On the Suboptimality of Orthogonal Transforms for Single- or Multi-Stage Lossless Transform Coding

David Mary, Dirk T.M. Slock
Institut EURÉCOM [1]
2229, Route des Crêtes, BP 193, 06904 Sophia Antipolis, France
Tel: +33 4 9300 2916/2606; Fax: +33 4 9300 2627
e-mail: {David.Mary,Dirk.Slock}@eurecom.fr

Abstract

Orthogonal transforms are compared with the causal transform in lossless transform coders. For single-stage lossless coding, it was shown in [1] that the integer-to-integer implementation of the best orthogonal decorrelating transform, the KLT, leads to lower compression performance than its causal counterpart. In this work, we pursue this analysis in the framework of a multi-stage lossless coding scheme, which yields a low resolution (lossy) signal, and an error signal. This scheme allows one to choose the respective bitrates of both complementary signals, depending for example on the bandwidth of the transmission link. We show that the causal approach allows one to code the data (almost) without causing any excess bitrate as compared with a single-stage coder, whereas for orthogonal transforms, the price paid for the multiresolution approach is a bitrate penalty of 0.25 bit per sample. This excess bitrate is due to a "gaussianization effect" of the transforms. Also, the approach based on the causal transform allows one to easily switch between a single- or a multi-stage compressor. Moreover, in the framework of interchannel redundancy removal, this approach allows one to easily fix the distortion and rate for both the low resolution and the error signal of each channel, by using different stepsizes in the quantization stage. Any of the channels may, as a particular case, be chosen to be directly losslessly coded. Finally, a side advantage of the causal approach is that entropy coding of the error signal is made very simple since for odd quantization stepsizes, the discrete error sources are uniformly distributed, so that the optimal codewords have the same length, and fixed rate coding is optimal.

1 Introduction

Suppose one disposes of a discrete vectorial source $\underline{x}$ whose samples are $\underline{x}_k$. This source may for example be composed of N scalar audio signals x_i, in which case

[1] Eurécom's research is partially supported by its industrial partners: Ascom, Swisscom, Thales Communications, STMicroelectronics, CEGETEL, Motorola, France Télécom, Bouygues Telecom, Hitachi Europe Ltd. and Texas Instruments.

$\underline{x}_k = [x_{1,k}...x_{N,k}]^T$, or by the samples of the same scalar source, in which case $\underline{x}_k = [x_k \; x_{k-1}...x_{k-N+1}]^T$. There are many ways of losslessly coding this source. Vector entropy coding is known to be asymptotically optimal w.r.t. to the block length, but requires to estimate the joint probabilities of the vectors. Such a coding procedure is thus very complex and not adapted to signals (such as audio) which present long term correlations. Some approaches have thus been proposed, which divide into two steps the coding procedure: firstly, a transform is applied to each block in the aim of decorrelation, and scalar entropy coding of the transform components is secondly realized. Lossless implementation of the DCT (Discrete Cosine Transform) has been for example described in [2], of the DFT (Discrete Fourier Transform) in [3], of the orthogonal Karhunen-Loève transform (KLT) in [4], and a comparison with the causal LDU (based on a Lower-Diagonal-Upper factorization of the covariance matrix $R_{\underline{x}\underline{x}}$) transform was proposed in [1]. In this framework, the vectorial signal $\underline{x}$ gives rise to N transform signals from which the decoder is able to losslessly recover the original signal $\underline{x}$.

Besides this single-stage or "one-shot" compression approach, another way of lossless coding is to first apply a lossy coding scheme (which yields a first stream of N components $\underline{y}^q$), and then to separately code the error signal $\underline{e}$, resulting in a two-stage structure, see Figure 1. Usually, $\{Q\}$ are integer rounding operators [5, 6],

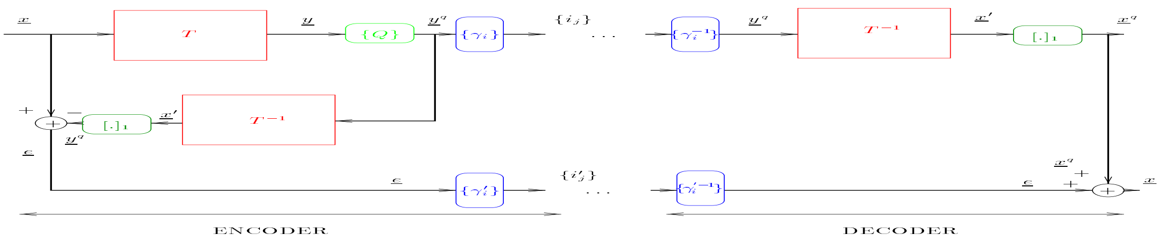

Figure 1: Two-Stage Lossless Transform Coding. $\{Q\}$ denotes uniform scalar quantizers, $\{\gamma_i\}$ and $\{\gamma_i'\}$ scalar entropy coders, and $[.]_1$ rounding operators.

