On Universal Quantization by Randomized Uniform/Lattice Quantizers

Ram Zamir and Meir Feder*

Dept. of Electrical Engineering - Systems
Faculty of Engineering
Tel-Aviv University
Tel-Aviv, 69978, ISRAEL

Abstract

Uniform quantization with dither, or lattice quantization with dither in the vector case, followed by a universal lossless source encoder (entropy coder), is a simple procedure for universal coding with distortion of a source that may take continuously many values. The rate of this universal coding scheme is examined, and we derive a general expression for it. An upper bound for the redundancy of this scheme, defined as the difference between its rate and the minimal possible rate, given by the rate distortion function of the source, is derived. This bound holds for all distortion levels. Furthermore, we present a composite upper bound on the redundancy as a function of the quantizer resolution which leads to a tighter bound in the high rate (low distortion) case.

Key Words: Uniform and Lattice Quantization, Randomized Quantization, Universal Coding, Rate-Distortion Performance

*Meir Feder was also supported by The Andrew W. Mellon Foundation, Woods Hole Oceanographic Institution

1 Introduction and Summary of Results

A. The Randomized Uniform/Lattice Quantizer

Uniform quantization with dither, or more generally lattice quantization with dither, followed by a universal lossless source encoder (entropy encoder) is a simple procedure for universal coding with distortion of a source that may take continuous real values. This procedure is universal since it does not depend on the source statistics. Due to the dither, the distortion in this procedure is independent of the source value. In this paper we consider the rate performance, i.e. the entropy, of the uniform (lattice) randomized quantizer, as compared with the optimal rate given by the rate-distortion function of the source.

The uniform quantizer with dither, or the randomized quantizer, is defined as follows. The code points of the uniform quantizer are $\{0, \pm\Delta, \pm 2\Delta, \ldots\}$. The quantizer function $Q : R \to R$ is such that,

$$Q(x) = i\Delta, \quad \text{for} \quad i\Delta - \Delta/2 \le x \le i\Delta + \Delta/2$$

Let Z be a random variable, distributed uniformly in the interval $[-\Delta/2, \Delta/2]$. The universal quantizer with dither represents a source value x as,

$$u = Q(x+z) - z \tag{1}$$

where z is a sample of the r.v. Z. It is easy to show that for any differece distortion measure, $\delta(x-u)$, e.g., distortion measure of the form $|x-u|^r$, the average error of this quantizer is independent of x, i.e.,

$$E_z\{\delta(Q(x+z) - z - x)\} = \frac{1}{\Delta}\int_{-\Delta/2}^{\Delta/2} \delta(\alpha) d\alpha = \epsilon \tag{2}$$

The generalization of the uniform quantizer to the vector case is the lattice quantizer. The code points of a K-dimensional lattice quantizer form a K-dimensional lattice $\{L_K\}$. The quantizer $Q_K(\cdot)$ maps every vector $\underline{x} \in R^K$ into the nearest lattice point $\underline{l}_i \in L_K$; the region of all K-vectors mapped into a lattice point $\underline{l}_i \in L_K$ is the Voronoi region,

$$V(\underline{l}_i) = \{\underline{x} \in R^K : \|\underline{x} - \underline{l}_i\| \le \|\underline{x} - \underline{l}_j\|, \ \text{for all} \ j \ne i\}$$

Clearly, the uniform scalar quantizer is the special case, Q_1, of the lattice quantizer. Now let $\underline{Z}$ be a K-dimensional random variable uniformly distributed over the basic cell of the lattice, V_0; this cell is the Voronoi region of the lattice point $\underline{0}$. The lattice quantizer with dither represents a source vector $\underline{x} \in R^K$ by a vector,

$$\underline{u} = Q_K(\underline{x} + \underline{z}) - \underline{z}, \qquad \underline{u} \in R^K \tag{3}$$

where $\underline{z}$ is a sample from the random vector $\underline{Z}$. As for the scalar case, one can easily show that for a difference distortion measure, the average distortion is independent of the value, $\underline{x}$.

The quantizers above can be used to encode a source vector, $\underline{x} \in R^n$, as follows. In the scalar case, a dither is added to each source component and the result is then quantized component by component. In the lattice case we assume that K divides n and the input is considered as a concatenation of n/K K-dimensional vectors, quantized independently. In both cases the entropy coder will then take into account the statistical properties of the entire n-dimensional vector.

