Reversible Integer-to-Integer Wavelet Transforms for Image Compression: Performance Evaluation and Analysis

Michael D. Adams, *Student Member, IEEE*, and Faouzi Kossentini, *Senior Member, IEEE*

Abstract

In the context of image coding, a number of reversible integer-to-integer wavelet transforms are compared on the basis of their lossy compression performance, lossless compression performance, and computational complexity. Of the transforms considered, several were found to perform particularly well, with the best choice for a given application depending on the relative importance of the preceding criteria. Reversible integer-to-integer versions of numerous transforms are also compared to their conventional (i.e., non-reversible real-to-real) counterparts for lossy compression. At low bit rates, reversible integer-to-integer and conventional versions of transforms were found to often yield results of comparable quality. Factors affecting the compression performance of reversible integer-to-integer wavelet transforms are also presented, supported by both experimental data and theoretical arguments.

Keywords

Reversible integer-to-integer wavelet/subband transforms, image coding/compression.

I. Introduction

There has been a growing interest in reversible integer-to-integer wavelet transforms for image coding applications [1–12]. Such transforms are invertible in finite-precision arithmetic (i.e., reversible), map integers to integers, and approximate the linear wavelet transforms from which they are derived. Due largely to these properties, transforms of this type are extremely useful for compression systems requiring efficient handling of lossless coding, minimal memory usage, or low computational complexity. Furthermore, these transforms are particularly attractive for supporting functionalities such as progressive lossy-to-lossless recovery of images [1, 6, 8], lossy compression with the lossless reproduction of a region of interest [9, 10], and strictly lossy compression with minimal memory usage [7]. Due to applications like these, we can see that there is a clear need to consider not only the lossless compression performance of a particular reversible integer-to-integer wavelet transform, but also its lossy compression performance.

In this paper, several reversible integer-to-integer wavelet transforms are compared on the basis of their lossy compression performance, lossless compression performance, and computational complexity. By using the results of such an evaluation, not only are we able to build systems with improved compression efficiency, but we can also better

Manuscript received January 4, 1999; revised December 12, 1999. This work was supported by both the Natural Sciences and Engineering Research Council of Canada, and Image Power Inc. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Antonio Ortega.

The authors are with the Signal Processing and Multimedia Group, Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, B.C., Canada V6T 1Z4 (e-mail: mdadams@ieee.org; faouzi@ece.ubc.ca).

understand the tradeoffs between compression efficiency and computational complexity. Reversible integer-to-integer versions of transforms are also compared to their conventional (i.e., non-reversible real-to-real) counterparts for lossy compression. The objective, in this case, is to quantify any performance degradation associated with the introduction of the reversible and integer-to-integer properties. If these properties do not adversely affect image quality, this would provide a compelling argument for the use of reversible integer-to-integer transforms in strictly lossy compression systems in order to reduce memory and computational requirements. If, however, compression performance is negatively impacted, it would be useful to have some quantitative measure of this degradation. Finally, factors affecting the compression performance of reversible integer-to-integer wavelet transforms are discussed, supported by both experimental data and theoretical arguments. Through the insight such information provides, one can hope to design new and more effective transforms.

The remainder of this paper is structured as follows. Section II begins by introducing the reversible integer-to-integer wavelet transforms considered in this study. This is followed, in Section III, by a brief comparison of the various transforms in terms of their computational complexity and memory requirements. Then, in Section IV, the transforms are employed for both lossy and lossless coding in a state-of-the-art image compression system and the results compared. Of the transforms considered, several were found to perform particularly well, with the best choice for a given application depending on the relative importance of lossy compression performance, lossless compression performance, and computational complexity. At low bit rates, reversible integer-to-integer and conventional versions of transforms were found to often yield results of comparable quality. In Section V, the lossy and lossless coding results are analyzed in more detail. Finally, we conclude with a brief summary of our results and some closing remarks in Section VI.

II. Transforms

The reversible integer-to-integer wavelet transforms considered in this study are based on the lifting framework [13], or a generalization thereof [6, 7], and were constructed using the technique described in [11]. In all, twelve transforms known to be effective for image coding were evaluated as listed in Table I. All of these transforms are strictly 1D in nature and based on 2-channel filter banks. The SPB and SPC transforms are based on filter banks with FIR analysis filters and IIR synthesis filters, not all of which have linear phase. All of the other transforms, however, are based on filter banks with linear-phase FIR filters. In the case of transforms based on FIR filter banks, we use the notation x/y to indicate that the underlying filter bank has lowpass and highpass analysis filters of lengths x and y, respectively. The notation $(\tilde{N}, N)$ is also employed to indicate that the analyzing and synthesizing wavelet functions have $\tilde{N}$ and N vanishing moments, respectively.

Since the transforms under consideration are 1D in nature, images are handled by transforming the rows and columns in succession. Unlike in the case of conventional (linear) versions of transforms, however, the order in which rows and columns are transformed is important. That is, the inverse transform must operate on rows and columns in the reverse order from that used in the forward transform; otherwise, invertibility cannot be guaranteed.

The forward transformation equations for each of the transforms are given in Table II. The input signal, lowpass subband signal, and highpass subband signal are denoted as $x[n]$, $s[n]$, and $d[n]$, respectively. For convenience, we also define the quantities $s_0[n] \triangleq x[2n]$ and $d_0[n] \triangleq x[2n+1]$. The inverse transformation equations can be trivially deduced from the forward transformation equations, and thus are not given here.

When transforming finite-length signals, it is necessary to adopt some strategy for handling filtering at the signal boundaries. In image coding applications, symmetric extension [20] is often employed for this purpose. To use symmetric extension, however, the underlying filter bank must preserve signal symmetry. Although most of the transforms considered in this paper are derived from linear filter banks with linear-phase filters, this is not sufficient to ensure compatibility with symmetric extension. Due to the very coarse quantization introduced in order to force a filter bank to map integers to integers, the symmetry-preserving property is often lost.

In our study, one of two techniques for handling finite-length signals is employed, depending on the transform at hand. In both instances, the method employed results in a nonexpansive transform for signals of arbitrary length. In cases where the underlying reversible integer-to-integer filter bank preserves signal symmetry, symmetric extension is used. More specifically, this strategy is employed in the case of the 5/3, 9/7-M, 5/11-C, 5/11-A, 13/7-T, 13/7-C, and 9/7-F transforms. In the case of the other transforms, a simple signal extension technique is employed—namely, the various signals being processed are extended by the repetition of their first and last sample values. Periodic extension is avoided due to its poor performance in image coding applications.

Reversible integer-to-integer wavelet transforms approximate their parent linear transforms. For this reason, the characteristics of the parent transforms are of great interest. In the case of linear wavelet transforms, a number of characteristics have been shown to potentially influence a transform's effectiveness for coding purposes, including the number of vanishing moments of the analyzing and synthesizing wavelet functions, the regularity of the analyzing and synthesizing scaling functions, coding gain, and the frequency selectivity of the analysis filters. These parameters are given for each transform in Table III. Also of great importance is the shape of the synthesizing scaling and wavelet functions. It is the shape of these functions that ultimately determines the signals for which a transform is effective. In addition, the shape of these functions determines the nature of the artifacts obtained in lossy signal reconstructions. The synthesizing scaling and wavelet functions for each transform are shown in Figs. 1–12.

III. Computational Complexity and Memory Requirements

All of the reversible transforms considered in this study are calculated using only fixed-point arithmetic. In particular, only integer addition/subtraction, multiplication, and division operations are required. Since all filter coefficients are approximated by dyadic rational numbers, all division operations can be implemented as bit shifts. For the purposes of computational complexity analysis, we assume that the two's complement representation for integers is used. By assuming a two's complement representation, we can take advantage of the fact that $\lfloor x/2^N \rfloor$ is equivalent to the arithmetic right shift of x by N bits.

Assuming a one-level wavelet decomposition (in one dimension), the number of addition, multiplication, and shift operations required per two input samples for each transform is given in Table IV. In each case, two sets of numbers are provided. The first set shows the number of operations in the straightforward implementation of the transform. The second set, shown in parentheses, indicates the number of operations required if all multiplications are converted to shift and add operations (via Booth's algorithm [21]). This complexity metric is of particular interest for hardware implementations or software implementations on architectures where integer multiplications are more costly than addition and shift operations. In the case where both sets of numbers are the same, only one is given. Symmetry in filter coefficients was considered in the calculation of these values.

Note that although a number of the transforms have equal operation counts, some require more multiplications than others. In particular, note that the 5/3, 2/6, 5/11-C and 5/11-A transforms are truly multiplierless (i.e., their underlying lifting filters all have coefficients that are strictly powers of two). Evidently, the 5/3 and 2/6 transforms require the least computation, followed by the SPB, 9/7-M, 2/10, 5/11-C, and 5/11-A transforms as a group, and then the 6/14, SPC, 13/7-T, and 13/7-C transforms as a group, and lastly the 9/7-F transform which requires the most computation.

As mentioned previously, reversible integer-to-integer versions of transforms have the potential for simpler implementations than their conventional counterparts. In the conventional case, all operands are real-valued. In the reversible integer-to-integer case, however, many operands are integers. Since, in both cases, approximately the same number of bits are required for the integer part of numbers, it follows that more bits require processing in the conventional case. Thus, reversible integer-to-integer transforms can offer reduced computational complexity through a smaller wordsize for arithmetic operations.

The amount of memory needed by a transform-based coder can be strongly influenced by the amount of memory required to store transform coefficients. In many image coding systems, transform coefficients are represented as real numbers. Typically, in such systems, 32 bits of memory are needed to store each coefficient. Transforms which map integers to integers can improve significantly upon this situation. Although some intermediate results are still real-valued, such results are only needed transiently. All stored values, including the final transform coefficients, are integers. In fact, experimental results suggest that for many reversible integer-to-integer wavelet transforms, 16-bit integers provide sufficient dynamic range for representing the transform coefficients (at least, for 8-bpp images). For example, in this study, 16-bit integers proved to be sufficient for all of the test images that were 10 bpp or less. By using 16 bits of memory per coefficient instead of 32, memory requirements can be reduced by half.

