A LATTICE VECTOR QUANTIZER FOR GENERALIZED GAUSSIAN SOURCES

Feng Chen, Zheng Gao, and John Villasenor

Electrical Engineering Department,
University of California, Los Angeles CA 90024
fchen@icsl.ucla.edu, villa@icsl.ucla.edu

ABSTRACT

A fixed-rate lattice vector quantizer for generalized Gaussian (GG) sources is presented. By using the contour of counstant probability of a generalized Gaussian source to bound the integer lattice, this vector quantizer in the limit of high dimension achieves optimal boundary gains for GG sources. Low-complexity techniques for rate control, optimal codevector search and enumeration are developed. This coder acheives 8.60 dB for a generalized Gaussian source with shape parameter of 0.5 at 1 bit/sample (compared with a rate distortion bound of 9.23 dB) and 32.62 dB for the 512X512 "lena" image at 0.25 bpp without the use of entropy coding. By using two small lookup tables, the nearest codeword search operation can be performed multiply-free and the overall quantization process costs one multiply and a few adds per sample.

I. INTRODUCTION

Lattice vector quantization (LVQ) offers the possibility of high coding efficiency without the need to generate and store a codebook [1]. Considerable attention has been given in the recent literature to pyramidal vector quantization (PVQ) in which quantization is performed to an integer lattice lying on a pyramidal shell [2]. Such an approach is ideally suited for Laplacian-distributed vectors of high dimension, which tend to lie on a pyramidal shell of an appropriate radius. In the present paper we are concerned with LVQ of subband image data which, in contrast with the Laplacian model used in PVQ, are best described by a generalized Gaussian distribution with a shape parameter ν in the range $0 < \nu < 1.0$. The surface of constant probability for these data is concave and can be described by $\sum_{i=0}^{L-1} |x_i|^\nu = K^\nu$, where x_i is the ith component of the vector, L is the dimension, and K is the radius, or l_ν-norm. We have developed a fixed-rate lattice vector quantizer with a truncation boundary that matches the surface of constant probability of the generalized Gaussian source and shown that this leads to significant coding gain when compared with a LVQ with a pyramidal boundary.

II. RATE CONTROL, CODEWORD SEARCH AND ENUMERATION

Figure 1 shows the overview of our fixed-rate lattice vector quantizer. After wavelet transformation, subband data are group into vectors of dimension L. A vector is first multiplied by a source scaling factor to achieve the relative source/lattice scaling minimizing the overall distortion. The l_ν-norm of the vector is then computed to determine whether it is inside or outside the truncation boundary. If inside the boundary, only a simple rounding operation is needed to find the nearest codevector, otherwise an optimal truncation algorithm is applied to find the nearest codevector on or within the boundary. The codevector is then assigned a unique binary index which is transmitted over a communication channel.

A crucial step in the design of a fixed-rate lattice vector quantizer is to determine a way to control the bit rate R used for the quantization. Since the bit rate and the number of lattice points N in an L-dimensional vector space are related by the equation $R = (\log_2 N)/L$, an equivalent task is to find out a way to compute the number of integer lattice points bounded by the truncation boundary. The exact number of points $N(L,K)$ enclosed by a L-dimensional truncation boundary with L_ν-norm K can be computed using the following recursion formula:

$$N(L,K) = N(L-1,K) + 2\sum_{i=1}^{\text{int}(K)} N(L-1,(K^\nu - i^\nu)^{1/\nu}) \quad (1)$$

This formula is a generalization of the result presented by Fischer in [2] for the calculation of the number of lattice points lying on a pyramidal shell. Equation (1) can be best understood by considering the L-dimensional lattice as being composed of an odd number of $L-1$ dimensional lattices. The largest of these $L-1$ dimensional lattices is obtained by setting the first coordinate x_0 to be zero. The number of points inside this lattice can be represented by the first term on the right side of the above equation. A pair of smaller $L-1$ dimensional lattices with $x_0 = \pm 1$ will have L_ν-norm $(K^\nu - 1)^{1/\nu}$ and each contain $N(L-1, (K^\nu - 1)^{1/\nu})$ lattice points. One continues to increment x_0 until it reaches the last lattice point within the boundary, located at $x_0 = \pm\text{int}(K)$, and the number of lattice points in each case is $N(L-1, (K^\nu - x_0^\nu)^{1/\nu})$. Equation (1) gives the exact solution by combining exact solutions from lower dimensional lattices with arbitrary floating number radii, and becomes computational prohibitive when K becomes large. As a solution, one can restrict the value of each $|x_i|^\nu$ to be an integer multiple of some quantity δ, defined as the "resolution" of the boundary. A similar approach was used in [3]. This approximation allows the recursion of equation (1) to be calculated by tabulating all the values of $N(L, K)$ for successively higher dimensions and larger radii. Smaller values of δ will more accurately approximate the exact boundary, but will require larger memory for the storage of the table. In practice, we found that a resolution of 0.1 is sufficient for our lattice vector quantizer.

