Reversible Wavelet and Spectral Transforms for Lossless Compression of Color Images

N. Strobel, S. K. Mitra, and B. S. Manjunath
Signal and Image Processing Laboratory
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
Email: strobel@nt.e-technik.uni-erlangen.de

Abstract

A new lossless image compression method for progressive-resolution transmission of color images is proposed. It is based on spatial and spectral transforms. Reversible wavelet transforms are performed across the red, green, and blue color components first. Then adaptive spectral transforms are applied to associated color subbands. Simulation results indicate that the proposed algorithm can achieve bit rates that are about 20% lower than results obtained with comparable lossless image compression techniques supporting progressive-resolution transmission.

1 Introduction

Recent years have seen a tremendous increase in the generation, transmission, and storage of color images. Although significant progress has been made for lossy image compression [1, 2], fewer advances have been reported for lossless techniques.

Color images are normally represented as RGB spectral components. As a result, two sources of redundancy exist — *spatial redundancy* and *spectral redundancy*.

Traditional color image compression methods typically apply a spectral decorrelation across the color components first. Then spatial transforms are employed to decorrelate individual spectral bands further [3, 4]. If spectral and spatial transforms are carried out independently, their order is insignificant and can be reversed. This offers the opportunity to apply different spectral transforms to associated color subbands sharing the same scale and orientation. The result is an effective lossless image compression algorithm.

This paper is organized as follows. First, Section 2 briefly outlines how reversible wavelet transforms are used to spatially decorrelate color components. Section 3 then explains the spectral decorrelation technique. In Section 4, experimental results are presented. Concluding remarks are given in Section 5.

2 Spatial Decorrelation of Color Bands

For lossless (or reversible) image compression, it is important to represent transform coefficients with integer numbers. As a result, for progressive-resolution transmission applications, reversible wavelet transforms (RWTs), such as the S-transform [5] or the SP-transform [6], are often used.

The one-dimensional S-transform, for example, reduces the necessary word length by making intelligent use of rounding operations. It successively decimates some input sequence $s_k[n]$ at resolution $r = 2^{-k}$ into truncated average or coarse versions $s_{k+1}[n]$ and associated difference or detail signals $d_{k+1}[n]$ at resolution $r = 2^{-k-1}$. Applying the S-transform to an input signal $s[n] = s_0[n]$ at resolution $r = 1$, we obtain resolutions that are negative powers of two only, i.e., $r = 2^{-k}$, $k > 0$. Further information about the notions of resolution and scale of discrete-time signals can be found in Rioul [7]. Note that a wavelet transform can also be viewed as a pyramid decomposition scheme. Then the subscript k indicates the pyramid level. Due to the downward truncation, denoted by $\lfloor . \rfloor$, the maximum number of bits required for local averages does not change. The detail signals (wavelet coefficients), which are composed of positive and negative integers, on the other hand, require a signed representation. Their number precision, thus, exceeds the original storage format. The significantly lower entropies of difference signals, however, compensate for their longer internal word lengths, since they facilitate the use of efficient coding methods [8].

A two-dimensional S-transform can be obtained by applying the 1-D S-transform sequentially to the rows and columns of a color image. In this case (truncated) average (or LL) bands at successively lower resolutions

0-8186-8821-1/98 $10.00

are recursively computed by

$$s_{k+1}[m,n,l] = \lfloor \frac{\bar{s}_k[2m,l] + \bar{s}_k[2m+1,l]}{2} \rfloor, \quad (1)$$

where the (integer) averages $\bar{s}_k[r,l]$ along row r and in color channel l are computed via

$$\bar{s}_k[r,l] = \lfloor \frac{s_k[r,2n,l] + s_k[r,2n+1,l]}{2} \rfloor. \quad (2)$$

Wavelet coefficients follow as associated directional differences.

