Linear filtering of image subbands for low complexity postprocessing of decoded color images

Ulug Bayazit

Isik University, Buyukdere Caddesi, Maslak 34398, Istanbul, TURKEY

ABSTRACT

In [1], image adaptive linear minimum mean squared error (LMMSE) filtering was proposed as an enhancement layer color image coding technique that exploited the statistical dependencies among the luminance/chrominance or Karhunen Loeve Transform (KLT) coordinate planes of a lossy compressed color image to enhance the red, blue, green (RGB) color coordinate planes of that image. In the current work, we propose the independent design and application of LMMSE filters on the subbands of a color image as a low complexity solution. Towards this end, only the coordinates of the neighbors of the filtered subband coefficient, that are sufficiently correlated with the corresponding coordinate of the filtered subband coefficient, are included in the support of the filter for each subband. Additionally, each subband LMMSE filter is selectively applied only on the high variance regions of the subband. Simulation results show that, at the expense of an insignificant increase in the overhead rate for the transmission of the coefficients of the filters and with about the same enhancement gain advantage, subband LMMSE filtering offers a substantial complexity advantage over fullband LMMSE filtering.

Keywords: Postprocessing, linear filtering, color image coding, enhancement layer, subband, wavelet transform

1. INTRODUCTION

In order to achieve high compression performance in color image coding, statistical dependencies among the RGB color coordinate planes must be exploited. These dependencies are largely, but not completely in the form of linear dependencies. Subjecting the RGB coordinate planes to a signal independent luminance/chrominance or a signal dependent KLT linear coordinate transformation results in largely or completely decorrelated coordinate planes which can be efficiently compressed by bit allocation. Yet, as noted in [2], these decorrelated planes still possess statistical dependencies among them such as large transitions that occur at the same spatial locations.

In the literature, a family of algorithms relies on vector quantization of the three RGB color coordinates to exploit the statistical dependencies among them for the color palette design problem. In [3] and [4], the initial palette was designed by assigning more quantization levels to the regions of the color space with larger numbers of pixels. This palette was later improved by the LBG (Linde-Buzo-Gray, [5]) algorithm. The same palette design problem was approached in [6] by designing a tree-structured vector quantizer that split the clusters at each tree node along the principal vector direction.

Recently, vector quantization of color coordinates of wavelet coefficients has been proposed in [7] and [8] as an attractive way to compress color images. In [8], magnitude ordering information for vectors of color coordinates of wavelet coefficients is coded by linking these vectors by the Spatial Orientation Trees (SOT) of Set Partitioning in Hierarchical Trees (SPIHT, [9]) algorithm. Significant vectors are further refined by multistage or lattice vector quantization.

A simpler, yet, effective color image compression method based on SPIHT is proposed in [10] which decorrelates the color coordinate planes by a luminance/chrominance or KLT transformation, and independently codes each luminance/chrominance or KL plane by SPIHT. A similar idea is utilized in [11] to independently code the 4:2:0 YUV planes by a cousin of SPIHT that forms the basis of a low implementation complexity variant of JPEG2000. In [10] and [11], the rate scalable codecs efficiently allocate bits to transform coordinate planes, but the magnitude dependencies among the coordinate planes are not exploited.

The SOT of EZW (Embedded Zerotree Wavelet, [12]) is extended to link up the YUV planes in [13] to take advantage of the magnitude dependencies among the transform coordinate planes. While the original SOT of [12] describes spatial relations for each of the YUV planes, each chrominance node of the SOT of [13] is also a child node of the luminance node at the same location in the wavelet pyramid. Thus, since a luminance plane wavelet transform coefficient usually has a larger magnitude than a chrominance plane wavelet transform coefficient at the same spatial location, zerotrees (see [12] for a definition) in luminance components are likely to be accompanied by zerotrees in chrominance components and these zerotree symbols can be efficiently jointly coded or coded as a single zerotree symbol. Similarly, in [14], SOT's of SPIHT are extended to link up the LL bands of luminance and chrominance planes. Such an approach yields an advantage over the simple one of [10] at low bit rates where no symbols need to be explicitly transmitted for the chrominance zerotrees in most cases. In [15], the magnitude ordering of KLT coordinates is incorporated into SPIHT for a better coding performance than that in [14].

As an alternative to joint coding of the transform coordinate planes, [1] proposed to enhance the coordinates of a pixel by applying a LMMSE filter with support extending across all coordinate planes as well as over the spatial neighbors of the pixel. This LMMSE filtering based enhancement layer coding technique was demonstrated to successfully exploit the statistical dependencies among the transform coordinate planes and thereby efficiently reduce the quantization noise. In traditional image processing applications, the support of the filter ([16, pg. 67]) used to transform the decoded luminance/chrominance (or KLT) coordinate planes to get reconstructed RGB color coordinate planes only covers the filtered pixel, and does not cover the sites of its neighbors. The work in [1] showed that inclusion of the correlated neighbors of the filtered pixel in the filter's support is essential for the enhancement layer coding gain obtained by the LMMSE filter.

