) = \begin{cases} 1 & \text{if } x \leq a \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

We notice immediately that this image function is translationally invariant, satisfying $f_e(x, a) = f_e(x-t, a-t)$ for any t . Therefore, if we assume that all positions a for the edge are equally probable, we may fix the position where the approximation error is evaluated. Without loss of generality,




Figure 4: Top: sampling points (p_i, w_i) of a one-dimensional sampling pattern with six samples optimized to approximate the $(\frac{1}{3}, \frac{1}{3})$ Mitchell-Netravali filter. Bottom: the resulting $\hat{g}_0(a)$ function (solid line) and the reference function $g_0(a)$ (dashed line). Function $\hat{g}_0(a)$ depends only on the sampling pattern, whereas function $g_0(a)$ depends only on the reference reconstruction filter. Note that the optimal sampling points in the top figure are not equidistantly distributed.

ality, the rest of the analysis assumes that the reference filter and the sampling pattern are centered at $x = 0$.

The reference value $g(x)$ is computed by filtering the image function with the reference filter. This is done by convolving the image function with the impulse response $h(x)$ of the reference filter, denoted $g(x) = f(x) \otimes h(x)$. For our image function $f_e(x, a)$, the convolution integral reduces to the integral of $h(x)$ from $x - a$ to positive infinity. We can thus define the reference function $g_0(a)$ for $x = 0$ as

$$g_0(a) = \int_{-a}^{\infty} h(t) dt. \quad (2)$$

When a sampling pattern is used, the reconstructed value $\hat{g}(x)$ is a weighted sum of image functions at the sampling points. Let p_i and w_i denote the position and weight of the i th sampling point, with index i running from 0 to $N - 1$, N being the number of sampling points. For image function $f_e(x, a)$, $\hat{g}(x)$ is simply the sum of weights w_i of the sampling points that satisfy $x + p_i \leq a$. We can thus write $\hat{g}_0(a)$ for $x = 0$ as

$$\hat{g}_0(a) = \sum_{i=0}^{N-1} \begin{cases} w_i & \text{if } p_i \leq a \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Figure 4 illustrates the relationship of a sampling pattern, $g_0(a)$ and $\hat{g}_0(a)$. We are now ready to define the error metric for one-dimensional case. By integrating over the squared difference of $g(x)$ and $\hat{g}(x)$ over all edge positions a , we obtain

$$E_1 = \sqrt{\int_{-\infty}^{\infty} (\hat{g}_0(a) - g_0(a))^2 da}, \quad (4)$$

which is the root sum square error of the reconstructed esti-




Figure 5: a) An edge in two dimensions has two parameters, the angle of the edge θ and the distance from origin a . b) Projecting the sampling points of a two-dimensional sampling pattern on the normal vector \mathbf{n} of the edge results in a one-dimensional sampling pattern.

mate when every possible one-dimensional single-edge image function is considered with equal probability. It can be further noticed that E_1 is finite only if the weights w_i of the sampling pattern sum to one.

This derivation was based on a geometrical interpretation of antialiasing. Alternatively, the same result can be derived from the signal processing point of view by first converting the sampling pattern into a filter with discontinuous impulse response. The value of E_1 is then equal to the root sum square difference between the *unit step responses* of the sampling pattern and the reference filter.

It should be noted that even though the definition of E_1 bears some resemblance to the L_2 discrepancy measure in one dimension, there are fundamental differences between the two. The standard discrepancy measure takes into account only the number of sampling points that fall inside an interval, therefore implicitly assuming box-shaped reconstruction filter and equal weights for all samples. Given that the samples may have distinct weights and that the reference reconstruction filter $g(a)$ is known *a priori*, E_1 metric provides more accurate estimate of the approximation error in the filtering context.

4.2. Error Metric in Two Dimensions

In two dimensions, an edge can be both oriented and positioned arbitrarily. We denote the angle of the edge by θ and the distance of the edge from the center of the coordinate system by a , as illustrated in Figure 5a.

The two-dimensional error analysis can be reduced to one dimension by investigating one angle at a time. In this method, the sampling points of a pattern are projected to the normal vector \mathbf{n} corresponding to θ , as illustrated in Figure 5b. We denote the error of a projected pattern in angle θ by $E_1(\theta)$.

