



ELSEVIER

Contents lists available at ScienceDirect

Signal Processing: *Image Communication*journal homepage: www.elsevier.com/locate/image

Application of reversible denoising and lifting steps to DWT in lossless JPEG 2000 for improved bitrates



Roman Starosolski*

Institute of Computer Science, Silesian University of Technology, Akademicka 16, 44-100 Gliwice, Poland

ARTICLE INFO

Article history:

Received 12 June 2015

Received in revised form

29 August 2015

Accepted 18 September 2015

Available online 9 October 2015

Keywords:

Reversible denoising and lifting step

Discrete wavelet transform

Denoising

Lifting technique

Lossless image compression

JPEG 2000.

ABSTRACT

In a previous study, we noticed that the lifting step of a color space transform might increase the amount of noise that must be encoded during compression of an image. To alleviate this problem, we proposed the replacement of lifting steps with reversible denoising and lifting steps (RDLS), which are basically lifting steps integrated with denoising filters. We found the approach effective for some of the tested images. In this study, we apply RDLS to discrete wavelet transform (DWT) in JPEG 2000 lossless coding. We evaluate RDLS effects on bitrates using various denoising filters and a large number of diverse images. We employ a heuristic for image-adaptive RDLS filter selection; based on its empirical outcomes, we also propose a fixed filter selection variant. We find that RDLS significantly improves bitrates of non-photographic images and of images with impulse noise added, while bitrates of photographic images are improved by below 1% on average. Considering that the DWT stage may worsen bitrates of some images, we propose a couple of practical compression schemes based on JPEG 2000 and RDLS. For non-photographic images, we obtain an average bitrate improvement of about 12% for fixed filter selection and about 14% for image-adaptive selection.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

1.1. Overview

Advantageous properties of a lifting scheme [1] (e.g., perfect reconstruction and in-place operation) made it useful to construct significant transforms for the lossless image compression domain, including discrete wavelet transform (DWT) and various color space transforms [2]. The scheme decomposes a transform into a series of lifting steps that are reversible even if implemented using finite precision integer numbers.

In a previous study [3], we noticed that the lifting step of a color space transform might increase the amount of noise that must be encoded during compression of an image. It is

known that Red, Green, and Blue primary color components of the RGB color space are highly correlated for natural images [2]. A common approach to RGB color image compression is to independently compress the image components obtained using a color space transform from the RGB to a less correlated color space. An unwanted side effect of color space transform performed using lifting steps is that removing correlation contaminates the transformed components with noise from other components. Although such transform does not change the overall image information content, the amount of noise in a component considered individually is increased. Because of the independent compression of transformed image components, the overall amount of noise that has to be encoded increases. To remove correlation without increasing noise in components, we proposed the replacement of lifting steps with reversible denoising and lifting steps (RDLS); basically, the new steps are lifting steps integrated with denoising filters. We found the approach to be effective for some of the tested images. In this study, we

* Tel.: +48 322372151.

E-mail addresses: rstarosolski@polsl.pl, rstaros@gmail.com

apply RDLS to DWT in lossless JPEG 2000 coding [4–6], evaluate RDLS effects on bitrates of various types of images, and show practical compression schemes exploiting RDLS.

The remainder of this paper is organized as follows. In Sections 1.2 and 1.3, we briefly characterize the DWT used in lossless JPEG 2000 and the RDLS. The main contribution of this study is described in Section 2, where a new RDLS–DWT transform obtained by the application of RDLS to DWT is introduced along with a heuristic for selecting denoising filters for RDLS–DWT. We evaluate RDLS–DWT on classic test images, images with impulse or Gaussian noise added, and a recent, large, and diverse set of images. Section 3 describes the experimental procedure and the test data used. Section 4 presents the results and discussion. Section 5 summarizes the research.

1.2. Discrete wavelet transform in reversible JPEG 2000 coding

For brevity, we describe here only the lifting-based 5×3 kernel DWT that is used in lossless JPEG 2000 compression of grayscale images, reduced to essentials. For further details as well as for more general characteristics of DWT, JPEG 2000, and the lifting scheme, the reader is referred to [1,4–6].

Using the lifting scheme [1], the one-dimensional DWT (1D DWT) transforms in-place a discrete signal $S = s_0 s_1 s_2 \dots s_{l-1}$ of finite length l into two subbands:

- a low-pass filtered signal L representing the original signal's low-frequency features; and
- a high-pass filtered signal H containing high-frequency features that, along with the low-pass signal, allows the perfect reconstruction of the original signal.

S is transformed in 3 steps. First, in the prediction step, we perform the high-pass filtering of odd samples (hereafter, the parity of sample or pixel is determined by its location and not its value) by applying the lifting step Eq. (1) to each of them:

$$s_x \leftarrow s_x - \lfloor (s_{x-1} + s_{x+1})/2 \rfloor, \tag{1}$$

where the floor symbol $\lfloor v \rfloor$ denotes the greatest integer not exceeding v . Note that the amount by which the value of the odd sample is changed depends only on even samples. Another lifting step is then applied to each even sample (update step):

$$s_x \leftarrow s_x + \lfloor (s_{x-1} + s_{x+1} + 2)/4 \rfloor. \tag{2}$$

Here, even samples are updated based directly on the already transformed odd samples. Finally, in the reorder step, we reposition even samples to the lower half of the original signal, preserving their ordering (sample s_x is moved to $s_{\lfloor x/2 \rfloor}$), and odd samples are moved to the upper half. We obtain separate subbands L and H , respectively. The two-dimensional DWT (2D DWT) transform of an image is obtained by first applying 1D DWT to each image column, which results in L and H subbands of the image. Then by applying 1D DWT to each row, we obtain the 1-level DWT transform, consisting of LL and HL subbands (transformed from L subband) and LH and HH subbands (from H subband); see Fig. 1A–C. Higher-level DWT transforms that provide multiresolution image representation are obtained by Mallat decomposition [7]. The $t+1$ -level transform is obtained by applying the 1-level transform to the LL subband of the t -level transform (Fig. 1D).

In lossless JPEG 2000, the transformed image is encoded in a complex and flexible manner [4–6]. For the remainder of this study, it is noteworthy that each subband is compressed independently of the others using a context-adaptive entropy coder.

1.3. Reversible denoising and lifting step

The lifting technique is also used in other domains; for example, the reversible color space transforms are built as sequences of lifting steps [2]. Such transforms remove correlations among color image pixel components, but, as a consequence, a transformed component is contaminated with noise from more than one original component. In [3], we integrated denoising into lifting steps to avoid noise propagation while preserving other transform properties (i.e., reversibility and removing correlation). The lifting step that was modified was generalized as the following:

$$c_x \leftarrow c_x \oplus f(c_1, \dots, c_{x-1}, c_{x+1}, \dots, c_m), \tag{3}$$

where c_n is the n -th component of the pixel and m is the number of components. The step is reversible provided that the function f is deterministic and the operation \oplus is reversible, that is, an inverse operation \otimes exists such that $c = a \oplus b \Leftrightarrow a = c \otimes b$.

Based on the lifting step (Eq. (3)), we constructed a reversible denoising and lifting step (RDLS, Eq. (4)) by simply denoising arguments of function f :

$$c_x \leftarrow c_x \oplus f(c_1^d, \dots, c_{x-1}^d, c_{x+1}^d, \dots, c_m^d), \tag{4}$$

where c_n^d is the denoised n -th component of the pixel, which is obtained by using a denoising filter, and different filters may be used for different components. Eq. (3) implies that the function f may use any arguments, except for c_x . In Eq. (4), we denoise these arguments. For denoising them we may use any component of any pixel,




Fig. 1. 1-level 2D DWT of an image (A–C) and 3-level 2D DWT (D).

except for c_x of the pixel to which the RDLS-modified lifting step is being applied; otherwise, this c_x would be indirectly used by f . Obviously, the denoising filter must be deterministic. A color space transform may be performed for each pixel independently of other pixels. The RDLS sequence, constructed based on a color space transform, is a transform of whole image components, not of the color space, since denoising of a pixel's component requires access to the same component of (at least) neighboring pixels. For the RDgDb color space transform [8] (also known as $A_{2,1}$ [9]) and a very simple denoising filter (linear smoothing filter), we found that the proposed method was effective for images in optical resolutions of acquisition devices. As noted in [3], the RDLS method is general in nature and is applicable to other lifting-based transforms; in this study, we apply it to DWT in lossless JPEG 2000. An example of RDLS application to DWT is presented in Section 2.2.

