

On the Accuracy and Complexity of Rate-Distortion Models for Fine-Grained Scalable Video Sequences

CHENG-HSIN HSU and MOHAMED HEFEEDA
Simon Fraser University

Rate-distortion (R-D) models are functions that describe the relationship between the bitrate and expected level of distortion in the reconstructed video stream. R-D models enable optimization of the received video quality in different network conditions. Several R-D models have been proposed for the increasingly popular fine-grained scalable video sequences. However, the models' relative performance has not been thoroughly analyzed. Moreover, the time complexity of each model is not known, nor is the range of bitrates in which the model produces valid results. This lack of quantitative performance analysis makes it difficult to select the model that best suits a target streaming system. In this article, we classify, analyze, and rigorously evaluate all R-D models proposed for FGS coders in the literature. We classify R-D models into three categories: analytic, empirical, and semi-analytic. We describe the characteristics of each category. We analyze the R-D models by following their mathematical derivations, scrutinizing the assumptions made, and explaining when the assumptions fail and why. In addition, we implement all R-D models, a total of eight, and evaluate them using a diverse set of video sequences. In our evaluation, we consider various source characteristics, diverse channel conditions, different encoding/decoding parameters, different frame types, and several performance metrics including accuracy, range of applicability, and time complexity of each model. We also present clear systematic ways (pseudo codes) for constructing various R-D models from a given video sequence. Based on our experimental results, we present a justified list of recommendations on selecting the best R-D models for video-on-demand, video conferencing, real-time, and peer-to-peer streaming systems.

Categories and Subject Descriptors: H.1.1 [**Models and Principles**]: Systems and Information Theory

General Terms: Performance

Additional Key Words and Phrases: Multimedia streaming, fine-grained scalable coding, rate-distortion models

ACM Reference Format:

Hsu, C.-H. and Hefeeda, M. 2008. On the accuracy and complexity of rate-distortion models for fine-grained scalable video sequences. *ACM Trans. Multimedia Comput. Commun. Appl.* 4, 2, Article 15 (May 2008), 22 pages. DOI = 10.1145/1352012.1352019 <http://doi.acm.org/10.1145/1352012.1352019>

1. INTRODUCTION

Video streaming on the Internet is increasingly getting very popular. The best-effort service offered by the Internet, however, poses unique challenges for high-quality video streaming. These challenges include heterogeneity and bandwidth variability in network channels between streaming servers and

This work was partially supported by the National Science and Engineering Research Council (NSERC) of Canada.

Authors' address: Simon Fraser University, 250-13450 102nd Ave, Surrey BC V3T AO3, Canada; email: {cha16, mhefeeda}@cs.sfu.ca.

M. Hefeeda is also affiliated with Mansoura University, Egypt.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2008 ACM 1551-6857/2008/05-ART15 \$5.00 DOI 10.1145/1352012.1352019 <http://doi.acm.org/10.1145/1352012.1352019>

ACM Transactions on Multimedia Computing, Communications and Applications, Vol. 4, No. 2, Article 15, Publication date: May 2008.

clients. These challenges require streaming systems to support bitrate scalability and error resiliency. Traditional streaming systems partially cope with these challenges using either multilayer or multi-description encoding of streams. These solutions, however, provide limited (coarse-grain) rate scalability: clients receiving incomplete layers or descriptions cannot use them to enhance display quality. These solutions also suffer from poor error resiliency, because the loss or corruption of a few bits render the entire layer useless.

In contrast to traditional multilayer video coding, fine granularity scalability (FGS) coding has been proposed to provide finer bitrate scalability and better error resiliency [Radha et al. 2001; Li 2001]. An FGS encoder compresses video data into two layers: a base layer that provides basic quality, and a single enhancement layer that adds incremental quality refinements proportional to the number of bits received. Arbitrary truncation (at the bit level) of the enhancement layer to achieve a target bitrate is possible, and more importantly, it does not require complex or resource-intensive operations from the streaming servers or their proxy caches. This in turn enables streaming servers to scale to larger and more heterogeneous sets of clients.

Given the flexibility of controlling the bitrate provided by FGS encoders and the constraints on and the variability of the channel bandwidth, researchers seek to optimize the quality of received video streams. A common method in the literature to achieve such quality-optimized systems is through the use of rate-distortion (R-D) models. R-D models are functions that describe the relationship between the bitrate and expected level of distortion in the reconstructed video stream. Knowing the R-D models enables us, for example, to determine the required bitrate to achieve a target quality, to optimally allocate a given bandwidth among frames, and to prioritize bits within the same frame. Clearly, the accuracy of the R-D models directly impacts the performance of streaming systems using them. In addition, the time cost of constructing various R-D models for a given sequence may prefer, or even dictate, a certain model over others. For example, in real-time streaming systems, building an R-D model should be fast enough to cope with the timing constraints of the video stream.

Due to the increasing importance and adoption of FGS coding systems, several studies have proposed R-D models for them [Dai and Loguinov 2003; Dai et al. 2004; Sun et al. 2005; Dai et al. 2006]. Each of these studies conducted limited performance evaluation, just enough to show the merits of the proposed model. The relative performance of different R-D models has not yet been thoroughly analyzed. Moreover, the time complexity of each model is not known, nor is the range of bitrates in which the model produces valid results. In addition, previous studies do not provide enough specifications to enable implementing and using the proposed R-D models. Finally, the importance of R-D models stems from their usefulness for different streaming systems. Because of the lack of quantitative performance analysis of different R-D models, it is currently difficult to select the model that best-suits a target streaming system.

In this article, we classify, analyze, and rigorously evaluate all R-D models proposed for FGS coders that we are aware of. We classify R-D models into three categories: analytic, empirical, and semi-analytic. Analytic models abstract the characteristics of an input video sequence by the probability distribution of its DCT coefficients. This distribution is then used in mathematical equations that describe the encoding/decoding processes. To derive the final R-D equation, typically several simplifying assumptions are made, which sometimes compromise the accuracy of the analytic models. Empirical R-D models directly measure the actual distortion by decoding the video sequence at many sampling bitrates. We call the third category semi-analytic models because they are *inspired* by analytic models, but they do not use mathematical derivations to develop R-D models. Rather, each semi-analytic model proposes a parametrized function that is thought to approximate the actual R-D function. The parametrized function takes the shape of an analytically derived function, but in a much simpler form. The parameters of the function are estimated using curve-fitting from a few actual rate-distortion data points.

We analyze the R-D models by following their mathematical derivations, scrutinizing the assumptions made, and explaining when the assumptions fail and why. In addition, we implement all R-D models, a total of eight, and evaluate them using a large and diverse set of carefully chosen video sequences. In our evaluation, we consider various source characteristics, diverse channel conditions, different encoding/decoding parameters, different frame types, and several performance metrics including accuracy, range of applicability, and time complexity of each model. We also present clear systematic ways (pseudo codes) for constructing various R-D models from a given video sequence. Finally, our experimental results enable us to provide guidelines on selecting the most suitable R-D model for a target streaming system. We present a justified list of recommendations on choosing the best R-D models for video-on-demand, video conferencing, real-time, and peer-to-peer streaming systems.

The rest of this article is organized as follows. In Section 2, we review previous works on R-D modeling. In Section 3, we review the main concepts of fine-grained scalable coders. We also describe how distortion is measured, and provide an overview on how analytic rate-distortion models are derived. Section 4 presents and analyzes various analytic R-D models, and Section 5 describes empirical and semi-analytic models. Our evaluation study is presented in Section 6. We conclude the paper in Section 7, where we also present a list of recommendations on choosing the best R-D model for a target streaming system.

2. RELATED WORK

Many R-D models have been proposed for non-scalable and layered coders. The pioneering work by Hang and Chen [1997] analyzes a simple coder that consists of a uniform quantizer followed by an ideal entropy coder. The authors propose a simple analytic R-D model, and a parameterized R-D model that considers some characteristics of real coders. Although this model provides useful theoretical insights, the experimental study in Lin and Ortega [1998] indicates that it is not very accurate. More recent R-D models use refinements based on observations from actual implementations. For instance, the ρ -domain model [He et al. 2001; He and Mitra 2001] leverages the high correlation between bitrates and percentage of nonzero-quantized coefficients.

To achieve higher accuracy, empirical R-D models have been proposed [Ding and Liu 1996; Chiang and Zhang 1997; Lin and Ortega 1998]. In Ding and Liu [1996], a rate-quantization model is proposed as $R(\Delta) = \alpha + (\beta/\Delta^\gamma)$, where α , β , and γ are estimated using actual measurements. In Lin and Ortega [1998], a piecewise cubic R-D model is proposed. The R-D model in Chiang and Zhang [1997] utilizes the high correlation of R-D curves among consecutive frames by using the quadratic relationship: $R(D) = aD^{-1} + bD^{-2}$, where a and b are estimated using empirical samples from previous frames.

All of the preceding models were proposed for non-scalable and layered coders that are quite different from the recent FGS coders. Various studies have shown that these models are not directly applicable to FGS coders. For example, Dai and Loguinov [2003] investigate the accuracy of the uniform quantization model of Hang and Chen [1997] and the quadratic model of Chiang and Zhang [1997] when applied to FGS-encoded sequences. The experiments indicate that these two models are not applicable to FGS coders. In addition, Zhang et al. [2003] experimentally show that the piecewise cubic model is inaccurate for FGS coders. Furthermore, because the ρ -domain model has no distortion-quantization model, applying ρ -domain model to FGS coders requires bitplane truncations and direct distortion computations. This process is similar to pure empirical approaches, which have high computational complexity. Therefore, the ρ -domain model may not be applicable to FGS coders.