The major drawback of the design of an LMMSE filter with a large (spatial) support is the estimation of a large correlation matrix from the decoded image data samples with high complexity. The number of elements in the correlation matrix grows with the square of the size of the filter support. Secondly, since the original image data needed for the LMMSE filter design is available at the encoder, but not at the decoder, the method proposed in [1] requires the coding and transmission of each LMMSE filter coefficient to the decoder as overhead. Even though the marginal return (ratio of the decrease in distortion to the increase in rate) at a given rate due to the application of the method is significantly larger than that yielded by any color image coding method, highest marginal returns are obtained for small filter supports which cover only the most correlated neighbors.

Motivated by the theoretical basis in [17], we pursue independent application of an LMMSE filter to each subband of a color image in our current work. In [17], it was shown that independent (optimal) p'th order linear prediction from each subband of a signal yields a performance superior to p'th order (optimal) linear prediction from the fullband. On the other hand, since the design or application computational complexities are predominantly proportional to the number of data samples used for the design or application of a p'th order linear filter, (optimal) p'th order linear prediction from each subband has roughly the same aggragate design or application computational complexity as (optimal) p'th order linear prediction from the fullband. However, for our implementation of an enhancement layer coding technique, transmission of the coefficients of a p'th order linear filter for each subband requires too much overhead rate. A p'th order filter for each subband is deemed unnecessary since lower order filters for each subband still yield a performance advantage to p'th order fullband filtering (see "super-optimal" prediction error filter of [17]).

We judiciously select a filter support for each subband that is smaller in size than the filter support in the fullband and thereby aim to achieve low computational complexity at the expense of a graceful sacrifice in the enhancement layer coding gain. Since, as pointed out in [1], the correlations between the coordinates of the filtered pixel and the coordinates of the neighboring pixel is an important determinant of the LMMSE filtering gain, we include in the filter support of each subband only those decoded luminance/chrominance coordinates of neighboring subband coefficients that are sufficiently correlated with the corresponding decoded coordinates of the filtered subband coefficient.

Next section describes the design of the optimal linear filter from the subband coefficients of a color image. A mechanism for reducing complexity while gracefully lowering the enhancement gain by means of varying a correlation threshold is explained in Section 3. Prediction of the effect of quantization of the coefficients of the filters prior to their design is studied in Section 4 . The simulation results of Section 5 provide a ground for comparing the proposed

subband LMMSE filtering approach with the fullband LMMSE filtering approach of [1] on the complexity (run time) and performance (reconstruction PSNR) scales.

2. LMMSE FILTERING OF YUV COORDINATES (OF SUBBAND COEFFICIENTS) FOR ENHANCED RECONSTRUCTION OF RGB COORDINATES

Let the original RGB coordinate planes and the decoded YUV luminance and chrominance coordinate planes of an image subband be given. Without loss of generality to other coordinate systems such as YCrCb, YIQ or KLT, we shall be using the YUV coordinate system throughout the remainder of this paper. The original value of the RGB coordinate $\zeta \in \{R, G, B\}$, and the decoded value of the YUV coordinate $\gamma \in \{Y, U, V\}$ of the filtered subband coefficient at spatial location (i, j), are denoted by $x_\zeta(i, j)$ and $r_\gamma(i, j)$, respectively. The optimal linear filter g_ζ^*, used to enhance coordinate plane ζ, satisfies

$$g_\zeta^* = \arg\min_{g_\zeta} \sum_{i,j} \left\| x_\zeta(i, j) - g_\zeta^T u(i, j) \right\|^2 \tag{1}$$

where $u(i, j)$ is the K dimensional vector of 3 decoded YUV coordinates of the filtered coefficient and $K - 3$ decoded YUV coordinates of its neighbors. Specifically

$$u(i, j) = \left[u_Y^T(i, j)\, u_U^T(i, j)\, u_V^T(i, j) \right]^T$$

In the above, each K_γ dimensional subvector $u_\gamma^T(i, j)$, $\gamma \in \{Y, U, V\}$ is defined as an ordered sequence of $\{r_\gamma(m, n) : (m - i, n - j) \in \aleph_\gamma\}$, the decoded values of the YUV coordinate γ of the coefficient at (i, j) and $K_\gamma - 1$ neighboring coefficients in $\aleph_\gamma$, a spatial neighborhood around (i, j). We have $K = K_Y + K_U + K_V$.

For each RGB coordinate ζ, the minimization in Eqn. (1) is achieved by requiring the error $x_\zeta(i, j) - g_\zeta^T u(i, j)$ to be orthogonal to $u(i, j)$ in the statistical sense. This yields K regression equations in K unknowns expressed succinctly as

$$R_{uu}\, g_\zeta^* = r_{x_\zeta u}$$

where, for an image of height N pixels and width M pixels,

$$R_{uu} = \frac{1}{NM} \sum_{i,j} u(i, j) u^T(i, j)$$

is the $K \times K$ correlation matrix estimate and

$$r_{x_\zeta u} = \frac{1}{NM} \sum_{i,j} x_\zeta(i, j) u(i, j)$$

is the K dimensional cross-correlation vector estimate.

