

Protection of Regions of Interest Against Data Loss in a Generalized Multiple Description Framework

Agnieszka C. Miguel

Eve A. Riskin

`{amiguel,riskin}@ee.washington.edu`

*Dept of EE, University of Washington
Seattle WA, 98195-2500*



UWEE Technical Report
Number UWEETR-2002-0007
January 2002

Department of Electrical Engineering
University of Washington
Box 352500
Seattle, Washington 98195-2500
PHN: (206) 543-2150
FAX: (206) 543-3842
URL: <http://www.ee.washington.edu>

Protection of Regions of Interest Against Data Loss in a Generalized Multiple Description Framework

Agnieszka C. Miguel
Eve A. Riskin

`{amiguel,riskin}@ee.washington.edu`

Dept of EE, University of Washington
Seattle WA, 98195-2500

University of Washington, Dept. of EE, UWEETR-2002-0007

January 2002

Abstract

We present a simple and efficient scheme for protecting a region of interest (ROI) in an image sent across lossy communication networks. To increase the probability that the ROI is received with high quality, we extend the unequal loss protection framework of MD-SPIHT [1] by adding more redundancy to the ROI than to other parts of the image. The ROI is therefore heavily protected at the expense of lower protection in the background. In addition, the ROI is coded to a higher bit rate than the rest of the image. This method can be used for packet erasure networks or in the generalized multiple description coding framework.

1 Introduction

Regions of interest (ROIs) are spatial areas in an image that are more important to the end user than other areas. For example, small regions of medical images may be sufficient for diagnostic purposes. Targets can be identified and tracked using only small sections of surveillance images. In database browsing, an ROI can be used to recognize and screen out an image.

Due to increased demand for ROI coding, the emerging image and video compression standards, such as JPEG-2000, MPEG-4, and MPEG-7 require object-oriented coding and motion estimation. Most ROI research concentrates on designing methods for coding the ROI separately from the rest of the image [2]; at a higher spatial resolution such as in foveated image compression [3]; or at a higher bit rate [4, 5, 6, 7, 8, 9] than the rest of the image. With the recent popularity of zerotree based coding, some of these methods have been extended to the Embedded Zerotree Wavelet (EZW) [10] and the Set Partitioning in




Figure 1: Example of tree selection and their coding rates in one description.

Hierarchical Trees (SPIHT) [11] compression algorithms. However, none of the current ROI coding methods addresses protection of the ROI in the presence of data loss. Rogers and Cosman [5] implement two ROI coding schemes in their Packetized Zerotree Wavelet (PZW) coding method, but they do not provide more loss protection for the ROI than for any other part of the image. In this paper, we propose a scheme for sending compressed images with regions of interest across lossy communication networks such that if data loss occurs, the ROI is received with high quality.

Most of the networks used today are based on an exchange of packets of data. When congestion occurs because of data moving from a higher to lower capacity link, the network is often forced to erase some of the packets. If retransmitting packets is not desirable, the decoder must recover meaningful information from the received packets alone. Such packet erasure networks are related to generalized multiple description (MD) coding schemes where the encoder produces multiple descriptions of the source data, each exactly filling one packet [12]. Recently, the generalized MD framework has received considerable attention and many new algorithms have been proposed [12, 13, 14, 15, 1].

We concentrate on networks that do not provide preferential treatment to packets and instead erase them randomly. Usually, however, data that we transmit vary in importance. If the network is unable to transmit all of the data, then we would like the discarded packets to affect the least important data more than the most important data. In [1] we presented an algorithm for using SPIHT in a generalized multiple description framework (MD-SPIHT). To combat description loss, redundancy is added to the original data during the compression process. Unequal loss protection is implemented by varying the amount of redundancy with the importance of data.

In this work, we extend MD-SPIHT to the protection of ROIs. If severe data loss occurs, a high quality ROI can still be reconstructed. The next section describes the MD-SPIHT algorithm originally proposed in [1]. Its extension to ROI coding and protection is introduced in Section 3. Section 4 presents results. Finally, Section 5 contains future research directions and conclusions.