The complexity to find the nearest codeword depends on whether a vector falls within or outside the lattice. For all vectors, an initial scaling step using a predetermined factor must be performed at the cost of 1 multiply/sample to achieve the relative source/lattice scaling that minimizes the overall (granular + overload) distortion. Determination of the l_ν norm costs one addition and one table lookup per sample, (the νth powers of integers are stored in a small lookup table) and then a single compare is used to assess whether the vector is within the lattice bound. For those vectors within the lattice, each component must be rounded to the nearest integer. Overload vectors lying outside the truncation boundary must be truncated to a code vector within the boundary. Defining "optimal" truncation to mean quantization to the lattice codeword that lies closest in the Euclidean sense to the overload vector, one can optimally truncate by first quantizing to the nearest Z lattice point on an unbounded integer lattice. Then the increase in distortion ΔD and reduction in l_ν norm Δl_ν that would result from decrementing each of the quantized components is determined. The component with the smallest ratio of $\Delta D/\Delta l_\nu$ is decremented. For practical values of δ, $1/\Delta l_\nu$ is restricted to a very small set of possible values (6 values for $\delta = .1$), and the division by Δl^ν can be implemented using an average of 1 add and 2 shifts. In subsequent iterations $\Delta D/\Delta l_\nu$ needs to be recomputed only for the component that was decremented in the previous iteration, so the marginal cost of multiple iterations is very low. The first iteration requires 5 additions per sample; subsequent steps, if needed (after optimal scaling most of the overload vectors will lie very close to the truncation boundary), cost a factor of $1/N$ less. This steepest descent process is iterated until the quantized vector moves within the lattice bound.

The average quantization complexity depends on the frequency with which overload vectors occur, which in turn is a function of bit rate and dimension. Table 1 gives the typical complexity on a per sample basis for overload and in-lattice vectors, and also gives measured complexity results for a range of bit rates and dimension. As the examples in the table illustrate, the percentage of overload vectors decreases with bit rate, leading to a decrease in the complexity. If the bit rate is held constant and the dimension is increased, the complexity increases, although at a very slow rate. For example, quadrupuling the dimension at 1 bit/sample leads to only a 30% increase in the number of multiplies/sample. The low complexity of lattice VQ stands in strong contrast to unconstrained vector quantizers in which the complexity to find the nearest codeword can increase in proportion to (or even exponentially with) bit rate and dimension. Table 1 also gives the probability of overload and the ratio $\text{MSE}_{gran}/\text{MSE}_{ovld}$ of the granular to overload distortion for each of the examples listed. Because the bit rates used in the example are low to moderate, the probability of overload is quite high, reaching 90% for 1 bit/sample and 64 dimensions. The $\text{MSE}_{gran}/\text{MSE}_{ovld}$ numbers also follow the expected increase with bit rate as described in [7]. This underscores the importance of an effective truncation scheme for the low bit rate case where most of the distortion is due to overload vectors.

Enumeration of the lattice points within the boundary is accomplished by ordering vectors according to the absolute value of the components, with lower-indexed components taking precedence over higher-indexed components. Vectors with pair-wise identical absolute values for all components are distinguished using a binary enumeration scheme. Since the number of code vectors $N(l, k)$ is known and tabulated for all $l \le L$ and $k \le K$, where l is an integer and k is an floating number of resolution δ, the method in [2] and [3] can be easily extended to perform vector encoding for this LVQ with optimal truncation boundary for generalized Gaussian sources.