ensuring that the $\underline{y}^q$ are discrete valued and can efficiently be entropy coded. A simple improvement to this approach can however be brought by introducing scalar quantizers with stepsizes greater than unity. Depending on these stepsizes, the rate dedicated to code the low resolution version $\underline{x}^q$ of $\underline{x}$ can then be made lower, at the price indeed of allowing a greater distortion for the low resolution signal. Since the rate-distortion trade-off is then more flexible than in the unity stepsize case, the advantage of this scheme is that, in the case of variable transmission bandwidth, an approximative version of the signal of interest can be quickly obtained, independently of the error signals. The original signal can then in any case be recovered by adding the error signals. In this framework, a transform is firstly applied to a block of signal, the decorrelated components $\underline{y}$ are then quantized by means of uniform scalar quantizers, and further entropy coded. By inverting the transform and taking the integer part of the resulting reconstructed value $\underline{x}'$, the error signal $\underline{e}$ can be generated by substraction : $\underline{e} = \underline{x} - \underline{x}^q$, and further entropy coded. The decoder generates then $\underline{x}^q$ in the same way , and recovers $\underline{x}$ by $\underline{x} = \underline{x}^q + \underline{e}$. Note that the rounding operations

are necessary: since T is a linear transform, $\underline{x}' = T^{-1}\underline{y}^q$ is generally not integer valued. In this framework, this work compares the compression performance of orthogonal and causal transforms (in which case the general multiresolution coding scheme is slightly different, see Section 3). A genralization of the two-stage structure to M stages is analyzed for the causal transform in [7].
Let us now denote by $r_{1-shot}(\underline{x})$ the bitrate dedicated to the losslessly code the source $\underline{x}$ with a single-stage lossless coder. The main question adressed in this work is: Is there a price to pay, in terms of rates, by using multiresolution approach ? Or in other words, will the overall bitrate $r_{LR}(\underline{y}) + \overline{r}(\underline{e})$ be more than $r_{1-shot}(\underline{x})$, and if yes, how much ? As will be seen in the next sections, orthogonal transforms suffer, among other drawbacks, from some rate penalty, whereas the causal transform does not. In the following, the KLT will be used as a benchmark for orthogonal transforms, but as will be underlined, the conclusions of this analysis can be generalized to other orthogonal transforms.
The rest of the paper is organized as follows. Section 2 states somes definitions and notations, recalls the main characteristics of the causal transform and some results about the "one-shot" compression. Section 3 describes the proposed two-stage coding structures and analyzes the statistics of the error signals. Section 4 is dedicated to the analysis of the bitrates in the case of Gaussian signals and Section 5 deals with non-Gaussian probability density functions (p.d.f.s). Section 6 considers the particular case where lossless transform coding is used to remove intrachannel redundancies, and the last section presents some simulations results.

2 Single-Stage Structure

Suppose we dispose of a vectorial source $\underline{x}$, which is obtained by some discretization (quantization) process from a continuous-amplitude source $\underline{x}^c$ (for notation convenience, the time index k will be often omitted). We assume in this work very high resolution, that is, x is integer valued, and $\sigma^2_{x_i} >> 1$. The rounding operation may be defined as follows. A component x^c_i of $\underline{x}^c$ can then be written as $x^c_i = sign(x^c_i) \times k + sign(x^c_i) \times \delta$, where k is a positive integer, and δ belongs to $[0, 1[$. The rounded value obtained from x^c_i and denoted by $[x^c_i]_1$ is then defined by

$$[x^c_i]_1 = round(x^c_i) \quad = \quad \left\{ \begin{array}{l} sign(x^c_i) \times k + sign(x^c_i) \times 1 \text{ if } \delta \geq 0.5, \\ sign(x^c_i) \times k \text{ if } \delta < 0.5. \end{array} \right. \tag{1}$$