Recently several studies have proposed R-D models for FGS coders. Most of these models are inspired by the R-D models of non-scalable and layered coders described above. For example, the square root model [Dai and Loguinov 2003] for FGS coders is a generalization of the R-D model in Hang and Chen [1997]. The linear rate-quantization function proposed in the ρ -domain model [He and Mitra 2001] is

employed in the FGS logarithm model [Dai et al. 2004]. The piecewise cubic model [Lin and Ortega 1998] is transformed into the FGS piecewise linear model [Zhang et al. 2003].

To the best of our knowledge, the performance and complexity of FGS R-D models have not been *rigorously* studied before. Only limited comparisons among some of the models were conducted. In Dai and Loguinov [2003] and Dai et al. [2004], the authors compared their new R-D models against non-scalable and layered R-D models for accuracy. They analyzed four sequences that are encoded with fixed coding parameters. More importantly, the comparison was conducted only at bitplane boundaries. Since FGS-encoded streams can be truncated at arbitrary bit positions, the accuracy of the R-D models across all bitrates should be considered. In Sun et al. [2005], the presented model was not compared against other models. Instead, different curve fitting schemes of the proposed model were compared. Similarly, the study in Zhang et al. [2003] only compared two alternatives: piecewise linear and piecewise exponential models.

We believe that our comparative study is rigorous because: (i) We compare the FGS R-D models against each other; (ii) We use a rich set of performance metrics for accuracy, applicable bitrate range, and time complexity; (iii) We choose wide ranges of sampling bitrates that match common usage patterns of various streaming applications; and (iv) We use a diverse set of video sequences encoded with various parameters.

3. BACKGROUND

In this section, we review the main concepts of fine grained scalable coders. We also describe how distortion is measured, and provide an overview on how analytic rate-distortion models are derived.

3.1 Fine Granularity Coding Systems

Video streaming over a dynamic and diverse environment like the Internet requires greater flexibility than that is provided by layered coding systems. Therefore, fine granularity scalability (FGS) has been designed to cover a wide range of bitrates at fine (bit level) steps. An FGS coding system encodes the video into two layers: a base layer and an enhancement layer. The base layer supports a single bitrate that delivers the basic video quality. The enhancement layer improves upon the quality of the base layer with a gain that is proportional to the number of *bits* received.

An FGS coding system may, in general, adopt any transform coders for its base layer and any FGS technique for its enhancement layer. Because of its simplicity and good coding efficiency [Li 2001; Radha et al. 2001], the bitplane coding defined in the MPEG-4 standard is the most recognized FGS technique nowadays. The basic idea of bitplane coding is to represent DCT coefficients as binary digits. Binary digits at the same significant position form a bitplane. Bitplanes are ordered by their significance, where a more significant bitplane contains more information per bit. Losing a bit in a higher significant bitplane results in a larger distortion than losing a bit in a lower significant bitplane. Each bitplane is separately run-length coded using the symbol format (RUN, LAST). These symbols are then coded by a variable-length coder (VLC) into a bitstream [ISO/IEC 14496-2 2004].

To illustrate how bitplane coding works, we present a simple example. This example also draws analogies between bitplane coding and uniform quantization. Both techniques are used to meet resource constraints by dropping less important details.

In bitplane coding, each coefficient is coded individually, therefore, we concentrate on a single coefficient for simplicity. Table I shows that we need $Z = 4$ bitplanes to code an original coefficient of value 15 (0x0F) in its binary format. Suppose the FGS decoder decodes $z \leq Z$ bitplanes. For example, if $z = 2$, the decoder decodes the first two significant bitplanes, and the reconstructed level is 12 (0x0C) instead of 15. Clearly, the reconstruction level is a function of z : We get no information if $z = 0$; and we get a perfect reconstruction if $z = Z$.

Table I. Bitplane Coding Example

DCT Coefficient	Original	Reconstructed				
	0x0F	0x00	0x08	0x0C	0x0E	0x0F
Bp #1	1	0	1	1	1	1
Bp #2	1	0	0	1	1	1
Bp #3	1	0	0	0	1	1
Bp #4	1	0	0	0	0	1
Quantization step, Δ		2^4	2^3	2^2	2^1	2^0
#decoded bitplanes, z		0	1	2	3	4

Notice that bitplane coding is essentially a quantizer with a uniform quantization step: $\Delta(z) = 2^{(Z-z)}$ (see the last two rows of Table I). This is because the information in the last $(Z - z)$ bitplanes is not seen by the decoder, as if this information were lost during quantization with a step of $2^{(Z-z)}$. We also observe that the reconstruction levels are defined by masking out the last $(Z - z)$ bitplanes. Hence, we can write the reconstruction level of an input coefficient x as:

$$L(x) = \begin{cases} \lfloor x/\Delta \rfloor \times \Delta, & x \geq 0; \\ \lceil x/\Delta \rceil \times \Delta, & x < 0, \end{cases} \quad (1)$$

where Δ is the quantization step.¹

FGS coding provides three major advantages to streaming applications: (1) better support for heterogeneous environments through bit-level scalability; (2) better error resiliency because partially corrupted bitstreams can be decoded; and (3) better streaming server scalability due to the separation of encoding and streaming processes.

3.2 Measuring Distortion and Rate-Distortion (R-D) Models

In the literature, distortion is commonly measured in terms of the mean square error (MSE) between the luminance values of pixels in the original and reconstructed frames. Chrominance components are usually ignored, because the human visual system is less sensitive to them compared to the luminance components, and they only occupy about 10% of the bitrate [Sullivan and Wiegand 1998]. A related quality measure is the Peak Signal-to-Noise Ratio (PSNR), which is given by: $\text{PSNR} = 10 \log_{10} \frac{255^2}{\text{MSE}}$ dB. PSNR is used in video quality comparison.

Distortion is clearly related to the quantization step Δ , because reconstruction levels depend on Δ according to (1). In addition, in the previous section we established the similarity between uniform quantization in traditional coding and bitplane coding. Hence, following Hang and Chen [1997], we can write the distortion (in MSE) as:

$$D(\Delta) = 2 \sum_{i=0}^N \int_{i\Delta}^{(i+1)\Delta} f(x)(x - i\Delta)^2 dx, \quad (2)$$

where $f(x)$ is a symmetric probability density function describing the distribution of the DCT coefficients. We call (2) the basic distortion-quantization (D-Q) function, because: (i) it relates the quantization step with the expected distortion, and (ii) it is typically the start point for developing analytic rate-distortion models.

R-D models are important for streaming applications that strive to optimize rendered quality in environments where channel conditions vary dynamically. R-D models are derived in three steps. First,

¹We note that Eq. (1) has been mentioned in the literature in different forms, such as Figure (5) in Dai et al. [2003] and Eq. (4) in Sun et al. [2005].

a probability density function $f(x)$ is assumed for the distribution of DCT coefficients. This density function is called the source model. The assumed density function is then substituted in (2), which yields—after approximation and manipulation—the D-Q function of the model. The second step in deriving R-D models is to derive a relationship between the bitrate and the quantization step. This is based on insights from the encoding/decoding processes. This relationship is called the R-Q function. Finally, the D-Q and R-Q functions are solved together to obtain the R-D model. We present and evaluate several R-D models in the next two sections.

4. ANALYTIC RATE-DISTORTION MODELS

Rate-distortion (R-D) models are functions that describe the relationship between bitrate and expected level of distortion in a reconstructed video stream. In this section, we present three analytic R-D models that are explicitly designed for the FGS enhancement layer. For each analytic model, we present the adopted source model. Then, we derive the R-D model in three steps: (i) we derive a distortion-quantization (D-Q) function, which relates the quantization step with the expected distortion; (ii) we derive a rate-quantization (R-Q) function, which relates the quantization step with the bitrate of the resulting bitstream; and (iii) we solve the D-Q and R-Q functions together either analytically or numerically to obtain the R-D function. Finally, we present our own implementation of each model. In Section 6, we experimentally evaluate and compare these R-D models.

4.1 Square Root Model

The square root model assumes that DCT coefficients follow a two-component Laplacian mixture (LM) distribution [Dai and Loguinov 2003]. A Laplacian mixture function is a linear combination of several Laplacian density functions. A two-component Laplacian mixture function is given as: $f_{LM}(x) = p_1 \frac{\lambda_1}{2} e^{-\lambda_1|x|} + p_2 \frac{\lambda_2}{2} e^{-\lambda_2|x|}$, where p_1 , p_2 and λ_1 , λ_2 are parameters that need to be estimated. Like other finite mixture distributions, the parameters of Laplacian mixture are not easily estimated using log-likelihood functions. Instead, researchers adopt numerical methods such as the expectation-maximization (EM) method Martinez and Martinez [2002, Section 8.4]. EM methods are widely used for: (1) incomplete observations, and (2) finite mixture models. We have designed and implemented an EM estimator based on the multidimensional estimator for Laplacian mixture distributions proposed in Park and Lee [2004]. The details of this estimator are given in Hsu and Hefeeda [2006b].

To derive a D-Q function, the square root model simplifies the basic D-Q function in (2) using the *high-resolution hypothesis*. The high-resolution hypothesis states that if bin sizes are adequately small relative to the variation rate of the density function, the density function—in each bin—can be approximated by a uniform density [Mallet and Falzon 1998]. This hypothesis leads to a new D-Q function:

$$D(\Delta) = 2 \sum_{m=0}^N \sum_{n=m\Delta}^{(m+1)\Delta-1} (n - m\Delta)^2 f_{LM}(n). \quad (3)$$

Unlike (2), the D-Q function in (3) uses a series of sub-bins to approximate the integral, and it assumes a constant density in each sub-bin. After substituting the LM source model, we get the D-Q function with the LM parameters p_1 , p_2 , λ_1 , and λ_2 .