In Eq. (2), the sample $s_k[m,n,l]$ describes a color pixel at row m, column n, and spectral band l, observed at resolution $r = 2^{-k}$. The red, green and blue color bands are specified by indices $l \in \{1,2,3\}$, respectively. For brevity, the matrix of all pixels in the l-th color band at resolution 2^{-k}, $0 \le k \le K$, is denoted as

$$\mathbf{s}_k[l] = \{s_k[m,n,l] \mid 0 \le m < \frac{M}{2^k}, 0 \le n < \frac{N}{2^k}\}. \quad (3)$$

The parameters M and N are image height and width, respectively. For simplicity, it is assumed that $M = N = 2^K$. This facilitates a K-level wavelet transform.

The rounding operations introduce a nonlinearity into the S-transform which produces a noteworthy side-effect. Since both row and column averages are truncated, fractional parts are always discarded. As a result, the transform becomes biased, i.e., integer scaling coefficients at progressively lower resolutions get increasingly smaller than the true local averages. Although this is a minor side-effect of the S-transform, the balanced rounding or BR-transform offers a simple yet effective solution to this problem. It compensates for the roundoff error by rounding up along image rows while truncating along image columns [9].

3 Spectral Decorrelation of Subband Channels

Figure 1 illustrates the proposed method. There, a reversible wavelet transform is first applied to each color band $\mathbf{s}_k[l]$ at resolution 2^{-k}. This yields three transform matrices

$$\mathbf{S}_{k+1}[l] = \{S_{k+1}[u,v,l] \mid 0 \le u < \frac{M}{2^k}, 0 \le v < \frac{N}{2^k}\}.$$

where $l \in \{1,2,3\}$. Applying a reversible spectral transform (ST) to $\mathbf{S}_{k+1}[l]$, $0 \le l \le 3$, we obtain associated prediction errors denoted as

$$\mathbf{e}_{k+1}[l] = \{e_{k+1}[u,v,l] \mid 0 \le u < \frac{M}{2^k}, 0 \le v < \frac{N}{2^k}\}.$$

The combination of inverse spectral transform (ST^{-1}) and inverse reversible wavelet transforms (RWT^{-1}) finally reconstructs the original RGB color channels exactly.

For a particular color l, the transform matrix $\mathbf{S}_{k+1}[l]$ can either be considered as an ensemble of transform coefficients or be viewed as a collection of four oriented subbands. Adopting the second point of view, we describe a color subband with orientation η at resolution 2^{-k-1}, as

$$\mathbf{S}^{(\eta)}_{k+1}[l] = \{S^{(\eta)}_{k+1}[u,v,l] \mid 0 \le u < \frac{M}{2^{k+1}}, 0 \le v < \frac{N}{2^{k+1}}\},$$

with $1 \le l \le 3$ and $\eta \in \{\mathrm{LL}, \mathrm{LH}, \mathrm{HL}, \mathrm{HH}\}$. The letters stand for low (L) and high (H) bands corresponding to the separable application of lowpass or highpass filters along the rows and columns, respectively. Consequently, the LL-band of $\mathbf{S}_{k+1}[l]$ is called $\mathbf{S}^{(\mathrm{LL})}_{k+1}[l]$, the LH-band represented by $\mathbf{S}^{(\mathrm{LH})}_{k+1}[l]$, the HL-band referred to as $\mathbf{S}^{(\mathrm{HL})}_{k+1}[l]$, and the HH-band finally denoted $\mathbf{S}^{(\mathrm{HH})}_{k+1}[l]$.

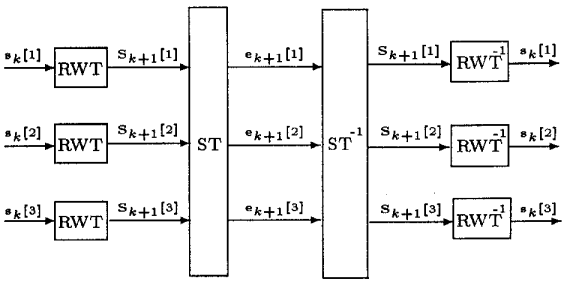

Figure 1: Proposed color decorrelation method.

