



Reconstruction-Based Subband Image Coding for UDP Transmissions over the Internet

XIAO SU

Computer Engineering Department, San Jose State University, San Jose, CA 95192-0180

BENJAMIN W. WAH

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

Received May 2, 2001; Revised March 4, 2002; Accepted March 28, 2002

Abstract. Quality-delay trade-offs can be made in transmitting subband-coded images in the Internet by 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 primarily by TCP today, we study in this paper the use of UDP to deliver multi-description reconstruction-based subband-coded images and the reconstruction of missing information at the receiver based on information received. We first determine empirically the interleaving factors that should be used in order to keep the probability of unrecoverable packet losses sufficiently small. Next, we propose a joint sender-receiver approach for designing transforms in multi-description subband coding. In the receiver, we use a simple interpolation-based reconstruction algorithm, as sophisticated concealment techniques cannot be employed in practice. In the sender, we design an *optimized reconstruction-based subband transform* (ORB-ST), with an objective of minimizing the mean squared error, assuming that some of the descriptions are lost and that the missing information is reconstructed by simple averaging at the destination. Experimental results show that our proposed ORB-ST performs well in real Internet tests, and UDP delivery of MDC images is an attractive alternative to TCP delivery.

Keywords: error concealment, interpolation-based reconstruction, multi-description coding (MDC), real-time multimedia in the Internet, reconstruction-based subband image coding, single-description coding (SDC), TCP, UDP, world-wide web

1. Introduction

As subband image coding emerges as the core technology in the JPEG2000 standard [1], more images transmitted in the World Wide Web will be coded using this technique. The delivery of subband-coded images in the current Internet can be done by the slow and reliable TCP or by the fast and unreliable UDP.

Quality and delay are two key performance measures to evaluate the delivery of 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 by using coarse grained timeout periods

(500 ms) and exponential backoffs of sending rates when congestion happens. This is evident in Web surfing, as one frequently experiences long stalls in downloading images. On the other hand, UDP delivery incurs shorter end-to-end delays but cannot be used for sending coded images because dependencies may render these images non-decodable when losses happen.

To address the need 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 using UDP. To prevent the propagation of losses across packet boundaries, we also study multi-description coding that defines units for error concealment and reconstruction.

Existing schemes on error concealment are performed either entirely in the receiver side, or in both the sender and the receiver sides.

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 [2]. Other approaches minimize the variations along edge directions or local geometric structures [3–5]. 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. They are usually more effective because receivers can convey information on losses and reconstruction methods to senders in order for senders to better adapt its control. There are two popular ways to facilitate such recovery: *layered coding* and *multi-description coding* (MDC).

In *layered coding* [6], data is partitioned into a base layer and a few enhancement layers. The base layer contains visually important image data that can be used to produce outputs of acceptable quality, whereas the enhancement layers contain complementary information that allows higher-quality image data to be generated. 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 when network conditions worsen. Layered coding has been popular with ATM networks but may not be applicable in the Internet for two reasons. First, the current Internet does not provide priority delivery service for different layers, although ongoing efforts of defining the Diffserv model [7] may make it possible in the future. Second, when the packet-loss rate is high and part of the base

layer is lost, it is hard to reconstruct the lost bit stream since no redundancy is present.

In contrast, *multi-description coding* (MDC) divides image data into equally important streams in such a way that the decoding quality using any subset is acceptable, and that better quality is obtained by more descriptions. It is assumed in MDC that the probability of losing all the descriptions is small. For subband coded images, MDC was first implemented by designing scalar quantizers [8] in which two side-scalar quantizers were applied to produce two descriptions. In order to minimize reconstruction errors when both descriptions were received, it then mapped a proper subset of index pairs formed from side quantizers to central-quantizer intervals. The difficulties with this approach are that optimal index assignments are hard to achieve in real time, and that suboptimal approaches, such as A2 index assignment [8], introduce a large overhead in bit rate [9]. A recent approach along this line produces two descriptions by choosing one index assignment per subband and by encoding explicitly this choice as map bits [10]. However, the authors pointed out that the extra complexity involved in choosing the optimal index assignment for each subband was not worthwhile for the marginal improvement in image quality. Further, the complicated decoding algorithms make it infeasible for real-time delivery.

In short, existing error concealment techniques either rely on the inadequate capability of 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. First, to reduce the effects of packet losses, we experimentally derive the number of descriptions that are needed so that the probability of packet losses that are not recoverable is sufficiently small. Based on the number of descriptions, 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. We call packets that carry related descriptions an *interleaved set* and the number of descriptions, the *interleaving factor*. Moreover, each packet may carry more than one descriptions from different segments. For example, with an interleaving

factor of two, packet 0 may carry the first description of the first two segments, whereas packet 1 may carry the other description of the first two segments. We assume that packets in an interleaved set are transmitted sequentially, one after another.

Next, we design multi-description coders at senders using 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 interpolation from the descriptions received. We have adopted a simple reconstruction algorithm at receivers in order to facilitate fast playback. This approach leads 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 help guide the design of ORB-ST in Section 3. Section 4 describes our experimental results in the Internet to evaluate 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 Champaign (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 connections, 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 TCP and UDP echo port of each of the destinations from the site in Champaign. In determining the number of packets to be sent, we assume that a

512-by-512 image is compressed at 8:1 ratio and sent in 512-byte packets, leading to 64 packets sent.

From the packets echoed back, we recorded their sequence numbers and sending and arrival times, and determined packet losses based on the sequence numbers recorded. Finally, we determined the loss rate and the cumulative distribution function (CDF) of burst lengths.

Bouncing messages off the TCP and UDP echo ports of a remote server are meant to emulate a “hypothetical” TCP connection and a UDP path that are twice as long as the path from the source to the remote computer. However, their timing results are not comparable directly because TCP and UDP echo ports are implemented differently. An UDP echo port reflects every incoming packet immediately after it is received, whereas a packet sent to a TCP echo port traverses two virtual links, each of which employs window-based flow control to accommodate packet losses and buffer overflows. Hence, a TCP echo port has a receiver window and a dependent sender window, since the echo port can only send a packet after all its previous packets have been received. To avoid such reassembly delay at the TCP echo port, we would like each TCP packet to traverse a single virtual path as that of a UDP packet. To this end, we modified the Linux kernel in order to encapsulate TCP echo packets in UDP ones and sent them to the UDP echo port of the remote server. In this way, encapsulated TCP packets will be echoed immediately as they are received at the echo server.

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 high losses 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, in sending the same amount of data by TCP, we encapsulated the TCP packets in UDP ones and sent them to the same remote UDP echo ports. The pacing and retransmissions of TCP packets were controlled by the TCP protocol,

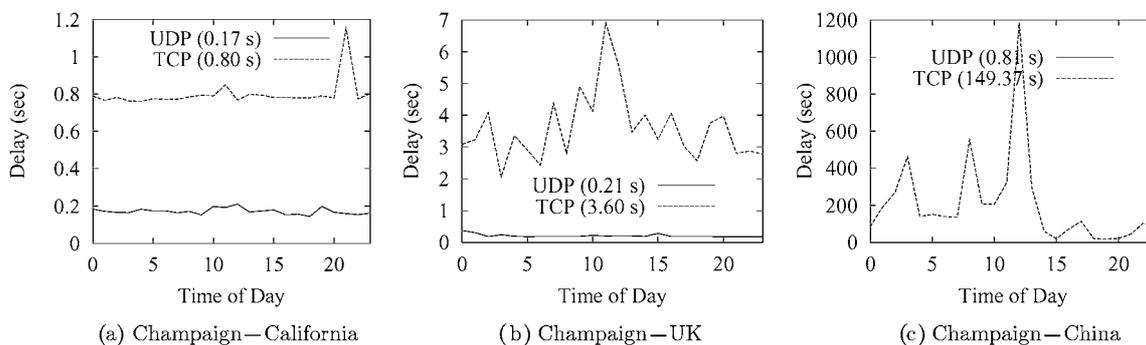


Figure 1. Round-trip delays of sending 64 UDP packets and the same data in TCP packets encapsulated in UDP ones to the UDP echo port of three remote computers. The experiments were carried out at the beginning of each hour for a 24-hour period on April 8, 2001.

based on round-trip-delay estimates between the source and the echo servers.

The graphs in Fig. 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 Champaign and UK (Fig. 1(b)), UDP delays range from 0.2 to 0.4 seconds, whereas TCP delays range from 2 to 7 seconds. In terms of speed, TCP transmissions took between 4 to 6 times longer than UDP transmissions for the California site, 7 to 30 times longer for the UK site, and 5 to 184 times longer for the China site. 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 one minute in order to download an image from China when using TCP.

2.3. Loss Behavior of UDP Transmissions

Although UDP delivery is much faster, they suffer from losses that may lead to large degradations in image

quality or render images non-decodable. In order to send images using UDP, we need to understand its loss behavior and conditions under which losses can be concealed.