The optimal filter g_ζ^* can be determined by computing the matrix inverse R_{uu}^{-1} or by transforming the matrix $\left[R_{uu} \mid r_{x_\zeta u} \right]$ to $\left[I \mid g_\zeta^* \right]$ by Gaussian elimination. This operation takes place at the encoding end since $r_{x_\zeta u}$ can only be estimated at the encoding end from the original color image (subband) data.

As suggested in Figure 1, the computed coefficients of the subband LMMSE filter of plane ς (components of vector g_ς^*) are conveyed to the decoding end, where the filter is applied to get the enhanced reconstruction of the subband of plane ς. The design and application of the subband LMMSE filters are not coupled, and can be carried out fully in parallel. From Figure 1, we also note that the proposed subband LMMSE filtering approach does not require additional DWT or IDWT filtering operations, if it is to be applied at the rear end of a wavelet based image codec.

The above procedure was originally used in [1] to design an LLMSE filter on the fullband image where $u(i, j)$ was the K dimensional vector of decoded YUV coordinates of pixels in a square region centered at the filtered coefficient. In designing and applying LMMSE filters on image subbands, we employ only a subset of the decoded YUV coordinates of the neighboring coefficients in the square region. We discuss this next.

3. DETERMINATION OF SUBBAND FILTER SUPPORTS

Usually, edges that are not parallel to the direction of application of the high-pass filter used to obtain a subband can be captured in that subband. Therefore, one would expect very little correlation to exist between the decoded YUV coordinate of the filtered subband coefficient and the corresponding decoded YUV coordinate of the neighboring subband coefficient along the direction of application of the high-pass filter. Hence, the inclusion of the decoded YUV coordinates of such neighboring coefficients can be expected to contribute very little to the enhancement gain of the LMMSE filter. However, decoded YUV coordinates of neighboring coefficients along other directions may also be excluded from the filter support if edges along those directions do not exist in the image (or are not captured in the image subband). We employ an adaptive filter support determination strategy by including the decoded coordinate γ_0 of the neighboring coefficient at location $(i-i_0, j-j_0)$ in a square region centered at the subband coefficient at location (i, j) if its covariance

$$\rho_\zeta(i_0, j_0, \gamma_0) = \frac{1}{NM}\sum_{i,j}(x_\zeta(i,j) - \bar{x}_\zeta)(r_{\gamma_0}(i-i_0, j-j_0) - \bar{r}_{\gamma_0})$$

with the corresponding coordinate ζ of the coefficient exceeds a threshold. Increasing this threshold serves to trade-off linear filtering gain for lower computational complexity.

In Figures 2 and 3, we display two instances of 3x3 spatial masks determined for the LMMSE filters of each of the 7 subbands of a two-level decomposition of color Lenna image for two different threshold values. Since the correlations between the YUV coordinates of a filtered coefficient and either of its two neighbors on its opposite sides are equal, neighboring coefficient coordinates are included in pairs in the filter supports. Typically, all coordinates of all neighbors in the lowest frequency subbands are retained in the filter supports even at very high thresholds. The decoded YUV coordinates of the filtered coefficient in each subband are not subjected to correlation thresholding and are unconditionally retained in the filter support.

High frequency subbands typically exhibit vast regions of low data variance which correspond to the smooth regions of the image, and small regions of high data variance which corresponds to the edges and textured regions in the image. As noted in [18], the LMMSE filter yields very little enhancement gain in the low variance regions of a subband. Therefore, to save modest complexity at the expense of very little sacrifice from enhancement gain, the LMMSE filter is designed only on data collected from the high variance regions and is applied to the same data in return.

4. EFFECT OF FILTER COEFFICIENT QUANTIZATION ON THE ENHANCEMENT GAIN

Once the coefficients of filters are determined by solving sets of regression equations, these coefficients need to be conveyed to the decoder as side information. It is of interest to know at what precision these coefficients need to be represented. In this section, we analyse the sacrifice from the enhancement gain due to the quantization of the coefficients of the LMMSE filters. For simplicity of analysis, we would like to relate the mean squared distortion of the enhanced subband to the mean squared distortion of the filter coefficients.

Let the vector of quantized coefficients of the filter for enhancing coordinate ζ be denoted as $\hat{g}_\zeta$ so that

$$\hat{g}_\zeta = g_\zeta + e_\zeta$$

where e_ζ is the vector of quantization error values in the coefficients of the filter. The enhanced coordinate ζ of pixel at location (i, j) is given by

$$\tilde{x}_\zeta(i, j) = g_\zeta^T u(i, j)$$

when the filter coefficients are at full (32 bit floating) precision. The enhanced coordinate ζ of pixel at location (i, j) is given by

$$\tilde{\tilde{x}}_\zeta(i,j) = \hat{g}_\zeta^T u(i,j)$$

when the filter coefficients are quantized for a low precision representation. By linearity of the filters, the difference may be written as

$$\tilde{\tilde{x}}_\zeta(i,j) - \tilde{x}_\zeta(i,j) = e_\zeta^T u(i,j)$$

By taking squares and expectation of both sides

$$E\left[(\tilde{\tilde{x}}_\zeta(i,j) - \tilde{x}_\zeta(i,j))^2\right] = E\left[u^T(i,j) e_\zeta e_\zeta^T u(i,j)\right]$$