One might even be able to use a smaller number of bits to represent transform coefficients in some cases. Using a smaller number of bits, however, does increase the risk of numerical overflow. Fortunately, the invertibility of lifting-based transforms can be made to hold even in the presence of overflow (e.g., by exploiting the properties of modulo integer arithmetic as in [12]). Provided that occurrences of overflow are infrequent, lossless compression performance is not seriously affected. In the case of lossy compression, however, one needs to be much more cautious about overflow. Overflow has the potential to cause very disturbing artifacts, especially if the overflow impacts coefficients that affect large areas in the reconstructed image.

IV. Experimental Results

As part of our study, various experimental results were obtained regarding the compression performance of the transforms under evaluation. In the sections that follow, we briefly present these results without any analysis. A detailed analysis of the results will be provided in Section V.

A. Evaluation Methodology

For evaluation purposes, the JPEG-2000 verification model (VM) software (version 0.0) was employed. Although some support for reversible integer-to-integer wavelet transforms was provided in this software, much of the functionality required for our analysis was not present. Therefore, in order to facilitate our testing, the original transform-related code was replaced with new and more flexible code. This new transform code is described at some length in [22].

The JPEG-2000 VM software employed in this work handles both lossy and lossless compression using transform-based methods. In the lossy case, a subband transform is applied to the original image, the transform coefficients are quantized using trellis coded quantization (TCQ) [23], and the resulting TCQ indices are entropy coded to yield the compressed bit stream. In the lossless case, a reversible integer-to-integer subband transform is applied to the original image, and the transform coefficients are then bitplane and entropy coded.

The various transforms under consideration were employed for decorrelation purposes in the image coder described above, and the results examined. Both lossy and lossless compression were considered in our analysis. Using the various transforms, each of the test images was compressed in a lossy manner at several bit rates (i.e., 0.0625, 0.125, 0.25, 0.5, 1, and 2 bpp), in addition to being compressed losslessly. In the lossless case, compression performance was measured in terms of the final bit rate. In the lossy case, compression performance was assessed using both the PSNR and subjective image quality metrics. Since more than one definition of PSNR exists, we note that the one employed in this study is

$$\text{PSNR} = 20\log_{10}\left(\frac{2^P - 1}{\sqrt{\text{MSE}}}\right) \tag{1}$$

where

$$\text{MSE} = \frac{1}{MN}\sum_{i=0}^{M-1}\sum_{j=0}^{N-1}(\hat{x}[i,j] - x[i,j])^2,$$

$x[\cdot]$ is the original image with dimensions $M \times N$ and P bpp, and $\hat{x}[\cdot]$ is the reconstructed image.

The test data used in our study consisted of the grayscale images listed in Table V, all of which were taken from the JPEG-2000 test set. These images vary considerably in size and cover a wide variety of content. For each set of experimental results, the same number of decomposition levels was employed for all transforms (typically six). Since the goal of the subjective evaluations was only to compare artifacts obtained from the transforms themselves and not tiling artifacts, the use of tiling was avoided whenever possible. More specifically, tiling was used only in the case of the `cmpnd2` image due to its large size, where a tile size of 2048×2048 was employed.

In order to perform the subjective testing, several regions from the various lossy reconstructed images were selected to be scrutinized. Small regions of interest were used (instead of whole images) for two reasons: 1) Many of the test images were too large to be displayed at their original resolution on the computer screen, and displaying the images at reduced resolution hides most of the artifacts introduced by compression. 2) If the whole image is used, there might be too much information for a human observer to analyze, making an accurate and meaningful comparison difficult. Since the human eye often cannot distinguish between a high bit rate lossy reconstruction of an image and the original, the subjective testing was restricted to the lower bit rates under consideration. As most applications typically utilize compression ratios from 16:1 to 128:1, the bit rates corresponding to these compression ratios were of most interest. The participants involved in the subjective testing came from diverse backgrounds, some having a strong grounding in image processing and/or image coding, while others were end users with no special technical training. Two types of subjective evaluations were employed in our study as described below.

The first type of subjective evaluation served to compare the image quality obtained with different reversible integer-to-integer wavelet transforms. For a specific test image and bit rate, a particular region of the lossy reconstructed images obtained with the various transforms was shown under magnification. The original image was also shown for reference purposes. Then, a human observer was asked to assign a simple integer ranking to the different results, with a value of one corresponding to best result and larger values indicating poorer results. In instances where a distinction in quality could not be made between two or more results, these results could be assigned the same rank. This process was repeated for various test images, bit rates, regions of interest within the image, and observers. The rankings were then used to determine the most effective transform from a subjective image quality standpoint.

The second type of subjective evaluation was used to compare the quality of the lossy reconstructed images obtained with reversible integer-to-integer and conventional versions of the same transform. For a specific test image, bit rate, and transform, a particular region of interest in the lossy reconstructed images obtained with the reversible integer-to-integer and conventional versions of a transform was shown under magnification. Again, the original image was also shown for reference purposes. Then, the observer was asked to rank the two reconstructed images relative to one another. The ranking was done on an integer scale from -2 to 2. The values -2, -1, 0, 1, 2 correspond to the result obtained with the reversible integer-to-integer version of the transform being significantly worse than, slightly worse than, the same as, slightly better than, and significantly better than the result obtained with the conventional version, respectively. This process was repeated for numerous transforms, images, bit rates, and observers.

B. Lossless Compression Performance

Each of the test images in Table V was compressed in a lossless manner using the various transforms under evaluation, resulting in the bit rates shown in Table VI. For each image, the best result has been highlighted. Clearly, no single transform performs best for all of the images. In order to provide additional insight into the results of Table VI, relative differences were also computed. For each image, the best transform was used as a reference and the relative differences for the remaining transforms were calculated. Often, the difference between the best and worst transforms is less than 2%. For the images where this difference is more significant (i.e., greater than 2%), the relative performance numbers have been included in Table VII.

As one might expect, image content is an important factor influencing transform effectiveness. If one considers only natural imagery, some particularly useful observations can be made. In addition to the lossless results for each image, Table VI provides the average bit rate for each transform taken over all of the natural 8-bpp images. For images of this type (e.g., `cafe`, `gold`, `woman`), the 5/11-C, 5/11-A, and 13/7-T transforms generally perform best, followed closely by the 9/7-M and 13/7-C transforms. The 5/3 transform also fares reasonably well considering its very low computational complexity. For images with a relatively greater amount of high-frequency content (e.g., `cmpnd2` and `target`), the 5/3 transform tends to yield the best results, often by a significant margin.

In [11] and [2], Calderbank et al. compare the lossless compression performance of several reversible integer-to-integer wavelet transforms also considered in our study (i.e., the 2/6, 5/3, SPB, 9/7-M, 5/11-C, and 9/7-F transforms). The image coder employed in their evaluation was based on techniques described in [1]. Generally, the results presented by Calderbank and his colleagues coincide with our findings. For example, both studies found that (for lossless compression): 1) the 5/3 transform outperforms the 2/6 transform for the vast majority of images; 2) the 9/7-F transform consistently performs very poorly; and 3) the 5/11-C transform fares amongst the best for natural imagery. Interestingly, the 13/7-T transform, which we found to be particularly effective for natural imagery, was described in [11] and [2], but was not included in their performance evaluation.

In the context of a reversible embedded image coder based on the EZW coding scheme [24], lossless compression results similar to those above have also been obtained [6, 25]. Thus, these results appear to be consistent across most types of lossless coders, especially those employing bitplane coding of the transform coefficients.

C. Objective Lossy Compression Performance

Each of the test images in Table V was compressed in a lossy manner at several bit rates using the various transforms under evaluation. The image quality was then measured using the PSNR metric as defined by (1). The results for a representative subset of the test images are given in Table VIII. For each test image and bit rate pair, the best result has been highlighted.

For low bit rates (i.e., compression ratios of 16:1 or greater), the 9/7-F, 6/14, 13/7-C, and 13/7-T transforms are consistently the best performers in approximately that order. The SPC and SPB transforms fare the worst. As the bit rate increases, lossy compression performance becomes increasing influenced by the lossless performance characteristics of the transforms. These results will be revisited later in Section V-B for analysis purposes.

In applications where lossy image reconstructions are processed by a computer system (e.g., some scientific applications), the use of a distortion metric based on the mean-squared error (i.e., PSNR) can be quite appropriate. Unfortunately, the PSNR can be a poor measure of image quality for applications in which images are to be viewed by humans. Hence, a subjective evaluation, the results of which are presented in the next section, was necessary.

D. Subjective Lossy Compression Performance

Since the PSNR does not necessarily correlate well with image quality as perceived by the human visual system, we supplemented our objective results with a subjective evaluation. This subjective testing was undertaken as specified in Section IV-A. In order to make the testing more practical, the evaluation was divided into two stages. The same testing strategy was used in each stage. Only the transforms under evaluation in each case differed.

In the first stage, the 5/3, 2/6, SPB, 9/7-M, 2/10, SPC, 13/7-T, and 9/7-F transforms were compared. After obtaining the rankings for 14 sets of images using 12 observers, these results were averaged together to yield the overall rankings given in Table IX. From these results, we see that the transforms fall into three distinct performance categories. The 9/7-F, 5/3, 13/7-T, and 9/7-M transforms perform best, followed by the 2/6 and 2/10 transforms, followed by the SPB and SPC transforms. These numerical results generally agree with comments made by participants during the subjective tests—namely, that the transforms fall into three distinct performance categories.