Figure 2 depicts typical cumulative distribution functions (CDFs) of lengths of bursty UDP packet losses for connections to the three chosen sites measured at 10 pm their local time on April 8, 2001. (Due to space limitation, we do not show the CDFs at other times.) The graphs show that the number of consecutive UDP packet losses is usually very small. For the Champaign-California and Champaign-UK connections, all losses are of burst length 2 or less, and for the Champaign-China connection, losses are of burst length 5 or less.

The results imply the use of small interleaving factors to convert bursty losses to random losses. However, the CDF of burst lengths alone is not sufficient to determine the interleaving factor because a bursty loss larger than the interleaving factor may span two interleaved sets and can be recovered from partial information received in each interleaved set. In general, an

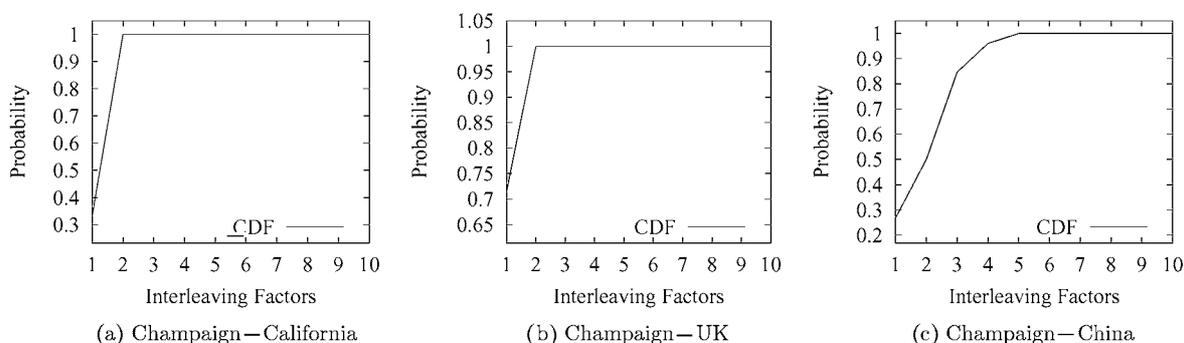


Figure 2. Cumulative distribution function (CDF) of consecutive packet losses in connections to the three chosen sites at 10 pm their local time on April 8, 2001.

interleaving factor i allows reconstructions by interpolation of a bursty loss of $i - 1$ packets or less if losses are from the same interleaved set, or of length in the range $[i, (2i - 2)]$ when losses are from different interleaved sets. 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 the total number of packets lost is n_s (independent of i), where:

$$n_s = \sum_{j=1}^i j \times m_j^i. \quad (1)$$

Given that all the packets in an interleaved set are lost, the conditional probability that the content of a packet cannot be recovered by reconstruction using interleaving factor i , can be derived from (1) as follows:

$$Pr(\text{fail} | \text{loss}, i) = \frac{i \times m_i^i}{n_s}. \quad (2)$$

$Pr(\text{fail} | i)$, the probability that a packet cannot be reconstructed in the stream received for interleaving factor i , can be computed as follows:

$$\begin{aligned} Pr(\text{fail} | i) &= Pr(\text{fail} | \text{loss}, i) \times Pr(\text{loss}) \\ &= \frac{i \times m_i^i}{n_s} \times \frac{n_s}{n_p} = \frac{i \times m_i^i}{n_p}. \end{aligned} \quad (3)$$

Figure 3 shows that $Pr(\text{fail} | \text{loss}, i)$ drops quickly with increasing interleaving factor i . For the Champaign-California connection, all its lost packets are recoverable using an interleaving factor of 2. For the Champaign-UK connection, the failure probability

can be upper bounded by 5% and is negligible most of the time using an interleaving factor of 2. For the Champaign-China connection, the loss rate can be as high as 50% for some part of the day, but the probability of not being able to reconstruct a lost packet is held under 5% and is zero most of the time using an interleaving factor of 4.

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 effectively by interleaving and reconstruction using an interleaving factor of four for most Internet destinations. Of course, it is expected that the reconstructed image will have lower quality. Such trade-offs between quality and delay may be acceptable since users, when downloading an image from the Web, may not want to wait for minutes to see a perfect image but may prefer to have a slightly degraded one in much shorter time.

In the next section we discuss coding methods to help recover lost information.

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 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 [11].

Figure 4 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

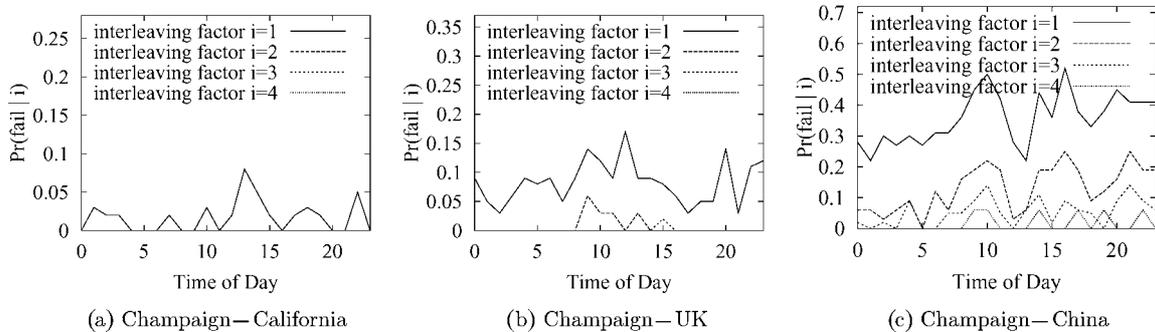


Figure 3. $Pr(\text{fail} | i)$, probability of bursty losses that cannot be recovered, conditioned on interleaving factor i , at different times on April 8, 2001.

several stages: the subband transformation stage that divides image data into components with different frequency contents, the quantizer that causes a controlled loss of information based on frequency information, and an optional entropy coder that removes residual redundancies among quantized symbols.

In a subband transform system, one can represent filtering operations as equivalent linear transformations in the spatial domain. Here, we use \hat{H} to denote $H_0(z)$ and $H_1(z)$ with down-sampling, and similarly \hat{G} for $G_0(z)$ and $G_1(z)$ with up-sampling. Our goal is to find a new analysis system \hat{H}' ($H'_0(z)$ and $H'_1(z)$ with down-sampling) in order to minimize reconstruction error \mathcal{E}_r after average interpolation, based on fixed quantization Q , inverse quantization IQ , and the original synthesis system, \hat{G} ($G_0(z)$ and $G_1(z)$ with up-sampling), where:

$$\mathcal{E}_r = \left\| \underbrace{\text{Interpolate}(\hat{G}(IQ(\vec{c})))}_{\text{decompression+reconstruction}} - \vec{x} \right\|^2. \quad (4)$$

In order to keep our decoders standard-compliant so that existing decoders at receivers can be used, we assume fixed inverse quantization IQ and synthesis system \hat{G} .

$$\hat{G} = \begin{pmatrix} \vdots & \vdots \\ \cdots & g_0(-2) & g_1(-2) & g_0(0) & g_1(0) & g_0(2) & g_1(2) & g_0(4) & g_1(4) & \cdots \\ \cdots & g_0(-3) & g_1(-3) & g_0(-1) & g_1(-1) & g_0(1) & g_1(1) & g_0(3) & g_1(3) & \cdots \\ \cdots & g_0(-4) & g_1(-4) & g_0(-2) & g_1(-2) & g_0(0) & g_1(0) & g_0(2) & g_1(2) & \cdots \\ \cdots & g_0(-5) & g_1(-5) & g_0(-3) & g_1(-3) & g_0(-1) & g_1(-1) & g_0(1) & g_1(1) & \cdots \\ \vdots & \vdots \end{pmatrix} \\ = \begin{pmatrix} \vdots & \vdots & \vdots & \vdots & \vdots \\ \cdots & S_{-1} & S_0 & S_1 & S_2 & \cdots \\ \cdots & S_{-2} & S_{-1} & S_0 & S_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix}, \quad \text{where } S_i = \begin{pmatrix} g_0(2i) & g_1(2i) \\ g_0(2i-1) & g_1(2i-1) \end{pmatrix}. \quad (6)$$

With quantization in place, the minimization of \mathcal{E}_r becomes an integer optimization problem, where \vec{c} in (4) 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. Specifically, the objective to be optimized in the approximation is:

$$\mathcal{E}_r = \left\| \text{Interpolate}(\hat{G}(\vec{c})) - \vec{x} \right\|^2. \quad (5)$$

The resulting transform is called optimized reconstruction-based subband transform (ORB-ST). We derive in the following ORB-ST based on partitioning image data into two descriptions and in Section 3.2, extensions to four descriptions.

3.1. ORB-ST for Two Descriptions

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 . Since the derivations are similar, we only show that for \vec{c}_1 .

