Progressive Transmission and Compression of Images

A. B. Kiely
Communications Systems and Research Section

We describe an image data compression strategy featuring progressive trans-
mission. The method exploits subband coding and arithmetic coding for compres-
sion. We analyze the Laplacian probability density, which closely approximates
the statistics of individual subbands, to determine a strategy for ordering the com-
pressed subband data in a way that improves rate-distortion performance. Results
are presented for a test image.

I. Introduction

An image data compression system that uses progressive transmission is one that allows a user to
reconstruct successively higher fidelity versions of an image as data are received. The goal of progressive
transmission is thus not only efficient overall compression, but efficient compression at every step.

If the data rate available for image transmission is unexpectedly low, or if the volume of compressed
data exceeds expectations, the available rate will be used to its full extent, to provide nearly the highest-
fidelity image possible given the rate constraint. Alternatively, if the available rate exceeds expectations,
it will be possible to send higher-resolution images than originally planned. In this sense, progressive
transmission strategies are robust with respect to the available data rate.

In situations where the reverse channel can be used, data compression can be combined with advanced
communications strategies to increase the volume of data returned. The use of retransmission schemes is
one example of such a strategy [8]. Progressive transmission gives another method, because it provides
the ability to quickly view low- or medium-resolution previews of an image, making efficient browsing
possible. A user can decide whether to continue the transmission of the full image or proceed to the next
image. This makes more efficient use of the channel because images of less value are not transmitted at
full resolution. This is particularly beneficial to deep-space missions with high data volume and severe
rate constraints.

II. Progressive Transmission and Rate-Distortion Theory

For a given source and distortion metric, we are interested in the trade-off between the rate (usually
measured in average number of bits/sample) and distortion. Some of the rate-distortion functions of
interest include

(1) The rate-distortion limit. This is the theoretical minimum distortion as a function of the average number of bits used to describe the source [2]. This describes the optimum performance obtainable, ignoring constraints such as speed and complexity, and is computable only if the source statistics are known.

(2) The rate-distortion performance achievable by a particular technique.

(3) The rate-distortion performance "progressively" achievable by a particular progressive transmission technique. This is the performance obtained by measuring the rate and distortion of the reconstructed versions of the data at each stage of the transmission.

In theory, we can select a "target" rate-distortion limit point that represents the rate and distortion when all of the compressed data are transmitted. Ideally, using progressive transmission, at every point in the transmission we would like to be as close as possible to the rate-distortion limit. Unfortunately, it is not always possible to meet the target and have all of the points leading up to the target lie along the rate-distortion limit [4].

Equitz and Cover showed that it is possible to meet the rate-distortion limit progressively when solutions to the rate-distortion problem can be written as a Markov chain [4]. Such a source is called "successively refinable." Whether or not a source is successively refinable depends not only on the source statistics but also on the distortion metric selected. For example, a Laplacian source is successively refinable with respect to mean absolute error but, in general, not with respect to other metrics [4]. Even if a source is successively refinable, the optimal strategy for refinement may be unwieldy. Fortunately, even when a source is not successively refinable, the penalty for progressive transmission may be small.

An example follows: Under mean square error (MSE) distortion, a memoryless Gaussian source with variance σ^2 has a rate-distortion limit $MSE = \sigma^2 2^{-2R}$, where R is the rate in bits [2, p. 99]. In a practical source coding system, we might be willing to sacrifice performance in exchange for the simplicity of using, say, a 6-bit uniform scalar quantizer, where each quantizer interval is identified by a 6-bit codeword, and each codeword is compressed by arithmetic coding. For the Gaussian source, this produces a different rate-distortion curve parameterized by the quantizer step size. Both rate-distortion functions are shown in Fig. 1.

Suppose we wish to transmit 24 independent and identically distributed Gaussian samples quantized using the 6-bit uniform quantizer, and we select a target rate of 4 bits/sample, which corresponds to
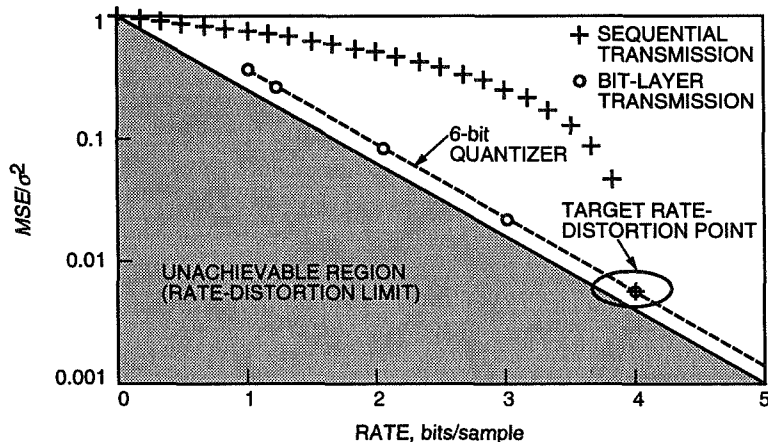

Fig. 1. **Some rate-distortion functions associated with a quantized Gaussian source.**

$MSE/\sigma^2 \approx 5.56 \times 10^{-3}$. Figure 1 shows the progressive rate-distortion performance obtained if we transmit the arithmetic encoded 6-bit quantized samples sequentially. Another option is to reorder the codeword bits before arithmetic coding so that the 24 most-significant bits of the samples are compressed and transmitted first, followed by the next most significant bit from each codeword, and so on. Figure 1 illustrates that although this bit-layer strategy ultimately achieves the same target rate-distortion point as sequential transmission, its progressive performance is superior.

Although in this article we focus on the MSE distortion measure, image quality is subjective and depends on the application. Neither a thousand words nor a single scalar metric can accurately describe the "quality" of an image.

III. Subband Coding

A progressive quantization and compression scheme, such as the technique illustrated in the example of Section II, can be used in combination with a decorrelating transform. In particular, we use a subband coding stage prior to quantization and compression. Unlike block-transform-based methods, such as the discrete cosine transform used in the Joint Photographic Experts Group (JPEG) algorithm, subband coding does not suffer from blockiness artifacts even when used in progressive transmission. For a general description of image compression via subband coding, see [9].