B. Previous Results

In [1],[2], the rate of these quantizers, in encoding a general n-dimensional source, was compared to the rate achieved by the optimal, entropy-constrained, vector quantizer (ECVQ). This optimal n-dimensional ECVQ can be designed, e.g., using the technique of [5]. It was shown, for example, that for a block of length n, emitted from any source, $\underline{X}$,

$$\frac{1}{n}H\left(Q_1(\underline{X}+\underline{Z}) - \underline{Z} \mid \underline{Z}\right) = \frac{1}{n}H\left(Q_1(\underline{X}+\underline{Z}) \mid \underline{Z}\right) \leq \frac{1}{n}H\left(Q^n_{opt}(\epsilon;\underline{X})\right) + \frac{1}{2}\log\frac{4\pi e}{12} \tag{4}$$

where $Q_1(\cdot)$ is the randomized scalar quantizer having a step size Δ, $\epsilon = \Delta^2/12$ is its mean square error, $Q^n_{opt}(\epsilon;\underline{X})$ is the optimal n-dimensional ECVQ for the source $\underline{X}$ having a mean square error ϵ, $\underline{Z}$ is the dither random vector, distributed uniformly, which can be generated simultaneously by the transmitter and the receiver, and $H(\cdot)$ is the entropy. Throughout the paper, the logarithm base is 2, and the entropy is measured in bits. Note that $\frac{1}{2}\log\frac{4\pi e}{12} \approx 0.754$ bits. For notational short-hand we will sometimes use $H(Q_1 \mid \underline{Z})$ for $H(Q_1(\underline{X}+\underline{Z}) \mid \underline{Z})$ and $H(Q^n_{opt})$ for $H(Q^n_{opt}(\epsilon;\underline{X}))$.

In [1], the bound in (4) has also been generalized to the case of lattice quantizer with dither, and a mean square error distortion measure, to yield the bound,

$$\frac{1}{n}H(Q_K \mid \underline{Z}) \leq \frac{1}{n}H(Q^n_{opt}) + \frac{1}{2}\log 4\pi e G_K \tag{5}$$

where G_K is the generalized second moment of the lattice, [6], i.e.

$$G_K = \frac{1}{K} \frac{\int_{V_0} \|\underline{x} - \hat{\underline{x}}\|^2 d\underline{x}}{\left(\int_{V_0} d\underline{x}\right)^{1+2/K}} \quad \underline{x}, \hat{\underline{x}} \in R^K \tag{6}$$

and $\hat{\underline{x}}$ is the centroid of the polytope V_0. Note that $G_1 = 1/12$, and the minimum value of $G_K \to 1/2\pi e$ as $K \to \infty$, so that the minimum value of $1/2 \log 4\pi e G_K \to 0.5$. In [2], bounds similar to (5) have been derived for other distortion criteria.

C. Summary of New Results

The important observation made in this paper is that the rate of the randomized uniform/lattice quantizer can be written as the mutual information between the input and the output of the additive noise channel shown in Figure 1. The input to this auxliary channel is the source to be quantized and the additive noise is uniformly distributed over the mirror image of the quantizer basic cell, i.e., almost as the dither. In other words, denoting the additive noise by $\underline{N}$, where its p.d.f. $f_{\underline{N}}(\underline{n}) = f_{\underline{Z}}(-\underline{n})$, the source input vector by $\underline{X}$ and the output by $\underline{Y} = \underline{X} + \underline{N}$, then,

$$H(Q \mid \underline{Z}) = I(\underline{X}; \underline{Y}) = H(\underline{X} + \underline{N}) - H(\underline{N}) \tag{7}$$

As will be shown, this general result allows us to calculate universal upper bounds on the rate of the randomized quantizers.

The interesting quantity for universal coding with distortion, is the excess rate over the rate-distortion function of that coding scheme. We thus define the redundancy of the randomized lattice quantizer, of dimension K, at the distortion level ϵ, as,

$$\rho_{K,n}(\epsilon) = \frac{1}{n} H(Q_K | \underline{Z}) - R_n(\epsilon) \tag{8}$$

where $R_n(\epsilon)$ is the rate distortion function of an n-dimensional random vector,

$$R_n(\epsilon) = \inf_{\{f(\underline{u}/\underline{x}) : \frac{1}{n}\overline{\delta} \leq \epsilon\}} \frac{1}{n} I(\underline{X}; \underline{U}) \tag{9}$$

and where $f(\underline{u}/\underline{x})$ is a conditional p.d.f of the representation, $\underline{u}$, given the source, $\underline{x}$, the term $\frac{1}{n}\overline{\delta} = \frac{1}{n} E_{\underline{u},\underline{x}}\{\delta(\underline{u} - \underline{x})\}$ is the average distortion, per source symbol, between the source and its representation, and $I(\underline{X}; \underline{U}) = I(X_1, \ldots, X_n; U_1, \ldots, U_n)$ is the mutual information between the random vectors $\underline{X}$ and $\underline{U}$.

One of the main results of the paper is that

$$\rho_{K,n}(\epsilon) \leq C \tag{10}$$

where $C = C(\delta(\cdot), K)$ is the *capacity* of the channel of Figure 1 under the constraint that the input $\underline{X}$ satisfies

$$\frac{1}{n} E\{\delta(\underline{X})\} \leq \epsilon$$

where $\epsilon = \frac{1}{n} E\{\delta(\underline{N})\}$ is the average distortion of the noise. This upper bound is true for all sources and for all distortion levels. We note that, e.g., for a square error distortion measure, this capacity, at dimension K, can be bounded by

$$C \leq \frac{1}{2} \log 4\pi e G_K \tag{11}$$

This is the same value achieved by Ziv, [1]; however, we actually improve upon Ziv's result since we have bounded the difference between the rate of the universal quantizer and the rate-distortion function, while in [1] the upper bound was for the excess rate as compared to the optimal n-dimensional quantizer, whose rate is always greater than the rate-distortion function. We will also observe that our upper bound is reachable, and thus it cannot be improved by a constant bound that holds for all distortion levels. As will be seen the bound is reached in a case that can be described as "low resolution quantization".