2. Proposed approach

2.1. RDLS-modified DWT

The application of RDLS to the 5×3 kernel 2D DWT is straightforward when we notice the analogy between a color space transform applied to a color image and a DWT of a grayscale image. In the former case, we apply a color space transform as a sequence of lifting steps to pixels consisting of color space components; all pixels constitute a color image. In the latter, we first apply a 1D DWT transform (many equivalent sequences of lifting steps are possible) to image columns consisting of pixels; all columns constitute a grayscale image. Then we repeat the process for the obtained rows. Thus, in order to obtain a RDLS-modified DWT transform (RDLS-DWT), we simply replace the DWT lifting steps with RDLS constructed based on them. The RDLS-modified prediction lifting step (Eq. (1)) becomes

$$s_x \leftarrow s_x - \left[(s_{x-1}^d + s_{x+1}^d) / 2 \right], \quad (5)$$

where s_n^d is a denoised sample s_n , and the RDLS-update is the following:

$$s_x \leftarrow s_x + \left[(s_{x-1}^d + s_{x+1}^d + 2) / 4 \right]. \quad (6)$$

But how do we select the denoising filter? Naturally, the filter should be matched to the noise characteristics of the data being filtered. We assume that the noise characteristics of the original image as well as of each of its subbands (including those that are temporarily created and then further transformed – see Fig. 1B), are invariant, i.e., the same denoising filter may be used while computing all samples of a specific subband. For example, when computing H subband by applying Eq. (5) for each odd sample in each column of the original image, we use the same filter for denoising all even samples. However, the noise characteristics of different subbands differ, so a different filter may be used when computing the HL subband.

In this study, for the denoising of a sample of a specific parity, we use samples of the same parity only. In the

above case of computing the H subband, for denoising a particular even sample of a particular column, we may use all even samples of all columns. Generally, during computation of a specific subband, we perform the denoising based on samples not belonging to this subband. For the process to be reversible, as noted in Section 1.3, while applying the RDLS to a specific sample s_x , we may use all signal samples for denoising, except the s_x . However, certain practical problems would arise if we used samples from the subband being computed. We assume that the subband characteristics are invariant; this assumption applies to the already computed subbands and to the original signal. Obviously, it cannot be valid for a subband during its computation, as during the computation only some subband samples are already transformed. Thus, for the denoising, we should either only use untransformed samples from this subband, or we should employ a sophisticated denoising filter that is able to take advantage of both transformed and untransformed samples. Second, actual JPEG 2000 implementations apply the lifting steps to a given subband in the same order when performing the inverse transform, as is applied during a forward transform; if we used samples from the subband being computed, then the orders should be exactly the opposite.

A special case of the denoising filter is the filter that does not alter the sample being denoised, denoted here as the None filter. Using the None filter in RDLS makes it a regular lifting step. Since samples may be noise-free, when performing the filter selection, we should be able to select the None filter.

For the t -level RDLS-DWT, we need to select denoising filters for $6t$ subbands. One of the possibilities explored in this study is to define a set of filters from which we select the filter for each subband individually, based on actual image compression effects. Subbands are encoded independently but are interdependent. For instance, selecting a filter for some subband that results in a worse bitrate of that subband may help to find a better filter for another subband and therefore may result in an overall better bitrate. Since testing the compression effects of all the combinations of denoising filters would be too complex even for just a few filters considered, we employ a filter selection heuristic, as described in Section 2.3. The resulting filter selection should be transmitted to the decoder as side information along with the compressed image.

In this study, we apply RDLS to the 5×3 kernel DWT because this is the wavelet used in lossless JPEG 2000 coding. DWT with different kernels, as well as other transforms that are implemented using lifting steps, may be modified in a similar way, i.e., by replacing lifting steps (Eq. (3)) with RDLSs (Eq. (4)). However, finding denoising filters for other transforms may be more complicated. For example, for the t -level 9×7 kernel transform, we would have to select $12t$ filters, since 1-level 9×7 kernel 1D DWT is done using 4 steps involving lifting (as opposed to 2 such steps in case of 5×3 kernel).

2.2. RDLS-DWT example

Assume that while computing the level 1 of unmodified 2D DWT, we perform the prediction step for image rows by

applying 1D DWT to each of them. The input signal is an image that contains a smoothly changing background and salt-and-pepper noise (Fig. 2A). We start the forward transform from the 3rd sample in the 1st image row $s_{3,1}$, and while reconstructing the image, we apply inverses of the lifting steps in reverse order as compared to forward transform. In contrast, an actual JPEG 2000 codec would start from $s_{1,1}$ and use the same order for forward and inverse transform. The prediction lifting step (Eq. (1)) for $s_{3,1}$ is $s_{3,1} \leftarrow s_{3,1} - \lfloor (s_{2,1} + s_{4,1})/2 \rfloor$, so $s_{3,1}$ becomes 28 (Fig. 2B). When reconstructing $s_{3,1}$, the signal contents will be identical to just after performing a forward lifting step for this sample (Fig. 2B). Lifting steps are trivially invertible, the inverse of Eq. (1) is $s_x \leftarrow s_x + \lfloor (s_{x-1} + s_{x+1})/2 \rfloor$, and in our case $s_{3,1} \leftarrow s_{3,1} + \lfloor (s_{2,1} + s_{4,1})/2 \rfloor$, so $s_{3,1}$ is reconstructed to 55.

For RDLS–DWT, we use a median filter with a 3×3 sample window. The RDLS-modified prediction lifting step (Eq. (5)) applied to $s_{3,1}$ is $s_{3,1} \leftarrow s_{3,1} - \lfloor (s_{2,1}^d + s_{4,1}^d)/2 \rfloor$, and $s_{3,1}$ becomes 1 (Fig. 2C). Note that denoising of a sample is performed using samples of the same parity only, so for denoising $s_{2,1}$ and $s_{4,1}$ during a step performed for image rows, we use only samples from even columns (Fig. 2D) and obtain $s_{2,1}^d = 54$ and $s_{4,1}^d = 55$. During the inverse transform, we apply the inverse of Eq. (5), that is, $s_x \leftarrow s_x + \lfloor (s_{x-1}^d + s_{x+1}^d)/2 \rfloor$. To $s_{3,1}$, we apply $s_{3,1} \leftarrow s_{3,1} + \lfloor (s_{2,1}^d + s_{4,1}^d)/2 \rfloor$, where the same $s_{2,1}^d$ and $s_{4,1}^d$ are obtained as during the forward transform, since they are results of denoising of exactly the same samples (Figs. 2C and 2D). Thus, $s_{3,1}$ is reconstructed to 55.




Fig. 2. Example of DWT and RDLS–DWT. A – original signal, B – signal after applying the prediction lifting step to sample $s_{3,1}$, C – signal after applying the RDLS-modified prediction lifting step to sample $s_{3,1}$, D – locations of samples used for denoising of $s_{2,1}$ (■) and $s_{4,1}$ (□).

It is worth noting that in the above example, RDLS reduced the effect of noise on the quality of prediction. Without noise, if $s_{4,1}$ was 56, the regular prediction would result in setting $s_{3,1}$ to 0; contaminating $s_{4,1}$ with impulse noise resulted in a much greater change for unmodified DWT than for RDLS–DWT due to efficient denoising of $s_{4,1}$ exploited by the latter.

2.3. Filter selection heuristic

Our filter selection heuristic consists of the greedy steps described below, in which step B may be repeated certain number of iterations. With a slight abuse of terminology, we also use the term JPEG 2000 compression for JPEG 2000 with DWT replaced by RDLS–DWT, even though it is not compliant with the JPEG 2000 standard.

- (A) For each of the filters, perform JPEG 2000 compression of an image, using this filter in RDLS steps for all subbands at all levels. Then for all subbands at all transform levels, select the filter that resulted in the best overall bitrate.
- (B) For each transform level a (starting from level 1) and for each subband b (at specific level analyzed in the $H, L, HL, HH, LL,$ and LH order), try to find a better denoising filter by checking for each filter (except for the one already selected) the bitrate obtained using this filter for subband b at level a , while the filters selected so far are used for other subbands.

It is noteworthy that the heuristic selects a filter instead of the one already selected only if such a change improves the overall bitrate. The None filter that turns RDLS into a regular lifting step may be selected in step A for all subbands; thus, the bitrate cannot be worsened by the heuristic compared to unmodified DWT.