The projection to one dimension does not affect the reconstructed value produced by a sampling pattern, since the

sampling points do not move from one side of the edge to another. Considering the filtering, the projection is equivalent to applying the reference reconstruction filter in the direction of the normal, which is the most reasonable choice when given a one-dimensional reference reconstruction filter such as the Mitchell-Netravali filter.

There is an important additional point in creating an error metric for two-dimensional antialiasing. It has been known for a long time that the visible jaggedness of a non-antialiased edge depends on the angle of the edge. Naiman [Nai98] conducted a psychovisual study, measuring the distances from which the jaggedness became apparent for edges with various slopes. We utilize this data to weight the one-dimensional error metric according to the angle-dependent acuity factors. This is essential for evaluating the perceptual quality of a pattern instead of purely theoretical quality.

Additionally, our tests indicate that evaluating the total sum of E_1 over all angles is not a good metric for the two-dimensional case, even though it seems to be the obvious choice. Considering only the sum of the one-dimensional errors neglects the presence of badly antialiased angles that are highly noticeable in test images and animations. As a solution, we measure the maximum weighted one-dimensional error over all angles instead. Denoting the acuity factor for an angle θ as $W(\theta)$ [Nai98], we can write the two-dimensional error metric E_2 as

$$E_2 = \max[W(\theta)E_1(\theta)], \quad (5)$$

where the maximum is taken over all angles θ . In practice, we evaluate $E_1(\theta)$ at a discrete set of angles for which the acuity factors $W(\theta)$ are given and take the maximum.

It can be argued that if it is better to take the maximum over all $E_1(\theta)$ in two dimensions, it might also be better to also define E_1 in terms of maximum difference between g_0 and \hat{g}_0 , instead of the integral in Equation 4. According to our tests, this option is not reasonable, since the maximum difference between g_0 and \hat{g}_0 over all a and θ tells far too little about the overall antialiasing quality of a pattern. This is because the single worst-case (a, θ) pair is encountered very rarely when rendering an image. However, taking the maximum in E_2 is reasonable, since when e.g. drawing a circle, we are bound to hit the worst-case θ at some point.

Until now we have assumed that the sampling pattern has only a single reconstruction kernel. If there are multiple reconstruction kernels in the pattern, we perform the analysis for every distinct reconstruction kernel and take the maximum E_2 of all kernels.

5. Optimizing the Quality of a Pattern

Quality estimation of sampling patterns is valuable in itself, but in order to find good patterns an optimization method is

needed. In this section we first discuss the problems in minimizing E_2 and then derive a heuristic optimization algorithm that performs significantly better than standard approaches.

Because E_2 is defined as the maximum of weighted E_1 errors (Equation 5), it becomes surprisingly hard a function to minimize. Also, when we have multiple reconstruction kernels in $P(s, r, n)$ with $s > 1$, we define E_2 as the maximum of E_2 over every reconstruction kernel. As a consequence, we end up with a formidable number of distinct $W(\theta)E_1(\theta)$ terms, whose maximum defines the resulting E_2 .

We first tried to apply the standard *simulated annealing* algorithm in the minimization. Unfortunately, it was found to practically stop converging after an initial phase of rapid convergence. This observation was quite puzzling, but it can be explained by the following reasoning.

After the optimization has reached a point where multiple $W(\theta)E_1(\theta)$ terms are very close to the maximum, any random mutation to the sampling point positions and weights is most likely to hoist at least one of them above the current maximum, no matter how small the mutation is. Thus, in order to decrease E_2 in this predominant situation, the modification to the sampling points must be carefully determined, so that *all* terms that are close to the maximum are decreased simultaneously.

The sensitivity to the direction of mutation renders optimization strategies based on random mutations unsuitable for minimizing E_2 . Therefore, we developed a heuristic algorithm based on the method of *steepest descent*. Since steepest descent is a local optimization method, the optimization must be performed for many initial configurations in order to find a good minimum.