The second stage of subjective testing included some of the best transforms from the first stage and several new transforms, resulting in the following set of transforms: 5/3, 9/7-M, 5/11-C, 5/11-A, 6/14, 13/7-T, and 13/7-C. After conducting the subjective tests with 14 sets of images and 12 observers, the rankings were averaged together with the results shown in Table X. In this round of tests, the general consensus amongst the test participants is that all of the transforms are very close in terms of subjective quality, with the possible exception of the 6/14 transform which performs slightly worse. In other words, although the average rankings may differ by a large amount in this stage of the testing, it is important to note that even large disparities correspond to very minute differences in subjective image quality.

Considering the results from both stages of the subjective testing, we conclude that the 9/7-F, 5/3, 13/7-T, 13/7-C, 5/11-C, 5/11-A, and 9/7-M transforms have the best performance, followed by the 6/14, 2/6, and 2/10 transforms, followed by the SPB and SPC transforms. We will again revisit these results in Section V-B where they will be examined in more detail.

E. Reversible Integer-to-Integer Versus Conventional Transforms for Lossy Compression

Using both reversible integer-to-integer and conventional versions of the transforms listed in Table I, the test images in Table V were compressed in a lossy manner at several bit rates. The reconstructed images obtained with the two different versions of each transform were then compared. The difference in the PSNR for a representative subset of the images is given in Table XI. A negative value corresponds to poorer performance in the reversible integer-to-integer transform case. As is evident from these results, the performance is almost always degraded in terms of PSNR in the reversible integer-to-integer transform case, although the degradation is often very small. These numerical results also suggest some other important trends that will be discussed later in Section V-A.

Subjective testing was also undertaken to compare the quality of the lossy image reconstructions obtained with reversible integer-to-integer and conventional versions of transforms. This testing was performed as described previously in Section IV-A. Due to the time consuming nature of subjective testing, only a subset of the transforms listed in Table I was used, namely the 5/3, 2/6, SPB, 9/7-M, 2/10, SPC, 13/7-T, and 9/7-F transforms. After averaging the scores obtained using 7 sets of images (per transform) and 10 observers, the results in Table XII were obtained. These numerical results indicate the relative performance of the reversible integer-to-integer and conventional versions of the transforms. A negative value corresponds to better performance in the conventional transform case. From the experimental results, we can see that in most cases there is only a small bias favoring the conventional version of a transform. Thus, there is a very slight performance penalty associated with using reversible integer-to-integer transforms in place of their conventional counterparts. In most practical applications, however, this difference in performance is not likely to be of concern. Again, we will revisit the results for analysis purposes in Section V-A.

V. Analysis of Experimental Results

Having briefly presented the various experimental results, we now examine these results in more detail in the sections that follow. To begin, we discuss the factors that affect the relative lossy compression performance of reversible integer-to-integer and conventional versions of transforms. Then, factors influencing the compression performance of reversible integer-to-integer wavelet transforms are presented.

A. *Reversible Integer-to-Integer Versus Conventional Transforms for Lossy Compression*

In Section IV-E, we presented results comparing the lossy compression performance of reversible integer-to-integer and conventional versions of transforms. Examining these results more closely, we see that several factors affect the performance of reversible integer-to-integer transforms relative to their conventional counterparts: 1) whether the underlying filter bank has IIR filters, 2) the number of lifting steps in the filter bank, 3) the quantization function used in the lifting steps, 4) the depth of the image (in bpp) being compressed, and 5) the bit rate used for compression of the image. Each of these factors is discussed in detail below.

Impact of IIR Filters

In the calculation of reversible integer-to-integer wavelet transforms, results are quantized to integer values in numerous places. If the underlying filter bank has IIR filters, this very coarse quantization occurs inside feedback paths. Thus, quantization errors can accumulate indefinitely[1]. For this reason, there is a much bigger difference between the reversible integer-to-integer and conventional versions of a transform in the case where the transform is associated with a filter bank having IIR filters. In the case of the transforms studied, the SPB and SPC transforms are associated with such filter banks. It is evident that these two transforms suffer the greatest degradation in performance as a result of the reversible and integer-to-integer properties.

Figs. 13 and 14 demonstrate the impact of having IIR filters. Figs. 13(a) and 13(b) show typical results in the case of transforms based on FIR filter banks. In this case, there is little perceptible difference between the results obtained with the reversible integer-to-integer and conventional versions of the same transform. Figs. 14(a) and 14(b) demonstrate typical results obtained in the case of transforms based on filter banks with IIR filters. In this case, the reversible integer-to-integer version of the transform yields much poorer results due to increased ringing.

Number of Lifting Steps

Clearly, the number of lifting steps also influences the compression performance. Since each lifting step further causes the approximation error to increase, transforms with fewer lifting steps will tend to perform better, all other things being equal. Consider, for example, the 5/3, 5/11-C, and 5/11-A transforms. All three transforms have the same first two lifting steps. The 5/3 transform, however, has only two lifting steps, while the other two transforms have three. At least in terms of PSNR performance, the 5/3 transform typically has less degradation in performance than the 5/11-C and 5/11-A transforms.

Quantization Function

The function used to quantize results to integer values also affects the difference in performance between reversible integer-to-integer transforms and their conventional counterparts. In this study, we chose to use a biased floor function (i.e., $Q(x) = \lfloor x + \frac{1}{2} \rfloor$) for this purpose. One could conceivably remove the bias of one-half, however. Such a change would reduce computational complexity at the expense of introducing systematic error. The resulting performance degradation at low bit rates is typically small. At higher bit rates, however, the degradation can be more significant, especially in the case of lossless compression.

[1] Moreover, it may even be difficult to show that the resulting nonlinear system is stable.

Depth of Image

As the dynamic range of the signal grows, the errors introduced by quantizing various results to integer values are relatively smaller. For 8-bpp images there is a clear difference in PSNR performance between the reversible integer-to-integer and conventional versions of a transform, especially at higher bit rates. In the case of 12-bpp images (e.g., `sar2`), however, there is little difference in performance provided that the transforms are based on FIR filter banks.

Bit Rate

As the bit rate is lowered, the difference in compression performance between reversible integer-to-integer and conventional versions of a transform decreases. At a sufficiently low bit rate, the quantization of transform coefficients becomes so coarse that any errors introduced by quantizing intermediate results to integers are masked by quantization of the coefficients themselves.

B. Factors Affecting Compression Performance

In our study, three factors were found to affect the compression performance of a reversible integer-to-integer wavelet transform: 1) the performance characteristics of its parent linear transform, 2) how closely it is able to approximate its parent linear transform, and 3) the dynamic range of the coefficients produced by the transform. The first two factors affect both lossy and lossless compression performance. The last factor is typically only important in the case of lossless compression performance. Each of these factors will be discussed in turn in the remainder of this section. Table XIV lists several parameters for each transform relevant to compression performance. Comments about these parameters will also be made at the appropriate point in the paragraphs that follow.

Parent Linear Transform

The first point above has been covered in depth by many researchers, and techniques for the design of linear wavelet transforms are plentiful in the literature. Some design considerations include: time and frequency resolution, number of vanishing moments, regularity, and coding gain. The effectiveness of a particular transform is also to some extent signal dependent.

Section IV-C presented the PSNR lossy compression performance results for the various transforms. The 9/7-F, 6/14, 13/7-C, and 13/7-T transforms were found to be most effective. Looking at parameter G_6 in Table III, we can see that these four transforms also have the parent linear transforms with the highest coding gain. Thus, the coding gain of the parent linear transform is of great importance, as one might suspect.

In Section IV-D, the subjective lossy compression performance results for the various transforms were presented. To a large extent, the subjective performance of these transforms can be explained by examining the synthesizing scaling and wavelet functions of their parent linear transforms as shown in Figs. 1–12. From these plots, we can see that the 9/7-F, 13/7-T, 13/7-C, 5/11-C, 5/11-A, and 9/7-M transforms have very smooth synthesizing scaling and wavelet functions, leading to their favorable performance. In the case of the 5/3 transform, the scaling and wavelet functions have discontinuities in their first derivatives. Since there are no oscillations or jump discontinuities in these functions, however, the 5/3 transform is still able to perform very well. In fact, the 5/3 transform is especially attractive as it often introduces less pronounced ringing artifacts than the other transforms due to its short filters. Although sometimes very good at preserving high frequency textures, the SPB and SPC transforms perform worst, due largely to their tendency to introduce ringing artifacts. This behavior can be attributed to the fact that both transforms are associated with filter banks having IIR synthesis filters. The 6/14 transform often introduces more ringing artifacts than some of the other transforms, leading to its relatively poor performance. This ringing is associated with the greater oscillatory nature of the synthesizing wavelet and scaling functions. In the case of the 2/6 transform, blockiness often results. This behavior can be attributed to the relatively rougher synthesizing scaling and wavelet functions associated with this transform. Lastly, the 2/10 transform often results in more pronounced ringing artifacts, leading to its poorer performance. This behavior can be attributed to the larger support of the synthesizing scaling and wavelet functions.

Approximation Behavior

The approximation characteristics of reversible integer-to-integer wavelet transforms have not been given much serious attention to date. This is, however, an important factor to consider. Clearly, even if a reversible integer-to-integer transform is derived from a linear transform with desirable properties, the extent to which these properties are present in the reversible integer-to-integer transform depends entirely on how well it approximates its parent transform. The predominant source of error is the quantization of intermediate results to integers. This error depends on the quantization function employed, the number of lifting steps in which this quantization is performed, and whether recursive filtering structures are used.

The number of lifting steps associated with each transform is given in Table XIV. In the case of lossless compression, most of the best transforms tend to have only two or three lifting steps. All other things being equal, as the number of lifting steps increases so too does the approximation error. In the lossy case, the approximation error is less critical. Unless the approximation error is very large, it is usually masked by the error introduced by transform coefficient quantization. Notice in particular that the 9/7-F transform has the most lifting steps of the transforms considered, and it also has arguably the poorest lossless compression performance.