The heuristic is quite time-consuming. The RDLS–DWT computational time complexity alone is higher than that of the unmodified DWT due to the noise filtering employed. However, subbands are encoded independently; thus, for example, changing the denoising filter for the level-2 subband neither affects the compression results of level-1 subbands $LH, HL,$ and HH nor requires their repeated denoising. There are two main elements of heuristic computational time complexity: RDLS–DWT subband transforms and encoding of transformed samples. The number of symbols that need to be encoded during a single iteration of step B of the heuristic for the t -level transform is

$$(4 - 4^{1-t})(f - 1)p, \tag{7}$$

where f is the number of denoising filters, and p is the number of pixels in the image. The number of RDLS per single iteration of step B is

$$\left(\frac{86}{9} - \frac{1}{3}2^{5-2t}t - \frac{43}{9}2^{1-2t}\right)(f - 1)p. \tag{8}$$

Note that the t -level DWT (or RDLS–DWT if the filters for the latter are already selected) is done using the

following number of lifting steps (or RDLS):

$$\left(\frac{8}{3} - \frac{1}{3}2^{3-2t}\right)p. \quad (9)$$

To allow straightforward comparison of the complexities of various variants of the heuristic and of compression schemes involving the heuristic and the unmodified JPEG 2000, we express the heuristic computational time complexity as relative to unmodified JPEG 2000. A few simplifications are necessary to this aim. Increasing the transform level, except for a couple of lowest levels, does not influence the heuristic or transform complexity significantly. Assuming the infinite transform level, Eq. (7) becomes $4(f-1)p$; thus, a single iteration of step B requires encoding of roughly $4(f-1)$ times more symbols than the unmodified JPEG 2000. By dividing Eq. (8) by Eq. (9), we find that the number of RDLS per iteration of heuristic step B is greater than the number of lifting steps per unmodified DWT $(43/12)(f-1)$ times, which we approximate with $4(f-1)$. All in all, the estimated complexity of the heuristic (step A and i iterations of step B) is roughly

$$f + 4(f-1)i \quad (10)$$

times greater than the complexity of the JPEG 2000 compression with additional subband denoising due to replacing lifting steps with RDLS.

3. Materials and methods

3.1. Denoising methods and denoising filters

In the research, the following denoising methods and filters were used:

- None – No denoising. RDLS becomes a regular lifting step.
- Smoothing – We used 12 simple low-pass linear averaging filters with 3×3 and 5×5 sample windows. The filtered sample was calculated as a weighted arithmetic mean of samples from the window. The number of samples used was smaller than the window size at the image edges. The weight of the window center point ranged between 1 and 1024 (even powers of 2 only), while the others' weights were fixed to 1. Smoothing filters (3×3 sample window) were found to be effective for the RDLS-modified color space transform and certain natural images [3].
- Median – Two median filters were used with 3×3 and 5×5 windows. The Median filter and other nonlinear filters described below are effective at removing impulse noise.
- RCRS-1 – Two filters (3×3 and 5×5 windows), which belong to a general family of rank-conditioned rank selection (RCRS) filters [10]. RCRS-1 filters replace a sample with the window median if the sample is greater than or smaller than all other samples in the window.
- RCRS-2 – Two filters (3×3 and 5×5 windows) that replace a sample with the second greatest window

sample value if the sample is greater than the median and the greatest; or, if it is smaller than the median and the smallest, they replace a sample with the second smallest window sample value. A practical advantage of this filter, as opposed to the Median and RCRS-1, is that it may be computed without sorting the window samples.

The effects of applying the above filters to a sample image are presented in Fig. 3; for examining filtering methods on the Reader's own images, a free implementation was prepared [11].

3.2. Test data

We used two diverse, publicly available datasets and additional sets of noisy images; all images were of 8-bit precision.

- GS2 – This is the classic and still commonly used set (“Grayscale Set 2”) of 12 images from the University of Waterloo Fractal Coding and Analysis Group [12] (Fig. 4). The set contains grayscale, natural photographic images (among others, “lena” and “peppers”) as well as others (computer generated, mixed-content, dithered, and satellite); image sizes vary from 464×352 to 672×496 .
- CT2g – This set contains images that are green components of images from a “CT2” set [13]. The CT2 is a recent, large set of color images, which was used in the research on reversible color space transforms [9,14]. It contains 746 images taken from different sources; image sizes vary from 180×117 to 6600×5100 . The set was divided into subsets: Photo with 499 photographs and No-photo with 247 non-photographic, computer-generated, or mixed-content images. The latter were further divided into images that were better compressed by JPEG 2000 without DWT (No-photo (a) – 81 images) and with the unmodified DWT (No-photo (b) – 166 images). The Photo images were not divided this way since for only one Photo image it would be better to skip the DWT stage of JPEG 2000.
- SP001_GS2, SP002_GS2, SP005_GS2, SP010_GS2, and SP020_GS2 – These are sets of GS2 images with salt-and-pepper impulse noise added; 1%, 2%, 5%, 10%, and 20% of image pixels, respectively, were replaced by black or white pixels (Fig. 5).
- WG002_GS2, WG005_GS2, WG010_GS2, WG020_GS2, and WG050_GS2 – These are sets of GS2 images with added white Gaussian noise with standard deviations 2, 5, 10, 20, and 50, respectively (Fig. 5).

3.3. Experimental procedure and implementations

We used the IRIS-JP3D JPEG 2000 part 10 (JP3D) [15] reference software developed by Tim Bruylants from Vrije Universiteit Brussel (VUB) and the Interdisciplinary Institute for BroadBand Technology (IBBT), version 1.1.1 [16], which is downward compatible with the basic JPEG 2000 standard. In this codec, it was relatively easy to implement RDLS. In the experiments, except for replacing JPEG 2000 lifting steps with RDLS and setting the transform level, we




Fig. 3. Effects of filtering a sample image (top left panel); image sizes were 32x30 pixels, and cw is the weight of the window center point of the Smoothing filter.

used the default codec settings. Particularly, a whole image was compressed as a single tile. As the 3-level DWT transform results in bitrates close to the default 5-level, we used the 3-level decomposition. To switch off the DWT transform stage in JPEG 2000 coding, we invoked the codec with the 0-level DWT setting.

For comparison, we also used JPEG-LS [17,18] and HEVC (H.265) [19,20] compression algorithms. JPEG-LS is a standard of the JPEG committee for primarily lossless compression of still images, and it originates from the LOCO-I algorithm. We used the SPMG/UBC JPEG-LS implementation, version 2.2 [21], with default codec settings. HEVC is the most recent video compression standard developed by MPEG and VCEG

committees, and it allows lossless compression of individual still images. We used the HEVC Test Model (HM) reference software, version 16.6 [22], with the following set of encoder options: `InputChromaFormat=400`, `FrameRate=1`, `FramesToBeEncoded=1`, `Profile=monochrome`, `IntraPeriod=1`, `QuadtreeTULog2MaxSize=5`, `TransquantBypassEnableFlag=1`, `CUTransquantBypassFlagForce=1`, and `ConformanceWindowMode=1`.

Groups of noisy images were created using `pgmtool`, version 0.72 [11].

The compression ratio, or bitrate r , expressed in bits per pixel (bpp) is calculated as $r=8e/p$, where p is the number of pixels in the image, and e is the total size in Bytes of the