It must be emphasized that there is no guarantee that the global minimum is found by employing the optimization algorithm. Particularly with sampling patterns with more than a few samples, it may be quite improbable to hit an initial configuration that converges to the global optimum. All optimized E_2 values in the results section are obtained empirically and should therefore be regarded as (relatively tight) upper bounds for the minimum E_2 .

In the next sections, we first derive the basic tools for the optimization in one dimension, and then describe the two-dimensional optimization algorithm.

5.1. Optimization in One Dimension

We begin by describing an iterative method that mutates a given one-dimensional *seed pattern* into a locally optimal sampling pattern according to Equation 4. A locally optimal sampling pattern has the property that no small perturbation of either the positions or weights of the sampling points results in a smaller error.

We treat E_1 (Equation 4) as a function of positions p_i and weights w_i of the sampling points, and our aim is to find such

positions and weights that E_1 is minimized. We exploit the partial derivatives of E_1 with respect to p_i and w_i , which can be determined analytically. This enables us to modify the positions and weights of the sampling points always towards the direction of decreasing E_1 .

The partial derivative of the error with respect to the position p_i of a sampling point is

$$\frac{\partial E_1}{\partial p_i} = \frac{w_i}{2\sqrt{E_1}} \left(2g_0(p_i) - 2 \sum_{j=0}^{i-1} w_j - w_i \right), \quad (6)$$

which assumes that the one-dimensional sampling points are sorted according to ascending p , making the summation cover all sampling points with $p < p_i$.

The partial derivative of E_1 with respect to w_i is more involved, as the mutation of one weight w_i affects every other weight as well due to the normalization to $\sum_{i=0}^{N-1} w_i = 1$. The normalization ensures that $\hat{g}_0(a)$ (Equation 3) reaches value 1 at infinity, thus making E_1 finite. After somewhat laborious derivation, we get

$$\frac{\partial E_1}{\partial w_i} = -\frac{1}{\sqrt{E_1}} \left(\sum_{j=0}^{N-2} \sum_{k=0}^j w_k H_j - \sum_{j=i}^{N-2} (w_i + 1) H_j \right), \quad (7)$$

where $H_j = (p_{j+1} - p_j) \sum_{k=0}^j w_k - \int_{p_j}^{p_{j+1}} g_0(a) da$,

where it is again assumed that the sampling points are sorted according to ascending p .

The one-dimensional optimization step is simple variant of steepest descent, replacing the line search for optimal step size with a constant step size. First, the positions are updated by setting $p'_i = p_i - c(\partial E_1 / \partial p_i)$, where c is a positive constant that defines the size of the step taken at each round of iteration. Then the weights are updated, $w'_i = w_i - c(\partial E_1 / \partial w_i)$, and normalized.

To reach a minimal E_1 in one-dimensional optimization, we would continue the iteration until the improvement in E_1 became sufficiently small. Figure 4 shows an example of an optimized one-dimensional pattern.

5.2. Optimization in Two Dimensions

To minimize Equation 5, we must consider $W(\theta)E_1(\theta)$ separately for each discrete θ . When a two-dimensional sampling pattern is projected into one dimension as explained in Section 4.2, the one-dimensional partial derivatives of sample positions (Equation 6) describe the magnitude of gradient along the direction of the normal vector corresponding to θ . Since E_2 is the maximum of $W(\theta)E_1(\theta)$, we must concentrate on the terms that are most likely to contribute to the maximum.

Performing the minimization only at the single $W(\theta)E_1(\theta)$ term that defines the current E_2 usually increases the error in other angles, causing the iteration to

oscillate instead of converging. Therefore we introduce an error threshold t and two step sizes c_0 and c_1 , and simultaneously optimize at the angles satisfying $W(\theta)E_1(\theta) < tE_2$ with $c = c_0$ and at the rest of the angles with $c = c_1$. Our implementation uses values $t = 0.99$ and $c_1 = 100c_0$, which were found to result in rapid convergence. This heuristic introduces a small update for the relatively small error terms and a large update for the error terms that are close to the maximum. As a result, each step of the two-dimensional iteration mutates the sampling pattern so that the decrease of large error terms is favored at the cost of increase in small error terms.

