Multiple-Description Coding by Dithered Delta-Sigma Quantization

Jan Østergaard Ram Zamir

Abstract

In this paper we address the connection between the multiple-description (MD) problem and Delta-Sigma quantization. Specifically, we exploit the inherent redundancy due to oversampling in Delta-Sigma quantization, and the simple linear-additive noise model resulting from dithered lattice quantization, in order to construct a symmetric MD coding scheme. We show that the use of feedback by means of a noise shaping filter makes it possible to trade off central distortion for side distortion. Asymptotically as the dimension of the lattice vector quantizer and order of the noise shaping filter approach infinity, we show that the symmetric two-channel MD rate-distortion function for the memoryless Gaussian source and MSE fidelity criterion can be achieved at any resolution. This realization provides a new interesting interpretation for the information theoretic solution. The proposed design is symmetric in rate by construction and there is therefore no need for source splitting.

1 Introduction

Delta-Sigma analogue to digital (A/D) conversion is a technique where the input signal is highly oversampled before being quantized by a low resolution quantizer. The quantization noise is then processed by a noise shaping filter which reduces the energy of the so-called in-band noise spectrum, i.e. the part of the noise spectrum which overlaps the spectrum of the input signal. The end result is high bit-accuracy (A/D) conversion even in the presence of imperfections in the analogue components of the system, c.f. [1].

The process of oversampling and use of feedback to reduce quantization noise is not limited to A/D conversion of continuous-time signals but is in fact equally applicable to, for example, discrete time signals in which case we will use the term Delta-Sigma quantization. Hence, given a discrete time signal we can apply Delta-Sigma quantization in order to discretize the amplitude of the signal and thereby obtain a digital signal. It should be clear that the process of oversampling is not required in order to obtain a digital signal. However, oversampling leads to a controlled amount of redundancy in the digital signal. This redundancy can be exploited in order to achieve a certain degree of robustness towards a partial loss of information of the signal due to quantization and/or transmission of the digital signal over error-prone channels.

In the information theory community the problem of quantization is usually referred to as a source coding problem whereas the problem of reliable transmission is referred to as a channel coding problem. Their combination then forms a joint source-channel

Jan Østergaard (janoe@ieee.org) is with the Information and Communication Theory Group, Delft University of Technology, Delft, The Netherlands. Ram Zamir (zamir@eng.tau.ac.il) is with the Department of Electrical Engineering-Systems, Tel Aviv University , Tel Aviv, Israel. This research is partly supported by the Dutch Technology Foundation STW.

coding problem. The multiple-description (MD) problem [2], which has recently received a lot of attention, is basically a joint source-channel coding problem. The MD problem is concerned with lossy encoding of information for transmission over an unreliable K-channel communication system. The channels may break down resulting in erasures and a loss of information at the receiving side. Which of the $2^K - 1$ non-trivial subsets of the K channels that is working is assumed known at the receiving side but not at the encoder. The problem is then to design an MD system which, for given channel rates, minimizes the distortions due to reconstruction of the source using information from any subsets of the channels. Currently, the achievable MD rate-distortion region is only completely known for the case of two channels, squared-error fidelity criterion and the memoryless Gaussian source [2, 3].

Practical symmetric MD lattice vector quantization (MD-LVQ) based schemes for two descriptions have been introduced in [4, 5], which in the limit of infinite-dimensional lattices and under high-resolution assumptions, approach the symmetric MD rate-distortion bound. An extension to $K \geq 2$ descriptions was presented in [6, 7]. Asymmetric MD-LVQ allows for unequal side distortions as well as unequal side rates and was first considered in [8, 9] for the case of two descriptions and extended in [10] to the case of $K \geq 2$ descriptions. Common for all of the designs [4–10] is that a central quantizer is first applied on the source after which an index-assignment algorithm maps the reconstruction points of the central quantizer to reconstruction points of the side quantizers, which is an idea that was first presented in [11].

