Figure Captions

Figure 1: Quantizer distortion as a function of source standard deviation for the Leech, E8, and Z lattices for a generalized Gaussian source with a shape parameter of $\nu = 0.5$. For high σ the distortion converges to the lattice G number, but at lower σ the densest lattices do not give the lowest distortion.

Figure 2. Normalized total (granular + overload) distortion curves for the Z^{24} and Leech lattices at .59 bits/sample for a source with a shape parmeter of $\nu = 0.5$. The minimum of the curve identifies the standard deviation to which a source should be scaled prior to quantization. The distortion that will result after dequantization is the product of the normalized distortion D_n and the source variance.

Figure 3. Distortion of the Z^{24} lattice relative to the Leech lattice (Fig. 3a) and of the Z^8 lattice relative to the E8 lattice (Fig. 3b) as a function of coding rate for a generalized Gaussian source with a shape parameter of 0.5.

References

[1] A. Gersho and R.M. Gray, *Vector Quantization and Signal Compression,* Boston: Kluwer, 1992.

[2] J.H. Conway and N.J.A. Sloane, "Fast quantizing and decoding algorithms for lattice quantizers and codes," *IEEE Trans. Inform. Theory,* vol. IT-28, pp. 227-232, 1982.

[3] J.H. Conway and N.J.A. Sloane, "On the structure of vector quantizers," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 157-166, 1982.

[4] M. Vedat Eyuboglu and G. David Forney, Jr., "Lattice and Trellis Quantization with Lattice- and Trellis-Bounded Codebooks – High-Rate Theory for Memoryless Sources," *IEEE Transactions on Information Theory,"* vol. 39, pp. 46-59, 1993

[5] J.H. Conway and N.J.A. Sloane, *Sphere Packings, Lattices and Groups*, New York: Springer-Verlag, 1988.

[6] T.R. Fischer, "A Pyramid Vector Quantizer," *IEEE Trans. Inform. Theory,* vol. IT-32, pp. 568-583, 1986.

[7] J.H. Conway and N.J.A. Sloane, "Soft decoding techniques for codes and lattices, including the Golay code and the Leech lattice," *IEEE Trans. Inform. Theory*, vol. IT-32, pp. 41-50, 1986.

[8] J. Leech, "Notes on sphere packing," *Can. J. Math.*, vol. 19, pp. 251-267, 1967.

[9] T.D. Lookabaugh and R.M. Gray, "High-resolution quantization theory and the vector quantizer advantage," *IEEE Transactions on Information Theory,* vol.35, pp. 1020-33, 1989.

at a coding rate of 0.25 bpp, the Z lattice resulted in a reconstructed image with 30.97 dB PSNR, while the Leech lattice gave 30.58 dB. The majority of the distortion improvement using the Z lattice occured in the outermost subbands where the shape parameters are lower.

III. Conclusions

Despite the lower G number offered by the E_8, Leech, and other high-density lattices, the Z lattice offers superior performance for low-bit-rate quantization of generalized Gaussian sources when the shape parameter is low. This occurs because for low shape parameters the Voronoi regions of the Z lattice are better-suited to the anisotropicity of the generalized Gaussian PDF, which places a stronger emphasis on vectors occuring along coordinate axes than on those occuring in off-axis directions. This effect is most significant at lower bit rates when most of the quantization occurs to lattice points in the immediate vicinity of the origin. These results show that for low-bit-rate applications such as wireless image transmission it is possible to choose a low complexity quantizer using the Z lattice while simulataneously maximizing reconstructed image quality.

sigma increases from 0. The transition for the Leech lattice occurs more rapidly than the transition for the Z and E_8 lattices, with the result that for values of $\sigma < 1.5$ the Z lattice offers superior distortion performance to the denser E_8 and Leech lattices.

An explanation for the distortion properties occuring at low source standard deviations is furnished by considering the relation between low-variance generalized Gaussian PDF and the lattice characteristics in the neighborhood of the origin. When σ is low, a high fraction of the vectors will lie within the central Voronoi region and will be quantized to zero; the majority of those vectors lying outside the central Voronoi region will lie in Voronoi regions that share a boundary with the central Voronoi region. The nature of the PDF when ν and σ are low is such that vectors lying near the coordinate axes have higher probability of occurence than those lying at off-axis locations. Therefore, the distance between lattice points along the axes becomes more important than the distance between lattice points averaged over all directions. In fact, one can attribute the excellent performance of the Z lattice at low variances to the high degree of anisotropicity of the lattice. The relatively high variation in the distance from a cell center to a cell boundary as a function of direction in the Z lattice comes at the expense of packing efficiency, but it permits a correctly oriented Z lattice to give excellent performance for generalized Gaussian sources.

In order to assess the importance of these effects it is necessary to consider not only the granular distortion, but in addition the overload distortion that arises when a bounded lattice is used. Figure 2 shows the normalized total distortion D_n that would be experienced under quantization of a $\nu = .5$ source with standard deviation σ_s. The lattices used in this comparison are the Z^{24} and Leech lattices at .59 bits/sample; we have also performed similar calculations for other bit rates and for comparisons of the Z^8 and E_8. The minimum of this curve identifies optimum relative scaling that should be performed between a source and a lattice of a given type and size (e.g. number of points). This figure is derived by first calculating the total (=granular+overload) distortion, and then normalizing it by the source variance. The horizontal axis has also been normalized and represents the *scaled* source standard deviation σ; in other words it represents a source with standard deviation σ_s that has been multiplied by the constant σ/σ_s. The vertical axis gives the normalized distortion that results, with the absolute distortion then given by the product $D_n\sigma_s^2$. For the limiting trivial case where the source is multiplied by zero prior to quantization, the distortion will of course be equal to the variance of the source (i.e. $D_n = 1$) regardless of the bit rate or lattice used. As the scaled standard deviation approaches infinity, for any finite bit rate all of the source vectors will be truncated and result in a normalized distortion that again approaches 1.