As mentioned before, we must start the optimization from multiple *seed patterns* to find a good local minimum. The likelihood of converging to a small E_2 can be increased by constructing the seed patterns using known good statistical distributions such as jittered grid.

6. Results and Discussion

The error metric can be used with any reconstruction filter, and the sampling patterns acquired by optimization depend on the selected filter. The results in this section were measured by using the $(\frac{1}{3}, \frac{1}{3})$ Mitchell-Netravali filter [MN88] as the reference reconstruction filter. It is acknowledged that better filters exist for suppressing aliasing patterns more efficiently in pathological cases, such as the lower left corner of Figure 8a. However, such filters tend to produce excessively blurry images, and thus we favored higher overall quality at the expense of some postaliasing artifacts.

We first present antialiased test images that illustrate the quality of optimized patterns in different situations. Next, we investigate how the image quality is affected when the radius of the reconstruction filter is limited below its true spatial extent. Then, we evaluate the quality of selected low-cost sampling patterns and high-quality sample distributions according to our metric, and compare these to the best sampling patterns found in the respective pattern families. Finally, we present a number of good low-cost patterns in different pattern families, and examine how the selection of a pattern family and sample count affects the error.

6.1. Linear and Curved Discontinuities

Since the proposed error metric is based on the analysis of the approximation error with linear discontinuities, the patterns that are optimized to minimize the error can be expected to perform well in these situations. A comparison between Poisson disk patterns and optimized patterns is presented in Figure 6, where the test scene consists of an infinite checkerboard.

Although the proposed error metric only estimates the pattern's ability to approximate linear discontinuities, it reasonably well estimates the antialiasing quality for discontinu-

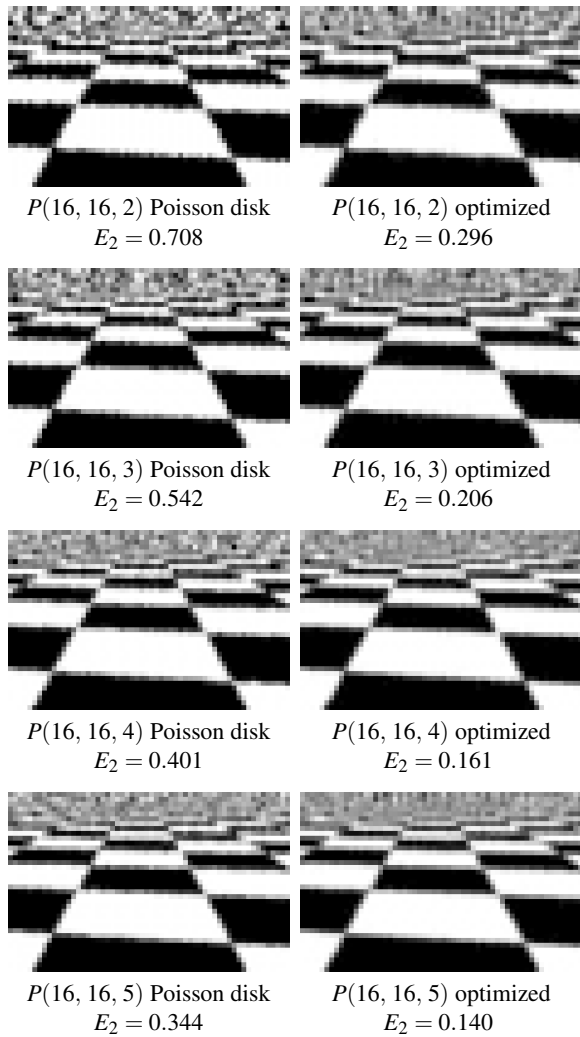


Figure 6: Comparison of randomly generated Poisson disk sample distributions and optimized patterns. All patterns are in $P(16, 16, n)$ family, corresponding to a 4×4 pixel repeating pattern and 4×4 pixel reconstruction kernels. The number of samples taken per pixel ranges from 2 to 5.

ities with more complicated shapes. Figure 7 shows the results of antialiasing a test image containing curved discontinuities of various sizes. It can be seen that even though curved discontinuities are not explicitly taken into account by the error metric, the optimized pattern performs somewhat better than a statistically distributed pattern.

