TECHNICAL RESEARCH REPORT

Institute for Systems Research

Entropy-Constrained Trellis Coded Quantization: Implementation and Adaptation

by C-C. Lee and N. Farvardin

The Institute for Systems Research is supported by the National Science Foundation Engineering Research Center Program (NSFD CD 8803012), Industry and the University

TR 93-36

Entropy-Constrained Trellis Coded Quantization: Implementation and Adaptation[†]

Cheng-Chieh Lee and *Nariman Farvardin*

Electrical Engineering Department
and
Institute for Systems Research
University of Maryland
College Park, MD 20742

Abstract

Entropy-constrained trellis coded quantization (ECTCQ) of memoryless sources is known to be an efficient source coding technique in the rate-distortion sense. We develop an ECTCQ scheme that employs a symmetric reproduction codebook. The symmetry of the reproduction codebook, while essentially costs no performance loss, is exploited to reduce the memory requirement in entropy coding the ECTCQ output. In practice, a buffer of finite, and preferably small, size is needed to interface the variable-length codewords to the fixed-rate channel. An adaptive ECTCQ (A-ECTCQ) scheme, which uses a buffer-state feedback to control the quantizer characteristics to avoid buffer overflow/underflow, is studied in this work. The choice of encoding delay is an important issue in A-ECTCQ, as too long a delay will adversely impact the performance of the feedback control. We propose a pathwise-adaptive ECTCQ (PA-ECTCQ) that solves the encoding delay problem. Simulation results indicate that, while the buffer overflow/underflow problems of the PA-ECTCQ can be practically eliminated, the overall quantization distortion is increased only negligibly over theoretical performance predictions. Our experiments also suggests that PA-ECTCQ is robust with respect to source mismatch.

Index Terms: Trellis coded quantization; buffer-instrumented adaptation; entropy coding; robust quantization.

[†]This work was supported in part by National Science Foundation grants NSFD MIP-91-09109 and NSFD CD-88-03012.

I Introduction

At high rates, optimum uniform vector quantization (VQ) offers the *granular gain* and the *boundary gain* over optimum uniform scalar quantization [1]. For memoryless uniform sources, the granular gain is the only advantage that vector quantizers can exploit over scalar quantizers. Trellis coded quantization (TCQ) is an effective method with modest complexity to realize the granular gain [2]. For encoding the memoryless uniform source using an 8-state trellis, TCQ achieves a mean squared-error (MSE) within 0.5 dB of the rate-distortion bound at all positive integer rates.

For memoryless sources with smooth nonuniform probability density functions (p.d.f.'s), the occurrences of the reproduction symbols of a TCQ are generally not equally probable. This indicates that the performance of the TCQ can be improved, in the rate-distortion sense, by entropy coding the TCQ output. Fischer and Wang have recently combined the TCQ and entropy encoding techniques to develop the entropy-constrained TCQ (ECTCQ) scheme [3]. For encoding memoryless sources with smooth p.d.f.'s, ECTCQ can provide a performance similar to that obtained by standard TCQ for encoding the memoryless uniform source. In this paper, we develop an ECTCQ scheme whose reproduction codebook satisfies a symmetry constraint. Based on this symmetry property of the reproduction codebook and a suitable partition of the codebook subject to the trellis structure, the sequence of the reproduction symbols can be converted to the sequence of the so-called *primary reproduction levels*, and vice versa [4]. It will be shown that for binary encoding the reproduction symbols, our scheme results in a reduction in memory requirement as compared to Fischer and Wang's scheme. This memory reduction comes essentially with no rate-distortion performance loss.

The binary codewords produced by encoding the reproduction symbols generally have a variable-length nature. In practice, transmission of the variable-length code-

words over a fixed-rate channel necessitates the use of a buffer of finite, and preferably small, size. To avoid the buffer overflow/underflow problem in interfacing the binary codewords of the entropy-coded scalar quantizer (ECSQ) to the channel, a buffer-state feedback control was studied in [5]. This approach is adapted to ECTCQ, as we generalize our ECTCQ scheme and propose an adaptive buffer-instrumented ECTCQ (A-ECTCQ) scheme. In A-ECTCQ, the choice of the encoding delay is an important issue. This issue has been solved by a technique called *pathwise adaptation*, leading to a pathwise-adaptive ECTCQ (PA-ECTCQ).

It is often desirable that the encoding scheme be robust when the actual source has a distribution different from what is assumed in the design. We confine our attention to the class of generalized Gaussian sources and show that PA-ECTCQ is reasonably robust under different mismatch conditions.

This paper is organized as follows. In Section II, we describe the ECTCQ with a symmetric reproduction codebook and the conversion between the reproduction symbols and the primary reproduction levels. The memory requirement in binary entropy coding of the reproduction levels is presented and compared to that of Fischer and Wang's scheme in Section III. The operations of A-ECTCQ and PA-ECTCQ are described in Section IV. In Section V, simulation results for the performance of PA-ECTCQ under source mismatch are given. Summary and conclusions are provided in Section VI.

II Entropy-Constrained Trellis Coded Quantization with Symmetric Reproduction Codebook

The key techniques used in trellis coded quantization (TCQ) [2] are *reproduction codebook expansion*, *reproduction codebook partitioning*, and *trellis labeling*. Based on codebook partitioning and labeling the trellis branches with the partitioned subsets, the trellis structure prunes the expanded number of quantization levels down to the

desired encoding rate.

In the sequel, we confine our attention to a specific type of TCQ in which the trellis is specified by a rate 1/2 convolutional encoder[1]. The reproduction codebook is sequentially partitioned into 4 subsets:D_0, D_1, D_2, and D_3. The trellis state is determined uniquely with the content of the convolutional encoder shift register. Corresponding to an input bit to the convolutional encoder, there are two possible outgoing trellis transitions. The next state is uniquely determined by the current state and the trellis transition. The two output bits are used to address one of the four partitioned subsets, which serves as a label to the corresponding trellis transition.