To derive the R-Q function for this model, the bitrate is first expressed as a function of number of decoded bitplanes z . Then z is replaced by the quantization step Δ according to the relationship derived in the previous section: $\Delta = 2^{Z-z}$, where Z is the total number of bitplanes. The square root model proposes the following quadratic $R(z)$ function:

$$R(z) = c_1 z^2 + c_2 z + c_3, \quad (4)$$

```

SqrtModel
/* Inputs:
*   FGS-encoded frame  $f$ 
*   Frame period  $T$  (in sec)
*/
/* Output:
*   Square root R-D model of the frame, specified by  $p_1, p_2, \lambda_1, \lambda_2$  and  $c_1, c_2, c_3$ .
*/
1.  $\langle Z; l_1, l_2, \dots, l_Z \rangle = \text{extractBitplaneInfo}(f)$ ;
2.  $\mathbb{C} = \text{getDCTcoefficients}(f)$ ;
3. Use expectation-maximization method to estimate  $p_1, p_2, \lambda_1, \lambda_2$ 
   of the Laplacian mixture distribution using  $\mathbb{C}$ ;
4.  $R = 0; \mathbb{R}_z = \{(0, 0)\}$ ;
5. for  $i = 1$  to  $Z$  do
6.    $R = R + l_i/T$ ;
7.    $\mathbb{R}_z = \mathbb{R}_z \cup \{(R, i)\}$ ;
8. endfor
9. Curve-fit elements of  $\mathbb{R}_z$  to  $R(z) = c_1 z^2 + c_2 z + c_3$  to determine  $c_1, c_2, c_3$ ;

```

Fig. 1. Pseudocode for the square root R-D model.

where c_1, c_2, c_3 are parameters that need to be estimated. To justify this function, Dai et al. empirically show that second order polynomials are sufficient for typical $R(z)$ functions [Dai and Loguinov 2003]. Furthermore, they prove that an actual $R(z)$ function changes convexity up to once, thus a second order polynomial is enough [Dai et al. 2006]. The proof, however, assumes that DCT coefficients follow a Laplacian, not Laplacian mixture, distribution. To estimate c_1, c_2, c_3 , a set of (R, z) values is needed. This set is easily computed at bitplane boundaries: At each bitplane boundary $z = 1, 2, \dots, Z$, the bitrate is computed as $\sum_{i=1}^z l_i/T$, where l_i is the size of bitplane i and T is the frame period.

Given that $\Delta = 2^{Z-z}$, we get the R-Q function:

$$R(\Delta) = c_1(Z - \log_2 \Delta)^2 + c_2(Z - \log_2 \Delta) + c_3. \quad (5)$$

Computing Δ from (5) and substituting (5) in the D-Q function, we obtain the $D(R)$ function of the square root model. The resulting equation is fairly complicated, therefore, not shown here.

We have implemented the square root model as follows. First, we estimate the $p_1, p_2, \lambda_1, \lambda_2$ parameters of the Laplacian mixture density function from the DCT coefficients of the considered frame. We use an expectation-maximization estimator. Then, we extract the size of each bitplane in the enhancement layer of the frame. Using these sizes, we construct a set \mathbb{R}_z of (R, z) elements at bitplane boundaries as described above. Then, we curve-fit the elements of \mathbb{R}_z to (4) and obtain c_1, c_2, c_3 . The square root R-D model is completely specified once we determine the source model parameters $p_1, p_2, \lambda_1, \lambda_2$ and the constants c_1, c_2, c_3 . The pseudocode in Figure 1 summarizes our implementation of the square root model.

4.2 Logarithm Model

The logarithm model employs a Laplacian mixture (LM) source model [Dai et al. 2004], same as the square root model. We use the same parameter estimator used for the square root model.

To derive a D-Q function, the set of enhancement layer coefficients \mathbb{C} is partitioned into two subsets: \mathbb{C}_0 and $\mathbb{C}_{\bar{0}}$. \mathbb{C}_0 contains coefficients falling in the interval $(-\Delta, \Delta)$, while $\mathbb{C}_{\bar{0}}$ contains all other coefficients. Note that the bitplane coding quantizes all coefficients in \mathbb{C}_0 to the quantization index with zero reconstruction level. Hence, \mathbb{C}_0 coefficients are called zero-quantized. We define $N = |\mathbb{C}|$ and $M = |\mathbb{C}_{\bar{0}}|$ to be the number of total and zero-quantized coefficients, respectively.

Dai [2004] reveals that \mathbb{C}_0 coefficients are responsible for the majority of distortion at low and medium bitrates. Therefore, the logarithm model strives to accurately compute the distortion caused by \mathbb{C}_0 , and to roughly approximate the distortion related to $\mathbb{C}_{\bar{0}}$. We define D_0 and $D_{\bar{0}}$ as the distortion contributed by \mathbb{C}_0 and $\mathbb{C}_{\bar{0}}$, respectively. We separately analyze each of them as follows.

Since all coefficients in \mathbb{C}_0 have zero reconstruction level, D_0 is given by: $D_0(\Delta) = \sum_{c_i \in \mathbb{C}_0} |c_i|^2$. On the other hand, $D_{\bar{0}}$ is approximated as: $D_{\bar{0}}(\Delta) = M \Delta^2 / 12$, where $\Delta^2 / 12$ is the average distortion of a single coefficient if a uniform quantizer is employed under the high resolution hypothesis. The total distortion D is computed as the normalized sum of D_0 and $D_{\bar{0}}$:

$$D(\Delta) = \frac{D_0 + D_{\bar{0}}}{N} = \frac{1}{N} \sum_{c_i \in \mathbb{C}_0} |c_i|^2 + \frac{M}{N} \frac{\Delta^2}{12}. \quad (6)$$

Although the total number of coefficients N is constant, the number of coefficients in \mathbb{C}_0 is a function of Δ . We define ζ as the percentage of nonzero-quantized coefficients. That is, $\zeta = M/N$, and it is a function of Δ . The logarithm model utilizes a linear R-Q function that was proposed for non-scalable coders in He and Mitra [2002]. The applicability of this linear R-Q function to fine-grained scalable coders was validated through simulations in Dai et al. [2004]. The linear R-Q model states that the bitrate is a linear function of the percentage of nonzero-quantized coefficient, that is: $R = \gamma \zeta$, where γ is the slope parameter that needs to be estimated. To estimate γ , a set of (R, ζ) values is needed. This set is constructed at bitplane boundaries: At each bitplane boundary $z = 1, 2, \dots, Z$, the bitrate is computed as $\sum_{i=1}^z l_i / T$, and the percentage of nonzero-quantized coefficients is computed from the DCT coefficients using a quantization step of 2^{Z-z} .

The authors of the logarithm model proposed the following approximation [Dai et al. 2004]:

$$\begin{aligned} D(R) &= a_1 - (a_2 \log^2 R + a_3 \log R + a_4)R, \text{ where} \\ a_1 &= \frac{2p_1}{\lambda_1^2} + \frac{2p_2}{\lambda_2^2}, \\ a_2 &= \frac{11p_1}{12\lambda_1^2\gamma} + \frac{11p_2}{12\lambda_2^2\gamma}, \\ a_3 &= \frac{(-24 - 22 \log \gamma)p_1}{12\lambda_1^2\gamma} + \frac{(-24 - 22 \log \gamma)p_2}{12\lambda_2^2\gamma}, \\ a_4 &= \frac{p_1(24 + 24 \log \gamma + 11 \log^2 \gamma)}{12\lambda_1^2\gamma} + \frac{p_2(24 + 24 \log \gamma + 11 \log^2 \gamma)}{12\lambda_2^2\gamma}. \end{aligned} \quad (7)$$

We have implemented the logarithm model as follows. We first estimate the $p_1, p_2, \lambda_1, \lambda_2$ parameters of the Laplacian mixture density using an expectation-maximization method. Then, we extract the size of each bitplane in the enhancement layer of the frame. Using these sizes and the set of coefficients \mathbb{C} , we construct a set \mathbb{R}_ζ of (R, ζ) elements at bitplane boundaries as described above. Then, we curve-fit the elements of \mathbb{R}_ζ to $R = \gamma \zeta$ and obtain γ . Last, we compute a_1, a_2, a_3, a_4 using (7). The logarithm root R-D model is completely specified once we determine the constants a_1, a_2, a_3, a_4 . The pseudocode in Figure 2 summarizes our implementation of the logarithm model.