6.2. Support of Reconstruction Filter

The radius of the reconstruction filter limits the extent to which a pattern can approximate the reference filter. As Mitchell-Netravali filter has a radius of 2, samples from a

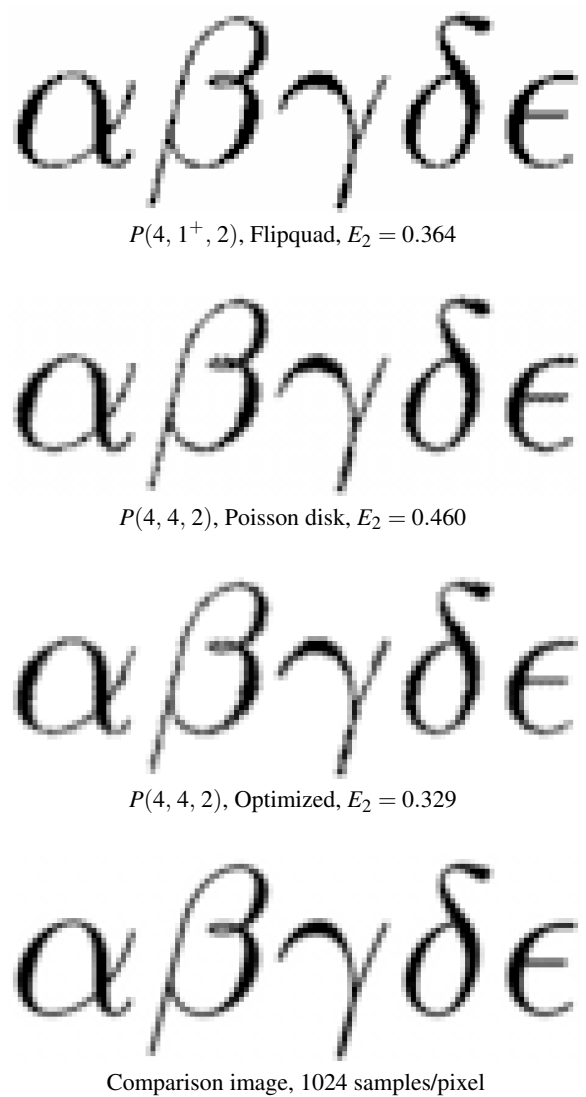


Figure 7: Test image with curved discontinuities of various sizes. All images except the comparison image have 2 samples/pixel. Flipquad suffers from its ability to produce only 3 distinct intermediate shades, and Poisson disk pattern exhibits slightly more noise than the optimized one.

4×4 pixels area are theoretically needed in reconstruction. Especially in real-time (hardware) rendering, it is favorable to keep the reconstruction filter as small as possible, thus reducing the number of pixels that are accessed during the filtering. Our results show that the $(\frac{1}{3}, \frac{1}{3})$ Mitchell-Netravali filter can be satisfactorily approximated by using the samples from only a 2×2 pixel area, corresponding to pattern families $P(s, 4, n)$.

Figure 8 illustrates the effect of a limited reconstruction filter radius with dense sampling patterns. It also shows the




Figure 8: Measured reconstruction errors for dense sampling patterns having different reconstruction filter sizes and sample weighting schemes. All images are sampled with 1024 samples/pixel in a jittered grid. The root mean square error (RMSE) of the approximated images are computed using brightness scale from 0 to 1. In the difference images, white corresponds to no error and black corresponds to 10% absolute error. From left to right: a) 100×100 pixel reference image filtered with full 4×4 reconstruction filter and Mitchell-Netravali weights; b) error with 1×1 truncated Mitchell-Netravali filter; c) as in (b) but with our optimized sample weights; d) error with 2×2 truncated Mitchell-Netravali filter; e) as in (d) but with optimized sample weights.

