SIBGRAPI'91

IV Simpósio Brasileiro de Computação Gráfica e Processamento de Imagens

PYRAMID-BASED DITHERING

Luiz Velho
(IMPA/CNPq)

Página em branco na versão original impressa.

Pyramid-based Dithering

Luiz Velho

IMPA - Instituto de Matemática Pura e Aplicada
Rio de Janeiro, RJ

University of Toronto
Department of Computer Science
Toronto, Ontario
lvelho@vis.toronto.edu

Abstract

This paper introduces a new halftoning algorithm that combines the probabilistic and deterministic approaches to the problem. The proposed method uses an image pyramid representation in order to achieve a global and local minimization of the quantization error.

Introduction

The representation of continuous-tone images on graphics display devices with a discrete number of gray levels necessarily implies on the quantization of the pixel intensities to the nearest displayable gray value. When the intensity resolution of the device is compatible with the tonal range of the image, the quantization error is usually negligible. In this case. the discretization process amounts to a simple truncation of the intensity values. On the other hand, when the graphics device has a restricted intensity resolution, special processing is often necessary in order to alleviate the quantization artifacts. The worst situation occurs on bilevel devices that can display only two intensities (e.g. a pixel has a binary state. "on" or "off"). This extreme case is by far the most important, because of the widespread use of bilevel graphic displays.

The process of generating a representation of a gray scale image adapted to a particular graphics device is called *digital halftoning*, in analogy with its counterpart in the traditional printing industry. The digital halftoning process consists of several steps, that generally include the following operations: *tone scale adjustment*, *sharpening* and *dithering* [Uli87].

Dithering is the central part of the halftoning process. It basically performs the discretization step in such a way that the displayed image is a "good approximation" of the original image. A formal statement of the problem is as follows: Given a continuous-tone image represented by an $m \times n$ array A of reals, $A_{ij} \in [0,1]$, generate a binary image represented by an $m \times n$ array B of integers, $B_{ij} \in \{0,1\}$ such that the two images are as close as possible in the perceptual sense.

All dithering algorithms use essentially the same strategy. They determine the state of the binary image elements by comparing the intensities of the original image elements with carefully chosen threshold values. If their intensity is less than the threshold, the output pixel is set to 0, otherwise it is set to 1.

The key to understanding dithering is the notion of "perceptual similarity", that establishes the correct framework for the problem. In this context, the dithering method explores the integration properties of human vision to accomplish its goals. More specifically, since the eye integrates all luminous stimuli within a solid angle of about 2 degrees, the quantization error for a given pixel can be distributed to its neighbors in a way that the local perceptual error is minimized. In other words, if the error in one pixel compensates the errors in nearby pixels then the mean quantization error should be close to zero.

There are two basic approaches to distributing the quantization error over the image. One is probabilistic and the other is deterministic. In the probabilistic approach the average error is made statistically very small, while in the deterministic approach the error is calculated and minimized explicitly.

The random perturbation, the ordered dither and the colored noise inversion algorithms are examples of probabilistic dithering methods. The random perturbation algorithm [Hal89] uses as the threshold a random variable with uniform probability distribution. This is equivalent to modulating the image intensity with white noise. Although this method is not really practical for bilevel quantization, it is very effective in eliminating contour effects for quantization to more levels. The ordered dither algorithm [Lim69] tiles the image with a table of pseudo-random thresholds. The threshold values are determined to avoid the introduction of low frequency noise. This method produces periodic patterns that impart a "computery" look to the picture. The colored noise inversion [GR90] employs a neural network to generate an image with desirable statistical characteristics. This method tends to trade tonal for spatial resolution.

The Floyd-Steinberg, the Witten-Neal and the dot diffusion algorithms are examples of deterministic dithering methods. The Floyd-Steinberg [FS75] algorithm computes the quantization error for each pixel and propagates it to its neighbors to the right and below.

This method often produces ghosts because the error is not evenly distributed. The Witten-Neal [WN82] algorithm propagates the error along a Peano curve, avoiding the directional artifacts of the previous method. The dot diffusion algorithm [Knu87] employs a table of pixel classes that determines the order in which the quantization error is propagated.

Pyramid Dithering

We propose a new dithering algorithm that combines the probabilistic and deterministic approaches to the problem. Our method is based on a hierarchical representation, known as an *image pyramid*. This type of representation is very common in image processing and graphics applications [TP75] [Wil83]. The image pyramid is a multi-resolution representation that allows the data to be manipulated at different levels of detail. The pyramid concept can be implemented either by a collection of variable size image arrays or by a quad-tree data structure [Sam90]. At the bottom level the image is stored in its full resolution and normally each intermediate level has half of the resolution of the previous one.

