



## An improved parallel contrast-aware halftoning<sup>\*</sup>

Ling-yue LIU<sup>†1</sup>, Wei CHEN<sup>1</sup>, Tien-tsin WONG<sup>2</sup>, Wen-ting ZHENG<sup>†‡1</sup>, Wei-dong GENG<sup>1</sup>

(<sup>1</sup>State Key Lab of CAD & CG, Zhejiang University, Hangzhou 310027, China)

(<sup>2</sup>Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China)

<sup>†</sup>E-mail: liulingyue@zju.edu.cn; wtzheng@cad.zju.edu.cn

Received May 22, 2013; Revision accepted Oct. 18, 2013; Crosschecked Nov. 18, 2013

**Abstract:** Digital image halftoning is a widely used technique. However, achieving high fidelity tone reproduction and structural preservation with low computational time cost remains a challenging problem. This paper presents a highly parallel algorithm to boost real-time application of serial structure-preserving error diffusion. The contrast-aware halftoning approach is one such technique with superior structure preservation, but it offers only a limited opportunity for graphics processing unit (GPU) acceleration. Our method integrates contrast-aware halftoning into a new parallelizable error-diffusion halftoning framework. To eliminate visually disturbing artifacts resulting from parallelization, we propose a novel multiple quantization model and space-filling curve to maintain tone consistency, blue-noise property, and structure consistency. Our GPU implementation on a commodity personal computer achieves a real-time performance for a moderately sized image. We demonstrate the high quality and performance of the proposed approach with a variety of examples, and provide comparisons with state-of-the-art methods.

**Key words:** Digital image halftoning, Error diffusion, GPU, Quantization, Space-filling curve

**doi:**10.1631/jzus.C1300142

**Document code:** A

**CLC number:** TP391

### 1 Introduction

Halftoning denotes a technique that converts a continuous-tone image into a binary tone image. In the past decade, digital image halftoning has played important roles in various applications such as digital holography, desktop publishing, image compression, and non-photorealistic rendering. Yet, achieving high-quality image halftoning with minimal computational time remains a challenging task.

An ideal halftoning algorithm should be capable of simultaneously reproducing both the tone and the local structures of the input image. In addition, it is desirable that the result should possess the blue-noise property (Ulichney, 1988). Early efforts, such as the

ordered dithering (Bayer, 1973) and error diffusion (Floyd and Steinberg, 1976) approaches, were devoted to maintaining tone consistency. The well-known blue-noise aware halftoning methods (Mitsa and Parker, 1992; Ostromoukhov, 2001) faithfully reproduce grey level ramps without introducing visually significant artifacts, but over-blur fine details.

More sophisticated structure-aware methods (Pang *et al.*, 2008; Chang *et al.*, 2009; Li and Mould, 2010) have enabled recent halftoning algorithms to achieve structure preservation to a great extent. The structure-aware halftoning (SAH) approach (Pang *et al.*, 2008) addresses the deficiencies of previous edge-enhancement halftoning approaches, such as the failure to detect weak edges or improper emphasis of details. It takes an iterative optimization procedure with the consequence of a very low performance. Its corresponding parallel version was proposed by Wu *et al.* (2013). The structure-aware error diffusion (SAED) algorithm presented in Chang *et al.* (2009) improves the performance, but results in low structural similarity. A recent contrast-aware halftoning (CAH)

<sup>‡</sup> Corresponding author

<sup>\*</sup> Project supported by the National Key Technology R&D Program of China (No. 2012BAH35B03), the National High-Tech R&D Program of China (No. 2012AA12090), the National Natural Science Foundation of China (Nos. 61232012 and 81172124), and the Zhejiang Provincial Natural Science Foundation (No. LY13F020002)

© Zhejiang University and Springer-Verlag Berlin Heidelberg 2013

approach (Li and Mould, 2010) employs an adaptive contrast-aware mask to preserve better texture details with faster speed. However, the basic and variant CAH approaches take sequential processing orders, which make them not suitable for being straightforwardly parallelized on GPUs or multi-core CPUs. Also, its blue-noise property is not near-optimal.

With the rapid improvements in the performance of GPU, many researchers have been investigating ways of bringing the computational power of GPU into broader practical applications. Although there do exist block-based error diffusions (Chang and Allebach, 2003; Li and Allebach, 2005), which avoid block-boundary artifacts and are well suited for parallel application, technical conflicts make them incompatible with the state-of-the-art structure-aware methods (Chang *et al.*, 2009; Li and Mould, 2010). To enable GPU implementation of structure-preserving error diffusion, the basic idea is to apply a structure-preserving algorithm to the image blocks without introducing visible blocking artifacts. To achieve this goal, our method uses a suite of novel approaches: (1) a parallelizable error-diffusion halftoning framework that handles the contrast-aware method (Li and Mould, 2010) with a block-wise scheme; (2) a novel quantization technique and an alternative traversal order that eliminate blocking artifacts.

## 2 Related work

There has been much work devoted to digital halftoning. Here we review only that most relevant to our research.

The pioneering halftoning methods handle each pixel individually by means of a screening or dithering process. One of those, the ordered dithering algorithm (Bayer, 1973), can be implemented in parallel, but produces redundant regular patterns that are unacceptable for tone reproduction. Mitsa and Parker (1992) introduced the blue-noise mask to make the results possess blue-noise characteristics. In general, the visual effects generated by these pointwise methods are inferior to those generated by error-diffusion techniques (Ostromoukhov, 2001; Zhou and Fang, 2003).

To obtain high-fidelity reproduction, a class of iterative optimization methods has been studied. For

instance, Analoui and Allebach (1992) used direct binary search (DBS) based optimization to compute a halftoning image with minimum error-metric value against the input image. Likewise, Neuhoff *et al.* (1997) used a 2D least-squares solution to minimize the difference between the perceived continuous-tone and halftone images based on a printer model and a visual perception model. Recently, the mean structure similarity index measurement (MSSIM) (Wang *et al.*, 2004) has been introduced to characterize perceived changes in structural information variation. Subsequently, the structural similarity measure (SSIM) leads to an objective function (Pang *et al.*, 2008) that accounts not only for tone similarity but also for structure similarity. All these iterative methods can generate halftones that are visually pleasing and isotropic, and have few uncorrelated patterns. Their main disadvantage lies in their high computational cost.