Pattern	Family	E_2	Best E_2
2×2 uniform grid	$P(1, 1, 4)$	0.698	0.359
rotated grid	$P(1, 1, 4)$	0.439	0.359
Quincunx	$P(1, 1^+, 2)$	0.518	0.359
Flipquad	$P(4, 1^+, 2)$	0.364	0.359
Fliptri	$P(4, 1^+, 1.25)$	0.636	0.560

Table 2: E_2 for common antialiasing patterns and for the best pattern found in the same family. The $(\frac{1}{3}, \frac{1}{3})$ Mitchell-Netravali filter was used as the reference reconstruction filter.

decrease in error when the weights are optimized instead of taken directly from the impulse response of the Mitchell-Netravali filter. Taking the weights from the impulse response results in truncation of the reference filter, and leads to a surprisingly bad approximation, especially with a 1×1 pixel reconstruction filter. With a 2×2 pixel reconstruction filter the error is considerably smaller, and acceptable in practice. Figure 9 includes a comparison of optimized sampling patterns having 4×4 and 2×2 pixel reconstruction kernels.

6.3. Comparison of Common Antialiasing Patterns

We now evaluate the quality of some common low-cost antialiasing patterns according to our metric. Table 2 gives the error E_2 for a number of antialiasing patterns. For reference, the error of the best pattern we found in the same family is shown. The uniform grid and the Quincunx pattern perform rather poorly compared to the best optimized patterns. On the contrary, the Flipquad pattern is very good, the best pattern having similar structure with slightly shifted sample positions.

Pattern	Min	Avg	Max	Opt W	Opt W+P
jittered grid	0.423	0.622	0.907	0.242	0.169
Poisson disk	0.346	0.473	0.782	0.221	0.161

Table 3: Comparison of high-quality sample distributions with 4 samples/pixel. Columns Min, Avg and Max show the minimum, average and maximum E_2 for 1000 randomly generated patterns. Opt W is the resulting E_2 when the sample weights of the best generated pattern are optimized. Opt W+P shows E_2 when both the weights and the positions are optimized.

We also consider two high-quality sample distributions, jittered grid and Poisson disk. In these cases, the samples are spread according to the desired distribution and the filtering is performed by weighting the samples with the impulse response of the reference reconstruction filter. In the measurements we have used patterns from the $P(16, 16, 4)$ family, resulting in a pattern that loops in 4×4 pixel tiles with 4 samples/pixel. A 4×4 pixel area of samples is used for reconstruction, yielding a total of 64 sampling points per output pixel.

Table 3 shows the minimum, average and maximum errors for 1000 patterns generated according to the distributions. The patterns with smallest E_2 were then optimized, first by optimizing only the sample weights and then by optimizing both the positions and the weights. It can be seen that the initially best Poisson disk pattern converged to slightly smaller E_2 after optimization. This confirms that a different local optimum was reached from the two initial configurations, but the difference in final E_2 values (0.008) is quite small compared to the difference in initial E_2 values (0.077).




Figure 9: a) Close-up of a test image antialiased with a Poisson disk pattern having Mitchell-Netravali weights and 4×4 pixel reconstruction filter. b) Optimized sample positions and weights. c) Optimized sample positions and weights, and a 2×2 pixel reconstruction filter. All images were generated using 4 samples/pixel.

Therefore, the optimization is not exceedingly sensitive to the quality of the initial seed pattern.

The values of E_2 seem high when compared to the low-cost patterns, and indeed the discrepancy metrics that disregard psychovisual acuities would rank Poisson disk and jittered grid much higher. The Poisson disk and jittered grid give fairly equal *unweighted* errors $E_1(\theta)$ for all angles, whereas low-cost patterns such as Flipquad and rotated grid concentrate on the perceptually most important horizontal, vertical and diagonal features.

It is also somewhat surprising how much the optimization of the sample weights decreases the error. The reason is that choosing the weight of a sampling point is not a purely local decision but should depend on the positions and weights of the other sampling points. Therefore the traditional weighting of simply evaluating the impulse response of the reference filter at the sampling point results in a greater error than the use of optimized weights.

The rightmost column in Table 3 shows the E_2 when both the positions and the weights are optimized. It can be seen that using a Poisson disk or jittered distribution as a seed pattern gives almost equally good results. Test images generated using a Poisson disk distribution and two optimized patterns are shown in Figure 9.

