

Coding and Transmission of Subband Coded Images in the Internet*

Benjamin W. Wah and Xiao Su

Department of Electrical and Computer Engineering
and the Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

ABSTRACT

Subband-coded images can be transmitted in the Internet using either the TCP or the UDP protocol. Delivery by TCP gives superior decoding quality but with very long delays when the network is unreliable, whereas delivery by UDP has negligible delays but with degraded quality when packets are lost. Although images are delivered currently over the Internet by TCP, we study in this paper the use of UDP to deliver multi-description reconstruction-based subband-coded images. First, in order to facilitate recovery from UDP packet losses, we propose a joint sender-receiver approach for designing *optimized reconstruction-based subband transform* (ORB-ST) in multi-description coding (MDC). Second, we carefully evaluate the delay-quality trade-offs between the TCP delivery of SDC images and the UDP and combined TCP/UDP delivery of MDC images. Experimental results show that our proposed ORB-ST performs well in real Internet tests, and UDP and combined TCP/UDP delivery of MDC images provide a range of attractive alternatives to TCP delivery.

Keywords: single-description coding (SDC), multi-description coding (MDC), real-time multimedia in the Internet, reconstruction-based subband image coding, TCP, UDP, World-Wide Web

1. INTRODUCTION

Quality and delay are two key performance measures to evaluate the delivery of coded images. Previously, high quality in delivery is considered more important because image data is not real time in nature and is generally sent using a reliable transport protocol like TCP.

With the advent of the World Wide Web, trade-offs between quality and delay in transferring image data may need to be changed. Oftentimes, when there are multiple images to be transferred from a Web server, users may prefer to see (slightly) degraded images in (much) faster turnaround time than to wait for a long time to see high-quality images. TCP delivery in such cases is not desirable because it incurs intolerable long delays. On the other hand, UDP delivery incurs shorter end-to-end delays but cannot be used for sending conventionally coded images because packet losses may render these images not decode-able.

To transfer coded images with shorter end-to-end delays, our goal in this paper is to design schemes for reconstructing lost information when image data is subband coded and sent by UDP.

Existing error-concealment schemes are performed in the sender, or in the receiver side, or in both sides.

Research supported by a grant from Mindspeed Technologies.

E-mail: {wah, xiao-su}@manip.crhc.uiuc.edu, URL: <http://manip.crhc.uiuc.edu>

Sender-side error concealments of coded images consist of two popular methods: *layered coding* and *multiple description coding*.

*Layered coding*¹ partitions images into a base layer with visually important image data and a few enhancement layers. In networks with priority support, the base layer is normally assigned a higher priority so that it has a larger chance to be delivered error free. However, layered coding is not applicable in the Internet for two reasons. First, the current Internet does not provide priority delivery service for different layers. Second, when part of the base layer is lost, the decoding quality becomes very poor, because the lost bit stream cannot be concealed.

In contrast, *multi-description coding* (MDC) divides image data into equally important streams. For subband coded images using MDC, scalar quantizers^{2,3} have been designed in order to produce two descriptions using two side-scalar quantizers. These schemes, however, have very complicated encoding and decoding algorithms that make them infeasible for low-delay transmissions.

Receiver-based recovery is usually formulated as heuristic optimizations based on the smoothness assumption of image pixels. One approach formulates spatial smoothness constraints into convex sets and derives a solution iteratively.⁴ Other approaches minimize the variations along edge directions or local geometric structures.^{5,6} Besides being computationally expensive, mistakes in detection of image structures may yield annoying artifacts and blurred edges.

Sender-receiver-based schemes require senders and receivers to cooperate in error concealments. The first kind of approaches, *joint source channel coding* (JSCC),^{7,8} minimize transmission errors by jointly designing the quantizer and the channel coder, according to a given channel-error model and feedbacks from receivers. They, however, are hard to apply in the Internet, since the Internet does not have a well-defined channel model. The second approach, interleaving with reconstruction,⁹ is a very simple and efficient way to generate multiple descriptions. However, this approach may be deficient because it generates multiple descriptions in the sender side, without considering the correlations between the descriptions generated and the reconstruction operation at the receiver side.

In short, existing error concealment techniques either rely on the inadequate capability of senders or receivers to do reconstruction, or make certain assumptions about transmission channels in designing encoders. None of them considers the reconstruction process performed at receivers.

In this paper, we study a joint sender-receiver-based coding and reconstruction scheme for the delivery of multi-description coded images by UDP. In our system, we *interleave* adjacent pixels of an image into multiple descriptions, decompose each description into segments so that each segment fits in a packet, code each segment using a nonredundant error-concealment coding scheme, and transmit the packets to the destination.