Given a source sequence, the Viterbi algorithm is used to search among the allowed sequences of the reproduction symbols for the one that has the minimum cumulative squared-error from the input sequence. In the standard TCQ [2], each reproduction symbol is specified by the input bit used to identify the trellis transition and the remaining bits used to address one element from the subset labeling the associated trellis transition.

For a source with a nonuniform p.d.f., the occurrences of the reproduction symbols are generally not equally probable. The redundancy, which is defined as the difference between the TCQ encoding rate and the entropy of the reproduction symbols, can be removed by entropy encoding the reproduction symbols. This fact has been the motivation for the entropy-constrained TCQ (ECTCQ) scheme [3] developed by Fischer and Wang.

For a memoryless uniform source, it was shown in [2] that an 8-state TCQ provides a signal-to-noise ratio (SNR) within about 0.5 dB of the rate-distortion bound at all positive integer rates. With a 256-state trellis, this gap is reduced to 0.21 dB. For memoryless sources with nonuniform p.d.f.'s, the 8-state ECTCQ of Fischer and

[1]The more general case in which the trellis is specified by a rate $\tilde{m}/\tilde{m}+1$ convolutional encoder can be found in [2].

Wang [3] also achieves a performance within 0.5 dB of the rate-distortion bound for rates above 1.5 bits/sample[2]. In [6], Marcellin has developed a slight modification of the scheme in [3] which results in performance within 0.5 dB of the rate-distortion bound at rates even below 1.5 bits/sample. In both of the two schemes in [3] and [6] the rate is measured in terms of the entropy of the ECTCQ output and a specific implementation of an entropy coder is not considered.

We now limit our interest to a specific type of reproduction codebook which is symmetric about the origin, $\{\pm a_1, \pm a_2, \cdots, \pm a_J\}$ where $a_1 < a_2 < \cdots < a_J$ and $2J$ is the reproduction codebook size. The reproduction codebook is partitioned into D_0, D_1, D_2, and D_3. Starting from the leftmost point $(-a_J)$, we assign the reproduction symbols to D_0, D_1, D_2, and D_3 in order. Define $B_0 = D_0 \cup D_2$ and $B_1 = D_1 \cup D_3$. We observe that no two elements in B_0 have the same magnitude, therefore each element can be identified by its magnitude. We define the magnitude of a reproduction symbol as the associated *primary reproduction level.* This is also true for B_1. The primary codebook is the collection of the primary reproduction levels. It can be shown that

- The allowed reproduction symbols for a given trellis state are from either B_0 or B_1;
- Due to the symmetry of the reproduction codebook, B_0 and B_1 have a common primary codebook.

The sequence of primary reproduction levels (converted from the sequence of reproduction symbols by taking the magnitude of each reproduction symbol) is independent of the trellis state. The primary reproduction levels can be thought of as the output of a discrete memoryless source (DMS). Besides, the sequence of primary reproduction levels along with the knowledge of the initial trellis state can be used to recover the sequence of reproduction symbols [4].

[2]Such ECTCQ scheme does not work for lower rates with the same high performance because 1 bit/sample is used to specify the trellis path.

We have developed an ECTCQ scheme, subject to the codebook symmetry constraint, that employs entropy coding on the primary reproduction levels. The ECTCQ design algorithm (see [7]) is essentially the same as that in [3] except for the added codebook symmetry constraint. This algorithm is used to design 8-state ECTCQs for rates from 0.1 to 8.0 in steps of 0.1 bits/sample for a unit-variance memoryless Gaussian source. Fig. 1 illustrates the performance of the designed ECTCQs. For comparison, the rate-distortion function of the Gaussian source and the performances of the Huffman-coded ECSQs designed in [8] are also included. In this case, the ECTCQs provide a 1.5-2.0 performance improvement over the coded ECSQs for rates above 2 bits/sample. The performance of the 8-state ECTCQ lies within 0.5 dB of the rate-distortion bound for all encoding rates we considered. Thus the ECTCQ with codebook symmetry constraint performs almost as well as Fischer and Wang's ECTCQ at high rates, indicating that the imposition of the symmetry constraint on the reproduction codebook essentially costs no performance loss. At low bit rates, the proposed scheme performs better than Fischer and Wang's ECTCQ and very close to Marcellin's ECTCQ [6].

III Binary Entropy Coding of the ECTCQ Output

In this section, we discuss the memory requirement of translating the reproduction symbols of ECTCQ into a binary sequence for transmission. Assume the ECTCQ has N trellis states and the reproduction codebook has $2J$ symbols. (In our ECTCQ scheme, the primary codebook has J reproduction levels.)

We first consider the problem of encoding the sequence of the reproduction symbols into a binary sequence using an L-th order Huffman encoder. In this case, the codebook size of the overall Huffman encoder is used as an indication of the required memory for implementation. In our ECTCQ scheme, the primary reproduction levels

are independent of the trellis state and can be considered as output of a DMS. The primary codebook size is J. Therefore the overall Huffman encoder codebook size is J^L.

In Fischer and Wang's ECTCQ scheme [3], a block of L bits is used to specify a sequence of L subset indices. There are 2^L possible sequences of subset indices corresponding to the 2^L possible blocks. Since each subset has $J/2$ reproduction symbols, $2^L(J/2)^L = J^L$ binary codewords are needed to address blocks of L reproduction symbols allowed from the given trellis state. For N states, the overall Huffman codebook size is upper bounded by NJ^L. When N is very large, two or more trellis states are likely to allow a common sequence of L subset indices. To reduce the overall codebook size, common blocks of L reproduction symbols can be encoded by a common portion of the codebook. In this case, another upper bound to the overall codebook size is $(2J)^L$. Therefore the required memory for implementing the Huffman encoder is $\min(N, 2^L)J^L$. Thus, our proposed ECTCQ encoder enjoys a memory reduction of the order of $\min(N, 2^L)$.