Similarly, an uniform quantizer with non unity stepsize Δ associates then to x^c_i a quantized value $[x^c_i]_\Delta$ by computing $[x^c_i]_\Delta = round(\frac{x^c_i}{\Delta}) \times \Delta$. In the case where $\underline{x}^c$ is a vector, $[\underline{x}^c]_\Delta$ will denote quantization of each component x^c_i.
In order to compute the analysis of the different rates, we will use the relation of differential to discrete entropy: $H(x_i) + \log_2 \Delta \rightarrow h(x^c_i)$ as $\Delta \rightarrow 0$, where H denotes the discrete entropy of the discrete source x_i, obtained by uniform quantization with stepsize Δ from the continuous amplitude source x^c_i with differential entropy h. For vectors, a similar relation can be derived, see [4, 1].
We now recall some results concerning single-stage compression of a vectorial source $\underline{x}$ by means of integer-to-integer transforms.

2.1 Lossless Implementation of the Transforms

In the causal case [8, 9], the vector $\underline{x}$ is decorrelated by means of a lower triangular transform L. The transform vector $\underline{y}$ is $L\underline{x} = \underline{x} - \overline{L}\underline{x}$, where $\overline{L}\underline{x}$ is the reference vector. the components y_i are the prediction errors of x_i with respect to the past values of $\underline{x}$, the $\{x_{1:i-1}\}$, and the optimal coefficients $-L_{i,1:i-1}$ are the optimal prediction coefficients. It follows that $R_{\underline{xx}} = L^{-1} R_{\underline{yy}} L^{-T}$, which represents the LDU factorization of $R_{\underline{xx}}$. In the unitary case, $R_{\underline{xx}} = V^{-1} \Lambda V^{-T}$, where Λ is the diagonal matrix of the eigenvalues of $R_{\underline{xx}}$. (Note that in both cases, $\det R_{\underline{yy}} = \det R_{\underline{xx}}$, since both tranforms are unimodular.)

However, since the resulting components y_i are generally not integer, the transform cannot be used as is for lossless coding. A lossless implementation of the LDU transform can be obtained as follows, see Figure 2. In this case, the transform signals

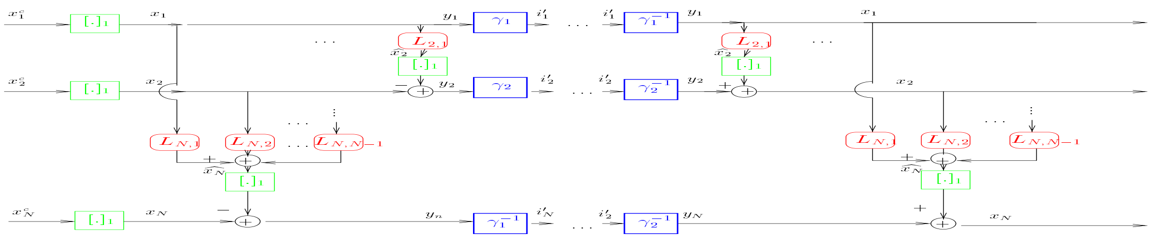

Figure 2: Lossless "One-Shot" Implementation of the LDU Transform.

are obtained by $y_{i,k} = x_{i,k} - [\widehat{x_{i,k}}]_1 = x_{i,k} - [\overline{L}_{i,1:i-1}\underline{x}_{1:i-1,k}]_1$, where $\widehat{x_{i,k}}$ is the estimate of $x_{i,k}$ based on the previous samples of $\underline{x}_k$. The signals y_i are then entropy coded (bitstreams $\{i'_j\}$). At the decoder, each component x_i is losslessly recovered by $x_i = y_i + [\widehat{x_{i,k}}]_1$.

Many lossless implementations of orthogonal transforms have been studied recently, see for example [4, 2, 3]. Concerning the KLT, the integer-to-integer approximation of the optimal linear orthogonal decorrelating transform is based on the factorization of the unimodular matrix into a product of triangular matrices, cascaded with rounding operations ensuring the inversibility of the global transform [4].

