

Protecting regions of interest in medical images in a lossy packet network

Toby Wu ^{*a}, Agnieszka C. Miguel ^a, Eve A. Riskin ^{**a}, Alexander E. Mohr ^b, Richard E. Ladner ^b, Scott Hauck ^a

^a Electrical Engineering, University of Washington, Box 352500, Seattle, WA 98195

^b Computer Science and Engineering, University of Washington, Box 352350, Seattle, WA 98195

ABSTRACT

We present two methods for protecting a region of interest (ROI) in a compressed medical image transmitted across a lossy packet network such as the Internet or a wireless channel. We begin with a high quality wavelet-based coder, the Set Partitioning in Hierarchical Trees (SPIHT) algorithm, which orders data progressively by coding the globally important information first. We then compress the ROI to a higher quality than the rest of the image by scaling the wavelet coefficients corresponding to the ROI. This approach moves ROI information earlier in the bit stream. Finally, we add more redundancy to the ROI than to the rest of the image by two techniques. With MD-SPIHT, we repeat wavelet coefficient trees corresponding to the ROI and code them to higher bit rates than the background trees. With ULP-FEC, we use forward error correction (FEC) in an unequal loss protection framework. We find that both methods increase the probability of receiving high quality ROI in the presence of packet loss.

Keywords: Region of Interest, Set Partitioning in Hierarchical Trees, unequal loss protection, packet loss, network, Multiple Description network, forward error correction, medical image, wavelet coder, wireless channels

1. INTRODUCTION

Networks such as the Internet rely on an exchange of data packets. In traversing the network, a packet is sent from router to router until it arrives at its destination. Yet when congestion occurs, the network is forced to discard packets, leading to the loss of data. These networks do not sort packets preferentially, but instead discard packets at random. If it is undesirable to retransmit the lost packets, the decoder must recover meaningful information from the received packets.

The data that we transmit usually vary in importance. For instance, if we transmit an image of a forest, the *global* data that allow us to recognize the trees are more important than data that show the shape of a few leaves. If the network is unable to transmit all of the forest data, then we would like it to discard the part describing the leaves and retain the part that allows recognition of the trees.

Spatial regions in images that are most important to the end user are called regions of interest (ROIs). 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. In addition, small regions of medical images may be sufficient for diagnostic purposes. For example, if a physician at a remote village transmits her patient's medical image to a city hospital, *local* data that allows a hepatologist (liver specialist) in the city to diagnose the liver are more important than data that show the small intestines.

*Toby Wu (tobywu@u.washington.edu); **Eve A. Riskin (riskin@isdl.ee.washington.edu)

In [1] we presented an algorithm for using a progressive encoder called Set Partitioning in Hierarchical Trees (SPIHT) in a multiple description network (MD-SPIHT). In multiple description (MD) coding schemes, the encoder produces several descriptions of the source data, each exactly filling one packet. To combat packet loss, redundancy is added to the original data before it is compressed in the wavelet domain. In [2], we introduced another algorithm for using SPIHT in lossy packet networks where redundancy is added to the data after the compression process in a forward error correction framework (ULP-FEC). In both algorithms, unequal loss protection is implemented by varying the amount of redundancy with the importance of data.

In this paper, we extend MD-SPIHT and ULP-FEC to the protection of ROIs. To increase the probability that the ROI is received with high quality when data loss occurs, we add 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.

2. BACKGROUND

2.1 Set Partitioning in Hierarchical Trees

We use Set Partitioning in Hierarchical Trees (SPIHT) [3] as our compression algorithm. The data structures and coding blocks used by SPIHT are wavelet coefficients grouped into trees. SPIHT provides a progressive ordering of data that enables us to determine which data are most important to the image quality. In general, progressive image encoders produce a bit stream where the earlier bits are refined by the later bits. For instance, if a user wishes to compress a portrait of a face, bits that allow us to recognize the person are more important than bits that show the detailed texture of the person's face. High quality wavelet-based algorithms, such as SPIHT, encode the low quality approximation of the image using the earlier bits while the later bits serve to further refine the image. Thus by progressively ordering the data, SPIHT sends the *globally* most relevant information first. In this paper, we use SPIHT to compress a medical image although other progressive encoders such as JPEG 2000 would suffice as well.