Arithmetic coding [9] has been shown to be superior to the better-known Huffman coding. The codebook has $2J$ reproduction symbols. In our proposed ECTCQ the sequence of primary reproduction levels are encoded using a codebook with only J levels. Besides, since the primary reproduction levels can be considered as the output of a DMS, no adaptation in the arithmetic coding is required.

Here, we have shown that our ECTCQ scheme provides a memory reduction in binary encoding of the reproduction symbols using either Huffman or arithmetic coding.

Fig. 2 illustrates the coding redundancy using Huffman and arithmetic codes for the ECTCQs designed for the unit-variance memoryless Gaussian source. The arithmetic coding approach apparently has a smaller coding redundancy than the Huffman coding approach. An L-th order Huffman code is used. L is chosen such

that $J^L \leq 2048$. The value of J depends on the encoding rate. Typically, larger rates require larger values of J. The arithmetic-coded ECTCQ has provided, to our knowledge, the best distortion-rate (actual coding rate) performance for encoding memoryless sources with smooth density functions.

IV Adaptive Buffer-Instrumented ECTCQ

To transmit the binary codewords produced from encoding the primary reproduction levels of the ECTCQ over a communication channel (operating at a fixed rate of R bits/sample), a channel buffer is needed to interface the variable-length codewords to the fixed-rate channel. Since the codeword lengths are independent variables, the variance of their sum grows without bound, hence implying that the buffer must eventually overflow or underflow. This will result in loss of codeword synchronization and large distortions.

A particular entropy-coded quantization approach utilizing an adaptive buffer-state feedback control scheme was described in [5]. To avoid buffer overflow or underflow, this scheme attempts, through feedback control of the quantization process, to maintain the buffer occupancy at the desired "half-full" state. Here, we develop an adaptive buffer-instrumented ECTCQ (A-ECTCQ) based on similar ideas.

The ECTCQ and an Lth-order Huffman encoder are first designed to operate at the nominal output rate, R bits/sample. Assume the channel buffer length is B bits. During each encoding cycle, the codeword produced from encoding a block of L primary reproduction levels is delivered as input to the buffer. Simultaneously, LR bits are released for transmission over the channel. At the nth encoding cycle, the channel buffer-state of the previous cycle, denoted by Z_{n-1}, is used to adapt the quantizer characteristics in order to encode the input samples to a binary codeword with length m_n such that, at the end of this cycle, the buffer-state $Z_n = Z_{n-1} + m_n - LR$ is

as close to $B/2$ as possible. Instead of actually modifying the quantizer characteristics, the following simpler alternative is considered. For most useful sources, the p.d.f.'s are decreasing in magnitude of the argument. Generally, input samples of smaller (larger) magnitudes are encoded into codewords of shorter (longer) lengths. Based on such consideration, the input samples are pre-divided by the feedback signal $f(Z_{n-1} - B/2)$ where $f(\cdot)$ is a monotone increasing function with $f(0) = 1$. An example of such function, which was used in [10] and is used throughout this work, is

$$f(x) = \begin{cases} \frac{1}{1-\gamma/2}; & x > K, \\ \frac{1}{1-\gamma x/2K}; & 0 \le x \le K, \\ 1 + \gamma x/2K; & -K \le x \le 0, \\ 1 - \gamma/2; & x < -K, \end{cases} \tag{1}$$

where $0 \le \gamma \le 2$ and $K \le B/2$ is a specified positive integer. The roles that γ and K play in determining the feedback characteristics are described in [10] and are omitted here.

There is an inherent encoding delay associated with the TCQ system. In the absence of buffer adaptation, a larger encoding delay always yields a better quantization performance. In A-ECTCQ encoding, however, a long delay will deteriorate the performance of the feedback control. Specifically, if the encoding delay is too large, the feedback control cannot adjust the magnitudes of the intermediate samples which have already been encoded in the trellis path (whose associated binary codewords are soon to be delivered to the channel buffer). Therefore, in A-ECTCQ, the choice of encoding delay becomes an important issue.

We have solved this issue by an approach which we call *pathwise adaptation.* In this approach, each trellis path in the trellis decoder is assigned a buffer occupancy counter (assuming they are sent through a pseudo-buffer) which records the buffer-state of the associated sequence of the primary reproduction levels. Since each trellis path holds the occupancy information on its own buffer state, the magnitudes of the input samples can be pathwise pre-divided by the feedback function in (1) depending on its

own buffer-state occupancy. In the encoding process, we search among all survivor trellis paths (whose pseudo buffers neither underflow nor overflow) for the one with the smallest distortion (more precisely, the Lagrangian functional used in ECTCQ encoding algorithm). If no survivor trellis path exists, then buffer underflow/overflow occurs. We refer to this scheme as pathwise-adaptive ECTCQ (PA-ECTCQ).

In Fig. 3, we illustrate the histogram of buffer occupancy of A-ECTCQ with encoding delays of 30, 50, and 80 samples for a memoryless Gaussian source with $R = 3$ bits/sample, $L = 2$, $B = 192$, $K = 96$, $\gamma = 2.0$. In A-ECTCQ, a large encoding delay results in a wider dispersion in the buffer occupancy distribution, which suggests a high probability of buffer overflow/underflow. In our experiment, the buffer overflow/underflow begins to occur when the encoding delay is 90. In PA-ECTCQ, buffer-state histograms corresponding to the same delays, which seem indistinguishable from each other, have a relatively narrow dispersion in the buffer occupancy distribution. This indicates that the pathwise adaptation has practically solved the aforementioned encoding delay problem and the PA-ECTCQ is more effective in avoiding the buffer underflow/overflow problem than A-ECTCQ. The A-ECTCQ and PA-ECTCQ performance (in normalized MSE) as a function of encoding delay is displayed in Fig. 4. We observe that, in PA-ECTCQ, larger encoding delays correspond to a better performance, while this is not true for A-ECTCQ.