Performing a K-level wavelet transform on $\mathbf{s}_0[1]$, the red band of the input color image with size $2^K \times 2^K$, we obtain a total of $3K+1$ oriented red subbands. They comprise $3K$ channels with wavelet coefficients and one lowpass coefficient representing the mean of the red spectral band. Applying the same RWT to the remaining green and blue color bands, we finally get sets of associated red, green, and blue subbands which can be effectively spectrally decorrelated.

Since there are potentially as many different spectral transforms for a K-level wavelet transform of color images as there are different subbands, it is normally no longer possible to switch the order of the spatial and spectral transforms. Instead, we gain the opportunity to apply an adaptive spectral decorrelation method. It is based on interband prediction as presented next.

3.1 Interband Prediction Coefficients

For three color bands, at most a two-band predictor is needed to predict the third subband from the remaining two. After interband prediction, each color subband coefficient in the l-th band is replaced by its difference with respect to the linear combination of the remaining spectral neighbors. The two-band prediction $\hat{\mathbf{S}}_{k+1}^{(\eta)}[l]$ for the l-th color subband at resolution 2^{-k-1} and orientation η can be compactly expressed as

$$\hat{\mathbf{S}}_{k+1}^{(\eta)}[l] = \alpha_1 \mathbf{S}_{k+1}^{(\eta)}[i] + \alpha_2 \mathbf{S}_{k+1}^{(\eta)}[j] + \bar{\mathbf{S}}_{k+1}^{(\eta)}[l], \qquad (4)$$

with $i \neq j \neq l$, and $i, j, l \in \{1, 2, 3\}$. $\mathbf{S}_{k+1}^{(\eta)}[i]$ and $\mathbf{S}_{k+1}^{(\eta)}[j]$ are the neighboring color subbands, while $\bar{\mathbf{S}}_{k+1}^{(\eta)}[l]$ refers to the mean of the l-th color subband. This third-order prediction model gives rise to the following procedure for computing integer subband residuals

$$\mathbf{e}_{k+1}^{(\eta)}[l] = \{e_{k+1}^{(\eta)}[u, v, l] \,|\, 0 \leq u < \frac{M}{2^{k+1}}, 0 \leq v < \frac{N}{2^{k+1}}\}$$

at spectral location l:

1. Compute prediction:

$$\begin{aligned} \hat{\mathbf{S}}_{k+1}^{(\eta)}[l] &= \left[\alpha_1 \left(\mathbf{S}_{k+1}^{(\eta)}[i] - \left[\bar{\mathbf{S}}_{k+1}^{(\eta)}[i]\right]_R\right)\right. \\ &+ \left. \alpha_2 \left(\mathbf{S}_{k+1}^{(\eta)}[j] - \left[\bar{\mathbf{S}}_{k+1}^{(\eta)}[j]\right]_R\right)\right]_R \\ &+ \left[\bar{\mathbf{S}}_{k+1}^{(\eta)}[l]\right]_R, \end{aligned}$$

   where $\bar{\mathbf{S}}_{k+1}^{(\eta)}[.] = \mathrm{E}\{S_{k+1}^{(\eta)}[.]\}$. For high-frequency subbands (wavelet coefficients), we have

$$\bar{\mathbf{S}}_{k+1}^{(\eta)}[i] = \bar{\mathbf{S}}_{k+1}^{(\eta)}[j] = \bar{\mathbf{S}}_{k+1}^{(\eta)}[l] = 0.$$

   Integer values are enforced by using the rounding operator $[.]_R$.