4.3 Generalized Gaussian Function Model

The generalized Gaussian function (GGF) model uses 64 distributions—of the same family but with different parameters—as the source model [Sun et al. 2005]. Each distribution models the set of coefficients of the same frequency component, and leads to an R-Q and a D-Q function for that frequency


```

LogModel
/* Inputs:
*   FGS-encoded frame  $f$ 
*   Frame period  $T$  (in sec)
*/
/* Output:
*   Logarithm R-D model of the frame, specified by  $a_1, a_2, a_3, a_4$ .
*/
1.  $\langle Z; l_1, l_2, \dots, l_Z \rangle = \text{extractBitplaneInfo}(f)$ ;
2.  $\mathbb{C} = \text{getDCTcoefficients}(f)$ ;
3. Use expectation-maximization method to estimate  $p_1, p_2, \lambda_1, \lambda_2$ 
   of the Laplacian mixture distribution using  $\mathbb{C}$ ;
4.  $R = 0; \mathbb{R}_\zeta = \{(0, 0)\}$ ;
5. for  $i = 1$  to  $Z$  do
6.    $R = R + l_i/T$ ;
7.    $\mathbb{C}_\zeta = \text{getQuantizedCoeff}(\mathbb{C}, 2^{Z-i})$ ;
8.    $\zeta = |\mathbb{C}_\zeta|/|\mathbb{C}|$ ;
9.    $\mathbb{R}_\zeta = \mathbb{R}_\zeta \cup \{(R, \zeta)\}$ ;
10. endfor
11. Curve-fit elements of  $\mathbb{R}_\zeta$  to  $R(\zeta) = \gamma\zeta$  to determine  $\gamma$ ;
12. Use Eq. (7) to compute  $a_1, a_2, a_3, a_4$ ;

```

Fig. 2. Pseudocode for the logarithm R-D model.

component. Recall that a frame is divided into 8x8 pixel blocks, and the DCT is applied on each block. This results in a set of 64 DCT coefficients per block, each belonging to a different frequency. Combining 64 R-Q functions gives a frame-level R-Q function, and aggregating 64 D-Q functions results in a frame-level D-Q function.

Zero-mean generalized Gaussian functions, $f_{GGF}(x)$, have been used to model DCT coefficients in the literature [Muller 1993]. The parameters of generalized Gaussian functions are often estimated using one of the two methods [Joshi and Fischer 1995]: (i) a maximum likelihood method, and (ii) a moment method. In the first method, a nonlinear likelihood equation was shown to be a maximum likelihood estimator if a unique root exists [Varanasi and Aazhang 1989]. Furthermore, this likelihood equation has a unique root when the sample size approaches infinity. When we implemented this method, though, the equation did not have a unique root in some cases. This is because of the relatively small sample size of DCT coefficients. In contrary, the moment method constantly results in a single estimation. In addition, experimental results show that it suffers only marginal loss of accuracy [Joshi and Fischer 1995]. Hence, we adopt the moment estimator for its simplicity and robustness. The details of this estimator are given in Hsu and Hefeeda [2006b].

To derive a D-Q model, the GGF model assumes a zero-mean source distribution (denoted by f_{ZM}) and the high-resolution hypothesis. Sun et al. substitute $f_{ZM}(x)$ for $f(x)$ in Equation (2); after manipulation and simplification, the D-Q model is given by [Sun et al. 2005]:

$$D(\Delta) = \Delta^2/3. \quad (8)$$

This $D(\Delta)$ is different from the classic uniform quantizer's $D(\Delta) = \Delta^2/12$ approximation, because of the different reconstruction levels: A bitplane coder uses the floor function rather than the bin midpoints for reconstructions (as discussed in Section 3.1).

To derive its R-Q function, the GGF model assumes an ideal variable-length coder (VLC) whose coding efficiency can arbitrarily approach the source entropy. Therefore, the GGF model employs the source

entropy to estimate the bitrate as:

$$R(\Delta) = - \int_{-\infty}^{\infty} f_{ZM}(x) \log_2 f_{ZM}(x) dx - \log_2 \Delta, \quad (9)$$

where the R is in bits per pixel. Notice that, this R-Q model imposes high computational complexity if the integral $\int_{-\infty}^{\infty} f_{ZM}(x) \log_2 f_{ZM}(x) dx$ is not analytically solvable. Unfortunately, replacing the generalized Gaussian function source model for f_{ZM} produces a function that requires a numerical integration.

The R-D function is derived by combining the D-Q model in (8) and the R-Q model in (9). We get $D_i(R_i)$ for component i as:

$$D_i(R_i) = 2^{-2R_i-2} \int_{-\infty}^{\infty} f_{GGF}(x) \log_2 f_{GGF}(x) dx / 3. \quad (10)$$

Next, we consider the aggregation of these 64 $D_i(R_i)$ functions for a frame-level $D(R)$ model. Note that a given bitrate R is achieved by truncating a coded stream. This truncation decides an R_i for each component i ($i = 1, 2, \dots, 64$). These 64 components, however, are not independently coded. Therefore, decomposing R into R_i is nontrivial. To address this difficulty, Sun et al. [2005] propose an exhaustive search on the quantization step Δ from 1 to 64. For each Δ , a frame-level R-D mapping is computed by composing 64 D_i 's (given by (8)) and 64 R_i 's (given by (9)). We follow this approach to compute R-D estimates for Δ 's. We then construct a piecewise linear R-D model using these estimations to cover intermediate Δ for completeness.

We have implemented the GGF model as follows. We first estimate the parameters of the generalized Gaussian function for frequency component i ($i = 1, 2, \dots, 64$) using a moment estimator. Then, we compute a set \mathbb{D}_R of (D, R) elements at integer quantization steps Δ from 1 to 64 as described above. The GGF R-D model is completely specified once we construct a piecewise linear function using the elements of \mathbb{D}_R . Because of the space limitations, we give the pseudocode of the GGF model in Hsu and Hefeeda [2006a].

5. EMPIRICAL AND SEMI-ANALYTIC MODELS

In this section, we present two empirical methods: (i) a pure empirical model for determining accurate R-D functions which we use for comparisons, and (ii) an efficient piecewise linear model. Then, we present three semi-analytic models that are inspired by the three analytic models described in the previous section. Semi-analytic models do not use mathematical derivations to develop R-D models. Rather, each semi-analytic model proposes a parametrized function that is thought to approximate the actual R-D function. The parametrized function takes the shape of an analytically derived function, but in a much simpler form. The parameters of the function are estimated using curve-fitting from a few actual rate-distortion data points.

5.1 Empirical Model

The pure empirical approach is a brute force method to derive R-D functions. It chooses many sampling bitrates and decodes the video sequence at each of them. Then, for each sample bitrate, it computes the distortion as the difference between the original and reconstructed video sequences. Interpolation is used to extend the discrete R-D mapping into a continuous function. The empirical approach requires a tremendous amount of computational power and storage space. Nevertheless, it produces accurate R-D curves. We use results produced by the empirical approach as the baseline for comparing various R-D models. In our implementation of this empirical model, we choose six sampling bitrates in each bitplane, including the bitplane boundaries. And we compute the distortion at these bitrates.

We note that constructing empirical R-D models for FGS-encoded sequences does *not* require reencoding of the sequence. We only need to truncate the bitstream at appropriate bit positions. This is

in contrast to constructing empirical R-D models for non-scalable or layered-scalable sequences, which requires reencoding a video sequence at each sampling bitrate. Reencoding is time consuming because coding systems employ many optimization procedures at the encoding time. Therefore, while empirical R-D models for FGS-encoded sequences are expensive to construct, they are much less expensive and more feasible than empirical R-D models for non-scalable or layered-scalable sequences [Zhang et al. 2003].

5.2 Piecewise Linear Model

Traditional empirical R-D models often employ exponential interpolation between sampling points. Zhang et al. [2003] found that exponential interpolations do not accurately track the actual R-D curves of FGS-encoded sequences. The experiments in Zhang et al. [2003] reveal that when sampling bitrates are located in the same bitplane, using linear interpolation produces smaller deviation than using exponential interpolation. However, if the sampling bitrates are from different bitplanes, applying exponential interpolation is more accurate. To implement the piecewise linear model, we choose sampling points at all bitplane boundaries and compute the distortion at these points. Then we connect these points with line segments.

5.3 Semi-Analytic Square Root Model

This model is based on the analytic square root model in Section 4.1. However, instead of using the complicated D-Q function, it uses the following *heuristic* function:

$$PSNR(z) = d_1 z^2 + d_2 z + d_3, \quad (11)$$

where PSNR is the quality, not the distortion, and d_1, d_2, d_3 are parameters that need to be estimated. This function is justified by the fact that the quality (in terms of PSNR) is a monotonically increasing function of the number of transmitted bitplanes. Furthermore, it has been shown in Dai et al. [2006] that this function does not change its convexity more than once. Hence, a quadratic function is a good approximation for it.

Combining this simplified D-Q function with the $R(z)$ function in (4) yields what we call the semi-analytic square root (sSqrt) R-D model:

$$PSNR(R) = c_1 R + c_2 \sqrt{R} + c_3, \quad (12)$$

where c_1, c_2 need to be estimated from empirical R-D samples, and $c_3 = 10 \log_{10}(255^2/\sigma^2)$ for a source with variance σ^2 . To implement this model, we determine the bitrate and PSNR values at all bitplane boundaries. Then, we do curve fitting to obtain c_1, c_2, c_3 .

5.4 Semi-Analytic Logarithm Model

This model is based on the analytic logarithm model in Section 4.2. However, instead of using parameters of the Laplacian mixture density function, this model employs curve fitting to compute a_1, a_2, a_3, a_4 in (7). This leads to what we call the semi-analytic logarithm (sLog) R-D model. To implement this model, we determine the bitrate and distortion values at all bitplane boundaries. Then we do curve fitting to get a_1, a_2, a_3, a_4 .