2.2 Set Partitioning in Hierarchical Trees with ROI coding

Previous techniques code the ROI either separately from the rest of the image [4]; at a higher spatial resolution such as in foveated image compression [5]; or at a higher bit rate than the rest of the image [6,7,8,9,10,11]. With the recent popularity of zerotree based coding, some of these methods have been incorporated into Packetized Zerotree Wavelet (PZW) algorithm [7], Embedded Zerotree Wavelet (EZW) algorithm [14] and SPIHT [3]. Most zerotree ROI coding schemes involve assigning the ROI a higher bit rate than the rest of the image by multiplying the wavelet coefficients corresponding to the ROI by some large factor. This way, the locally important information is sent in the earlier parts of the compressed bit stream. This method requires both the encoder and the decoder to know the ROI parameters and the multiplication factor, but does not involve any modifications to the original algorithm. However, none of the current ROI coding methods addresses the protection of the ROI in the presence of data loss. In this work we use the zerotree ROI coding scheme in SPIHT to code a medical image. We add more redundancy to the earlier parts of the bit stream and therefore protect ROI data more than the background.

2.3 Multiple Description SPIHT

The MD-SPIHT algorithm [1] adds controlled amounts of redundancy to the original data before the compression process. It adds more redundancy to the earlier parts of the bit stream than to the later parts in the form of repeated wavelet coefficient trees. To increase the resilience of SPIHT to data loss in a MD framework, we first modify SPIHT in manner similar to PZW [7]. Spatially dispersed wavelet coefficient

trees are grouped together, coded with SPIHT, and transmitted in one packet. We use arithmetic coding within each packet and code each packet to the same bit rate. Next, we assign each tree a set of redundant trees which is coded to a successively lower bit rate. Coding a tree to a lower bit rate corresponds to sending the most important data. The original trees and their redundant trees are located in many different packets and the most important part of each tree is sent in more packets than other parts.

As shown in Fig. 1, four original trees (1-4) are sent in four packets (P1-P4) with two sets of redundant trees. The completely filled trees correspond to the original trees while the partially filled trees represent redundant trees coded to successively lower bit rates. If some packets are lost, the most important parts of the original trees in those lost packets will be recovered because redundant trees coded at lower bit rates will be present in the received packets. In Fig. 1, for instance, if two packets are lost (P1 and P4), the first and fourth trees are partially recovered from the received packets.

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 packet 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 packets is lost. We use the extended BFOS algorithm [12], which allocates bits among various tree copies such that all allocations lie on the lower convex hull of the rate-distortion curve.

2.4 Unequal Loss Protection - Forward Error Correction

A second algorithm called Unequal Loss Protection - Forward Error Correction (ULP-FEC) [2] adds unequal amounts of redundancy to the bit stream but after the compression process. The redundancy is in the form of systematic Reed-Solomon codes, which are a form of forward error correction (FEC). The ULP-FEC algorithm dynamically chooses the length and the content of the each message fragment. We use a number of message fragments equal to the number of bytes in each packet. We add FEC to each message fragment to protect against packet loss such that the fragment and the FEC form a stream.

As shown in Fig. 2(a), a 32 byte message and 10 bytes of FEC are partitioned into seven streams with each stream containing one byte from each of the six packets. We add more FEC to the earlier parts of the message and less to the later parts because SPIHT generates a progressive bit stream where the earlier parts of the message contain the information most vital to image quality. In the event of data loss, we can recover the entire stream as long as the number of packets lost is less than or equal to the number of FEC bytes in that stream. Figures 2(a) to 2(c) illustrate the recovery sequence when one packet out of six is lost and five are properly received. The first six streams can be decoded because they contain at least one FEC byte. The last stream cannot be recovered because it contains no FEC byte.

