

Lattice Quantization with Side Information

Sergio D. Servetto

Laboratoire de Communications Audiovisuelles
Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland.
URL: <http://lcawww.epfl.ch/~servetto/>

Abstract

We consider the design of lattice vector quantizers for the problem of coding Gaussian sources with uncoded side information available only at the decoder. The design of such quantizers can be reduced to the problem of finding an appropriate sublattice of a given lattice codebook. We study the performance of the resulting quantizers in the limit as the encoding rate becomes high, and we evaluate these asymptotics for three lattices of interest: the hexagonal lattice A_2 , the Gosset lattice E_8 , and the Leech lattice Λ_{24} . We also verify these asymptotics numerically, via computer simulations based on the lattice A_2 . Surprisingly, the lattice E_8 achieves the best performance of all cases considered.

1 Introduction

1.1 Rate Distortion with Side Information

Let $\{(X_n, Y_n)\}_{n=1}^{\infty}$ be a sequence of independent drawings of a pair of dependent random variables X and Y , and let $D(x, \hat{x})$ denote a single-letter distortion measure. The problem of rate distortion with side information at the decoder asks the question of how many bits are required to encode the sequence $\{X_n\}$ under the constraint that $\mathbf{E}D(x, \hat{x}) \leq d$, assuming the side information $\{Y_n\}$ is available to the decoder but not to the encoder [4, Ch. 14.9]. This problem, first considered by Wyner and Ziv in [13], is a special case of the general problem of coding correlated information sources considered by Slepian and Wolf [11], in that one of the sources ($\{Y_n\}$) is available *uncoded* at the decoder. But it also generalizes the setup of [11], in that coding is with respect to a fidelity criterion rather than noiseless.

In [12, 13], Wyner and Ziv derive the rate/distortion function $R^*(d)$ for this problem, for general sources and general (single letter) distortion metrics. Under similar assumptions, a more general function $R(d_1, d_2)$ was considered by Heegard and Berger in [7], for the case when there is uncertainty on whether side information is available at the decoder or not. In this work however we restrict our attention only to Gaussian sources (in which $Y_n = X_n + Z_n$, where Z is also Gaussian and independent of X), and mean squared error (MSE) distortion. This case is of special interest because, under these conditions, it happens that $R^*(d) = R_{X|Y}(d)$, the conditional rate/distortion

function *assuming Y is available at the encoder* [13]. We are intrigued by the fact that there exist coding methods which can perform as well as if they had access to the side information at the encoder, even though they don't. Our main goal in this paper then is to construct a family of quantizers which realizes these promised gains.

The design of quantizers for the problem of rate distortion with side information –also referred to in this work as Wyner-Ziv Encoding– was first considered recently by Zamir, Verdú and Shamai, where they present design criteria for two different cases: Bernoulli sources with Hamming metric, and jointly Gaussian sources with mean squared error metric [10, 15]. The key contribution presented in that work is a constructive mechanism for, given a codebook, using the side information at the decoder to reduce the amount of information that needs to be encoded to identify codewords, while at the same time achieving essentially the distortion of the given codebook. However, the authors do not present any concrete examples on the application of their technique to a particular codebook, the first of which is then worked out by Pradhan and Ramchandran in [8], where they design and thoroughly analyze the performance of trellis codes based on the codebook partitioning ideas of [10, 15].

1.2 Lattice Quantizers for Wyner-Ziv Encoding

High-rate quantization theory provides much of the motivation to consider lattices [6]. Under an assumption of fine quantization, the performance of an n -dimensional quantizer Λ whose Voronoi cells are all congruent to a polytope P is given by

$$d = C(P) \cdot e^{-2(\mathcal{H}(\Lambda, p_X) - h(p_X))}, \quad (1)$$

where p_X is the joint source distribution in n dimensions, \mathcal{H} is the discrete entropy induced on the codebook Λ by quantization of the source p_X , h is the differential entropy, and

$$C(P) = \frac{\frac{1}{n} \int_P \|x - \hat{x}\|^2 dx}{\left(\int_P dx\right)^{1+\frac{2}{n}}}$$

is the normalized second moment of P (using MSE as a distortion measure) [5, 14].