Another category of halftoning methods employs error diffusion processes. The well-known Floyd-Steinberg error diffusion process (Floyd and Steinberg, 1976) significantly outperforms ordered dithering (Bayer, 1973), but still suffers from heavy artifacts. Later, it was found that an ideal halftone image has a blue-noise spectrum (Ulichney, 1988), which drove research attention to spectral analysis. Ostromoukhov (2001) addressed several qualitative problems in Floyd-Steinberg error diffusion using an off-line optimization process for optimal diffusion coefficients. Li and Allebach (2004) optimized not only the diffusion coefficients but also the quantized thresholds derived from a human visual system model. Zhou and Fang (2003) applied a variable threshold modulation to the variable-coefficient error-diffusion algorithm (Ostromoukhov, 2001) to improve quality in the mid-tones. Note that these methods perform well in tone reproduction, but are inefficient in respect of structure preservation. Meanwhile, parallelizing conventional tone-aware halftoning methods like the ones presented in Ostromoukhov (2001) and Zhou and Fang (2003) is very complicated.

To preserve the appearance of textured regions in an image, edge enhancement techniques have been used. Eschbach and Knox (1991) used an image-dependent threshold to modulate the edge enhancement influence under an error diffusion framework. However, edge enhancement halftoning techniques

always suffer from deficiencies in edge detection operators. Chang *et al.* (2009) first introduced structure preservation into the conventional error diffusion scheme to obtain a fast structure-preserving algorithm, but failed to convey the same high-level structural similarity as the structure-aware halftoning approach (Pang *et al.*, 2008). More recently, Li and Mould (2010) proposed contrast-aware weights to diffuse the errors, and obtained a stable prime structure and contrast preservation. Our approach advances the contrast-aware scheme by incorporating it into a block-wise parallelizable pipeline.

In addition, much research has been devoted to performance improvement. Although ordered dithering can be implemented in a parallel fashion, its quality is far from satisfactory. The dot-diffusion method (Knuth, 1987) leverages the Floyd-Steinberg error diffusion algorithm, and retains both the sharpness of error diffusion methods and the parallelism of ordered dithering methods. Another tone-aware parallelism scheme is the block-wise process mode (Chang and Allebach, 2003; Li and Allebach, 2005) in which an image is handled based on small blocks. Its main challenge is the elimination of artifacts such as boundary artifacts and periodic patterns. Li and Allebach (2005) divided an image into two groups of interlaced blocks, to eliminate boundary artifacts by specifying a spiral scanning path, and employed a training framework for tone-dependent error diffusion to reduce self-engendered diagonal artifacts. Chang and Allebach (2003) introduced a reduced-LUT (look-up table) tone-dependent error diffusion algorithm together with a novel scanning path. However, it suffers from some diagonal artifacts. Inspired by parallel Poisson disk sampling (Wei, 2008), our approach takes an improved block-wise error diffusion

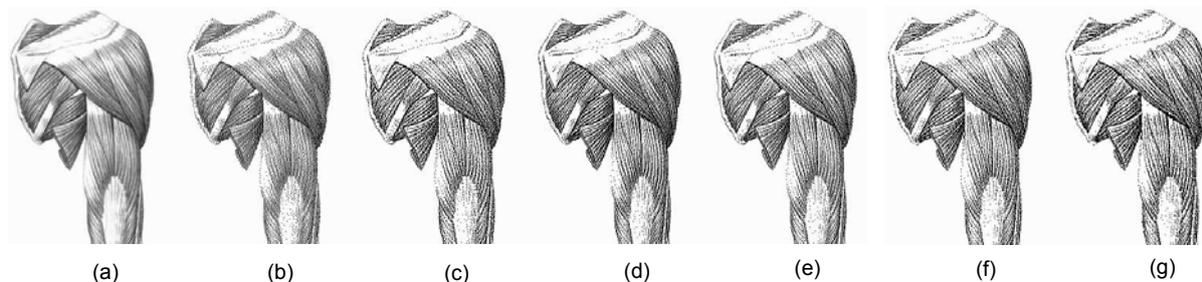
pipeline, and uses an innovative quantization technique to address the artifact problem. Parallel structure-aware halftoning (Wu *et al.*, 2013) is presented to boost the very slow convergence process of SAH, and halftone images can be generated within seconds on GPUs. Our proposed parallel algorithm produces high quality structure-preserving bitonal images in much less time.

Space-filling curves, in particular, have been used for traversing images in digital halftoning. Witten and Neal (1982) traversed image pixels along Peano curves, yielding a very low cumulative error in the resultant halftone images. Velho and Gomes (1991) used space-filling curves to generate aperiodic patterns of evenly distributed dots without directional artifacts. Zhang and Webber (1993) merged dot diffusion with space-filling curves. Wong and Hsu (1995) designed a special scheme that completes the entire process by selective precipitation and adaptive clustering. By contrast with the aforementioned methods, our method incorporates space-filling curves into the block-wise processing pipeline, to eliminate the artifacts that may be caused by traversal paths such as the spiral path.

Our approach compares favorably with existing solutions in terms of quality and performance, as indicated in Fig. 1. More details are given in Section 4.

### 3 Our approach

Achieving high quality and efficient halftoning is essentially a multi-target optimization problem. Each of the previous solutions addresses parts of the problem, yet also is inefficient in certain aspects. Embracing the advantages of these methods with a



**Fig. 1 Result comparison for the Arm image (200×307)**

(a) The input; (b) Ostromoukhov error diffusion (Ostromoukhov, 2001); (c) SAH (Pang *et al.*, 2008); (d) SAED (Chang *et al.*, 2009); (e) The basic CAH (Li and Mould, 2010); (f) The variant CAH (Li and Mould, 2010); (g) Ours

novel mechanism to overcome their shortcomings is not a trivial task.

We take further the newest quality-oriented approach (Li and Mould, 2010) by incorporating it into a new block-wise parallelizable pipeline. Straightforward integration will, however, lead to visible artifacts arising from both elements of the treatment. Consequently, we introduce two techniques to eliminate these artifacts, a new quantization scheme and a new internal traversal order for each block.

The entire pipeline is as illustrated in Fig. 2. First, an input image is subdivided into four groups of blocks, which are processed alternately. The blocks of each group are handled in parallel, and the pixels in a block are traversed along a specific space-filling curve (Section 3.3). Each pixel is quantized with a novel quantization model (Section 3.2). Subsequently, a contrast-aware mask is built for the quantized pixel, and applied to its unprocessed neighbors (Section 3.1). In this way, error diffusion is accomplished for both the internal regions and the boundary of each block.