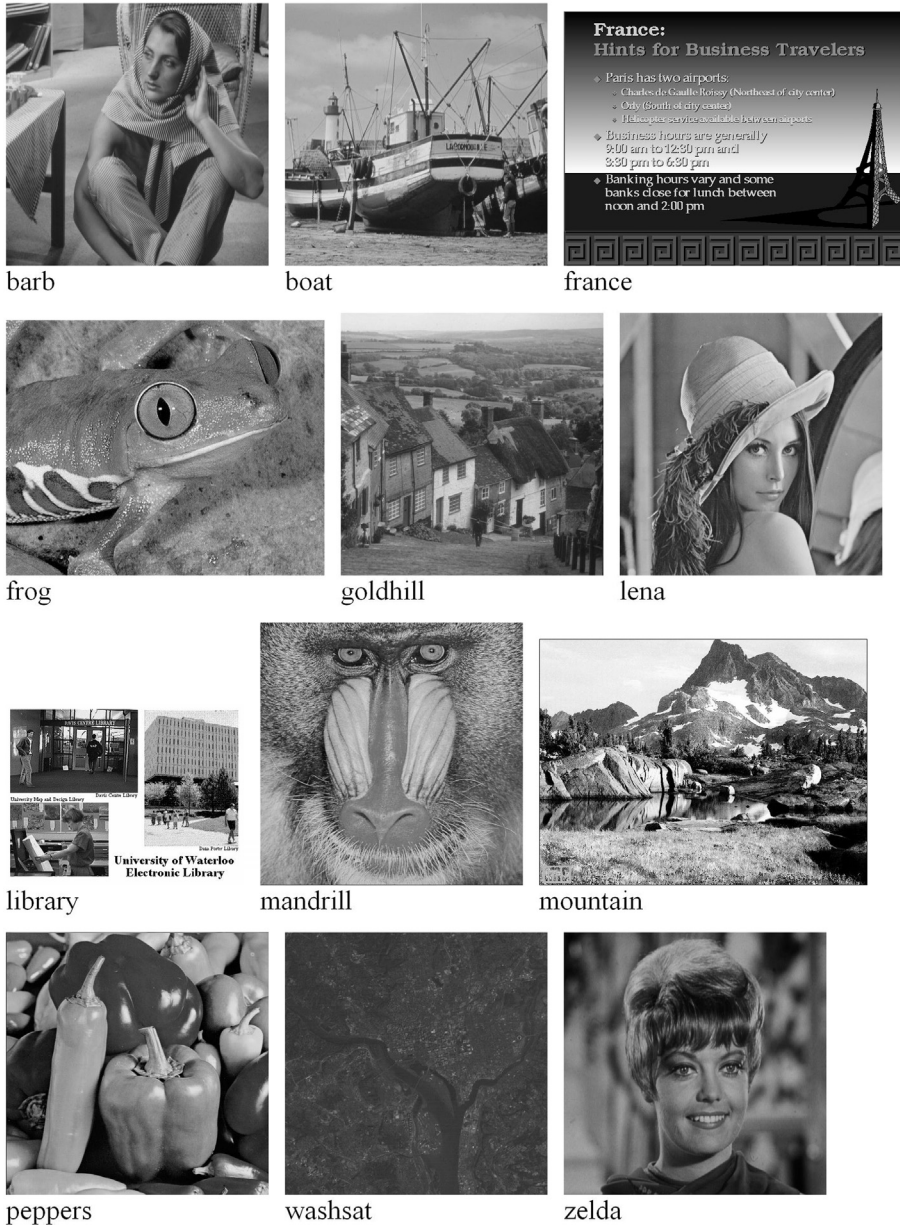


Fig. 4. The GS2 set.

compressed image, including the compressed file format header. Hence, smaller r means better compression.

We introduced modifications to JPEG 2000, and then we evaluated the modification effects by analyzing bitrate changes with respect to the bitrate of the reference method, that is, of unmodified JPEG 2000, expressed in percentages of the reference method bitrate. Due to the number of test-sets and to the size of the greatest one (CT2g), except for initial observations using GS2 set, we report averaged bitrate changes for a set or for its specific subset rather than results for individual images. Instead of averaged bitrate changes, we could use other common measures. However, the 2 popular ones described below seem inadequate for CT2g images. We could measure how many times the bitrate obtained using

the proposed method would be smaller than the bitrate of the reference method. However, the results would be biased toward effects obtained for non-photographic images that were better compressed without the DWT stage, as the bitrates of some of them were 2 to nearly 5 times smaller than the reference method bitrates (which we express as 50% to 80% improvement). Instead of calculating the average of bitrate changes, the average absolute bitrate could be calculated and then compared to the reference method average bitrate. This would result in virtually ignoring the effects obtained for images compressed by the unmodified JPEG 2000 really well and would result in a bias toward effects on poorly compressible images. The unmodified JPEG 2000 bitrate of many non-photographic images was below 1 bpp;




Fig. 5. Fragment (128x128 pixels) of the “lena” image from the GS2 set with added white Gaussian noise (middle row, wg – noise standard deviation) and with added salt-and-pepper impulse noise (bottom row, sp – percentage of image pixels replaced by black or white pixels).

improving this by 50% would affect the averaged bitrate less than improving the bitrate of the image that was compressed to 6 bpp by 10%.

4. Results and discussion

4.1. Initial observations

When the 3 iterations of step B and all the filters presented in Section 3.1 were used, the filter selection heuristic time complexity was roughly 235 times greater than the complexity of JPEG 2000 with subband denoising. Therefore, we initially focused on RDLS effects for the small GS2 set and also analyzed the results for individual images as well as for GS2 images with impulse and Gaussian noise added. Apart from observing RDLS effects on lossless JPEG 2000 bitrates of various types of images, the aim was to check if all the filters are necessary and if the number of iterations may be lowered, thus speeding up further evaluation using the large CT2g set. Results for individual GS2 images as well as averaged results for GS2 set and for sets of noisy images are presented in

Table 1. We report lossless JPEG 2000 bitrates obtained for the 3-level unmodified DWT and bitrate changes relative to it obtained for no DWT and for RDLS–DWT with filters selected by up to 3 iterations of step B of the heuristic. For 0 iterations, that is, applying the same filter for all subbands, in case of individual images, we also report the selected filter.

Let’s focus on cases when the RDLS bitrate improvement may be practically useful. Since the filter selection heuristic allows only the bitrate improvement, and we start with a large number of filters, a small bitrate improvement might be to some extent accidental. We assume that the bitrate improvement below 1% is unjustified in practical applications when obtained at the cost of a large increase in compression process complexity. Fig. 6 presents cases from Table 1 when the improvement exceeds 1%. For the GS2 images, the improvement exceeds 1% for 4 images (actually, RDLS–DWT is for them at least 3.18% better than DWT). These images are not continuous tone photographic images—“france” is a computer slide containing sharp lines and filled with gradient, “frog” is a dithered photo, “library” is composed of images and text, and finally, “washsat” is a satellite image of a highly sparse histogram [23]. Other images seem

photographic, except for “mountain,” which is probably a dithered photo. For images with Gaussian noise added, the improvement is small, which for photographic images and small noise standard deviation could be expected from the GS2 results, since unaltered natural images contain such noise [24]. Interestingly, adding Gaussian noise to images, bitrates of which were significantly improved by RDLS, considerably decreases RDLS improvements—only for “france” for the 2 lowest levels of noise added and for “library” with the lowest level of noise the RDLS improvement exceeds 1%. Probably, the nonlinear denoising filters used ceased to be effective in the presence of Gaussian noise. We see that all groups of impulse noise images (actually, all such images) are compressed significantly better with RDLS. Section 4.3 explains the poor DWT performance on images containing impulse noise or sharp lines. Except for the highest noise level, the RDLS improvement is generally greater when the impulse noise level is higher, while for Gaussian noise, no such dependency can be observed. Moreover, for some GS2 images, it is better not to use DWT; in such cases, RDLS–DWT results in bitrates between bitrates of JPEG 2000 without DWT and with DWT. There are only 3 such images, and the opposite case may be observed for individual images with impulse noise added. If skipping DWT improves the JPEG 2000 bitrate of these images, then in most cases, applying RDLS–DWT results in greater improvement. We take advantage of these observations in Sections 4.4 and 4.5.

4.2. Selection of heuristic parameters

Table 1 shows for individual images, that if a single filter, used for all subbands in RDLS–DWT, significantly improves the GS2 image bitrate, then it is the Median filter. However,

greater improvements were obtained when different filters were applied in different RDLS–DWT steps. To reduce the heuristic complexity while preserving high bitrate improvements, for GS2 and noisy images we analyzed which filters were used when the RDLS bitrate improvement was over 1%. The 4 most frequently used were Median 5×5 (for 49.7% of all RDLS–DWT steps), Median 3×3 (21.3%), RCRS-2 3×3 (18.7%), and RCRS-1 3×3 (4.8%). It is noteworthy that the Median 5×5 filter was used in nearly three-fourths of the RDLS–DWT update steps (74.3%). For the remainder of this study, we decided to keep these 4 filters along with the None filter. Even though the None filter was rarely used (0.2%), the heuristic could result in bitrate worsening without it. Smoothing filters were used in 0.0% to 1.1% of the RDLS–DWT steps, RCRS-1 5×5 in 0.2% of steps, RCRS-2 5×5 in 2.6%. We also limited the number of heuristic step B iterations to 2. Although compared to a single iteration, only the “france” image from the GS2 set is noticeably better compressed when we allow 2 iterations, more such cases may happen for the CT2g set with over 700 images. The heuristic complexity for 2 iterations and 5 filters is roughly 37 times greater than the complexity of JPEG 2000 with subband denoising. Table 2 presents the results for the reduced set of filters and 2 iterations.