To design multi-description coders at senders, we adopt a joint sender-receiver approach, instead of using previous approaches that design coders independent of reconstruction methods. The coder at a sender applies an *optimized reconstruction-based subband transform* (ORB-ST) that minimizes the reconstruction error, when some of the descriptions are lost and reconstructed using average interpolations from the descriptions received. We have adopted a simple reconstruction algorithm at receivers in order to facilitate fast playback.

Since the proposed MDC can generate coded streams resilient to packet losses, we deliver them using the unreliable but fast UDP or combined TCP/UDP protocols. Our proposed approach can lead to good reconstruction quality with small end-to-end delays but, as expected, degraded decoding quality when compared to the TCP delivery of single-description coded (SDC) images.

This paper is organized as follows. Section 2 studies end-to-end delays and packet-loss patterns of Internet transmissions. The statistics helps guide the design and evaluation of ORB-ST in Section 3.




Figure 1. Round-trip delays of sending 64 UDP packets and the same data in TCP packets encapsulated in UDP ones to the UDP echo port of three remote computers. The experiments were carried out at the beginning of each hour for a 24-hour period on April 8, 2001.

Section 4 further evaluates delay and quality trade-offs of our proposed approach. Finally, Section 5 concludes the paper.

2. TRANSMISSION DELAYS AND LOSS BEHAVIOR IN THE INTERNET

We study in this section the end-to-end delays of both TCP and UDP delivery in the Internet and the loss behavior of UDP delivery.

2.1. Experimental Setup for Statistics Collection

From a site in Urnaba (cw.crhc.uiuc.edu), we chose three destination sites in our experiments. The first one is a domestic site (daedalus.cs.berkeley.edu) representing a low-loss connection, the second to the United Kingdom (www.uea.ac.uk), representing a medium-loss connection, and the last to China (www.shmu.edu.cn), representing a high-loss connection.

Since we have no control of these destination computers, we carried out our experiments by sending packets to the echo port of each of the destinations. To ensure fair delay comparisons between TCP and UDP, we modified the Linux kernel so that TCP echo packets were encapsulated in UDP ones and sent to the UDP echo port of the remote server. In this way, encapsulated TCP packets would be echoed immediately when they were received at the echo server, just as UDP packets. In determining the number of packets sent, we assume that a 512-by-512 image is compressed at an 8 : 1 ratio and sent in 512-byte packets, leading to 64 packets sent per image.

2.2. Comparisons of End-to-End Delays in TCP and UDP Transmissions

Figure 1 shows the end-to-end delays of sending 64 UDP packets and 64 TCP packets encapsulated in UDP packets to the UDP echo port of three remote servers. The experiments were carried out at the beginning of each hour for a 24-hour period on April 8, 2001.

To avoid overflow of receiver buffers when all 64 packets were dumped to a remote UDP echo port simultaneously, we sent them in three batches, each consisting of 20 packets and separated by 20 ms. The 20-ms delay was the minimum chosen in such a way that longer delays did not lead to lower average loss rate. (Such a choice is, of course, not TCP friendly.) In contrast, the pacing and retransmissions of TCP packets were controlled by the TCP protocol itself.

The graphs in Figure 1 show that the end-to-end response times of UDP delivery have far less variations and are shorter than those of TCP delivery. For example, for transmissions between Urnaba and UK




Figure 2. $Pr(fail|i)$, probability of bursty losses that cannot be recovered, conditioned on interleaving factor i , at different times on April 8, 2001.

(Figure 1b), UDP delays range from 0.2 to 0.4 seconds, whereas TCP delays range from 2 to 7 seconds. In terms of average speed, TCP packets are one to two orders slower than UDP ones. The long delay in TCP delivery is attributed to TCP's coarse grained timeouts and congestion avoidance algorithms. In practice, this means that we may need to wait for over two minutes in downloading an image from China when using TCP.

2.3. Loss Behavior of UDP Transmissions

Although UDP delivery is much faster, it suffers from losses that may lead to large degradations in image quality or render images not decode-able. In order to send images using UDP, we need to understand its loss behavior and conditions under which losses can be concealed.

It is well known that interleaving can convert bursty losses to random losses and facilitate recovery by exploiting inherent redundancies in image data. In order to conceal bursty losses most of the time, it is necessary to choose interleaving factor i so that the probability of packets that are not recoverable using i is small enough.