To avoid the difficulty of designing efficient index-assignment algorithms it was suggested in [12] that the index assignments of a two-description system can be replaced by successive quantization and linear estimation. More specifically, the two side descriptions can be linearly combined and further enhanced by a refinement layer to yield the central reconstruction. The design of [12] suffers from a rate loss of 0.5 bit/dim. at high resolution and is therefore not able to achieve the MD rate-distortion bound. Recently, however, this gap was closed by Chen et al. [13] who recognized that the rate region of the MD problem forms a polymatroid, and showed that the corner points of this rate region can be achieved by successive estimation and quantization. The design of Chen et al. is inherently asymmetric in the description rate since any corner point of a non-trivial rate region will lead to asymmetric rates. To symmetrize the coding rates, it is necessary to break the quantization process into additional stages, which is a method known as "source splitting" (following Urbanke and Rimoldi's rate splitting approach for the multiple access channel). When finite-dimensional quantizers are employed, there is a space-filling loss due to the fact that the quantizer's Voronoi cells are not completely spherical and as such each description suffers a rate loss. The rate loss of the design given in [13] is that of $2K - 1$ quantizers because source splitting is performed by using an additional $K - 1$ quantizers besides the conventional K side quantizers. In comparison, the designs based on index assignments suffer from a rate loss of only that of K quantizers (actually, when using index assignments, the space-filling loss is that of K quantizers having spherical Voronoi cells [4–7, 10]). An interesting open question is: can we avoid both the complexity of the index assignments and the loss due to source splitting in symmetric MD coding?

Inspired by the works presented in [12–14], we present a two-channel MD scheme based on two times oversampled dithered Delta-Sigma quantization, which is inherently symmetric in the description rate and as such there is no need for source splitting. The rate loss when employing finite-dimensional quantizers (in parallel) is therefore given by that of two quantizers. Asymptotically as the dimension of the vector quantizer and order of the noise shaping filter approach infinity, we show that the symmetric two-channel MD rate-distortion function for the memoryless Gaussian source and MSE fidelity criterion can be achieved at any resolution. It is worth emphasizing that our design is not limited to two descriptions but, in fact, an arbitrary number of descriptions can be created simply by increasing the oversampling ratio. However, in this paper, we focus on the case of two descriptions.

In the Delta-Sigma quantization literature, there seems to be a consensus of avoiding long feedback filters. We suspect this is mainly due to the fact that the quantization error in traditional Delta-Sigma quantization is a deterministic non-linear function of the input signal, which makes it difficult to perform an exact system analysis. Thus, there might be concerns regarding the stability of the system. In our work we use dithered (lattice) quantization, so that the quantization error is a stochastic process, independent of the input signal, and the whole system becomes linear. This linearization is highly desirable, since it allows an exact system analysis for any filter order and at any resolution. For finite filter order, it can be shown that the optimal filter coefficients are found by solving a set of Yule-Walker equations. The case of infinite filter order, which we will focus on in this paper, has a very simple solution, which (for large lattice dimension) guarantees that the proposed scheme achieves the symmetric two-channel MD rate-distortion function [2, 3].

To gain some insight into why this solution is asymptotically optimal, observe that the Delta-Sigma quantization structure resembles the nature of the optimum test channel that achieves the two-channel MD rate-distortion region [2, 3]. This channel (as shown in Fig. 3) has two additive noise branches $Y_1 = X + N_1$ and $Y_2 = X + N_2$, where the pair (N_1, N_2) is *negatively* correlated. At high resolution conditions and symmetric rates and distortions, the side reconstructions $\hat{X}_1$ and $\hat{X}_2$ become $\hat{X}_1 = Y_1$ and $\hat{X}_2 = Y_2$, while the central reconstruction $\hat{X}_c$ becomes a simple average, i.e. $\hat{X}_c = (\hat{X}_1 + \hat{X}_2)/2$. We may view the negatively correlated additive noises as adjacent samples of "high pass noise", and the averaging operation of the central reconstruction as "lowpass filtering". Intuitively, for a fixed side distortion the central distortion is reduced by shaping the spectrum of the noise to be away from the source band (the source component in Y_1 and Y_2 is the same which amounts to a lowpass signal). Thus, Delta-Sigma quantization provides a time-invariant filter version of this double branch test channel. This is further addressed in Section 3.1.

2 Dithered Delta-Sigma quantization

Let X be an i.i.d. zero-mean unit-variance Gaussian random process. Furthermore, let x denote a realization of X and let boldface letters indicate vectors.