Our error diffusion is expressed by the following equations:

$$B(x, y) = \begin{cases} 1, & \text{if } I'(x, y) > T_i \text{ and } I'(x, y) > T_j, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

$$E(x, y) = B(x, y) - I'(x, y). \quad (2)$$

Here,  $I'(x, y)$  is the modified intensity of pixel  $(x, y)$ , ranging from 0 to 1, and its quantized value is  $B(x, y)$ .  $T_i$  and  $T_j$  belong to two arbitrary different elements of the set  $\{t, T(x, y), G(x, y)\}$ .  $t$  represents the constant threshold 0.5.  $T(x, y) = \text{BN}(x \bmod M, y \bmod N)$  is the threshold of the blue-noise mask for pixel  $(x, y)$ . BN denotes the conventional blue noise mask (Mitsa and Parker, 1992) at the size of  $M \times N$ , and  $G(x, y)$  is the Gaussian weighted average of the original pixels around  $(x, y)$ .

In the following, we highlight the integration of the contrast-aware mask and the block-wise process mode. Then, we describe this new quantization model that can be used to eliminate the block-boundary artifacts (Chang and Allebach, 2003) occurring in most conventional block-wise processing methods. Next, we show how to avoid tonal artifacts using a space-filling curve guided traversal path.

### 3.1 Block-wise error diffusion with the contrast-aware mask

Traditional error diffusion approaches traverse the input image with a sequential path such as a serpentine path or raster scan path. Nevertheless, the key component of error diffusion is the distribution of the quantization error to the neighboring unprocessed pixels in accordance with the principle of the spatial integration of the human vision system (HVS).

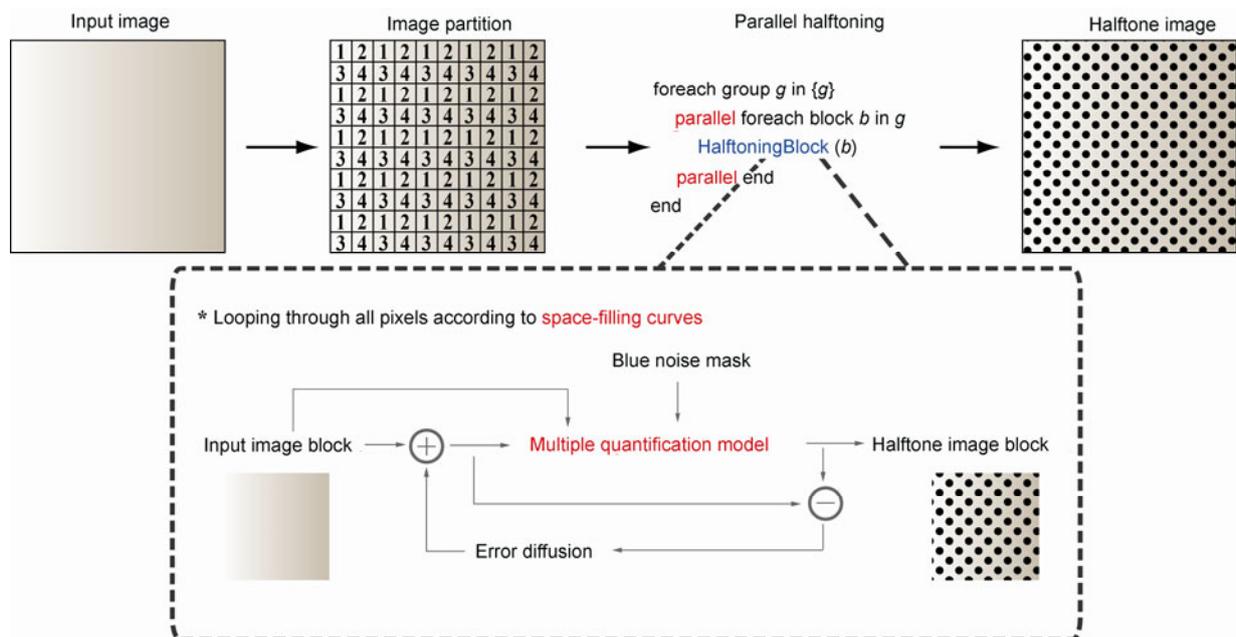


Fig. 2 The conceptual overview of our approach

The block-based error diffusion techniques (Chang and Allebach, 2003; Li and Allebach, 2005) effectively improve both tone reproduction and parallelism by decomposing the image into blocks. These blocks are then processed in a specific order, such as the interleaved order (Chang and Allebach, 2003). Handling the input image based on image blocks effectively provides parallelism and consequently improves the performance. However, this scheme has not been used in structure-preserving halftoning methods. The main reason is that block-wise error diffusion approaches generally require an offline tone-aware optimization process, and neglect structure preservation.

In our approach, the input image is divided into four groups of blocks to circumvent race conditions (Fig. 2). Within each block of each group, we employ the contrast-aware mask (Li and Mould, 2010) to accomplish error distribution after quantization is accomplished (Eq. (1)). The contrast-aware mask is a circular mask centered on the quantized pixel. It is computed with respect to the quantization error  $E(x, y)$  and the pixel intensities beneath this mask:

$$H(x, y) = \begin{cases} \frac{I'(x, y)}{R^k}, & \text{if } E(x, y) > 0, \\ \frac{1 - I'(x, y)}{R^k}, & \text{otherwise,} \end{cases} \quad (3)$$

where  $I'(x, y)$  is the intensity value of the pixel at  $(x, y)$ ,  $R$  denotes the offset from  $(x, y)$  to the center pixel, and  $k$  is a positive exponent. After all the weights  $H(x, y)$  of the mask are calculated, we normalize the weights by making the transformation  $\tilde{H}(x, y) = \frac{H(x, y)}{\sum_{(i, j)} H(i, j)}$ .