As explained above, the synthesis system consisting of up-sampling, $G_0(z)$ and $G_1(z)$ is equivalent to a linear transform, \hat{G} , in the spatial domain. Let $g_0(i)$ ($i = -N_1, -N_1 + 1, \dots, M_1 - 1, M_1$) and $g_1(j)$ ($j = -N_2, -N_2 + 1, \dots, M_2 - 1, M_2$) denote the filter coefficients of G_0 and G_1 , respectively. Then \hat{G} can be written in the following form:

Then after synthesis filtering, output \vec{y}_1 is calculated as:

$$\vec{y}_1 = \hat{G} \vec{c}_1. \quad (7)$$

The set of interpolated pixels, \vec{z}_1 , is obtained by inserting even-numbered columns as the average of columns from \vec{y}_1 , with the boundary column

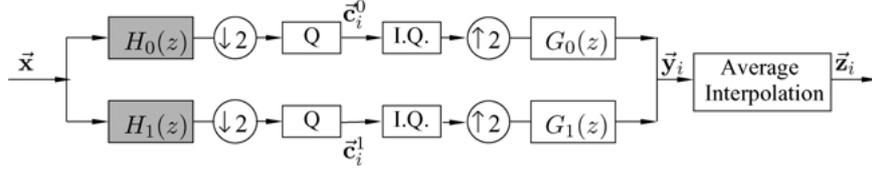


Figure 4. Basic building blocks of a modified codec. (The shaded block is our proposed ORB-ST.)

duplicated:

$$\begin{aligned} \bar{\mathbf{z}}_1 &= \begin{pmatrix} \vdots & \vdots & \vdots & \vdots & \vdots \\ \cdots & 1 & 0 & 0 & 0 & \cdots \\ \cdots & 0.5 & 0.5 & 0 & 0 & \cdots \\ \cdots & 0 & 1 & 0 & 0 & \cdots \\ \cdots & 0 & 0.5 & 0.5 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix} \bar{\mathbf{y}}_1 \\ &= \begin{pmatrix} \vdots & \vdots \\ \cdots & B_0 & B_1 & \cdots \\ \cdots & & B_0 & B_1 & \cdots \\ \cdots & & & B_0 & B_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix} \bar{\mathbf{y}}_1 = \mathbf{U} \bar{\mathbf{y}}_1, \end{aligned} \quad (8)$$

where

$$B_0 = \begin{pmatrix} 1 & 0 \\ 0.5 & 0.5 \\ 0 & 1 \\ 0 & 0.5 \end{pmatrix} \quad \text{and} \quad B_1 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0.5 & 0 \end{pmatrix}.$$

The distortion between the original and the received and reconstructed pixels becomes:

$$\mathcal{E}_r = \|\mathbf{U}\hat{\mathbf{G}}\bar{\mathbf{c}}_1 - \bar{\mathbf{x}}\|^2 = \|\mathbf{P}\bar{\mathbf{c}}_1 - \bar{\mathbf{x}}\|^2. \quad (9)$$

Since the linear system of equations $\mathbf{P}\bar{\mathbf{c}}_1 = \bar{\mathbf{x}}$ is an over-determined one, there exists at least one least-square solution $\bar{\mathbf{c}}_1$ that minimizes (9), according to the theory of linear algebra [12]. Specifically, the solution $\bar{\mathbf{c}}_1$ with the smallest length $|\bar{\mathbf{c}}_1|^2$ can be found by first performing SVD decomposition of matrix \mathbf{P} :

$$\mathbf{P} = \mathbf{S}[\text{diag}(w_j)]\mathbf{D}^T, \quad j = 1, 2, \dots, \frac{n}{2}, \quad (10)$$

where \mathbf{S} is an $n \times \frac{n}{2}$ column-orthogonal matrix, $[\text{diag}(w_j)]$, 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:

$$\bar{\mathbf{c}}_1 = \mathbf{D}[\text{diag}(1/w_j)] \mathbf{S}^T \bar{\mathbf{x}}. \quad (11)$$

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

To derive the ORB-ST transform for $\bar{\mathbf{c}}_2$, simply replace B_0 and B_1 in (8) by:

$$B_0 = \begin{pmatrix} 0 & 0.5 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix} \quad B_1 = \begin{pmatrix} 0.5 & 0 \\ 1 & 0 \\ 0.5 & 0.5 \\ 0 & 1 \end{pmatrix}.$$

The rest of the steps are similar.

In our derived ORB-ST, \mathbf{P} in (9) is a product of two matrices, \mathbf{U} and $\hat{\mathbf{G}}$, both of which are block circulant matrices, as evident in (6) and (8). Hence, \mathbf{P} is also block circulant. It is easy to verify that the least-square solution of a block circulant matrix is still block circulant [13]. This implies that our derived ORB-ST still takes the form of a linear filter.

When both $\bar{\mathbf{c}}_1$ and $\bar{\mathbf{c}}_2$ are available at the receiver, we can derive an inverse transform to achieve perfect reconstruction. Let \mathbf{T}_1 and \mathbf{T}_2 represent ORB-ST for $\bar{\mathbf{c}}_1$ and $\bar{\mathbf{c}}_2$, respectively.

$$\begin{aligned} \mathbf{T}_1 &= (\bar{\mathbf{a}}_1 \quad \bar{\mathbf{a}}_2 \quad \dots \quad \bar{\mathbf{a}}_{\frac{n}{2}})^T \\ \mathbf{T}_2 &= (\bar{\mathbf{b}}_1 \quad \bar{\mathbf{b}}_2 \quad \dots \quad \bar{\mathbf{b}}_{\frac{n}{2}})^T, \end{aligned}$$

where $\bar{\mathbf{a}}_i$ and $\bar{\mathbf{b}}_i$ are the transpose of row vectors in each transform. First, we interleave row vectors of \mathbf{T}_1 and \mathbf{T}_2 to give a combined transform, \mathbf{T} , as follows.

$$\mathbf{T} = (\bar{\mathbf{a}}_1 \quad \bar{\mathbf{b}}_1 \quad \bar{\mathbf{a}}_2 \quad \bar{\mathbf{b}}_2 \quad \dots \quad \bar{\mathbf{a}}_{\frac{n}{2}} \quad \bar{\mathbf{b}}_{\frac{n}{2}})^T.$$

The inverse of T is the transform that will be applied at the receiver when both descriptions are available. In practice, perfect reconstruction is not always achievable due to errors caused by truncations of floating point numbers as well as quantization.

3.2. Handling Burst Lengths of Four

As described in Section 2, bursty losses of length greater than two are likely for transcontinental connections. In this section we describe two ways to handle cases with a maximum burst length of four. We do not describe methods to handle longer burst lengths because such cases are infrequent.

We can partition image data into four descriptions by interleaving the original frame \vec{z} in the horizontal direction into two streams, \vec{z}_{h_1} and \vec{z}_{h_2} , and then by interleaving and transformations in the vertical direction to get two additional descriptions. In a different way, we can also get four descriptions by first partitioning in the vertical direction and then in the horizontal direction. The four descriptions, \vec{z}_{h_1, v_1} , \vec{z}_{h_1, v_2} , \vec{z}_{h_2, v_1} and \vec{z}_{h_2, v_2} , are then sent in distinct packets to the receiver.

First, assume that only one out of three interleaved descriptions, say Description 1 (\vec{z}_{h_1, v_1}), is received. The remaining three descriptions can be restored as follows:

$$\hat{z}_{i,j} = \begin{cases} \frac{(z_{i,j-1} + z_{i,j+1})}{2} & \hat{z}_{i,j} \in \vec{z}_{h_2, v_1} \\ \frac{(z_{i-1,j} + z_{i+1,j})}{2} & \hat{z}_{i,j} \in \vec{z}_{h_1, v_2} \\ \frac{(z_{i-1,j-1} + z_{i-1,j+1} + z_{i+1,j-1} + z_{i+1,j+1})}{4} & \hat{z}_{i,j} \in \vec{z}_{h_2, v_2}, \end{cases} \quad (12)$$

where $z_{i,j}$ is the value of the pixel in row i and column j . The transformed values of Description 1 in order to achieve the optimal reconstruction in (12) can be derived as outlined in Section 3.1. In a similar way, we need to derive transformations when zero, two, or three descriptions are lost. Since it is impossible to know the specific loss pattern for an interleaved set until it is received at the receiver and it will be either overly optimistic or overly pessimistic if one loss pattern is selected a priori, the method is impractical for use in the Internet.

Second, we can carry out the following operations based on the inverse flow of the interleaving process in order to reconstruct any missing descriptions.

- (a) If one out of the four interleaved descriptions is received, say \vec{z}_{h_1, v_1} , then \vec{z}_{h_1, v_2} can be reconstructed optimally by taking averages along the vertical direction of pixels from \vec{z}_{h_1, v_1} . By taking averages along the horizontal direction, \vec{z}_{h_2, v_1} and \vec{z}_{h_2, v_2} can then be recovered.
- (b) If two out of the four interleaved descriptions are received, then there are two possible scenarios. If the lost descriptions are from the same horizontally interleaved group, say \vec{z}_{h_1, v_1} and \vec{z}_{h_1, v_2} , then they can be reconstructed by averaging of their horizontal neighbors. If the lost descriptions do not belong to the same horizontally interleaved description, say \vec{z}_{h_1, v_1} and \vec{z}_{h_2, v_2} , then they can be reconstructed optimally by taking averages of their respective vertical neighbors.
- (c) If three out of the four descriptions are received, then the lost description can be reconstructed by taking averages along the vertical direction.