A two-band subband decomposition uses high-pass and low-pass digital filters to decompose a data sequence into high and low subbands, each containing half as many points as the original sequence. This is done independently on horizontal and vertical lines of the image. Because image signal energy is usually concentrated in the lower frequencies, the lowest subband may be repeatedly decomposed. Figure 2 illustrates a two-stage decomposition of the test image.

The filters used for the tests in this article[1] are eighth-order quadrature mirror filters from [1, p. 267]. Because quadrature mirror filters are orthogonal, the MSE introduced by quantization in the transform domain is equal to the MSE of the reconstructed image. We can use this fact to improve the progressive rate-distortion performance by carefully selecting the order in which blocks of information from each subband are transmitted. At each step, we transmit the set of compressed bits giving the largest reduction in MSE per bit. Nonorthogonal filters have the potential to offer improved performance for the same complexity [7]; however, when such filters are used, the analysis in the subsequent sections no longer applies.

The filters are implemented using circular convolution, i.e., each data block is periodically extended before filtering [6]. At the edges of data blocks, this often produces high-frequency components that are processed separately from the rest of the subband data because they are not as easily compressed. Since these components form a small fraction of the image, the penalty for inefficient compression is small. An elegant alternative to circular convolution is to use interpolation at the edges [10]. Since the improvement in rate-distortion performance is small, this alternative is not investigated in this article, though it could be implemented without significant change in the progressive transmission strategy described here.

Each subband generally has a probability density function (PDF) quite close to Laplacian.[2] Figure 3 illustrates an empirical PDF for one of the subbands of the test image. In the following sections, we will analyze a Laplacian source to determine a strategy for the transmission order of the subband data.

---

[1] For this article, we are more concerned with the overall compression and transmission strategy than with the selection of a particular filter.

[2] The lowest subband may be less well behaved and can often be compressed more efficiently after taking differences, although this process makes progressive transmission more difficult. Some performance improvement may be possible by revising the coding strategy used for the lowest subband.
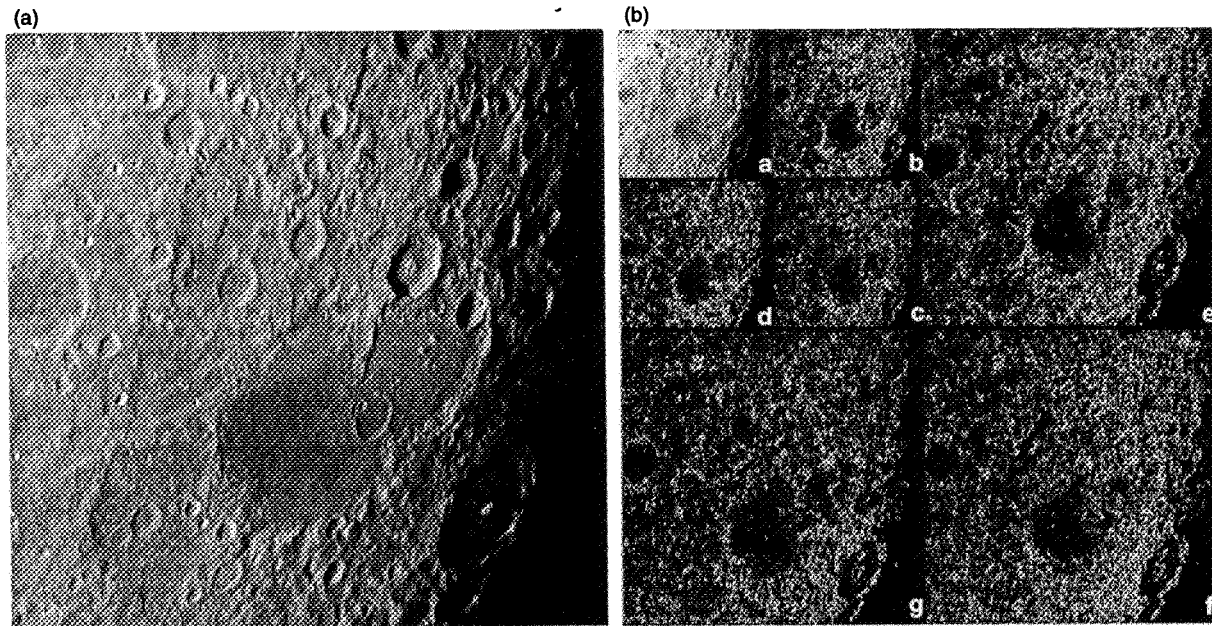
Fig. 2. Two-stage decomposition: (a) original image and (b) image decomposed into subbands labeled a–g (contrast enhanced).
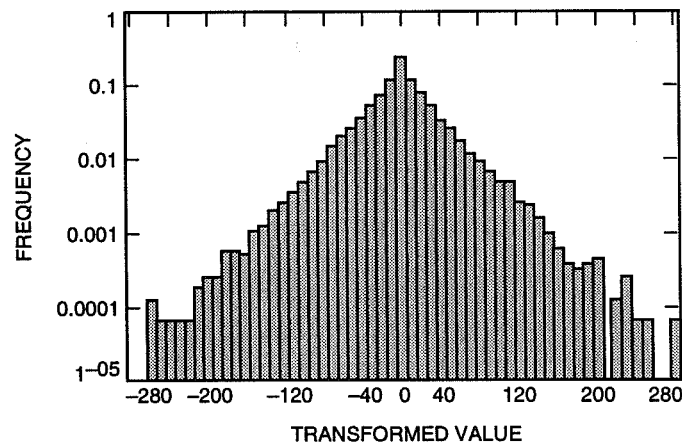

Fig. 3. Empirical PDF of the subband labeled "b" in Fig. 2.

IV. A Block-Oriented Progressive Transmission Scheme

A simple technique for progressive transmission using subband coding is to transmit the subbands in order from lowest to highest frequency [9]. Moderate quality images can be reconstructed even when only the lowest subband has been received, because signal energy in images is usually more concentrated in the lower frequencies. However, the ability to transmit each subband progressively provides an added dimension that can improve the progressive rate-distortion performance because we can switch between the subbands during transmission. In this section, we describe a progressive transmission method for an individual subband.