Then the quantization error of the mask center is distributed to its unprocessed neighbors:

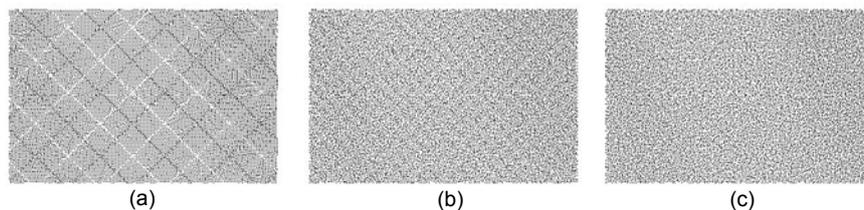
$$I'(x, y) = I(x, y) + E(x, y) \cdot \tilde{H}(x, y). \quad (4)$$

Obvious seams will arise if we simply employ the conventional quantization model and traversal path (e.g., the spiral path) to parallelize the basic contrast-aware technique (Fig. 3a). Our method effectively addresses this problem (Fig. 3c), and meanwhile offers efficient parallelism.

### 3.2 Multiple quantization model

The kernel of a halftoning algorithm is the conversion of a grey value to a binary value, called quantization. When the partitioning strategy is introduced into conventional error diffusion, the original serial continuity near block boundaries can be destroyed. The single threshold  $t$  is not able to cope adequately with this situation. In this work we introduce some other meaningful thresholds to avoid the relative instability existing in the classical quantization model. One is the blue-noise mask, and the other is the corresponding Gaussian filtered value. The grainy appearance of white-noise patterns is directly related to the fact that they possess energy at very low frequencies. Compared with white-noise patterns, blue-noise patterns are aperiodic and radially symmetric. Benefiting from this uncorrelated structure without low-frequency graininess, blue-noise patterns are more pleasing to the human vision system.

Billotet-Hoffman and Bryngdahl (1983) replaced the fixed threshold with conventional ordered dither arrays in the classical Floyd-Steinberg error diffusion algorithm. The resulting halftone images are



**Fig. 3 Results for a smooth grey image (266×173) by the block-wise parallelization of the basic contrast-aware technique with the conventional quantization model and the spiral path (a), our multiple quantization model and the spiral path (b), and the multiple quantization model and the space-filling path (c)**

It is apparent that (a) shows serious annoying seams. The slight diagonal artifacts shown in (b) may appear when the block size is 8×8. These artifacts are eliminated by introducing the space-filling path, even when the block size is smaller

not very different from those images obtained using traditional ordered dithering methods. The primitive ordered dithering (Bayer, 1973) performs quantization by comparing the input image with a threshold matrix in a pixelwise manner. However, its result contains visually periodic patterns. Following the blue-noise concept introduced by Ulichney (1988), the periodic patterns are eliminated in the blue-noise dithering approach (Mitsa and Parker, 1992), which employs a blue-noise mask. The blue-noise mask has the structure of an ordered-dither mask and blue-noise-producing properties. Similar to Billotet-Hoffman and Bryngdahl's idea, the blue-noise mask is employed in our multiple quantization model. In addition, the application of this mask removes the scanning and start-up artifacts that appear in the results from the error diffusion techniques.

The threshold  $t$  determines the quantized value by the current intensity of the pixel, not its initial intensity. The blue-noise threshold arrays actually introduce visually insensitive randomness into the quantization procedure. If the results given by both thresholds are in contradiction with each other, the Gaussian-filtered value of the neighborhood is used to maintain a close visual perception to the original image. If the pixel value is larger than this weighted average, the output is more likely to be 1, and vice versa.

The proposed multiple quantization model leads to sharp structures without distracting artifacts. Meanwhile, using the new model eliminates the start-up artifacts that are commonly produced by block-based error diffusion approaches. More importantly, this model is irrelevant to the error diffusion methods employed, and thus can be combined with arbitrary error diffusion approaches, especially block-wise error diffusion approaches. Essentially, these advantages are mainly due to the favorable randomness from the use of the blue-noise mask, which uses image information to guide the quantization.

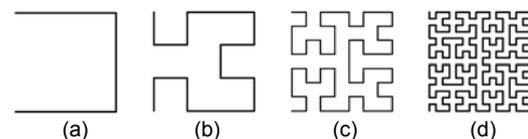
### 3.3 Traversal path based on the space-filling curve

The halftoning process of pixels in each block requires a traversal path in the block. The most popular order is the spiral scanning path (Chang and Allebach, 2003; Li and Allebach, 2005; Zhou *et al.*, 2009) but it may still yield distracting diagonal artifacts. Conventional tone-dependent error diffusion

methods (Li and Allebach, 2004) are insensitive to the traversal path, but insufficient to characterize human sensitivity to textures. The most recent approach that uses a contrast-aware mask to perform error diffusion (Li and Mould, 2010), can produce high quality results but is incompatible with the block-wise process mode, because it is very sensitive to the traversal path. Although a dynamic priority technique can partially address this problem, maintaining dynamic priority greatly influences the efficiency, and also impairs the quality of tone reproduction.

To approximate the halftoning process on the entire image by the block-based scheme, we must guarantee that the cumulative error existing within each block is limited, so that the cumulative error from all blocks is sufficiently small and does not influence the quality. The space-filling curves that have been used for halftoning (Witten and Neal, 1982; Velho and Gomes, 1991; Zhang and Webber, 1993; Wong and Hsu, 1995) can effectively avoid repetitive patterns.

A space-filling curve is initially invented to obtain a subjective and continuous map from a unit interval to a unit square. The Hilbert curve (Fig. 4) used in our approach is the first geometric construction for such a map. Other space-filling curves include the Moore curve, Peano curve, Lebesgue curve, and Sierpinski curve. In the field of digital halftoning, the space-filling curves are first used by Velho and Gomes (1991) to generate aperiodic patterns of clustered dots. In our method, all blocks of the four groups are scanned with the same Hilbert curve. The results shown in Fig. 3 reveal that the resulting halftone images are favorably free from the interference of the block-wise process mode even when the block size is very small.



**Fig. 4 Hilbert curves at different levels**  
(a) Level 1; (b) Level 2; (c) Level 3; (d) Level 4

### 3.4 Implementation details

We have implemented both a CPU-based and an unoptimized Compute Unified Device Architecture

(CUDA) based version of the proposed approach. In particular, we map four adjacent image blocks within different groups into one CUDA thread. A CUDA block contains a number of such image-block units. Because each image block has the same traversal path, the corresponding program control logic and memory addressing are relatively simple compared with conventional block-wise error diffusion approaches.

There are two types of masks in our approach. The contrast-aware mask is dependent on the underlying image, and has to be generated for each pixel. In contrast, the blue-noise mask is constant, and can be preprocessed. In the entire pipeline, both the block size and contrast-aware mask size influence the quality and GPU-oriented parallelism. Thus, appropriate values have to be used. According to our experiments, setting the block size and contrast-aware mask size to be  $8 \times 8$  and  $5 \times 5$  respectively can obtain pleasing results and high parallelism.