The basic strategy of the pyramid dithering method is to assign the total image intensity in a hierarchical fashion, distributing the quantization residuals at each level proportionally to the local error. More specifically, consider an image pyramid of depth K where each element $P_{i,j}$ at level k is the sum of the four corresponding elements of the level $k-1$ immediately below:

$$P^k_{i,j} = \sum_{i'=2i}^{2i+1} \sum_{j'=2j}^{2j+1} P^{k-1}_{i',j'}$$

We call P^k the parent node and $P^{k-1}_{i',j'}$ the children nodes of P^k.

Given, without loss of generality, a gray scale image of size $2^{K+1} \times 2^{K+1}$, its pyramid will have at the top level $k = K$ only one element that represents the total sum of the intensities of its elements. If we assume that the non-linearities of the bilevel quantization mapping have been taken care of in the tone-scale adjustment operation, we can say that the best binary approximation of the gray-scale image should have a total integer intensity equal to $\lfloor P^K + 0.5 \rfloor$. That is, the two images are *globally* as close as possible. In order to refine this global approximation to local neighborhoods we start at the top level of the pyramid and build a binary approximation of the image by recursively assigning integer intensities from each parent node to its children.

The computation of an individual element is as follows: Let V and W be the real and integer intensities of the current element E. We are given the intensities v_i of the children nodes S_i, $i = 1 \ldots 4$, and wish to determine the intensities w_i such that their sum is exactly W, and each integer intensity w_i is proportional to the real intensity v_i. To do this we divide W in four integer parts proportional to v_i, then we distribute the residual according to the quantization error of each child node. More specifically, we break w_i into two components: d_i and r_i. Let $d_i = \lfloor \frac{v_i}{V} W \rfloor$, and then assign the residual $W - \sum_{i=1}^{4} d_i$, in a round-robin fashion to the variables r_i choosing $r_i = 1$ with relative probability $P(e_i)$. Where $e_i = (\frac{v_i}{V} W) - d_i$ is the quantization error of S_i without the residual distribution. Note that the first component is deterministic and the second component is probabilistic. Note also that as we descend the pyramid towards the base (e.g. closer to the final binary image) the importance of the probabilistic component increases. This is because the relative size of the residual increases as the absolute magnitude of v_i decreases.

The pyramid-based dithering algorithm is shown bellow:

The structure of the main program is:

```
Read gray-scale image.
Build image pyramid.
Perfom pyramid dither.
Write binary image.
```

The pyramid dither procedure is:

```
pyramid_dither(P)
        image_pyramid P;
{
        /* compute integer intensity w of the root */
        P[K].w = floor(P[K].v + 0.5);

        for (each level k = K to 2)
               for (each element E of P[k])
                      /* transmit intensity w from
                         node E to its children S */
                      dither_element(E, S);
}
```

The processing of a node is:

```
dither_element(E, S)
      parent node E;
      children nodes S;
{
      /* compute the 1st part and the error */
      for (each node in S) {
            S.d = floor((S.v / E.v) * E.w);
            S.e = (S.v / E.v * E.w) - S.d;
       }
       /* compute the residual */
       R = E.w - sum_of(S.d);

       /* distribute the residual */
       while (R > 0) {
             for (each node in S) {
                   if (random() < S.e && S.r == 0) {
                         S.r = 1;
                         R -= 1;
                   }
                   if (R == 0) break;
             }
       }
       /* assign the final integer intensity */
       for (each node in S)
             S.w = S.d + S.r;
}
```

Examples

We selected three images to determine the effectiveness of our dithering algorithm. They are examples of the main test cases used to evaluate this kind of application. The first image is a gray-scale ramp. Its purpose is to reveal how the method handles a smooth transition of gray levels. The second image is a computer generated ray-traced picture. It is representative of a synthesized noise free image. The third image is a digitized photograph of a face. Its an example of a real scene captured through a camera.

Figures 1, 2 and 3 show respectively the three images described above dithered with the pyramid-based method. They are printed using a 300 dpi laser printer at the expanded resolution of 75 dpi. This means that each pixel is replicated 4 times in the X and Y directions. The results obtained are very good. The method generates aperiodic patterns without any directional artifacts. It is also very successful in capturing the fine details of the images.

Conclusions

In summary, we presented a new dithering method based on a hierarchical image representation. It employs a pyramidal data structure to recursively distribute the total integer intensity to smaller neighborhoods starting from the entire image down to the pixel level. The resulting binary image is the best global and local approximation of the original gray scale image, in the sense that the sum of intensities over variable size areas differs only by an integer truncation value. Moreover, the integer approximation is statistically proportional to the quantization error.

The method combines the deterministic and probabilistic approaches to the dithering problem in a original way. This allows the explicit minimization of the quantization error and, at the same time, the generation of anisotropic dither patterns. The method can be applied equally well to the general problem of the quantization to N bits.