Another result considers bounds on the redundancy, $\rho_{K,n}(\epsilon)$, as a function of the distortion, ϵ. We were motivate to derive these results by the upper bounds, tighter than (10) and (11), available for deterministic uniform and lattice quantizers, in cases which can be characterized as "high resolution", see [3] and [4]. For a square error distortion measure the bounds on the redundancy that we came up with can be written as,

$$\rho_{K,n}(\epsilon) \leq \frac{1}{2} \log 2\pi e G_K + \frac{1}{2} \log \frac{P(\underline{X}+\underline{N})}{P(\underline{X})} \tag{12}$$

where the first term is the bound of [3],[4], which holds for deterministic quantizers under the "high resolution" assumption made there, and the second term involves the entropy power, $P(\underline{X})$, of the source, and the entropy power, $P(\underline{X}+\underline{N})$, of the source with the additive noise of Figure 1. Note that at high resolution, $P(\underline{X}+\underline{N}) \approx P(\underline{X})$, and thus the upper bound of [3] and [4], proved there for deterministic quantizers, holds for randomized quantizers too.

The paper is organized as follows. In section 2 we derive the general result on the rate of the uniform/lattice randomized quantizers and in section 3 we derive the universal bound, C, on the redundancy, starting with scalar quantization and memoryless sources and then generalize to arbitrary n-dimensional sources and lattice quantizers. In section 4 we derive the bound on the redundancy that depends on the distortion and then emphasize the high resolution case. A few experimental results and insights are presented in section 5. Summary and suggestion for further research will conclude the paper.

2 The Rate of the Randomized Uniform/Lattice Quantizer

The main observation made in this paper is that the rate of the randomized quantizer can be expressed via a simple formula that involves the mutual information between input and output of the channel of Figure 1. This result is initially presented for the scalar case, and it is then extended to the vector case and lattice quantizers.

A. Scalar Case - The Rate of the Randomized Uniform Quantizer

The scalar quantizer with dither generates an output, given by (1), which is a function of the source value, x, sampled from the source X, and the dither value, z, sampled from the dither random variable Z. In a randomized quantization procedure the dither value is pseudo-noise known to the receiver. Thus, we are interested in the conditional entropy of the quantizer output, given Z, i.e., $(Q(X+Z)-Z \mid Z) = H\left(Q(X+Z) \mid Z\right)$. For notational simplicity, as above, this entropy will be denoted $H(Q|Z)$. The following theorem provides an expression for $H(Q|Z)$:

Theorem 1 *The entropy of the randomized uniform quantizer, with step-size Δ, is given by*

$$H(Q|Z) = I(X;Y) = H(X+N) - H(N) = H(Y) - \log\Delta \tag{13}$$

where $Y = X+N$, X is the source, N is a random variable independent of X distributed uniformly over $[-\Delta/2, \Delta/2]$, $Y = X+N$, and $I(X;Y)$ is the mutual information between X and Y.

The proof of this theorem is given in Appendix A.

The expression for the rate of the quantizer, as provided by this theorem, depends on the distortion via the uniform quantizer step size, Δ. For any given difference distortion measure, $\delta(\cdot)$, we can use the relation,

$$\epsilon = E\{\delta(Z)\} = \frac{1}{\Delta}\int_{-\Delta/2}^{\Delta/2} \delta(\alpha)d\alpha = \phi(\Delta) \tag{14}$$

and get, by substituting $\Delta = \phi^{-1}(\epsilon)$, the rate as a function of the average distortion, ϵ.

When we consider a vector source, $\underline{X} = X_1 \ldots X_n$, but still use a scalar uniform quantizer and i.i.d. dither values $\underline{Z} = Z_1 \ldots Z_n$, it is easy to show, by a straight-forward extension of theorem 1, that the rate, per symbol, of the scalar randomized quantizer, denoted $H(Q_1 \mid \underline{Z})$ is given by,

$$\frac{1}{n}H(Q_1 \mid \underline{Z}) = \frac{1}{n}I(\underline{X};\underline{Y}) = \frac{1}{n}H(\underline{Y}) - \log\Delta \tag{15}$$

where $\underline{Y} = \underline{X} + \underline{N}$ is the output of the auxiliary channel with input $\underline{X}$. The the quantizer step size now defines the distortion per symbol, ϵ, via,

$$\epsilon = \frac{1}{n \cdot \Delta^n} \int_{D^n} \delta(\underline{\alpha}) d\underline{\alpha} = \frac{1}{\Delta} \int_{-\Delta/2}^{\Delta/2} \delta(\alpha_i) d\alpha_i \tag{16}$$

where D^n is the n dimensional cube, $\{\underline{u} \in R^n : -\Delta/2 \leq u_i \leq \Delta_2\}$, and where the last equality holds when the distortion measure satisfies $\delta(\underline{x} - \underline{u}) = \sum_{i=1}^{n} \delta(x_i - u_i)$, i.e., it is a single symbol distortion measure.