The signal $\boldsymbol{x}$ is oversampled by a factor of two to produce the oversampled signal $\boldsymbol{a}$. It follows that $\boldsymbol{a}$ is a redundant representation of the input signal, which can be obtained simply by inserting a zero between every sample of $\boldsymbol{x}$ and apply an interpolating (ideal lowpass) filter $h(z)$ as shown in Fig. 1. At the other end of the system we apply an anti-aliasing filter $h_a(z)$ and downsample by two in order to get back to the original sample rate. After being oversampled, the signal is then quantized using entropy-coded dithered (lattice) quantization (ECDQ) [15]. ECDQ relies upon subtractive dither, which makes sure that the quantization error E is an i.i.d. zero-mean random process of variance σ_E^2. Furthermore, the quantization error is independent of the input signal and it can be assumed that the rate (or entropy) of the quantized variables is given by the conditional entropy $H(Q_L(\boldsymbol{X}+\boldsymbol{Z})|\boldsymbol{Z})$ of the L-dimensional dithered quantizer Q_L (where the conditioning is with respect to the dither sequence $\boldsymbol{Z}$). It is known that this conditional entropy is equal to the mutual information over the additive dither channel $\boldsymbol{Y}=\boldsymbol{X}+\boldsymbol{E}$ where $\boldsymbol{E}$ (the channel's noise) is distributed as $-\boldsymbol{Z}$, see [15] for details. The rate of the quantizer is therefore given by $I(\boldsymbol{X};\boldsymbol{Y}) = h(\boldsymbol{X}+\boldsymbol{E}) - h(\boldsymbol{E})$, where $I(\cdot,\cdot)$ denotes the mutual information and $h(\cdot)$ denotes the differential entropy. In the quadratic Gaussian case, if optimal pre and post filters are used, the rate redundancy over the rate-distortion function $R(D)$ of a Gaussian source satisfies [16]

$$\frac{1}{L}H(Q_L(\boldsymbol{X}+\boldsymbol{Z})|\boldsymbol{Z}) \leq R(D) + \frac{1}{2}\log_2(2\pi e G_L), \tag{1}$$

where G_L is the dimensionless normalized second moment of the L-dimensional lattice quantizer Q_L [17]. The quantity $2\pi e G_L$ is the space-filling loss of the quantizer and $\frac{1}{2}\log_2(2\pi e G_L)$ is the divergence of the quantization noise from Gaussianity. It follows that it is desirable to have Gaussian distributed quantization noise in order to make G_L as small as possible and thereby drive the rate of the quantizer towards $R(D)$. Fortunately, it is known that, there exists lattices where $G_L \to 1/2\pi e$ as $L\to\infty$ and the quantization noise of such quantizers becomes asymptotically (in dimension) Gaussian distributed in the divergence sense [18].

From the preceding arguments it is clear that we would like to use high-dimensional quantizers. However, at first sight, it might appear as the sequential scalar nature of Delta-Sigma quantization prevents the use of anything but scalar quantizers. That this is not so will soon become clear. But before going into more details about this issue we will first introduce the dithered Delta-Sigma quantization system, which is sketched in Fig. 1. As previously mentioned, an i.i.d. source sequence $\boldsymbol{x}$ is upsampled by a factor of two to yield the redundant sequence $\boldsymbol{a}$. The sequence $\boldsymbol{a}$ is then sequentially quantized on a sample by sample basis and the quantization error e_k of the kth sample is then input to the feedback filter $c^*(z) = \sum_{i=1}^{p} c_i z^{-i}$. At this point we also introduce the pth order noise shaping filter $c(z)$, which is defined as

$$c(z) \triangleq \sum_{i=0}^{p} c_i z^{-i}, \tag{2}$$

where $c_0 = 1$ so that $c(z) = 1 + c^*(z)$. The purpose of $c^*(z)$ is to predict the in-band noise component $\tilde{e}_k$ (after the synthesis) based on the past p noise samples $e_{k-1}, e_{k-2}, \dots, e_{k-p}$ as shown in Fig. 1. It follows that $\tilde{e}(z) = c^*(z)e(z)$ or equivalently $\tilde{e}_k = \sum_{i=1}^{p} c_i e_{k-i}$.

It is known that the additive noise model is exact for ECDQ and we can therefore represent the quantization operation as an additive noise term. Thus, the output of the quantizer is given by $\hat{a}_k = a_k + e_k + \tilde{e}_k$. The reconstruction error in the oversampled domain is then given by $\epsilon_k = \hat{a}_k - a_k$. Furthermore, ϵ_k is obtained by passing e_k through the noise shaping filter $c(z)$. To see this, notice that the output is $\hat{a}(z) = a(z) + e(z) + c^*(z)e(z)$ and the reconstruction error is therefore given by $\epsilon(z) = c(z)e(z)$.