3. METHODOLOGY

3.1 Multiple Description SPIHT with ROI coding

In this work, we extend MD-SPIHT to code images with ROIs [13]. We protect ROIs more than other parts of the image so that if data loss occurs, high quality ROIs can still be reconstructed. We repeat redundant trees corresponding to the ROI more often than the non-ROI redundant trees. In addition, we code the ROI trees to a higher bit rate than non-ROI trees. As a result, the ROI trees have higher probability of being received because they are present in many more packets than the non-ROI trees.

A simple example illustrating the function of MD-SPIHT 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 gray. In the first case (Fig. 3(a)), two levels of redundancy are allocated to all trees. As a result, when 4 descriptions are lost, the ROI trees may be unrecoverable. In the next case (Fig. 3(b)),

corresponding to the application of MD-SPIHT in ROI coding, the second level of redundancy is allocated exclusively to ROI trees. If 4 descriptions are again lost, the two ROI trees are recovered.

Using the schemes proposed by Shapiro [8], Rogers and Cosman [7], and Atsumi and Farvardin [6], the wavelet coefficients corresponding to the ROI are multiplied by a scale factor. The magnitude of the scaling factor used reflects the level of ROI prioritization over other regions in the image. Scaling the ROI coefficients causes them to be coded much earlier. Again, because MD-SPIHT adds more redundancy to the earlier parts of the bit stream than to later parts, the ROI will have a higher probability of being received intact in the event of data loss. This implementation requires sending the ROI position and scaling factor as a side information.

3.2 Unequal Loss Protection - Forward Error Correction with ROI coding

We next extend ULP-FEC to code images with ROIs. Again, the wavelet coefficients corresponding to the ROI are scaled by a large factor, but the resulting bit stream is protected with unequal amounts of FEC instead of incorporating redundant trees. A simple example illustrating the function of ULP-FEC is shown in Fig. 4. The image consists of 23 bytes sent in six streams. The amount of ROI information contained in each stream is proportional to the number of lines under the bytes in the stream. More FEC is distributed to the earlier streams than to the later streams. In the first case (Fig. 4(a)), the ROI information is distributed equally throughout all streams. The ROI information in the last stream will be unrecoverable if two packets are lost. In the next case (Fig. 4(b)), corresponding to the proposed method, the wavelet coefficients belonging to the ROI are scaled by a large factor. As a result, the majority of ROI information is sent in the earlier streams. As shown in the figure, more ROI information can be recovered with FEC if two packets are lost. Therefore in the presence of data loss, a higher quality ROI can be reconstructed.

4. RESULTS

4.1 Multiple Description SPIHT

We demonstrate our results on a 256 x 256 Magnetic Resonance Image (MRI) of a brain shown in Fig. 5(a) using MD-SPIHT. The chosen ROI, which comprises 22% of the total image, is outlined in Fig. 5(b). We set the number of descriptions to eight and the total bit rate to 1.5 bpp.

We study the case when only 5 packets are received and a ROI scaling factor of 16 is used. For no ROI scaling and no redundancy, the ROI cannot be fully recovered (Fig. 5(c)). Scaling the ROI without 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. 5(d)). When redundancy is added but no ROI scaling is done, the redundancy is spread to all parts of the image. The ROI is visible but its quality is low (Fig. 5(e)). Finally, the ROI is scaled and redundancy is added. Because the redundancy is highly concentrated in the ROI area, the ROI is fully received and of high quality (Fig. 5(f)).

4.2 Unequal Loss Protection - Forward Error Correction

Next, we demonstrate our results on the same 256 x 256 Magnetic Resonance Image (MRI) of a brain shown in Fig. 5(a) using ULP-FEC. We set the number of packets to eight and the total bit rate to 1.5 bpp.