To further demonstrate the influence of the number of lifting steps on compression performance, we losslessly compressed several images using two different versions of the 13/7-T transform. The first version of the transform is the one given in Table II which has two lifting steps. The second version has four lifting steps. In each case, we measured the lossless bit rate and average approximation error. As can be seen from the results in Table XIII, the lossless compression performance is better for the version of the transform with two lifting steps. Note that the approximation error is also lower in this case as well. This observation suggests that, all other things being equal, transforms with two lifting steps (i.e., interpolating transforms) may offer the most promise for lossless compression[2].

Dynamic Range

In the case of lossless compression, all bits of the transform coefficients must be encoded. Thus, as the dynamic range of the transform coefficients increases, so too does the number of bits required to encode these coefficients. For this reason, dynamic range is an important factor affecting lossless compression performance. Typically, the coefficients obtained by the application of a transform have a greater dynamic range than the original samples. This dynamic range growth is, therefore, of concern. The worst-case dynamic range growth of a particular transform is an approximate function of the 1-norm of its analysis filters. The 1-norms of the lowpass and highpass analysis filters determine the

[2] Incidentally, using similar reasoning, it would seem that reversible integer-to-integer M-band wavelet transforms (where $M > 2$) [6] are not likely to be as effective for lossless coding as 2-band transforms.

worst-case growth in the lowpass and highpass channels, respectively. As the 1-norms increase, the worst-case dynamic range growth also increases. For reference purposes in what follows, the analysis filter 1-norms for each transform are given in Table XIV.

Since, in the case of a wavelet transform, the lowpass subband signal is successively decomposed, the worst-case dynamic range growth associated with the lowpass channel can potentially have the most effect on the dynamic range of the transform coefficients. To date, much attention has been focused on worst-case lowpass-channel dynamic range growth. For example, it is often cited that the 2/6 and 2/10 transforms are particularly well suited to lossless compression as they have no dynamic range growth in the lowpass channel. While the no lowpass-channel growth attribute can help contribute to good lossless compression performance, it is not good to place too much emphasis on this property.

The first problem with the above philosophy is that it is based on a worst-case argument. The class of signals that lead to transform coefficients having the maximum possible magnitudes are rarely encountered in practice. The worst case and actual dynamic ranges are thus quite different. Therefore, one must instead consider the typical distribution of coefficient magnitudes as opposed to the worst-case scenario. There is, however, a more fundamental problem with placing too much emphasis on lowpass-channel dynamic range growth. Namely, reduced growth in the lowpass channel often comes at the expense of increased growth in the highpass channel. For images, however, the 1D transform is applied twice at each level of wavelet decomposition: once horizontally and once vertically. This yields a subband structure like that shown in Fig. 15. This means that one quarter of all transform coefficients (i.e., the HH_0 band) are the result of two highpass filtering stages. Similarly, three quarters of all coefficients have been obtained by highpass filtering in at least one direction. Thus, if the lowpass-channel growth is made smaller at the cost of increased highpass-channel growth, this can, at some point, become counterproductive. Although the lower frequency bands may have coefficients with a smaller dynamic range, the number of coefficients in these bands is relatively small, and their influence is less significant.

Consider, for a moment, the 5/3 and 2/6 transforms, both being comparable in complexity. Unlike the 5/3 transform, the 2/6 transform has no growth in the lowpass channel. The 5/3 transform, however, has a lower highpass-channel growth. (See Table XIV.) In Section IV-B, the 5/3 transform was shown to outperform the 2/6 transform for lossless compression on the vast majority of the test images. In Table XV, for three of the test images, we show the average coefficient magnitude for the HL_0, LH_0, and HH_0 bands individually and all bands together, and the lossless bit rate obtained with each transform. For the `gold` and `target` images, where the 5/3 transform yields a notably lower bit rate, the HL_0, LH_0, and HH_0 bands have coefficients with lower average magnitudes in the 5/3 case. The average coefficient magnitude is also lower for both images with the 5/3 transform. In the case of the `cmpnd1` image, where both transforms yield approximately the same bit rate, the average coefficient magnitudes for the HL_0, LH_0, and HH_0 bands are much closer. Since the coefficient magnitudes are closer in this instance, the lower average magnitudes for the lower frequency bands in the 2/6 transform case become a dominating factor influencing performance.

From the above argument, we can conclude that the amount of growth in the lowpass channel is not necessarily a good predictor of a transform's effectiveness for lossless compression. In fact, in Section IV-B, we saw that the best transforms for lossless compression all have worst-case growth in the lowpass channel. Clearly, one must also consider the amount of growth in the highpass channel.

Examining the 13/7-T and 13/7-C transforms, we observe that they have very similar lifting filter coefficients. The 13/7-T transform, however, has analysis filters with smaller 1-norms. In Section IV-B, the 13/7-T transform was found to often outperform the 13/7-C transform for lossless compression. Again, this helps to illustrate that, all other things being equal, the transform with analysis filters having smaller 1-norms will typically perform better for lossless compression. In the case of the 9/7-F transform, another factor helps to contribute to its poor lossless compression performance. This transform is the only transform of those considered with a lowpass analysis filter having a DC gain greater than one. (See Table XIV.) This larger gain contributes to increased growth in the dynamic range of the transform coefficients.

VI. Conclusions

Several reversible integer-to-integer wavelet transforms have been compared on the basis of their lossy compression performance, lossless compression performance, and computational complexity. For lossless compression, image content is an important factor influencing transform effectiveness. For smooth images, the 5/11-C, 5/11-A, and 13/7-T transforms are most effective, while the 5/3 transform performs best for images with a greater amount of high frequency content. In the case of lossy compression, both objective (i.e., PSNR) and subjective performance measures were employed. In terms of PSNR performance, the 9/7-F, 6/14, 13/7-C, and 13/7-T transforms yield the best results. At the same time, the 9/7-F, 5/3, 9/7-M, 5/11-C, 5/11-A, 13/7-T, and 13/7-C transforms achieve the most favorable subjective ratings. Of the transforms considered, the 5/3 and 2/6 transforms have the lowest computational complexity while the 9/7-F transform has the highest. Unfortunately, the 9/7-F transform, although very good for lossy compression, is significantly higher in complexity than all of the other transforms, and performs poorly for lossless compression.

Obviously, no one transform has superior lossy and lossless compression performance for all classes of images in addition to low computational complexity. Therefore, tradeoffs are involved, and the most appropriate choice of transform depends on the needs of the application at hand. For example, in the case of applications requiring low computational complexity, the 5/3 transform is particularly attractive. It performs reasonably well for both lossy and lossless compression, and has the lowest computational complexity of all of the transforms evaluated.

Reversible integer-to-integer and conventional versions of the same transforms were also compared. Using reversible integer-to-integer versions of transforms instead of their conventional counterparts for lossy compression was found to only introduce a slight degradation in subjective image quality in many cases. When transforms associated with FIR filter banks are employed, the performance penalty is almost negligible. At low bit rates, the difference in PSNR performance is also small.

By examining the performance characteristics of the various reversible integer-to-integer wavelet transforms, we were also able to determine several factors that affect the compression performance of such transforms. By considering these factors in the future, we can hope to design new and more effective transforms.

Acknowledgment

The authors would like to express their gratitude to the anonymous reviewers for their useful comments and thank all of the participants in the subjective testing for their time and effort.

References

[1] A. Said and W. A. Pearlman, "An image multiresolution representation for lossless and lossy compression," *IEEE Trans. on Image Processing*, vol. 5, no. 9, pp. 1303–1310, Sept. 1996.

[2] A. R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo, "Lossless image compression using integer to integer wavelet transforms," in *Proc. of IEEE International Conference on Image Processing*, Santa Barbara, CA, USA, Oct. 1997, vol. 1, pp. 596–599.

[3] M. J. Gormish, E. L. Schwartz, A. F. Keith, M. P. Boliek, and A. Zandi, "Lossless and nearly lossless compression of high-quality images," in *Proc. of SPIE*, San Jose, CA, USA, Mar. 1997, vol. 3025, pp. 62–70.

[4] N. Memon, X. Wu, and B.-L. Yeo, "Improved techniques for lossless image compression with reversible integer wavelet transforms," in *Proc. of IEEE International Conference on Image Processing*, Chicago, IL, USA, Oct. 1998, vol. 3, pp. 891–895.

[5] F. Sheng, A. Bilgin, P. J. Sementilli, and M. W. Marcellin, "Lossy and lossless image compression using reversible integer wavelet transforms," in *Proc. of IEEE International Conference on Image Processing*, Chicago, IL, USA, Oct. 1998, vol. 3, pp. 876–880.

[6] M. D. Adams, "Reversible wavelet transforms and their application to embedded image compression," M.A.Sc. thesis, Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada, Jan. 1998, Available from `http://www.ece.ubc.ca/~mdadams`.

[7] M. D. Adams and F. Kossentini, "Performance evaluation of several reversible integer-to-integer wavelet transforms in the JPEG-2000 baseline system (VM release 0.0)," ISO/IEC JTC 1/SC 29/WG 1 N866, June 1998, Available from `http://www.ece.ubc.ca/~mdadams`.

[8] A. Zandi, J. D. Allen, E. L. Schwartz, and M. Boliek, "CREW: Compression with reversible embedded wavelets," in *Proc. of IEEE Data Compression Conference*, Snowbird, UT, USA, Mar. 1995, pp. 212–221.

[9] D. Nister and C. Christopoulos, "Lossless region of interest with a naturally progressive still image coding algorithm," in *Proc. of IEEE International Conference on Image Processing*, Chicago, IL, USA, Oct. 1998, vol. 3, pp. 856–860.

[10] E. Atsumi and N. Farvardin, "Lossy/lossless region-of-interest image coding based on set partitioning in hierarchical trees," in *Proc. of IEEE International Conference on Image Processing*, Chicago, IL, USA, Oct. 1998, vol. 1, pp. 87–91.

[11] A. R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo, "Wavelet transforms that map integers to integers," *Applied and Computational Harmonic Analysis*, vol. 5, no. 3, pp. 332–369, July 1998.