2.2 Bitrates for the "One-Shot" Structures

Because of its triangular structure, the LDU transform is naturally well suited for factorizations involving lifting steps and roundings ($N-1$ for an N-transform). This is not the case of noninteger- valued orthogonal transforms, in which case the number of rounding operations decrease the coding performance. It was shown in [1] that the minimum rate required to losslessly code the transform signals can be related to the mutual information between the x_i. For L and V, these rates are

$$\begin{aligned} r_{1-shot,L} &= \tfrac{1}{N}\sum_{i=1}^{N} H(y_i, L) \approx \frac{1}{2}\log_2 2\pi e \det R_{\underline{xx}}^{\frac{1}{N}} - \frac{1}{24N\log 2\sigma_{x_1}^2}, \\ r_{1-shot,V} &= \tfrac{1}{N}\sum_{i=1}^{N} H(y_i, V) \approx \frac{1}{2}\log_2 2\pi e \det R_{\underline{xx}}^{\frac{1}{N}} + \frac{1}{2N\log 2}\sum_{i=1}^{N}\frac{d_i}{\lambda_i} > r_{1-shot,L}, \end{aligned} \tag{2}$$

where V denotes a KLT of $R_{\underline{x}\underline{x}}$, and d_i are positive quantities which depend on the coefficients of V. Thus, for single-stage coders, the best linear decorrelating orthogonal transform is slightly less efficient than the causal one. For other transforms such as DCT, DFT..., the compression performance will most probably be still worse, since their decorrelation efficiency is less than that the KLT, and they are square matrices with non-integer coefficients. In the next section, these approaches are compared for a two-stage structure.

3 Two-Stage Stucture

3.1 Orthogonal Transforms

As stated in introduction, the vectorial source $\underline{x}$ can be losslessly coded by means of a two-stage structure, yielding a low resolution version $\underline{x}^q$, and an error signal $\underline{e}$.

In the case of orthogonal transforms (KLT, DCT...), the coding scheme is represented by Figure 1. At the encoder, the transform signals are obtained by applying the transform T to $\underline{x}$. The resulting components y_i are then quantized, scalar entropy coded, and transmitted to the decoder (bitstreams $\{i_j\}$). The low resolution version $\underline{x}^q$ of $\underline{x}$ is then computed as follows at both encoder and decoder (assuming no transmission errors). The transform $T^{-1} = T^T$ is applied to $\underline{y}^q$. Since the result $\underline{x}'$ is generally not integer-valued, rounding operations are necessary to ensure that the error signals $e_i = x_i - x_i^q$ are integer and can be scalar entropy coded (bitstreams $\{i'_j\}$). The original signal $\underline{x}$ is recovered at the decoder by $\underline{x} = \underline{x}^q + \underline{e}$.

We may now consider the statistics of the error signals. Let $\underline{q}$ denote the quantization noise in the transform domain ($q_i = y_i^q - y_i$). In the signal domain, the quantization noise resulting on $\underline{x}'^q$ is $T^{-1}\underline{q}$. The covariance matrix $E\underline{x}'\underline{x}'^T = R_{\underline{q}'\underline{q}'}$ equals then, under high resolution assumption, $R_{\underline{x}\underline{x}} + T^{-1}R_{\underline{q}\underline{q}}T^{-T}$, which shows that the variance of the quantization noise on a component $x_i^{q'}$ is $(T^{-1}R_{\underline{q}\underline{q}}T^{-T})_{ii}$. In case of signal dependent transforms T such as KLT, chhosing all stepsizes equal ensure that this variance is $\frac{\Delta^2}{12}$. Then the rounding operations increase the variances on the components x_i^q, which can be approximated as $\frac{1}{N}E||\underline{x} - \underline{x}^q||^2 \approx \frac{1}{N}E||\underline{x} - \underline{x}'^q||^2 + \frac{1}{12}$. Thus, the distortion is indeed fixed by the quantization stepsize Δ, and is the same on all the components, $\frac{\Delta^2+1}{12}$. The error signal is now $\underline{e} = \underline{x} - \underline{x}^q = [\underline{x} - \underline{x}'^q]_1 = [T^{-1}\underline{q}]_1$. Thus, each e_i is a discretized mixture of N random variables (r.v.s), wich, as shown by high resolution quantization theory, are uniform if Δ is small in comparison with the variances $\sigma_{y_i}^2$. Since the convolution of N uniform r.v.s tends quickly to a Gaussian, the error signals e_i may be approximated as continuous Gaussian r.v.s with variances $\frac{\Delta^2}{12}$, discretized with stepsize unity. The minimum distortion is now obtained by setting $\Delta = 1$, resulting in a distortion of $\frac{2}{12} = \frac{1}{6}$ on each component. Thus, this scheme does not offer the simple mean of switching from the two-stage to the "one-shot" coder by only setting the quantization stepsizes to 1. Since the e_i are nearly Gaussian, the probability that an error occurs for a general Δ can be approximated with the error function:

$$P(e_i \neq 0) = P(|e_i| \geq \frac{1}{2}) \approx 1 - erf(\sqrt{\frac{3}{2}}\frac{1}{\Delta}). \tag{3}$$

For $\Delta = 1$, we obtain $P(e_i \neq 0) \approx 0.08$, which means that one out of twelve samples should be corrected at the decoder to ensure the losslessness. The question of the rate dedicated to code $\underline{e}$ is examined in the next section.