The progressive transmission scheme we use is bit-wise arithmetic coding [5]. A source sample is quantized, and each quantizer interval is identified by a b-bit codeword. Figure 4 illustrates the codeword
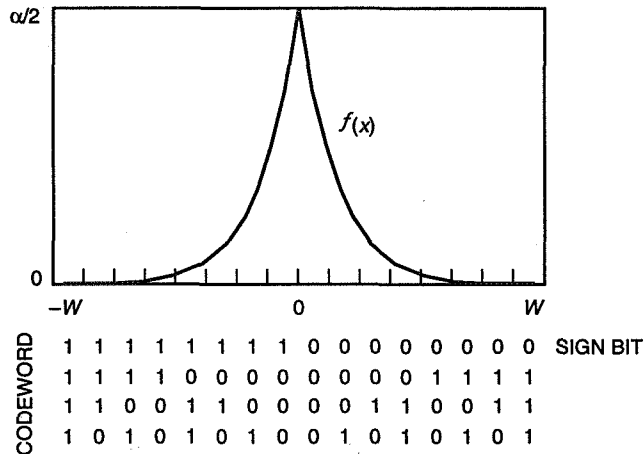
Fig. 4. Codeword assignment for a 4- bit quantizer.

assignment for a 4-bit uniform scalar quantizer along with the PDF $f(x)$ for a Laplacian source with variance $2/\alpha^2$. We can see that the bit assignment scheme is progressive—deleting i bits from the end of each codeword produces the same effect as using a lower resolution $b - i$ bit quantizer. The codeword assignment scheme also ensures that, for a Laplacian source, a zero is more likely in every bit position, with the exception of the "sign" bit.

The codewords corresponding to the quantized coefficients of a subband are grouped together. The ith bit layer, which consists of the ith bit from each codeword in the group, is compressed using the block-adaptive binary arithmetic encoder described in [5]. Each layer is compressed independently, i.e., the arithmetic encoder uses an estimate of the unconditional probability of a zero at each layer.

Slight modification of the quantizer can improve the progressive rate-distortion performance. Figure 5 shows three different quantizer options (illustrated for 3-bit quantizers). We can evaluate the rate and MSE distortion of each quantizer using the results of Appendix A. In each case, the quantizer range $[-W, W]$, step size Δ, and number of bits b are related by

$$\Delta = W2^{1-b}$$

A continuous rate-distortion performance curve for a quantizer can be obtained by fixing b and varying Δ. Figure 6 shows several such curves for the quantizer illustrated in Fig. 5(c). Decreasing Δ increases rate and lowers distortion until a minimum distortion point is reached. After this point, the distortion increases and the quantizer becomes inefficient. Ordinarily this inefficient region is not shown.

When we use a quantizer progressively, the range is fixed, but the number of bits b increases, reducing step size by half for each successive bit layer transmitted. The points in Fig. 6 illustrate progressive rate-distortion performance. We observe in the figure that as b becomes large, the distortion does not approach zero, but is dominated by the contribution from the overload regions. As this happens, the progressive rate-distortion performance flattens, as points are lying along the inefficient regions of the rate-distortion curves. This floor is easy to find and depends only on the range and the Laplacian parameter α. We find from Eq. (A-1) in Appendix A that

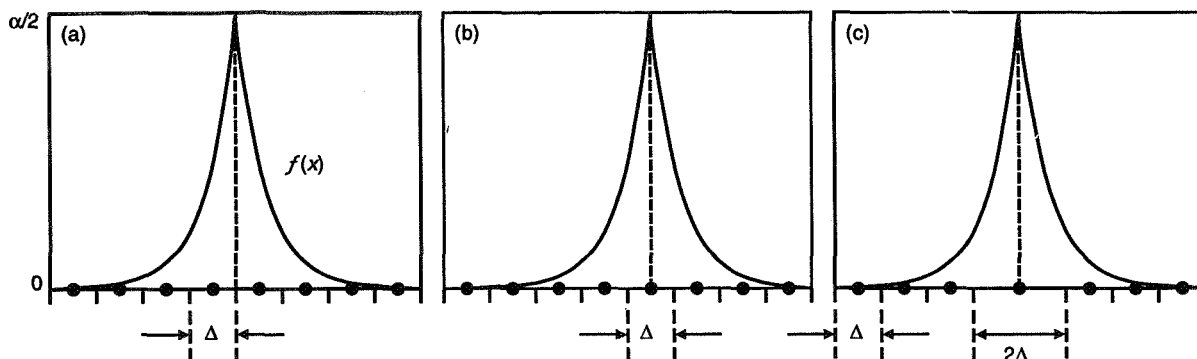$$\lim_{b \to \infty} \frac{MSE}{\sigma^2} = e^{-\alpha W}$$

Fig. 5. Different quantizer options: (a) uniform symmetric with bin boundary at origin, (b) uniform with reconstruction point at origin, and (c) almost uniform with enlarged center bin.
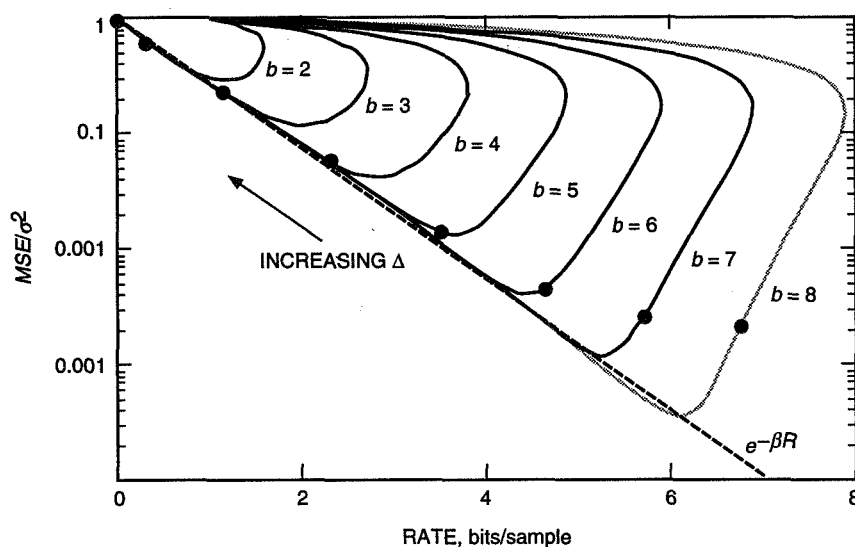

Fig. 6. Rate-distortion performance curves for the quantizer of Fig. 5(c).

when the reconstruction point of each interval is the interval midpoint.[3] In a real application, the range of possible values is limited by instrument dynamic range and the filters chosen, so the MSE floor can be avoided by making the quantizer range sufficiently large.