Note that overflow or underflow may occur when handling intensity adjustment. In these circumstances, the pixel is clamped to 0 or 255, and the residual error is added to the next pixel along the traversal path. In addition, the errors from isolated pixels should be included in the residual error.

## 4 Results and analysis

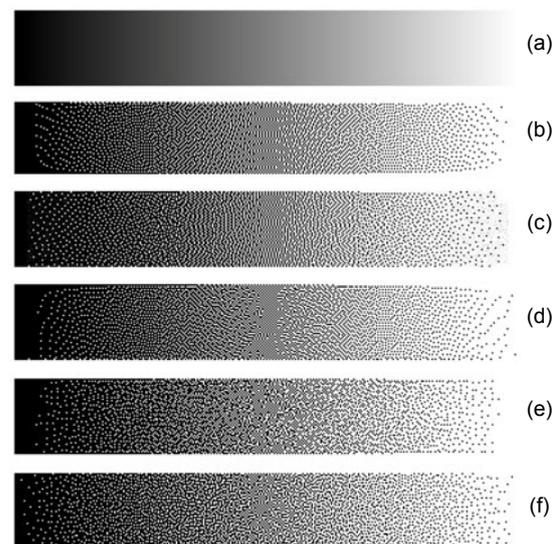
We have implemented the proposed approach and conducted several experiments on a rich variety of images characterized by either exquisite tone gradation or abundant textures. Below we verify our approach in terms of quality and efficiency, and compare it with representative approaches. Due to the special characteristics of the halftone image, the results are best viewed in the original size.

### 4.1 Quality

#### 4.1.1 Tone consistency

We measure the tone consistency of various input images by the peak signal-to-noise ratio (PSNR). PSNR is applied to the Gaussian blurred image and the corresponding halftoned image. It is well known that the Ostromoukhov method (Ostromoukhov, 2001) performs well in terms of tone reproduction. Table 1 compares the PSNR results obtained using our ap-

proach and the Ostromoukhov method. Our approach achieves comparable quality. Additionally, the comparisons for the Ramp image (Fig. 5) and the constant greyness images (Figs. 5–7) show that our method can improve the tonal quality of shadow and highlight regions. This is mainly due to the use of the multiple quantization model; meanwhile, using the space-filling curve as the traversal path helps eliminate negative effects caused by integration of the contrast-aware mask into the block-wise process mode. However, some artifacts still exist in the mid-tones (Fig. 5), as in other structure-aware methods. More effort is needed to improve mid-tone quality.



**Fig. 5 Result comparison for the Ramp image (330×50)**  
(a) The input; (b) Floyd-Steinberg ED; (c) SAH; (d) The basic CAH; (e) The variant CAH; (f) Ours

**Table 1 The PSNR measurement for various images**

Image	PSNR	
	Ours	Ostromoukhov
Mole	38.7652	41.0261
Portrait	37.7834	40.1502
Ribbon	27.8874	30.7509
Road	38.0954	40.0772
Arm	22.8009	22.6017
Bat	33.3089	34.7745
Cat	38.3482	33.7599
Knee	38.7506	30.7424
Snail	20.1646	19.9946

Results for the Ostromoukhov method come from Li and Mould (2010)

#### 4.1.2 Structure preservation

The mean structure similarity measure (MSSIM) has been widely used to measure the degree of structure preservation. Figs. 8 and 9 show the relationship between the block and mask sizes and the MSSIM measurement. Two facts are observed: the MSSIM increases with the increase of the mask size; the block size has no significant influence on structure preservation.

Table 2 lists the MSSIM measurements for different structure-preserving halftoning approaches. The comparison indicates that our approach outperforms other structure-preserving approaches on the representative test images. Although the variant CAH method achieves higher values, it costs much more time than the basic CAH method (as shown later in Table 4). Subjective observation on Figs. 1 and 5 makes the structural awareness of our method more convincing, without too obvious an influence of the block-wise process mode.

**Table 2 The MSSIM measurement for various images**

Image	MSSIM				
	Ours	Basic CAH	SAH	SAED	Ostro-moukhov
Mole	11.46	10.63	10.11		6.16
Portrait	29.82	28.77	27.45		18.61
Ribbon	33.73	32.92	28.51	32.37	28.36
Road	32.15	30.03	29.06		18.79
Arm	56.06	55.11	54.79	55.10	38.45
Bat	27.85	25.56	26.78		16.17
Cat	12.53	11.42	12.30		6.73
Knee	47.99	45.63	45.34	43.59	29.51
Snail	45.84	43.86	44.75	42.22	38.03

Results of other approaches come from Chang *et al.* (2009) and Li and Mould (2010). For the SAED method, only the results for the Ribbon, Arm, Knee, and Snail images are available

#### 4.1.3 Contrast similarity

To further verify the contrast awareness of our method, we use the contrast PSNR (CPSNR) to measure contrast similarity. This metric was proposed by Matković *et al.* (2005) and introduced by Li and Mould (2010) to evaluate contrast preservation. Table 3 shows that our approach has a comparable measurement to CAH methods in this test.

#### 4.1.4 Blue-noise property

Containing the blue-noise property has proven to

be essential to identify the quality of a halftoning method. An ideal blue-noise binary pattern can be demonstrated by the radially averaged power spectrum (RAPS) or the frequency spectrum of the halftoned image (Ulichney, 1987).