Let the total number of packets sent be n_p and the interleaving factor be i . Over all the interleaved sets, assuming that losses of j consecutive packets, $j \leq i$, happen m_j^i times, then $Pr(fail | i)$, the unconditional probability that a packet cannot be reconstructed in the stream received using interleaving factor i , can be computed as follows:

$$Pr(fail | i) = Pr(fail | loss, i) \times Pr(loss) = \frac{i \times m_i^i}{n_s} \times \frac{n_s}{n_p} = \frac{i \times m_i^i}{n_p}. \quad (1)$$

Figure 2 shows that $Pr(fail | i)$ drops quickly with increasing interleaving factor i . In summary, the probability of not able to reconstruct a lost packet is under 5% with an interleaving factor of 2 for connections to California and UK, and with an interleaving factor of 4 to China.

Based on the statistics collected, we conclude that UDP delivery can be one to two orders faster than TCP delivery, and that bursty losses in UDP delivery can be concealed by interleaving and by reconstruction using an interleaving factor of four for most Internet destinations.

3. ORB-ST FOR CONCEALING BURSTY LOSSES

Although interleaving and interpolations are effective for concealing bursty losses, simple coding of interleaved streams may not work well because the original coding algorithm may not be the best for




Figure 3. Basic building blocks of a modified codec. (The shaded block is our proposed ORB-ST.)

reconstructing lost streams. In this section we propose a new optimized reconstruction-based subband transform (ORB-ST) that takes into account the reconstruction process at receivers. A different derivation of an optimized reconstruction-based DCT transform for video coding can be found elsewhere.¹⁰

In the following, we first derive ORB-ST based on partitioning image data into two descriptions. Its extension to four descriptions can be found in our previous work¹⁰ and is omitted here. Next, we compare the quality of images transformed by ORB-ST and the original subband transform (ST) by conducting experiments in the Internet.

3.1. ORB-ST for Two Descriptions

Figure 3 shows the basic building blocks in our proposed subband image coding system. It is based on existing state-of-the-art image codecs that consist of several stages: a subband transformation, a quantizer and an optional entropy coder.

Assume that each row of the original image, \vec{x} of size n , is transformed into \vec{c}_1 of size $\frac{n}{2}$ and \vec{c}_2 of size $\frac{n}{2}$, corresponding to the descriptions of odd-numbered and even-numbered pixels. Here, \vec{c}_i , $i = 1, 2$, is an interleaved vector of components from \vec{c}_i^0 and \vec{c}_i^1 , where \vec{c}_i^j is the output from subband j , and subbands are ordered from low to high frequency.

Our objective is to find \vec{c}_1 and \vec{c}_2 in order to minimize \mathcal{E}_r , the reconstruction error between the reconstructed output \vec{z}_i and the input \vec{x} . If we consider quantization, the minimization of \mathcal{E}_r becomes an integer optimization problem, because \vec{c}_i takes integer values. Such optimizations are computationally prohibitive in real time. In the following, we derive an approximate solution that does not take into account quantization effects. Since the derivations are similar, we only show that for \vec{c}_1 .

As the synthesis system, consisting of up-sampling, $G_0(z)$, and $G_1(z)$, is equivalent to a linear transform \hat{G} in spatial domain, the output \vec{y}_1 , after synthesis filtering, can be calculated as:

$$\vec{y}_1 = \hat{G} \vec{c}_1. \quad (2)$$

The set of interpolated pixels, \vec{z}_1 , obtained by inserting even-numbered columns as the average of columns from \vec{y}_1 , with the boundary column duplicated, can also be expressed as a linear transform of \vec{y}_1 as follows:

$$\vec{z}_1 = \mathbf{U} \vec{y}_1 = \mathbf{U} \hat{G} \vec{c}_1. \quad (3)$$

Hence, the distortion between the original and the reconstructed pixels becomes:

$$\mathcal{E}_r = \left\| \mathbf{U} \hat{G} \vec{c}_1 - \vec{x} \right\|^2 = \left\| \mathbf{P} \vec{c}_1 - \vec{x} \right\|^2. \quad (4)$$

Since the linear system of equations $\mathbf{P} \vec{c}_1 = \vec{x}$ is an over-determined one, there exists at least one least-square solution \vec{c}_1 that minimizes (4), according to the theory of linear algebra.¹¹ Specifically, the solution \vec{c}_1 with the smallest length $|\vec{c}_1|^2$ can be found by first performing SVD decomposition of matrix \mathbf{P} :

$$\mathbf{P} = \mathbf{S} [\text{diag}(w_j)] \mathbf{D}^T, \quad j = 1, 2, \dots, \frac{n}{2}, \quad (5)$$

