

Window-Extent Tradeoffs in Inverse Dithering

Shiufun Cheung & Robert A. Ulichney
Compaq Computer Corporation, Cambridge Research Laboratory
Cambridge, Massachusetts, U.S.A

Abstract

A video-rate color inverse-dithering system based on edge-sensitive low-pass filtering is introduced. The twin requirements of low complexity and low memory necessitate tradeoffs in system design, especially in the extent of the processing window. Our investigation shows that a single-line window is sufficient for achieving good results. We also investigate whether the separate color channels (RGB) should be processed independently. Our study suggests that processing the color channels jointly does not provide significant benefits.

Introduction

Limitations in hardware, video memory, and communication bandwidth can necessitate the reduction of amplitude resolution, or pixel bit depth, in digital images. Dithering is the process by which the illusion of a large number of color or grayscale levels is created using only a small number of actual amplitude levels [1]. Artifacts associated with low bit depth, such as false contours, are thereby mitigated.

Although the dithered version of images is appropriate for storage or transmission purposes, perceptual quality of printed or displayed images can be enhanced by restoring the dithered images to an approximation of their original pixel bit-depth. Such image-restoration algorithms are collectively termed inverse dithering. Previously suggested approaches to the problem of inverse dithering include wavelet decomposition [2], MAP estimation [3], and adaptive algorithms [4,5]. It is worth noting that with the exception of [5] all the approaches were targeted towards the special case of inverse dithering in which bitonal images are restored to full grayscale.

We introduce in this paper a video-rate *color* inverse-dithering system designed for implementation on a high-end graphics accelerator chip for which space is at a premium. The twin requirements of low complexity and low memory inevitably lead to tradeoffs in system design. In the following, emphasis will be placed on the tradeoff issues in one particular aspect of the system—the extent of the

processing window, both in the spatial sense, and in the sense of crossing color channels.

Inverse Dithering

The reason dithering is effective in preserving an illusion of continuous color or grayscale is that visual acuity decreases dramatically at high spatial frequencies, so dither patterns are actually perceived as their local macro averages. Given this known nature of dithering, an algorithm based on adaptive low-pass filtering is considered for the proposed inverse dithering system. This approach is particularly attractive because it is simple and can be easily optimized to achieve low complexity.

Edge-Sensitive Low-Pass Filtering

It is known that the application of a low-pass filter to a dithered image can produce the conflicting effects of removing the graininess caused by the dither and blurring the object edges. The former effect is useful; the latter is undesirable. The blurring effect is especially objectionable with synthetically generated images, such as those most commonly displayed on a computer screen.

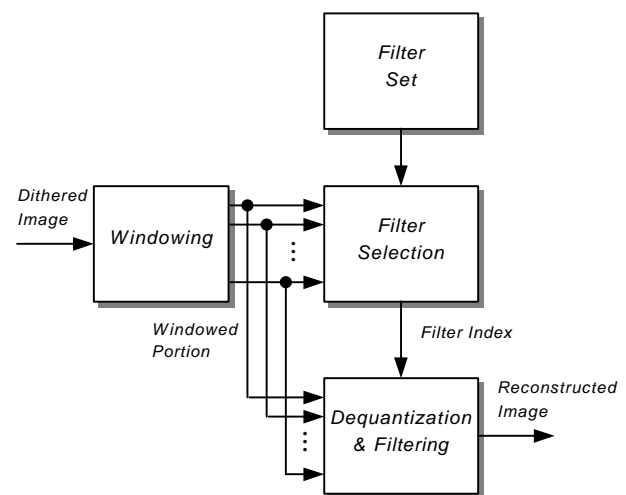


Figure 1: High-Level Block Diagram of Inverse Dithering System.

The proposed system remedies this problem by using an edge-sensitive adaptive mechanism. This approach is illustrated in Figure 1, which is a high-level block diagram of the system. Such adaptive algorithms, by their nature, necessitate pixel-by-pixel processing. For each pixel, processing is limited to a windowed portion of the image, which is typically the area immediately surrounding the pixel. In the filter selection module, edge detection is performed. From the result, a filter is selected with the largest region of support that does not overlap an object edge. The pixels in the window are mapped to the higher bit depth by a process we refer to as dequantization [6], and subsequently filtered with the chosen filter to produce the reconstructed pixel value.