First, we investigate the influence of the ROI scaling factor on the reconstruction quality in ULP-FEC. We allocate redundancy for a 37.5% expected loss rate and ROI scaling factors ranging from 1 to 16. Figures 6(a) and 6(b) show the reconstruction quality for the ROI and whole image, respectively. We see that the higher the ROI scaling factor, the lower the ROI distortion as expected. The opposite is true for the overall image quality. However, we can obtain a substantial improvement in the ROI quality with only moderate degradation of the background. We want to choose a scale factor that exhibits high ROI quality when few

packets are lost and graceful degradation in ROI quality with greater packet loss. A scale factor of 16 satisfies all these criteria. It exhibits high ROI quality when up to 3 packets are lost and graceful degradation in ROI quality when more packets are lost.

In addition, we investigate the effect of the expected loss rate on the reconstruction quality in ULP-FEC. We use a ROI scaling factor of 16 and vary the expected loss rate from 5% to 45%. Figures 7(a) and 7(b) show the reconstruction quality for the ROI and whole image, respectively. With an expected loss rate of 25%, for instance, we observe that high ROI quality is obtained when few packets are lost and graceful degradation in ROI quality occurs when we lose additional packets.

4.3 Multiple Description SPIHT and Unequal Loss Protection - Forward Error Correction

Finally, we compare MD-SPIHT with ULP-FEC. The results for two expected packet loss rates of 5% and 25% and an ROI scaling factor of 16 are shown in Fig. 8. When all eight packets are received, notice that ULP-FEC designed for 25% expected loss rate and MD-SPIHT designed for 5% expected loss rate both produce images with nearly identical PSNR values: the amount of redundancy allocated by the two schemes is approximately the same. ULP-FEC gives better ROI quality for loss rates lower or equal to the rate it has been designed for. MD-SPIHT gives better ROI quality at higher than expected loss rates. Therefore, ULP-FEC is more suited for protecting data transmitted through communication channels where the loss rate is as expected. On the other hand, it is preferable to use MD-SPIHT in wireless channels where the loss rate is difficult to estimate and where environmental changes can cause the loss rate to deviate from what is expected.

5. CONCLUSIONS

We have shown two algorithms, MD-SPIHT and ULP-FEC, for protecting a region of interest from data loss. Both algorithms add varied amounts of redundancy to the bit stream. To increase the probability that the ROI is received with high quality, most of the redundancy is concentrated in the region of interest. The ROI is therefore heavily protected at the expense of lower protection in the background. The results for both algorithms show that high ROI quality is achieved even when only a few packets are received.

The ROI protection schemes proposed in this work are not limited to a single ROI per medical 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. The basic principles of this method can be used for protecting ROIs in medical images transmitted in any lossy packet network.

ACKNOWLEDGEMENTS

The authors would like to thank Professor William Pearlman for the SPIHT source code.

REFERENCES

- [1] A. C. Miguel, A. E. Mohr, and Eve A. Riskin, "SPIHT for generalized multiple description coding," in *Proceedings of ICIP*, volume 3, pages 842-846, 1999.
- [2] A. E. Mohr, Eve A. Riskin, and Richard E. Ladner, "Unequal Loss Protection: Graceful Degradation Over Packet Erasure Channels Through Forward Error Correction," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 819-828, June 2000.
- [3] A. Said and W.A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. on CSVT*, vol. 6, no. 3, pp. 243-250, June 1996.