2 Multiple Description SPIHT

Our MD-SPIHT algorithm [1] adds controlled amounts of redundancy to the original data during the compression process. We use SPIHT because it allows for easy determination of the data importance for the overall quality of an image. By progressively ordering the data, SPIHT sends the globally most relevant information first. To protect data that are more important to the overall image quality more heavily than less important data, MD-SPIHT adds more redundancy to the earlier parts of the bit stream than to the later parts.




Figure 2: Coding example: four original trees (1-4) are sent in four descriptions (D1-D4) with two sets of redundant trees. If two descriptions are lost (D2 and D3), the second and third trees are partially recovered from their copies in the received descriptions.

To increase the resilience of SPIHT to data loss, we first modify it in a manner similar to PZW [5]. Spatially dispersed wavelet coefficient trees are grouped together, coded with the SPIHT algorithm, and transmitted as one description. We use arithmetic coding within each description and code each description to the same bit rate.

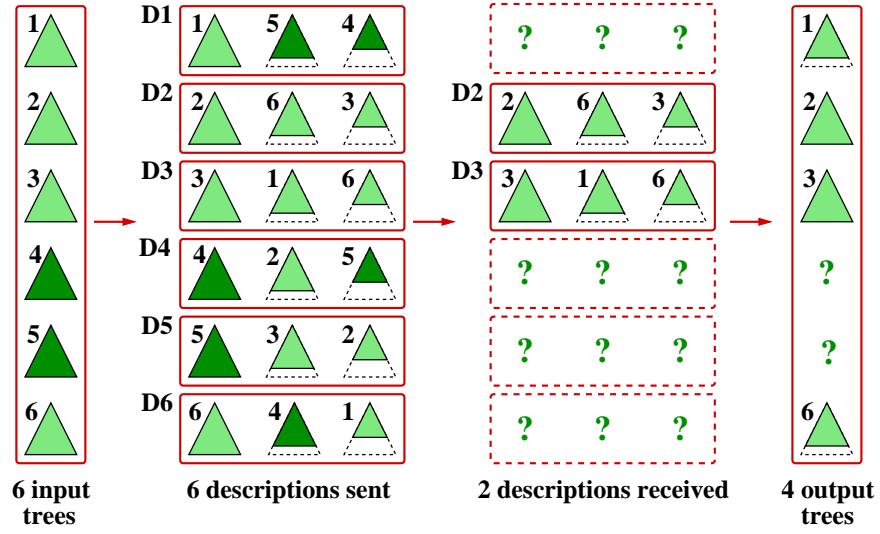
We add redundancy to the image data before it is compressed. This redundancy simply consists of repeated wavelet coefficient trees. Each description has one set of fully coded original trees and multiple sets of partially coded redundant trees. Trees in any one description correspond to different spatial locations in the image. Each set of redundant trees is coded to a successively lower bit rate (see Fig. 1). Because coding a tree to a lower bit rate corresponds to sending only the most important data, the more important parts of each tree are sent in more descriptions than the less important parts. In this way, we vary the amount of redundancy according to the importance of data to the overall image quality.

Information about one tree is located in many different descriptions. As shown in the coding example in Fig. 2, if some descriptions are erased, the most important parts of the original trees in those descriptions will be recovered because their copies at lower bit rates will be present in the other descriptions that are received.