V AECTCQ Performance to Source Mismatch

In this section we provide a brief discussion on the robustness of PA-ECTCQ[3] to source mismatch. We will confine our attention to memoryless *generalized* Gaussian sources with p.d.f.

$$p(x;\alpha,\beta) = \frac{\alpha\eta(\alpha,\beta)}{2\Gamma(1/\alpha)} \exp\{-[\eta(\alpha,\beta)|x|]^{\alpha}\} \tag{2}$$

[3]Since PA-ECTCQ has consistently outperformed A-ECTCQ in our experiments, we have not included any result for A-ECTCQ in this section.

where

$$\eta(\alpha,\beta) = \beta^{-1}\left[\frac{\Gamma(3/\alpha)}{\Gamma(1/\alpha)}\right]^{1/2},$$

$\alpha > 0$ describes the exponential rate of decay, β^2 is the source variance σ^2, and $\Gamma(\cdot)$ is the Gamma function.

Suppose that the nominal source p.d.f. $p_0(x)$ is chosen within this class by selecting the parameter $\alpha = \alpha_0$ and variance $\sigma^2 = \sigma_0^2$. We present in this section the performance of the PA-ECTCQ in the presence of some specific source mismatch conditions.

A. Mismatch to Scale

Mismatch with respect to scale occurs when the actual density $p(x)$ differs only in the source variance. More specifically, the actual distribution is described by $p(x) = (\sigma_0/\sigma)p_0((\sigma_0/\sigma)x)$. We define $\zeta = (\sigma/\sigma_0)^2$ to represent the scale mismatch parameter.

The performance (in normalized MSE) of PA-ECTCQ operating at $R = 3$, $L = 2$, $K = B/2$, and encoding delay of 100 samples[4] for a memoryless Gaussian source as a function of ζ is illustrated in Fig. 5. The corresponding performance of the matched Lloyd-Max quantizer, matched ECTCQ as well as the rate-distortion bound are included. For large B and γ, the PA-ECTCQ performance is insensitive to a rather wide range of source scale mismatch. For a fixed buffer length B, the buffer overflow/underflow immunity can be improved with an increased value of γ, but at the cost of some increased distortion.

We observed in our experiments that the mismatch in this case results in an increase (decrease) in the average occupancy level from the design point $Z_0 = B/2$ for $\zeta > 1$ ($\zeta < 1$). This phenomenon is very similar to that in AECQ under the same source mismatch [10] and can be explained likewise.

[4]An encoding delay of 100 samples will be assumed as the PA-ECTCQ performance seems to saturate for larger delays.

B. *Mismatch to Shape*

In this case, the source has the same variance σ_0^2 assumed in the nominal design. But the exponential decay parameter α may be different from the nominal decay parameter α_0.

In Fig. 6, we illustrate the normalized MSE as a function of α for PA-ECTCQ designed under a nominal Gaussian assumption ($\alpha_0 = 2.0$) at $R = 3.0$ bits/sample,$L = 2$, $K = B/2$ and $\gamma = 2.0$. For reference, we also include the performance of the Lloyd-Max quantizer and ECTCQ, both matched to each value of α, and the corresponding rate-distortion bound. As can be observed from Fig. 6, the performance of PA-ECTCQ can be improved by going to larger values of B. For small values of α, the PA-ECTCQ performance is far away from the rate-distortion bound. In this case, as in [10], the performance of PA-ECTCQ may be improved by using a larger reproduction alphabet. We must mention that the adaptation we have considered in this work only modifies the ECTCQ reproduction levels. The PA-ECTCQ performance may be further improved if the entropy-coder is also made adaptive.

The average buffer occupancy state under shape mismatch is always shifted somewhat to the left. As was explained in [10] for the AECQ case, this is because the Gaussian ($\alpha_0 = 2$) p.d.f. has the highest differential entropy for a given variance. The histograms for broad-tailed distributions ($\alpha < 2.0$) generally have wider buffer-state dispersions. For more narrow-tailed distributions ($\alpha > 2.0$) the dispersion decreases and the buffer occupancy histogram is highly peaked. Explanations of such phenomena can also be found in [10].

VI Summary and Conclusions

We have proposed an ECTCQ scheme which uses a symmetric reproduction codebook. The binary codewords are obtained from encoding the primary reproduction levels,

converted from the reproduction symbols. Our simulations show that the performance of our ECTCQ scheme is the same as Fischer and Wang's ECTCQ scheme, indicating that the imposition of the symmetry constraint on the reproduction codebook costs no performance penalty. When noiseless codes (for example, Huffman and arithmetic codes) are employed to encode the sequence of the reproduction symbols into a binary stream for transmission, our scheme enjoys a reduction in the required implementation memory.

In order to combat buffer overflow/underflow in interfacing the variable-length binary codewords to the fixed-rate channel, we have proposed an adaptive ECTCQ scheme. The conflict in choosing the encoding delay in the adaptive scheme is solved by a pathwise adaptation approach. The PA-ECTCQ scheme practically eliminates the buffer overflow/underflow problem at the cost of a negligible increase in overall quantization distortion. Within the class of memoryless sources with generalized Gaussian distribution, our simulations also demonstrate that the PA-ECTCQ performance is robust in the presence of scale-mismatch and to a lesser extent for shape-mismatch.