- [4] J. G. Apostolopoulos and J.S. Lim, "Critically sampled wavelet representations for multidimensional signals with arbitrary regions of support," *Proceedings of ICASSP*, vol. 3, pp. 1509-12, 1998.
- [5] U. Rauschenbach, "The rectangular fish eye view as an efficient method for the transmission and display of large images," in *Proceedings of ICIP*, vol. 3, pp.115-119, 1999.
- [6] 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.
- [7] 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.
- [8] J. M. Shapiro, "Apparatus and method for emphasizing a selected region in the compressed representation of an image." US Patent # 5563960, October 1996.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] E. A. Riskin, "Optimum bit allocation via the generalized BFOS algorithm," in *IEEE Transactions on Information Theory*, vol. 37, pp. 400-402, March 1991.
- [13] A. C. Miguel and Eve A. Riskin, "Protection of Regions of Interest Against Data Loss in a Generalized Multiple Description Framework," *Proceedings of DCC*, pp. 562, March 2000.
- [14] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3445-3462, December 1993.

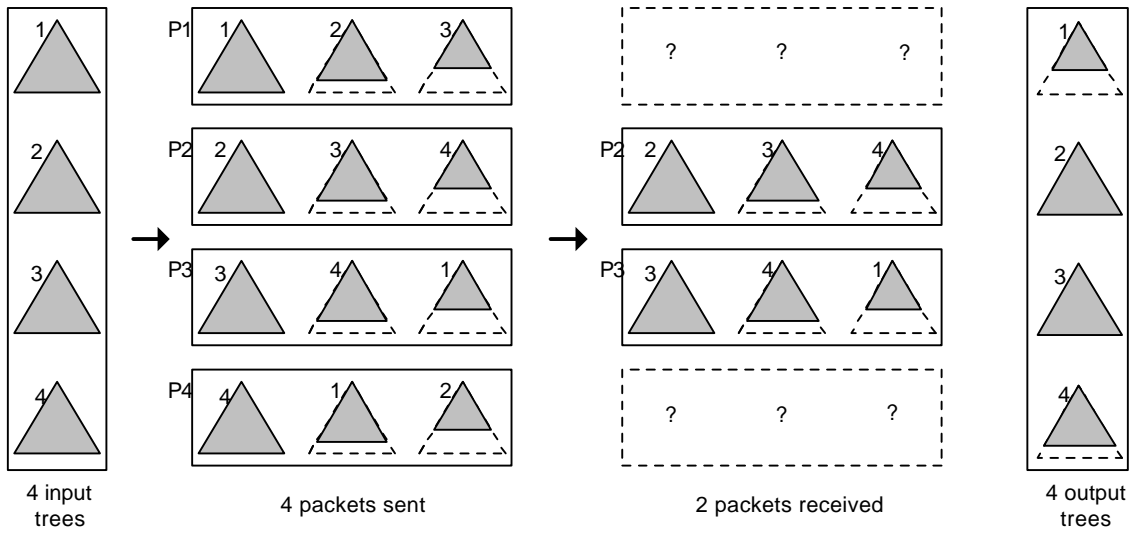


Figure 1: MD-SPIHT coding example. Four original trees (1-4) are sent in four packets (P1-P4) with two sets of redundant trees. The completely filled trees correspond to the original trees while the partially filled trees represent redundant trees coded to successively lower bit rates. If two packets are lost (P1 and P4), the first and fourth trees are partially recovered from the received packets.