Figure 7 compares the performance of each quantizer option of Fig. 5. The rate-distortion curves shown correspond to 8-bit quantizers. The progressive performance points are obtained by transmitting bit-layers with a target rate of 6 bits/sample. In Appendix A, we derive rate-distortion functions used to compare the quantizers.

Quantizer (a), a symmetric uniform quantizer, has poor performance at low bit rates, because the sign bit is incompressible; hence, rates below 1 bit/sample are not achievable. Quantizer (b) is obtained by shifting quantizer (a) by $\Delta/2$, so that a reconstruction point is present at the origin, as done in [5]. When Δ is large, low rates are achievable, and the performance obtained by varying Δ is close to the rate-distortion limit. However, when this quantizer is used progressively and the step size is small, the entropy of the sign bit approaches 1. Thus, the progressive performance is generally poor at low bit rates.

---

[3] This quantity is reduced by half when the centroid of each quantizer interval is used (see Eq. (A-2) in Appendix A).
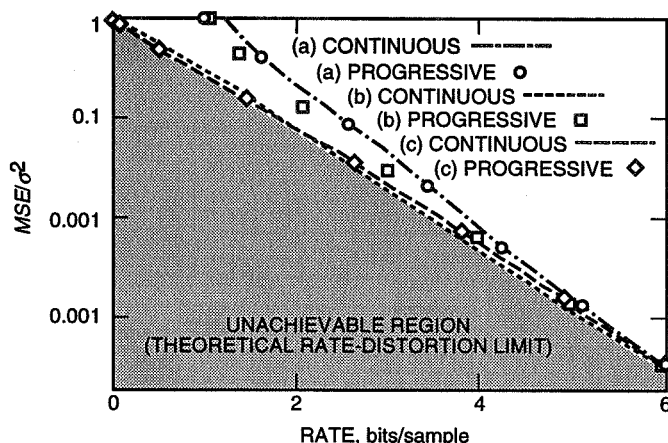
Fig. 7. Rate-distortion performance of different 8-bit quantizers. The continuous curves are obtained by varying the quantizer step size Δ while keeping the number of bits fixed. The progressive performance points are obtained by calculating the rate-distortion performance after the transmission of each bit layer.

To overcome these problems, we use quantizer (c) with range $[-W, W]$, which effectively combines the two center regions of quantizer (a), in combination with a different transmission order. We do not initially transmit the sign bit layer, but rather we begin by transmitting the next layer (after arithmetic coding). Then, for each codeword where a "1" appeared (i.e., each quantized value in the range $[-W, -W/2] \cup [W/2, W]$), the sign bit is transmitted. Then the next layer is transmitted, followed by the sign bits for each codeword representing a quantized value in the range $[-W/2, -W/4] \cup [W/4, W/2]$. This continues at each layer. Finally, if the quantized value lies in the range $[-W2^{1-b}, W2^{1-b}]$ (the two centermost quantization points), the sign bit is never transmitted. In progressive reconstruction, for any quantized value for which no sign bit was received, the origin is used as the reconstruction point. The advantage of this transmission order is that we delay the transmission of the sign bits, which are not easily compressible. Figure 7 shows that this technique offers improved progressive performance compared to quantizers (a) and (b) and is, in fact, close to the rate-distortion limit for a Laplacian source with MSE distortion, which is computed in Appendix B.

V. Ordering the Compressed Subband Data

In the previous section, we described a progressive transmission strategy that can be used for each subband. This strategy allows us to interrupt transmission of a subband to transmit data from another subband. In this section, we describe a method for choosing the order of transmission of subband data designed to improve the progressive rate-distortion performance.

At any point in transmission, we wish to transmit the next (compressed) layer of bits from the subband that gives the largest reduction in MSE per transmitted bit. Thus, we select the bit layer from the subband with a rate-distortion curve whose slope is minimum. Using the analytical expressions for rate and distortion of the quantizer (see the Appendices) is rather intractable, so we use the approximation

$$MSE \approx \sigma^2 e^{-\beta R} = \frac{2}{\alpha^2} e^{-\beta R}$$

where R is the rate in bits, and $\beta = 1.3$. This approximation is shown in Fig. 6. Using this approximation, omitting some algebraic details, we have the following subband selection strategy: Transmit the next bit layer from the subband with parameter α and rate R that minimizes $\ln \alpha + \beta R/2$.

To apply this strategy, we need to estimate the Laplacian parameter α for each subband or, equivalently, to estimate the mean absolute value $1/\alpha$. For each subband, we update the estimate at each bit layer transmitted. The estimate is created by keeping track of the frequency of sign bits transmitted at each bit layer, which is a quantity already computed in the compression stage.
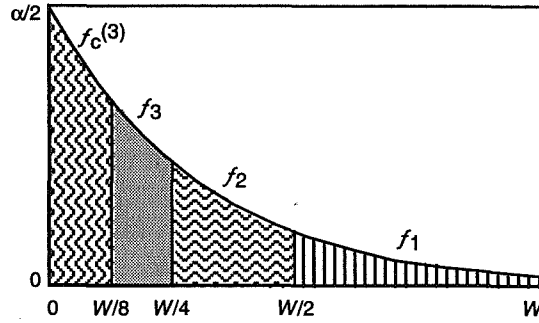

Fig. 8. Frequencies used to estimate the Laplacian parameter α at the third bit layer.