5.5 Semi-Analytic Generalized Gaussian Function Model

This model is based on the analytic GGF model in Section 4.3. However, instead of numerically computing the integration in (9), this model utilizes a *heuristic* function [Sun et al. 2005]:

$$PSNR(R) = g_1 R + g_2 - \frac{g_2 - D_b}{1 + g_3 R}, \quad (13)$$

Table II. List of Test Video Sequences

Class	Sequences	Resolution	# of Frames
Low	<i>Akiyo, Mother, Foreman, Paris, Bridge-close, Bridge-far, Hall-monitor, Highway, Tempete</i>	352×288	300
	<i>Claire, Salesman, Grandma, Carphone</i>	176×144	300
Medium	<i>Bus</i>	352×288	150
	<i>Mobile, Container</i>	352×288	300
	<i>Coastguard</i>	352×288	300
High	<i>Garden</i>	352×240	115
	<i>Tennis</i>	352×240	112
	<i>Football</i>	352×240	125

where PSNR is the quality, D_b is the base layer only quality, and g_1, g_2, g_3 are parameters that need to be estimated. This function was proposed after empirically observing the actual R-D data and the R-D estimations produced by the analytic GGF model. To implement this model, we determine the bitrate and PSNR values at all bitplane boundaries. Then, we do curve fitting to derive g_1, g_2, g_3 .

6. EVALUATION OF RATE-DISTORTION MODELS

In this section, we present an extensive experimental study to evaluate the performance of the R-D models described in the previous sections. We first describe the video sequences used in the experiments and why we chose them. Then, we present our experimental setup and the performance metrics considered. Then, we present the results of comparing analytic models, followed by results of comparing semi-analytic models.

6.1 Selection of Test Video Sequences

Choosing a representative set of video sequences is a critical step in evaluating and analyzing the performance of R-D models. A homogeneous set of sequences may produce biased comparison results, because some models may perform exceptionally well under certain sequences.

We use two key features to characterize video sequences: spatial complexity and temporal complexity. To quantify these complexities, we adopt the neighborhood difference metric [Adjeroh and Lee 2004]. This metric captures the amount of variations between a block and its neighbor by accumulating the differences among corresponding transform coefficients in the two blocks. The neighborhood metric approaches 0 when neighboring blocks are similar, and increases as the similarity between the two blocks decreases. The spatial complexity is measured by averaging all neighborhood differences in the same frame, whereas the temporal complexity is measured by averaging neighborhood differences between adjacent frames.

To form a set of test sequences, we considered twenty video sequences from various sources [Video Traces Research Group 2006; Center for Image Processing Research 2006]. Table II lists the considered sequences. We then classify sequences based on their spatial and temporal complexities. Sequences that have high spatial and temporal complexities are what we call high-complexity sequences. In contrast, sequences with low spatial and temporal complexities are called low-complexity sequences. Sequences that have either high temporal or spatial complexity, but not both, are called medium-complexity sequences.

Out of these twenty sequences, we chose seven representative sequences with different complexities: *Akiyo, Mother, Foreman, Mobile, Bus, Garden, and Football*. We believe that these seven test sequences form a diverse enough set to examine the performance of R-D models.

1. Choose a sequence from the test sequences (described in Section 6.1);
2. FGS-encode the sequence with base layer rates R_b at 8, 16, 32, 64, and 128 Kbps;
3. **for** each R-D model $model \in \{Sqrt, Log, GGF, Linear, sSqrt, sLog, sGGF\}$ **do**
4. **for** each base layer rate R_b **do**
5. **for** each frame f in the sequence **do**
6. Construct the set \mathbb{R}_s of sampling bitrates by choosing equally-spaced K samples from each bitplane;
7. **for** each sampling rate $R_s \in \mathbb{R}_s$ **do**
8. Measure actual distortion $D^{Emp}(R_s)$ by comparing reconstructed and original frames;
9. Estimate distortion using the R-D model $D^{model}(R_s)$;
10. Compute the distortion deviation as: $\varepsilon_f(R_s) = |D^{emp}(R_s) - D^{model}(R_s)|$;
11. **endfor**
12. Compute the average distortion deviation per frame;
13. Maintain the average distortion deviation for each frame type: I, P, B;
14. **endfor**
15. Compute the average and maximum distortion deviations across all frames;
16. Measure the total running time;
17. Compute the normalized range of applicability;
18. **endfor**
19. **endfor**
20. Repeat steps 1—19 for another video sequence;

Fig. 3. Our procedure for rigorously evaluating the performance of various R-D models.

6.2 Experimental Setup and Performance Metrics

6.2.1 Software Used and Implemented. We use the MPEG-4 Reference Software Version 2.5 [ISO/IEC 14496-5 2004] developed by Microsoft as an experimental package for the MPEG-4 standard. It is implemented in C++ and contains three major executables: *encoder*, *decoder*, and *fgs_server*. The *encoder* is a configurable MPEG-4 encoder that can compress a raw video file into two bitstreams: base layer and enhancement layer. Each bitstream is stored in a separate file. The *decoder* is FGS-enabled, that is, it can process an incomplete enhancement layer and produces a raw video file with proportional quality improvements. The *fgs_server* is a utility to trim the enhancement layer according to a given target bitrate. It does so by calculating the number of allowed bits in each frame at the target bitrate. These bits, which represent the truncated bitstream, are then saved in a new file. We have instrumented the reference software to extract various characteristics of an input video sequence. For instance, we collect transform coefficients, number of bitplanes, and size of each bitplane in the enhancement layer. This information is used to estimate parameters of the R-D models.

We have implemented all R-D models described in this paper, a total of eight models. Specifically, we have implemented the three analytic R-D models: square root (denoted by Sqrt in the plots), logarithm (Log), and generalized Gaussian (GGF_F); the pure empirical model (Emp), the piecewise linear (Linear), and the three semi-analytic R-D models: semi-analytic square root (sSqrt), semi-analytic logarithm (sLog), and semi-analytic generalized Gaussian ($sGGF_F$). All models are implemented in Matlab.

6.2.2 Selection of Sampling Bitrates. We evaluate the R-D models across a wide range of bitrates, not only at bitplane boundaries. This is important because streaming applications using FGS-encoded sequences are allowed to truncate bitstreams at arbitrary *bit* positions. It is also critical to emphasize on the range of bitrates that most applications will heavily use. To achieve these two goals, we propose the following sampling scheme. Suppose Z is the total number of bitplanes, l_z is the size of bitplane z ($z = 1, 2, \dots, Z$), and T is the frame period. We choose K equally-spaced bitrates in each bitplane.

Specifically, we form the set of sampling bitrates \mathbb{R}_s as

$$\mathbb{R}_s = \{0\} \cup \{R_s^{z,k} \mid z = 1, 2, \dots, Z \text{ and } k = 1, 2, \dots, K\}, \quad (14)$$

where $R_s^{z,k} = \frac{\sum_{u=1}^{z-1} l_u + (l_z/K)(k)}{T}$.

The set \mathbb{R}_s achieves the above two goals, because it covers the whole bitrate range from base layer only ($R_s = 0$) to full quality ($R_s = R_s^{Z,K}$). Furthermore, since the most significant bitplanes are typically small in size, gaps between sampling bitrates will be small. Therefore, the accuracy of the considered R-D model will be higher in these significant bitplanes. The R-D performance of significant bitplanes is crucial for two reasons. First, significant bitplanes have high per-bit reduction rates in distortion. This implies more sampling bitrates are needed to capture the rapid distortion variations. Second, significant bitplanes are streamed more often owing to the embedded property of FGS-encoded streams. The embedded property states that a lower-rate bitstream is always a prefix of a higher-rate one.

6.2.3 Performance Metrics. We consider the following performance metrics for each R-D model: accuracy, range of applicability, and time complexity.

We consider model accuracy at different granularities: for individual sampling bitrates, for every frame, and for the whole sequence. We measure accuracy as the difference between the quality (in PSNR) predicted by the considered R-D model and the actual quality at a given bitrate. Specifically, for every sampling bitrate $R_s \in \mathbb{R}_s$, we compute the absolute distortion deviation (or error) as:

$$\varepsilon_f(R_s) = |D^{emp}(R_s) - D^{model}(R_s)|, \quad (15)$$

where $D^{emp}(R_s)$ is the actual distortion computed by comparing the reconstructed and original frames, and $D^{model}(R_s)$ is the distortion estimated by the R-D model for frame f at bitrate R_s . For the frame level, we compute the average distortion deviation over all sampling bitrates. We also compute the average distortion deviation for different frame types (I, P, and B). For the sequence level, we compute the average and maximum deviations across all frames.

Ideally, an R-D model should provide valid distortion estimations for all sampling bitrates R_s . Unfortunately, some R-D models only support a subset of the whole bitrate range. To capture this important aspect, we use the range of applicability metric. The lower bound of this range is set to 0 and the upper bound is set as the smallest bitrate beyond which the R-D model fails to provide valid—finite and nonnegative—results. To enable cross-sequence comparisons, we normalize the range of applicability by dividing it by the maximum bitrate of the sequence. The maximum bitrate of the sequence occurs when all bitplanes are transmitted.

Time complexity is also an important metric especially for streaming applications that use R-D models in real-time, for example, video conferencing. We measure the running time to construct R-D models. To mitigate the effect of clock precisions and interruptions from the operating system, we measure the running time for all frames in a video sequence and report the average per-frame running time. All experiments are conducted on a 2.8 GHz Pentium 4 workstation.

6.2.4 Evaluation Procedure. We evaluate the R-D models on the set of video sequences described in Section 6.1. We encode each sequence at 30 fps with the following structure: 6 P-frames between I-frames and 2 B-frames between successive P-frames. We FGS-encode every video sequence at several different bitrates for the base layer. The rate of the base layer directly impacts the number and size of bitplanes in the enhancement layer. Therefore, we are effectively evaluating the models under different relative sizes of the base and enhancement layers. This is important because streaming systems work in different network environments, and they typically adjust the base layer rate based on these conditions. In addition, for every video sequence encoded at a given base layer rate, we decode the enhancement layer at different rates. This is done by truncating the bitstream at the appropriate bit positions.