In short, the second method can be generalized easily to 2^m -way interleaving, $m > 0$. It is flexible because the transformation at the sender does not depend on the loss pattern at the receiver. For this reason, we have adopted this approach in our experiments.

4. Experimental Results

In this section, we first compare the performance of ORB-ST and the original subband transform (ST) in two scenarios: a synthetic scenario under controlled losses and real Internet tests. We then study trade-off between delay and quality on UDP delivery of MDC streams and TCP delivery of SDC streams.

We carried out our experiments using four test images: *barbara*, *goldhill*, *peppers*, and *lena*, and evaluated the reconstruction quality by the *peak signal-to-noise ratio* (PSNR):

$$PSNR = 10 \log \frac{255^2}{\sum_i (x_i - \hat{y}_i)^2}, \quad (13)$$

where x_i and \hat{y}_i are, respectively, the original and the reconstructed pixel values.

4.1. Reconstruction Quality under Controlled Losses

In this section, we study the reconstruction quality under controlled-loss scenarios for ORB-ST and the

Table 1. Reconstruction quality of frames in PSNR (dB) when transformed by ORB-ST and ST along the horizontal direction and only one of the descriptions is received under two-descriptions coding. Gain is defined as the difference in PSNR between ORB-ST and ST, and boxed numbers represent positive gains.

Image	Quant. effects	Bit rate	Odd received			Even received			Both received		
			ST	ORB-ST	Gain	ST	ORB-ST	Gain	ST	ORB-ST	Gain
<i>barbara</i>	No	–	25.21	26.67	1.46	25.16	26.62	1.46	perfect reconstruction		
<i>goldhill</i>		–	32.58	34.05	1.47	32.64	34.12	1.48	perfect reconstruction		
<i>peppers</i>		–	31.60	33.24	1.64	31.23	32.69	1.46	perfect reconstruction		
<i>lena</i>		–	34.11	35.83	1.72	34.19	35.93	1.74	perfect reconstruction		
<i>barbara</i>	Yes	0.25	23.41	24.25	0.84	23.38	24.24	0.86	25.78	25.67	–0.11
		0.50	24.35	25.31	0.96	24.30	25.26	0.96	27.94	27.82	–0.12
		1.0	24.92	25.98	1.06	24.87	25.93	1.06	32.88	32.80	–0.08
<i>goldhill</i>	Yes	0.25	27.88	28.33	0.45	27.90	28.34	0.44	28.70	28.64	–0.06
		0.50	29.55	30.21	0.66	29.54	30.22	0.66	31.33	31.17	–0.16
		1.0	31.09	32.14	1.05	31.09	32.16	1.07	34.89	34.65	–0.24
<i>peppers</i>	Yes	0.25	29.61	29.92	0.31	29.39	29.62	0.23	30.50	30.51	0.01
		0.50	30.53	31.23	0.70	30.22	30.89	0.67	31.67	31.68	0.01
		1.0	31.06	32.17	1.11	30.70	31.73	1.03	34.49	34.30	–0.19
<i>lena</i>	Yes	0.25	29.91	30.63	0.72	29.97	30.73	0.76	30.91	30.99	0.08
		0.50	32.12	33.18	1.06	32.19	33.27	1.08	34.59	34.44	–0.15
		1.0	33.33	34.70	1.37	33.41	34.79	1.38	37.71	37.50	–0.21

original ST. To isolate the effects due to ST and ORB-ST, we first eliminate quantization losses by removing quantization and dequantization in the coding process.

The first four rows of Table 1 compare the reconstruction quality of frames transformed by ST and by ORB-ST, assuming that image data is divided into two descriptions along horizontal directions, and that quantization effects are ignored. Since data in one stream is received without loss, it is not further segmented to create additional synchronization points in the stream. Results along the vertical direction are similar and are not shown. When either the odd-numbered or even-numbered description is received, the ORB-ST transformed frames have consistently better quality (1.46–1.74 dB or 81–85% of the reconstruction error) than that of the ST transformed frames. When both descriptions are available, we omit the results in the table because perfect reconstruction is achieved and the PSNR values for both ST and ORB-ST are infinite.

Similarly, the top four rows of Table 2 present the results of dividing image data into four descriptions

by recursive 2-way interleaving without quantization. It shows that ORB-ST transformed frames have better quality in all cases except in Case IV. However, Case IV corresponds to losses of burst length one and should be infrequent when using four descriptions.

Next, we show results on reconstruction quality after including quantization effects.

The image codec used is an implementation based on SPIHT [14] that was downloaded from <http://qccpack.sourceforge.net>. SPIHT is a state-of-the-art subband-based image codec that generates an embedded bit stream. Using an embedded codec, its encoder can stop the encoding process at any point, thus allowing a target bit rate to be met exactly. Its high performance is due to the exploitation of self-similarity within wavelet coefficients among different subbands, and the ordering of transform coefficients by magnitude and ordered bit-plane transmissions, both resulting in bit streams ordered by importance. Throughout the experiments, we have used the Daubechies 9/7 filters [15] in ST and the ORB-ST derived from them.

Table 2. Reconstruction quality in PSNR (dB) when image data is divided into four descriptions by recursive 2-way interleaving. Case I represents the case in which three out of the four interleaved descriptions were lost; II, the case in which two descriptions, both from the same horizontal group, were lost; III, the case in which two descriptions, each from a different horizontal group, were lost; IV, the case in which one out of the four interleaved descriptions was lost; and V, the case in which none of the four descriptions was lost. Boxed numbers represent positive gains.

Image	Bit rate	Case I			Case II			Case III			Case IV			Case V		
		ST	ORB-ST	Gain	ST	ORB-ST	Gain	ST	ORB-ST	Gain	ST	ORB-ST	Gain	ST	ORB-ST	Gain
<i>barbara</i>	-	25.05	25.82	0.77	25.21	26.67	1.46	27.12	27.93	0.81	29.66	29.31	-0.35	24.34	24.15	-0.19
<i>goldhill</i>	-	30.01	31.33	1.32	32.58	34.05	1.47	31.90	32.45	0.55	34.34	35.41	1.07	26.89	26.91	0.02
<i>peppers</i>	-	30.07	31.27	1.20	31.60	33.24	1.64	30.83	31.92	1.09	33.04	34.10	1.06	31.31	31.29	-0.02
<i>lena</i>	-	32.86	34.25	1.39	34.10	35.83	1.73	36.12	37.43	1.31	36.23	36.41	0.18	28.02	27.96	-0.06
<i>barbara</i>	0.25	22.70	23.39	0.69	22.66	23.45	0.79	23.34	23.66	0.32	23.95	23.78	-0.17	28.81	28.81	0.00
	0.50	23.61	24.46	0.85	23.57	24.62	1.05	24.66	25.32	0.66	25.90	25.71	-0.19	30.93	30.85	-0.08
	1.0	24.48	25.34	0.86	24.52	25.64	1.12	26.06	26.84	0.78	28.01	27.45	-0.56	34.14	33.96	-0.18
<i>goldhill</i>	0.25	26.23	26.79	0.56	26.55	26.91	0.36	26.60	26.91	0.31	26.75	26.97	0.22	28.41	28.60	0.19
	0.50	27.63	28.42	0.79	28.31	28.89	0.58	28.45	28.68	0.23	28.84	29.04	0.20	32.03	31.94	-0.09
	1.0	28.86	29.92	1.06	30.22	31.14	0.92	30.14	30.46	0.32	31.15	31.54	0.39	35.73	35.42	-0.31
<i>peppers</i>	0.25	27.70	28.20	0.50	28.08	28.53	0.45	28.12	28.42	0.30	28.44	28.65	0.21	28.02	28.78	-0.03
	0.50	29.06	29.71	0.65	29.71	30.36	0.65	29.63	30.05	0.42	30.31	30.61	0.30	30.93	30.85	-0.08
	1.0	29.64	30.59	0.95	30.67	31.78	1.11	30.27	31.02	0.75	31.57	32.19	0.62	34.14	33.96	-0.18
<i>lena</i>	0.25	27.94	28.46	0.52	28.02	28.51	0.49	28.22	28.55	0.33	28.33	28.58	0.25	28.41	28.60	0.19
	0.50	30.24	31.05	0.81	30.51	31.29	0.78	31.29	31.55	0.26	31.58	31.61	0.03	32.03	31.94	-0.09
	1.0	31.98	33.10	1.12	32.67	33.86	1.19	33.96	34.58	0.62	34.32	34.38	0.06	35.73	35.42	-0.31

The bottom parts of Tables 1 and 2 show the reconstruction quality for two-description and four-description coding after incorporating quantization in the coding process. The three bit rates tested correspond to transmitting 16, 32 and 64 512-byte packets. When some descriptions are lost, the quality of reconstructed frames transformed by ORB-ST is better than that by ST for all cases under 2-description coding, and for most cases under 4-description coding, with a few exceptions for *barbara* and *goldhill*. However, when all the descriptions are received, the quality of ORB-ST transformed frames is, in general, not as good as that of original ST transformed ones. As explained before, cases with no loss and those in which three out of the four descriptions were received should occur infrequently when MDC was used. Moreover, since degradations are not as high as gains when losses happen, we should expect an improvement in average quality.