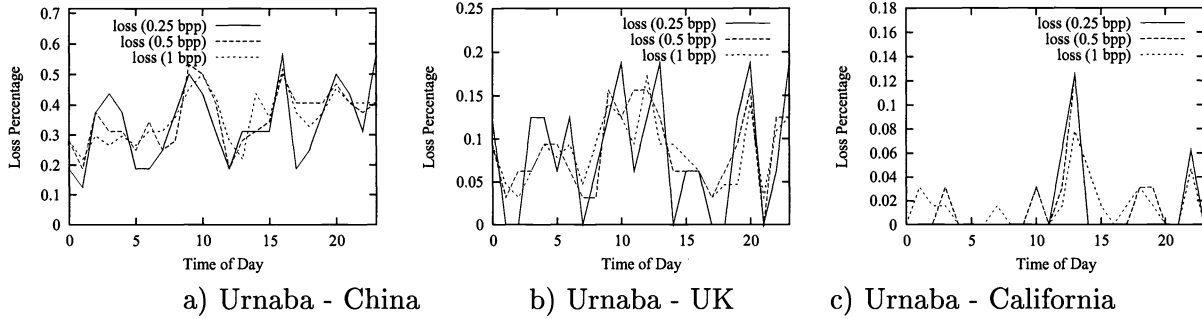


Figure 4. Loss rates of 16-, 32- and 64-packet transmissions from Urnaba to three remote sites.

where \mathbf{S} is an $n \times \frac{n}{2}$ column-orthogonal matrix, $[diag(w_j)]$, an $\frac{n}{2} \times \frac{n}{2}$ diagonal matrix with positive or zero elements (singular values), and \mathbf{D} , an $\frac{n}{2} \times \frac{n}{2}$ orthogonal matrix. Then the least-square solution can be expressed as:

$$\vec{c}_1 = \mathbf{D} [diag(1/w_j)] \mathbf{S}^T \vec{x}. \quad (6)$$

In the above diagonal matrix $[diag(1/w_j)]$, element $1/w_j$ is replaced by zero if w_j is zero. Therefore, ORB-ST is a product of three matrices: \mathbf{D} , $[diag(1/w_j)]$, and \mathbf{S}^T .

To derive the ORB-ST transform for \vec{c}_2 , simply replace \mathbf{U} by the matrix corresponding to interpolation of odd-numbered columns, and the rest of the steps are similar.

When both \vec{c}_1 and \vec{c}_2 are available at the receiver, we can apply an inverse transform to achieve perfect reconstruction, where the inverse transform is derived by inverting the matrix that is made up of interleaved row vectors from ORB-ST for \vec{c}_1 and \vec{c}_2 . In practice, perfect reconstruction is not always achievable due to errors caused by truncations of floating point numbers as well as quantization.

3.2. Comparisons of ORB-ST and ST

To evaluate our proposed scheme, we built a prototype and tested the quality of frames reconstructed by linear interpolation of adjacent pixels received when the original frame was either ST transformed or ORB-ST transformed. For a fair comparison under the same traffic conditions, we did trace-driven simulations by applying reconstructions on the trace of packets received in real Internet transmissions (see Section 2).

In order to packetize coded descriptions in such a way that each packet can be decoded independent of other packets, we first divide an image into equal-size segments, code each segment using the same bit rate, and place every description from each coded segment into a distinct packet. Since we have used suboptimal strategies of dividing segments into equal sizes and coding each using the same bit rate, we expect losses in image quality when compared to MDC without segmentation.

Figure 4 plots the loss rates of traces over a 24-hour period when sending 16, 32, and 64 packets at the beginning of each hour to the remote UDP echo ports of servers in China, UK and California, respectively. Figure 5 compares the reconstruction quality of sending two test images, *barbara* and *peppers*, using the traces obtained, when each image was coded at, respectively, 0.25 bpp, 0.5 bpp, and 1 bpp and put into 16, 32, and 64 packets for transmission.

For the Urnaba-China connection, ORB-ST outperforms ST at all bit rates for both images, with an average of 0.31-0.38 dB gain for the 0.25-bpp case, 0.44-0.48 dB gain for the 0.5-bpp case, and 0.61-0.88 dB gain for the 1-bpp case. Quality gains improve with increasing bit rates when there is less quantization noise. When entire interleaved sets were lost at certain hours, they were simply filled in by the average of image pixels, leading to significant quality degradations (such as hours 9, 11, 17 and 19 at 1 bpp).