The structure of the algorithm is suitable to a parallel implementation. since the processing of nodes at a given level of the pyramid depends only on a self-contained set of data. The use of a standard data structure makes possible to take advantage of special purpose hardware.

References

[FS75] R. Floyd and L. Steinberg. An adaptive algorithm for spatial gray scale. *SID Symposium*, 36–37, 1975.

[GR90] R. Geist and R. Reynolds. Colored noise inversion in digital halftoning. *Proceedings of Graphics Interface*, 31–38, 1990.

[Hal89] Roy Hall. *Illumination and Collor in Computer Generated Imagery*. Springer-Verlag, New York, 1989.

[Knu87] Donald Knuth. Digital halftones by dot diffusion. *ACM Transactions on Graphics*. 6(4):245–273, Oct. 1987.

[Lim69] J. Limb. Design of dither waveforms for quantized visual signals. *Bell Systems Technical Journal*, 48(7):2555–2582, 1969.

[Sam90] H. Sammet. *Applications of Spatial Data Structures: Computer Graphics, Image Processing and GIS*. Addison-Wesley, Reading, MA, 1990.

[TP75] S. Tanimoto and T. Pavidlis. A hierarchical data structure for picture processing. *Computer Graphics and Image Processing*, 4(2), 1975.

[Uli87] R. Ulichney. *Digital Halftoning*. MIT Press, Cambridge, Ma. 1987.

[Wil83] Lance Williams. Pyramidal parametrics. *Computer Graphics*. 17(3), 1983.

[WN82] H. Witten and M. Neal. Using peano curves for bilevel display of continuous tone images. *IEEE Computer Graphics and Applications*. 47–52, May 1982.

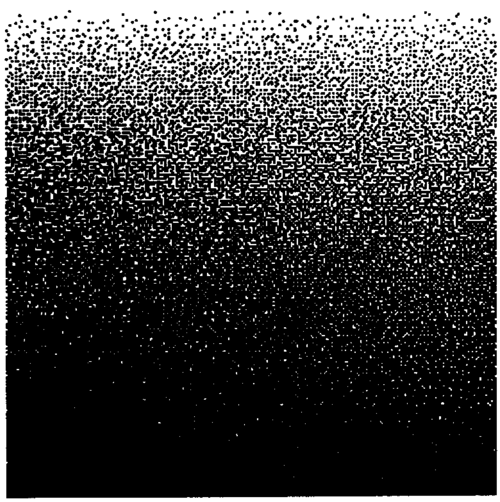

Figure 1: Gray scale ramp

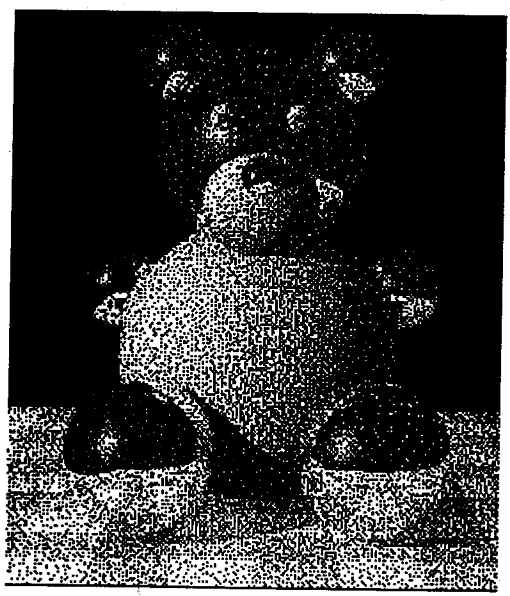

Figure 2: Computer generated image

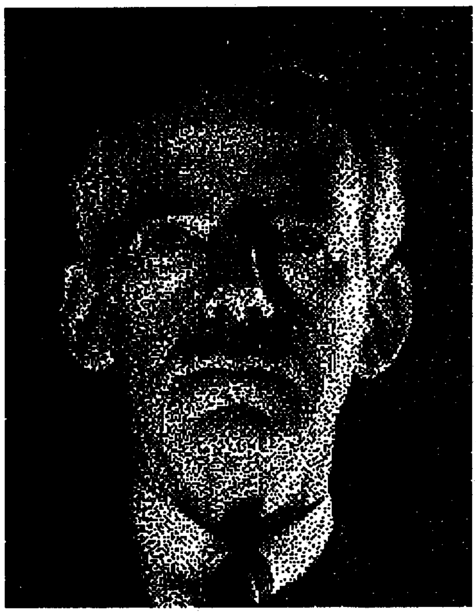

Figure 3: Digitized photograph