References

[1] V. Eyuboğlu and G. D. Forney, Jr., "Lattice and trellis quantization with lattice- and trellis-bounded codebooks—high-rate theory for memoryless sources," *IEEE Trans. Inform. Theory*, vol. IT-39, pp. 46–59, January 1993.

[2] M. Marcellin and T. Fischer, "Trellis coded quantization of memoryless and Gauss-Markov sources," *IEEE Trans. Commun.*, vol. COM-38, pp. 82–93, Jan. 1990.

[3] T. R. Fischer and M. Wang, "Entropy-constrained trellis-coded quantization," *IEEE Trans. Inform. Theory*, vol. IT-38, pp. 415–426, March 1992.

[4] R. Laroia, "On optimal shaping of multidimensional constellations — an alternative approach to lattice-bounded (Voronoi) constellations," *Submitted to IEEE Trans. Inform. Theory*, November 1991.

[5] N. Farvardin and J.W. Modestino, "Adaptive buffer-instrumented entropy-coded quantizer performance for memoryless sources," *IEEE Trans. Inform. Theory*, vol. IT-32, pp. 9–22, Jan. 1986.

[6] M. W. Marcellin, "On entropy-constrained trellis-coded quantization," *IEEE Trans. Communications*, to appear.

[7] C. Lee and N. Farvardin, "Entropy-constrained trellis coded quantization: Implementation and adaptation," *In preparation*, 1993.

[8] N. Tanabe and N. Farvardin, "Subband image coding using entropy-coded quantization over noisy channels," *IEEE J. Select. Areas Commun.*, vol. SAC-10, pp. 926–943, June 1992.

[9] J. Rissanen, "Generalized Kraft inequality and arithmetic coding," *IBM J. Res. Develop.*, pp. 198–203, May 1976.

[10] J.W. Modestino, D.D. Harrison, and N. Farvardin, "Robust adaptive buffer-instrumented entropy-coded quantization of stationary sources," *IEEE Trans. Commun.*, vol. COM-38, pp. 859–867, June 1990.

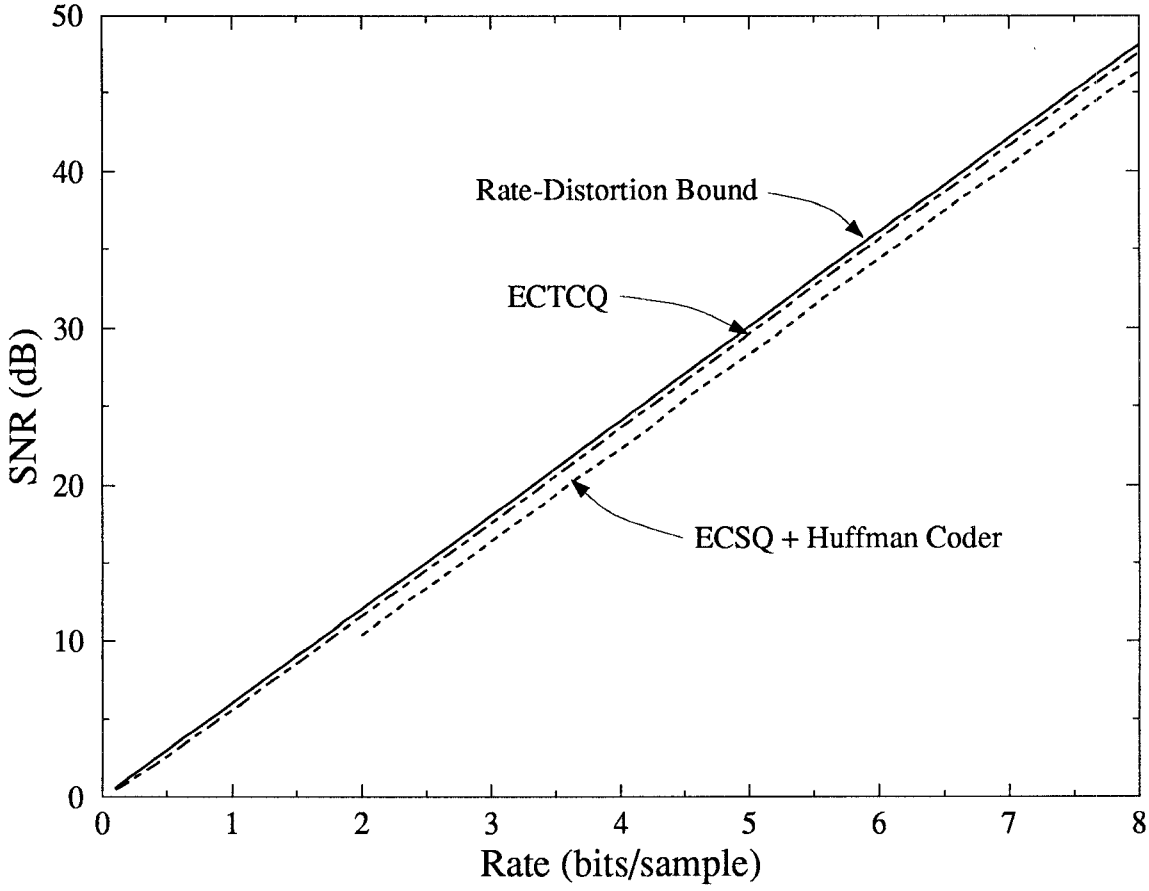

Figure 1: Performance (SNR in dB) of ECTCQ for a memoryless Gaussian source.