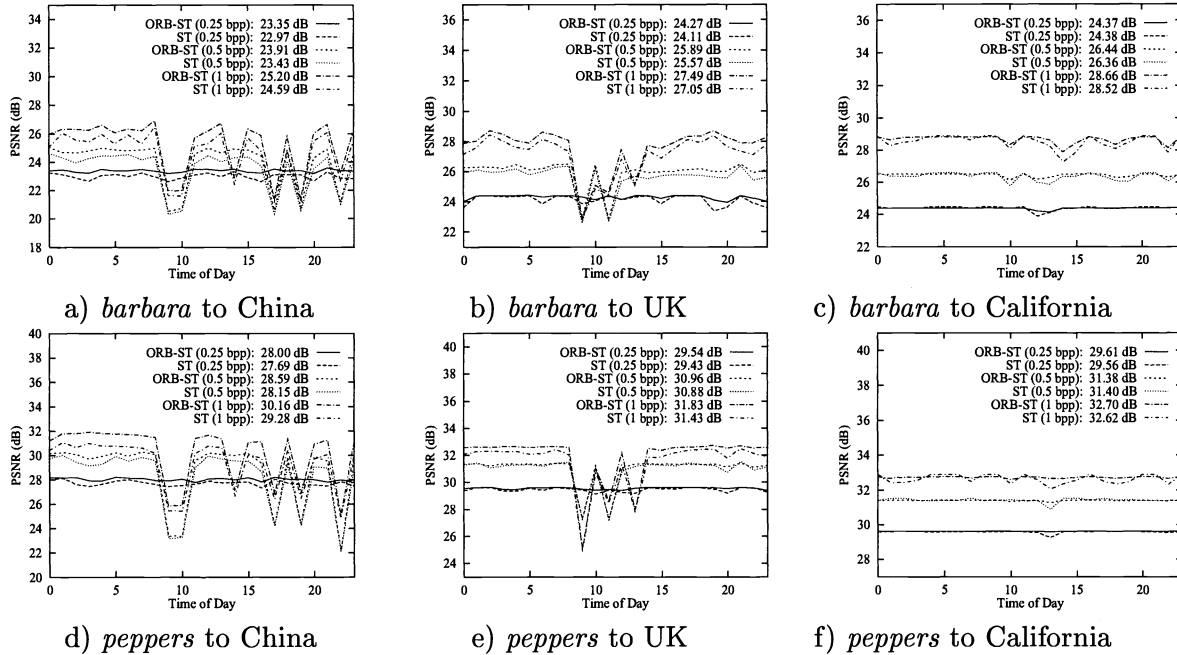


Figure 5. Comparisons of reconstruction quality for image *barbara* and *peppers* over a 24-hour period for the transmissions from Urnaba to three remote sites, when each image was coded at, respectively, 0.25 bpp, 0.5 bpp, and 1 bpp, and placed in 16, 32, and 64 packets for transmission.

For the Urnaba-UK connection, the average reconstruction quality based on ORB-ST is 0.16-0.44 dB better for image *barbara*, and 0.08-0.40 dB better for image *peppers*, depending on their coded bit rates. For the Urnaba-California connection, the reconstruction quality of the two schemes are comparable. In these two cases, the gain of performing ORB-ST is, in general, not as much as in the Urnaba-China connection because the gain is offset by degradations when all the descriptions are received under low loss rates.

These results lead us to conclude that ORB-ST is more suitable for the delivery of images over unreliable channels than the original ST.

4. EVALUATION OF DELAY-QUALITY TRADE-OFFS

4.1. TCP and UDP Delivery of Images

In this section, we evaluate the delay-quality trade-offs between the UDP delivery of MDC images and the TCP delivery of SDC images. Figure 6 shows such trade-offs at 12 noon local time of the remote server using five modes of delivery: a) TCP delivery of SDC image data, b) TCP delivery of MDC data in which the image is not segmented, c) TCP delivery of MDC data in which the image is segmented, d) UDP delivery of MDC ST-coded and segmented image data, and e) UDP delivery of MDC ORB-ST-coded and segmented image data. Results at other times are similar and are not shown.

In Figure 6, the two curves and one point related to TCP delivery was obtained by assuming that each image was coded in 1 bpp and transmitted in 64 packets. Based on the statistics collected, we calculated the average arrival times of the first i , $i = 1, 2, \dots, 64$, packets and evaluated the quality of the corresponding packets after decoding them by the SPIHT decoder. The times in each curve include both end-to-end delays and decoding times.

The two points related to UDP delivery were obtained under 1 bpp and included end-to-end delays, decoding time, and reconstruction time when losses happened.