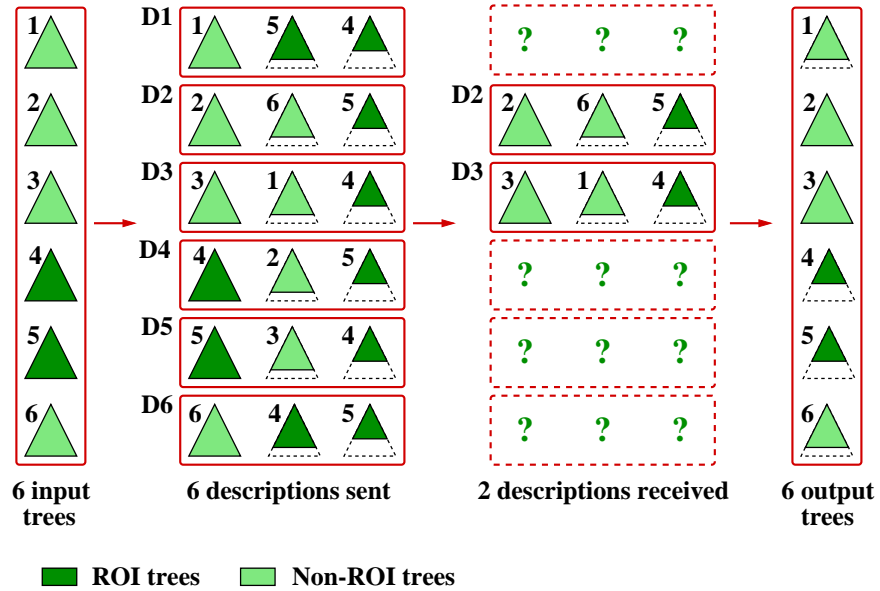
To optimize the amounts of redundancy assigned to each wavelet coefficient tree, we use an allocation algorithm based on rate-distortion trade-offs. It minimizes the expected distortion of the received data subject to a description loss model. We assume that the network behavior can be modeled by an estimator that outputs a probability mass function indicating the likelihood that a particular number of descriptions is lost. We use the extended BFOS algorithm [16], which allocates bits among various tree copies such that all allocations lie on the lower convex hull of the rate-distortion curve. See [1] for more details.

3 Multiple Description SPIHT with ROI Coding

In this work we extend MD-SPIHT to code images with regions of interest. Our goal is to protect ROIs more than other parts of the image so that if data loss occurs, a high quality



(a)



(b)

Figure 3: Coding Example: Six original trees (1-6) are sent in six descriptions (D1-D6). Two of the original trees (4 and 5) contain ROI coefficients. (a) All trees are allocated two levels of redundancy. If only descriptions D2 and D3 are received, the ROI trees can not be recovered. (b) The second level of redundancy is allocated exclusively to ROI trees. If only descriptions D2 and D3 are received, both ROI trees are partially recovered from their redundant copies in the received descriptions.

ROI can still be reconstructed. For this reason, we increase the amount of redundancy allocated to the ROI. Redundant trees corresponding to the ROI are copied more often and are assigned higher bit rates than the non-ROI redundant trees. As a result, the ROI trees have higher probability of being received because they are present in many more descriptions than the non-ROI trees.

A simple coding example illustrating our approach is shown in Fig. 3. A total of 6 trees are sent in 6 descriptions. Two of these trees (4 and 5) belong to an ROI and are shown in black. The non-ROI trees are shown in light gray. In the first case (Fig. 3(a)), two levels of redundancy are allocated to all trees as is done in the original MD-SPIHT algorithm. As a result, when 4 descriptions are lost, the ROI trees may be unrecoverable. In the next case (Fig. 3(b)), corresponding to the proposed algorithm, the second level of redundancy is allocated only to the ROI trees. If 4 descriptions are again lost, the two ROI trees are recovered.

Similar to other ROI coding methods, we code the ROI to a higher bit rate than the rest of the image. Using the schemes proposed by Shapiro [6], Rogers and Cosman [5], and Atsumi and Farvardin [4], the wavelet coefficients corresponding to the ROI are scaled by a large factor. As a result, these coefficients are coded much earlier than the rest of the image in the SPIHT algorithm. This method requires sending the ROI position and the scaling factor as a side information. The value of the scaling factor corresponds to the level of ROI prioritization over the rest of the image.