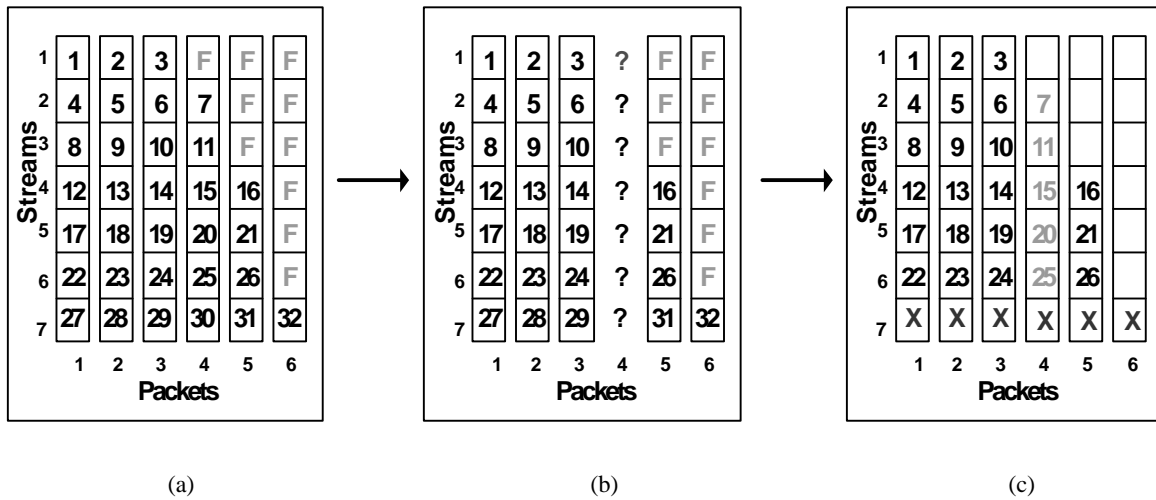
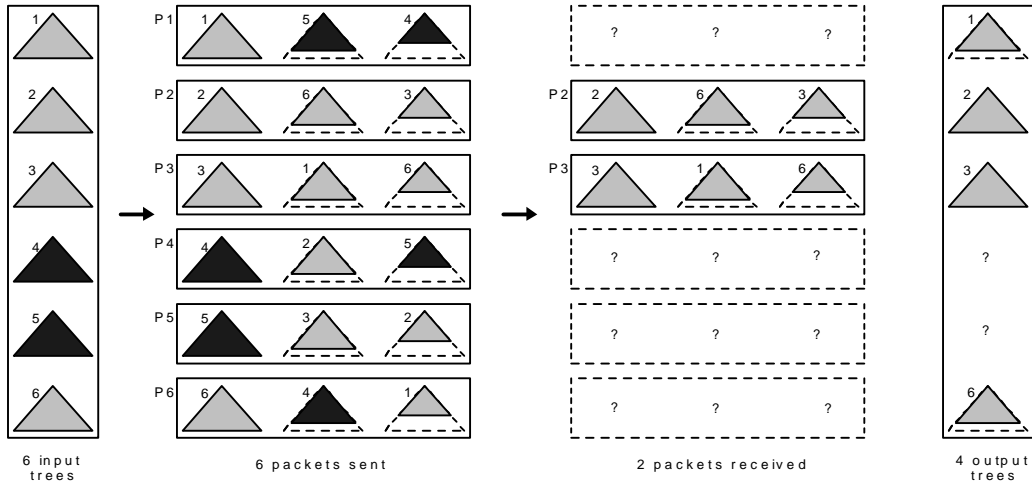
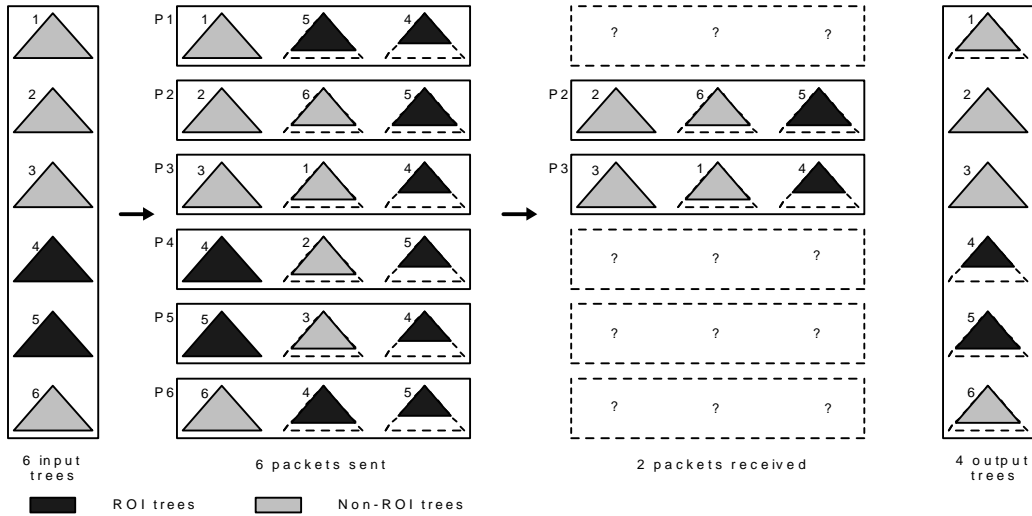