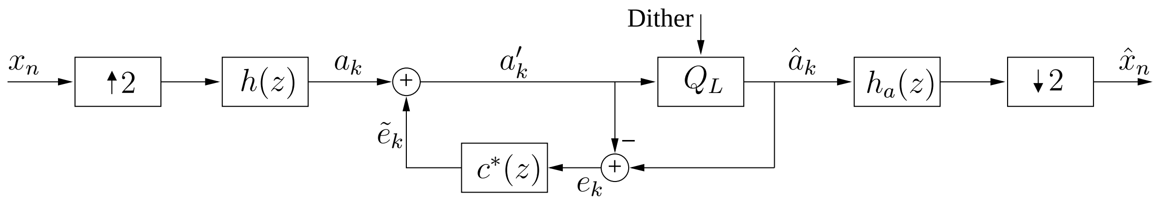

Figure 1: *Dithered Delta-Sigma quantization.*

As previously mentioned, the rate R of the ECDQ is given by the mutual information between the input and the output of the quantizer. Thus, the rate (per sample) is given by

$$R = I(A'_k; \hat{A}_k) = I(A'_k; A'_k + E_k), \tag{3}$$

where if A_k and E_k are Gaussian (where E_k is independent of of the present and past samples of A'_k by the dithered quantization assumption), then we get

$$R = \frac{1}{2}\log_2\left(1 + \frac{\text{Var}(A'_k)}{\sigma_E^2}\right), \tag{4}$$

where $\text{Var}(A'_k)$ denotes the variance of the random variable A'_k. At high resolution, $\text{Var}(A'_k) + \sigma_E^2 \approx \sigma_X^2$ which implies that

$$R \approx \frac{1}{2}\log_2\left(\frac{\sigma_X^2}{\sigma_E^2}\right). \tag{5}$$

We will now address the issue of high-dimensional quantization but first let us consider the scalar case, i.e. $L = 1$. The input to the quantizer is $a'_k = a_k + \sum_{i=1}^{p} c_i e_{k-i}$ and the output is $\hat{a}_k = a_k + \sum_{i=0}^{p} c_i e_{k-i}$. Since a'_k is a scalar the input to the quantizer is a scalar and the quantizer depicted in Fig. 1 is therefore a scalar quantizer. To justify the use of high-dimensional vector quantizers we will consider a setup involving L independent sources. These sources can, for example, be obtained by demultiplexing the scalar process X into L independent parallel i.i.d. processes $X^{(l)} = \{X_{nL+l-1}\}, \forall n \in \mathbb{Z}$ and $l = 1, \dots, L$.[1] In this case the nth sample of the lth process $X^{(l)}$ is identical to the $(n \times L + l - 1)$th sample of the original process X. In the case where $L = 2$ we have two independent scalar processes, where $X^{(1)}$ consists of the even samples of X and $X^{(2)}$ consists of the odd samples of X. The processes $X^{(1)}$ and $X^{(2)}$ are each upsampled by a factor of two so that we obtain the two processes $A^{(1)}$ and $A^{(2)}$, which each are input to a Delta-Sigma quantization system. Hence, in this case, two coders are operating in parallel and instead of a single sample a'_k we have the pair of independent samples $(a_k'^{(1)}, a_k'^{(2)})$. This makes it possible to apply two-dimensional ECDQ on the vector formed by cascading the pair of scalars. If L coders are operating in parallel, we can form the set of L independent samples $(a_k'^{(1)}, a_k'^{(2)}, \dots, a_k'^{(L)})$ and make use of L-dimensional ECDQ on the vector $(a_k'^{(1)}, a_k'^{(2)}, \dots, a_k'^{(L)})$. In general, we will allow L to become large so that, according to (1) and the paragraph that follows just below (1), the rate loss can be made arbitrarily small. Thus, for large L, E_k will be approximately Gaussian distributed.

[1] Notice that the delay between two consecutive samples of the lth process will be that of L input samples.