The original SPIHT algorithm orders the data progressively and sends the *globally important* information first. By adding more redundancy to the earlier parts of the bit stream than to the later parts, MD-SPIHT in [1] protects data that are globally important to the image quality more than less important data. In this work, scaling the ROI causes it to be coded in the earlier parts of the bit stream. This way, *locally important* information is sent first. Again MD-SPIHT adds more redundancy to the earlier parts of the bit stream but this time it results in assigning more redundancy to the localized ROI than to the background. Therefore, the ROI has higher probability of being received intact and its quality will be higher than the quality of the background, which is consistent with the existing ROI coding methods.

4 Results

As in [4], we demonstrate our results on a 512×512 gray scale image of a little girl where the girl's face is an 80×80 region of interest. We use a 5-level DWT decomposition to obtain a total of 64 wavelet coefficient trees. As seen in Fig. 4, there are 9 wavelet coefficient trees that correspond to the ROI. We set the number of descriptions to 8 and the total bit rate to 1.0 bpp.

First, we investigate the trade-off between the quality of the ROI and the background for different ROI scaling factors. As Fig. 5 shows, for every 5 dB increase in ROI quality, the overall image PSNR drops only by about 0.5 dB. Therefore, for this image, we can achieve a significant improvement in the ROI quality with only a slight degradation of the background. For the presentation of our results, we choose ROI scaling factors of 16 and 1, corresponding to medium and no prioritization of the ROI.

Following [1], we add 35% redundancy for the case of the ROI scaled by 16. Fig. 6




Figure 4: The original little girl image with the region of interest and location of the 64 wavelet coefficient trees.

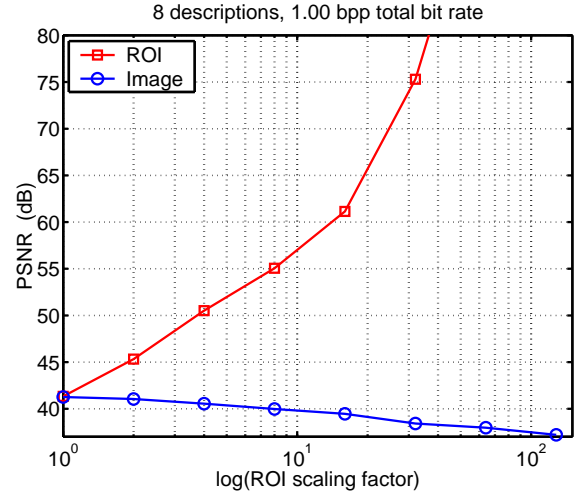


Figure 5: PSNR results for the ROI and the whole image for different ROI scaling factors and the image in Fig. 4.

shows the quality of the little girl image as 1, 2, 3, and 4 out of 8 descriptions are received. Note that the ROI is the first section of the image to appear.

A further justification of our method can be seen in Fig. 7. We show 4 cases in which 5 out of 8 descriptions are received. For no ROI scaling and 0% redundancy, the ROI is still not fully received (Fig. 7(a)). Scaling the ROI without adding any redundancy causes the received parts of the ROI to be of higher quality than in the previous case, but its missing sections render the ROI almost unusable (Fig. 7(b)). With no ROI scaling, the 35% redundancy is spread to all parts of the image. The ROI is visible but its quality is low (Fig. 7(c)). Finally, the ROI is scaled and redundancy is added. Because of the scaling, the redundancy is highly concentrated in the ROI area. Therefore, the ROI is fully received and of high quality (Fig. 7(d)).

Fig. 8 plots the PSNR in the ROI for the four scenarios described here. If all 8 descriptions are received, the ROI will be best for the case of ROI scaling and 0% redundancy, as expected. However, if any descriptions are lost, the ROI will have the highest quality when it is scaled and redundancy is added.

As can be seen in Fig. 9, most of the 35% redundancy allocated to the image when the ROI is scaled by 16 is concentrated in the ROI area. The concentration of redundancy in the ROI ensures that it will be received even when high loss occurs. When the ROI is not scaled, the redundancy is spread evenly throughout the image. Note that even within the ROI, the redundancy is allocated unequally; the upper three trees, which contain the lowest number of ROI wavelet coefficients, are assigned less redundancy than the other six ROI trees.