Figure 2: ULP-FEC coding example. (a) A message of 32 bytes of data (1-32) and ten bytes of FEC (F) is divided into seven streams and sent in six packets. Each row is a stream and each column contains a packet. (b) If one packet is lost (P4), (c) all the lost bytes in streams two through six can be recovered with FEC bytes.

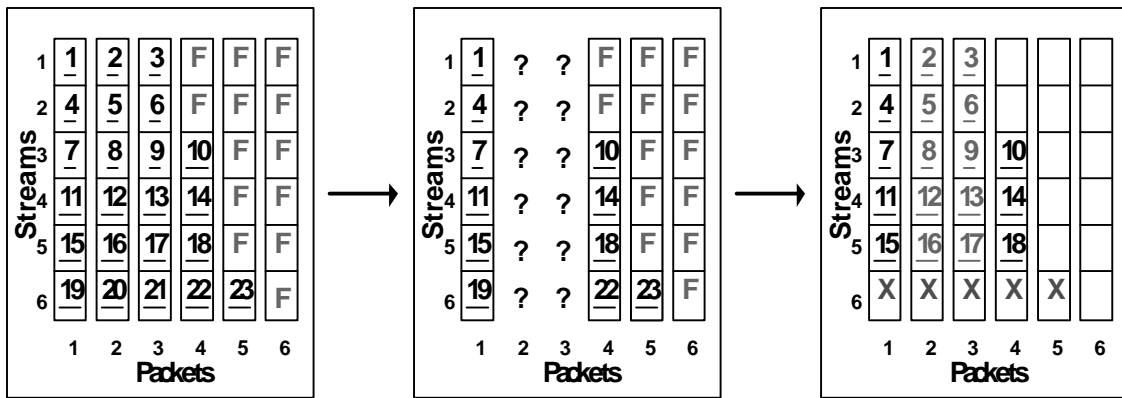


(a)

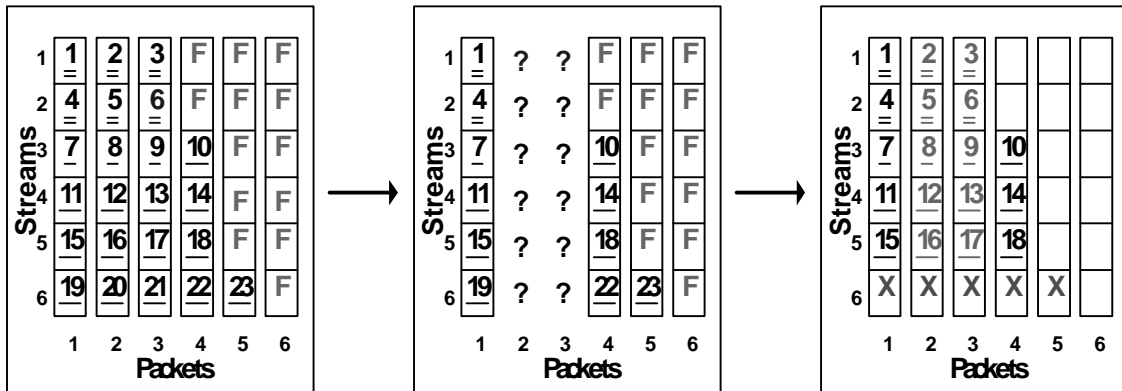


(b)