$$= E_u\left[u^T(i,j) E_e\left[e_\zeta e_\zeta^T \mid u(i,j)\right] u(i,j)\right]$$

where the inner expectation is over the realizations of the random vector e_ζ given $u(i,j)$, the vector of decoded values of luminance and chrominance coordinates of the pixels in the filter's support region, and the outer expectation is over $u(i,j)$. We can assume that e_ζ is independent of $u(i,j)$ and is zero mean. If we further assume that the quantization error in each filter coefficient is independent of the quantization errors in others and let the quantization errors for all coefficients have the same variance, we get

$$E_e\left[e_\zeta e_\zeta^T \mid u(i,j)\right] = E_e\left[e_\zeta e_\zeta^T\right] = \sigma_{e_\zeta}^2 I$$

Hence

$$E\left[(\tilde{\tilde{x}}_\zeta(i,j) - \tilde{x}_\zeta(i,j))^2\right] = \sigma_{e_\zeta}^2 E\left[u^T(i,j) u(i,j)\right]$$

In the above expression, the quantity, $E\left[u^T(i,j) u(i,j)\right]$ is the expected energy of the coefficients of the subband in the support region of the filter for the enhancement of coordinate plane ζ. If the filter coefficients are uniformly quantized with a quantization bin interval of length Δ, then the loss in enhancement gain can be approximately expressed as

$$E\left[(\tilde{\tilde{x}}_\zeta(i,j) - \tilde{x}_\zeta(i,j))^2\right] \approx \frac{\Delta^2}{12} E\left[u^T(i,j) u(i,j)\right] \quad (2)$$

Eqn. (2) can be used to obtain an estimate of the loss in enhancement gain due to fixed precision representation of the filter coefficients. This might be useful since computing the second order moment of Eqn. (2) has considerably less complexity than linear filtering the frames at full-precision to get an exact figure for what the enhancement gain would have been without any loss. Coefficients of different coordinate planes and/or subbands might need to be represented at different precisions if the goal is to jointly minimize the overhead rate expenditure and the loss in enhancement gain. However, in general, such a joint optimization and variable precision representation of the coefficients yields very little rate-distortion advantage when compared to fixed precision representation, since the overhead rates of interest are negligible. Therefore, joint minimization of overhead rate expenditure and the loss in enhancement gain by variable precision representation of filter coefficients is not pursued in this work.

5. EXPERIMENTAL RESULTS

For the simulations, the YUV planes of the 512x512 color image Lenna, and 1600x960 color image MIT000 (first frame of MIT color sequence) have been coded by the SPIHT algorithm of [9] as mentioned in [1]. The decoded YUV coordinate planes of these images are transformed by means of the simple standard YUV-RGB transformation filter of [16, pg. 67] to get the reconstructed RGB planes, and the fullband LMMSE filter of [1] and the subband LMMSE filter proposed in this paper to get the enhanced RGB coordinate planes. The subband LMMSE filters are designed and applied only on the high variance regions of the subbands of a two-level wavelet decomposition of the decoded YUV planes. A subband coefficient is deemed to belong to a high variance region if the energy of the coordinates in a 3x3 square region centered at the coefficient exceeds the experimentally determined threshold value of 270.

The results reported in Table 1 were obtained on Lenna coded at a rate of 0.2345bpp. When compared to the standard filtering, the fullband LMMSE filtering yields 0.1-0.4 dB improvements in the RGB planes at the expense of 0.0034bpp overhead rate for the transmission of the coefficients. The 7 subband LMMSE filtering with the correlation

threshold of Section 3 set at 0.2 yields gains in a similar range at the expense of 0.0073bpp overhead rate. The aggregate marginal returns (the ratio of the total distortion reduction to the total overhead rate) for fullband filtering and subband filtering are 3200. and 2054., respectively. The marginal return of the SPIHT codec is approximately only 366. at this rate.

In Table 1, the reported design and application run times represent the average time required to determine the LMMSE filters for all subbands of RGB planes, and the average time required to filter all coded subbands of RGB planes using these filters, respectively. These averages were obtained over 3 runs on a P4 3.0Ghz processor system with 1Gb RAM and exclude the time for DWT and IDWT that are part of the SPIHT codec. The source code for subband LMMSE filtering makes multiple calls to the same function called once for each ζ by the source code for fullband LMMSE filtering. Therefore, it is legitimate to compare complexities based on run times. The number of elements in the correlation matrix being proportional to the square of the size of the filter support gains the subband LMMSE filtering approach a distinct advantage in design complexity due to few coordinates of neighboring coefficients exceeding the correlation threshold in high frequency subbands.

Table 2 shows similar results for MIT000. Here, the aggregate marginal return is significantly larger for LMMSE filtering approaches, since the number of image pixels is considerably larger whereas the number of overhead bits is approximately the same.