3 Multiple-description coding

In this section we show that the sequential dithered Delta-Sigma quantization system, which is shown in Fig. 1, can be regarded as an MD coding system.[2] For example, in the case of an oversampling ratio of two, each input sample leads to two output samples and we have in fact a two-channel MD coding system as shown in Fig. 2, where we have replaced the dithered quantizer with its additive noise model. In this case the first description is given by the even output samples and the second description by the odd output samples. The filter $h_p(z)$ corrects the phase of the second description and the post filters α and β are described in Section 3.3. The distortion due to reconstructing using both descriptions

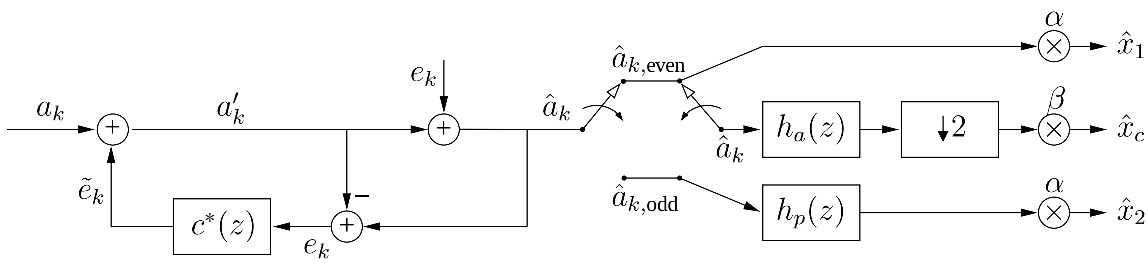

Figure 2: *Two-channel MD coding based on dithered Delta-Sigma quantization.*

is traditionally called the central distortion d_c and the distortion due to reconstructing using only a single description is called the side distortion d_s.

3.1 New interpretation of Ozarow's double branch test channel

We now show that the proposed Delta-Sigma quantization scheme, when analyzed in the frequency domain, leads to a new interpretation of Ozarow's double branch test channel shown in Fig. 3. In addition, this frequency interpretation reveals that the role of the noise

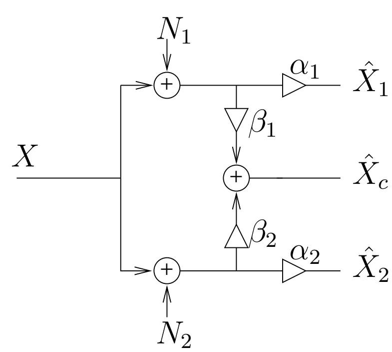

Figure 3: *The MD optimum test channel of Ozarow [3]. At high resolution* $\alpha_i = 1$ *and* $\beta_i = 1/2, i = 1, 2$ *so that* $\hat{X}_1 = Y_1, \hat{X}_2 = Y_2$ *and* $\hat{X}_c = \frac{1}{2}(\hat{X}_1 + \hat{X}_2)$.

shaping filter is not simply to shape away the quantization noise from the in-band spectrum,

---

[2] Recent related works include that of Boufounos and Oppenheim who considered Delta-Sigma quantization with deterministic quantization and finite order noise shaping for frame expansions [19]. Boufounos and Oppenheim also addressed the case of erasures in quantized frame expansions, where the transmitter is aware of the erasures [20].

as is the case in traditional Delta-Sigma quantization, but rather to delicately control the tradeoff between the in-band noise versus the out-of-band noise. This tradeoff is done while keeping the coding rate fixed, which, at least at high resolution, is equivalent to keeping σ_E^2 fixed.

The power spectrum S_X of the i.i.d. process X is constant over the complete interval $-\pi$ to π. Now recall that we assume ideal sinc interpolation when resampling. As such, since we upsample by a factor of two, the power spectrum S_A of the upsampled signal A ranges from $-\pi/2$ to $\pi/2$. This is illustrated in Fig. 4.

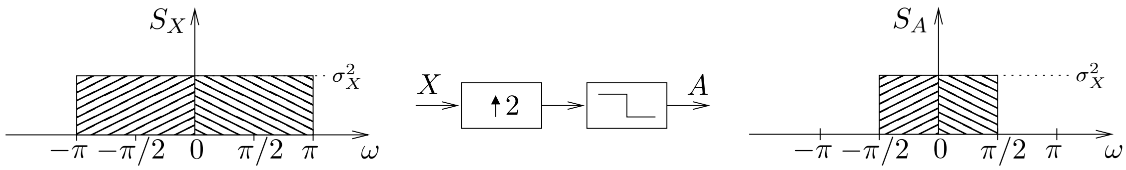

(a) Spectrum of X (b) Oversampling by two (c) Spectrum of A

Figure 4: *The power spectrum of (a) the input signal and (c) the oversampled signal. (b) illustrates the oversampling process where the input signal is first upsampled by two and then filtered by an ideal half-band lowpass filter.*

The quantization operation, which is based on ECDQ, adds a white noise signal E to the oversampled signal A, after which the feedback filter ensures that the noise gets shaped appropriately. At the decoder we apply the anti-aliasing filter (ideal lowpass filtering) and then downsample. Hence, the central distortion is given by the energy of the quantization noise that falls within the in-band spectrum. The inclusion of a noise shaping filter makes it possible to shape away the quantization noise from the in-band spectrum and thereby reduce the central distortion. By increasing the order of the noise shaping filter it is possible to reduce the central distortion accordingly.