The tables show less gain for cases with quantization when compared to those without quantization. As pointed out in Section 3, these degradations were caused by the lossy quantization process in which it made certain changes to the transformed pixels that were not invertible. Although an inverse process exists for ORB-ST, quantization errors in the coded bit stream make it hard to achieve perfect reconstruction in practice. The tables also show improved gains in quality with increasing bit rates.

4.2. Tests in the Internet

To further evaluate our proposed schemes, we built a prototype (see Fig. 5) 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).

Our trace-driven simulations involved a sender process and a receiver process. The sender process was responsible for coding an image, packetizing it, and mapping packet losses to the losses of coded descriptions. The receiver process was in charge of decompressing coded streams, deinterleaving them, and performing reconstruction by linear interpolation. When an entire interleaved set is lost, it is filled by the average of image pixels.

An important consideration that is different from the experiments carried out under controlled losses is in the packetization of images. In our coding schemes, each of the coded descriptions is too large to fit into a single packet and must be decomposed into distinct packets. If we simply divide a single coded description into multiple packets, then the loss of a packet can render subsequent packets useless because the decoder has a mismatch in its control paths when decoding significance maps. Further, there is no synchronization units, such as GOBs of H.263, in a coded image. Hence, we need to decompose heuristically an image into segments in such a way that a coded segment can fit in a packet, that each description from the coded segment can be decoded independent of other coded segments, and that the total bit rate is unchanged from that of SDC. As a simple solution, we divided an image into equal-size segments in such a way that each coded

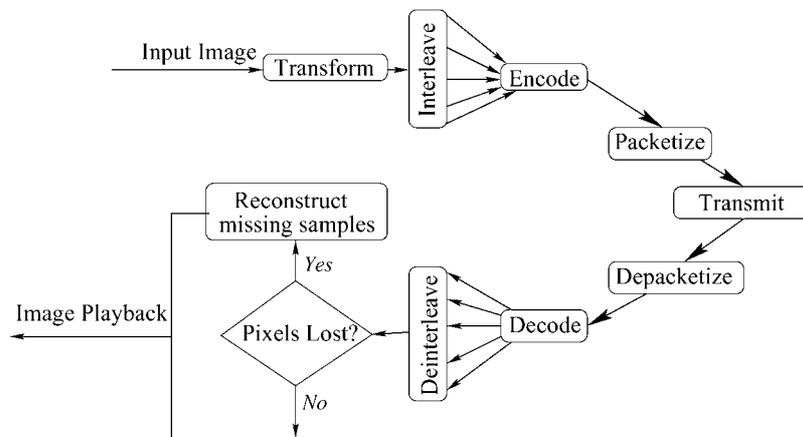


Figure 5. Components of our image-transmission prototype.

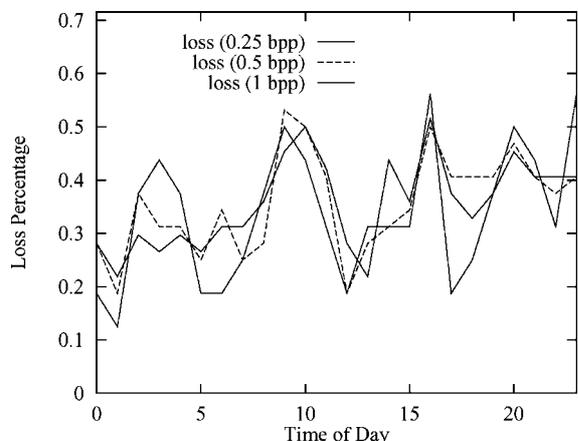


Figure 6. Loss rates of 16-, 32- and 64-packet transmissions between Champaign and China.

segment would fit in a single packet, and that each segment was coded using the same bit rate. Since our strategies of using equal-size segments and the coding of each using the same bit rate are suboptimal, we expect losses in image quality when compared to MDC without segmentation.

Figure 6 (resp. Figs. 8 and 10) 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 China (resp. UK and California). Figure 7 (resp. Figs. 9 and 11) compares the reconstruction quality of sending four test images 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 Champaign-China connection, ORB-ST outperforms ST at all bit rates for all four images, with

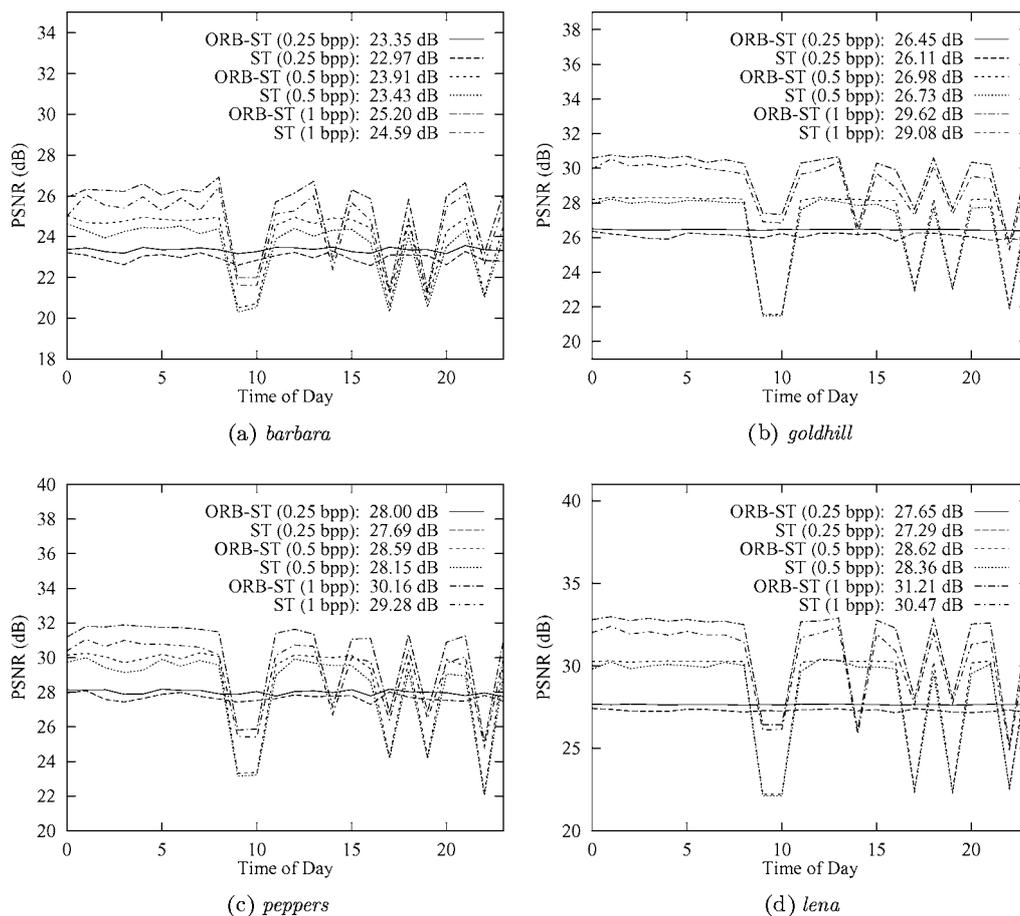


Figure 7. Comparisons of reconstruction quality over a 24-hour period for the Champaign to China connection, when each image was coded at respectively, 0.25 bpp, 0.5 bpp and 1 bpp, and placed into 16, 32 and 64 packets for transmission.

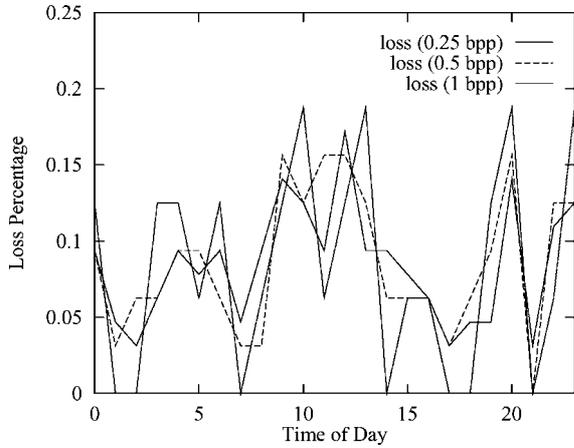


Figure 8. Loss rates of 16-, 32- and 64-packet transmissions between Champaign and UK.

an average of 0.31 to 0.38 dB better for the 0.25-bpp case, 0.25 to 0.48 dB better for the 0.5-bpp case, and 0.54 to 0.86 dB better for the 1-bpp case. Quality gain improves with increasing bit rates when there is less quantization noise. When entire interleaved sets were lost at certain hours, quality degraded significantly (such as hours 9, 11, 17 and 19 at 1 bpp).