In cases when the RDLS bitrate improvement obtained using the set of 19 filters is high, the reduction of the filter set to 5 filters has almost no influence on RDLS performance. Both for such individual GS2 images and for averaged improvements for groups of images with impulse noise added, due to simplifying the heuristic, the RDLS improvement is worsened by no more than 0.05 percentage points. The negligible performance drop when reducing the number of filters to 5 suggests that further reduction of the filter set

Table 1

RDLS effects on lossless JPEG 2000 bitrates of individual GS2 images and averaged effects for GS2 set and sets of noisy images.

Image(s)	r_{DWT}	Δr_{NO-DWT} (%)	Step A filter	$\Delta r_{RDLS(19,0)}$ (%)	$\Delta r_{RDLS(19,1)}$ (%)	$\Delta r_{RDLS(19,2)}$ (%)	$\Delta r_{RDLS(19,3)}$ (%)
barb	4.6586	19.59	None	0.00	−0.25	−0.25	−0.25
boat	4.4041	17.39	None	0.00	−0.09	−0.09	−0.10
france	2.0294	41.30	Median 3×3	−7.04	−23.32	−25.30	−25.30
frog	6.2546	−17.35	Median 5×5	−2.87	−3.18	−3.18	−3.18
goldhill	4.8338	12.05	Smoothing $3 \times 3, 2^{10}$	−0.02	−0.13	−0.13	−0.13
lena	4.3137	17.65	Smoothing $3 \times 3, 2^{10}$	0.00	−0.09	−0.09	−0.09
library	5.6892	−8.90	Median 5×5	−3.87	−5.10	−5.11	−5.11
mandrill	6.1075	3.71	None	0.00	−0.11	−0.12	−0.12
mountain	6.6983	−5.62	Smoothing $3 \times 3, 2^{10}$	−0.02	−0.10	−0.10	−0.10
peppers	4.6158	15.74	None	0.00	−0.07	−0.08	−0.08
washsat	4.4308	6.77	Median 5×5	−0.87	−5.89	−5.90	−5.90
zelda	3.9951	26.99	Smoothing $3 \times 3, 2^{10}$	−0.01	−0.29	−0.29	−0.29
GS2	4.8359	10.78		−1.22	−3.22	−3.39	−3.39
SP001_GS2	5.4584	−1.31		−9.71	−11.22	−11.26	−11.26
SP002_GS2	5.8882	−6.93		−13.74	−15.03	−15.13	−15.13
SP005_GS2	6.7766	−14.73		−18.95	−20.57	−20.63	−20.63
SP010_GS2	7.5945	−18.44		−21.83	−22.77	−22.78	−22.79
SP020_GS2	8.3557	−18.58		−20.47	−21.11	−21.13	−21.13
WG002_GS2	5.2614	8.43		−0.26	−0.55	−0.55	−0.56
WG005_GS2	5.7360	5.03		−0.09	−0.19	−0.19	−0.19
WG010_GS2	6.3022	2.17		−0.03	−0.12	−0.12	−0.13
WG020_GS2	7.0183	0.39		−0.04	−0.14	−0.15	−0.15
WG050_GS2	8.0824	−1.90		−0.31	−0.35	−0.35	−0.35

r_{DWT} – JPEG 2000 bitrate for 3-level unmodified DWT, Δr_{NO-DWT} – bitrate change obtained by skipping the DWT stage of JPEG 2000, $\Delta r_{RDLS(19,i)}$ – bitrate change obtained by using RDLS–DWT and all denoising filters presented in Section 3.1, adaptively selected in i iterations of heuristic step B. For $\Delta r_{RDLS(19,0)}$, in case of individual images we report the filter selected in step A of the heuristic.

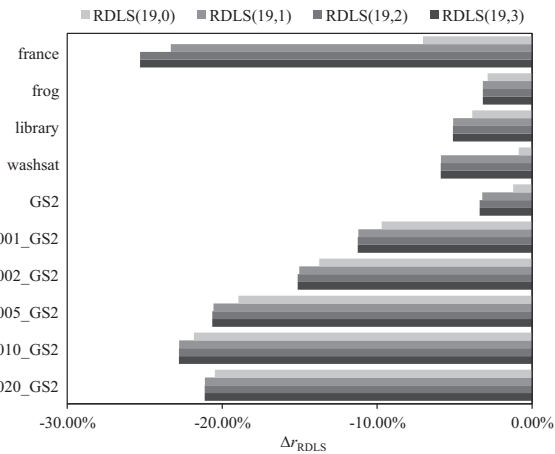


Fig. 6. RDLS effects on lossless JPEG 2000 bitrates of individual GS2 images and averaged effects for the GS2 set and sets of noisy images in cases when the bitrate improvement exceeds 1%. Δr_{RDLS} – bitrate change or average bitrate change due to RDLS as compared to lossless JPEG 2000 bitrate for 3-level unmodified DWT.

may result in still good RDLS bitrate improvements. Therefore, in subsection 4.5, we additionally consider using only the None filter and the Median 5×5 filter. Interestingly, in one case of the SP005_GS2 images, we obtained greater bitrate improvement (by 0.01 percentage points), which shows that heuristic results are not optimal.

4.3. RDLS effect on impulse noise and sharp lines

Fig. 7 shows that in an image with impulse noise added, a single noisy pixel results in the distortion of all subbands of the DWT transformed image. Furthermore, the noisy pixel's neighborhood, originally unaffected by noise, is now altered in all or most subbands, depending on the noisy pixel position in the original image.

The neighborhood is affected since both 1D DWT lifting steps—prediction (Eq. (1)) and update (Eq. (2))—alter the sample based on its neighborhood of opposite parity. Since the update is performed after the prediction, regardless of whether the noisy sample is even or odd (and will end up as part of L or H subband, respectively), both subbands are contaminated. JPEG 2000 encodes subbands independently, so the information on noise appearance has to be encoded 4 times for the 1-level 2D DWT or more for higher levels, increasing the bitrate. Sharp lines affect the DWT similarly, resulting in altered line neighborhoods due to the way 2D DWT is built by using 1D DWTs. For instance, the horizontal line in an image affects the 1D DWT done for the image column the same way as the single impulse noise pixel. The DWT is generally expected to decompose an image into subbands that contain low- and high-frequency image features. For impulse noise and sharp lines, the low- and high-frequency image features are duplicated rather than separated into subbands.

The RDLS–DWT also fails to decompose low- and high-frequency features (Fig. 7, right-hand column). Noisy pixels during prediction and update are just left where they were, and then during the reorder step of transform, they are placed in subbands that correspond to their parities.

However, due to denoising, it almost completely avoids generating the unnecessary low-frequency information on noise and propagating the noise into all subbands, allowing improvement of the bitrate.

4.4. RDLS effect on bitrates of typical images

Table 3 reports lossless JPEG 2000 bitrates obtained for the 3-level unmodified DWT and bitrate changes relative to it obtained for no DWT and for RDLS–DWT with a reduced set of 5 filters selected by up to 2 iterations of step B of the heuristic. The results are reported for the CT2g set and its subsets. Fig. 8 presents the results for individual images.

The overall lossless JPEG 2000 bitrate improvement of CT2g images due to RDLS–DWT is significant (roughly 4.8%); it mainly results from RDLS effects on bitrates of non-photographic images, as photographic image bitrates are improved by less than 1% on average. The bitrate improvement for non-photographic images is 13.29% on average, and greater improvements are observed for images that are better compressed with the skipped DWT stage of JPEG 2000. In such cases, RDLS–DWT results in average bitrates that are better than those of JPEG 2000 compression with both unmodified DWT and without DWT.

A single iteration of step B of the parameter selection heuristic results in bitrate improvements close to the improvements obtained with 2 iterations (worse by up to 0.12 percentage points for one of the subsets) and much better than those obtained after step A only (by up to 4.48 percentage points). The increased complexity of parameter selection by using 2 iterations seems practically unjustified: for 2 iterations, it is roughly 37 times greater than the complexity of JPEG 2000 with subband denoising, while for 1 iteration, it is 21 times. The initial selection of denoising filters was arbitrary; the number of filters was reduced based on the results obtained for other images. We could probably obtain larger improvements without such increase in complexity by finding better denoising filters.