(a) Image: 20.93 dB, ROI: 28.25 dB

(b) Image: 24.38 dB, ROI: 34.79 dB



(c) Image: 26.36 dB, ROI: 38.54 dB

(d) Image: 27.57 dB, ROI: 40.47 dB

Figure 6: Image and ROI quality at 35% redundancy, 1.0 bpp total bit rate, and ROI scaling factor of 16. Images show 1, 2, 3, and 4 descriptions received out of 8.

5 Conclusions and Future Work

We have shown a simple and effective algorithm for protecting regions of interest from data loss. Varied amounts of redundancy are added to packetized SPIHT in the generalized multiple description framework. To increase the probability that the ROI is received, most of the redundancy is concentrated in the region of interest. The ROI is therefore heavily protected at the expense of lower protection in other areas of the image. The results presented show that high ROI quality is achieved even when only a few descriptions are received.

Future work will focus on incorporating the results of Goyal *et al.* [17] on quantization effects in overcomplete frame expansions. When multiple wavelet coefficient trees from the same spatial location are received, MD-SPIHT uses only the tree with the highest bit rate for image reconstruction. When using an appropriate frame expansion and consistent reconstruction techniques [17, 14], the additional trees will contribute useful information



(a) Image: 20.89 dB, ROI: 23.49 dB

(b) Image: 20.90 dB, ROI: 21.27 dB



(c) Image: 34.14 dB, ROI: 35.06 dB

(d) Image: 30.18 dB, ROI: 42.67 dB

Figure 7: Image and ROI quality for 5 out of 8 descriptions received. The total bit rate is 1.0 bpp. (a) ROI scaling factor: 1, redundancy: 0%, (b) ROI scaling factor: 16, redundancy: 0%, (c) ROI scaling factor: 1, redundancy: 35%, (d) ROI scaling factor: 16, redundancy: 35%.

and improve the PSNR when many descriptions are received.

The ROI protection scheme proposed here is not limited to a single ROI per image. In fact, multiple ROIs can be created in one image and assigned different levels of importance. When redundancy is added, it will be concentrated in these regions in an unequal manner.

Finally, the ROI protection algorithm shown here can be extended beyond MD-SPIHT. The basic principles of this method can be used for protecting ROIs in images transmitted through any lossy packet network and in any generalized MD framework.