For the Champaign-UK connection, the average reconstruction quality based on ORB-ST is better than that of ST most of the time, with only two exceptions: *goldhill* and *lena* at 0.5 bpp. For the Champaign-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 Champaign-China connection because the gain of performing ORB-ST is offset by degradations

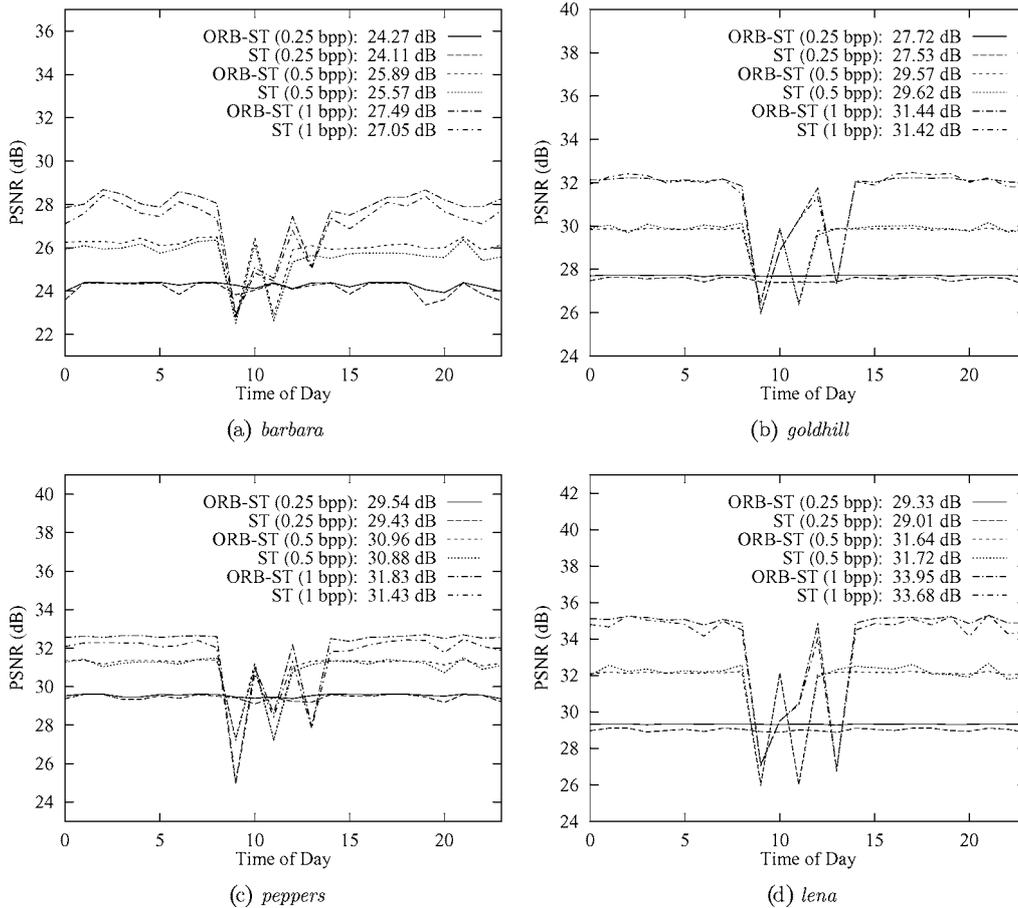


Figure 9. Comparisons of reconstruction quality over a 24-hour period for the Champaign to UK connection, when each image was coded at, respectively, 0.25 bpp, 0.5 bpp and 1 bpp, and placed into 16, 32 and 64 packets for transmission.

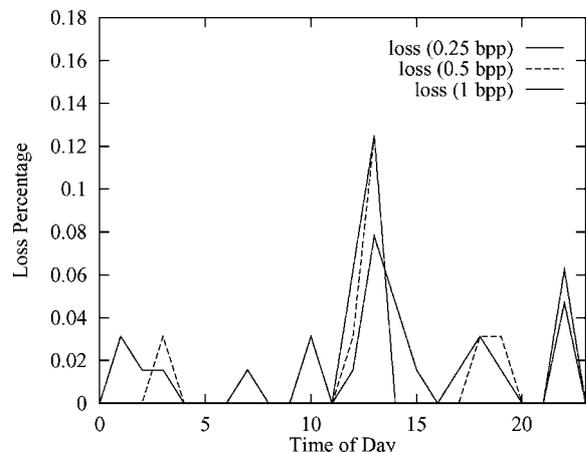


Figure 10. Loss rates of 16-, 32- and 64-packet transmissions between Champaign and California.

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.

4.3. Quality-Delay Trade-offs between TCP and UDP Delivery of Images

As described in Section 2, we have reduced the probability of unrecoverable UDP packet losses to less than 5% by properly choosing interleaving factors according to network conditions. Yet we still have degradations in quality when compared to the TCP delivery of SDC data due to the reconstruction process and MDC. We examine in detail such degradations in this section.

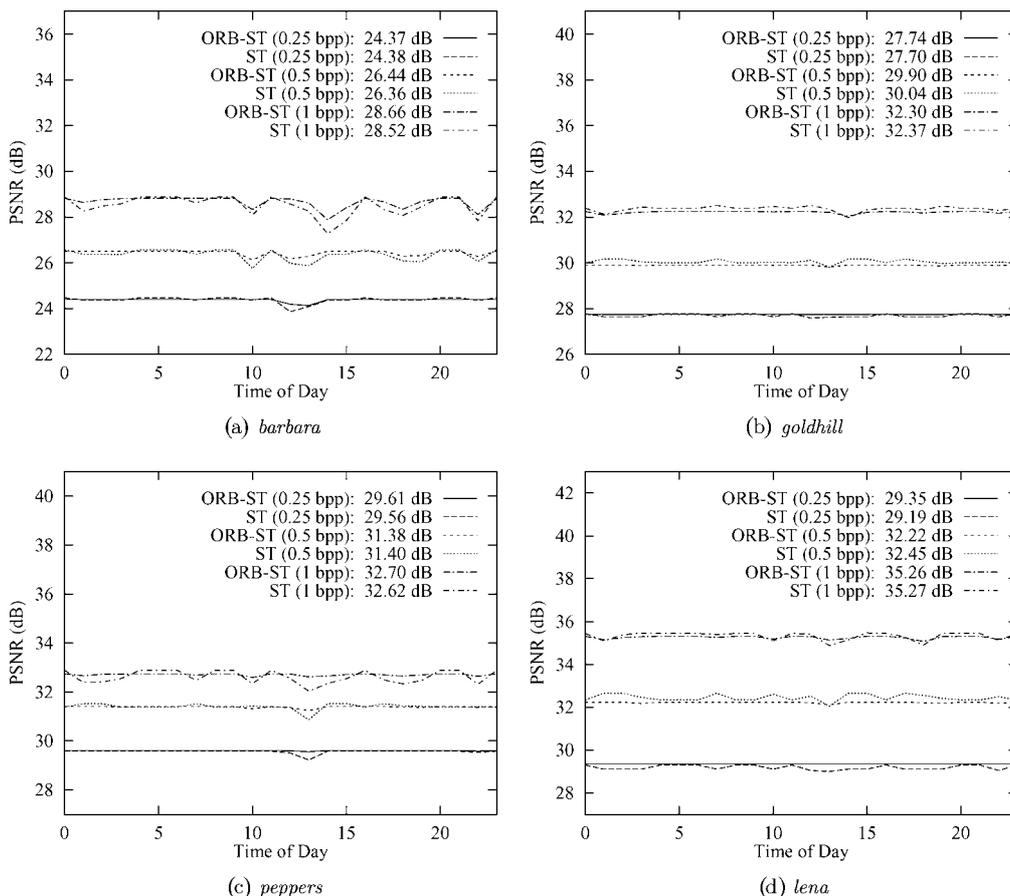


Figure 11. Comparisons of reconstruction quality over a 24-hour period for the Champaign to California connection, when each image was coded at, respectively, 0.25 bpp, 0.5 bpp and 1 bpp, and placed into 16, 32 and 64 packets for transmission.