6.4. Good Low-cost Patterns

Figure 10 shows the smallest E_2 found for different pattern families with respect to the number of samples/pixel. The patterns were generated by optimizing hundreds of random seed patterns. The corresponding patterns are shown in Figure 11.

7. Conclusions and Future Work

We have presented a new error metric for evaluating the antialiasing quality of sampling patterns. With dozens or




Figure 10: Error E_2 of the best patterns found in different pattern families with various number of samples/pixel. In $P(4, 1^+, n)$ family the errors are evaluated for 0.25 increments in n .




Figure 11: The best patterns found for different pattern families, the size of a sampling point corresponding to its weight. From top to bottom: $P(1, 1, n)$, $P(4, 1^+, n)$, $P(1, 4, n)$, $P(4, 4, n)$. Columns from left to right correspond to 2 – 6 samples/pixel. Only one of the four reconstruction kernels is shown for each $P(4, 4, n)$ pattern in the bottom row. The numerical values of the sample positions and weights are available at <http://www.tml.tkk.fi/~samuli/patterns>.

hundreds of samples per pixel, even image quality connoisseurs have great difficulty in observing differences between the sampling patterns. In contrast, the choice of antialiasing pattern is crucial when only a small number of samples can be afforded per pixel. This is the typical situation with graphics hardware architectures and other real-time rendering systems. Unlike earlier methods, our metric is applicable to the analysis of such low-cost patterns. We have also de-

scribed a method for constructing patterns that are almost optimal with respect to the new error metric that takes the reconstruction filter and psychovisual acuities into account.

Increasing the size of the reconstruction filter from 1×1 to 2×2 pixels enables the sampling pattern to more accurately approximate a wider reference filter. This is a possible way of improving the image quality in hardware rendering systems, requiring no extra memory and only limited additional complexity to the filtering hardware.

Currently we use L_2 norm the one-dimensional error E_1 and L_∞ norm for the two-dimensional error E_2 . While our tests suggest that these two error metrics should use a different norm, it is not obvious whether the current choices are the best ones. The L_∞ norm gives perfectly plausible results in most cases, but some extremely low cost patterns such as Fliptri currently rank higher than they probably should [HAML05]. This is because only the maximum angular error is taken into account, and the high average error is disregarded. This could perhaps be solved by employing L_p norm with $2 < p < \infty$ in E_2 . Furthermore, it would be interesting to derive the best sampling patterns for other reference reconstruction filters as well.

Acknowledgements

We would like to thank Ville Miettinen, Tomas Akenine-Möller, Lauri Savioja, Janne Kontkanen, and Jaakko Lehtinen for their helpful comments. This work was partially supported by grant #108222 from the Academy of Finland, and the National Technology Agency of Finland, Anima Vitae, Bitboys, Hybrid Graphics, Nokia, and Remedy Entertainment.