Figure 3: MD-SPIHT with ROI coding example: Six original trees (1-6) are sent in six packets (P1-P6). Two of the original trees (4 and 5) contain ROI coefficients. (a) All trees are allocated two levels of redundancy. If only packets P2 and P3 are received, the ROI trees cannot be recovered. (b) The second level of redundancy is allocated only to ROI trees. If only packets P2 and P3 are received, both ROI trees are partially recovered from their redundant copies in the received packets.




(a)




(b)


Figure 4: ULP-FEC with ROI coding example: Twenty-three bytes (1-23) are sent in six streams. The amount of ROI information contained in each stream is proportional to the number of lines under the bytes in the stream. Earlier streams receive more FEC protection. (a) When the ROI is not scaled, ROI information is distributed equally throughout all streams. If more than one packet is lost, ROI information contained in the last stream cannot be recovered. (b) The ROI is scaled and the bulk of ROI information is contained in the first and second streams. Thus, more ROI information can be recovered compared to not scaling the ROI if no more than 3 packets are lost.




(a) Original MRI brain image.




(b): Region of Interest outlined.




(c) Image: 20.41 dB, ROI: 23.39 dB



(d) Image: 18.51 dB, ROI: 23.47 dB




(e) Image: 32.78 dB, ROI: 32.08 dB




(f) Image: 25.93 dB, ROI: 37.52 dB

Figure 5: MD-SPIHT results. (a) Original image. (b) Region of interest outlined. Image and ROI quality for 5 out of 8 packets received. The total bit rate is 1.5 bpp. (c) ROI scaling factor: 1, with no redundancy, (d) ROI scaling factor: 16, with no redundancy, (e) ROI scaling factor: 1, with redundancy, (f) ROI scaling factor: 16, with redundancy.




(a)

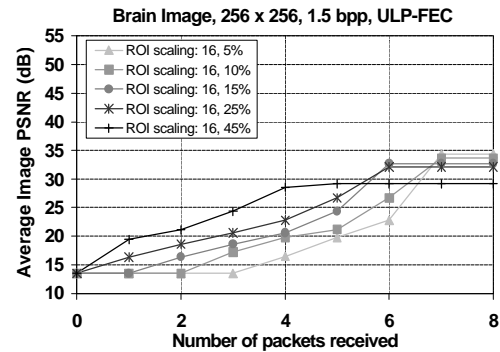


(b)

Figure 6: ULP-FEC results for the brain image coded to 1.5 bpp in eight packets for five different ROI scaling factors and an expected packet loss rate of 37.5%. (a) Results for the ROI. (b) Results for the image.

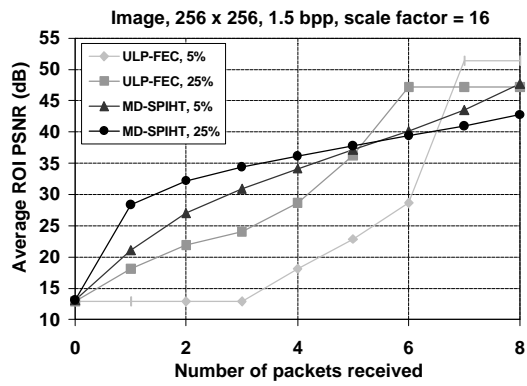


(a)




(b)

Figure 7: ULP-FEC results for the brain image coded to 1.5 bpp in eight packets with five different expected packet loss rates and a constant scaling factor of 16. (a) Results for the ROI. (b) Results for the image.



(a)



(b)

Figure 8: Comparison between MD-SPIHT and ULP-FEC at a constant ROI scale factor of 16 for two different expected packet loss rates. Results for the (a) ROI and (b) image.