3.2 Causal Transform

The two stage-causal structure may be described by Figure 3. The transform signals

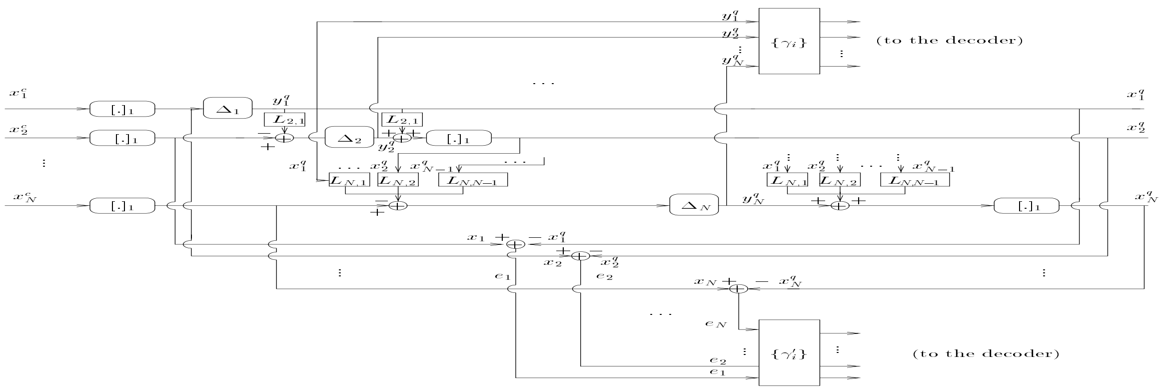

Figure 3: Encoder of the Two-Stage Lossless Coding in the Causal Case.

are computed by substracting the optimal estimate of x_i based on the past *quantized* samples $x^q_{1:i-1}$, and by quantizing with some stepsize Δ_i the resulting error prediction. The reason for computing the prediction by means of quantized data is that we are interested in a low resolution signal which can be computed *independently* of the error signals. Thus, only the available x^q_i at the decoder should be used to compute the remaining x^q_j, $j > i$. As will be commented in the rate analysis, prediction based on quantized data is slightly less efficient than that based on original data, though this difference will be shown to be negligible in most of the cases. Each error signal is thus computed by

$$e_i = x_i - x^q_i = x_i - [y^q_i + \overline{L}_{i,1:i-1}\underline{x}^q_{1:i-1}]_1 = [x_i - \overline{L}_{i,1:i-1}\underline{x}^q_{1:i-1} - y^q_i]_1 = [y_i - y^q_i]_1. \tag{4}$$

Thus, the errors e_i are now the discretized version of the quantization error in the transform domain, which takes values in the interval $[-\frac{\Delta_i}{2}; \frac{\Delta_i}{2}]$. Concerning the statistics of e_i, three cases should be considered. Firstly, if $\Delta = 1$, it can be checked that fixing all stepsizes to 1 yields the coding scheme of Figure 2. If now Δ_i is an odd integer greater than 1, the rounding definition (1) yields equally likely errors (with probabilities $p_i = \frac{1}{\Delta_i}$), and belonging to $\{-\frac{\Delta_i - 1}{2}, -\frac{\Delta_i - 1}{2} + 1, ..., \frac{\Delta_i - 1}{2}\}$.
If finally Δ_i is even, all the errors are equally likely except from + and $-\frac{\Delta_i}{2}$, which are, due to (1), twice less likely than the other ones (for example, $+\frac{\Delta_i}{2}$ occurs only for

positive values of $\widehat{x}_i$). Thus, the values $0, \pm 1 ... \pm \frac{\Delta_i}{2} - 1$ are equaly likely with probabilities $\frac{1}{\Delta_i}$, and $\pm \frac{\Delta_i}{2}$ have probabilityies $\frac{1}{2\Delta_i}$. These remarks lead to a probability of nonzero error which is

$$P(e_i \neq 0) = P(|e_i| \geq \frac{1}{2}) = 1 - \frac{1}{\Delta_i} \; \forall \Delta_i. \quad (5)$$