References

- [AM03] AKENINE-MÖLLER T.: *An Extremely Inexpensive Multisampling Scheme*. Tech. rep., Chalmers University of Technology, August 2003.
- [AMS03] AKENINE-MÖLLER T., STRÖM J.: Graphics for the Masses: A Hardware Rasterization Architecture for Mobile Phones. *ACM Transactions on Graphics*, 22, 3 (2003), 801–808.
- [BTB91] BOUVILLE C., TELLIER P., BOUATOUCH K.: Low sampling densities using a psychovisual approach. In *Eurographics 91* (1991), pp. 167–182.
- [CA87] CHEN D. S., ALLEBACH J. P.: Analysis of error in reconstruction of two-dimensional signals from irregularly spaced samples. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 35, 2 (1987), 173–180.
- [Coo86] COOK R. L.: Stochastic sampling in computer graphics. *ACM Transactions on Graphics* 5, 1 (1986), 51–72.
- [CPC84] COOK R. L., PORTER T., CARPENTER L.: Distributed ray tracing. In *Computer Graphics (Proceedings of ACM SIGGRAPH 84)* (1984), vol. 18, ACM, pp. 137–145.
- [Cro77] CROW F. C.: The aliasing problem in computer-generated shaded images. *Communications of the ACM* 20, 11 (1977), 799–805.
- [DEM96] DOBKIN D. P., EPPSTEIN D., MITCHELL D. P.: Computing the discrepancy with applications to supersampling patterns. *ACM Transactions on Graphics* 15, 4 (1996), 354–376.
- [DM93] DOBKIN D. P., MITCHELL D. P.: Random-edge discrepancy of supersampling patterns. In *Graphics Interface '93* (May 1993), pp. 62–69.
- [DN02] DEERING M., NAEGLE D.: The SAGE graphics architecture. *ACM Transactions on Graphics* 21, 3 (2002), 683–692.
- [DW85] DIPPÉ, M. A. Z., WOLD E. H.: Antialiasing through stochastic sampling. In *Computer Graphics (Proceedings of ACM SIGGRAPH 85)* (1985), vol. 19, ACM, pp. 69–78.
- [Gla95] GLASSNER A. S.: *Principles of Digital Image Synthesis*. Morgan Kaufmann, 1995.
- [GW00] GOSS M. E., WU K.: *Study of Supersampling Methods for Computer Graphics Hardware Antialiasing*. Tech. rep., HP Labs, 12 2000.
- [Hal70] HALTON J. H.: A retrospective and prospective survey of the monte carlo method. *SIAM Review* 12 (1970), 1–63.
- [HAML05] HASSELGREN J., AKENINE-MÖLLER T., LAINE S.: A Family of Inexpensive Sampling Schemes. *to appear in Computer Graphics Forum* (2005).
- [HDK01] HILLER S., DEUSSEN O., KELLER A.: Tiled blue noise samples. In *Proceedings of VISION, MODELING, AND VISUALIZATION 2001* (2001), IOS Press, pp. 265–271.
- [Kel01] KELLER A.: *Strictly Deterministic Sampling Methods in Computer Graphics*. Tech. rep., Mental Images, 2001.
- [KH01] KELLER A., HEIDRICH W.: Interleaved sampling. In *Proceedings of the Eurographics Workshop on Rendering* (2001), Springer Verlag, pp. 269–276.
- [KK02] KOLLIG T., KELLER A.: Efficient multidimensional sampling. *Computer Graphics Forum* 21, 3 (2002), 557–563.
- [Mat99] MATOUSEK J.: *Geometric Discrepancy – An Illustrated Guide*, vol. 18 of *Algorithms and Combinatorics*. Springer, 1999.
- [MF92] MCCOOL M., FIUME E.: Hierarchical poisson disk sampling distributions. In *Proceedings of the con-*

- ference on Graphics interface '92 (1992), Morgan Kaufmann, pp. 94–105.
- [Mit87] MITCHELL D. P.: Generating antialiased images at low sampling densities. In *Computer Graphics (Proceedings of ACM SIGGRAPH 87)* (1987), ACM, pp. 65–72.
- [MN88] MITCHELL D. P., NETRAVALI A. N.: Reconstruction filters in computer-graphics. In *Computer Graphics (Proceedings of ACM SIGGRAPH 88)* (1988), ACM, pp. 221–228.
- [Mol91] MOLNAR S.: *Efficient Supersampling Antialiasing for High-Performance Architectures*. Tech. rep., Dept. of Computer Science, University of North Carolina at Chapel Hill, April 1991.
- [Nai98] NAIMAN A. C.: Jagged edges: when is filtering needed? *ACM Transactions on Graphics* 17, 4 (1998), 238–258.
- [NVI01] NVIDIA: *HRAA: High-Resolution Antialiasing Through Multisampling*. Tech. rep., 2001.
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics* 23, 3 (2004), 488–495.
- [Pav90] PAVLIDIS T.: Comments on "Stochastic Sampling in Computer Graphics". *ACM Transactions on Graphics* 9, 2 (1990), 233–236.
- [Shi91] SHIRLEY P.: Discrepancy as a quality measure for sampling distributions. In *Eurographics 91* (1991), pp. 183–193.
- [Yel83] YELLOTT J. I.: Spectral consequences of photoreceptor sampling in the rhesus retina. *Science* 221, 4608 (1983), 382–385.