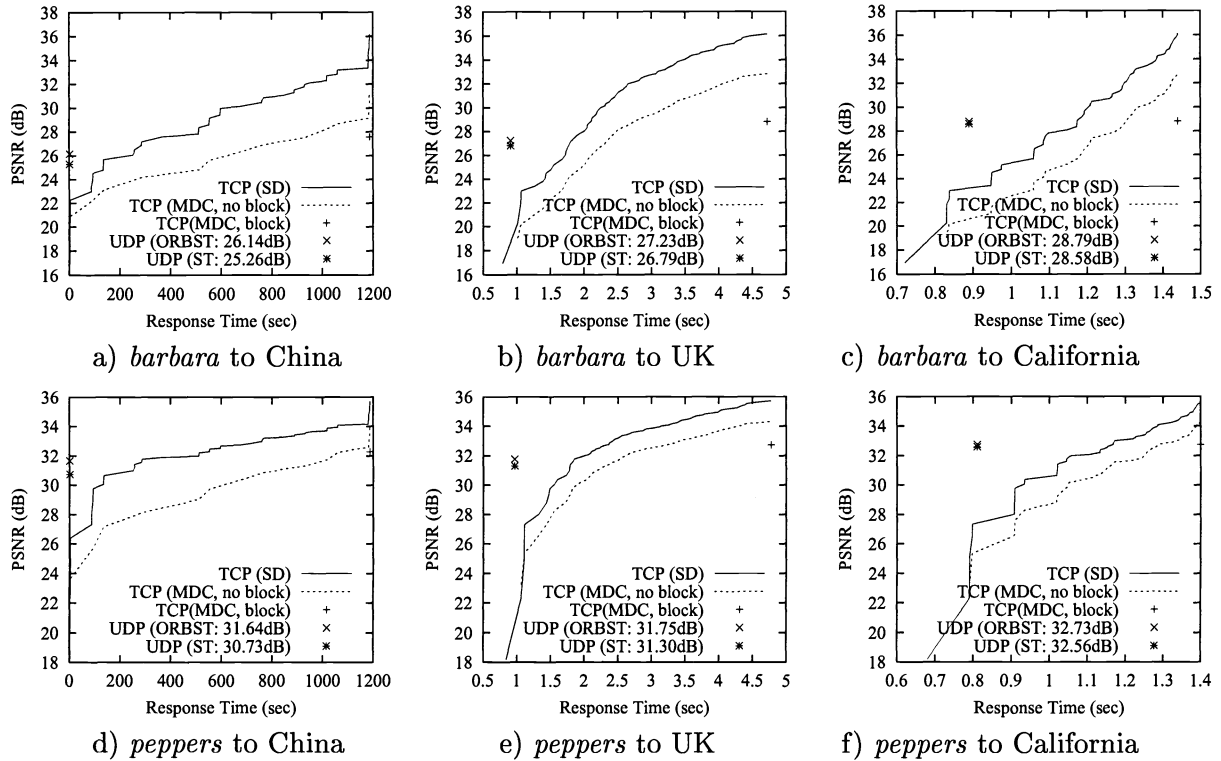


Figure 6. Delay-quality trade-offs between TCP delivery of SDC image data for the transmissions to three remote sites at 12 noon their local time. (The behavior at other times are similar and are not shown.)

The graphs show that the UDP delivery of MDC images is an attractive alternative to the TCP delivery of SDC images when the delay that an end user can tolerate is small and when absolute quality is not critical. The graphs show that, without exception, TCP delivery leads to poorer quality using the same amount of time required by UDP delivery.

The graphs also illustrate three factors that cause the degradation in quality by several dBs between the TCP delivery of SDC images and the UDP delivery of MDC images.

First, MDC alone causes between 1-3.5 dB loss in PSNR and is the price paid for improved error resilience. This is illustrated by the difference between the top two curves in each graph that shows the quality of TCP delivery of SDC images and that of MDC images. Such degradations happen because of reduced correlations when partitioning an image into multiple descriptions and the suboptimal fixed coding rate for each description.

Second, another 2-3.5 dB loss in PSNR is caused by the suboptimal strategies of using fixed-size segments in the segmentation of image data in each description and of using a fixed coding rate for each segment in order for the coded segment to fit in a 512-byte packet (the difference between the point on the right of the dotted line and the cross on the right of each graph).

Third, packet losses and reconstructions in the UDP delivery of segmented ST-MDC data lead to further degradations of up to 2 dB.