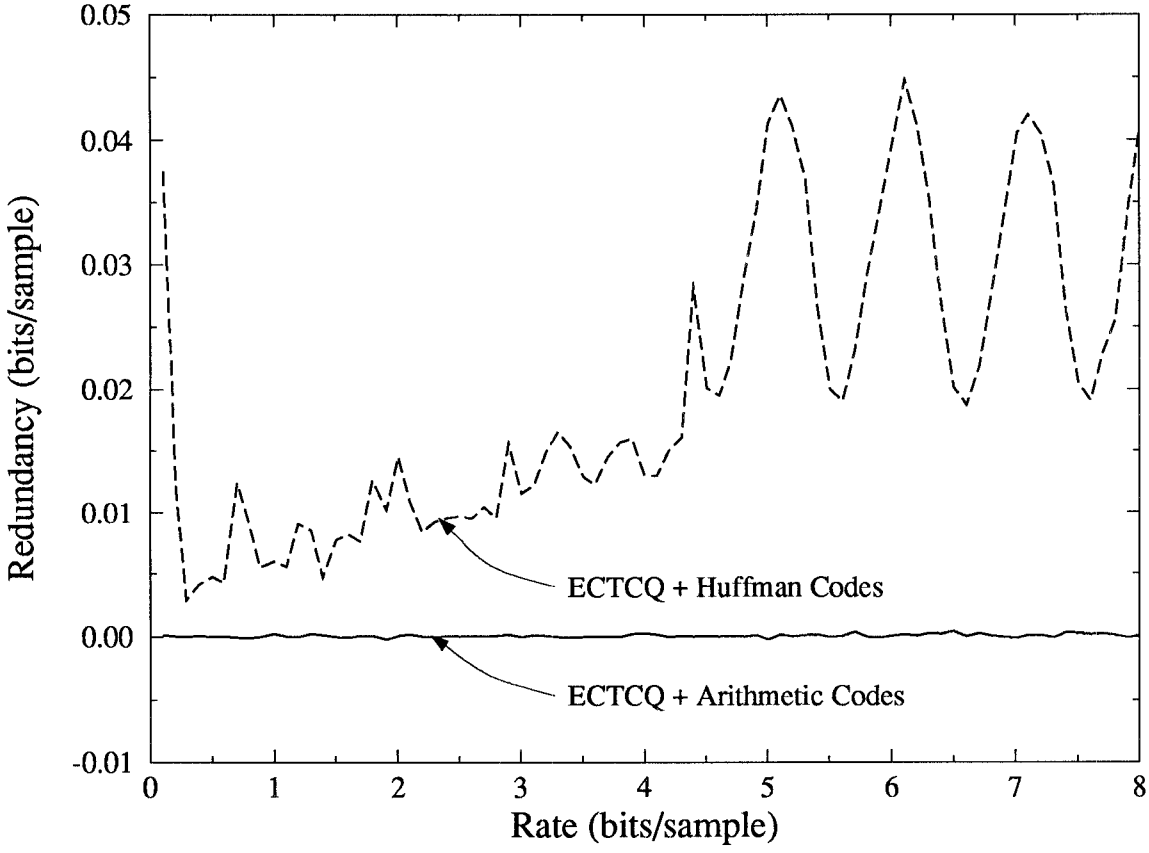

Figure 2: Coding redundancy of ECTCQ encoded by Huffman and arithmetic coder for a memoryless Gaussian source.

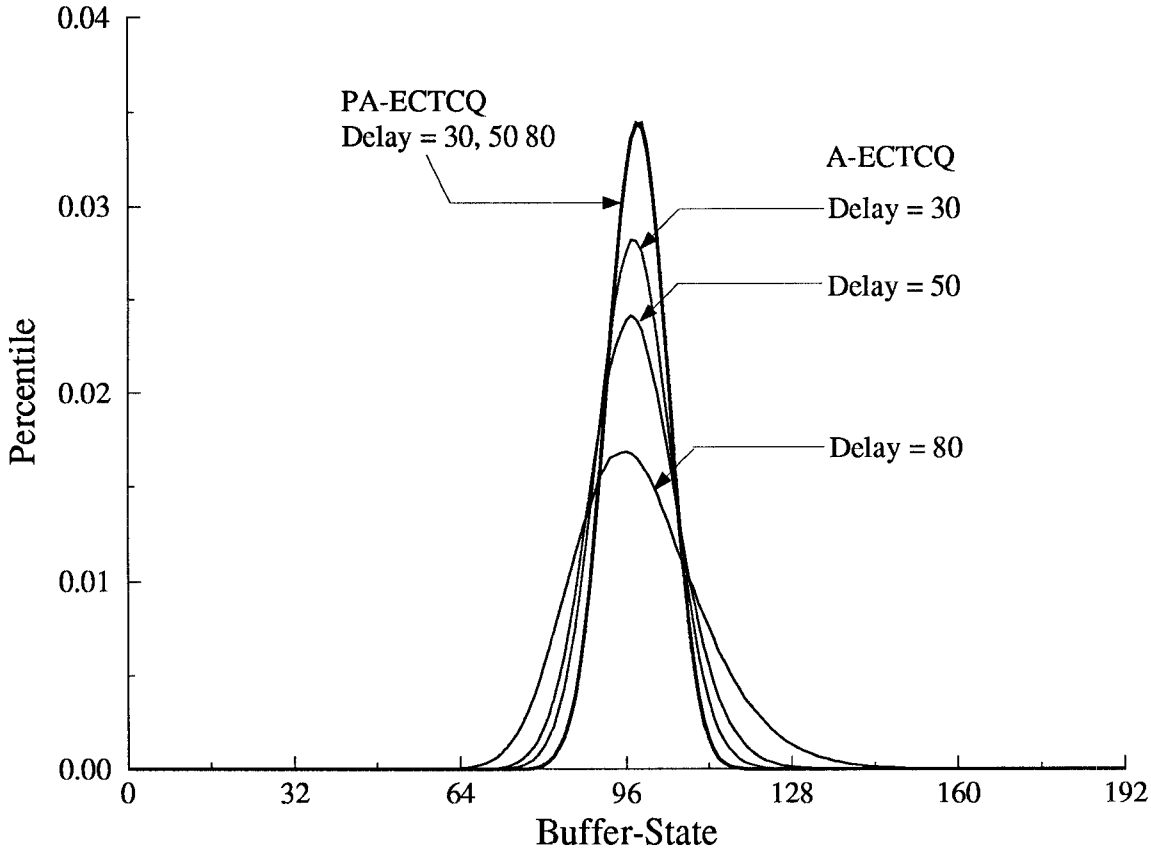

Figure 3: Histogram of the buffer state in A-ECTCQ and PA-ECTCQ for a memoryless Gaussian source; $R = 3$ bits/sample, $L = 2$, $B = 192$, $K = 96$, $\gamma = 2.0$.