In the problem of rate distortion with side information, for Gaussian sources and MSE distortion, the goal is to attain a distortion value d using $R_{X|Y}(d) < R_X(d)$ nats/sample. In (1) this means that, at fixed bit rate R_0 , we want to design quantizers that achieve distortion

$$d_0 \approx c_n \cdot e^{-2(R_0 - h(p_{X|Y}))}$$

when coding X , where $c_n \leq C(P)$ is the coefficient of quantization in n dimensions [5]. But since we do not have access to Y (we only know $p_{X|Y}$), using classical quantizers we can only attain a distortion value

$$d \approx c_n \cdot e^{-2(R_0 - h(p_X))} > d_0$$

(because $h(X|Y) < h(X)$), or equivalently, we need to use some extra rate $\rho \approx R_X - R_{X|Y}$ such that

$$d_0 \approx c_n \cdot e^{-2(R_0 + \rho - h(p_X))}.$$

What makes this problem interesting is that we are only allowed to use R_0 nats/sample, not $R_0 + \rho$. One way to do that has been proposed by Shamai, Verdú and Zamir in [10, 15], which consists of: (a) taking a codebook with roughly $2^{n(R_0+\rho)}$ codewords and distortion d_0 , (b) partitioning this codebook into 2^{nR_0} sets of size $2^{n\rho}$ each, (c) encoding only enough information to identify each one of the 2^{nR_0} sets, and (d) using the side information Y to discriminate among the $2^{n\rho}$ codewords collapsed into each set. In that construction however, two problems need to be dealt with. First, the codebook partitions need to be designed in a way such that the ability of the decoder to discriminate among codewords within each set is maximized. Second, the fact that partitions which “maximize discriminability” result in a roughly uniform distribution for the symbols to be encoded (we will see later why), thus sacrificing whatever gains may be possible due to entropy coding.

Lattice structures provide an intuitively appealing framework in which to design quantizers for the Wyner-Ziv problem, for a number of reasons:

- Certain interesting lattices have an algebraic structure which is useful in the design of the sought codebook partitions, as well as fast encoding algorithms which make them attractive from an implementation viewpoint [3, Ch. 20].
- The Asymptotic Equipartition Property (AEP) for stationary ergodic sources implies the existence of *typical sets* [4, Ch. 3]. These are sets of sequences of source symbols, which contain almost all the probability mass, and in which all sequences are roughly equally likely. This is particularly relevant for us, since it provides a way to cope with our inability to take advantage of whatever entropy coding gains may be possible for this particular source: since lattices are vector quantizers, we are effectively coding points in the typical set, for which no entropy coding is needed anyway due to their uniform distribution.
- Since we are only changing the encoding procedure to take advantage of the side information at the decoder, but not the shape of the quantization cells, we still need quantizers with cells having good second moment properties. This is a condition met by many known lattices [3, Ch. 2 & 21].
- As pointed out in [15], good codes for Wyner-Ziv encoding must be simultaneously good quantizers as well as good channel codes. And in many cases, the best known lattice quantizer is also the densest known sphere packing, which under certain conditions is equivalent to the channel coding problem [3, Ch. 3].

We should also mention among the reasons to consider lattices our wish to answer a challenge recently posed by Zamir and Shamai in [15]. They present an encoding procedure very closely related to the one we propose here, they argue the existence of good lattices to use with that procedure, they study their distortion performance, but they do not present any examples of concrete constructions: their paper concludes by saying that “*beyond the question of existence, it would be nice to find specific constructions of good nested codes*” (sic). Finding those specific constructions is one of the original contributions in this work.

1.3 Main Contributions and Organization of the Paper

In this paper we construct lattice quantizers for the problem of rate distortion with side information, for jointly Gaussian sources with MSE distortion. Our main contribution is that of presenting specific examples of pairs of nested lattices, the computation of their resulting distortion at high rates, and numerical simulations to verify the computed asymptotics in one simple case. Since our goal is to present a systematic and constructive approach, we have chosen to restrict our attention to lattices with a rich algebraic structure, like those studied in [3]. To design these quantizers we search for an appropriate sublattice of a given lattice, and in this search the algebraic structure of the lattice is helpful.