[12] H. Chao, P. Fisher, and Z. Hua, "An approach to integer wavelet transforms for lossless for image compression," in *Proc. of International Symposium on Computational Mathematics*, Guangzhou, China, Aug. 1997, pp. 19–38.

[13] W. Sweldens, "The lifting scheme: A new philosophy in biorthogonal wavelet constructions," in *Proc. of SPIE*, San Diego, CA, USA, Sept. 1995, vol. 2569, pp. 68–79.

[14] D. Le Gall and A. Tabatabai, "Sub-band coding of digital images using symmetric short kernel filters and arithmetic coding techniques," in *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing*, New York, NY, USA, Apr. 1988, vol. 2, pp. 761–764.

[15] J. D. Villasenor, B. Belzer, and J. Liao, "Wavelet filter evaluation for image compression," *IEEE Trans. on Image Processing*, vol. 4, no. 8, pp. 1053–1060, Aug. 1995.

[16] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Wellesley, MA, 1996.

[17] M. D. Adams and F. Kossentini, "Low-complexity reversible integer-to-integer wavelet transforms for image coding," in *Proc. of IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, Victoria, BC, Canada, Aug. 1999, pp. 177–180.

[18] M. D. Adams, I. Kharitonenko, and F. Kossentini, "Report on core experiment CodEff4: Performance evaluation of several reversible integer-to-integer wavelet transforms in the JPEG-2000 verification model (version 2.1)," ISO/IEC JTC 1/SC 29/WG 1 N1015, Oct. 1998, Available from `http://www.ece.ubc.ca/~mdadams`.

[19] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. on Image Processing*, vol. 1, no. 2, pp. 205–220, Apr. 1992.

[20] C. M. Brislawn, "Preservation of subband symmetry in multirate signal coding," *IEEE Trans. on Signal Processing*, vol. 43, no. 12, pp. 3046–3050, Dec. 1995.

[21] A. D. Booth, "A signed binary multiplication technique," *Quarterly Journal of Mechanics and Applied Mathematics*, vol. 4, pp. 236–240, 1951.

[22] M. D. Adams and F. Kossentini, "SBTLIB: A flexible computation engine for subband transforms," ISO/IEC JTC 1/SC 29/WG 1 N867, June 1998, Available from `http://www.ece.ubc.ca/~mdadams`.

[23] J. H. Kasner, M. W. Marcellin, and B. R. Hunt, "Universal trellis coded quantization," *IEEE Trans. on Image Processing*, vol. 8, no. 12, pp. 1677–1687, Dec. 1999.

[24] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. on Signal Processing*, vol. 41, no. 12, pp. 3445–3462, Dec. 1993.

[25] M. D. Adams and A. Antoniou, "Reversible EZW-based image compression using best-transform selection and selective partial embedding," Tech. Rep., Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada, 1998, Available from `http://www.ece.ubc.ca/~mdadams`.

Michael D. Adams (M'89) was born in Kitchener, Ontario, Canada, in 1969. He received the B.A.Sc. degree in computer engineering from the University of Waterloo, Waterloo, Canada, in 1993, and the M.A.Sc. degree in electrical engineering from the University of Victoria, Victoria, Canada, in 1998. From 1993 to 1995, Michael was a member of technical staff at Bell-Northern Research in Ottawa, Canada, where he developed real-time software for fiber-optic telecommunication systems.

Currently, Michael is working toward the Ph.D. degree in electrical engineering at the University of British Columbia, Vancouver, Canada. He is the recipient of a Natural Sciences and Engineering Research Council (of Canada) Postgraduate Scholarship. Also, he is a voting member of the Canadian delegation to ISO/IEC JTC 1/SC 29/WG 1, and has been an active participant in the JPEG-2000 standardization effort. His research interests include digital signal processing, wavelets, multirate systems, image coding, and multimedia systems.

Faouzi Kossentini (M'90-SM'98) received the B.S., M.S., and Ph.D. degrees from the Georgia Institute of Technology, Atlanta, in 1989, 1990, and 1994, respectively. During 1995, Dr. Kossentini had been working as a research scientist at Nichols Research Corporation, Huntsville, AL, USA. Since January 1996, he has been employed as an assistant professor and then as an associate professor in the Department of Electrical and Computer Engineering at the University of British Columbia, where he is involved in research in the areas of signal processing, communications and multimedia.

He has co-authored more than one hundred journal papers, conference papers, book chapters and patents. Dr. Kossentini has been active as a voting member, and recently as head of delegation, of the Canadian delegate to ISO/IEC JTC1/SC29, which is responsible for the standardization of coded representation of audiovisual, multimedia, and hypermedia information. In particular, he has participated in most current JBIG/JPEG and MPEG-4 standardization activities. He has also participated in most current ITU-T low bit rate video coding standardization activities. Dr. Kossentini is a Senior Member of the IEEE. He has served as a technical area coordinator and member of the technical program committee of ICIP-1997, and as a member of the technical program committee of ISCAS-1999. He is also the Vice General Chairman of ICIP-2000. Dr. Kossentini is currently an Associate Editor for the IEEE Transactions on Image Processing.

List of Figures

List of Tables

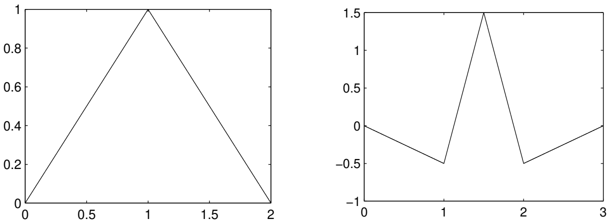

(a) (b)

Fig. 1. Synthesizing scaling and wavelet functions for the 5/3 transform. (a) Scaling function. (b) Wavelet function.

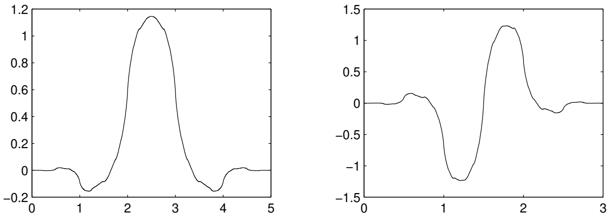

(a) (b)

Fig. 2. Synthesizing scaling and wavelet functions for the 2/6 transform. (a) Scaling function. (b) Wavelet function.

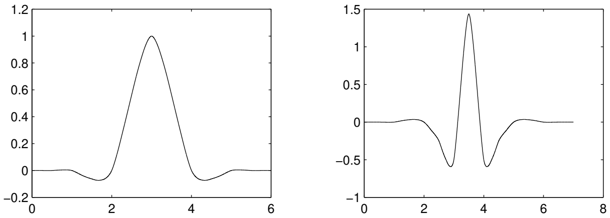

(a) (b)

Fig. 3. Synthesizing scaling and wavelet functions for the 9/7-M transform. (a) Scaling function. (b) Wavelet function.

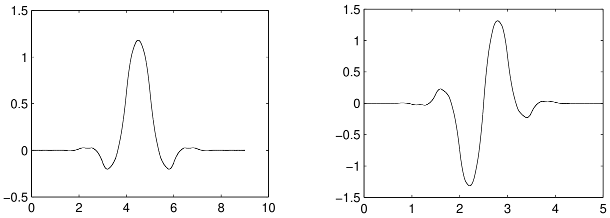

(a) (b)

Fig. 4. Synthesizing scaling and wavelet functions for the 2/10 transform. (a) Scaling function. (b) Wavelet function.

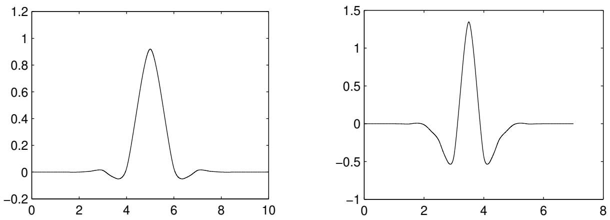

(a) (b)

Fig. 5. Synthesizing scaling and wavelet functions for the 5/11-C transform. (a) Scaling function. (b) Wavelet function.

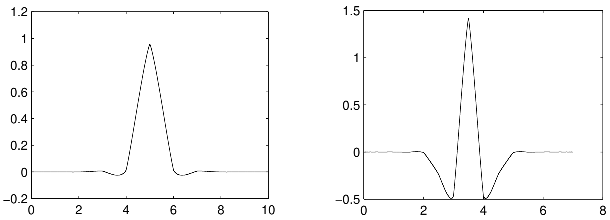

(a) (b)

Fig. 6. Synthesizing scaling and wavelet functions for the 5/11-A transform. (a) Scaling function. (b) Wavelet function.

TABLE I

TRANSFORMS

Name	Description / References
5/3	5/3, (2, 2), [14], equation (4.1) in [11]
2/6	2/6, (3, 1), [15], [8]
SPB	S+P Transform with Predictor B, IIR synthesis filters, [1]
9/7-M	9/7, (4, 2), [16], equation (4.2) in [11]
2/10	2/10, (5, 1), [3]
5/11-C	5/11, (4, 2), equation (4.6) in [11]
5/11-A	5/11, (2, 2), [17]
6/14	6/14, (3, 3), [18]
SPC	S+P Transform with Predictor C, IIR synthesis filters, [1]
13/7-T	13/7, (4, 4), [16], equation (4.5) in [11]
13/7-C	13/7, (4, 2), [18]
9/7-F	9/7, (4, 4), [19], section 4.4 in [11]

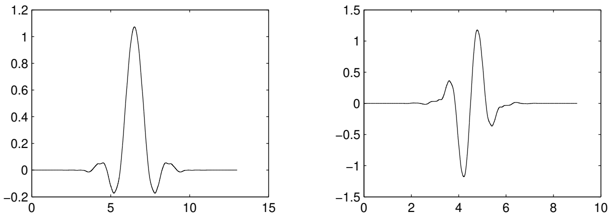

(a) (b)

Fig. 7. Synthesizing scaling and wavelet functions for the 6/14 transform. (a) Scaling function. (b) Wavelet function.