In Figures 4 and 5, we display the enhancement gain and the number of overhead bits for the two images as we vary the correlation threshold from ∞ to 0. In general, high aggregate marginal returns are observed when coefficients with high correlations are retained in the filter support and the aggregate marginal returns gracefully decline with the overhead rate. The highest aggregate marginal returns are not observed for the lowest overhead rates since the fixed rate transmission of a mask, that indicates which coordinates are retained in the supports of the filters, consumes a large part of the low overhead rates. Nevertheless, the curves staying well below the dotted lines drawn as references shows the advantage of selecting coefficients based on correlations.

In Table 3, the mean squared error values of reconstructed RGB planes, after enhancement with subband LMMSE filtering, are tabulated for the two cases in which the filter coefficients are lossy compressed at the precision reported, and not compressed (represented at full 32 bits precision). The small difference is well predicted by Eqn. (2), but is a little larger probably because the scalar quantizer used to compress the coefficients has a dead zone of size 2Δ whereas the other bins are of size Δ.

In Figure 6, the luminance planes of the middle part of the original color image Lenna, its reconstruction after coding and decoding at 0.2345bpp, and its enhanced reconstruction with the subband filtering method are displayed for visual comparison. The subband filters attenuate ringing and serrations due to the loss or incorrect reproduction of the high frequency content with the SPIHT based codec.

6. REFERENCES

[1] U. Bayazit, “Postprocessing Of Decoded Color Images By Adaptive Linear Filtering,” Signal Processing: Image Comm. **18**(2), 91-101 (Feb. 2003).

[2] N. Netravali, C.B. Rubinstein, "Luminance adaptive coding of chrominance signals," IEEE Trans. Comm., **27**(4), 703-710 (April 1979).

[3] G. Braudaway, "A procedure for optimum choice of a small number of colors from a large color palette for color imaging," *Electronic Imaging '87*, San Francisco, CA (1987).

[4] P. Heckbert, "Color image quantization for frame buffer display," Computer Graphics, **16**(3), 297-307 (July 1982).

[5] Y. Linde, A. Buzo and R. Gray, “An Algorithm for Vector Quantizer Design,” IEEE Trans. Comm., **28**(1), 84-95 (Jan. 1980).

[6] M. T. Orchard, C. A. Bouman, "Color Quantization of Images," IEEE Trans. Signal Proc., **39**(12), 2677-2690 (Dec. 1991).

[7] S. Mitra et al., "Efficient color image compression using integrated fuzzy neural networks for vector quantization," *IEEE Int’l Conf. on Syst., Man, and Cybernetics, Comp. Cybernetics and Simulation*, 1, 184 – 188 (Oct. 1997).

[8] D. Mukherjee, S. K. Mitra, "Arithmetic coded vector SPIHT with classified tree-multistage VQ for color image coding," , *IEEE 2. Workshop on Multimedia Signal Proc*., 444 – 449 (1998).
[9] A. Said, W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," IEEE Trans. Circuits and Syst. for Video Tech., **6**(3), 243-250 (June 1996).
[10] SPIHT FAQ.what method is used for color compression? http://www.cipr.rpi.edu/research/SPIHT/spiht6.html.
[11] W. A. Pearlman et al., "Efficient, Low-Complexity Image Coding with a Set-Partitioning Embedded Block Coder," IEEE Trans. Circuits and Syst. for Video Tech., **14**(11), 1219-1235 (Nov. 2004).
[12] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," IEEE Trans. Signal Proc., **41,** 3445-3462 (Dec. 1993).
[13] K. Shen, E. J. Delp, "Wavelet based rate scalable video compression," IEEE Trans. Circuits and Syst. for Video Technology, **9**(1), 109-122 (Feb. 1999).
[14] A. A. Kassim, W. S. Lee, "Embedded Color Image Coding Using SPIHT with Partially Linked Spatial Orientation Trees," IEEE Trans. Circuits Syst. for Video Tech., **13**(2), 203-206, (Feb 2003).
[15] C.-L. Yang et al., "A Novel Ordered-SPIHT for Embedded Color Image Coding," IEEE *Int'l Conf. Neural Net. and Signal Proc*., Nanjing, China, (Dec. 2003).
[16] W. K. Pratt, *Digital Image Processing*, 66-67, Wiley-Interscience, New York, (1991).
[17] S. Rao, W. A. Pearlman, "Analysis of Linear Prediction, Coding and Spectral Estimation from Subbands," IEEE Trans. on Info. Theory, **42,** 1160-1178 (July 1996).
[18] A. Nosratinia, "Optimal subband synthesis," *IEEE Int'l Conf. on Image Proc*., 641-644, Santa Barbara, CA (1997).

		Decoded image	Subband LMMSE filtering	Fullband LMMSE filtering
Total Rate (bpp)		0.2345	0.2418	0.2379
PSNR (dB)	Blue	29.31	29.88	29.68
	Green	30.72	30.94	30.84
	Red	32.07	32.39	32.41
Run time (msec.)	Design		1240	12375
	Application		406	609

Table 1: Enhancement results - Lenna(512x512).