Design of the Filter Set

The inverse dithering system uses a set of pre-defined low-pass filters with finite impulse responses. For the adaptive mechanism to be effective, filters of varying regions of support are necessary. Special filters are also needed for pixels near object edges. In the proposed implementation, filters are restricted to be one of two kinds:

- (a) Horizontally even symmetric filters where

$$h_i(x_0, y_0) = h_i(-x_0, y_0) \text{ for all } (x_0, y_0). \quad (1)$$

- (b) Part of a horizontally asymmetric pair, $h_i(x, y)$ and $h_j(x, y)$, which are defined to be

$$h_i(x_0, y_0) = h_j(-x_0, y_0) \text{ for all } (x_0, y_0). \quad (2)$$

The asymmetric filters are typically designed to be one-sided, meaning that their region of support are either restricted to $x \geq 0$ or $x \leq 0$. In general, the processing of most pixels would call for the even symmetric filters. Asymmetric filters are applied only to pixels near an object edge. For practical implementation purposes, the filters are placed in order of decreasing spatial extent to facilitate the selection process.

Edge Detection in Dithered Images

As mentioned earlier, the filter selection process is based on the detection of object edges within the window. However, the graininess introduced by dithering produces edge-like artifacts that can be easily confused with genuine object edges. Traditional edge detection algorithms will therefore fail.

For this system, methods that exploit our knowledge of dithering are devised as alternatives to traditional approaches. As a first step, the selection criterion is slightly modified to a stricter form: a filter is selected so that it has the largest region of support consistent with a dithered constant-color area.

There can be different stringency levels to the consistency requirement. For example, a basic rule would be to require that all pixel values in the region to be within one

amplitude level of each other, as is expected in a dither pattern. If the dither method is known, it is possible to make the requirement more precise by tabulating all possible patterns that can occur within the region.

Both the aforementioned methods were found to be effective in the inverse-dithering system.

Window-Extent Tradeoffs

The above description of the inverse dithering system, though brief, clearly shows that there is a certain amount of flexibility in its implementation. The filter set, for example, can vary widely from implementation to implementation. Aside from the few aforementioned constraints, the shape and support of the filters are entirely at the discretion of the designer.

It is easy to see how the size of the processing window and the filter sizes are related. The window extent must be large enough to cover the region of support of the largest filter. In the following section, we discuss the tradeoff issues involved in the determination of an appropriate window extent.

Single-Line vs. Multi-line Processing

The window extent is especially important in the vertical direction. Given that a digital image is generally scanned and stored from left to right and from top to bottom in a raster fashion, the memory buffer required for processing pixel values is dramatically reduced if the vertical window extent is shrunk from multiple lines to a single line. In other words, if the system can be limited to one-dimensional processing, its complexity and memory requirements can be significantly lowered. This is illustrated in Figure 2.

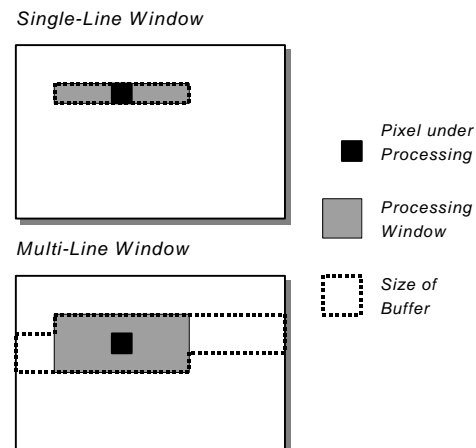


Figure 2: Buffering for single line vs. multiple line processing.