The definitely best denoising filter among the examined ones is Median 5×5 , which outperformed more sophisticated filters and the Median filter with a smaller window. It is the strongest among the nonlinear filters used, which suggests that a stronger denoising might further improve bitrates. Note also that we did not test larger windows or windows of different shapes. Since when computing H and L subbands during 2D DWT by applying 1D DWT to image columns, we perform denoising based on samples of a specific parity only in a vertical direction and on all samples in a horizontal direction; the actual area corresponding to the denoising filter window in the original image is rectangular. Since the correlation of pixels is proportional to the distance between them, the window should instead be circular or square.

4.5. Practical observations

Up to this point, we evaluated RDLS effects on DWT transform in lossless JPEG 2000 coding. In practice, we would be interested in the bitrate achievable by an algorithm. To improve it, we may consider the observation that for some

Table 2

Effects of RDLS with the reduced set of filters on lossless JPEG 2000 bitrates of individual GS2 images and averaged effects for GS2 set and sets of noisy images.

Image(s)	r_{DWT}	$\Delta r_{RDLS(19,3)}$ (%)	$\Delta r_{RDLS(5,2)}$ (%)
Barb	4.6586	-0.25	-0.21
Boat	4.4041	-0.10	0.00
France	2.0294	-25.30	-25.30
Frog	6.2546	-3.18	-3.18
Goldhill	4.8338	-0.13	-0.02
Lena	4.3137	-0.09	-0.01
Library	5.6892	-5.11	-5.12
Mandrill	6.1075	-0.12	-0.06
Mountain	6.6983	-0.10	-0.02
Peppers	4.6158	-0.08	-0.01
Washesat	4.4308	-5.90	-5.90
Zelda	3.9951	-0.29	-0.13
GS2	4.8359	-3.39	-3.33
SP001_GS2	5.4584	-11.26	-11.21
SP002_GS2	5.8882	-15.13	-15.13
SP005_GS2	6.7766	-20.63	-20.65
SP010_GS2	7.5945	-22.79	-22.79
SP020_GS2	8.3557	-21.13	-21.13
WG002_GS2	5.2614	-0.56	-0.50
WG005_GS2	5.7360	-0.19	-0.12
WG010_GS2	6.3022	-0.13	-0.04
WG020_GS2	7.0183	-0.15	-0.02
WG050_GS2	8.0824	-0.35	-0.09

r_{DWT} – JPEG 2000 bitrate for 3-level unmodified DWT, $\Delta r_{RDLS(f,i)}$ – bitrate change or average bitrate change obtained by using RDLS–DWT and f denoising filters, adaptively selected in i iterations of heuristic step B out of ($f=19$) all filters from Section 3.1 or ($f=5$) Median 5×5 , Median 3×3 , RCRS-2 3×3 , RCRS-1 3×3 , and None.

individual images, it is better to skip the DWT or RDLS–DWT stage of JPEG 2000. The cost of the bitrate improvement is also important. Therefore, we tested a couple more variants of compression, including allowing the DWT stage of JPEG 2000 to be skipped at the cost of a single extra execution of the compression process, using the most frequently selected Median 5×5 filter and None filter only, and replacing the heuristic with a fixed assignment of denoising filters to RDLS steps. In the latter case, we used the Median 5×5 filter for update steps and the None filter for prediction steps. Table 4 presents the results.

The application of RDLS to DWT transform in lossless JPEG 2000 coding results in significant bitrate improvements for non-photographic images. In most cases, these improvements are obtained at the cost of image-adaptive selecting of denoising filters using the heuristic. The only extra cost during decompression is due to denoising.

First, let's examine the results obtained without using the filter selection heuristic (last 3 rows in Table 4). Using the Median 5×5 filter for all RDLS-modified DWT update steps and leaving prediction lifting steps unmodified (row Δr_{RDLS_FIXED}), we improved the bitrates of non-photographic images by 8.64% on average. Since some images were better compressed when the RDLS–DWT stage was skipped, at the cost of an additional execution of the compression process, we may obtain greater improvement of 11.64% for non-photographic images ($\Delta r_{RDLS_FIXED, NO-DWT}$). However, this improvement should be compared to another improvement, which is obtained at the same cost of a double compression process execution. When for each image we pick a better

bitrate out of those obtained by using the unmodified DWT and by skipping the DWT stage ($\Delta r_{NO-DWT, DWT}$), the average improvement is 7.43%. Both of these cases of RDLS application worsen the bitrates of photographic images by about 0.5%. We may avoid the bitrate increase either by employing the filter selection heuristic that allows selecting the None filter or by directly checking the unmodified DWT bitrate. In the latter case, at the cost of 3 executions ($\Delta r_{RDLS_FIXED, NO-DWT, DWT}$), we obtain a compression scheme that is useful for practical applications. It significantly improves the bitrates of non-photographic images by almost 12% on average and does not worsen the bitrate of any image. In [3,9], it was shown that by using the entropy of prediction error of a transformed color image component, we might efficiently select the color space transform or the denoising filter for RDLS-modified color space transform without executing the compression process and independently of the actual image compression algorithm (JPEG 2000, JPEG-LS, and HD Photo/JPEG-XR [25,26] were used in these studies). In [9], only a small subset of image pixels (up to 10000) was found sufficient for close to optimum transform selection. Finding a similar estimator of DWT variants, including RDLS–DWT, unmodified DWT, and skipped DWT is an interesting field of future research.

The image-adaptive selection of denoising filter for each RDLS–DWT step allows greater bitrate improvements at a greater cost. Since for non-photographic images, supplementing the RDLS denoising filter selection heuristic with checking the effect of the compression with the skipped DWT stage results in noticeably greater improvements (by over 1 percentage point), we focus on results obtained this way. Using only the Median 5×5 denoising filter and the None filter at the cost of more than two times greater complexity ($\Delta r_{RDLS(2,1), NO-DWT}$), we obtain an average improvement of 14%, i.e., better by over 2 percentage points, than without using the heuristic ($\Delta r_{RDLS_FIXED, NO-DWT, DWT}$). The improvement is mainly due to larger improvements for images that are better compressed with the skipped DWT. Adding 3 adaptively selected denoising filters further increases the complexity three-fold, but the bitrate is further improved by below 0.5 percentage points only ($\Delta r_{RDLS(5,1), NO-DWT}$). Such cost may be too high for practical applications, but a fast estimation of RDLS–DWT effects could make the adaptive application of many filters useful.

4.6. Comparison to other algorithms

Table 5 compares average bitrate changes relative to bitrates of unmodified lossless JPEG 2000 obtained by one of the practical variants that exploit RDLS discussed in the previous subsection (RDLS(2,1), NO-DWT) and obtained by JPEG-LS and HEVC algorithms. It can be seen that JPEG-LS is the best for the CT2g set as well as for its subsets. The recent HEVC is in-between JPEG-LS and our RDLS variant for non-photographic images. By applying RDLS to DWT in lossless JPEG 2000 and by allowing to skip the DWT stage we obtained a compression scheme, results of which for No-photo images are closer to JPEG-LS than to unmodified JPEG 2000. However, JPEG-LS is better by roughly 7.5 percentage points, indicating that the further improvement of lossless JPEG 2000 bitrates should be possible in case of




Fig. 7. Effects of impulse noise and sharp lines on 1-level 2D DWT. Left-hand column – original images (32x32 pixels), middle column – DWT transformed images, right-hand column – RDLS-DWT transformed images (Median 3x3 filter used); transformed images are normalized to dynamic range of original ones.

Table 3

Averaged RDLS effects on lossless JPEG 2000 bitrates of CT2g images.