B. Vector Case - The Rate of the Randomized Lattice Quantizer

The most general case we consider is the quantization of an n-dimensional vector source using a randomized K-dimensional lattice quantization. As noted before, we assume that K divides n; the source vector is divided into n/K vectors of dimension K and each sub-vector is coded independently. The randomized lattice quantization is performed by first adding to each source sub-vector, $\underline{X}_K$, of dimension K, a dither, $\underline{Z}_K$, that is uniformly distributed over the lattice Voronoi cell, V_0, and then representing the result by the nearest lattice point.

In appendix A we also prove the following generalization of theorem 1. Let $\underline{N}_K$ be a random K-dimensional vector distributed uniformly over V_0^-, the reflection of V_0, i.e.

$$V_0^- = \{\underline{x} : -\underline{x} \in V_0\}$$

Then the rate of the randomized lattice quantizer for each sub-vector is given by,

$$H(Q_K|\underline{Z}_K) = I(\underline{X}_K; \underline{Y}_K) = H(\underline{Y}_K) - \log V \tag{17}$$

where $Y_K = \underline{X}_K + \underline{N}_K$ such that $\underline{X}_K$, $\underline{N}_K$ are independent, and $V = \int_{V_0^-} d\underline{\alpha} = \int_{V_0} d\underline{\alpha}$ is the volume of the K-dimensional Voronoi cell of the lattice.

When the entire n-dimensional source vector $\underline{X}$ is encoded, the average rate will be given by,

$$\frac{1}{n} H(Q_K|\underline{Z}) = \frac{1}{n} I(\underline{X}; \underline{Y}) = \frac{1}{n} H(\underline{Y}) - \frac{1}{K} \log V \tag{18}$$

where $\underline{Z}, \underline{X}, \underline{N}$ and $\underline{Y} = \underline{X} + \underline{N}$ are concatenations of n/K sub-vectors of dimension K, as defined above. Thus, if these subvectors are i.i.d., $H(\underline{Y}) = (n/K) \cdot H(\underline{Y}_K)$, $I(\underline{X}; \underline{Y}) = (n/K) \cdot I(\underline{X}_K; \underline{Y}_K)$, etc. The resullt, (18), however, holds also when the subvectors are dependent.

As in the scalar case, the average distortion per symbol, ϵ, can be expressed in terms of the volume V,

$$\epsilon = \frac{1}{n \cdot V^{n/K}} \int_{(V_0^-)^{n/K}} \delta(\underline{n}) d\underline{n} = \frac{1}{K \cdot V} \int_{V_0^-} \delta(\underline{n}_K) d\underline{n}_K \tag{19}$$

where $(V_0^-)^{n/K}$ is the n/K-fold cartesian product of the cell V_0^-, and the last equality holds when the distortion measure can be expressed as a sum of K-dimensional distortion measures.

3 Universal Upper Bound on the Redundancy of Randomized Quantizers

The expression for the randomized quantizer entropy, calculated in theorem 1 and its generalizations above, does not provide insight to its performance with respect to the optimal performance. The main result of this section, presented in theorem 2 below, is a derivation of an upper bound for the redundancy, i.e. to the difference between this entropy and the rate distortion function.

A. Scalar Uniform Quantizer

We start again with the scalar conditional entropy $H(Q \mid Z)$. Before we state the theorem, consider the auxiliary channel of Figure 1. Let $\delta(\cdot)$ be a distortion measure and let the input to the channel, X, be constrained so that,

$$E\{\delta(X)\} = \int \delta(x) f_X(x) dx \le \epsilon \tag{20}$$

where $\epsilon = \frac{1}{\Delta}\int_{-\Delta/2}^{\Delta/2} \delta(\alpha)d\alpha$, is the average distortion of the noise, N. The *capacity* of this channel, with the constraint (20), is given by,

$$C = \sup_{\{f_X(x): E\{\delta(X)\} \le \epsilon\}} I(X; X+N) \tag{21}$$

With the definition of this capacity we can state the upper bound:

Theorem 2 *The redundancy of the scalar uniform randomized quantizer satisfies*

$$\rho_{1,1}(\epsilon) = H(Q|Z) - R(\epsilon) \le C \tag{22}$$

where $\epsilon = \frac{1}{\Delta}\int_{-\Delta/2}^{\Delta/2} \delta(\alpha)d\alpha$*,* $\delta(\cdot)$ *is the distorsion measure,* Δ *is the quantizer step size,* C *is defined above, (21), and* $R(\epsilon)$ *is the rate distortion function.*