The tradeoffs associated with the use of different window extents were first investigated by performing an informal subjective study. Images were first dithered and

then inverse dithered with filter sets based on processing windows that have heights of 3 lines, 2 lines and 1 line. The filter widths were kept the same in all cases. The images were dithered from the original 24 bits to 16 bits, 12 bits and 8 bits. Test subjects would view the reconstructed images on high-resolution and low-resolution monitors. They were then asked to rate the perceptual quality. The rating scale was calibrated by nominally assigning the original 24-bit image on a high-resolution monitor the highest quality and the dithered 8-bit image on a low-resolution monitor the lowest quality. The results are shown in Figure 3. The horizontal lines in the figure denote the perceptual quality given to the various dithered images.

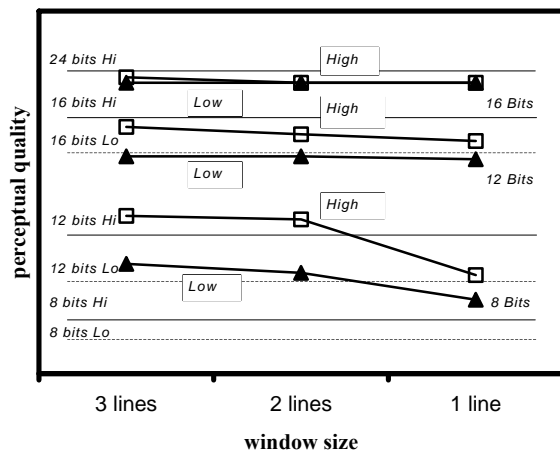


Figure 3: Comparison of perceptual quality of reconstructed images when different numbers of scan lines are used in inverse dithering (informal study).

Although this is not a formal subjective evaluation, the graphs indicate that degradation due to reduction of the window extent to one line is not substantial in most cases. The most objectionable artifact associated with single-line inverse dithering is the appearance of horizontal streaks across constant-color regions. This is because the dither patterns are designed in two dimensions. It is quite possible for consecutive lines in a dither pattern to have significantly different averages. When the inverse-dithering system only use single-line filters, the differences in the line averages show up as horizontal streaks in the reconstructed image. Figure 4 is a diagram showing how this artifact is formed.

One possible way to avoid this artifact is to use single-line or one-dimensional dithering. When the dither patterns are designed in one dimension, the line averages are guaranteed to be the same across all lines. However, for single-line dithering methods to achieve a homogeneous appearance, the one-dimensional dither pattern on each line has to be offset, or phase adjusted, from the lines above and below. Figure 5 is an example of a dither pattern generated by the combination of vertical phase adjustment and single-

line dithering. The choice of the vertical offset/phase vector is an interesting study in itself. More information on this and other aspects of one-dimensional dithering can be found in [7].




Figure 4: Artifact from the use of single-line processing. The dither pattern is a conventional one generated using recursive tessellation for a gray level of 1/8.




Figure 5: Single-Line Dithering. This dither pattern is for a gray level of 1/8. The vertical offset vector is [0 4 2 7 1 5 3 6].

Dependent vs. Independent Color Channels

The other issue under investigation is another aspect of window extent. In this case, we have to decide whether to extend the processing window to include the three color, namely the red, green and blue, channels, or to process the color channels as three independent monochromatic pictures.

The decision mainly affects the edge-detection mechanism in the filter-selection module. If the color channels are processed separately, filter selection for each channel will be independent. Therefore, at each pixel, it is possible for the filter selected for one channel to have a different region of support from the one selected for another channel. If, on the other hand, the color channels are dependent, the same filter will be selected for all three channels. This implies that the filter with the smallest region of support will be chosen, which is akin to saying an edge in one channel implies an edge in all channels.

Table 1 shows the edge-detection results from using both approaches. For three representative images, we enumerate the different kinds of pixels that are associated with object edges. The three kinds of edge pixels are described below.

- a) *Left of edge*. Possible filter region of support extends more than 4 pixels to the left and less than 1 pixel to the right.
- b) *Right of edge*. Possible filter region of support extends more than 4 pixels to the right and less than 1 pixel to the left.
- c) *Straddle on edge*. Possible filter region of support extends less than 1 pixel in both left and right directions.