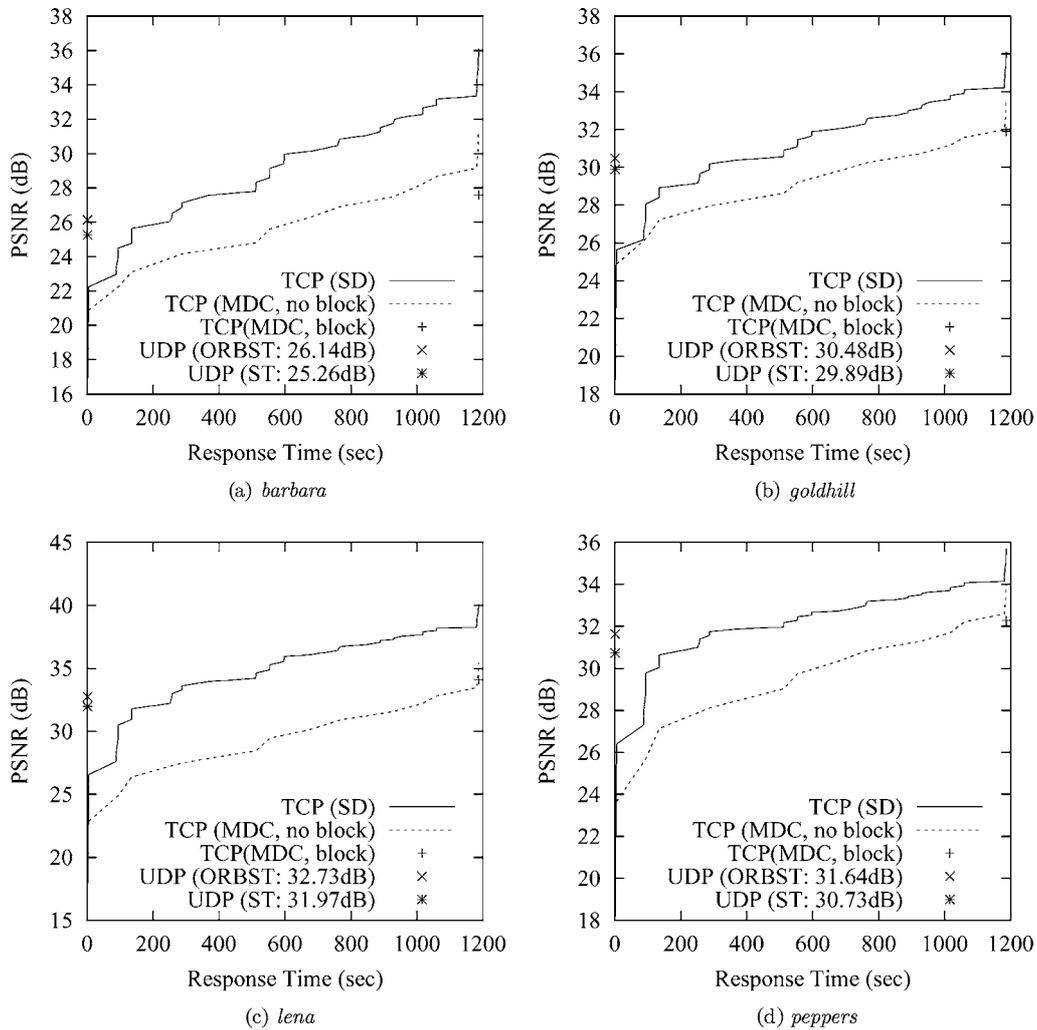


Figure 12. Quality-time trade-offs between TCP delivery of SDC image data and UDP delivery of MDC data for the Champaign-China connection at 12 noon Beijing Standard Time. (The behavior at other times are similar and are not shown.)

Figures 12–14 show the trade-offs between delay and quality at 12 noon local time of the remote server using five modes of delivery: TCP delivery of SDC image data, TCP delivery of MDC data in which the image is not segmented, TCP delivery of MDC data in which the image is segmented, UDP delivery of MDC ST-coded and segmented image data, and UDP delivery of MDC ORB-ST-coded and segmented image data. Results at other times are similar and are not shown.

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 each of the first 64 packets and then 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. Since packets may arrive out of order in UDP delivery and the algorithm needs to wait for all packets to arrive before decoding, it is not possible to generate the sequence of quality-delay points as in TCP.

The graphs show that the UDP delivery of MDC images is an attractive alternative to the TCP delivery of

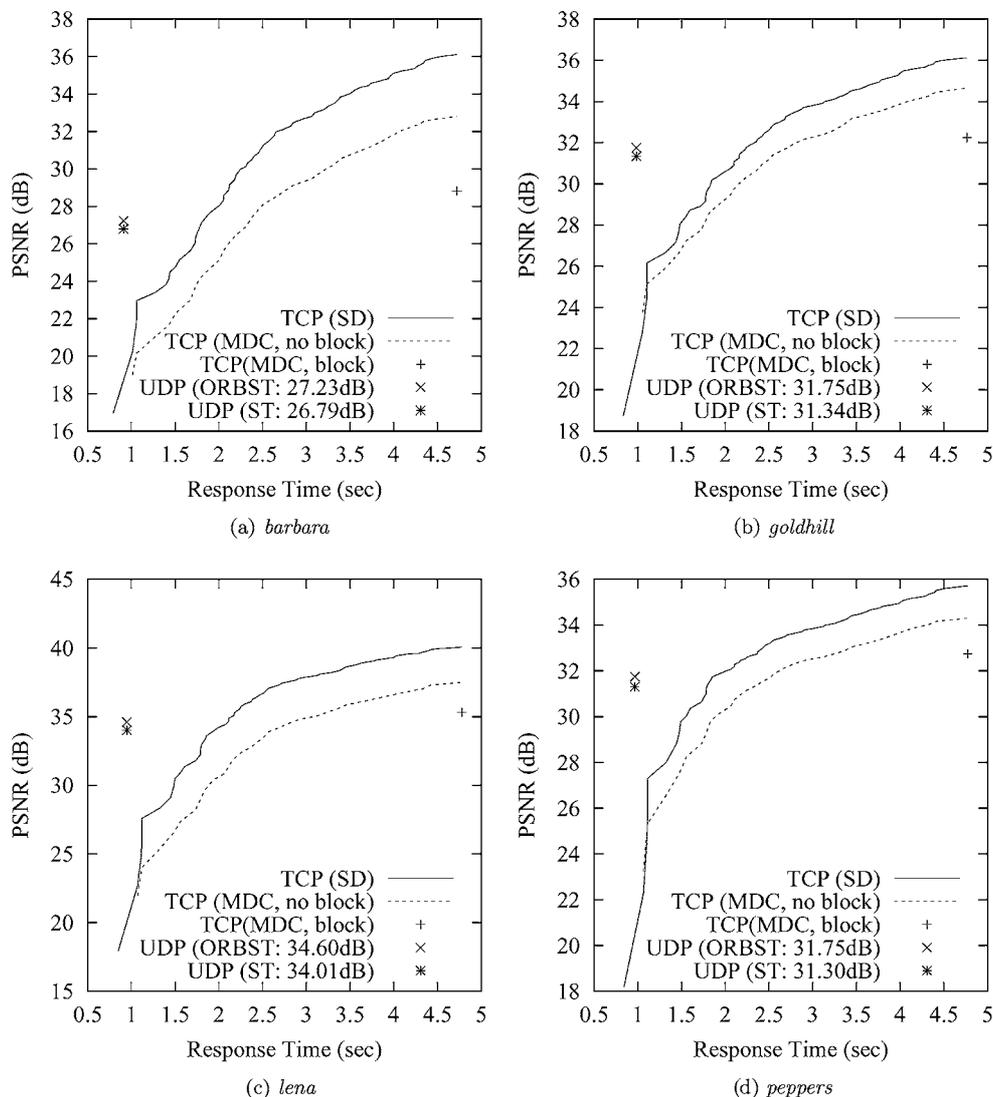


Figure 13. Quality-time trade-offs between TCP delivery of SDC image data and UDP delivery of MDC data for the Champaign-UK connection at 12 noon GMT. (The behavior at other times are similar and are not shown.)

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. Hence, UDP delivery is beneficial when one is interested to preview a coarser image, rather than waiting impatiently for the arrival of a perfect TCP transmitted image.