Proof: Let U be any random variable, which in general depends on the source X but is independent of the noise, N, in the auxiliary channel, that satisfies,

$$E\{\delta(X-U)\} \le \epsilon \tag{23}$$

We recall that the rate distortion function is the minimum of $I(X;U)$ over all U that satisfy (23). Now, from theorem 1, $H(Q|Z) = I(X;Y)$; thus we can write,

$$H(Q|Z) - I(X;U) = I(X;Y) - I(X;U) \tag{24}$$

where $Y = X + N$ is the output of the auxiliary channel. In Appendix B it is shown that,

$$I(X;Y) - I(X;U) = I(X;Y|U) - I(X;U|Y) \leq I(X;Y|U) \tag{25}$$

but, the noise N in the channel is independent of U; thus,

$$I(X;Y|U) = H(Y|U) - H(N) = H(Y-U|U) - H(N) \leq H(Y-U) - H(N) = I(X-U;Y-U) \tag{26}$$

Now, using (23), and recalling the definition of C, we get,

$$I(X-U;Y-U) \leq C \tag{27}$$

Combining (24), (25), (26) and (27), we get that $H(Q|Z) - I(X;U) \leq C$. Since this holds for all U that satisfy (23) it also true for the U that achieves the rate-distortion function, and the theorem is proved. $\square$

We note the following interesting observation. Let X^* be the random variable that achieves the capacity in (21). When X^* is encoded by the universal random quantizer, its rate achieves the upper bound of (22). To see this, note that $E\{\delta(X^*)\} \leq \epsilon$; thus its rate distortion, denoted $R(\epsilon; X^*)$, for distortion ϵ, is zero since at that distortion we can choose $U \equiv 0$, which satisfies the constraint, to achieve the rate distortion function. Since $H\left(Q(X^*+Z)|Z\right) = I(X^*;Y)$ where $Y = X^* + N$, and since X^* achieves the capacity,

$$H\left(Q(X^*+Z)|Z\right) - R(\epsilon; X^*) = H\left(Q(X^*+Z)|Z\right) = I(X^*; X^*+N) = C \tag{28}$$

In appendix C the capacity, (21), is investigated further. It is shown there, that, for a square error distortion measure, this capacity can be bounded by,

$$\frac{1}{2}\log\left(1+\frac{2\pi e}{12}\right) < C < \frac{1}{2}\log\frac{4\pi e}{12} \tag{29}$$

where the upper bound is achieved if the channel output, Y, is Gaussian. This bound, $\frac{1}{2}\log\frac{4\pi e}{12} \approx 0.754$ bits, is the same as the bound calculated in [1] for the difference between the entropy of the randomized quantizer and the optimal quantizer. We note, however, that since a Gaussian r.v. can be generated only as the sum of two independent Gaussian r.v's, this upper bound cannot be achieved, and so our bound, C, is strictly smaller then $\frac{1}{2}\log\frac{4\pi e}{12}$.

Appendix C also contain upper and lower bounds for this channel capacity for distortion measures of the form $|x-u|^r$.

Theorem 2 can be generalized, and we can claim that for any n-dimensional source, the rate of the scalar randomized quantizer having distortion ϵ is greater then the rate distortion function, for distortion measures as above, by no more than C bits per sample, i.e.,

$$\frac{1}{n}H(Q_1|\underline{Z}) - R_n(\epsilon) \leq C \tag{30}$$

where now

$$C = \sup_{\{f_{\underline{X}}(\underline{x}):\frac{1}{n}\int \delta(\underline{x})f_{\underline{X}}(\underline{x})d\underline{x}\leq\epsilon\}} \frac{1}{n}I(\underline{X};\underline{X}+\underline{N}) \tag{31}$$

is the power constrained capacity of the channel in Figure 1, and $R_n(\epsilon)$ is the n-dimensional rate-distortion function of the source defined in (9). When the distortion measure is a single symbol distortion measure, then since the components of the additive noise $\underline{N}$ are independent, i.e. the auxiliary channel is memoryless, the capacity (31) is exactly the capacity calculated in (21).

It is straight-forward to show this generalized result since all the variables in the proof of theorem 2 and in Appendix B can be n-dimensional random vectors. We also note that the source that achieves the bound, (30), is the memoryless source X^* described above.

B. Randomized Lattice Quantizers

Following the derivation of theorem 2 and Appendix B we can also generalize theorem 2 to the lattice case, and claim that for any n-dimensional source, the rate of the K-dimensional (where K divides n) randomized lattice quantizer, Q_K, having a distortion ϵ, is greater then the rate distortion function, by no more then C bits per sample, i.e.,

$$\frac{1}{n}H(Q_K|\underline{Z}) - R_n(\epsilon) \leq C \tag{32}$$

where C is the constrained capacity of the additive noise channel in which each K components of the noise vector, $\underline{N}$, are independent and are uniformly distributed over V_0^-, and the input p.d.f. is constrained to satisfy $\frac{1}{n}\int \delta(\underline{x})f_{\underline{X}}(\underline{x})d\underline{x} \leq \epsilon$, that is,

$$C = \sup_{\{f_{\underline{X}}(\underline{x})\,:\,\frac{1}{n}\int \delta(\underline{x})f_{\underline{X}}(\underline{x})d\underline{x}\leq\epsilon\}} \frac{1}{n}I(\underline{X};\underline{X}+\underline{N}) \tag{33}$$