Let f_i denote the frequency of sign bits transmitted at the ith layer, and $f_c^{(i)}$ denote the frequency of the center region at the ith layer. This is illustrated in Fig. 8. For a Laplacian distribution with parameter α, at the ith bit layer, we expect

$$f_j = \begin{cases} e^{-\alpha W/2}, & j = 1 \\ e^{-\alpha W 2^{-j}} - e^{-\alpha W 2^{1-j}}, & 1 < j \le i \end{cases}$$

$$f_c^{(i)} = 1 - e^{-\alpha W 2^{-i}}$$

which gives parameter estimates of

$$\hat{\alpha}_j = \begin{cases} -\dfrac{2}{W} \ln f_1, & j = 1 \\[2mm] -\dfrac{2^j}{W} \ln \left[\dfrac{1 \pm \sqrt{1 - 4f_j}}{2} \right], & j > 1, f_j \le 1/4 \\[2mm] \dfrac{2^j}{W} \ln 2, & j > 1, f_j > 1/4 \end{cases} \tag{1}$$

$$\hat{\alpha}_c^{(i)} = -\frac{2^i}{W} \ln \left(1 - f_c^{(i)} \right)$$

where, in Eq. (1), we choose the estimator closest to $\hat{\alpha}_c^{(i)}$. For our overall estimate of the mean absolute value at the ith bit layer, we take the weighted average of the above estimators:

$$\frac{1}{\hat{\alpha}} = \frac{f_c^{(i)}}{\hat{\alpha}_c^{(i)}} + \sum_{j=1}^{i} \frac{f_j}{\alpha_j} \tag{2}$$

If any of the f_i's are zero, we can assign any finite value to the corresponding $\hat{\alpha}_i$, since it will have no weight in Eq. (2).

Since the quantizer range may be quite large compared to the variance of a subband, it is common for the first few bit layers to consist entirely of zeros. For each subband, we first transmit an integer identifying the number of leading all-zero bit layers. This ensures that $\hat{\alpha}$ is always well defined.

VI. Results

Curve (a) in Fig. 9 illustrates the progressive rate-distortion performance on the test image obtained using the subband selection strategy described in Section V combined with the parameter estimates from Eq. (2). In the figure, we compare this performance to the more traditional progressive transmission method (transmitting the subbands sequentially from lowest to highest frequency), shown as curve (b), and the optimum transmission order (that obtained by testing every possible transmission order of the subband data), curve (c). We also show the performance of the graphic-in-line-format (GIF) compression technique, a progressive transmission method that has recently become popular in Web browsing software.

In Fig. 10, we show four intermediate reconstructed images obtained in progressive transmission. The rate-distortion points corresponding to these four images are shown in Fig. 9.
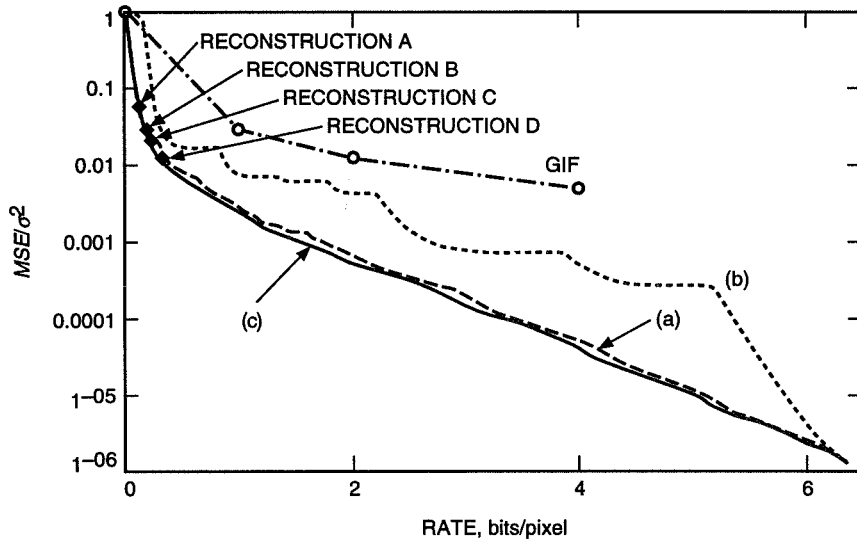

Fig. 9. Progressive rate-distortion performance on the test image using 12-bit quantizers on each subband. Curve (a) shows the performance obtained using the strategy described in Section V. Curve (b) is the performance obtained when we transmit the subbands in order from the lowest to the highest frequency. Curve (c) is the performance obtained if we choose the optimum subband bit layer at each step. The curve labled "GIF" corresponds to the performance obtained using the GIF progressive transmission technique. The images corresponding to the four labled reconstruction points are shown in Fig. 10.

VII. Conclusion

We have shown that careful selection of the quantization and compression strategy can significantly improve the progressive rate-distortion performance of a data compression system. The use of transforms such as subband coding can lead to further improvements, because they provide additional freedom in choosing the order in which the compressed data are transmitted.