Images	Count	r_{DWT}	Δr_{NO-DWT} (%)	$\Delta r_{RDLS(5,0)}$ (%)	$\Delta r_{RDLS(5,1)}$ (%)	$\Delta r_{RDLS(5,2)}$ (%)
Photo	499	3.9975	25.6	-0.02	-0.62	-0.64
No-photo	247	2.9162	18.8	-9.09	-13.18	-13.29
No-photo (a)	81	3.2882	-22.6	-23.89	-28.37	-28.45
No-photo (b)	166	2.7348	39.0	-1.87	-5.77	-5.89
All	746	3.6395	23.3	-3.02	-4.78	-4.83

r_{DWT} – JPEG 2000 bitrate for 3-level unmodified DWT; Δr_{NO-DWT} – bitrate change obtained by skipping the DWT stage of JPEG 2000; $\Delta r_{RDLS(5,i)}$ – bitrate change obtained by using RDLS-DWT and 5 denoising filters (Median 5 × 5, Median 3 × 3, RCRS-2 3 × 3, RCRS-1 3 × 3, and None), adaptively selected in i iterations of step B of the heuristic.

non-photographic images. The improvement of the RDLS variant is much greater for No-photo (a) images than for No-photo (b), both as percentage of unmodified JPEG 2000 bitrate and as compared to improvement obtained by JPEG-LS. Interesting observations may be made for Photo images. Here we obtain only a small fraction of improvement possible by use of JPEG-LS. Generally, our method gives only small improvements for photographic images, which probably could be improved, since we were selecting filters for non-photographic and noisy images. It may be also noted, that HEVC bitrates are for Photo images significantly worse than those of unmodified JPEG 2000.

The HEVC settings we applied might not be optimal; however, in [27] for lossless coding of video sequences, the HEVC bitrates were worse than bitrates of JPEG-LS by about 11%, which complies with our results.

5. Conclusions and future work

In this study, we applied RDLS, which is basically the integration of lifting steps with denoising filters, to DWT in JPEG 2000 lossless coding and evaluated RDLS effects on bitrates of various types of images. The RDLS application is




Fig. 8. RDLS effects on lossless JPEG 2000 bitrates of CT2g images. r_{DWT} – JPEG 2000 bitrate for 3-level unmodified DWT, $\Delta r_{RDLS(5,2)}$ – bitrate change obtained by using RDLS–DWT with 5 denoising filters (Median 5×5 , Median 3×3 , RCRS-2 3×3 , RCRS-1 3×3 , and None), adaptively selected by 2 iterations of step B of the heuristic. Top panel – photographic images; bottom panel – non-photographic images.

Table 4
Averaged effects of various compression variants on lossless JPEG 2000 bitrates of CT2g images.

Compression variants	Complexity	Photo	No-photo	No-photo (a)	No-photo (b)	All
r_{DWT}	1	3.9975	2.9162	3.2882	2.7348	3.6395
Δr_{NO-DWT}	1	25.59%	18.77%	-22.65%	38.98%	23.33%
$\Delta r_{NO-DWT, DWT}$	2	-0.01%	-7.43%	-22.65%	0.00%	-2.47%
$\Delta r_{RDLS(5,0)}$	5	-0.02%	-9.09%	-23.89%	-1.87%	-3.02%
$\Delta r_{RDLS(5,0), NO-DWT}$	6	-0.03%	-11.11%	-30.05%	-1.87%	-3.70%
$\Delta r_{RDLS(5,1)}$	21	-0.62%	-13.18%	-28.37%	-5.77%	-4.78%
$\Delta r_{RDLS(5,1), NO-DWT}$	22	-0.63%	-14.34%	-31.90%	-5.77%	-5.17%
$\Delta r_{RDLS(5,2)}$	37	-0.64%	-13.29%	-28.45%	-5.89%	-4.83%
$\Delta r_{RDLS(5,2), NO-DWT}$	38	-0.65%	-14.43%	-31.92%	-5.89%	-5.21%
$\Delta r_{RDLS(2,1)}$	6	-0.56%	-12.61%	-27.13%	-5.53%	-4.55%
$\Delta r_{RDLS(2,1), NO-DWT}$	7	-0.57%	-14.00%	-31.37%	-5.53%	-5.02%
Δr_{RDLS_FIXED}	1	0.53%	-8.64%	-17.29%	-4.42%	-2.51%
$\Delta r_{RDLS_FIXED, NO-DWT}$	2	0.52%	-11.64%	-26.42%	-4.42%	-3.51%
$\Delta r_{RDLS_FIXED, NO-DWT, DWT}$	3	-0.31%	-11.90%	-26.42%	-4.81%	-4.14%

r_{DWT} – lossless JPEG 2000 bitrate for 3-level unmodified DWT; $\Delta r_{variant_list}$ – bitrate change obtained by picking for each image the best bitrate out of bitrates obtained using the listed variants; DWT – unmodified DWT transform; NO-DWT – skipping the DWT stage of JPEG 2000; $RDLS(f,i)$ – RDLS–DWT using f denoising filters adaptively selected in i iterations of step B of the heuristic out of ($f=5$) Median 5×5 , Median 3×3 , RCRS-2 3×3 , RCRS-1 3×3 , and None or ($f=2$) Median 5×5 and None; $RDLS_FIXED$ – RDLS–DWT with fixed assignment of denoising filters to RDLS steps (Median 5×5 for update steps and None for prediction steps); Complexity – computational time complexity as roughly compared to the complexity of the JPEG 2000 compression process which may involve subband denoising or DWT stage skipping.

straightforward, as there is an analogy between the DWT of a grayscale image and a transform for which the RDLS technique was originally proposed—the color space transform of a color image. However, assuming that the DWT subbands’ characteristics are invariant, we must find proper denoising filters for each RDLS-modified DWT (RDLS–DWT) subband, that is, $6t$ filters for t decomposition levels. Subbands are interdependent, so we proposed a filter selection heuristic consisting of greedy steps A and B, where step B may be repeated given a number of iterations. The heuristic requires multiple executions of the JPEG 2000 compression process that may involve subband denoising.

In the experiments, we used simple low-pass linear averaging filters (Smoothing filters) that were effective in the RDLS-modified color space transform and nonlinear filters (including the Median filter) that were effective in removing impulse noise; 3×3 and 5×5 sample windows were used for both types of filters. First, for classic grayscale images (Waterloo GraySet2) and for images with impulse or Gaussian noise added, we found that practically useful bitrate improvements with RDLS were obtained for non-photographic images and for images with impulse noise added. In these cases, reducing the set of denoising filters to the 4 filters most frequently selected by the heuristic along with the None filter that turns RDLS into a regular lifting step and using 2 iterations of step B of the heuristic has almost no influence on bitrate improvements due to RDLS as compared to using all 19 filters and a greater number of iterations. All of the 4 most frequently selected filters are nonlinear, and the most frequently selected filter is the Median 5×5 filter. We also observed that some images are better compressed by JPEG 2000 if we skip the DWT stage of the algorithm; the unmodified DWT increases the amount of information that has to be

encoded when an image contains impulse noise or sharp lines.

Next, we evaluated RDLS effects on DWT by using a recent, large, and diverse set of photographic and non-photographic images (CT2g). Significant bitrate improvements due to RDLS were also observed for the latter; an average bitrate improvement of over 13% was obtained at the cost roughly 21 times greater than the cost of JPEG 2000 compression process with subband denoising. However, it is noteworthy that for about one-third of these images, the DWT transform worsens the lossless JPEG 2000 bitrate, so the bitrate may be improved by simply skipping the DWT stage. In such cases, RDLS–DWT usually results in average bitrates that are better than those of JPEG 2000 compression with both unmodified DWT and without DWT, but for some individual images, skipping DWT gives the best results. Therefore, in practice, including the possibility of skipping the DWT stage is justified.

Finally, from a practical viewpoint, we evaluated a couple of compression schemes by exploiting RDLS and also the possibility of skipping the DWT stage or using a fixed assignment of denoising filters to RDLS–DWT subbands, instead of time-consuming heuristic. A simple scheme based on a fixed filter assignment, at the cost of a triple compression process execution, significantly improves the bitrates of non-photographic images by almost 12% on average and does not worsen the bitrate of any image. Image-adaptive filter selection by using the heuristic results in an average improvement of 14% and with over two times greater cost. In both cases, only the Median 5×5 and None filters were used. Further bitrate improvements are possible by using more filters but at a cost that seems too high for practical applications. As compared to JPEG-LS, the proposed schemes for non-photographic images allow to obtain bitrates that are closer to JPEG-LS than to unmodified JPEG 2000; however, the better performance of the former indicates that the further improvement of lossless JPEG 2000 bitrates should be possible.

Bitrate improvements were obtained at the cost of multiple executions of the compression process and, in most cases, of multiple denoising of RDLS–DWT subbands. Finding a fast estimator of denoising filter selection effects to be used within or instead of filter selection heuristic is an interesting field of future research. Another field involves finding better denoising filters, as we found reasons for why the filter set used in this study might be not optimal. Some images are better compressed when the DWT/RDLS–DWT compression stage is

skipped, but similar to the denoise filtering that is most effective when applied to some subbands only, the optimum may be in-between skipping and applying the transform. Therefore, we are currently investigating the image-adaptive, partial transform skipping.