The bound in (32) is tight in the sense that we can find a source for which the difference between the quantizer rate and the rate distortion function is exactly C. As was the case for the scalar randomized quantizer, this is the source that achieves the capacity of the additive channel described above. This source is composed of independent and identically distributed K-dimensional sub-vectors, $\underline{X}_K$, whose probability density function satisfies

$$f_{\underline{X}_K}(\underline{x}_K) = \arg\max I(\underline{X}_K; \underline{X}_K + \underline{N}_K)$$

where the maximization is over all source p.d.f which satisfies $\frac{1}{K}\int_{R^K} f_{\underline{X}_K}(\underline{x}_K)\delta(\underline{x}_K)d\underline{x}_K \leq \epsilon$

Consider again, more closely, the case of the square error distortion measure, i.e. $\delta(\underline{x} - \underline{u}) = \|\underline{x} - \underline{u}\|^2$. In this case we have to find the power-constrained capacity of the channel described above. This capacity, for the general lattice case, can be bounded by, (see Appendix C)

$$\frac{1}{2}\log(1 + 2\pi e G_K) \leq C \leq \frac{1}{2}\log 4\pi e G_K \tag{34}$$

Now, we notice that as $K \to \infty$, the minimum value of G_K, i.e. the optimal second moment at dimension K, approaches $1/2\pi e$ (see e.g. [6]). Thus, as $K \to \infty$, the upper and lower bounds on the capacity approach each other and approach 0.5. This capacity bounds the redundancy of the appropriate lattice quantizer, and as noted above, it is indeed achievable and describes the performance of the quantizer for the region that can be characterized as "low resolution quantization". In the next section we will extend our results and examine the performance of the randomized quantizer in various quantization resolutions.

4 The Redundancy of the Randomized Quantizers at Various Distortion Levels

Deterministic uniform and lattice quantizers have been investigated by Gish and Pierce [3], and by Gersho [4],[11]. Unlike our results for the randomized quantizer, the rate, or the entropy, of the deterministic quantizers has been calculated only in the high resolution, or low distortion case, and for sources with "smooth" probability density function. Nevertheless, under these conditions, it was shown there that the optimal, entropy-constrained, quantizer becomes a uniform (or lattice) quantizer, and, in general, the results there imply that the rate of the deterministic lattice quantizer, of dimension K, satisfies,

$$\frac{1}{n}H(Q_K(\underline{X})) - R(\epsilon) \le \mu_K(\epsilon) \tag{35}$$

where for a square error distortion,

$$\lim_{\epsilon \to 0} \mu_K(\epsilon) = \frac{1}{2}\log 2\pi e G_K \tag{36}$$

and G_K is defined in (6). Note that this high resolution bound is better by 0.5 bit/sample then our universal bound derived in the previous section.

In this section we show that under the "high resolution" assumption we can get bounds similar to (35) for the randomized quantizers as well. Furthermore, we develop an alternative upper bound to the redundancy of the randomized quantizer, which expresses the redundancy variation as a function of the quantizer resolution; the "high resolution" results will be a special case of this bound.

To simplify the exposition we derive the results for the square error distortion criterion, where we get expressions that depends on the second moment of the lattice. A similar derivation for other distortion measures is straight-forward.

A. Alternative Upper Bound for the Quantization Redundancy

Consider again the expressions for the rate of the randomized uniform and lattice quantizers for a source $\underline{X}$. As derived in the previous sections we can write,

$$\frac{1}{n}H(Q_K|\underline{Z}) - R_n(\epsilon) = \frac{1}{n}\left[I(\underline{X};\underline{Y}) - I(\underline{X};\underline{U})\right] \tag{37}$$

where all the quantities in (37) have been defined above. Recall that $I(\underline{X};\underline{Y}) = H(\underline{Y}) - H(\underline{N})$, where for lattice quantizer of dimension K, $\frac{1}{n}H(\underline{N}) = \frac{1}{K}\log V$, and where V is the volume of the lattice's Voronoi cell, i.e. $V = \int_{V_0^-} d\underline{\alpha} = \int_{V_0} d\underline{\alpha}$. Since for a square error distortion measure $\epsilon = \frac{1}{K}\int_{V_0^-}\frac{1}{V}\|\underline{\alpha}\|^2 d\underline{\alpha} = \frac{1}{K}\int_{V_0}\frac{1}{V}\|\underline{\alpha}\|^2 d\underline{\alpha}$, we get from (6) that $V^{2/K} = \frac{\epsilon}{G_K}$. Thus,

$$\frac{1}{n}H(\underline{N}) = \frac{1}{2}\log\frac{\epsilon}{G_K} \tag{38}$$

Now, $I(\underline{X};\underline{U}) = H(\underline{X}) - H(\underline{X}|\underline{U})$. We also know that $\frac{1}{n}E\{\|\underline{X}-\underline{U}\|^2\} \leq \epsilon$. Thus we can write,

$$\begin{aligned}\frac{1}{n}H(\underline{X}|\underline{U}) &= \frac{1}{n}H(\underline{X}-\underline{U} \mid \underline{U}) \leq \frac{1}{n}H(\underline{X}-\underline{U}) \\ &\quad \text{and since } \tfrac{1}{n}E\{\|\underline{X}-\underline{U}\|^2\} \leq \epsilon, \\ &\leq \frac{1}{2}\log 2\pi e\epsilon \end{aligned} \tag{39}$$

where the last inequality was achieved by substituting the p.d.f of a Gaussian vector whose components are i.i.d, each with zero-mean and variance ϵ, which is the maximum entropy p.d.f under the second moment constraint.