When an effective progressive transmission scheme is used, the progressive rate-distortion performance may be good at every point in transmission. For image compression, this means that even if the rate
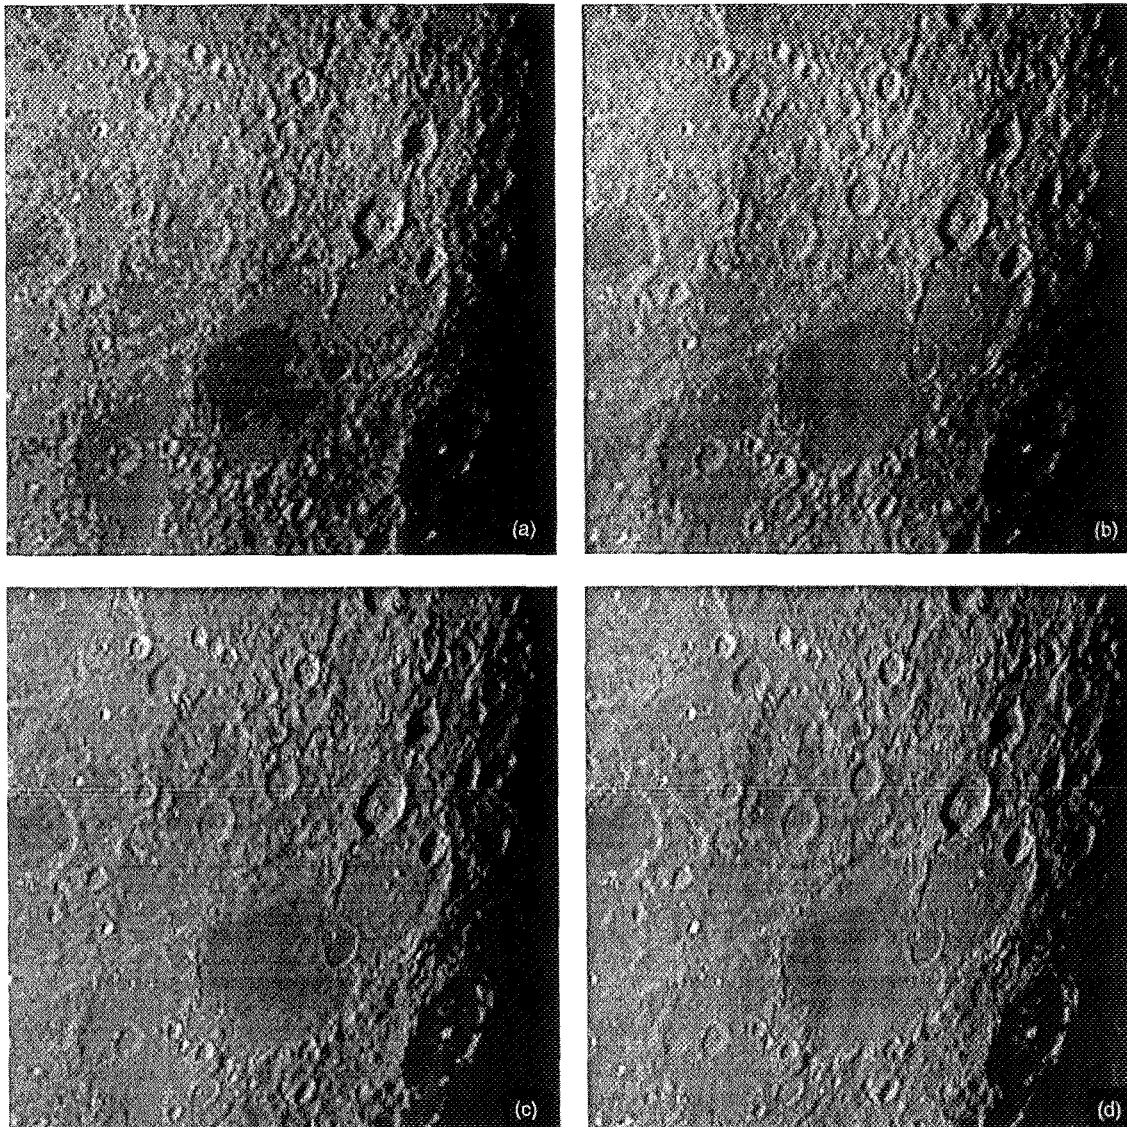
Fig. 10. Reconstructed images at various stages in the transmission. The savings compared to transmission of the original 8-bit image without compression are (a) 18.2, (b) 16.5, (c) 15.5, and (d) 13.9 dB, respectively.

constraints change during a mission, redesign of the compression strategy is not necessary. When the reverse channel can be used for browsing, the benefits of progressive transmission become even greater—we are not required to decide in advance the amount of resources devoted to a particular image.

It is interesting to consider the savings in dB that can be offered by lossy data compression. Suppose that 2:1 lossless compression is possible for the original image shown in Fig. 2. In this case, the lower-quality image of Fig. 10(d) could be obtained at a savings of approximately 10.9 dB. Alternatively, for the price of a single losslessly compressed image, we could obtain 12 images of quality comparable to that of Fig. 10(d).

References

[1] A. N. Akansu and R. A. Haddad, *Multiresolution Signal Processing*, San Diego, California: Academic Press, 1992.

[2] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression*, Englewood Cliffs, New Jersey: Prentice-Hall, 1971.

[3] S. Dolinar, "Maximum-Entropy Probability Distributions Under L_p-Norm Constraints," *The Telecommunications and Data Acquisition Progress Report 42-104, October–December 1990*, Jet Propulsion Laboratory, Pasadena, California, pp. 74–87, February 15, 1991.

[4] W. H. R. Equitz and T. M. Cover, "Successive Refinement of Information," *IEEE Trans. Inform. Theory.* vol. 37, pp. 269–275, March 1991.

[5] A. B. Kiely, "Bit-Wise Arithmetic Coding for Data Compression," *The Telecommunications and Data Acquisition Progress Report 42-117, January–March 1994*, Jet Propulsion Laboratory, Pasadena, California, pp. 145–160, May 15, 1994.

[6] A. B. Kiely and F. Pollara, "A Seismic Data Compression System Using Subband Coding," *The Telecommunications and Data Acquisition Progress Report 42-121, January–March 1995*, Jet Propulsion Laboratory, Pasadena, California, pp. 242–251, May 15, 1995.
http://edms-www.jpl.nasa.gov/tda/progress_report/42-121/121J.pdf

[7] E. Majani, "Low-Complexity Wavelet Filter Design for Image Compression," *The Telecommunications and Data Acquisition Progress Report 42-119, July–September 1994*, Jet Propulsion Laboratory, Pasadena, California, pp. 181–200, December 15, 1994.
http://edms-www.jpl.nasa.gov/tda/progress_report/42-119/119E.pdf

[8] F. Pollara and L. Ekroot, "Analysis of Automatic Repeat Request Methods for Deep-Space Downlinks," *The Telecommunications and Data Acquisition Progress Report 42-122, April–June 1995*, Jet Propulsion Laboratory, Pasadena, California, pp. 66–83, August 15, 1995.
http://edms-www.jpl.nasa.gov/tda/progress_report/42-122/122L.pdf

[9] M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, Bellingham, Washington: SPIE Press, 1991.