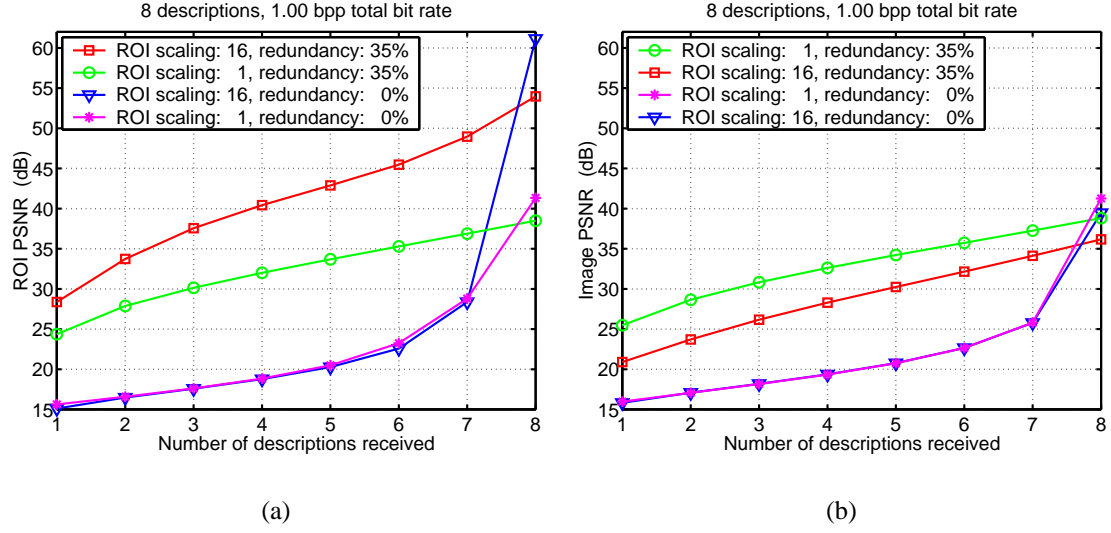


Figure 8: PSNR curves for (a) ROI and (b) image quality for 5 out of 8 descriptions received.




Figure 9: Redundancy allocation throughout the little girl image for the case of (a) the ROI scaled by 16 and (b) no scaling.

6 Acknowledgments

The authors would like to thank Professor William Pearlman for the SPIHT source code and Professor Nariman Farvardin for providing the little girl image.

References

- [1] A. C. Miguel, A. E. Mohr, and E. A. Riskin, "SPIHT for generalized multiple description coding," in *Proceedings of ICIP*, 1999.
- [2] J. G. Apostolopoulos and J. S. Lim, "Critically sampled wavelet representations for multi-dimensional signals with arbitrary regions of support," in *Proceedings of ICASSP*, vol. 3, pp. 1509–12, 1998.
- [3] U. Rauschenbach, "The rectangular fish eye view as an efficient method for the transmission and display of large images," in *Proceedings of ICIP*, 1999.
- [4] E. Atsumi and N. Farvardin, "Lossy/lossless region-of-interest coding based on set partitioning in hierarchical trees," in *Proceedings of ICIP*, vol. 1, pp. 87–91, 1998.
- [5] J. K. Rogers and P. C. Cosman, "Robust wavelet zerotree image compression with fixed-length packetization," in *Proceedings Data Compression Conference*, pp. 418–427, March–April 1998.
- [6] J. M. Shapiro, "Apparatus and method for emphasizing a selected region in the compressed representation of an image." US Patent # 5563960, October 1996.
- [7] D. Nister and C. Christopoulos, "Lossless region of interest with a naturally progressive still image coding algorithm," in *Proceedings of ICIP*, vol. 3, pp. 856–860, 1998.
- [8] T. Frajka, P. G. Sherwood, and K. Zeger, "Progressive image coding with spatially variable resolution," in *Proceedings of ICIP*, vol. 1, pp. 53–56, 1998.
- [9] L. Martel and A. Zaccarin, "Prioritization of region-of-interest (ROI) using embedded coding of wavelets coefficients," in *Proceedings of SPIE Signal Processing: Sensor Fusion, and Target Recognition VII*, vol. 3374, pp. 235–46, 1998.
- [10] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3445–3462, December 1993.
- [11] A. Said and W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 243–250, June 1996.
- [12] V. K. Goyal, J. Kovačević, R. Arean, and M. Vetterli, "Multiple description transform coding of images," in *Proceedings of ICIP*, vol. 1, pp. 674–678, October 1998.
- [13] W. Jiang and A. Ortega, "Multiple description coding via polyphase transform and selective quantization," in *Proceedings of SPIE: Visual Communications and Image Processing*, January 1999.
- [14] P. A. Chou, S. Mehrotra, and A. Wang, "Multiple description decoding of overcomplete expansions using projections onto convex sets," in *Proceedings Data Compression Conference*, pp. 72–81, March 1999.
- [15] A. E. Mohr, E. A. Riskin, and R. Ladner, "Generalized multiple description coding through unequal forward error correction," in *Proceedings of ICIP*, 1999.
- [16] E. A. Riskin, "Optimum bit allocation via the generalized BFOS algorithm," *IEEE Transactions on Information Theory*, vol. 37, pp. 400–402, March 1991.
- [17] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in \mathbf{R}^n : Analysis, synthesis, and algorithms," in *IEEE Transactions on Information Theory*, vol. 44, pp. 16–31, January 1998.