It is also interesting to understand what influences the side distortion. Recall that the side descriptions are constructed by using either all odd samples or all even samples of the output A. Hence, we effectively downsample A by a factor of two. It is important to see that this downsampling process takes place without first applying an anti-aliasing filter. Thus, aliasing is inevitable. It follows, that not only the noise which falls within the in-band spectrum contributes to the side distortion but also the noise that falls outside the in-band spectrum (i.e. the out-of-band noise) affects the distortion. Since, in traditional Delta-Sigma quantization, the noise is shaped away from the in-band spectrum as efficiently as possible, the out-of-band noise is likely to be the dominating contributor to the side distortion. We have illustrated this in Fig. 5.

It should now be clear that, in two-channel MD Delta-Sigma quantization, the role of the noise shaping filter is to trade off the in-band noise versus the out-of-band noise. For example, in the asymptotical case where the order of the noise shaping filter goes to infinity, it is possible to construct a brick-wall filter which has a power spectrum of $1/\delta$ in the passband (i.e. for $|\omega| \leq \pi/2$) and of δ in the stopband (i.e. for $\pi/2 < |\omega| < \pi$). In this case, the central distortion is proportional to $1/\delta$ whereas the side distortion is proportional to $1/\delta + \delta$. This situation, which is illustrated in Fig. 5(b), will be discussed in more detail in the next section.

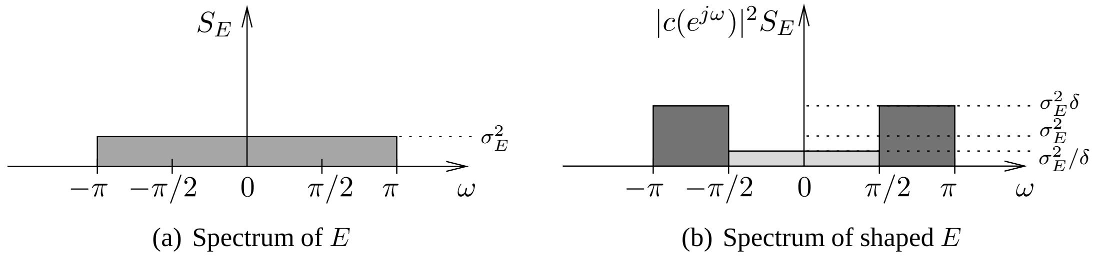

Figure 5: *The power spectrum of (a) the quantization noise (b) the shaped quantization noise. In (b) the energy of the lowpass noise spectrum (the bright region) corresponds to the central distortion and the energy of the full spectrum corresponds to the side distortion.*

Let us now revisit Ozarow's double branch test channel as shown in Fig. 3. In this model the noise pair (N_1, N_2) is negatively correlated (except from the case of no-excess marginal rates, in which case the noises are independent). Notice that this is in line with the above observations, since the highpass nature of the noise shaping filter causes adjacent noise samples to be negatively correlated. The more negatively correlated they are, the greater is the ratio of side distortion over central distortion. Furthermore, at high resolution, the filters in Ozarow's test channel degenerate and the central reconstruction is simply given by the average of the two side channels. This averaging operation can be seen as a lowpass filtering operation, which leaves the signal (since it is lowpass) and the in-band noise intact but removes the out-of-band noise.