4.2. Combined TCP/UDP Delivery of images

The delay-quality trade-offs studied previously only show two extreme cases of image transmission, either by TCP or by UDP. By inspecting the trade-off graphs, we see a promising hybrid approach that can hopefully give better trade-offs. For TCP delivery, quality improves very quickly in the beginning but

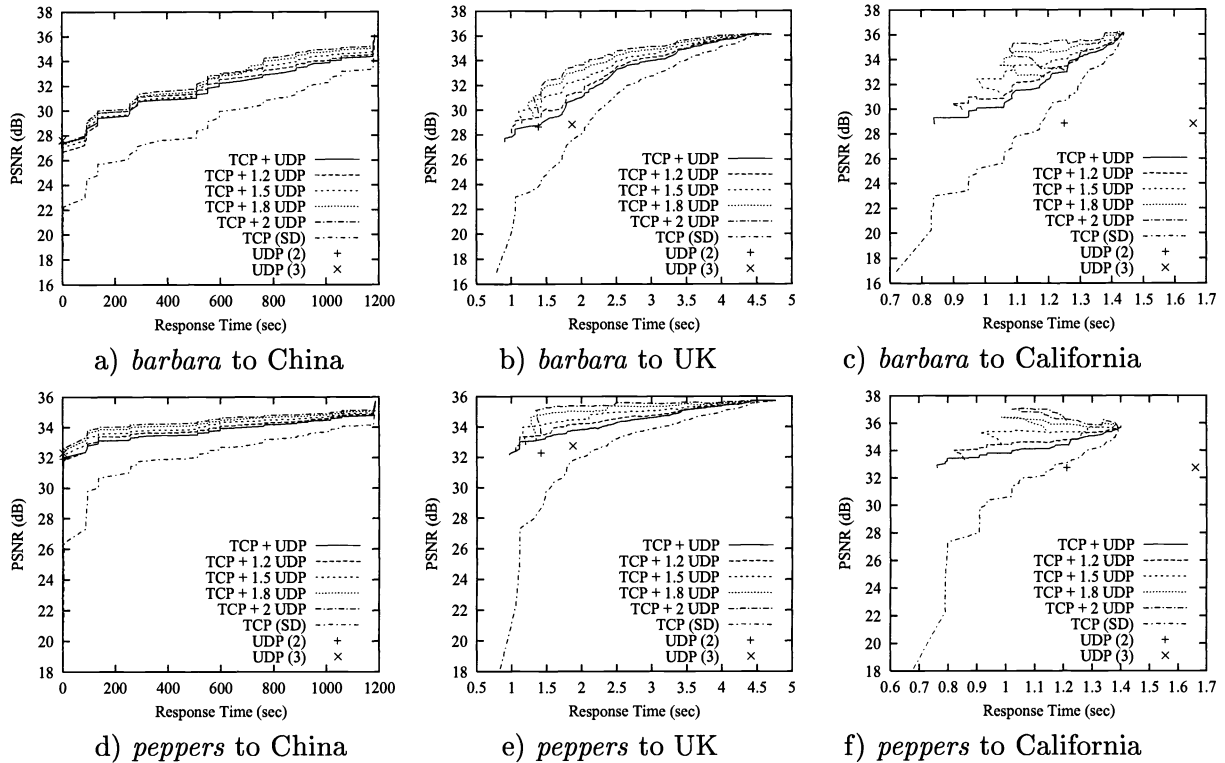


Figure 7. Delay-quality trade-offs of the combined TCP/UDP delivery of SDC/MDC data for *barbara* and *peppers* for the transmissions from Urnaba to three remote sites at 12 noon their local time.

saturates gradually when more packets are available. Since the first few packets delivered by TCP incur insignificant delays, we can transmit them by TCP and deliver the MDC residuals by UDP. In general, we can characterize this approach as follows:

$$\{x : \text{SDC by TCP}\} \cup \{\alpha(1 - x) : \text{MDC by UDP}\} \text{ where } 0 \leq x \leq 1 \text{ and } \alpha \geq 1.$$

In this approach, the first $x\%$ of the bit stream is coded by SDC and delivered by TCP, and the rest of the bit stream is coded either redundantly ($\alpha > 1$) or non-redundantly ($\alpha = 1$) by MDC and delivered by UDP.

The combined approach can, to some extent, reduce the three kinds of losses identified in Section 4.1, depending on the factor x chosen: since the first $x\%$ of the bit stream is coded in SDC and transmitted by TCP, it suffers from none of the above three kinds of losses. The larger the x is, the less degradation in quality one has to pay and the longer delay one has to wait. In fact, if both x and α are equal to one, this approach is reduced to pure TCP delivery of SDC images, leading to the best quality and the longest delay. Again, this approach involves delay-quality trade-offs.