		Decoded image	Subband LMMSE filtering	Fullband LMMSE filtering
Total Rate (bpp)		0.7918	0.7931	0.7923
PSNR (dB)	Blue	31.90	32.29	32.23
	Green	34.32	34.61	34.40
	Red	33.26	33.46	33.32
Run time (msec.)	Design		10917	73067
	Application		2771	3583

Table 2: Enhancement results - MIT000(1600x960).

		Lossy comp. ($\Delta = 0.0005$)	32 bit full precision	Predicted mse increase (Eqn. 2)
MSE	Blue	38.6928	38.6838	0.006778
	Green	22.7419	22.7357	0.006814
	Red	29.5756	29.5562	0.006804

Table 3: Comparison of lossy and lossless compression of filter coefficients for enhancement of MIT000 (1600x960) where correlation threshold was set at 0.2.

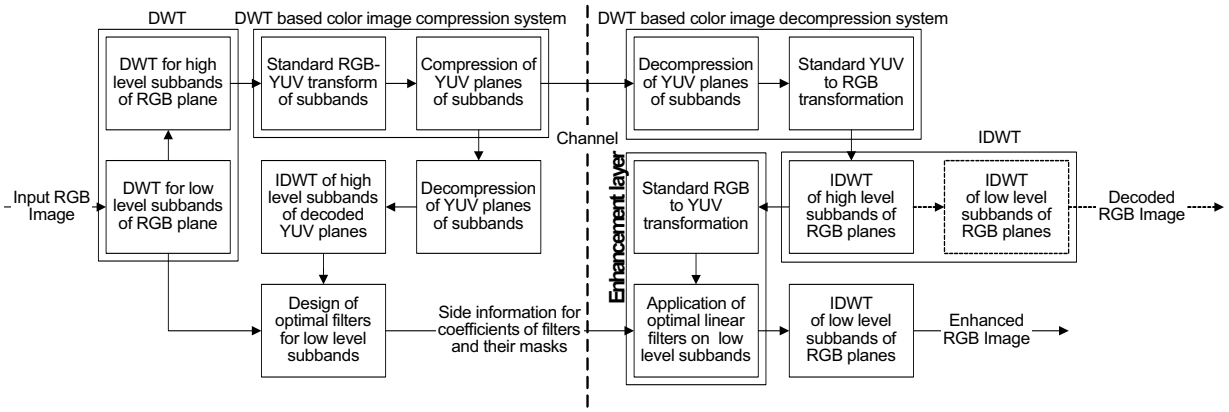

Figure 1: The low level subband filters are designed at the encoder (left) and transmitted to the decoder (right) for enhanced reconstruction.