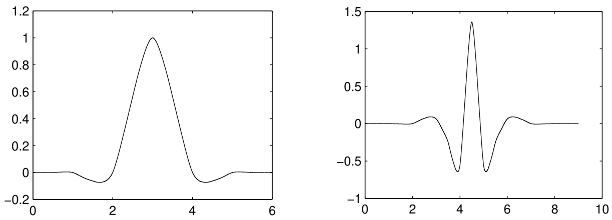

(a) (b)

Fig. 8. Synthesizing scaling and wavelet functions for the 13/7-T transform. (a) Scaling function. (b) Wavelet function.

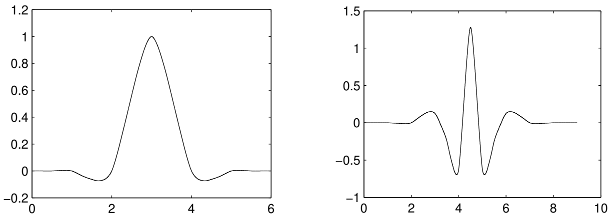

(a) (b)

Fig. 9. Synthesizing scaling and wavelet functions for the 13/7-C transform. (a) Scaling function. (b) Wavelet function.

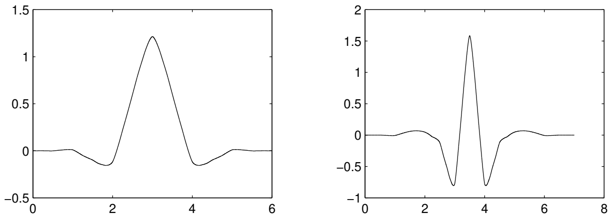

(a) (b)

Fig. 10. Synthesizing scaling and wavelet functions for the 9/7-F transform. (a) Scaling function. (b) Wavelet function.

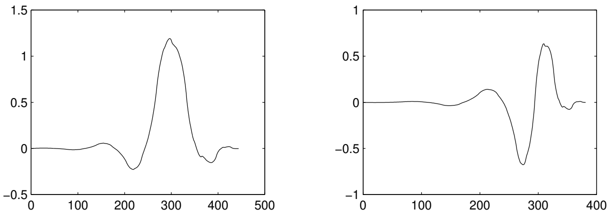

(a) (b)

Fig. 11. Synthesizing scaling and wavelet functions for the SPB transform. (a) Scaling function. (b) Wavelet function.

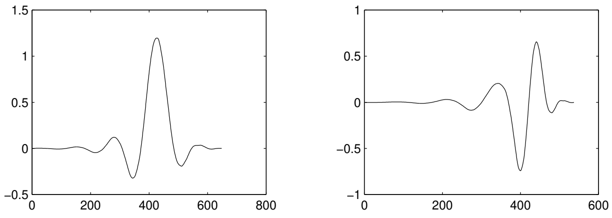

(a) (b)

Fig. 12. Synthesizing scaling and wavelet functions for the SPC transform. (a) Scaling function. (b) Wavelet function.

(a) (b)

Fig. 13. Part of the `cafe` image. Lossy reconstructed image at a bit rate of 0.25 bpp using the (a) conventional and (b) reversible integer-to-integer versions of the 5/11-A transform.

(a) (b)

Fig. 14. Part of the `cafe` image. Lossy reconstructed image at a bit rate of 0.25 bpp using the (a) conventional and (b) reversible integer-to-integer versions of the SPC transform.

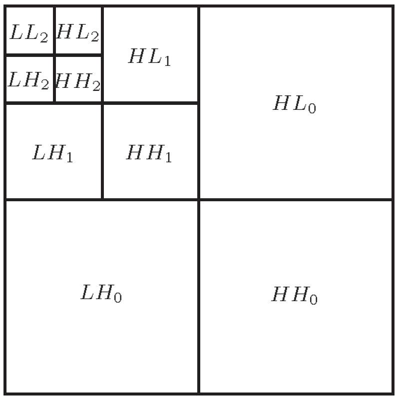

Fig. 15. Subband structure for the 2D wavelet transform formed by two 1D transforms.

TABLE II
FORWARD TRANSFORMS

5/3	$\begin{cases} d[n] = d_0[n] - \lfloor \frac{1}{2}(s_0[n+1] + s_0[n]) \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{4}(d[n] + d[n-1]) + \frac{1}{2} \rfloor \end{cases}$
2/6	$\begin{cases} d_1[n] = d_0[n] - s_0[n] \\ s[n] = s_0[n] + \lfloor \frac{1}{2} d_1[n] \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{4}(-s[n+1] + s[n-1]) + \frac{1}{2} \rfloor \end{cases}$
SPB	$\begin{cases} d_1[n] = d_0[n] - s_0[n] \\ s[n] = s_0[n] + \lfloor \frac{1}{2} d_1[n] \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{8}(-3s[n+1] + s[n] + 2s[n-1] + 2d_1[n+1]) + \frac{1}{2} \rfloor \end{cases}$
9/7-M	$\begin{cases} d[n] = d_0[n] + \lfloor \frac{1}{16}((s_0[n+2] + s_0[n-1]) - 9(s_0[n+1] + s_0[n])) + \frac{1}{2} \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{4}(d[n] + d[n-1]) + \frac{1}{2} \rfloor \end{cases}$
2/10	$\begin{cases} d_1[n] = d_0[n] - s_0[n] \\ s[n] = s_0[n] + \lfloor \frac{1}{2} d_1[n] \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{64}(22(s[n-1] - s[n+1]) + 3(s[n+2] - s[n-2])) + \frac{1}{2} \rfloor \end{cases}$
5/11-C	$\begin{cases} d_1[n] = d_0[n] - \lfloor \frac{1}{2}(s_0[n+1] + s_0[n]) \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{4}(d_1[n] + d_1[n-1]) + \frac{1}{2} \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{16}(s_1[n+2] - s_1[n+1] - s_1[n] + s_1[n-1]) + \frac{1}{2} \rfloor \end{cases}$
5/11-A	$\begin{cases} d_1[n] = d_0[n] - \lfloor \frac{1}{2}(s_0[n+1] + s_0[n]) \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{4}(d_1[n] + d_1[n-1]) + \frac{1}{2} \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{32}(s_1[n+2] - s_1[n+1] - s_1[n] + s_1[n-1]) + \frac{1}{2} \rfloor \end{cases}$
6/14	$\begin{cases} d_1[n] = d_0[n] - s_0[n] \\ s[n] = s_0[n] + \lfloor \frac{1}{16}(-d_1[n+1] + d_1[n-1] + 8d_1[n]) + \frac{1}{2} \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{16}(s_1[n+2] - s_1[n-2] + 6(-s_1[n+1] + s_1[n-1])) + \frac{1}{2} \rfloor \end{cases}$
SPC	$\begin{cases} d_1[n] = d_0[n] - s_0[n] \\ s[n] = s_0[n] + \lfloor \frac{1}{2} d_1[n] \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{16}(-8s[n+1] + 4s[n] + 5s[n-1] - s[n-2] + 6d_1[n+1]) + \frac{1}{2} \rfloor \end{cases}$
13/7-T	$\begin{cases} d[n] = d_0[n] + \lfloor \frac{1}{16}((s_0[n+2] + s_0[n-1]) - 9(s_0[n+1] + s_0[n])) + \frac{1}{2} \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{32}(-d[n+1] - d[n-2]) + 9(d[n] + d[n-1])) + \frac{1}{2} \rfloor \end{cases}$
13/7-C	$\begin{cases} d[n] = d_0[n] + \lfloor \frac{1}{16}(s_0[n+2] + s_0[n-1] - 9(s_0[n+1] + s_0[n])) + \frac{1}{2} \rfloor \\ s[n] = s_0[n] + \lfloor \frac{1}{16}(5(d[n] + d[n-1]) - (d[n+1] + d[n-2])) + \frac{1}{2} \rfloor \end{cases}$
9/7-F	$\begin{cases} d_1[n] = d_0[n] + \lfloor \frac{1}{128}(203(-s_0[n+1] - s_0[n])) + \frac{1}{2} \rfloor \\ s_1[n] = s_0[n] + \lfloor \frac{1}{4096}(217(-d_1[n] - d_1[n-1])) + \frac{1}{2} \rfloor \\ d[n] = d_1[n] + \lfloor \frac{1}{128}(113(s_1[n+1] + s_1[n])) + \frac{1}{2} \rfloor \\ s[n] = s_1[n] + \lfloor \frac{1}{4096}(1817(d_1[n] + d_1[n-1])) + \frac{1}{2} \rfloor \end{cases}$

TABLE III

TRANSFORM PARAMETERS. THE (A) PARAMETER VALUES, AND (B) PARAMETER DEFINITIONS

(a)

Transform	$\tilde{N}$	N	$\tilde{R}$	R	G_1	G_6	S_{H_0}	S_{H_1}
5/3	2	2	0.000	1.000	6.277	9.579	0.165	0.192
2/6	3	1	0.000	1.000	5.651	9.601	0.134	0.138
SPB	2	1	—	—	—	—	0.134	0.217
9/7-M	4	2	0.142	2.000	6.181	9.566	0.129	0.065
2/10	5	1	0.000	1.804	5.625	9.615	0.134	0.062
5/11-C	4	2	0.000	2.142	6.280	9.550	0.165	0.085
5/11-A	2	2	0.000	1.220	6.300	9.603	0.165	0.131
6/14	3	3	1.000	1.809	5.885	9.750	0.035	0.025
SPC	2	1	—	—	—	—	0.134	0.255
13/7-T	4	4	0.841	2.000	6.241	9.708	0.062	0.065
13/7-C	4	2	0.809	2.000	6.260	9.735	0.029	0.065
9/7-F	4	4	1.034	1.701	5.916	9.883	0.043	0.058