Decoding at different rates captures the heterogeneous nature of clients and the diverse communication channels over which they receive video sequences. Figure 4 summarizes the procedure that we follow in evaluating the performance of R-D models.

We believe that our experimental set up and evaluation procedure rigorously evaluate the R-D models, because they account for various source characteristics, diverse communication channel conditions, and different encoding/decoding parameters.

6.3 Results for Analytic Models

We first discuss accuracy of the analytic R-D models, followed by their applicable ranges and time complexities.

6.3.1 Accuracy. We plot the frame-level average distortion deviation in Figure 4. We present results of a sample sequence in each complexity class (from left to right: low, medium, and high complexity) in each vertical column of subfigures. Each horizontal row represents a different base layer rate R_b . We make a few observations on this figure. First, the GGF model produces the highest distortion deviation, at least 4 dB, in all cases. Second, the logarithm model performs better than the square root model in low complexity sequences. Furthermore, in medium and high complexity sequences, the logarithm model still outperforms the square root model for base layer bitrate $R_b > 32$ Kbps. These two models have comparable accuracy for $R_b = 32$ Kbps; the square root model performance is improved for $R_b < 32$ Kbps.

We also study the accuracy for different frame types. Plots of the average distortion deviation for I-, P-, and B-frames are given in the technical report Hsu and Hefeeda [2006a] because of space limitations. The results show the same relative accuracy. This indicates the frame type does not affect relative accuracy of these models.

We show the sequence-level average and maximum distortion deviations for different sequences that are coded at different base layer rates R_b in Figure 5 (more figures are given in Hsu and Hefeeda [2006a]), where the sequences are sorted in ascending order of their complexities. We observe that the logarithm model is the most accurate in terms of the average distortion deviation with at most 3 dB deviation, while the square root model produces as high as 6 dB and the GGF model results in as high as 7 dB. We note that, the GGF model constantly results in the highest deviation in all sequences. Comparing the square root model versus the logarithm model, we observe that: (i) The square root model results in higher distortion deviations for lower complexity sequences as shown in Figures 5(a) and 5(b). This is because, the high-resolution hypothesis introduces higher approximation errors when the variation rate of the density function is large. The DCT density function has high variability in two cases: (1) when the base layer is encoded at high bitrate, and (2) when the video sequence has low temporal and spatial complexity. Note that, although the logarithm model also assumes high-resolution hypothesis, the hypothesis is not applied to the zero-bin, which is the quantization bin with the largest density function variation rate. This explains why the accuracy of the square root model was worse than the logarithm model in both cases. (ii) The logarithm model produces higher distortion deviations for higher complexity sequences that are coded with low base layer rates R_b as shown in in Figure 5(b). We will further discuss this when analyzing R-D curves of each frame.

Then, we study the shape of R-D curves for a few frames. We can draw the following observations from Figure 6(a) (and from similar plots in Hsu and Hefeeda [2006a]):

—*The GGF model suffers very high deviations at low bitrates.* There are two main reasons behind these deviations: (i) The GGF model uses the high resolution hypothesis that does not hold when the quantization step is large; and (ii) The GGF model utilizes a very generic source model—a zero-mean source distribution. Note that the logarithm model does not suffer deviations at low bitrates because it only uses the high resolution hypothesis in nonzero bins where the variation in the density function

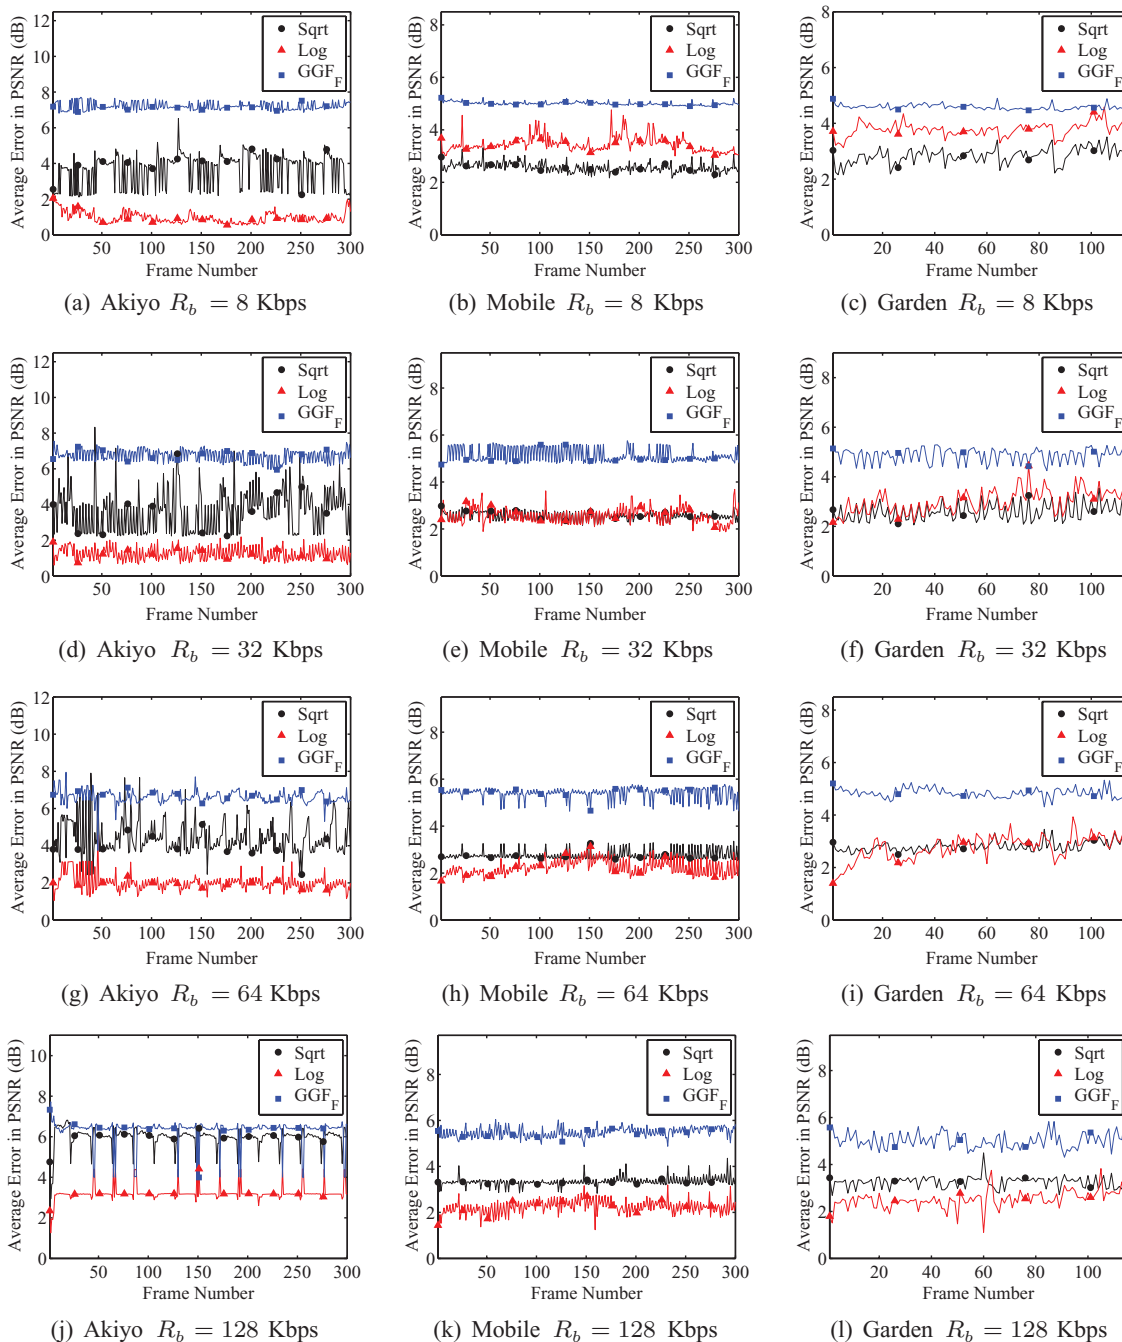


Fig. 4. The average distortion deviation of the R-D models across all frames of three video sequences of different complexities, where the base layers are coded with different bitrates.

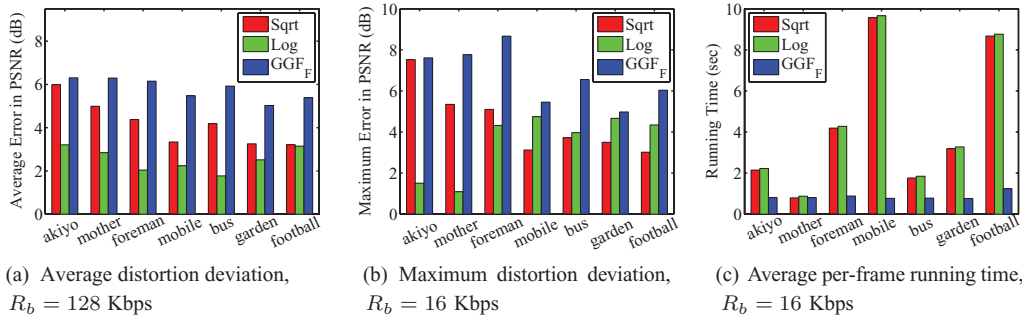


Fig. 5. The sequence-level performance of each R-D model across all frames of all video sequences of different complexities. Samples of (a) average deviation distortion, (b) maximum deviation distortion, and (c) average running time are shown.