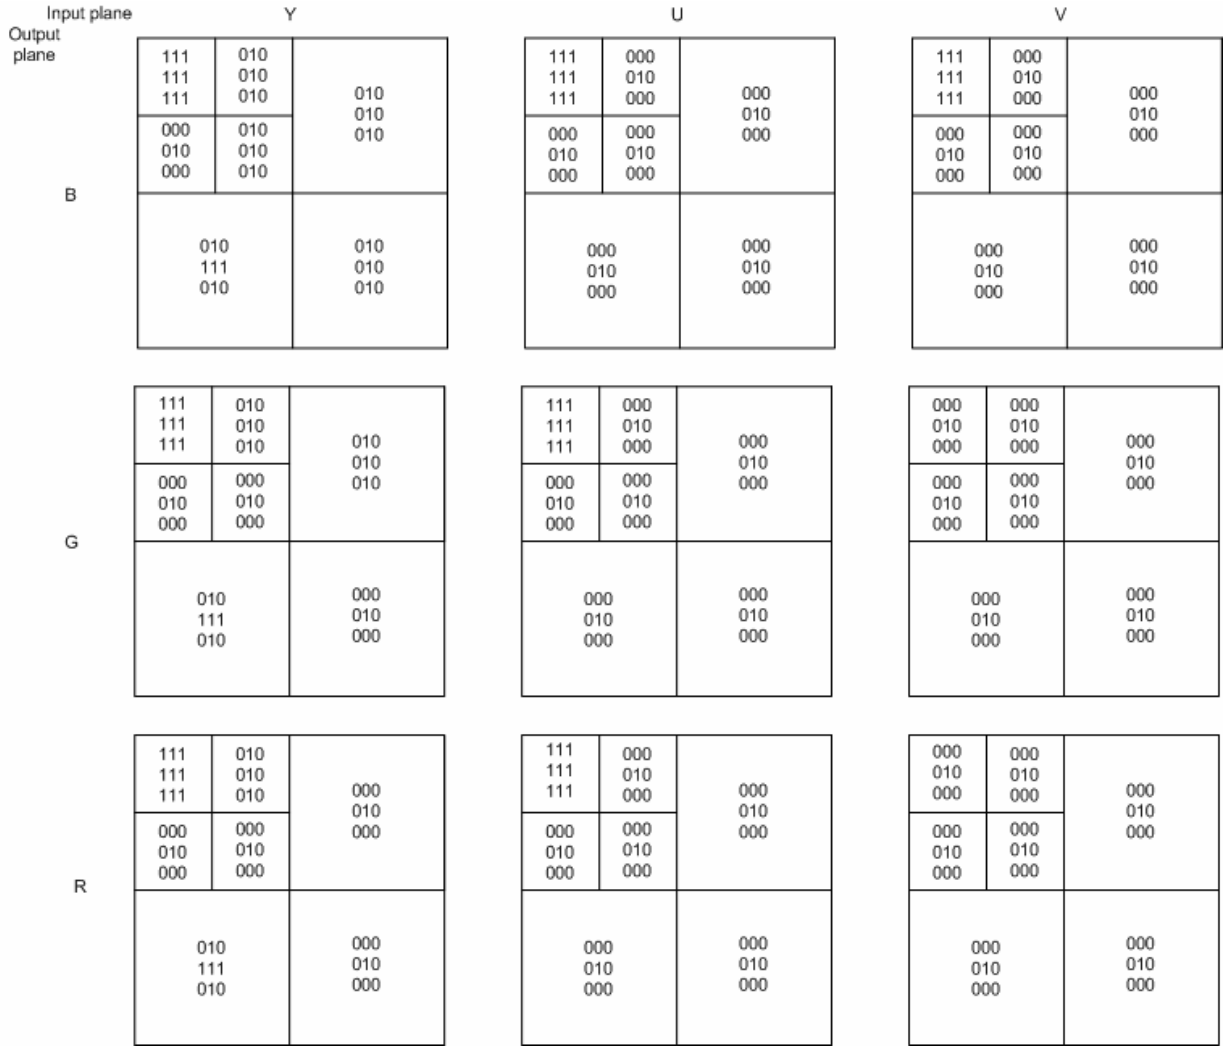

Figure 2: Spatial masks that indicate which coordinates (YUV) of which neighbors are included in the spatial support of each subband filter of each output color component (RGB) plane. (Lenna, correlation threshold=0.3)