Figure 4.a) plots the observed and theoretic probabilities of error in the orthogonal case and in the causal case as given by (3) and (5) respectively (for these simulations, all the quantizations stepsizes are equal, see simulations details in Section 7). As a conclusion, the causal transform allows one, on the one hand, to switch easily between either a single, or a two-stage structure, by simply fixing the stepsizes to 1. Moreover, the stepsizes Δ_i may in general be different, allowing one to choose a possibly different rate-distortion trade-off for each signal x_i^q. Also, any x_i can be chosen in the causal case to be directly losslessly coded, by setting the corresponding Δ_i to 1. On the other hand, the KLT does not benefit of these advantages because of a mixing effect of the quantization errors in the signal domain.
As shown by Figure 4, the probability that an error occurs is higher in the causal case than in the orthogonal case as soon as $\Delta > 1$. Does this preclude that the rate associated to the error signal is in the causal case higher than in the orthogonal case ? As will be shown in the next section, the answer is no because of the Gaussianity of the quantization error in the orthogonal case.

4 Analysis of the Rates

In this section we assume Gaussian signals for which close form expression for the rates can be obtained. (The case of non-Gaussian p.d.f.s will be discussed in the Section 6.) Moreover, we assume that all the quantization stepsizes are equal in both causal and unitary cases (though this, as stated in Section 3, not necessary for the LDU transform).

4.1 Low Resolution Versions

For the two transformations, one should compute $r_{LR_T} = \frac{1}{N}\sum_{i=1}^{N} H(y_i, T)$ $\approx \frac{1}{N}\sum_{i=1}^{N} h(\sigma^2_{y_i}, T) - \log_2 \Delta$, with $T = L, V$.For both transforms, the transform signals are indeed Gaussian. The variances $\sigma^2_{y_i}$ are in the orthogonal case the eigenvalues λ_i of $R_{\underline{xx}}$, which leads to

$$r_{LR_V} \approx \frac{1}{N}\sum_{i=1}^{N}\left(\frac{1}{2}\log_2 2\pi e \lambda_i - log_2 \Delta\right) = \frac{1}{2}\log_2 2\pi e (\det R_{\underline{xx}})^{\frac{1}{N}} - \log_2 \Delta. \quad (6)$$

In the causal case, the variances of the transform signals $\sigma^2_{y_i}$ are not exactly the optimal prediction error variances $\sigma^2_{y_i^0}$ of order $i-1$ based on $x_{1:i-1}$, because the prediction is computed by means of quantized samples. One shows [9] that $\sigma^2_{y_i} = \sigma^2_{y_i^0} + \frac{\Delta^2}{12}(\overline{LL}^T)_{ii}$. As in ADPCM, the prediction error variances are increased due to a quantization noise feedback. These results lead then to

$$\overline{r}_{LR_L} = \frac{1}{N}\sum_{i=1}^{N} H(y_i, L) \approx \frac{1}{2}\log_2 2\pi e(\det R_{\underline{xx}})^{\frac{1}{N}} - \log_2\Delta + \frac{\Delta^2}{24N\log 2}\sum_{i=1}^{N}\left(\frac{1}{\lambda_i} - \frac{1}{\sigma^2_{y_i^0}}\right). \quad (7)$$

Thus, for the same distortion $\frac{\Delta^2+1}{12}$ on each component x_i^q, the bitrate required to entropy code the low resolution version obtained by means of the causal transform should require an excess bitrate in comparison with the KLT. Simulations in Section 7 show however that this excess bitrate is negligible for many practical coding situations.

4.2 Error Signals

Concerning the rate $\overline{r}_T$ dedicated to the error signals, one can compute the entropies of the signals e_i by using the error analysis of Section 3. In the unitary case, each e_i can be seen as discretized Gaussian r.v. with variance $\frac{\Delta^2}{12}$. Thus, the bitrate $\overline{r}_V = \frac{1}{N}\sum_{i=1}^{N} H(e_i, V)$ can be written as

$$\overline{r}_V \approx \frac{1}{N}\sum_{i=1}^{N}\frac{1}{2}\log_2 2\pi e\frac{\Delta^2}{12} = \log_2\Delta + \underbrace{\frac{1}{2}\log_2\frac{\pi e}{6}}_{\approx 0.25\ bit}, \quad (8)$$

where 0.25 bit is the well known difference between Gaussian and uniform entropies. In the causal case we obtain, depending on the parity of Δ

$$\overline{r}_{L,even} = -\sum_{i=1}^{N} p_i\log_2 p_i = \log_2\Delta + \frac{1}{\Delta}\ ;\ \overline{r}_{L,odd} = -\sum_{i=1}^{N} p_i\log_2 p_i = \log_2\Delta. \quad (9)$$

Comparing (8) and (9), the rate which is required to code the error signal in the unitary case is ≈ 0.25 b/s more than in the causal case. Moreover, in the case of odd Δ, the error are uniformly distributed, which means that no compression is required for the bitrate to reach the entropy of the sources e_i, and a simple optimal coding procedure is simply to transmit the binary representation of the values e_i.