[10] J. R. Williams and K. Amaratunga, *A Discrete Wavelet Transform Without Edge Effects Using Wavelet Extrapolation*, IESL Technical Report 95-02, Massachusetts Institute of Technology, Cambridge, Massachusetts, January 28, 1995.
http://www-iesl.mit.edu/pub_docs/Wavelets/15_FWT_poly-94.html

Appendix A

Performance of a Scalar Quantizer on a Laplacian Source

In this appendix, we analyze the rate-distortion performance of a Laplacian source with PDF $f(x)$ using a scalar quantizer that is nearly uniform. Consider the symmetric b-bit quantizer (with $2^b - 1$ intervals) shown in Fig. A-1. The reconstruction point of each interval may be the midpoint or the centroid of the interval.
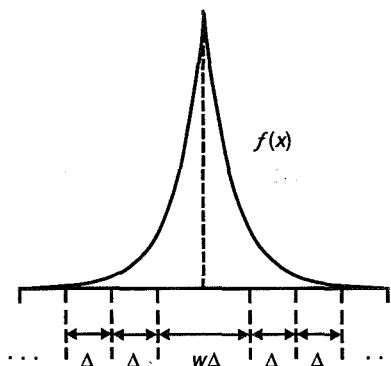

Fig. A-1. A symmetric scalar quantizer.

The normalized MSE obtained when this quantizer is used is

$$\frac{MSE}{\sigma^2} = \frac{\alpha^2}{2} \left[MSE_c + 2MSE_{ov} + 2 \sum_{i=0}^{2^{b-1}-3} MSE_i \right]$$

where MSE_c is the MSE contribution of the center region,

$$MSE_c = \frac{1}{\alpha^2} \left[2 - e^{-y} \left(y^2 + 2y + 2 \right) \right]$$

and $y = \alpha\Delta/2$. MSE_i is the MSE contribution of the region $[(w/2 + i)\Delta, (w/2 + i + 1)\Delta)$, $i = 0, 1, \cdots, 2^{b-1} - 3$,

$$MSE_i = \frac{e^{-y-si}}{2\alpha^2} \left[1 - e^{-s} + \frac{s^2}{1 - e^s} \right]$$

when the centroid is used, and

$$MSE_i = \frac{e^{-y-si}}{8\alpha^2} \left[(8 + s^2)(1 - e^{-s}) - 4s(1 + e^{-s}) \right]$$

when the midpoint is used, where $s = \alpha\Delta$. MSE_{ov} is the MSE of each overload region $\pm[(w/2 + 2^{b-1} - 2)\Delta, \infty)$,

$$MSE_{ov} = \frac{e^{s(2-M)}e^{-y}}{2\alpha^2}$$

when the centroid is used, and

$$MSE_{ov} = \left(\frac{s^2 - 4s + 8}{8\alpha^2}\right) e^{s(2-2^{b-1})}e^{-y}$$

when the midpoint is used.

Combining, we find that when midpoints are used,

$$\frac{MSE}{\sigma^2} = 1 - e^{-y}\left[\frac{1}{2}y(y+2) + \frac{1}{8}s(4-s) + \frac{s\left(1 - e^{s(2-2^{b-1})}\right)}{e^s - 1}\right] \qquad \text{(A-1)}$$

and when centroids are used,

$$\frac{MSE}{\sigma^2} = 1 - \frac{1}{2}e^{-y}\left[(y+1)^2 + \frac{s^2 e^s \left(1 - e^{s(2-2^{b-1})}\right)}{(e^s - 1)^2}\right] \qquad \text{(A-2)}$$

Next we compute the rate. The center region has probability $P_c = 1 - e^{-y}$; each overload region has probability $P_{ov} = (1/2)e^{-y}e^{s(2-2^{b-1})}$; and the region $[(w/2 + i)\Delta, (w/2 + i + 1)\Delta)$ has probability $(1/2)e^{-y-is}(1 - e^{-s})$. The entropy is

$$H = (e^{-y} - 1)\ln(1 - e^{-y}) + e^{-y}(y + \ln 2) + e^{-y}\left(1 - e^{s(2-2^{b-1})}\right)\left[\frac{s}{e^s - 1} - \ln(1 - e^{-s})\right] \text{ nats}$$

Using the bit-wise transmission strategy described in Section IV, the sign bit has entropy of 1 bit but is not transmitted when the sample value lies in the center region; this contributes e^{-y} bits to the rate. If j denotes the bit layer index, with $j = 1$ corresponding to the least significant layer and $j = b$ the sign bit layer, then the probability that the jth bit is a 1 is

$$e^{s-y}\left(\frac{e^{-s2^{j-1}} + e^{-s2^{b-1}}}{1 + e^{-s2^{j-1}}}\right)$$

so the rate of bit-wise transmission is

$$R = e^{-y} + \sum_{j=0}^{b-2} \mathcal{H}_2\left[e^{s-y}\left(\frac{e^{-s2^j} + e^{-s2^{b-1}}}{1 + e^{-s2^j}}\right)\right] \text{ bits}$$

where $\mathcal{H}_2(x) \triangleq -\log_2 x - (1-x)\log_2(1-x)$ is the binary entropy function.

100

These expressions are used to compare the different quantizers in Section IV. We can also observe the losses resulting from the bit-wise transmission strategy and using the midpoint rather than the centroid of each quantizer interval. In fact, these losses are negligible, as can be seen by comparing the rate-distortion curves in Fig. 7.

Appendix B

Rate-Distortion Limit For a Laplacian Source With MSE Distortion

In this appendix, we derive the rate-distortion limit of a Laplacian source using a scalar quantizer with MSE distortion metric. A Laplacian source with parameter α has PDF $f(x) = (\alpha/2)e^{-\alpha|x|}$ and variance $\sigma^2 = 2/\alpha^2$. For a discussion of the mean absolute error distortion metric for this source, see [3].