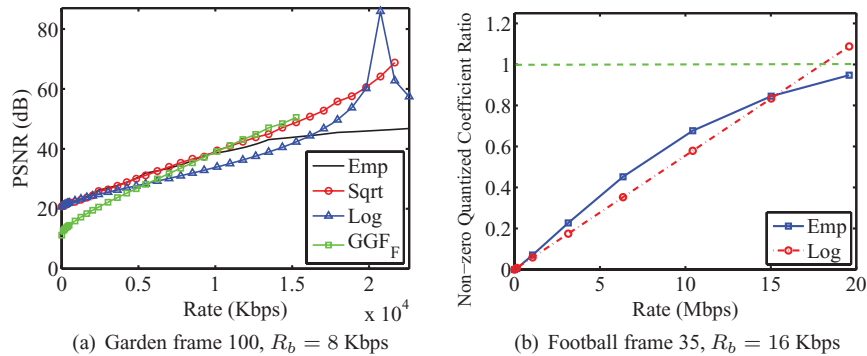


Fig. 6. (a) Sample R-D curves for four models applied to video sequences of different complexities; (b) The ζ -R curve shows strong linearity in the first four bitplanes, but not in other bitplanes. This inaccuracy results in erroneous R-D estimations.

is not significant, thus the assumption is relatively valid. In addition, while the square root model also uses the high resolution hypothesis, it employs a more elaborate source model—the Laplacian mixture distribution—than the simple source model adopted by the GGF model. This explains why the square root model performs better than the GGF model at low bitrates.

- The square root model over estimates the quality at high bitrates.* This can be explained by having $D(\Delta) = 0$ when $\Delta = 1$ in Equation (3). Consequently, the produced PSNR approaches infinity as Δ goes to 1. This results in exaggerated estimations of the video quality at high bitrates, which may mislead streaming applications using the square root model. However, since the human visual system is not sensitive to further quality improvement beyond a high PSNR level, applications may use a minimum MSE threshold with the square root model to avoid the misleading estimations.
- The logarithm model produces inaccurate quality for high-complexity sequences at high bitrate.* It even produces nondecreasing R-D curves in some cases, a sample result is shown in Figure 6(a). This is inaccurate because higher bitrates should always lead to better quality. We investigate the root cause as follows. The logarithm model utilizes a linear R-Q model which assumes that the percentage of nonzero quantized coefficients ζ is a linear function of the bitrate. Through extensive experiments, we find that this linearity is strong in the more significant bitplanes, especially in the first four bitplanes. Since low complexity sequences only have few bitplanes, the linear relationship is strong for these sequences. In contrast, for high complexity sequences coded at low base layer rates, this linearity starts to break. We show sample results in Figure 6(b), a more complete set of plots are

given in [Hsu and Hefeeda 2006a]. Figure 6(b) explains the abnormal R-D curves produced by the logarithm model at high bitrates observed in Figure 6(a).

In summary, the GGF model is the least accurate in all cases, while the logarithm model is the most accurate except for medium or high complexity sequences that are coded at base layer rates less than 32 Kbps. The square root model is not accurate for low complexity sequences coded at high base layer bitrates, because these coded streams: (i) have source distributions with large variation rates that contradict the high frequency hypothesis, and (ii) have a few bitplanes, therefore, the estimation inaccuracy in the least significant bitplanes dominates the overall performance. The performance of the logarithm model suffers in high complexity sequences coded at low base layer bitrates, because the R-Q relationship of these code streams deviates from the assumed linear R-Q model.

6.3.2 Range of Applicability. Plots of the range of applicability are given in Hsu and Hefeeda [2006a] due to space limitations. The results show that:

- The logarithm model supports all bitrates—from 0 to the full quality.
- The square root model is less applicable in the least significant bitplanes. The square root model produces PSNR that approaches infinity as Δ goes to 1. Thus, we cannot use it when Δ is close to 1, that is, near the full bitrate of the enhancement layer. For medium- and high-complexity sequences, the range of applicability is about 95% of the bitrate. However, this range is smaller (about 70%) for low-complexity sequences because the sequences have fewer bitplanes.
- The GGF model’s applicability is no more than 70% in all cases. Its R-Q function (in Equation (9)) defines the bitrate as a function of the quantization step, which is an integer value. Since the least quantization step is 1, the GGF model has an implicit maximum bitrate: $R(1)$. Note that, an FGS coder has lower coding efficiency than an ideal entropy coder. Hence, bitrate $R(1)$ is lower than the full quality bitrate. This leads to a range of bitrates that are not supported by the GGF model.

6.3.3 Time Complexity. We show a sample of the average running time for each frame in Figure 5(c) (more plots are given in Hsu and Hefeeda [2006a]). This figure shows that: (i) The GGF model is the most efficient; and (ii) The square root model is slightly faster than the logarithm model, because the latter requires a time-intensive estimation of its R-Q model parameter γ .

6.4 Results for Empirical and Semi-Analytic Models

We present and analyze the results for empirical and semi-analytic models. We consider one empirical model—the piecewise linear (Linear) model—and three semi-analytic models: the square root (sSqrt), the logarithm (sLog), and the GGF (sGGF_F). We implemented these models using the nonlinear least-squares fitting subroutine provided by Matlab.

6.4.1 Accuracy. Figure 7 shows a sample of the frame-level average distortion deviation of these models. We observe that: (i) the piecewise linear model produces negligible, always less than 0.2 dB, error in all cases; (ii) the square root model has up to 1 dB deviation and is more accurate than the GGF and the logarithm model; and (iii) the logarithm model results in the highest estimation error in almost all cases.

We also notice that the semi-analytic models result in more fluctuations of frame-level average error. This is because the nonlinear optimization methods may not lead to a global optimal solution, and sometimes they do not even terminate. Therefore, the methods are often associated with a maximum number of iterations. We set the iteration limit to be 1000 times. In contrast, the piecewise linear model is derived using a simple arithmetic—connecting two points with a line. This is not only more robust, but also much more efficient.

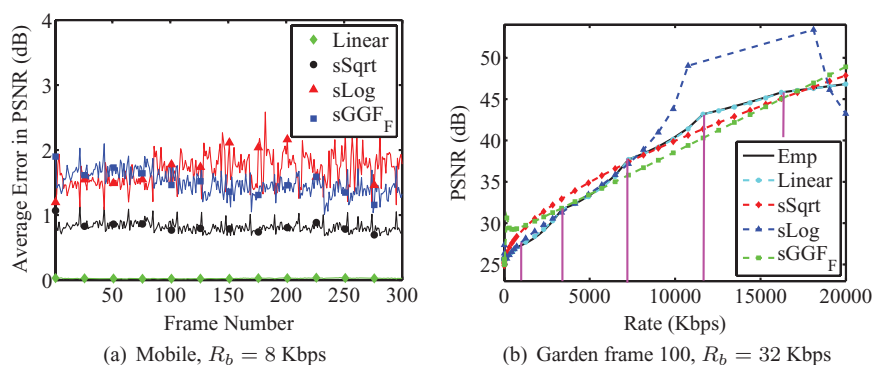


Fig. 7. Samples of the performance of the empirical and semi-analytic R-D models. (a) frame-level average distortion deviation; (b) frame-level R-D curve, where vertical lines indicate the bitplane boundaries.

Next, we study the shape of R-D curves in Figure 7(b). We see abnormal R-D curves caused by the imperfect curve fitting: for instance, the semi-analytic logarithm model produces a decreasing R-D segment in Figure 7(b). In addition, Figure 7(b) illustrates that while other models produce accurate estimations only at bitplane boundaries, the piecewise linear model results in accurate estimations in the entire bitplane. This explains the highest accuracy provided by the piecewise linear model.

6.4.2 Time Complexity. Our experiments [Hsu and Hefeeda 2006a] imply that while the piecewise linear model has negligible time complexity, the semi-analytic models require at least tens of milliseconds to terminate. We find that the logarithm model requires the longest running time, up to 600 milliseconds in the worst case, because of its high nonlinearity. The square root is the most efficient semi-analytic model.

7. CONCLUSIONS AND RECOMMENDATIONS FOR CHOOSING R-D MODELS

7.1 Conclusions

We investigated various rate-distortion (R-D) models for fine-grained scalable video sequences. We classified R-D models into three categories: analytic, empirical, and analytically-inspired empirical (or semi-analytic). We have analyzed the three analytical R-D models known in the literature: square-root (Sqrt), logarithm (Log), and generalized Gaussian function (GGF_F). We also presented three semi-analytic models based on the above analytic models: sSqrt, sLog, and (sGGF_F). In addition, we discussed how a pure empirical model can be implemented by measuring the distortion at equally-spaced sampling bitrates. Finally, we presented the piecewise linear model, which improves upon the pure empirical model by choosing sampling bitrates at bitplane boundaries and connecting these samples by line segments.

We presented systematic ways (pseudo codes) for constructing the R-D models from a given video sequence. We implemented all of the above eight R-D models in Matlab, and we evaluated them using a large and diverse set of carefully-chosen video sequences. In our evaluation process, we considered various source characteristics, diverse channel conditions, different encoding/decoding parameters, different frame types, and several performance metrics including accuracy, range of applicability, and time complexity of each model.

The findings from our extensive experimental study can be summarized as follows:

- (1) The GGF model shows large deviations from actual distortions. It also does not produce valid results at high bitrates. Therefore, we believe it will be of little use in practice.