For the synthetic image, in this case a screen capture of text in which object boundaries are clearly demarcated, the number of edge points, as expected, does not change whether the color channels are processed jointly or separately. However, for natural images, the number of edge points decreases when the color channels are independent. While this appears to be intuitive, more detailed analysis reveals a much more complex scenario, in which the changes in edge points depend largely on the nature of the images.

Table 1: Edge-Detection Results (Dependent vs. Independent Color Channels)

		Screen of Text	Female Portrait	Desert Formations
Dependent Color Channels				
Left of Edge		5.13%	13.20%	4.78%
Right of Edge		5.40%	8.45%	4.00%
Straddle on Edge		9.80%	13.46%	23.84%
Total Edge Pixels		20.34%	35.10%	32.62%
Independent Color Channels				
Red	Left	5.13%	11.75%	6.18%
	Right	5.40%	7.13%	5.20%
	Straddle	9.80%	6.83%	15.70%
	Total	20.34%	25.71%	27.08%
Green	Left	5.13%	10.16%	6.21%
	Right	5.40%	5.83%	5.35%
	Straddle	9.80%	4.71%	14.78%
	Total	20.34%	21.70%	26.34%
Blue	Left	5.13%	10.27%	6.45%
	Right	5.40%	6.04%	5.40%
	Straddle	9.80%	4.85%	14.49%
	Total	20.34%	21.17%	26.34%

However, for inverse dithering, the more important issue is the difference in perceptual quality. Subjective

observations show that extending the window to include all three color channels results in a reconstructed image of increased sharpness but also increased graininess. This is in line with our expectations because the more stringent criterion should lead to more filters with smaller regions of support being selected. This dubious improvement comes at an expense of increased complexity. The finalized system is therefore designed to process the color channels independently.

Conclusion

In this paper, we have introduced a color inverse dithering system based on edge-sensitive adaptive low-pass filtering. Our investigations have shown that using a window that has the height of one line and processes the color channels independently offers adequate perceptual quality.

References

1. R. Ulichney, *Digital Halftoning*, MIT Press, 1987.
2. Z. Xiong, M. T. Orchard and K. Ramchandran, "Wavelet-Based Approach to Inverse Halftoning," *Proc. of IS&T/SPIE Symposium on Electronic Imaging Science and Technology*, 1997.
3. R. L. Stevenson, "Inverse Halftoning via MAP Estimation," *IEEE Transactions on Image Processing*, vol. 5, no. 4, pp. 574—583, April 1997.
4. C. M. Miceli and K. J. Parker, "Inverse Halftoning," *Journal of Electronic Imaging*, vol. 1, no. 2, pp. 143—151, 1992.
5. A. C. Barkans, "Color Recovery: True-Color 8-Bit Interactive Graphics," *IEEE Computer Graphics and Applications*, vol. 17, no. 1, January/February 1997.
6. R. Ulichney and S. Cheung, "Pixel Bit-Depth Increase by Bit Replication," *Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts III*, Proc. SPIE vol. 3300, pp. 232-241, 1998.
7. R. A. Ulichney "One-Dimensional Dithering," *Int. Symposium on Electronic Image Capture and Publishing (EICP'98)*, Zurich, Switzerland, May 1998.

Acknowledgements

The authors would like to acknowledge the contributions of Mr. Giridharan Iyengar and Dr. Robert MacNamara to the development of the inverse dithering algorithm.


Email contact addresses:

Shiufun Cheung: cheung@crl.dec.com


Robert Ulichney: ulichney@crl.dec.com



Figure 6: Artifact from the use of single-line processing. The horizontal streaks are a result of different line averages in a two-dimensionally designed dither pattern. This artifact can be avoided by using single-line dither.




(a) 24-bit Original




(b) Dithered Image



(c) Reconstructed Image



(d)



(e)



(f)

Figure 7: Performance of Inverse Dithering System. The image in (a) was first dithered (each color channel from 8 bits to 3 bits) to form (b) and then reconstructed using the proposed inverse dithering system to form (c). The Images in (d), (e) and (f) are blowups of (a), (b) and (c) respectively. The blowup area is indicated by the blue square in (a).