The rest of this paper is organized as follows. In Section 2 we define and give some intuition on the structure of the proposed quantizers. In Section 3 we study the asymptotic performance of these quantizers at high rates, and we present some experimental results consistent with these asymptotics. Finally, we present conclusions in Section 4.

2 Structure of Wyner-Ziv Lattice Quantizers

2.1 Definitions

A *Wyner-Ziv Lattice Vector Quantizer* (WZ-LVQ) is a triplet $\mathcal{Q} = (\Lambda, \kappa, s)$, where:

- Λ is a lattice.
- $\kappa : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is a linear operator such that $\kappa u \cdot \kappa v = c u \cdot v$, and such that $\Lambda' \triangleq \kappa(\Lambda) \subseteq \Lambda$. Essentially, κ defines a *similar* sublattice of Λ .¹
- $s \in (0, \infty)$ is a scale factor that expands (or shrinks) Λ and Λ' .

Intuitively, the lattice Λ is the fine codebook, the one whose codewords are to be partitioned into equivalence classes. We choose to implement this partition by considering a sublattice $\Lambda' \subseteq \Lambda$, and then considering the resulting quotient group Λ/Λ' . Since the fine lattice is partitioned into $|\Lambda/\Lambda'|$ equivalence classes, the rate of the resulting quantizer is $\frac{1}{n} \ln(|\Lambda/\Lambda'|)$ nats/sample (n is the dimension of the lattice), irrespective of the scale factor s . And s is a constant that multiplies the generator matrices of the lattices considered, which is to be adjusted as a function of the correlation between the source X and the side information Y : small s leads to fine quantization but poor ability to discriminate among codewords within each partition of Λ , large s leads to good discrimination ability but coarse quantization.

The question of the existence of similar sublattices arose recently in connection with another vector quantization problem [9], and also in the study of symmetries of

¹Two lattices Λ_1 and Λ_2 (with generator matrices M_1 and M_2) are said to be similar when there exist a constant $c \neq 0$, a matrix U with integer entries and $|\det(U)| = 1$, and a real matrix B with $BB^\top = I$, such that $M_2 = cUM_1B$ [3]. Intuitively, similar lattices “look the same”, up to a rotation, a reflection, and a change of scale.

quasicrystals [1]. The subject is thoroughly covered in [2], where necessary (and in some cases sufficient) conditions are given for their existence.

2.2 Encoding/Decoding

For a lattice Λ of dimension n , let $Q_\Lambda : \mathbf{R}^n \rightarrow \Lambda$ be its *nearest neighbor* map (i.e., $Q_\Lambda(x) = \arg \min_{\lambda \in \Lambda} \|x - \lambda\|^2$), and let Λ/Λ' denote the corresponding quotient group. Let X denote a block of n source samples, and Y a block of n side information samples. The encoder and decoder are maps $f : \mathbf{R}^n \rightarrow \Lambda/\Lambda'$ and $g : \Lambda/\Lambda' \rightarrow \Lambda$, defined by

$$\tilde{X} = f(X) = Q_\Lambda(X) - Q_{\Lambda'}(X), \quad \hat{X} = g(\tilde{X}) = Q_{\Lambda'+\tilde{X}}(Y),$$

whose operation is illustrated in Fig. 1, with an example based on the lattice A_2 .




Figure 1: To illustrate the mechanics of the proposed quantizers (left: encoding, right: decoding). A sublattice similar to the base lattice is chosen (circled points), matched to how far X and Y are expected to be: in this example, with high probability X and Y are in neighboring Voronoi cells of the fine lattice. Then X is quantized twice (using the fine and the coarse lattice), and the difference between these two is sent to the decoder, as a representative of the set of all codewords collapsed into the same equivalence class. At the decoder, the entire class is recreated (all the points with a thick arrow in the right picture), and among these, the point closest to the side information Y is declared to be the original quantized value for X . Note that there is always a chance that a particular realization of the noise process may take Y too far away from X , in which case a decoding error occurs.