**Table 3 The CPSNR measurement for various images**

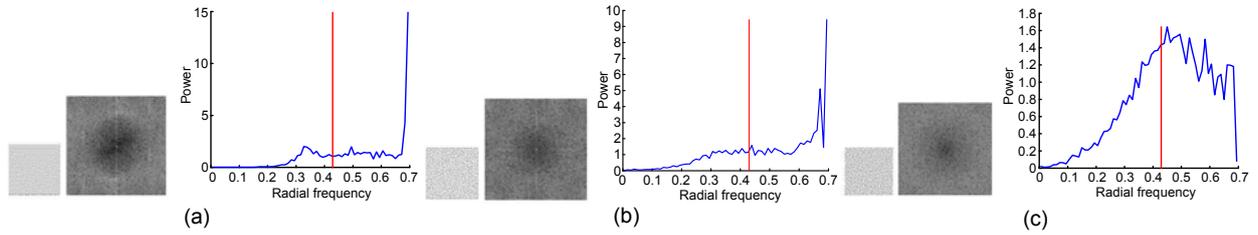
Image	CPSNR				
	Ours	Variant CAH	Basic CAH	Ostro-moukhov	FS error diffusion
Mole	6.50	6.61	6.43	6.27	6.16
Portrait	7.99	7.89	7.59	7.24	7.12
Ribbon	8.22	8.26	7.96	7.75	7.62
Road	7.83	7.57	7.34	7.04	6.92
Arm	11.05	10.93	10.68	10.24	10.21
Bat	7.28	7.21	6.91	6.45	6.38
Cat	6.37	6.51	6.26	6.29	6.01
Knee	10.00	9.73	9.42	8.89	8.88
Snail	9.12	9.02	8.74	8.58	8.50

Results of other approaches come from Li and Mould (2010). FS: Floyd-Steinberg

Fig. 6 shows the results for a smooth grey image with the CAH methods and our approach. When the basic CAH method is used, start-up artifacts are generated (Fig. 6a). To improve the blue-noise property of the contrast-aware scheme (Li and Mould, 2010), other techniques like the random tie-breaking of equal priorities are introduced. Although the improved variant CAH method achieves much better structural preservation, the blue-noise property is still unsatisfying, as is demonstrated by the frequency spectrum shown in Fig. 6b. In contrast, our result (Fig. 6c) shows the desired pattern. Another comparison is made between our approach and the structure-aware halftoning algorithm (Pang *et al.*, 2008), which attains the blue-noise property. The similar blue-noise patterns shown in Fig. 7 demonstrate that our approach possesses the blue-noise property just as the structure-aware halftoning algorithm does.

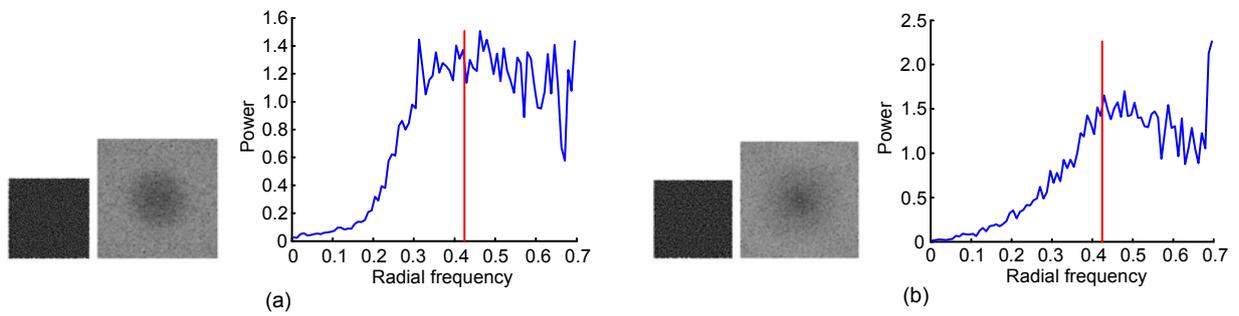
## 4.2 Performance

It is hard for us to have all the implementations and test them in the same hardware configuration. Below we list the performance statistics of each approach collected from the corresponding paper resource, and give a comparative analysis.



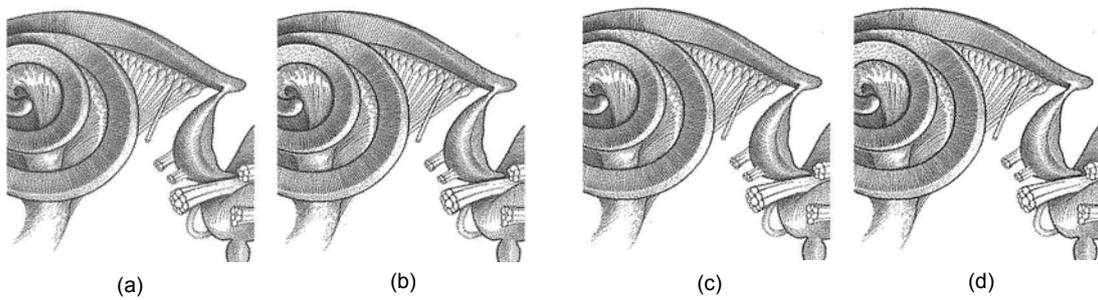
**Fig. 6 Results for a constant-greyness image (100×100, greyness=0.82) by the basic CAH method (a), the variant CAH method (b), and our approach (c)**

The principal frequency is plotted with a vertical line on the frequency axis. In each subfigure, from left to right: halftoned image, frequency spectrum, and radially averaged power spectrum density (RAPSDF)