2. Compute prediction error:

$$\mathbf{e}_{k+1}^{(\eta)}[l] = \mathbf{S}_{k+1}^{(\eta)}[l] - \hat{\mathbf{S}}_{k+1}^{(\eta)}[l].$$

   The error subband $\mathbf{e}_{k+1}^{(\eta)}[l]$ comprises differences between actual and predicted color subband coefficients at the same spatial locations.

3. Encode prediction error and include all necessary side information such as prediction coefficients. Then store or transmit it.

The prediction coefficients α_1 and α_2 are obtained by straightforward application of least-squares regression formulas. While the above derivation is for the two-band prediction scheme, it is easily specialized to single-band prediction.

3.2 Interband Prediction Order

Reversible linear prediction must be implemented such that it can be resolved based on the information already received. Since lossless prediction involves nonlinear rounding operations, color subband decorrelation must be carried out sequentially. To this end, an *anchor* band has to be specified first. It serves as a reference for predicting the second color subband. Finally the first two subbands provide the basis from which to predict the remaining third. A prediction order must be found such that the overall entropy after color subband prediction is minimized. If we restrict ourselves to single-band prediction, then this problem can be modeled into a graph-theoretic problem [10]. While such an approach holds some promise for multispectral images with hundreds of different bands, better color compression results are obtained when two-band prediction is also considered. Then a scheme for color subband decorrelation leads to $3! = 6$ different scenarios. For example, the green subband component can be used as an anchor to predict the red, then red and green can be employed to predict the blue subband coefficients. A solution to the problem of finding the best order out of the six possibilities is presented below. It determines the prediction order such that the approximated sum of subband entropies after prediction is smallest.

It can be shown that an approximation of the first-order entropy of a color subband is given by [11]

$$H(X) = \frac{1}{2} \log_2(\gamma_x \sigma_x^2). \qquad (5)$$

Equation (5) provides an entropy estimate for color subbands based on their shape factor, γ_x, and their variance σ_x^2.

To obtain error variances, we select the anchor subband first. Let it be denoted as $\mathbf{S}_{k+1}^{(\eta)}[i]$. Subtracting the associated rounded mean $\left[\bar{\mathbf{S}}_{k+1}^{(\eta)}[i]\right]_R$, we get the first error subband $\mathbf{e}_{k+1}^{(\eta)}[i]$. Next, the anchor band is used to predict the second color subband $\mathbf{S}_{k+1}^{(\eta)}[j]$. The resulting prediction error is called $\mathbf{e}_{k+1}^{(\eta)}[j]$. Third, $\mathbf{S}_{k+1}^{(\eta)}[i]$ and $\mathbf{S}_{k+1}^{(\eta)}[j]$ are combined to estimate $\mathbf{S}_{k+1}^{(\eta)}[l]$. This two-step prediction yields the difference band $\mathbf{e}_{k+1}^{(\eta)}[l]$. Finally the variances of the three error subbands are computed. They are:

$$\begin{aligned} \mathrm{var}\{\mathbf{e}_{k+1}^{(\eta)}[i]\} &= \mathrm{E}\{(\mathbf{S}_{k+1}^{(\eta)}[i] - \left[\bar{\mathbf{S}}_{k+1}^{(\eta)}[i]\right]_R)^2\}, \\ \mathrm{var}\{\mathbf{e}_{k+1}^{(\eta)}[j]\} &= \mathrm{E}\{(\mathbf{S}_{k+1}^{(\eta)}[j] - \left[\hat{\mathbf{S}}_{k+1}^{(\eta)}[j]\right]_R)^2\}, \text{ and} \\ \mathrm{var}\{\mathbf{e}_{k+1}^{(\eta)}[l]\} &= \mathrm{E}\{(\mathbf{S}_{k+1}^{(\eta)}[l] - \left[\hat{\mathbf{S}}_{k+1}^{(\eta)}[l]\right]_R)^2\}. \end{aligned}$$

Note that $\mathrm{var}\{e_{k+1}^{(\eta)}[i]\}$ is associated with a zero-mean color subband, while $\mathrm{var}\{e_{k+1}^{(\eta)}[j]\}$ and $\mathrm{var}\{e_{k+1}^{(\eta)}[l]\}$ result from prediction residuals.