The goal is to find the optimum (possibly infinite) scalar quantizer. Let the quantizer regions be $[t_i, t_{i+1})$, $i = \cdots, -1, 0, 1, 2, \cdots$. The MSE distortion of the quantizer is $MSE = \sum_{i=-\infty}^{\infty} MSE_i$, where MSE_i is the contribution to the MSE from the region $[t_i, t_{i+1})$ (assume $0 \leq t_i < t_{i+1}$; a similar analysis applies when this is not the case):

$$MSE_i = \int_{t_i}^{t_{i+1}} (x - c_i)^2 f(x)dx = \frac{e^{-\alpha t_i} - e^{-\alpha t_{i+1}}}{2\alpha^2} + \frac{(t_{i+1} - t_i)^2}{2(e^{\alpha t_i} - e^{\alpha t_{i+1}})} \tag{B-1}$$

Here c_i is the centroid of the region (this choice minimizes MSE):

$$c_i = \frac{1}{P_i} \int_{t_i}^{t_{i+1}} xf(x)dx = \frac{1}{\alpha} + \frac{t_i e^{-\alpha t_i} - t_{i+1} e^{-\alpha t_{i+1}}}{e^{-\alpha t_i} - e^{-\alpha t_{i+1}}}$$

and $P_i = (1/2)(e^{-\alpha t_i} - e^{-\alpha t_{i+1}})$ is the probability of the region. The total entropy associated with the quantizer is $H = \sum_{i=-\infty}^{\infty} -P_i \ln P_i$, nats.

Consider what happens when we vary the threshold t_i:

$$\frac{\partial MSE}{\partial t_i} = \frac{(t_i - t_{i-1})}{2(e^{\alpha t_{i-1}} - e^{\alpha t_i})^2} \left[2e^{\alpha t_{i-1}} + e^{\alpha t_i}(\alpha t_i - \alpha t_{i-1} - 2) \right]$$

$$- \frac{(t_{i+1} - t_i)}{2(e^{\alpha t_i} - e^{\alpha t_{i+1}})^2} \left[-2e^{\alpha t_{i+1}} + e^{\alpha t_i}(\alpha t_{i+1} - \alpha t_i + 2) \right]$$

and

$$\frac{\partial H}{\partial t_i} = \frac{\alpha}{2e^{\alpha t_i}} \ln \left[\frac{e^{-\alpha t_i} - e^{-\alpha t_{i+1}}}{e^{-\alpha t_{i-1}} - e^{-\alpha t_i}} \right]$$

So by the chain rule, the slope of the rate-distortion curve (when varying t_i) is $\partial MSE/\partial H = (\partial MSE/\partial t_i)/\partial H/\partial t_i$. If the quantizer is locally uniform with step size Δ, i.e., $t_{i-1} = t_i - \Delta$, $t_{i+1} = t_i + \Delta$, then (omitting algebraic manipulations),

$$\frac{\partial MSE}{\partial H} = \frac{(2-s)e^s - 2 - s}{\alpha^2(e^s - 1)}$$

where $s = \alpha\Delta$. By the Kuhn–Tucker conditions, since this expression depends only on the normalized step size s, the optimal scalar quantizer is uniform (and infinite) on both sides of the origin and has the same step size on both sides.

Next, we need to determine the behavior of the quantizer near the origin. Suppose there is an arbitrary (possibly different) offset on each side of the origin. By symmetry, setting these two offsets equal will give a solution to the Kuhn–Tucker conditions. So we want to know the width of the center region relative to the normalized step size s.

Let the width of the center region be $w\Delta$. The center region has probability mass $1 - e^{-y}$, where $y = sw/2$, and makes MSE contribution

$$MSE_c = 2\int_0^{w\Delta/2} x^2 f(x)dx = \frac{1}{\alpha^2}\left[2 - \left(y^2 + 2y + 2\right)e^{-y}\right]$$

The right-hand side regions are $[(w/2 + i)\Delta, (w/2 + i + 1)\Delta)$, $i = 0, 1, 2, \cdots$. Using $t_i = (w/2 + i)\Delta$ and $t_{i+1} = (w/2 + i + 1)\Delta$ in Eq. (B-1) yields

$$MSE_i = \frac{e^{-si-y}}{2\alpha^2}\left[1 - e^{-s} + \frac{s^2}{1 - e^s}\right]$$

$$P_i = \frac{1}{2}e^{-si-y}(1 - e^{-s})$$

so the rate contribution is $R_i = -P_i \ln P_i = (1/2)e^{-si-y}(1 - e^{-s})[y + si + \ln 2 - \ln(1 - e^{-s})]$. The total MSE is

$$MSE = MSE_c + 2\sum_{i=0}^{\infty} MSE_i = \frac{1}{\alpha^2}\left\{2 + e^{-y}\left[\frac{s^2}{(1 - e^s)(1 - e^{-s})} - (y+1)^2\right]\right\}$$

(Alternatively, this can be derived from Eq. (A-2) in Appendix A taking the limit as $b \to \infty$.) The entropy is

$$H = -P_c \ln P_c + 2\sum_{i=0}^{\infty}(-P_i \ln P_i) = e^{-y}\left\{\frac{s}{e^s - 1} + \ln\left[\frac{2(e^y - 1)}{1 - e^{-s}}\right]\right\} - \ln(1 - e^{-y})$$

Given these expressions for MSE and entropy in terms of s and y, we apply the Kuhn–Tucker conditions once more. When s is fixed and y varies, the resulting rate-distortion curve has slope

$$\frac{1 - \dfrac{s^2 e^s}{(e^s - 1)^2} - y^2}{2\left(\dfrac{s}{e^s - 1} + \ln\left[\dfrac{2(e^y - 1)}{1 - e^{-s}}\right]\right)}$$

When y is fixed and s varies, the resulting rate-distortion curve has slope

$$1 - s\left(\frac{1}{2} + \frac{1}{e^s - 1}\right)$$

By the Kuhn–Tucker conditions, these two quantities must be equal. We can use this fact to solve for y in terms of s numerically. We find that w varies between 1 (high rate limit) and 2 (low rate limit). The resulting rate-distortion limit is shown in Fig. 7.