5 Intrachannel Redundancy Removal

The coding schemes presented in Figures 1 and 3 can indeed be used to remove intrachannel redundancies, In this case, each data block is $\underline{x}_k = [x_k\ x_{k-1}...x_{k-N+1}]^T$. Again, we assume a Gaussian p.d.f. and equal quantization stepsize Δ for $\{Q\}$. By letting the block length grow to infinity, and using the asymptotic distribution of Toeplitz matrices [10], we get for the bitrates of the low resolution signals

$$\begin{aligned} r_{LR_V} &\approx \tfrac{1}{2}\log_2 2\pi e\ e^{\int_{-\frac{1}{2}}^{\frac{1}{2}}\log S_{xx}(f)df} - \log_2\Delta \\ r_{LR_L} &\approx r_{LR_V} + \tfrac{\Delta^2}{24\log 2}\left[\int_{-\frac{1}{2}}^{\frac{1}{2}} S_{xx}^{-1}(f)df - e^{-\int_{-\frac{1}{2}}^{\frac{1}{2}}\log S_{xx}(f)df}\right], \end{aligned} \quad (10)$$

where $S_{xx}(f)$ denotes the power spectral density of $\{x\}$.
The bitrates corresponding to the error signals (8) and (9) remain unchanged.

a) b)

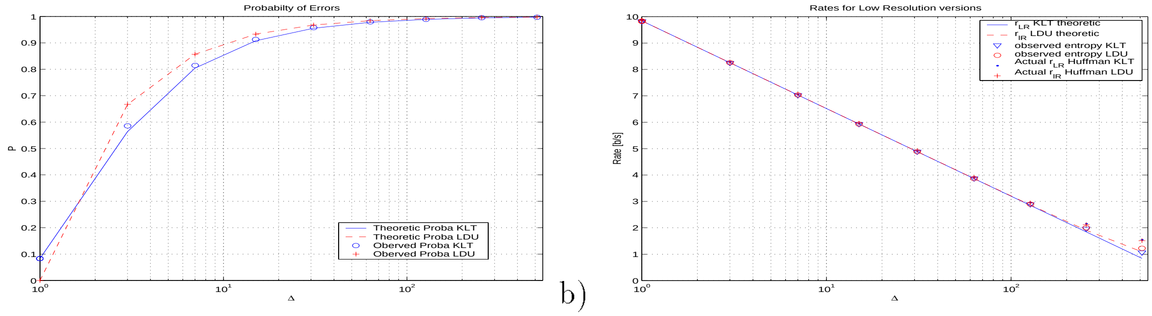

Figure 4: a) Error probability and b) rates for low resolution versions.

6 Case of Non-Gaussian p.d.f.s

In the orthogonal case, non-Gaussian p.d.f.s of the x_i should lead to Gaussian p.d.f.s of the y_i for relatively high N. This "gaussianization" effect is less pronounced in the causal case, since each transform variable is a linear combination of $i-1$ r.v.s only. Thus, the entropy of the low resolution version should also tend to (6) in the orthogonal case. In the causal case, the entropies $H(y_i^q, L)$ are $h(\sigma_{y_i}^2) - \log_2 \Delta$, where h will generally be some unknown function of the prediction error variances $\sigma_{y_i}^2$. Since the Gaussian p.d.f. maximizes the differential entropy for a given variance, $\sum_{i=1}^N H(y_i^q, L) < \frac{1}{2}\log_2 2\pi e \det R_{\underline{xx}}^{\frac{1}{N}} \approx \sum_{i=1}^N H(y_i^q, V)$. The rate $\overline{r}_{LR_L}$ will probably be less than in the Gaussian case (7), whereas due to the "gaussianization" effect, this decrease in rate should be less sensitive in the orthogonal case.
Concerning the error signal, if Δ is small in comparison with the variances of the transform signals, the analysis of the previous sections are still valid. The quantization errors in the transform domain are still uniform, leading in the signal domain to nearly Gaussian errors in the orthogonal case, and to nearly uniformly distributed errors in the causal case. Thus, the 0.25 bit suboptimality of the orthogonal transforms remains, irrespectively of the p.d.f.s of the sources.