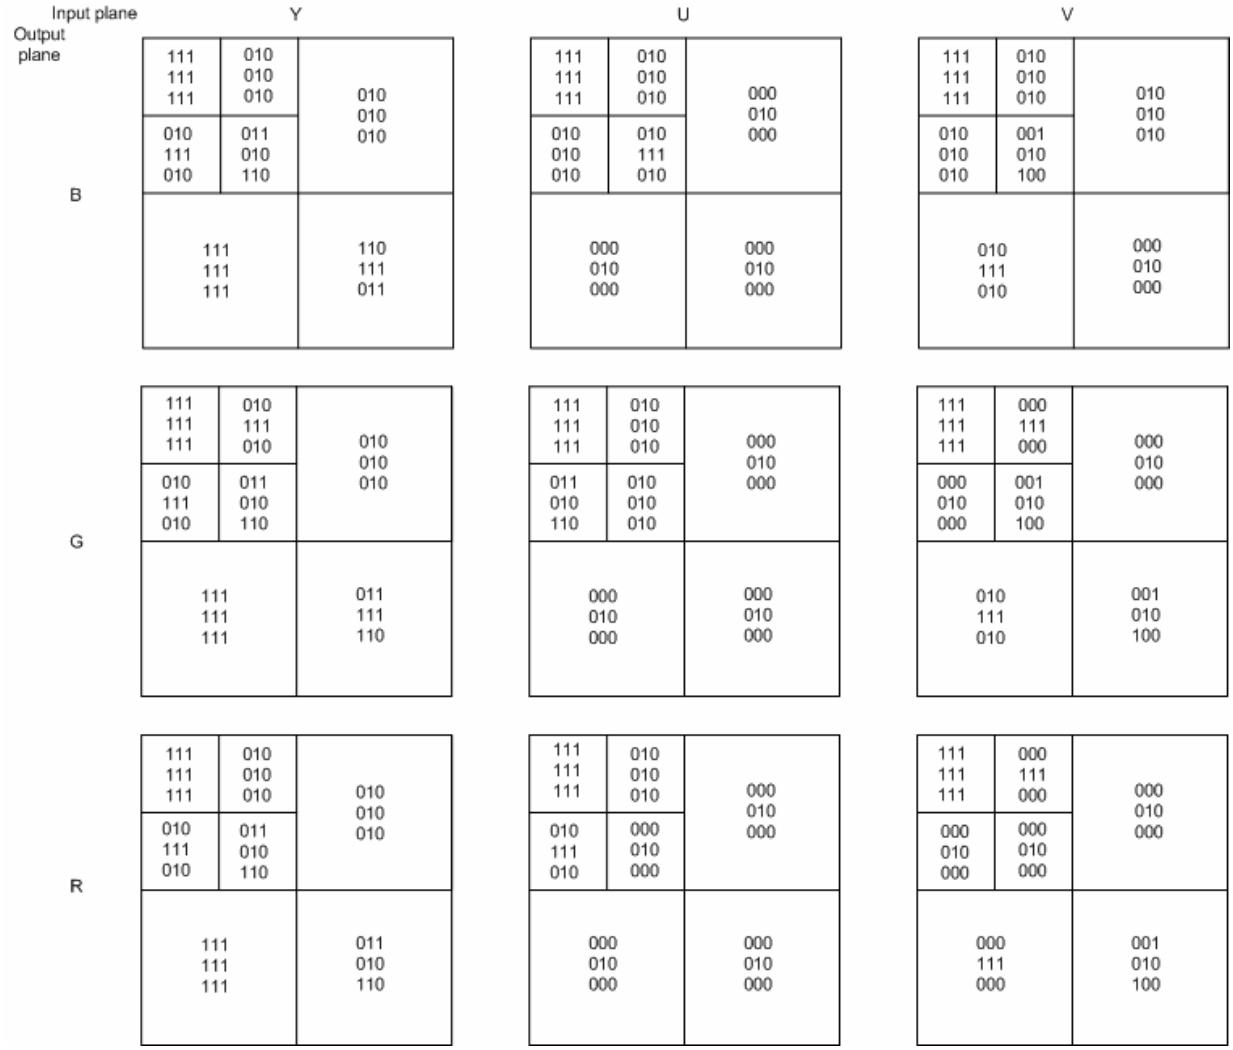

Figure 3: Spatial masks that indicate which coordinates (YUV) of which neighbors are included in the spatial support of each subband filter of each output color component (RGB) plane. (Lenna, correlation threshold=0.1)

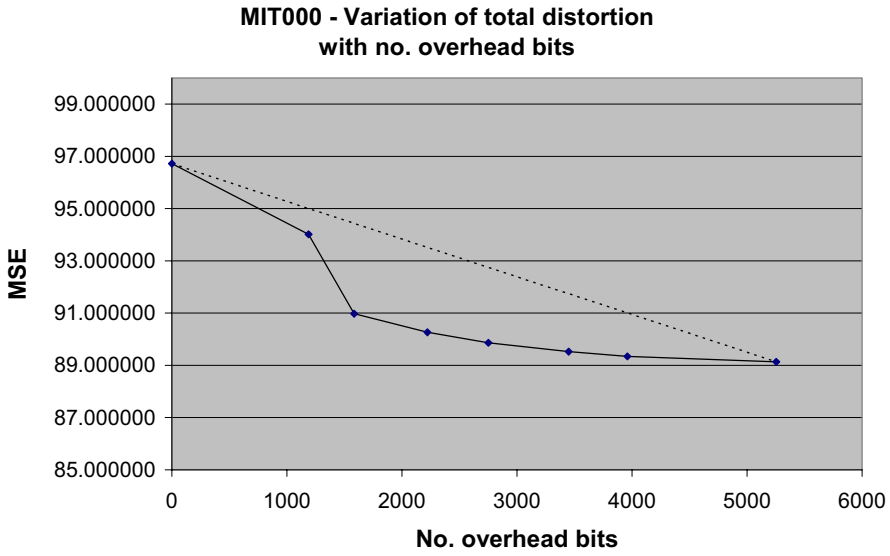

Figure 4: Enhanced reconstruction distortion vs. overhead rate parameterized by correlation threshold (low overhead rates correspond to high thresholds)

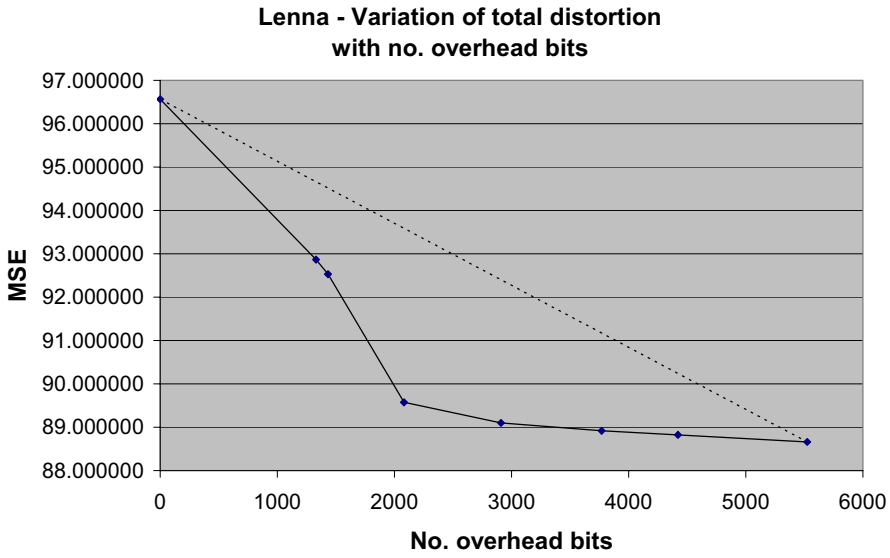

Figure 5: Enhanced reconstruction distortion vs. overhead rate parameterized by correlation threshold (low overhead rates correspond to high thresholds)

Figure 6: When compared to the original Lenna image shown at the top, the artifacts are more pronounced in the circled regions of the decoded Lenna image (total coding rate 0.2345bpp) shown on the bottom left than the corresponding regions of the Lenna image enhanced by subband linear filtering (total coding rate 0.2418bpp) shown on the bottom right. (Only the luminance component is shown).