Using (37), (38) and (39), we can bound the redundancy of the randomized lattice quantizer,

$$\frac{1}{n}H(Q_K|\underline{Z}) - R_n(\epsilon) \leq \frac{1}{n}\left[H(\underline{X}+\underline{N}) - H(\underline{X})\right] + \frac{1}{2}\log 2\pi e G_K \tag{40}$$

The quantity $\frac{1}{n}\left[H(\underline{X}+\underline{N}) - H(\underline{X})\right] = \frac{1}{n}\left[H(\underline{Y}) - H(\underline{X})\right]$ can be a measure of the quantizer resolution. As will be shown below, it tends to zero for high resolution quantization and smooth source density. We can write this quantity as,

$$\frac{1}{n}\left[H(\underline{Y}) - H(\underline{X})\right] = \frac{1}{2}\log\frac{P(\underline{Y})}{P(\underline{X})} \tag{41}$$

where $P(\underline{Y}) = \frac{1}{2\pi e}2^{2H(\underline{Y})}$ is the entropy power of $\underline{Y}$ and $P(\underline{X})$ is the entropy power of $\underline{X}$.

A summarized statement of this resolution-dependent bound is given in the following theorem:

Theorem 3 *The redundancy of the randomized lattice quantizer, for a square error distortion measure, is upper bounded by*

$$\rho_{K,n}(\epsilon) = \frac{1}{n}H(Q_K|\underline{Z}) - R_n(\epsilon) \leq \frac{1}{2}\log\frac{P(\underline{Y})}{P(\underline{X})} + \frac{1}{2}\log 2\pi e G_K \tag{42}$$

The proof was given in the derivation above. Note that the first term in the bound is a measure of the quantizer resolution, while the second term measures the second moment of the lattice quantizer.

B. High Resolution case

The high resolution assumption can be stated as follows. Given some small ϵ_r, let the quantizer be fine enough and the source p.d.f smooth enough so that for each $\underline{\xi} \in V_0$, where V_0 is the lattice basic cell,

$$f_{\underline{X}}(\underline{\alpha}) \approx f_{\underline{X}}(\underline{\alpha}+\underline{\xi}) \quad \text{or} \quad |f_{\underline{X}}(\underline{\alpha}) - f_{\underline{X}}(\underline{\alpha}+\underline{\xi})| < \epsilon_r \tag{43}$$

Under this condition we get,

$$f_{\underline{Y}}(\underline{\alpha}) \approx f_{\underline{X}}(\underline{\alpha}) \quad \text{or} \quad |f_{\underline{Y}}(\underline{\alpha}) - f_{\underline{X}}(\underline{\alpha})| < \epsilon_r \tag{44}$$

since $f_{\underline{Y}}(\underline{\alpha})$ is the average over all $\underline{\xi} \in V_0$ of $f_{\underline{X}}(\underline{\alpha}+\underline{\xi})$. We also assume that $|\int \log f_{\underline{X}}(\underline{x})|$ and $|\int \log f_{\underline{Y}}(\underline{y})|$ are finite, i.e. bounded by M. Now from (44), it is easy to show, by simple algebraic manipulations that $|H(Y) - H(X)| < \epsilon_r M$, which implies,

$$\frac{1}{n}\left[H(\underline{Y}) - H(\underline{X})\right] \approx 0 \tag{45}$$

and from (40), the redundancy of the lattice quantizer for the square error distortion will be,

$$\lim_{\epsilon_r \to 0}\left[\frac{1}{n}H(Q|\underline{Z}) - R_n(\epsilon)\right] \leq \frac{1}{2}\log 2\pi e G_K = \lim_{\epsilon_r \to 0} \mu_K(\epsilon_r) \tag{46}$$

which is the redundancy for the deterministic quantizer, (36).

C. Examples

The rate of the uniform/lattice randomized quantizer, can now be easily assessed, for practical usage, using the expressions and the bounds derived in this paper.

As a first example, let us consider the rate of the quantizer when it operates on a Gaussian, memoryless, source. The rate can be determined by (**??**), which gives an expression for the exact rate, or from the bound (42). We note, however, that for the Gaussian source, both expressions are identical. In Figure 2, we have plotted this rate, as a function of the distortion, and compared it to the rate-distortion function of the Gaussian source (the straight line, in the log-log scale). Note that this rate is greater by at most .75 bits than the rate-distortion function, at the low resolution case. This is slightly less than the bound 0.754, since the Gaussian distribution is not the one that achieves the capacity.