Conflict of interest statement

None.

Role of the funding source

The study sponsors were not involved in the study design, in the collection, analysis and interpretation of data, in the writing of the manuscript, or in the decision to submit the manuscript for publication.

Acknowledgments

We thank Professor Tilo Strutz for providing the classification of the CT2 set into photographic and non-photographic images and the anonymous reviewers for their constructive comments. This work was supported by the BK-266/RAU2/2014 Grant from the Institute of Informatics, Silesian University of Technology and by the POIG.02.03.01-24-099/13 Grant: GeCONil – Upper-Silesian Center for Scientific Computations.

References

- [1] I. Daubechies, W. Sweldens, Factoring wavelet transforms into lifting steps, *J. Fourier Anal. Appl.* 4 (1998) 247–269, <http://dx.doi.org/10.1007/BF02476026>.
- [2] H.S. Malvar, G.J. Sullivan, S. Srinivasan, Lifting-based reversible color transformations for image compression, in: *Proceedings of SPIE, Applications of Digital Image Processing XXXI*, 7073 (2008) 707307, <http://dx.doi.org/10.1117/12.797091>.
- [3] R. Starosolski, Reversible denoising and lifting based color component transformation for lossless image compression, arXiv:1508.06106 [cs.MM], available: (<http://arxiv.org/abs/1508.06106>), (accessed 28.08.15).
- [4] Information Technology – JPEG 2000 image coding system: core coding system, ISO/IEC International Standard 15444-1 and ITU-T recommendation T.800, 2004.
- [5] T. Bruylants, A. Munteanu, P. Schelkens, Wavelet based volumetric medical image compression, *Signal Process. Image Commun.* 31 (2015) 112–133, <http://dx.doi.org/10.1016/j.image.2014.12.007>.

Table 5

Comparison of RDLS-modified JPEG 2000 with JPEG-LS and HEVC.

Compression methods	Photo	No-photo	No-photo (a)	No-photo (b)	All
r_{DWT}	3.9975	2.9162	3.2882	2.7348	3.6395
$\Delta r_{RDLS(2,1), NO-DWT}$	–0.57%	–14.00%	–31.37%	–5.53%	–5.02%
$\Delta r_{JPEG-LS}$	–2.72%	–21.46%	–38.27%	–13.26%	–8.92%
Δr_{HEVC}	9.32%	–17.04%	–36.48%	–7.55%	0.59%

r_{DWT} – lossless JPEG 2000 bitrate for 3-level unmodified DWT; $\Delta r_{RDLS(2,1), NO-DWT}$ – average bitrate change obtained by picking for each image the best bitrate out of bitrates obtained by skipping the DWT stage of JPEG 2000 and by RDLS–DWT using Median 5×5 and None denoising filters adaptively selected in single iteration of heuristic step B; $\Delta r_{JPEG-LS}$ – average bitrate change obtained by using JPEG-LS instead of JPEG 2000; Δr_{HEVC} – average bitrate change obtained by using HEVC.

- [6] C. Christopoulos, A. Skodras, T. Ebrahimi, The JPEG2000 still image coding system: an overview, *IEEE Trans. Consum. Electron.* 46 (4) (2000) 1103–1127, <http://dx.doi.org/10.1109/30.920468>.
- [7] S.G. Mallat, A theory for multiresolution signal decomposition: the wavelet representation, *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (1989) 674–693, <http://dx.doi.org/10.1109/34.192463>.
- [8] R. Starosolski, New simple and efficient color space transformations for lossless image compression, *J. Vis. Commun. Image Represent.* 25 (5) (2014) 1056–1063, <http://dx.doi.org/10.1016/j.jvcir.2014.03.003>.
- [9] T. Strutz, Multiplierless reversible colour transforms and their automatic selection for image data compression, *IEEE Trans. Circuits Syst. Video Technol.* 23 (7) (2013) 1249–1259, <http://dx.doi.org/10.1109/TCSVT.2013.2242612>.
- [10] R.C. Hardie, K.E. Barner, Rank conditioned rank selection filters for signal restoration, *IEEE Trans. Image Process.* 3 (2) (1994) 192–206, <http://dx.doi.org/10.1109/83.277900>.
- [11] Free grayscale image transformation tool “pgmtool”, version 0.72, available: (<http://sun.aei.polsl.pl/~rstaros/imgtransf/pgmtool/index.html>), (accessed 28.07.15).
- [12] Set “Grayscale Set 2” of grayscale images from the University of Waterloo, Fractal Coding and Analysis Group, available: (<http://links.uwaterloo.ca/Repository.html>), (accessed 20.05.15).
- [13] Set “CT2” of color images from the Deutsche Telekom, Hochschule für Telekommunikation Leipzig, Institute of Communications Engineering, available: (<http://www1.hft-leipzig.de/strutz/Papers/Testimages/CT2/>), (accessed 20.05.15).
- [14] T. Strutz, A. Leipzig, Reversible Colour Spaces without Increased Bit Depth and Their Adaptive Selection, *IEEE Signal Process. Lett.* 22 (9) (2015) 1269–1273, <http://dx.doi.org/10.1109/LSP.2015.2397034>.
- [15] Information Technology – JPEG 2000 image coding system: extensions for three-dimensional data, ISO/IEC International Standard 15444-10 and ITU-T recommendation T.809, 2011.
- [16] IRIS-JP3D, JPEG 2000 part 10 reference software, version 1.1.1, available: (<http://www.irissoftware.be/>), (accessed 20.05.15).
- [17] M.J. Weinberger, G. Seroussi, G. Sapiro, The LOCO-I lossless image compression algorithm: principles and standardization into JPEG-LS, *IEEE Trans. Image Process.* 9 (8) (2000) 1309–1324, <http://dx.doi.org/10.1109/83.855427>.
- [18] Information technology – Lossless and near-lossless compression of continuous-tone still images – Baseline, ISO/IEC International Standard 14495-1 and ITU-T Recommendation T.87, 2006.
- [19] G.J. Sullivan, J.R. Ohm, W.J. Han, T. Wiegand, Overview of the high efficiency video coding (HEVC) standard, *IEEE Trans. Circuits Syst. Video Technol.* 22 (12) (2012) 1649–1668, <http://dx.doi.org/10.1109/TCSVT.2012.2221191>.
- [20] Information technology – High efficiency coding and media delivery in heterogeneous environments – Part 2: high efficiency video coding, ISO/IEC International Standard 23008-2 and ITU-T Recommendation H.265, 2015.
- [21] Signal Processing and Multimedia Group, University of British Columbia JPEG-LS implementation, version 2.2, available: (<http://www.stat.columbia.edu/~jakulin/jpeg-ls/mirror.htm>), (accessed 28.07.15).
- [22] HEVC Test Model (HM) reference software (ISO/IEC International Standard 23008-5 and ITU-T Recommendation H.265.2), version 16.6, revision 4573, available: (https://hevc.hhi.fraunhofer.de/svn/svn_HEVCSoftware/), (accessed 28.07.15).
- [23] R. Starosolski, Compressing high bit depth images of sparse histograms, in: Proceedings of AIP, International Electronic Conference on Computer Science IeCCS 2007 Part II, 1060, 2008, pp. 269–272, <http://dx.doi.org/10.1063/1.3037069>.
- [24] T. Bernas, R. Starosolski, R. Wójcicki, Application of detector precision characteristics for the denoising of biological micrographs in the wavelet domain, *Biomed. Signal Process. Control.* 19 (2015) 1–13, <http://dx.doi.org/10.1016/j.bspc.2015.02.010>.
- [25] S. Srinivasan, C. Tu, S.L. Regunathan, G.J. Sullivan, HD Photo: a new image coding technology for digital photography, in: Proceedings of SPIE Applications of Digital Image Processing XXX, 6696, 2007, p. 66960A, <http://dx.doi.org/10.1117/12.767840>.
- [26] Information technology – JPEG XR image coding system – Image coding specification, ISO/IEC International Standard 29199-2 and ITU-T Recommendation T.832, 2012.
- [27] Q. Cai, L. Song, G. Li, N. Ling, Lossy and lossless intra coding performance evaluation: HEVC, H.264/AVC, JPEG 2000 and JPEG LS, in: 2012 Asia-Pacific Signal & Information Processing.