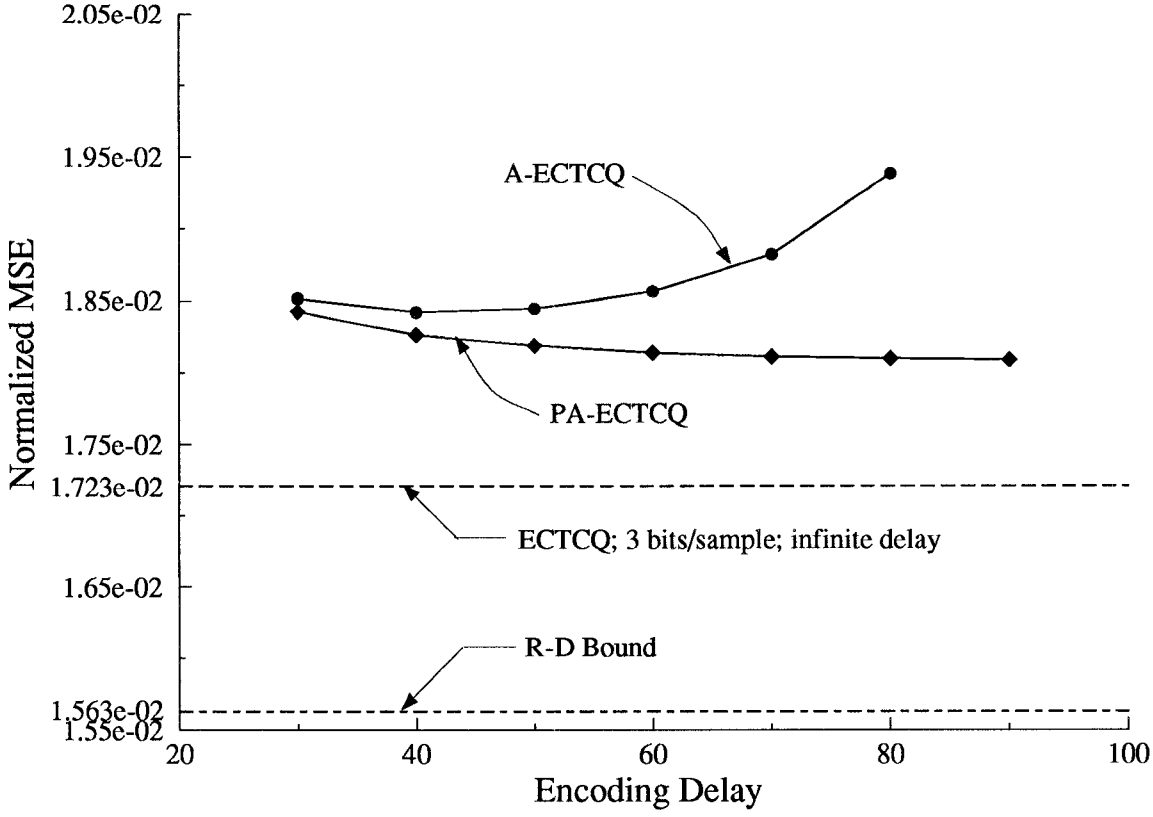

Figure 4: A-ECTCQ and PA-ECTCQ performance (in normalized MSE) vs encoding delay for a memoryless Gaussian source; $R = 3$ bits/sample, $L = 2$, $B = 192$, $K = 96$, $\gamma = 2.0$.