Now it should become clear why earlier we said that the proposed encoding procedure results in uniformly distributed symbols at the output of the encoder. Under a fine quantization assumption, the Voronoi cells of the sublattice will be small, and therefore the source density will be roughly constant within these cells (assuming a continuous source density, true in the Gaussian case considered here). But our partition of the fine lattice into equivalence classes results in a split of each one of these large cells into $|\Lambda/\Lambda'|$ smaller cells of equal area, and therefore equally likely.

3 Asymptotics of Wyner-Ziv Lattice Quantizers

In this section we study the distortion performance of the proposed quantizers, asymptotically as their encoding rate becomes high.

3.1 Derivation of the Distortion Equation

When applied to a WZ quantizer (Λ, Λ', s) , equation (24) of [5] takes the form:

$$\mathcal{D}_s = \left[(1 - p_e(\Lambda, \Lambda', s, \sigma_Z)) C(\Lambda) + p_e(\Lambda, \Lambda', s, \sigma_Z) \sigma_X^2 \right] \cdot e^{-2[\mathcal{H}(s\Lambda, p_X) - h(p_X)]}, \quad (2)$$

where:

- $X_n \sim \mathcal{N}(0, \sigma_X^2)$ is the source sequence, $Z_n \sim \mathcal{N}(0, \sigma_Z^2)$ is a noise sequence (independent of X_n), and $Y_n = X_n + Z_n$ is the side information at the decoder.
- p_e is the probability of a decoding error:

$$p_e(\Lambda, \Lambda', 1, \sigma_Z) = \frac{1}{(\sigma_Z \sqrt{2\pi})^n} \int_{V(\Lambda')} e^{x \cdot x / 2\sigma_Z^2} dx = \frac{\tau}{2} \operatorname{erfc} \left(\frac{\rho e^R}{\sigma_Z \sqrt{2}} \right) \quad (3)$$

where τ is the kissing number of Λ , and ρ is half the norm of a shortest vector in Λ [3, Ch. 3.1].^{2,3} By a simple change of variables we also get, for general s , $p_e(\Lambda, \Lambda', s, \sigma_Z) = p_e \left(\Lambda, \Lambda', 1, \frac{\sigma_Z}{s} \right)$.

- $C(\Lambda)$ is the normalized (and hence scale-independent) second moment of the Voronoi cells of the lattice Λ .
- $h(p_X)$ is the *differential* entropy of the source to encode.
- $\mathcal{H}(s\Lambda, p_X)$ is the *discrete* entropy induced on the points in $s\Lambda$ when this one is used as a vector quantizer for the source X . Under a fine quantization assumption, we have that

$$\mathcal{H}(s\Lambda) \approx h(p_X) - \frac{1}{n} \ln(\operatorname{Vol}(s\Lambda))$$

Equation (2) has a simple interpretation: if no decoding error is made, the MSE is that of the fine codebook; else, the error is bounded by the source variance. And using the above simplifications, it can be rewritten as:

$$\mathcal{D}_s = \sigma_Z^2 e^{-2R} \left\{ \left[1 - \frac{\tau}{2} \operatorname{erfc} \left(\frac{\rho s e^R}{\sigma_Z \sqrt{2}} \right) \right] C(\Lambda) + \left[\frac{\tau}{2} \operatorname{erfc} \left(\frac{\rho s e^R}{\sigma_Z \sqrt{2}} \right) \right] \sigma_X^2 \right\} \left(\frac{s e^R}{\sigma_Z} \right)^2 \det(\Lambda)^{\frac{1}{n}} \quad (4)$$

²The index $|\Lambda/\Lambda'|$ is given by $c^{n/2}$, where c is the norm of the similarity κ [2]. Hence, half the length of a shortest vector in Λ' is $\rho\sqrt{c} = \rho|\Lambda/\Lambda'|^{\frac{1}{n}} = \rho e^R$, and where $R = \frac{1}{n} \ln(|\Lambda/\Lambda'|)$ is the rate of the quantizer.

³The last equality holds only under the assumption that $\sigma_Z \ll \rho\sqrt{|\Lambda/\Lambda'|}$.