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 to 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 show 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 to 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

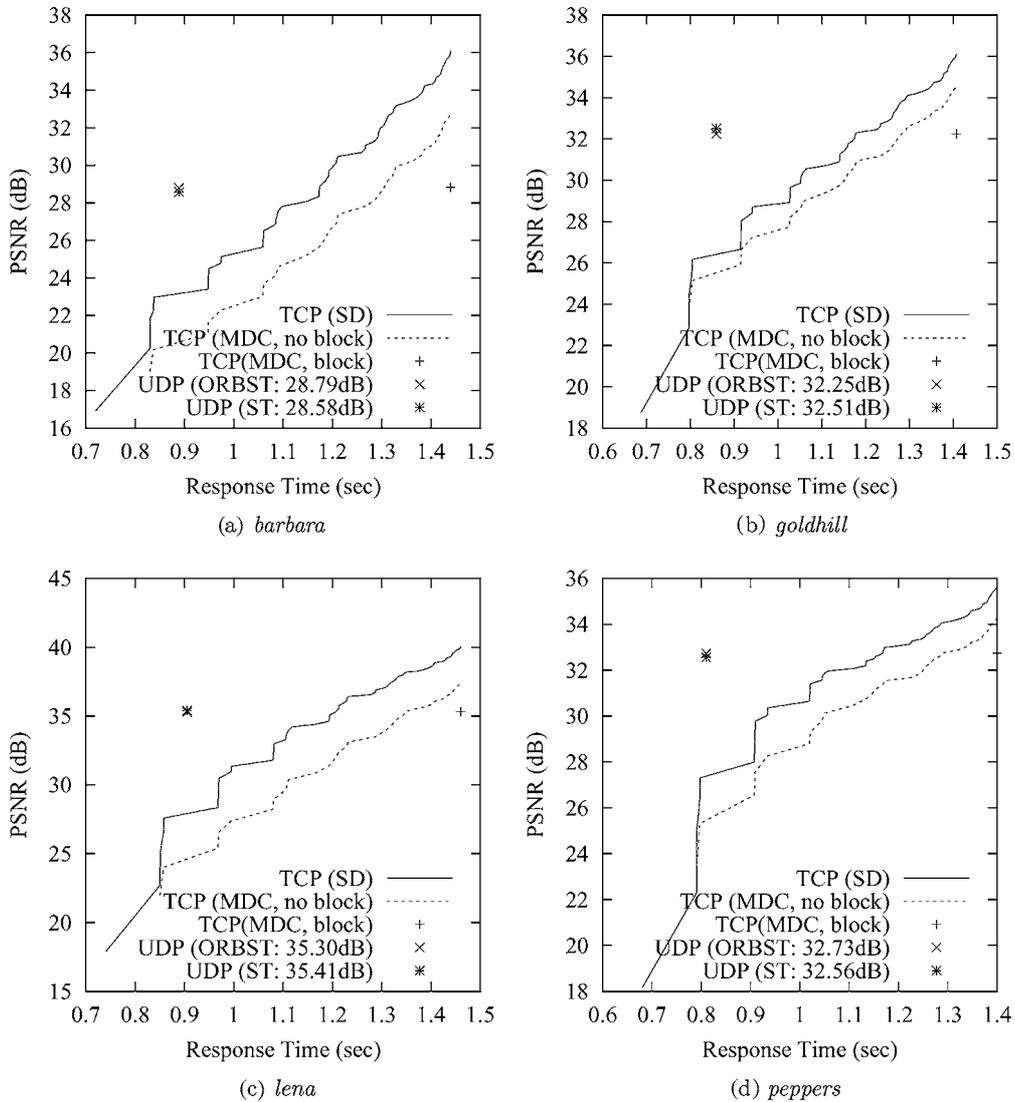


Figure 14. Quality-time trade-offs between TCP delivery of SDC image data and UDP delivery of MDC data for the Champaign-California connection at 12 noon PDT. (The behavior at other times are similar and are not shown.)

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). Note that in the TCP delivery of segmented MDC data, decoding cannot be done until all the packets in both descriptions have been received. The segmentation of an image before coding and packetization is necessary because image data in each description does not contain synchronization points and cannot be decoded when some packets are lost. We plan to study in the

future better strategies for allocating coding rates to segments.

Third, packet losses and reconstructions in the UDP delivery of segmented ST-MDC data lead to further degradations. These degradations are between 1 to 2 dB for the Champaign-China connection. Large degradations may also be caused by the loss of all the packets in an interleaved set. Degradations for the connection to UK are less than 1 dB, and those for the connection to California are negligible. Note that improvements due to ORB-ST when compared to ST in the UDP delivery



(a) *barbara* transferred by UDP to China using segmented MDC (27.85 dB)



(b) *barbara* transferred by TCP using SDC (36.10 dB)



(c) *goldhill* transferred by UDP using segmented MDC (30.59 dB)



(d) *goldhill* transferred by TCP using SDC (36.11 dB)

Figure 15. Images *barbara* and *goldhill* when transferred, respectively, by UDP using segmented MDC and by TCP using SDC.

of segmented MDC data are less than 1 dB (the two points on the left of each graph), as evaluated in the last subsection.

As an illustration, Fig. 15 depicts *barbara* (resp. *goldhill*) when it was delivered to UK (resp. China) by UDP and by TCP. Although they differ by 6–8 dB in PSNR, subjective quality differences are not significant.

The quality-delay trade-offs studied in this paper 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 we plan to study in the future. For TCP delivery, quality improves very quickly in the beginning but 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. The hybrid approach will give better quality than pure UDP delivery and shorter delays than pure TCP delivery.

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 a given amount of data 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 were lost, and the missing information is reconstructed using simple interpolation. In controlled and Internet transmission experiments, we carefully evaluate delay-quality trade-offs between TCP delivery of SDC images and UDP delivery of MDC images, and show 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.

Acknowledgment

This research is supported by a grant from Conexant Systems, Inc.

References

1. International Standards Organization, JPEG 2000 image coding system. Final Committee Draft of ISO International Standard 15444 Part 1.
2. 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*, vol. 8, no. 4, 1998, pp. 422–434.
3. W. Kwok and H. Sun, "Multi-Directional Interpolation for Spatial Error Concealment," *IEEE Trans. on Consumer Electronics*, vol. 39, no. 3, 1993, pp. 455–460.
4. J. Suh and Y. Ho, "Error Concealment Based on Directional Interpolation," *IEEE Trans. on Consumer Electronics*, vol. 43, no. 3, 1997, pp. 295–302.
5. 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*, vol. 9, no. 4, 1999, pp. 648–665.
6. 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*, vol. 9, 1999, pp. 244–254.
7. IETF. An informal management model for diffserv routers: IETF draft standard, July 2000.
8. V.A. Vaishampayan, Design of Multiple Description Scalar Quantizers," *IEEE Trans. on Information Theory*, vol. 39, no. 3, 1993, pp. 821–834.
9. Y. Wang and Q. Zhu, "Error Control and Concealment for Video Communications: A Review," *Proc. of the IEEE*, vol. 86, no. 5, 1998, pp. 974–997.
10. S.D. Servetto, K. Ramchandran, V.A. Vaishampayan, and K. Nahrstedt, "Multiple Description Wavelet Based Image Coding," *IEEE Trans. on Image Processing*, vol. 9, 2000, pp. 813–826.
11. X. Su and B.W. Wah, "Streaming Real-Time Video Data with Optimized Reconstruction-Based DCT and Neural-Network Reconstructions," *IEEE Trans. on Multimedia*, vol. 3, no. 1, 2001, pp. 123–131.
12. L. Rade and B. Westergren, *Mathematics Handbook for Science and Engineering*, Studentlitteratur Birkhauser, 1995.
13. P.J. Davis, *Circulant Matrices*, New York: John Wiley Sons, 1979.
14. A. Said and W.A. Pearlman, "A New Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 6, 1996, pp. 243–250.
15. M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image Coding Using Wavelet Transform," *IEEE Trans. on Image Processing*, vol. 1, 1992, pp. 205–220.



Xiao Su received her BS degree in Computer Science and Engineering from Zhejiang University, Hangzhou, China, in 1994, and her MS and Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 1997 and 2001, respectively.

Since fall 2002, she has been an assistant professor in Computer Engineering Department, San Jose State University. Prior to her appointment with San Jose State University, she worked in the Content Networking Division of Inktomi Corporation from 2001 to 2002. Her areas of interests include multimedia computing and networking, mobile computing and network security.
xsu@email.sjsu.edu



Benjamin W. Wah received his Ph.D. degree in computer science from the University of California, Berkeley, CA, in 1979. He is currently the Robert T. Chien Professor of Electrical and Computer Engineering and a Professor in the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory of the University of Illinois at Urbana-Champaign, Urbana, IL. Previously, he had served on the faculty of Purdue University (1979–85), as a Program Director at the National Science Foundation (1988–89), as Fujitsu Visiting Chair Professor of Intelligence Engineering,

University of Tokyo (1992), and McKay Visiting Professor of Electrical Engineering and Computer Science, University of California, Berkeley (1994). In 1989, he was awarded a University Scholar of the University of Illinois; in 1998, he received the IEEE Computer Society Technical Achievement Award; and in 2000, the IEEE Millennium Medal.

Wah's current research interests are in the areas of nonlinear search and optimization and multimedia signal processing.

Wah was the Editor-in-Chief of the *IEEE Transactions on Knowledge and Data Engineering* between 1993 and 1996, and is the Honorary Editor-in-Chief of *Knowledge and Information Systems*.

He currently serves on the editorial boards of *Information Sciences*, *International Journal on Artificial Intelligence Tools*, *Journal of VLSI Signal Processing, Parallel Algorithms and Applications*, *World Wide Web*, and *Neural Processing Letters*. He had chaired a number of international conferences and was the International Program Committee Chair of the IFIP World Congress in 2000. He has served the IEEE Computer Society in various capacities; among them include Vice President for Publications (1998 and 1999) and President (2001). He is a Fellow of the IEEE and the Society for Design and Process Science.

wah@manip.crhc.uiuc.edu