(b)

Parameter	Definition
$\tilde{N}$	number of vanishing moments of analyzing wavelet
N	number of vanishing moments of synthesizing wavelet
$\tilde{R}$	regularity of analyzing scaling function
R	regularity of synthesizing scaling function
G_1	coding gain for 1-level decomposition
G_6	coding gain for 6-level decomposition
S_{H_0}	frequency selectivity of analysis lowpass filter [17]
S_{H_1}	frequency selectivity of analysis highpass filter [17]

TABLE IV

COMPUTATIONAL COMPLEXITY

Transform	Adds	Shifts	Multiplies	Total
5/3	5	2	0	7
2/6	5	2	0	7
SPB	7 (8)	4 (3)	1 (0)	12 (11)
9/7-M	8 (9)	2 (3)	1 (0)	11 (12)
2/10	7 (10)	2 (6)	2 (0)	11 (16)
5/11-C	10	3	0	13
5/11-A	10	3	0	13
6/14	10 (11)	3 (5)	1 (0)	14 (16)
SPC	8 (10)	4 (5)	2 (0)	14 (15)
13/7-T	10 (12)	2 (4)	2 (0)	14 (16)
13/7-C	10 (12)	2 (4)	2 (0)	14 (16)
9/7-F	12 (26)	4 (18)	4 (0)	20 (44)

TABLE V
TEST IMAGES

Image	Size	Depth	Description
aerial2	2048×2048	8	aerial photograph
bike	2048×2560	8	collection of objects
cafe	2048×2560	8	outside of cafe
cats	3072×2048	8	cats
chart	1688×2347	8	color chart
cmpnd1	512×768	8	mixed text and graphics
cmpnd2	5120×6624	8	mixed text and graphics
cr	1744×2048	10	computer radiology
ct	512×512	12	computer tomography
elev	1201×1201	12	fractal-like pattern
finger	512×512	8	fingerprint
gold	720×576	8	houses and countryside
hotel	720×576	8	hotel
mat	1528×1146	8	mountains
mri	256×256	11	magnetic resonance
sar1	1024×1024	16	synthetic aperture radar
sar2	800×800	12	synthetic aperture radar
seismic	512×512	8	seismic data
target	512×512	8	patterns and textures
tools	1524×1200	8	collection of tools
txtur1	1024×1024	8	aerial photograph
txtur2	1024×1024	8	aerial photograph
us	512×448	8	ultrasound
water	1465×1999	8	water and land
woman	2048×2560	8	lady
xray	2048×1680	12	x-ray

TABLE VI

LOSSLESS COMPRESSION RESULTS

Image	Bit Rate (bpp)											
	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
aerial2	5.243	5.250	5.234	5.221	5.232	5.223	5.226	5.229	5.262	**5.217**	5.222	5.233
bike	4.415	4.437	4.417	4.406	4.427	4.407	4.401	4.429	4.454	**4.397**	4.402	4.456
cafe	5.194	5.215	5.191	5.177	5.203	5.178	**5.175**	5.217	5.228	**5.175**	5.186	5.220
cats	2.534	2.528	2.500	2.493	2.499	2.493	2.509	2.493	2.493	**2.491**	2.494	2.518
chart	**3.052**	3.095	3.088	3.103	3.115	3.094	3.063	3.176	3.156	3.104	3.117	3.225
cmpnd1	**2.142**	2.143	2.216	2.454	2.368	2.487	2.403	2.578	2.362	2.492	2.523	2.540
cmpnd2	**2.540**	2.605	2.655	2.753	2.733	2.770	2.686	2.855	2.790	2.775	2.797	2.906
cr	**3.216**	3.225	3.230	3.225	3.230	3.225	3.219	3.220	3.262	3.217	**3.216**	3.221
ct	2.886	2.890	2.739	2.654	2.710	2.663	2.784	2.727	**2.578**	2.664	2.677	2.778
elev	2.453	2.464	2.408	2.320	2.358	**2.317**	2.380	2.378	2.350	2.326	2.335	2.530
finger	5.470	5.550	5.423	5.321	5.415	5.323	5.381	5.380	**5.312**	5.328	5.337	5.372
gold	4.476	4.502	4.497	4.476	4.498	4.471	**4.467**	4.490	4.539	4.473	4.477	4.518
hotel	4.458	4.475	4.476	4.458	4.469	4.449	**4.440**	4.466	4.528	4.448	4.453	4.493
mat	2.994	3.102	3.110	3.036	3.116	2.984	**2.983**	3.173	3.164	3.049	3.081	3.270
mri	4.022	4.034	3.939	3.926	3.960	3.926	3.970	3.954	**3.879**	3.925	3.931	3.971
sar1	12.68	12.69	12.68	12.68	12.69	12.68	**12.67**	12.69	12.73	**12.67**	12.68	**12.67**
sar2	**7.610**	7.625	7.626	7.615	7.627	7.616	**7.610**	7.627	7.662	7.612	7.616	7.611
seismic	2.781	2.807	2.704	2.610	2.684	2.654	2.724	2.653	**2.577**	2.613	2.615	2.879
target	**2.129**	2.231	2.234	2.256	2.313	2.247	2.213	2.355	2.242	2.250	2.243	2.546
tools	5.281	5.293	5.272	5.260	5.283	5.263	5.261	5.289	5.308	**5.256**	5.263	5.286
txtur1	6.581	6.579	6.590	6.584	6.586	6.586	6.578	6.585	6.644	**6.577**	6.581	6.578
txtur2	5.427	5.426	5.430	5.421	5.433	5.425	**5.418**	5.440	5.486	**5.418**	5.425	5.435
us	**3.031**	3.054	3.078	3.169	3.111	3.141	3.104	3.211	3.162	3.167	3.192	3.473
water	**3.325**	3.334	3.368	3.357	3.354	3.338	**3.325**	3.333	3.436	3.337	3.331	3.408
woman	4.394	4.407	4.382	4.366	4.387	4.367	4.372	4.385	4.409	**4.362**	4.368	4.420
x_ray	4.408	4.418	4.408	4.404	4.414	4.407	4.405	4.414	4.435	**4.400**	4.403	4.410
Mean†	4.599	4.623	4.607	4.583	4.608	4.577	4.580	4.608	4.636	4.579	4.586	4.631

†Mean over natural 8-bpp images (i.e., aerial2, bike, cafe, cats, finger, gold, hotel, mat, tools, txtur1, txtur2, water, woman).

TABLE VII

LOSSLESS COMPRESSION RESULTS RELATIVE TO THE BEST TRANSFORM FOR EACH IMAGE

Image	Increase in Bit Rate (%)											
	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
chart	—	1.4	1.2	1.7	2.1	1.4	0.4	4.1	3.4	1.7	2.1	5.7
cmpnd1	—	0.0	3.5	14.6	10.6	16.1	12.2	20.4	10.3	16.3	17.8	18.6
cmpnd2	—	2.6	4.5	8.4	7.6	9.1	5.7	12.4	9.8	9.3	10.1	14.4
ct	11.9	12.1	6.2	2.9	5.1	3.3	8.0	5.8	—	3.3	3.8	7.8
elev	5.9	6.3	3.9	0.1	1.8	—	2.7	2.6	1.4	0.4	0.8	9.2
finger	3.0	4.5	2.1	0.2	1.9	0.2	1.3	1.3	—	0.3	0.5	1.1
mat	0.4	4.0	4.3	1.8	4.5	0.0	—	6.4	6.1	2.2	3.3	9.6
mri	3.7	4.0	1.5	1.2	2.1	1.2	2.3	1.9	—	1.2	1.3	2.4
seismic	7.9	8.9	4.9	1.3	4.2	3.0	5.7	2.9	—	1.4	1.5	11.7
target	—	4.8	4.9	6.0	8.6	5.5	3.9	10.6	5.3	5.7	5.4	19.6
us	—	0.8	1.6	4.6	2.6	3.6	2.4	5.9	4.3	4.5	5.3	14.6
water	—	0.3	1.3	1.0	0.9	0.4	—	0.2	3.3	0.4	0.2	2.5

TABLE VIII

LOSSY COMPRESSION RESULTS FOR (A) AERIAL2, (B) BIKE, (C) CAFE, (D) SAR2, AND (E) TARGET IMAGES

(a)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	24.40	24.47	24.05	24.46	24.55	24.41	24.45	24.66	23.63	24.60	24.64	**24.72**
0.1250	26.24	26.31	25.83	26.36	26.42	26.30	26.33	26.53	25.35	26.50	26.54	**26.59**
0.2500	28.15	28.25	27.49	28.35	28.37	28.26	28.25	28.53	26.91	28.50	28.55	**28.60**
0.5000	30.36	30.37	29.25	30.49	30.48	30.47	30.45	30.68	28.54	30.62	30.67	**30.69**
1.0000	33.30	33.30	31.89	33.35	33.31	33.34	33.36	**33.54**	30.97	33.50	33.53	33.50
2.0000	38.55	38.51	36.74	38.46	38.52	38.45	38.53	38.56	35.72	38.56	**38.58**	38.13

(b)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	22.49	22.63	22.16	22.61	22.70	22.59	22.52	22.88	21.60	22.78	22.85	**22.95**
0.1250	25.03	25.26	24.63	25.12	25.32	25.03	25.08	25.47	24.02	25.34	25.41	**25.58**
0.2500	28.27	28.34	27.84	28.31	28.37	28.21	28.31	28.56	27.17	28.57	28.62	**28.76**
0.5000	32.27	32.24	31.59	32.23	32.18	32.13	32.29	32.33	30.77	32.30	32.51	**32.57**
1.0000	36.97	36.87	35.60	36.85	36.74	36.82	36.95	36.87	34.76	37.04	**37.06**	36.83
2.0000	**42.77**	42.57	40.55	42.42	42.43	42.48	42.70	42.35	39.49	42.47	42.42	41.12