Equation (4) is a high-SNR approximation of equation (2), obtained by replacing the integral in (3) by its closed form expression under the assumption of small noise variance. Written in this form, issues that make the problem of rate distortion with side information interesting (and different from the classical rate distortion problem) become clear:

- Whereas in the classical lattice quantization problem specification of the coding rate uniquely determines the volume of the Voronoi cells of the lattice quantizer, in this problem it does not, it only specifies the index $|\Lambda/\Lambda'|$. This means that the scale parameter s remains to be determined.
- In determining the best value of s , two contradicting goals are pursued. On one hand, we want to make s as large as possible, because in this way p_e will be small and the distortion will be mostly due to quantization noise: that is, the construction must be a good *channel* code. On the other hand, we want to make s as small as possible, to achieve low quantization noise: that is, the construction must be a good *source* code.

3.2 Scale Selection

To complete our design procedure, we still need to specify criteria for the selection of the scale factor s , for which a natural choice is the minimization of \mathcal{D}_s as a function of s . At this point, we could proceed simply by solving $\partial\mathcal{D}/\partial s = 0$ and checking the appropriate sign-inversion condition around the solution found. But since we have not been able to do this analytically (only numerically), this solution is not particularly useful in terms of gaining understanding on this problem. To get some intuition then, we first plot \mathcal{D}_s –as defined by eqn. (4)–, for three lattices of interest: A_2 , E_8 and Λ_{24} . The resulting plot is shown in Fig. 2.

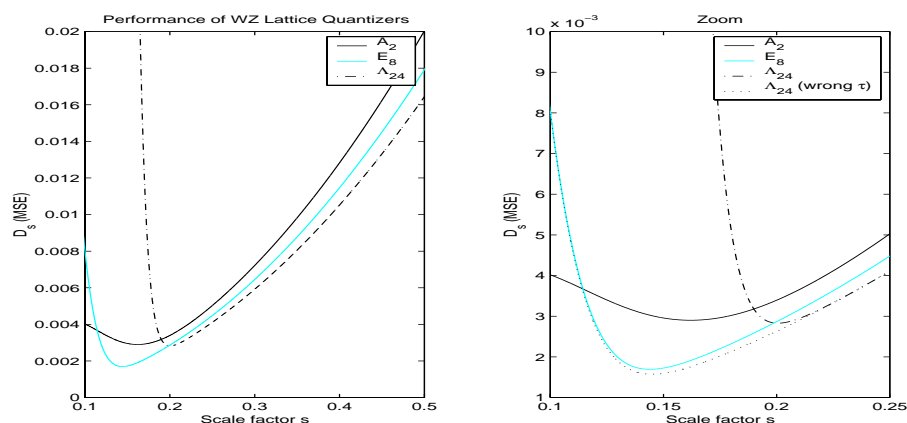


Figure 2: Distortion of the proposed WZ quantizers, for a fixed rate of 1 nat/sample and for a source with $\sigma_X^2 = 1$ and $\sigma_Z^2 = 0.01$, as a function of the scale factor s . The curve denoted “ Λ_{24} (wrong τ)” in the right plot refers to the performance that would be attained by the WZ quantizer based on the Leech lattice, if the kissing number of this lattice were 240 (like in the case of E_8) instead of 196560. All lattices are normalized to have determinant 1.

We observe some interesting properties of the MSE function:

- In all cases, there is a local minimum of the error. This minimum corresponds to a point s^* where, if $s < s^*$, the MSE is dominated by a high probability of error, whereas if $s > s^*$ it is dominated by quantization noise.
- In this problem, $(1 - p_e)C(\Lambda) + p_e\sigma_X^2$ plays a role analogous to that of the coefficient of quantization c_n in the classical problem. A remarkable property of c_n is that it depends only on the geometry of the quantizer, but is independent of the encoding rate and of the shape of the source [5, 14]. However, that property does *not* hold in the presence of side information, since p_e depends essentially on the noise variance σ_Z^2 .
- $\frac{e^{2R}}{\sigma_Z^2} \mathcal{D}_s(R)$ depends on s , R and σ_Z only through $s e^R / \sigma_Z$. Therefore, knowledge of the location and value of the local minimum s^* for one fixed value of R uniquely determines its location and value at all rates and noise levels.⁴

Based on the above considerations, we define the distortion of a WZ lattice quantizer $\mathcal{D}(R) = \min_s \mathcal{D}_s(R)$, and the optimal s^* as that value of the scale factor s which attains the minimum in the definition of \mathcal{D} . $\mathcal{D}(R)$ is plotted in Fig. 3, where it is compared to other applicable performance bounds.