3.2 Achieving the MD distortion product at high resolution

The two-channel MD rate-distortion region is completely characterized only in the case of memoryless Gaussian sources and MSE fidelity criterion [2, 3]. Based on the results of [3] it was shown in [21] that, at high resolution, the product of the central and side distortions of an optimal two-channel MD scheme satisfies

$$d_c d_s = \frac{\sigma_X^4}{4}\frac{1}{1-d_c/d_s}2^{-4R}. \tag{6}$$

Lemma 3.1. *At high resolution and asymptotically as $p \to \infty$ the distortion product given by (6) is achievable.*

Proof. The central distortion is equal to the total energy P_{d_c} of the in-band noise spectrum where

$$P_{d_c} = \frac{\sigma_E^2}{2\pi}\int_{-\pi/2}^{\pi/2} |c(e^{j\omega})|^2 d\omega. \tag{7}$$

The side distortion is equal to the energy P_{d_s} of the in-band noise spectrum of the side descriptions which contains aliasing due to the subsampling process. Since we downsample by two we have

$$P_{d_s} = \frac{\sigma_E^2}{4\pi}\int_{-\pi}^{\pi} |c(e^{j\omega/2})|^2 + |c(e^{j(\omega/2+\pi)})|^2 d\omega. \tag{8}$$

The noise shaping filter is part of a feedback loop and it must therefore be minimum phase. Since $c_0 = 1$ it can be shown that the spectrum of such a minimum phase filter satisfies

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} \log_e |c(e^{j\omega})|^2 d\omega = 0. \tag{9}$$

It follows that the area under $\log_e(|c(e^{j\omega})|^2)$ must be equally distributed above and below the 0 dB line. It is easy to see that if we let $|c(e^{j\omega})|^2 = 1/\delta$ for $|\omega| \leq \pi/2$ and $|c(e^{j\omega})|^2 = \delta$ for $\pi/2 < |\omega| < \pi$ where $0 < \delta \in \mathbb{R}$ then (9) is satisfied. Hence, for any $\delta > 0$ it follows from (8) that $d_s = \frac{1}{2}\sigma_E^2(\delta + \delta^{-1})$ and from (7) we see that $d_c = \frac{1}{2}\sigma_E^2/\delta$ which yields the distortion product $d_c d_s = \frac{\delta+\delta^{-1}}{4\delta}\sigma_E^4$.

Under high resolution assumptions the description rate R depends only upon the ratio σ_X^2/σ_E^2 and as such it is independent of the noise shaping filter. In this case $R \approx \log_2(\sigma_X^2/\sigma_E^2)$ which implies that $\sigma_E^4 \approx \sigma_X^4 2^{-4R}$. Finally, since $d_c/d_s = \delta^{-1}/(\delta + \delta^{-1})$ it follows that

$$\frac{1}{1 - d_c/d_s} = \frac{\delta + \delta^{-1}}{\delta}.$$

□

3.3 Optimal post filters for two descriptions at general resolution

Let $p \to \infty$ and let the side distortion be given by $\sigma_E^2(\delta + \delta^{-1})/2$ and the central distortion by $\sigma_E^2\delta^{-1}/2$. It can then be shown that the post filters $\alpha = \alpha_i$ and $\beta = \beta_i, i = 1, 2,$ are given by $\alpha = \frac{\sigma_X^2}{\sigma_X^2 + \sigma_E^2(\delta+\delta^{-1})/2}$ and $\beta = \frac{\sigma_X^2}{2\sigma_X^2 + \sigma_E^2\delta^{-1}}$. It can also be shown that the side distortion when using post filters is given by

$$d_s = \frac{\sigma_X^2\sigma_E^2(\delta + \delta^{-1})}{2\sigma_X^2 + \sigma_E^2(\delta + \delta^{-1})}, \tag{10}$$

and the central distortion when using post filters is given by

$$d_c = \frac{\sigma_X^2\sigma_E^2\delta^{-1}}{2\sigma_X^2 + \sigma_E^2\delta^{-1}}. \tag{11}$$

3.4 Achieving the symmetric two-channel rate-distortion function

Let us first recall the solution to the quadratic Gaussian MD problem as proven by Ozarow [3], i.e. the set of achievable distortions given the description rate R which is the union of all distortion pairs $(\bar{d}_c, \bar{d}_s)$ satisfying

$$\bar{d}_s \geq \sigma_X^2 2^{-2R} \tag{12}$$

and

$$\bar{d}_c \geq \frac{\sigma_X^2 2^{-4R}}{1 - (\sqrt{\Pi} - \sqrt{\Delta})^2}, \tag{13}$$

where $\Pi = (1 - \bar{d}_s/\sigma_X^2)^2$ and $\Delta = \bar{d}_s^2/\sigma_X^4 - 2^{-4R}$.