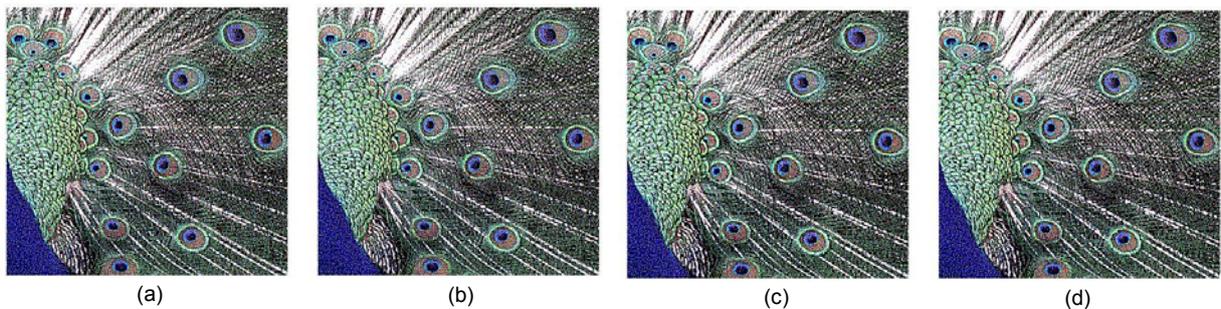


**Fig. 7 Results for a constant-greyness image (128×128, greyness=0.18) by the SAH algorithm (a) and our approach (b)**

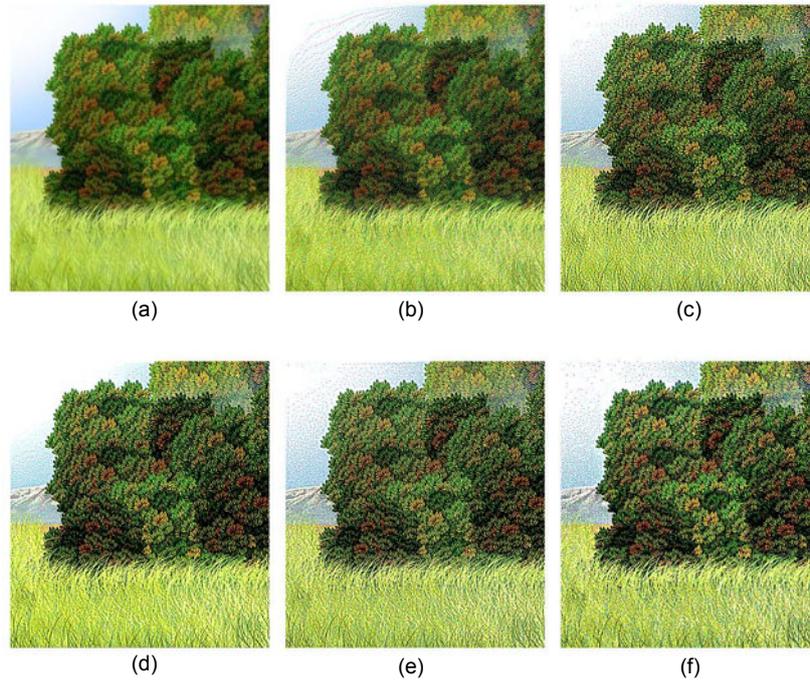
In each subfigure, from left to right: halftoned image, spectrum, and radially averaged power spectrum density (RAPSDF)



**Fig. 8 Results for the Snail image (383×399) with different block and mask sizes: (a) 8×8, 5×5; (b) 8×8, 7×7; (c) 30×30, 7×7; (d) 64×64, 9×9**



**Fig. 9 Results for a color image (361×344) with different block and mask sizes: (a) 16×16, 7×7; (b) 32×32, 7×7; (c) 32×32, 9×9; (d) 64×64, 9×9**



**Fig. 10 Result comparison for the ColorTree image (400×445)**

(a) The input; (b) Floyd-Steinberg error diffusion (Floyd and Steinberg, 1976); (c) SAH (Pang *et al.*, 2008); (d) The basic CAH (Li and Mould, 2010); (e) The variant CAH (Li and Mould, 2010); (f) Ours

**Contrast-aware halftoning:** Compared with the structure-aware halftoning algorithm (Pang *et al.*, 2008), both contrast-aware halftoning (Li and Mould, 2010) and structure-aware error diffusion (Chang *et al.*, 2009) approaches achieve higher performance, but still cannot be parallelized. The best tradeoff between performance and quality for the contrast-aware halftoning approach is reached by using a  $7 \times 7$  mask (Li and Mould, 2010). For a  $512 \times 512$  image, it takes 0.492 s with the basic CAH method and 2.955 s with the variant CAH method to perform halftoning on a PC equipped with an Intel Core Duo 3.0 GHz CPU and 3 GB host memory. It takes 0.27 s and 2.202 s for basic and variant CAH respectively if a  $5 \times 5$  mask is used. Table 4 summarizes the timing statistics.

**Table 4 Performance of the CAH approach\***

Mask size	Time (s)	
	Basic CAH	Variant CAH
$5 \times 5$	0.27	2.20
$7 \times 7$	0.49	2.96
$9 \times 9$	0.78	3.12

\* Only CPU is used

**Structure-aware error diffusion:** The running time of the structure-aware error diffusion approach (Chang *et al.*, 2009) includes two parts, local frequency analysis and error diffusion. While local frequency analysis can be performed in parallel, the error diffusion process has to be done in a sequential way. Thus, the entire pipeline cannot be parallelized. In particular, the error diffusion costs 0.22 s for a typical  $512 \times 512$  image. The timing of local frequency analysis depends on the kernel size. If the size is not adequately large, important structural details may be lost. For a typical kernel size  $16 \times 16$ , local frequency analysis takes 0.75 s, and the total time is 0.97 s (Table 5).

**Ours:** We have implemented a software-based halftoning system and an unoptimized GPU version of the proposed approach. The GPU implementation is tested on a PC equipped with an Intel Core Quad 2.83 GHz with 3.25 GB host memory, and an NVIDIA GeForce GTX 480 video card. Table 6 lists our performance statistics with different mask sizes. The performance is inversely proportional to the sizes of the image blocks and the contrast-aware mask. To achieve the best performance and quality balance, they are set to be  $8 \times 8$  and  $5 \times 5$  in our experiments.

When the block size is  $8 \times 8$ , a  $512 \times 512$  image contains  $64 \times 64$  blocks, which means that 1024 ( $32 \times 32$ ) GPU threads are needed to achieve full parallelism.

**Table 5 Performance of the SAED approach**

Kernel size	Time (s)	
	CPU only	CPU+GPU
$4 \times 4$	1.40	0.26
$8 \times 8$	2.49	0.45
$16 \times 16$	6.96	0.97
$32 \times 32$	26.53	15.31

**Table 6 Performance of our approach**

Mask size	Time (s)
$5 \times 5$	0.0238
$7 \times 7$	0.0313

Block size is  $8 \times 8$

Note that the current test platform for our approach provides limited computational resources for parallelism of GPU implementation because a GeForce GTX 480 video card contains 480 cores. Another test is carried out on the same PC equipped with an NVIDIA GeForce 9600 GT video card with 512 MB video memory; the performance becomes 0.09 s and 0.12 s, respectively. Even though the computational power of our test platform is similar to, or worse than, the ones used in the experiments of the aforementioned approaches, the timing statistics demonstrate the advantages of our approach in respect of parallelism. On one hand, the contrast- and structure-aware halftoning approaches cannot be directly implemented on GPU. On the other hand, although the local frequency analysis stage of the structure-aware error diffusion approach can be implemented in parallel, the second stage has to be processed sequentially, making it not suitable for full GPU implementation.

## 5 Conclusions

We present a novel parallelizable structure-preserving method. Experimental results on a variety of images demonstrate the efficiency of our approach. In particular, our approach gains higher structural preservation and fine blue-noise property, and achieves higher performance compared with the basic

version of the most recent contrast-aware halftoning method. Although the improved variant CAH method enhances structural preservation by introducing a priority scheme, its performance is greatly lowered.