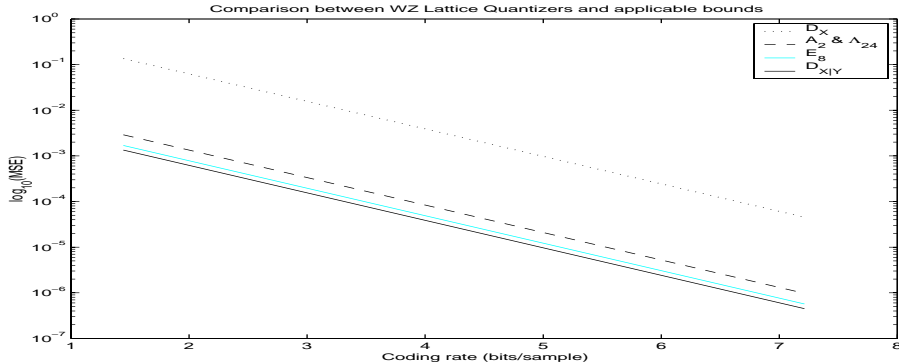


Figure 3: Predicted performance of the proposed lattice quantizers: (a) D_X : Shannon’s distortion/rate function for X ; (b) evaluation of \mathcal{D} for the lattices A_2 and A_{24} (negligible difference among these two); (c) evaluation of \mathcal{D} for E_8 ; (d) $D_{X|Y}$: Wyner-Ziv’s distortion/rate function. This plot is for $\sigma_X^2 = 1$, and $\sigma_Z^2 = 0.01$. All lattices are normalized to have determinant 1 (which in the case of A_2 involves scaling it so that its shortest vectors have norm $2/\sqrt{3}$).

A surprising result is that the quantizer based on the lattice E_8 comes much closer to the Wyner-Ziv bound than the 24-dimensional Leech lattice does, whose performance is a negligible improvement over that of the two-dimensional hexagonal lattice. This result can be explained however in terms of the kissing numbers of these lattices. The equivalence between the channel coding problem and the sphere

⁴This is particularly helpful in the numerical minimization of \mathcal{D}_s , since it means the minimization needs to be done only once, for a fixed rate and noise variance.

packing problem for lattices depends heavily on a high SNR assumption: when the variance of the noise is not negligible compared to the minimum separation between lattice codewords, the probability of error is dominated by the kissing number of the lattice. But in our minimization of \mathcal{D}_s , we essentially search for the smallest scale of the shortest lattice vectors which does not result in too high a probability of error, i.e., we reduce the lattice SNR as much as we can. Under these circumstances the Leech lattice, having many more neighbors at close distance than the Gosset lattice does, necessarily performs worse. This is further corroborated by the fact that, if the kissing number of Λ_{24} were identical to that of E_8 , then it would indeed attain the best performance, as it follows from the fact that the curve denoted “ Λ_{24} (wrong τ)” in Fig. 2 has the lowest local minimum; at higher SNRs though, the Leech lattice attains the least MSE, as was to be expected.

3.3 Numerical Simulations based on the Hexagonal Lattice

To conclude this section, we present preliminary results obtained in a computer implementation of the proposed WZ quantizer, based on the hexagonal lattice. Although we are currently working on implementing these quantizers using other lattices, we feel it is important for us to back our asymptotic analysis with at least one case of experimental validation, and these results –shown in Fig. 4– are relatively simple to obtain using the hexagonal lattice.




Figure 4: Computation of the distortion performance by Monte Carlo methods: 10^8 samples are quantized and dequantized, at bit rates in the range of about 1-6 bits/sample. Observe how the experimental curve approaches the theoretical curve as the bit rate increases.

C code implementing the WZ hexagonal quantizer, figures for a few large index sublattices of the hexagonal lattice, and the latest results are available from our webpage, at <http://lca.vwww.epfl.ch/~servetto/>.