Lemma 3.2. *The side distortion given by (10) and the central distortion given by (11) achieve the lower bound of Ozarow's MD rate-distortion function.*

Proof. The description rate, at general resolution, can be shown to be given by

$$R = \frac{1}{2}\log_2\left(\frac{\sigma_X^2 + \sigma_E^2(\delta + \delta^{-1})/2}{\sigma_E^2}\right). \tag{14}$$

By expressing Π and Δ via (10) and (14) it can be shown that (11) is identical to (13). □

References

[1] J. C. Candy and G. C. Temes, Eds., *Oversampling Delta-Sigma Data Converters: Theory, Design and Simulation*. IEEE Press, 1992.

[2] A. A. E. Gamal and T. M. Cover, “Achievable rates for multiple descriptions,” *IEEE Trans. Inf. Theory*, vol. IT-28, no. 6, pp. 851 – 857, November 1982.

[3] L. Ozarow, “On a source-coding problem with two channels and three receivers,” *Bell System Technical Journal*, vol. 59, pp. 1909 – 1921, December 1980.

[4] S. D. Servetto, V. A. Vaishampayan, and N. J. A. Sloane, “Multiple description lattice vector quantization,” in *Proc. Data Compression Conf.*, March 1999, pp. 13 – 22.

[5] V. A. Vaishampayan, N. J. A. Sloane, and S. D. Servetto, “Multiple-description vector quantization with lattice codebooks: Design and analysis,” *IEEE Trans. Inf. Theory*, vol. 47, no. 5, pp. 1718 – 1734, July 2001.

[6] J. Østergaard, J. Jensen, and R. Heusdens, “n-channel symmetric multiple-description lattice vector quantization,” in *Proc. Data Compression Conf.*, March 2005, pp. 378 – 387.

[7] ——, “n-channel entropy-constrained multiple-description lattice vector quantization,” *IEEE Trans. Inf. Theory*, vol. 52, no. 5, pp. 1956 – 1973, May 2006.

[8] S. N. Diggavi, N. J. A. Sloane, and V. A. Vaishampayan, “Design of asymmetric multiple description lattice vector quantizers,” in *Proc. Data Compression Conf.*, March 2000, pp. 490 – 499.

[9] ——, “Asymmetric multiple description lattice vector quantizers,” *IEEE Trans. Inf. Theory*, vol. 48, no. 1, pp. 174 – 191, January 2002.

[10] J. Østergaard, R. Heusdens, and J. Jensen, “n-channel asymmetric multiple-description lattice vector quantization,” in *Proc. IEEE Int. Symp. Information Theory*, September 2005, pp. 1793 – 1797.

[11] V. A. Vaishampayan, “Design of multiple description scalar quantizers,” *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 821 – 834, May 1993.

[12] Y. Frank-Dayan and R. Zamir, “Dithered lattice-based quantizers for multiple descriptions,” *IEEE Trans. Inf. Theory*, vol. 48, no. 1, pp. 192 – 204, January 2002.

[13] J. Chen, C. Tian, T. Berger, and S. S. Hemami, “A new class of universal multiple description lattice quantizers,” in *Proc. IEEE Int. Symp. Information Theory*, Adelaide, Australia, September 2005, pp. 1803 – 1807.

[14] N. S. Jayant, “Subsampling of a DPCM speech channel to provide two self-contained, half-rate channels,” *Bell Syst. Tech. Jour.*, p. 501 509, April 1981.

[15] R. Zamir and M. Feder, “On universal quantization by randomized uniform/lattice quantizer,” *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 428 – 436, March 1992.

[16] ——, “Information rates of pre/post-filtered dithered quantizers,” *IEEE Trans. Inf. Theory*, vol. 42, no. 5, pp. 1340 – 1353, September 1996.

[17] J. H. Conway and N. J. A. Sloane, *Sphere packings, Lattices and Groups*, 3rd ed. Springer, 1999.

[18] R. Zamir and M. Feder, “On lattice quantization noise,” *IEEE Trans. Inf. Theory*, vol. 42, pp. 1152 – 1159, 42 1996.

[19] P. T. Boufounos and A. V. Oppenheim, “Quantization noise shaping on arbitrary frame expansions,” *EURASIP Journal on Applied Signal Processing*, vol. 2006, pp. 1 – 12, 2006.

[20] ——, “Compensation of coefficient erasures in frame representations,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, vol. 3, May 2006, pp. 848 – 851.

[21] V. A. Vaishampayan, J.-C. Batllo, and A. Calderbank, “On reducing granular distortion in multiple description quantization,” in *Proc. IEEE Int. Symp. Information Theory*, August 1998, p. 98.