- (2) If the base layer is encoded at high bitrate (>32 Kbps), the logarithm model produces the most accurate distortion estimations.
- (3) If the base layer is encoded at lower bitrates (<32 Kbps) *and* the video sequence has high or medium complexity, the square root model is more accurate than the logarithm model. However, if the sequence has low complexity, the logarithm model produces better results with any base layer rate.
- (4) The reasons behind (2) and (3) can be explained as follows. If the base layer rate is high or the sequence has low complexity, fewer bitplanes will exist in the enhancement layer. In this case, the linear relationship between the bitrate and percentage of nonzero-quantized coefficients is accurate as indicated by our results. This relationship is employed by the logarithm model; it is why it produces good results in this case. On the other hand, the square root model produces higher deviations in the least-significant bitplanes. And since we have a few total bitplanes, the impact of these high deviations will dominate, and the performance of the square root model will suffer.
For high and medium complexity sequences and when the base layer rate is small, the converse of the above argument applies. That is, more bitplanes exist in the enhancement layer and the linear relationship used by the logarithm model becomes less accurate, reducing the accuracy of the logarithm model. At the same time, the effect of the deviations in the least-significant bitplanes of the square root model decreases as the number of bitplanes increases, making the square root model more accurate.
- (5) Regarding the range of applicability, the logarithm model always produces valid results at all possible bitrates. This is not the case for the square root model. As we explained in Section 4.1, the square root model greatly overestimates the quality for small quantization steps, approaching infinity for a quantization step of 1. Small quantization steps correspond to high bitrates. This overestimation of quality limits the applicability of the square root model to about 70% of the possible bitrates for low complexity sequences.
- (6) The time complexity of the square root model is smaller than that of the logarithm model. The logarithm model is slow because estimating the slope parameter of the linear relationship between the bitrate and percentage of nonzero-quantized coefficients is expensive. This is because the estimation process requires bitplane truncations and computing the corresponding percentage of the truncated coefficients.
- (7) In general, the empirical and semi-analytic R-D models are more accurate than the analytic R-D models.
- (8) The curve-fitting methods employed by the semi-analytic models are time-consuming and may occasionally result in inaccurate estimations for the model parameters.
- (9) Finally, the piecewise linear model is simple and fairly accurate, when compared to all other models.

7.2 Recommendations for Choosing R-D models

Our experimental results enable us to provide guidelines on selecting the most suitable R-D model for a target streaming system. We summarize our recommendations in the following.

—*Streaming Stored Video Sequences and Video-on-Demand Systems.* In these systems, the R-D model can be computed off-line and it is desired to be as accurate as possible. This is because a video sequence is usually streamed many times to clients with heterogeneous network connections. This implies that the cost of building an accurate R-D model will be amortized over many sessions. Therefore, we recommend using the piecewise linear model, where the model parameters will be computed once and stored in a meta file with every video sequence.

- Streaming in Real-Time and Video Conferencing*. Timing is critical in this case. In addition, real-time streaming usually happens for sports events, which have high complexity. Therefore, we recommend using the square root model because it has low time complexity and it is fairly accurate for medium and high complexity sequences. The square root model can further be accelerated by adopting a Laplacian, instead of Laplacian mixture, density function for modeling the DCT coefficients. Estimating the parameters of Laplacian density functions is straightforward and much simpler than the expectation-maximization method used with Laplacian mixture density functions. With Laplacian density, the square root R-D model is at least an order of magnitude faster than with Laplacian mixture density, while the accuracy of the resulting R-D model does not suffer much.
- Broadcasting News and Low-Complexity Video Clips*. News clips usually have low complexity. We recommend using the logarithm model in this case because it more accurate than the square root model. The cost of constructing a piecewise linear model may not be justified in this case because news clips are usually broadcast a few times.
- Streaming in Distributed and Peer-to-Peer Environments*. In these systems, there are potentially many servers. It would be very costly if we let every server independently construct the R-D model for the same video sequence. Also, video sequences are streamed many times. Therefore, we recommend that only one server, for example, the creator or the introducer of the sequence, builds a piecewise linear model for the sequence and attach it as a meta file with video data.

REFERENCES

- ADJEROH, D. AND LEE, M. 2004. Scene-adaptive transform domain video partitioning. *IEEE Trans. Multimedia* 6, 1, 58–69.
- CENTER FOR IMAGE PROCESSING RESEARCH. 2006. <http://www.cipr.rpi.edu/resource>.
- CHIANG, T. AND ZHANG, Y. 1997. A new rate control scheme using quadratic rate distortion model. *IEEE Trans. Circ. Syst. Video Techn.* 7, 1, 246–250.
- DAI, M. 2004. Rate-distortion analysis and traffic modeling of scalable video coders. Ph.D. thesis, Department of Electrical Engineering, Texas A&M University.
- DAI, M. AND LOGUINOV, D. 2003. Analysis of rate-distortion functions and congestion control in scalable Internet video streaming. In *Proceedings of ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'03)*. Monterey, CA.
- DAI, M., LOGUINOV, D., AND RADHA, H. 2003. Statistical analysis and distortion modeling of MPEG-4 FGS. In *Proceedings of IEEE International Conference on Image Processing (ICIP'03)*. Barcelona, Spain.
- DAI, M., LOGUINOV, D., AND RADHA, H. 2004. Rate-distortion modeling of scalable video coders. In *Proceedings of IEEE International Conference on Image Processing (ICIP'04)*. Singapore.
- DAI, M., LOGUINOV, D., AND RADHA, H. 2006. Rate-distortion analysis and quality control in scalable Internet streaming. *IEEE Trans. Multimedia* 8, 6, 1135–1146.
- DING, W. AND LIU, B. 1996. Rate control of MPEG video coding and recording by rate-quantization modeling. *IEEE Trans. Circ. Syst. Video Techn.* 6, 1, 12–20.
- HANG, H. AND CHEN, J. 1997. Source model for transform video coder and its application I: Fundamental theory. *IEEE Trans. Circ. Syst. Video Techn.* 7, 2, 287–298.
- HE, Z., KIM, Y., AND MITRA, S. 2001. Low-delay rate control for DCT video coding via ρ -domain source modeling. *IEEE Trans. Circ. Syst. Video Techn.* 11, 8, 928–940.
- HE, Z. AND MITRA, S. 2001. A unified rate-distortion analysis framework for transform coding. *IEEE Trans. Circ. Syst. Video Techn.* 11, 12, 1221–1236.
- HE, Z. AND MITRA, S. 2002. A linear source model and a unified rate control algorithm for DCT video coding. *IEEE Trans. Circ. Syst. Video Techn.* 12, 11, 970–982.
- HSU, C. AND HEFEEDA, M. 2006a. On the accuracy and complexity of rate-distortion models for fine-grained scalable video sequences. Tech. rep. TR 2006-12, Simon Fraser University. <http://nsl.cs.surrey.sfu.ca/projects/fgs/>.
- HSU, C. AND HEFEEDA, M. 2006b. Source models for fine-grained scalable video sequences. Tech. rep., Simon Fraser University.
- ISO/IEC 14496-2. 2004. Coding of audio-visual objects—part 2: Visual.

- ISO/IEC 14496-5. 2004. MPEG-4 Visual reference software.
- JOSHI, R. AND FISCHER, T. 1995. Comparison of generalized Gaussian and Laplacian modeling in DCT image coding. *IEEE Signal Process. Lett.* 2, 5, 81–82.
- LI, W. 2001. Overview of fine-granularity scalability in MPEG-4 video standard. *IEEE Trans. Circ. Syst. Video Techn.* 11, 3, 301–317.
- LIN, L. AND ORTEGA, A. 1998. Bit-rate control using piecewise approximated rate-distortion characteristics. *IEEE Trans. Circ. Syst. Video Techn.* 8, 4, 446–459.
- MALLET, S. AND FALZON, F. 1998. Analysis of low bit rate image transform coding. *IEEE Trans. Signal Process.* 46, 4, 1027–1042.
- MARTINEZ, W. AND MARTINEZ, A. 2002. *Computational Statistics Handbook with Matlab*, 1st ed. Chapman and Hall, Upper Saddle River, NJ.
- MULLER, F. 1993. Distribution shape of two-dimensional DCT coefficients of natural images. *IEEE Electron. Lett.* 29, 22, 1935–1936.
- PARK, H. J. AND LEE, T. W. 2004. Modeling nonlinear dependencies in natural images using mixture of Laplacian distribution. In *Proceedings of Advances in Neural Information Processing Systems (NIPS'04)*. Vancouver, Canada.
- RADHA, H., SCHAAR, M., AND CHEN, Y. 2001. The MPEG-4 fine-grained scalable video coding method for multimedia streaming over IP. *IEEE Trans. Multimedia* 3, 1, 53–68.
- SULLIVAN, G. AND WIEGAND, T. 1998. Rate-distortion optimization for video compression. *IEEE Signal Process. Magazine* 15, 6, 74–90.
- SUN, J., GAO, W., ZHAO, D., AND HUANG, Q. 2005. Statistical model, analysis and approximation of rate-distortion function in MPEG-4 FGS videos. In *Proceedings of SPIE International Conference on Visual Communication and Image Processing (VCIP'05)*. Beijing, China.
- VARANASI, M. AND AAZHANG, B. 1989. Parametric generalized Gaussian density estimation. *J. Acoust. Soc. Amer.* 86, 4, 1404–1415.
- VIDEO TRACES RESEARCH GROUP. 2006. <http://trace.eas.asu.edu/yuv>.
- ZHANG, X., VETRO, A., SHI, Y., AND SUN, H. 2003. Constant quality constrained rate allocation for FGS-coded video. *IEEE Trans. Circ. Syst. Video Techn.* 13, 2, 121–130.

Received August 2006; revised December 2006; accepted February 2007