To evaluate these trade-offs, we compared the following approaches for three chosen sites and two test images: a) TCP delivery of SDC images; b) combined TCP/UDP delivery of SDC and MDC images, with α set to 1, 1.2, 1.5, 1.8 and 2, respectively; and c) redundant UDP delivery of MDC images by sending y copies of UDP packets that contain MDC data, with y set to 2 and 3.

Figure 7 shows the delay-quality trade-offs using the above approaches to transmit *barbara* and *peppers* to China, UK, and California, respectively, at 12 noon their local time. (The behavior at other times are similar and are not shown.). We can see that the quality of the combined approach improves with increasing α , and that the combined approach can generate a range of trade-offs with better quality than

pure UDP delivery and with shorter delays than pure TCP delivery. Further, the redundant UDP delivery of MDC images does not appear to be an attractive approach due to its high redundancy, mediocre quality, and long delays.

Based on the above results, we conclude that the hybrid approach provides a range of delay-quality trade-offs between pure UDP delivery and pure TCP delivery. Based on such trade-offs, users can choose a suitable combination to suit their QoS requirements and available resources.

5. CONCLUSIONS AND FUTURE WORK

This paper studies delay-quality trade-offs in transferring subband-transformed (ST) images in the Internet. Our experiments reveal that delays using TCP to deliver an image are much longer than those using UDP, but that packet losses in UDP may lead to poor decoding quality if the image is single-description coded (SDC) and the losses cannot be concealed. To reduce the effects of packet losses, we propose to use multi-description coding (MDC) and determine experimentally the interleaving factors that should be used in order to keep the probability of unrecoverable losses sufficiently small. Next, we propose an optimized reconstruction-based subband transform (ORT-ST) that is designed to minimize distortions if some of the descriptions are lost, and the missing information is reconstructed using simple interpolation. In Internet transmission experiments, we carefully evaluated delay-quality trade-offs of several delivery strategies, and showed experimentally that ORB-ST is more suitable than ST for lossy transmissions. Our future work includes a study of better packetization strategies for UDP delivery and alternative coding and transmission approaches that can achieve better delay-quality trade-offs.

REFERENCES

1. I. Sodagar, H.-J. Lee, P. Hatrack, and Y.-Q. Zhang, "Scalable wavelet coding for synthetic/natural hybrid images," *IEEE Trans. on Circuits and Systems for Video Technology* **9**, pp. 244-254, Mar. 1999.
2. V. A. Vaishampayan, "Design of multiple description scalar quantizers," *IEEE Trans. on Information Theory* **39**, pp. 821-834, May 1993.
3. S. D. Servetto, K. Ramchandran, V. A. Vaishampayan, and K. Nahrstedt, "Multiple description wavelet based image coding," *IEEE Trans. on Image Processing* **9**, pp. 813-826, May 2000.
4. G. Yu, M. M. Liu, and M. W. Marcellin, "POCS-based error concealment for packet video using multiframe overlap information," *IEEE Trans. on Circuits and Systems for Video Technology* **8**, pp. 422-434, Aug. 1998.
5. J. Suh and Y. Ho, "Error concealment based on directional interpolation," *IEEE Trans. on Consumer Electronics* **43**, pp. 295-302, Aug. 1997.
6. W. Zeng and B. Liu, "Geometric-structure-based error concealment with novel applications in block-based low-bit-rate coding," *IEEE Trans. on Circuits and Systems for Video Technology* **9**, pp. 648-665, June 1999.
7. J. M. Danskin, G. M. Davis, and X. Song, "Fast lossy Internet image transmission," in *ACM Multimedia Conf.*, (San Francisco), Nov. 1995.
8. V. A. Vaishampayan and N. Farvardin, "Optimal block cosine transform image coding for noisy channels," *IEEE Trans. on Communications* **38**, pp. 327-336, Mar. 1990.
9. C. J. Turner and L. L. Peterson, "Image transfer and end-to-end design," in *Proc. SIGCOMM'92*, pp. 258-268, 1992.
10. X. Su and B. W. Wah, "Streaming real-time video data with optimized reconstruction-based DCT and neural-network reconstructions," *IEEE Trans. on Multimedia* **3**, pp. 123-131, Mar. 2001.
11. L. Rade and B. Westergren, *Mathematics Handbook for Science and Engineering*, Studentlitteratur Birkhauser, 1995.