Once the variances have been computed, entropies of their associated subbands are estimated. The sum of entropies of all three transformed subbands at resolution 2^{-k-1} and orientation η is called $H_{k+1}^{(\eta)}$. According to Eq. (5), it can be approximated by

$$H_{k+1}^{(\eta)} = \frac{1}{2}\Big(\log_2(\gamma_i\gamma_j\gamma_l) + \log_2(\mathrm{var}\{e_{k+1}^{(\eta)}[i]\}\,\mathrm{var}\{e_{k+1}^{(\eta)}[j]\}\,\mathrm{var}\{e_{k+1}^{(\eta)}[l]\})\Big).$$

The shape factors γ_i, γ_j, and γ_l are associated with the pdfs of $e_{k+1}^{(\eta)}[i]$, $e_{k+1}^{(\eta)}[j]$, and $e_{k+1}^{(\eta)}[l]$, respectively.

Each prediction order yields a different value for $H_{k+1}^{(\eta)}$. The best ordering is found by selecting the prediction sequence resulting in the smallest value for $H_{k+1}^{(\eta)}$. For simplicity, we assume that the product of shape factors remains constant regardless of the prediction sequence chosen. The underlying assumption is that the overall statistical character of the prediction errors remains the same regardless of the prediction order. Since the logarithm is monotonically increasing, we only need to compare products of error variances.

4 Experimental Results

Experiments are carried out using the six test images Barbara, Boats, Fruits, Girl, Goldhill, and Zelda. Each consists of three color channels in RGB format. Except Fruits, which is a 512×512 image, the spatial dimensions of the other images are 720×576. Three six-level different RWTs are performed. The RWTs used are associated with the S-transform and the TT-transform. After spatial and spectral decorrelation, the resulting color transform subbands are individually entropy encoded. To that end, a context-based arithmetic coder is employed. It is based on the implementation suggested in [6].

Experimental results are compared based on the average bit rate per pixel (bpp), simply denoted by R. R is based on the compressed file size and takes into account all the side information necessary to losslessly reconstruct the original image. It is defined as

$$R := \frac{\text{total file length incl. overhead}}{\text{no. pixels}}\ [\text{bpp}]. \qquad (6)$$

The bit rates obtained with the S-transform followed by spectral decorrelation is called $R_{\mathrm{ST}}^{(\mathrm{S})}$. Outcomes associated with the TT filter band are labelled $R_{\mathrm{ST}}^{(\mathrm{TT})}$. The results are shown in Table 1. There lowest bit rates are printed in bold. The second column represents the sum of the original color band entropies. It is called H_{orig}. The last column lists state-of-the-art lossless color compression results reported by Gormish et al. [12]. They are denoted as $R_{\mathrm{col}}^{(\mathrm{CREW})}$.

Table 1: Color Compression Results

Images	H_{orig}	$R_{\mathrm{ST}}^{(\mathrm{S})}$	$R_{\mathrm{ST}}^{(\mathrm{TT})}$	$R_{\mathrm{col}}^{(\mathrm{CREW})}$
Barbara	22.82	**9.664**	12.080	12.050
Boats	21.35	**8.503**	10.929	10.918
Fruits	8.74	**6.832**	7.182	7.181
Girl	21.73	**8.820**	11.115	11.113
Goldhill	22.22	**9.526**	12.540	12.554
Zelda	22.72	**8.286**	10.782	10.870

A comparison illustrates that the lowest bit rates are achieved with the simple S-transform.

Although this result might appear unusual, a look at Table 2 offers some insight. Shown are linear correlation coefficients among HH subband coefficients at resolution 2^{-1}. The subbands were computed using the S-transform. In Table 2, the parameter $\rho_{12}^{(\mathrm{HH})}$ refers to the correlation between the red and green subband coefficients in the HH band, $\rho_{23}^{(\mathrm{HH})}$ denotes the correlation between green and blue HH subband coefficients, and $\rho_{13}^{(\mathrm{HH})}$ specifies their red-blue counterpart.