7 Simulations

For the simulations, we generated 10^5 real Gaussian i.i.d. vectors with covariance matrix $R_{\underline{xx}} = HR_{AR1}H^T$. R_{AR1} is the covariance matrix of an AR(1) process with $\rho = 0.9$ and variance 10^5. H is a diagonal matrix whose ith entry is $(N-i+1)^{1/3}$, $N = 3$. The data are then rounded. A "one-shot" approach requires roughly 9 b/s to losslessly code these data. For the presented simulations, $\Delta =$ is odd. Figure 4b) compares the theoretic (expression (6) for the KLT, and (7) for the LDU) and observed bitrates for the low resolution signals. "Observed entropy" denotes the entropy of the transform signals as estimated of the whole set of data, and "Actual r_{LR} Huffman" is the average codewords length obtained by Huffman coding. Figure 5a) compares the theoretic (expression (8) for the KLT, and (9) for the LDU) and observed bitrates for the error signals. Finally, Figure 5b), which compares the theoretic and observed total rates for the two-stage coders in both approaches, shows that the best orthogonal approach is ≈ 0.25 bits suboptimal than its causal counterpart in most cases.

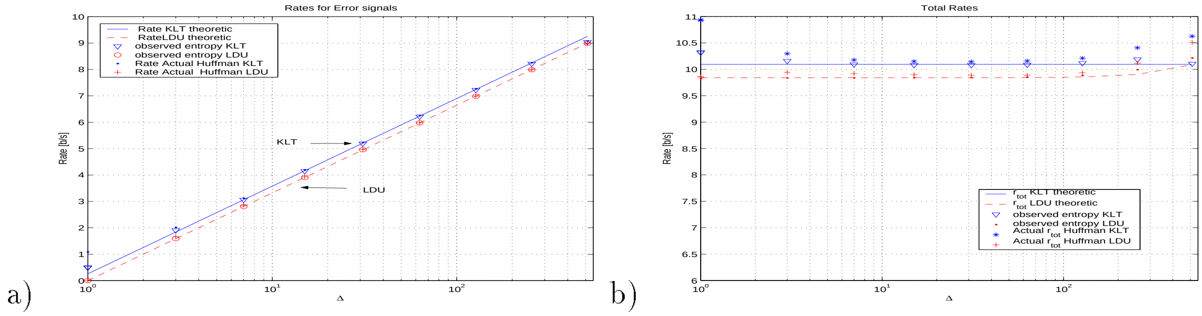

Figure 5: a)Rates for error signals and b) total rates.

References

[1] D. Mary and D. T. M. Slock, "A performance analysis of integer-to-integer transforms for lossless coding of vectorial signals," in *Proc. SMMSP'02*, (Toulouse, France), September 2002.

[2] Y. Wang, L. Yaroslavsky, M. Vilermo, and M. Vaananen, "A multichannel audio coding algorithm for inter-channel redundancy removal," in *110th AES Convention, Amsterdam*, May 2001.

[3] S. Oraintara, Y.-J. Chen, and T. Nguyen, "Integer fast fourier transform," *IEEE trans. on Signal Processing*, vol. 50, March 2002.

[4] V. Goyal, "Transform coding with integer-to-integer transforms," *IEEE trans. on Inf. Theory*, vol. 46, March 2000.

[5] M. Purat, T. Liebchen, and P. Noll, "Lossless transform coding of audio signals," in *102nd AES Convention, Munich*, 1997.

[6] M. Hans and R. Schafer, "Lossless Compression of Digital Audio," *IEEE Sig. Proc. Mag.*, vol. 18, Jul. 2001.

[7] D. Mary and D. T. M. Slock, "Multistage integer-to-integer multichannel prediction for scalable lossless coding," in *Asilomar Conf. on Signals, Systems and Computers*, (Pacific Grove, CA), Nov. 2002.

[8] S.-M. Phoong and Y.-P. Lin, "Prediction-based lower triangular transform," *IEEE trans. on Sig. Proc.*, vol. 48, July 2000.

[9] D. Mary and D. Slock, "Vectorial dpcm coding and application to wideband coding of speech," in *ICASSP 2001*, (Salt lake City), May 2001.

[10] R. M. Gray, "Toeplitz and Circulant Matrices: A review," tech. rep., Stanford University, June 1971. Revised Aug. 2002, Apr. 1993. Stanford, California 94305.