(c)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	18.77	18.89	18.17	18.78	18.88	18.72	18.77	18.99	17.61	18.92	18.97	**19.04**
0.1250	20.29	20.47	19.77	20.30	20.45	20.25	20.29	20.59	19.14	20.45	20.50	**20.62**
0.2500	22.75	22.90	22.10	22.75	22.89	22.72	22.77	23.03	21.39	22.94	22.98	**23.12**
0.5000	26.33	26.31	25.56	26.36	26.30	26.32	26.39	26.43	24.81	26.55	26.56	**26.64**
1.0000	31.41	31.30	30.59	31.42	31.26	31.37	31.47	31.38	29.76	31.62	31.62	**31.63**
2.0000	**38.62**	38.42	36.98	38.41	38.32	38.39	38.58	38.33	35.98	38.52	38.49	38.06

(d)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	22.21	22.33	21.77	22.22	22.34	22.18	22.21	22.41	21.31	22.32	22.34	**22.42**
0.1250	22.97	23.09	22.28	22.98	23.09	22.94	22.97	**23.18**	21.76	23.07	23.10	**23.18**
0.2500	23.99	24.12	23.06	23.95	24.13	23.93	23.97	24.25	22.44	24.09	24.13	**24.26**
0.5000	25.54	25.63	24.43	25.50	25.60	25.47	25.53	25.80	23.65	25.64	25.70	**25.82**
1.0000	28.22	28.44	27.08	28.08	28.35	28.12	28.20	28.54	26.03	28.32	28.38	**28.55**
2.0000	33.85	33.86	32.26	33.66	33.74	33.63	33.79	33.98	31.04	33.86	33.90	**34.08**

(e)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	17.57	17.30	17.04	17.55	17.34	17.46	17.54	18.17	16.41	18.12	18.36	**18.95**
0.1250	19.89	19.44	18.12	20.20	19.34	19.99	19.98	20.45	17.54	20.73	**21.05**	20.97
0.2500	23.28	23.01	22.75	23.64	23.18	23.36	23.41	24.76	22.31	24.38	**25.48**	24.92
0.5000	29.98	29.30	28.81	30.73	29.74	30.64	30.55	31.15	28.49	31.76	**32.51**	31.82
1.0000	41.38	40.03	37.99	41.25	39.79	41.04	40.99	40.68	37.69	41.70	**41.94**	40.21
2.0000	54.14	**54.79**	50.79	52.26	52.91	52.07	52.37	52.22	48.42	52.13	51.78	44.60

TABLE IX

OVERALL SUBJECTIVE RANKINGS FROM FIRST STAGE OF TESTING

Rank							
5/3	2/6	SPB	9/7-M	2/10	SPC	13/7-T	9/7-F
2.51	4.47	6.05	2.97	4.57	6.88	2.88	2.36

TABLE X

OVERALL SUBJECTIVE RANKINGS FROM SECOND STAGE OF TESTING

Rank						
5/3	9/7-M	5/11-C	5/11-A	6/14	13/7-T	13/7-C
1.84	2.30	2.34	2.15	2.75	1.77	1.83

TABLE XI

DIFFERENCE IN PSNR PERFORMANCE BETWEEN REVERSIBLE INTEGER-TO-INTEGER AND CONVENTIONAL VERSIONS OF TRANSFORMS FOR THE (A) AERIAL2, (B) BIKE, (C) CAFE, (D) SAR2, AND (E) TARGET IMAGES

(a)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	-0.03	-0.04	-0.54	-0.05	-0.04	-0.03	-0.03	-0.05	-0.69	-0.03	-0.03	-0.03
0.1250	-0.05	-0.06	-0.64	-0.06	-0.07	-0.06	-0.04	-0.06	-0.84	-0.05	-0.02	-0.08
0.2500	-0.14	-0.09	-0.96	-0.10	-0.10	-0.15	-0.07	-0.10	-1.29	-0.09	-0.09	-0.08
0.5000	-0.09	-0.14	-1.31	-0.13	-0.16	-0.13	-0.11	-0.15	-1.63	-0.15	-0.15	-0.16

(b)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	-0.02	-0.02	-0.52	-0.02	-0.03	-0.02	-0.02	-0.02	-0.70	-0.02	-0.02	-0.01
0.1250	-0.04	-0.06	-0.75	-0.05	-0.05	-0.05	-0.05	-0.05	-0.94	-0.04	-0.05	-0.02
0.2500	-0.08	-0.09	-0.71	-0.10	-0.10	-0.10	-0.08	-0.09	-0.87	-0.09	-0.10	-0.07
0.5000	-0.16	-0.21	-0.99	-0.20	-0.23	-0.21	-0.16	-0.22	-1.16	-0.36	-0.22	-0.24

(c)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	-0.01	-0.01	-0.62	-0.01	-0.01	-0.01	-0.01	-0.01	-0.84	-0.01	-0.00	-0.01
0.1250	-0.01	-0.02	-0.68	-0.01	-0.02	-0.01	-0.01	-0.01	-0.86	-0.01	-0.01	-0.01
0.2500	-0.02	-0.03	-0.82	-0.04	-0.03	-0.03	-0.03	-0.02	-1.00	-0.03	-0.02	0.02
0.5000	-0.05	-0.06	-0.91	-0.06	-0.07	-0.07	-0.06	-0.07	-0.94	-0.05	-0.07	-0.07

(d)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	-0.01	-0.00	-0.54	-0.00	-0.00	0.00	-0.00	0.00	-0.79	-0.00	0.00	-0.00
0.1250	-0.00	-0.00	-0.76	-0.00	-0.00	0.00	-0.00	-0.00	-1.05	-0.00	-0.00	-0.00
0.2500	0.03	-0.00	-0.97	0.00	-0.00	-0.00	0.00	-0.00	-1.24	0.00	-0.00	0.00
0.5000	-0.00	-0.03	-1.10	-0.00	-0.00	-0.00	-0.00	-0.00	-1.37	-0.00	-0.00	0.00
1.0000	-0.00	0.00	-1.15	-0.07	0.00	-0.00	-0.00	-0.00	-1.36	-0.00	0.07	-0.00
2.0000	-0.00	-0.01	-1.27	-0.00	-0.00	-0.00	-0.01	-0.00	-1.39	-0.00	0.00	0.00

(e)

Bit Rate	PSNR (dB)											
(bpp)	5/3	2/6	SPB	9/7-M	2/10	5/11-C	5/11-A	6/14	SPC	13/7-T	13/7-C	9/7-F
0.0625	0.01	-0.05	-0.34	-0.01	-0.03	0.01	0.02	0.08	-0.69	-0.01	-0.01	0.05
0.1250	0.00	0.03	-1.01	-0.00	-0.14	-0.03	-0.02	-0.06	-1.12	-0.00	-0.11	-0.07
0.2500	0.00	-0.35	-0.28	-0.02	-0.12	-0.03	-0.03	-0.01	-0.87	-0.02	-0.03	-0.03
0.5000	-0.05	-0.25	-1.20	-0.11	-0.14	-0.14	-0.13	-0.11	-1.55	-0.20	-0.17	-0.33

TABLE XII

DIFFERENCE IN SUBJECTIVE PERFORMANCE BETWEEN REVERSIBLE INTEGER-TO-INTEGER AND CONVENTIONAL VERSIONS OF TRANSFORMS

Rank							
5/3	2/6	SPB	9/7-M	2/10	SPC	13/7-T	9/7-F
-0.20	-0.41	-0.80	-0.29	-0.36	-0.77	-0.37	-0.24

TABLE XIII

INFLUENCE OF THE NUMBER OF LIFTING STEPS ON LOSSLESS COMPRESSION PERFORMANCE

Image	2 lifting steps		4 lifting steps	
	MAE†	Bit Rate (bpp)	MAE†	Bit Rate (bpp)
gold	0.409	4.473	0.540	4.479
target	0.216	2.250	0.295	2.311
woman	0.415	4.362	0.536	4.373

†mean absolute error

TABLE XIV

PARAMETERS OF TRANSFORMS AFFECTING COMPRESSION PERFORMANCE

Transform	# of Lifting Steps	$\|\|h_0\|\|_1$	$\|\|h_1\|\|_1$	$\|H_0(e^{j0})\|$	$\|H_1(e^{j\pi})\|$
5/3	2	1.500	2.000	1.000	2.000
2/6	3	1.000	2.500	1.000	2.000
SPB	3	1.000	2.750	1.000	2.500
9/7-M	2	1.500	2.250	1.000	2.000
2/10	3	1.000	2.781	1.000	2.000
5/11-C	3	1.500	2.219	1.000	2.000
5/11-A	3	1.500	2.109	1.000	2.000
6/14	3	1.250	2.797	1.000	2.000
SPC	3	1.000	3.125	1.000	2.750
13/7-T	2	1.625	2.250	1.000	2.000
13/7-C	2	1.750	2.250	1.000	2.000
9/7-F	4	1.698	2.109	1.230	1.625

TABLE XV

AVERAGE COEFFICIENT MAGNITUDES AND LOSSLESS BIT RATES FOR THE 5/3 AND 2/6 TRANSFORMS IN THE CASE OF THE (A) GOLD, (B) TARGET, AND (C) CMPND1 IMAGES

(a)

Transform	Avg. Coef. Magnitude				Bit Rate (bpp)
	LH_0	HL_0	HH_0	All	
5/3	4.489	4.259	3.785	5.056	4.476
2/6	4.841	4.796	5.488	5.287	4.502

(b)

Transform	Avg. Coef. Magnitude				Bit Rate (bpp)
	LH_0	HL_0	HH_0	All	
5/3	15.951	15.874	12.165	15.030	2.129
2/6	18.073	18.206	19.705	17.312	2.231

(c)

Transform	Avg. Coef. Magnitude				Bit Rate (bpp)
	LH_0	HL_0	HH_0	All	
5/3	8.834	7.695	7.218	8.887	2.142
2/6	8.610	7.788	9.836	8.641	2.143