One main concern of our approach is that both the capability of structural preservation and the degree of parallelism are dependent on the size of the contrast-aware mask. In other words, a smaller block size yields better parallelism, yet can lower the preservation of important structures. We plan to find a new way to optimize the sizes of the mask and the image blocks, and further improve the capability for structural preservation when the block size is very small.

In addition, our approach directly associates all blocks with the same given Hilbert curve. There may be more helpful treatments, such as different Hilbert curves for different groups, or each block just scanned with a randomly selected Hilbert curve. Digital color halftoning is a more advanced topic related to the interactions between color channels, color contrast, etc. Here we simply handle each color channel separately for a color image, and exclude those relationships (Figs. 9 and 10).

## References

- Analoui, M., Allebach, J.P., 1992. Model-based halftoning using direct binary search. *SPIE*, p.96-108.
- Bayer, B.E., 1973. An Optimum Method for Two-Level Rendition of Continuous-Tone Pictures. *IEEE Int. Conf. on Communications*, p.26-11-26-15.
- Billotet-Hoffmann, C., Bryngdahl, O., 1983. On the error diffusion technique for electronic halftoning. *Proc. Soc. Inf. Display*, **24**(3):253-258.
- Chang, J.H., Alain, B., Ostromoukhov, V., 2009. Structure-aware error diffusion. *ACM Trans. Graph.*, **28**(5):162:1-162:8. [doi:10.1145/1661412.1618508]
- Chang, T.C., Allebach, J.P., 2003. Memory efficient error diffusion. *IEEE Trans. Image Process.*, **12**(11):1352-1366. [doi:10.1109/TIP.2003.818214]
- Eschbach, R., Knox, K.T., 1991. Error-diffusion algorithm with edge enhancement. *J. Opt. Soc. Am. A*, **8**(12):1844-1850. [doi:10.1364/JOSAA.8.001844]
- Floyd, R.W., Steinberg, L., 1976. An adaptive algorithm for spatial grayscale. *Proc. Soc. Inf. Display*, **17**(2):75-77.
- Knuth, D.E., 1987. Digital halftones by dot diffusion. *ACM Trans. Graph.*, **6**(4):245-273. [doi:10.1145/35039.35040]
- Li, H., Mould, D., 2010. Contrast-aware halftoning. *Comput. Graph. Forum*, **29**(2):273-280. [doi:10.1111/j.1467-8659.2009.01596.x]
- Li, P.S., Allebach, J.P., 2004. Tone-dependent error diffusion. *IEEE Trans. Image Process.*, **13**(2):201-215. [doi:10.

- 1109/TIP.2003.819232]
- Li, P.S., Allebach, J.P., 2005. Block interlaced pinwheel error diffusion. *J. Electron. Imag.*, **14**(2):023007. [doi:10.1117/1.1900136]
- Matković, K., Neumann, L., Neumann, A., Psik, T., Purgathofer, W., 2005. Global Contrast Factor—a New Approach to Image Contrast. Proc. 1st Eurographics Conf. on Computational Aesthetics in Graphics, Visualization and Imaging, p.159-167.
- Mitsa, T., Parker, K.J., 1992. Digital halftoning technique using a blue-noise mask. *J. Opt. Soc. Am. A*, **9**(11):1920-1929. [doi:10.1364/JOSAA.9.001920]
- Neuhoff, D.L., Pappas, T.N., Seshadri, N., 1997. One-dimensional least-squares model-based halftoning. *J. Opt. Soc. Am. A*, **14**(8):1707-1723. [doi:10.1364/JOSAA.14.001707]
- Ostromoukhov, V., 2001. A Simple and Efficient Error-Diffusion Algorithm. Proc. 28th Annual Conf. on Computer Graphics and Interactive Techniques, p.567-572. [doi:10.1145/383259.383326]
- Pang, W.M., Qu, Y.G., Wong, T.T., Cohen-Or, D., Heng, P.A., 2008. Structure-aware halftoning. *ACM Trans. Graph.*, **27**(3):89:1-89:8. [doi:10.1145/1360612.1360688]
- Ulichney, R., 1987. Digital Halftoning. MIT Press, Cambridge, MA.
- Ulichney, R.A., 1988. Dithering with blue noise. *Proc. IEEE*, **76**(1):56-79. [doi:10.1109/5.3288]
- Velho, L., Gomes, J.D.M., 1991. Digital halftoning with space filling curves. *ACM SIGGRAPH Comput. Graph.*, **25**(4): 81-90. [doi:10.1145/127719.122727]
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, **13**(4):600-612. [doi:10.1109/TIP.2003.819861]
- Wei, L.Y., 2008. Parallel Poisson disk sampling. *ACM Trans. Graph.*, **27**(3):20:1-20:9. [doi:10.1145/1360612.1360619]
- Witten, I.H., Neal, R.M., 1982. Using Peano curves for bilevel display of continuous-tone images. *IEEE Comput. Graph. Appl.*, **2**(3):47-52. [doi:10.1109/MCG.1982.1674228]
- Wong, T.T., Hsu, S.C., 1995. Halftoning with Selective Precipitation and Adaptive Clustering. In: Paeth, A.W. (Ed.), Graphics Gems V (IBM Version). Morgan Kaufmann, Burlington, USA, p.302-313.
- Wu, H.S., Wong, T.T., Heng, P.A., 2013. Parallel structure-aware halftoning. *Multim. Tools Appl.*, **67**(3):529-547. [doi:10.1007/s11042-012-1048-6]
- Zhang, Y.F., Webber, R.E., 1993. Space Diffusion: an Improved Parallel Halftoning Technique Using Space-Filling Curves. Proc. 20th Annual Conf. on Computer Graphics and Interactive Techniques, p.305-312. [doi:10.1145/166117.166156]
- Zhou, B.F., Fang, X.F., 2003. Improving mid-tone quality of variable-coefficient error diffusion using threshold modulation. *ACM Trans. Graph.*, **22**(3):437-444. [doi:10.1145/882262.882289]
- Zhou, Y., Chen, C., Wang, Q., Bu, J.J., 2009. Block-Based Threshold Modulation Error Diffusion. Proc. 16th IEEE Int. Conf. on Image Processing, p.2033-2036. [doi:10.1109/ICIP.2009.5413974]