4 Conclusions

In this paper we presented an overview of our work on the design of lattice quantizers for the problem of rate distortion with side information. We showed how the nested codes studied by Zamir and Shamai in [15] can be implemented as a standard lattice

and a similar sublattice. We studied the asymptotic behavior of the distortion that results from applying our construction to three lattices, and we presented numerical simulations whose results are consistent with the predicted distortion at high rates.

Perhaps the most important conclusion that follows from our results is related to those properties which make certain lattices better than others for this problem. Specifically, we find that dense lattice packings with high kissing numbers (such as the Leech lattice) are not well suited to this problem. Instead, good quantizers with small kissing numbers are to be preferred.

An interesting question that still requires further work is that of finding a *concise* way of presenting the performance of our quantizers at high rates. For classical quantizers, one such representation is given by $e^{2R}\mathcal{D} = C(\Lambda)e^{2h}$, since premultiplication by e^{2R} makes the distortion a rate-independent quantity. But this trick does not work for WZ quantizers, since we saw in eqn. (4) that the optimal scale factor also depends on the encoding rate. We hope to be able to find one such short description by taking advantage of the fact that $\mathcal{D}_s(R)$ depends on s , R and σ_Z only through se^R/σ_Z (i.e., rate changes result only in a change of scale for the error function).

Acknowledgements. I would like to thank V. A. Vaishampayan and N. J. A. Sloane, for a number of (very educational) conversations on lattices and vector quantization; to S. Pradhan and K. Ramchandran, for an accessible explanation of their work which served as my introduction to this problem [8]; and to an anonymous reviewer, who referred me to the related information-theoretic work of Heegard and Berger [7].

References

- [1] M. Baake and R. V. Moody. Similarity Submodules and Semigroups. Preprint. Available from <http://solid13.tphys.physik.uni-tuebingen.de/baake/preprints.html>.
- [2] J. H. Conway, E. M. Rains, and N. J. A. Sloane. On the Existence of Similar Sublattices. *Canad. J. Math.*, 1999. To appear. Available from <http://www.research.att.com/~njas/>.
- [3] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*. Springer Verlag, 3rd edition, 1998.
- [4] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley and Sons, Inc., 1991.
- [5] A. Gersho. Asymptotically Optimal Block Quantization. *IEEE Trans. Inform. Theory*, IT-25(4):373–380, 1979.
- [6] R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Trans. Inform. Theory*, 44(6):2325–2383, 1998.
- [7] C. Heegard and T. Berger. Rate Distortion when Side Information May Be Absent. *IEEE Trans. Inform. Theory*, IT-31(6):727–734, 1985.
- [8] S. Pradhan and K. Ramchandran. Distributed Source Coding Using Syndromes (DISCUS): Design and Construction. In *Proc. IEEE Data Compression Conf. (DCC)*, Snowbird, UT, 1999.
- [9] S. D. Servetto, V. A. Vaishampayan, and N. J. A. Sloane. Multiple Description Lattice Vector Quantization. In *Proc. IEEE Data Compression Conf. (DCC)*, Snowbird, UT, 1999.
- [10] S. Shamai, S. Verdú, and R. Zamir. Systematic Lossy Source/Channel Coding. *IEEE Trans. Inform. Theory*, 44(2):564–579, 1998.
- [11] D. Slepian and J. K. Wolf. Noiseless Coding of Correlated Information Sources. *IEEE Trans. Inform. Theory*, IT-19(4):471–480, 1973.
- [12] A. D. Wyner. The Rate-Distortion Function for Source Coding with Side Information at the Decoder-II: General Sources. *Inform. Contr.*, 38:60–80, 1978.
- [13] A. D. Wyner and J. Ziv. The Rate-Distortion Function for Source Coding with Side Information at the Decoder. *IEEE Trans. Inform. Theory*, IT-22(1):1–10, 1976.
- [14] P. Zador. Asymptotic Quantization Error of Continuous Signals and the Quantization Dimension. *IEEE Trans. Inform. Theory*, IT-28:139–149, 1982.
- [15] R. Zamir and S. Shamai. Nested Linear/Lattice Codes for Wyner-Ziv Encoding. In *Proc. IEEE Inform. Theory Workshop*, Killarney, Ireland, 1998.