Table 2: Correlations Among Color Subbands

	$\rho_{12}^{(\mathrm{HH})}$	$\rho_{23}^{(\mathrm{HH})}$	$\rho_{13}^{(\mathrm{HH})}$
Barbara	0.993	0.993	0.998
Boats	0.999	0.999	0.999
Fruits	0.695	0.667	0.458
Girl	0.999	0.995	0.996
Goldhill	0.999	0.999	0.999
Zelda	0.999	0.998	0.999

For almost all images we find that spectrally adjacent gradients in neighboring color channels are extremely strongly correlated. It is precisely this effect that facilitates an effective spectral decorrelation yielding unexpectedly good lossless compression rates. Only "Fruits" stands out as an exception to the rule. It

has large areas of saturated colors with little spectral overlap.

5 Discussion and Conclusions

At this point a warning is in order. Spectral decorrelation seems only beneficial if there exists a substantial amount of correlation between the color components to start with. Otherwise, little improvement beyond spatial decorrelation should be expected. For the latter case, good reversible subband transforms are essential. They normally involve well-designed filter banks characterized by longer analysis highpass filters with higher stopband attenuation.

Fortunately, in many cases color bands are strongly correlated. If this is the case and lossless image compression for progressive-resolution transmission is desired, one could consider the simple S-transform combined with adaptive spectral prediction. Then bit rates are possible which are about 20% lower than what has been reported in the literature so far [12].

Acknowledgments

This work was funded in part by the Alexandria digital library project under NSF grant IR194–11330, and in part by a University of California MICRO grant with matching support from the Xerox Corporation. The authors also thank Prof. Tor Ramstad for useful criticism.

References

[1] M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, SPIE, Bellingham, WA, 1991.

[2] A. N. Netravali and B. G. Haskell, *Digital Pictures*, Plenum Press, New York, NY, 1995.

[3] W. K. Pratt, "Spatial transform coding of color images," *IEEE Transactions on Communications Technology*, 19(6):980–992, 1971.

[4] H. Gharavi and A. Tabatabai, "Sub-band coding of monochrome and color images," *IEEE Transactions on Circuits and Systems*, 35:207 – 214, February 1988.

[5] T. H. Wendler and D. Meyer-Ebrecht, "Proposed standard for variable format picture processing and a codec approach to match diverse imaging devices," in *Proceedings of the SPIE*, volume 318, pages 298–305, 1982.

[6] A. Said and W. A. Pearlman, "An image multiresolution representation for lossless and lossy compression," *IEEE Transactions on Image Processing*, 5:1303–1310, 1996.

[7] O. Rioul, "A discrete-time multiresolution theory," *IEEE Transactions on Signal Processing*, 41(8):2591–2606, 1993.

[8] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic coding for data compression," *Commun. ACM*, 30:520–540, June 1987.

[9] N. Strobel, S. K. Mitra and B. S. Manjunath, "Lossless compression of color images for digital image libraries," in *Proceedings of the 13th International Conference on Digital Signal Processing*, volume 1, pages 435–438, Santorini, Greece, 1997.

[10] S. R. Tate, "Band ordering in lossless compression of multispectral images," *IEEE Transactions on Computers*, 46(4):477–83, 1997.

[11] T. A. Ramstad and S. O. Aase and J. H. Husøy, *Subband Compression of Images: Principles and Examples*, Elsevier, Amsterdam, 1995.

[12] M. J. Gormish, E. Schwartz, A. Keith, M. Boliek and A. Zandi, "Lossless and nearly lossless compression for high quality images," in *Proceedings of the SPIE*, volume 3025, San Jose, CA, February 1997.