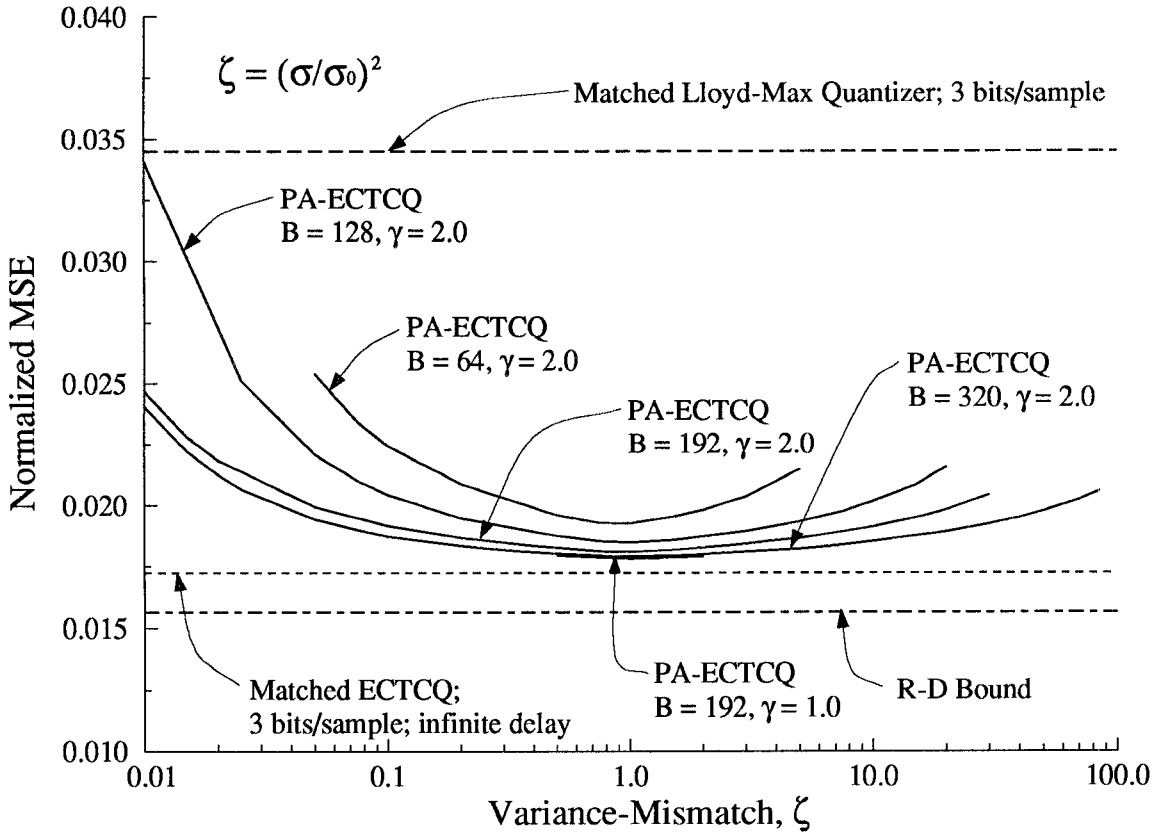

Figure 5: Normalized MSE of PA-ECTCQ under scale mismatch. Nominal source Gaussian ($\alpha_0 = 2.0$); $R = 3$ bits/sample, $L = 2$, $K = B/2$, encoding delay = 100.

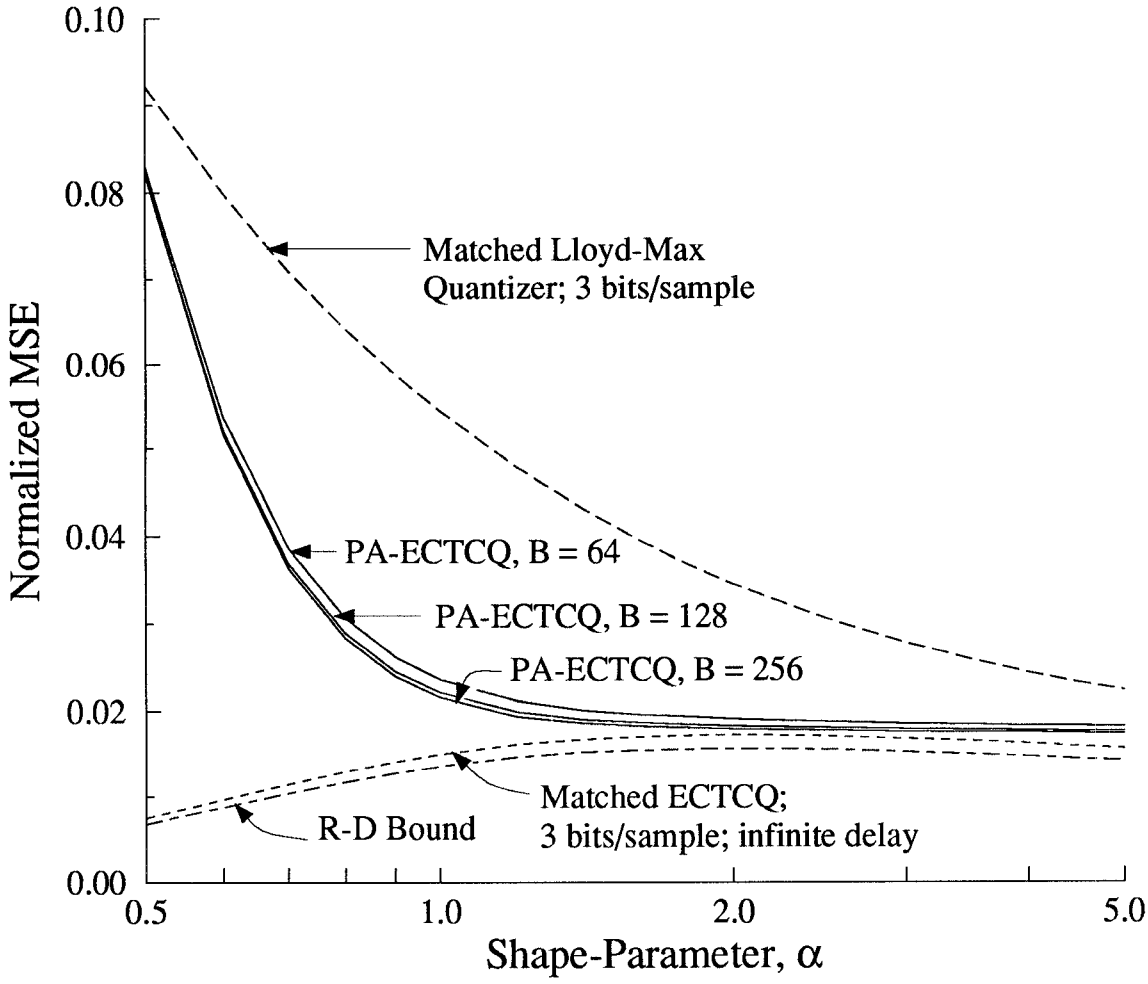

Figure 6: Normalized MSE of PA-ECTCQ under shape mismatch. Nominal source Gaussian ($\alpha_0 = 2.0$); $R = 3$ bits/sample, $L = 2$, $K = B/2$, $\gamma = 2.0$, encoding delay = 100.