The rate of the lattice quantizer as a function of the distortion, for encoding this memoryless Gaussian source, is depicted in Figure 3, where the rate is compared to the randomized scalar quantizer, and so we can see the vector advantage. The rate was again computed via (42), where we have assumed that $G_K = \frac{1}{2\pi e}$, i.e., the excess rate over the rate-distortion function was only the resolution measure, $\frac{1}{2}\log\frac{P(\underline{Y}}{P(\underline{X}}$.

Calculating explicitly the rate of the universal quantizer and the rate distortion function can be complicated even for simple sources. However, the resolution measure, that appears in the upper bound, (42), can be calculated relatively easy. In Figure 4 we present this measure, as a function of the quantizer step-size (normalized by the source variance, σ), for the uniform memoryless source (U), the memoryless Laplacian source and the memoryless Gaussian source. Note that since the low resolution case of $SNR = 1$ correspond to $\Delta/\sigma = \sqrt{12}$, the plots in Figure 4 correspond to a medium resolution case.

5 Summary and Further Research

In this paper we have found universal upper bounds for the difference between the rate of the randomized uniform and lattice quantizers and the rate distortion function of the source. This difference, defined as the redundancy of the quantizer, was upper bounded by two procedures. In the first procedure, we achieved a bound which hold for all distortion levels, and can be described as the capacity of the channel of Figure 1; however, this bound is tight only for "low resolution" quantization. The second procedure leads to a bound that depends on the allowable distortion; this bound converges to the results obtained for deterministic uniform/lattice quantizers for "high resolution" quantization.

An interesting problem, for further research, is to determine the difference between the rate of the randomized uniform/lattice quantizers and the rate of the optimal quantizer, $H(Q_{opt})$. As discussed in the paper, an upper bound for this difference is any upper bound we have derived above for the redundancy. However, we expect to find in a further research tighter bounds. Qualitatively, at high-resolution quantization, following claims by Gersho [4], the optimal quantizer becomes a uniform (or lattice) quantizer; thus the difference between their rate approaches zero. On the other hand, in the low-rate case, the optimal quantizer rate may be much closer to the rate distortion function then the uniform/lattice quantizer. For example, at high distortion $R(\epsilon) \approx H(Q_{opt}) \approx 0$, and in this case we can use our bounds for the redundancy.

To summarize, the rate (as a function of distortion) of the optimal quantizer is upper bounded by the rate of the uniform/lattice quantizer and lower bounded by the rate distortion function (both as a function of distortion), and the distance between these two curves has been determined in this paper. We wish, in a further research, to get a more precise estimate of the rate-distortion curve of the optimal quantizer, and thus to appreciate the necessity of using it.

6 Example: Memoryless Zero-Mean Gaussian Source

This simple example may clarify many of the results presented in the paper.

Figure Captions

.

Figure 1: The Auxliary Channel

Figure 2: The Rate of the Uniform Randomized Quantizer for a Gaussian Source as compared to the rate-distortion function

Figure 3: The Rate of the randomized lattice quantizer for a memoryless Gaussian source, as compared to the rate of the uniform randomized quantizer

Figure 4: The resolution measure for Uniform, Laplacian and Gaussian memoryless sources

References

[1] J. Ziv. On universal quantization. *IEEE Trans. Information Theory*, IT-31:344–347, May 1985.

[2] M. Gutman. On universal quantization with various distortion measures. *IEEE Trans. Information Theory*, IT-33, Jan. 1987.

[3] H. Gish and N. J. Pierce. Asymptotically efficient quantization. *IEEE Trans. Information Theory*, IT-14:676–683, Sept. 1968.

[4] A. Gersho. Asymptotically optimal block quantization. *IEEE Trans. Information Theory*, IT-25:373–380, July 1979.

[5] P. A. Chou, T. Lookabaugh, and R. M. Gray. Entropy constrained vector quantization. *IEEE Trans. Acoustics, Speech, and Signal Processing*, ASSP-37:31–42, Jan. 1989.

[6] J. H. Conway and N. J. A. Sloane. Voronoi regions of lattices, second moments of polytops, and quantization. *IEEE Trans. Information Theory*, IT-28:211–226, Mar. 1982.

[7] R. G. Gallager. *Information Theory and Reliable Communications.* Wiley, New York, N.Y., 1968.

[8] L. D. Davisson. Universal lossless coding. *IEEE Trans. Information Theory*, IT-19:211–226, Jan. 1987.

[9] R. E. Kriechevski and V. E. Trofimov. The performance of universal encoding. *IEEE Trans. Information Theory*, IT-27:199–207, March 1981.

[10] J. Ziv and A. Lempel. Compression of individual sequences via variable rate coding. *IEEE Trans. Information Theory*, IT-24:530–536, Sept. 1978.

[11] A. Gersho. On the structure of vector quantizers. *IEEE Trans. Information Theory*, IT-28:157–166, Mar. 1982.