Rate distortion characteristics can be obtained by creating a series of curves similar to Figure 2 for different lattice sizes and identifying the minimum distortion for each. Figure 3a shows the distortion of the Z^{24} lattice relative to the Leech lattice without entropy coding for bit rates between 0.5 and 4.5 bpp. We have also calculated the coding rates after entropy coding and found that the relative lattice performance is not significantly altered. As the figure shows, for bit rates of under 2.2 bpp, the Z lattice will outperform the Leech lattice, with the amount of the advantage increasing as the coding rate is decreased. At high bit rates, the shape advantage of the Z lattice disappears, and the space-filling advantage of the Leech lattice enables it to give better quantization than the Z lattice. Figure 3b shows the relative performance of the E_8 and Z^8 lattices, and again shows the superiority of the Z lattice for low bit rates.

In addition to the simulation studies that generated the data used to produce Figures 3a and 3b, we also performed LVQ on image subbands obtained using the wavelet transform. The experiment results were consistent with the results presented above; for example, for the 512 by 512 Lena image

I. Introduction

Lattice vector quantization (LVQ) [1, 2, 3, 4, 5, 6] allows efficient coding without the need to generate and store codebooks. A common lattice performance metric for coding applications is the G number, which is a measure of the mean square reconstruction error, normalized by the number of dimensions, that would be experienced in quantizing a uniformly distributed source. Both from a conceptual and a computational standpoint, the simplest lattice is the integer or Z-lattice, which has a G number of 1/12. Other lattices that have been commonly used in coding applications include the E_8 lattice [7] with a G number of .072, the 16-dimensional Barnes-Wall lattice (G = .068), and the Leech lattice [8] (G = .066).

While uniform, Gaussian and Laplacian distributions have been extensively considered in the literature on LVQ, less attention has been given to comparing the performance of different lattices on the important but less mathematically tractable generalized Gaussian distributions that are typically found in subband coded images. This becomes especially important in low-bit-rate applications where the source PDF has a significant effect on the resulting coding performance. As we show here, the relative performance of Leech, E_8, and Z lattices at low bit rates for image subbands and other generalized Gaussian sources is quite different from what occurs at higher bit rates for uniform sources, and is strongly dependent on the shape parameter of the generalized Gaussian source and on the scaling applied to the lattice during the quantization. These results are consistent with the observations of other researchers [9] that the source PDF can be of greater importance in determining the performance of a vector quantizer than the density of the lattice. While the idea that the "space-filling" advantage of VQ can be less important than the "shape" advantage is known, the exploration of this concept for generalized Gaussian sources at low bit rates for the Leech , E_8, and Z lattices is new. Of particular interest are the numerical results we present concerning the relative rate distortion characteristics of these lattices, and the resulting conclusions regarding the optimality of the Z lattice.

II. Distortion in Lattice Coding of Generalized Gaussian Sources

Most subband image data can described by a generalized Gaussian PDF, which has the form $p_X(x) = C_1 e^{-C_2|x|^\nu}$ where C_1, C_2 are constants that ensure unit area and parametrize the standard deviation σ. The variable ν is the shape parameter and typically takes on a value in the range from 0.3 to 0.8 for image subbands. We first condsider the effects of granular distortion using unbounded lattices. To allow a meaningful comparison between lattices, an initial scaling must be performed such that the the volume of the Voronoi cells is equivalent. This means for example, that to compare the Leech and Z lattices in 24 dimensions, a relative scaling must be applied so that a truncation boundary located at a large distance from the origin would enclose the same number of points for both lattices. Using these lattices, one can then perform quantization on a set of source vectors and observe the differences in the distortion that results. For a given source distribution and a given unbounded lattice, the only parameter that can be varied is the relative scaling of the lattice and source.

Figure 1 shows the distortion as a function of source standard deviation (assuming no scaling of the lattice) for a generalized Gaussian source with a shape parameter of $\nu = .5$ for the Z, E_8, and Leech lattices. The data for this figure were obtained by simulation. In the limiting cases of low and high σ the distortion converges to zero and the G number respectively. The most interesting region of the graph lies in the transition region where, the curves for the different lattices diverge as

Abstract

In lattice vector quantization, the distortion associated with a given lattice is often expressed in terms of the G number, which is a measure of the mean square error per dimension generated by quantization of a uniform source. Subband image coefficients, however, are best modeled by a generalized Gaussian distribution, leading to distortion characteristics that are quite different from those encountered for uniform, Laplacian, or Gaussian sources. We have calculated the distortion associated with Z, E_8, and Leech lattice quantization for coding of generalized Gaussian sources, and show that for low bit rates the Z lattice offers both the best performance and the lowest implementational complexity.

.

A Comparison of the Z, E_8, and Leech Lattices for Quantization of Low-shape-parameter Generalized Gaussian Sources

Zheng Gao, Ben Belzer, and John Villasenor

Electrical Engineering Department
University of California, Los Angeles
405 Hilgard Avenue
Los Angeles, CA 90024-1594

IEEE Signal Processing Letters

Vol. 2, pp. 197-199, October, 1995

Please address correspondence to:
John D. Villasenor
Electrical Engineering Department
University of California, Los Angeles
405 Hilgard Avenue
Los Angeles, CA 90024-1594
Tel: (310) 825-0228, FAX: (310) 206-8495
e-mail: villa@ee.ucla.edu