III. EXPERIMENTAL RESULTS AND DISCUSSION

The simulation results for a generalized Gaussian source with unit variance and shape parameter of $\nu = .5$ are presented in Table 2 in dB, where the dB value is $-10\log_{10}(MSE)$ In addition to indicating the simulation results, the table indicates the Shannon lower bound [2], the rate distortion bound calculated using the Blahut algorithm [5], and the asymptotic distortion [4]. Detailed derivations of the performance of this LVQ can be found in [6]. The simulation results are reported for dimension 8, 32, 64, 128, and 256. Each value in Table 2 represents the average of simulations using ten different sets of source vectors. The standard deviation for the values is approximately 0.05 dB. The simulation results given in Table 2 are also presented in graphical form in Figure 3. The figure gives the distance in dB from the Shannon lower bound, and conveys the relationship between the quantities of interest as a function of bit rate. Table 2 and Figure 2 contain several noteworthy features. As expected, the gap between the achieved distortion and the rate distortion bound narrows as the dimension increases. At low bit rates and high dimensions, the quantizer exceeds the asymptotic bound and performs (for 1 bit/sample, 128 dimensions) within 1.1 dB of the Shannon lower bound. As the bit rate approaches zero, the rate-distortion bound diverges from the Shannon lower bound because the distortion of the Shannon lower bound does not go to σ^2 for $R \rightarrow 0$. It is also evident from the positive slope of the simulation curves at higher bit rates that the optimally-bounded quantizer fails to realize the full 6.02 dB/bit gain in performance that is possible in theory. This is due to the finite dimension, which puts full benefit of the boundary gain beyond reach, and to the clustering of vectors within a shell. The optimal relative source/lattice scaling will place this shell near the truncation boundary of the lattice. As the bit rate is increased, the bounding-surface to volume ratio of the lattice decreases. Since the volume of the shell containing the vectors will be roughly proportional to the surface area of the truncation boundary, the number of codewords that are "wasted" gets proportionally higher with increasing bit rate. To verify the performance of our coder on actual images we also encoded the the 512x512 lena image and achieved (without entropy coding) a PSNR of 32.62 dB at 0.25 bits/pixel.

IV. CONCLUSION AND FUTURE WORK

We have shown the design of a low-complexity fixed-rate lattice vector quantizer for the quantization of generalized Gaussian sources. This LVQ matches the statistics of generalized Gaussian sources and is well suited to coding of image subband data, which are usually best modelled by generalized Gaussian sources with a shape parameter in the range $0 < \nu < 1.0$. The computational challenges associated with rate control, codeword search and enumeration for a polytopal bounded LVQ have been addressed and fast algorithms for the solutions have been proposed.

There are a number of possible directions for future work in this area. First, the concepts presented here could be extended to include product code quantization. This would certainly give improved results at high coding rates, but would be highly complex because the intersection of the Z lattice and the truncation boundary will not generally define a sublattice. It would also be of interest to extend this approach to other highly peaked sources. One possible approach would be to identify the shape parameter of the generalized Gaussian that most closely resembles the true (non-generalized Gaussian) source, and to analyze the loss in boundary gain due to the mismatch. Alternatively, the true contour of constant probability could be identified and used in techniques analogous to those presented here.

VI. ACKNOWLEDGEMENTS

This work was supported by Texas Instruments, Hughes Aircraft Company, and the Whitaker Foundation. We thank Erich Vogel for providing programming assistance.

1. REFERENCES

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

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

[3] R. Laroia and N. Farvardin, "A structured fixed-rate vector quantizer derived from a variable length scalar quantizer – Part I: Memoryless sources," *IEEE Trans. Inform. Theory,* vol. 35, pp. 851-867, 1991.

[4] H. Gish and J.N. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory,* vol. IT-14, pp. 676-683, Sept. 1968.

[5] R.E. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inform. Theory,* vol. IT-18, pp. 460-473, July 1972.

[6] F. Chen, Z. Gao and J.D. Villasenor, "Lattice vector quantization of generalized Gaussian sources," submitted to *IEEE Trans. of Inform. Theory.*

[7] M. Vedat Eyuboğlu and G. David Forney, Jr., "Lattice and trellis quantization with lattice- and trellis-bounded codebooks – high-rate theory for memoryless sources," *IEEE Trans. Inform. Theory,"* vol. 39, pp. 46-59, Jan. 1993.

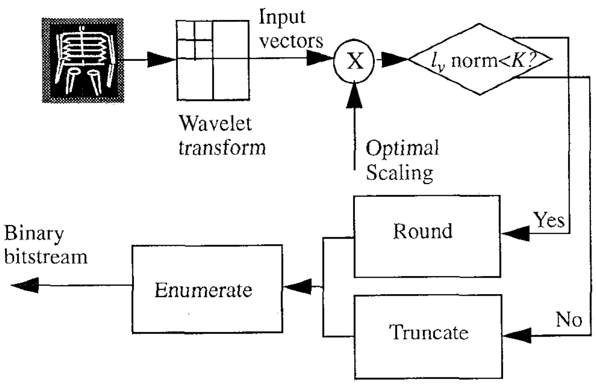

Figure 1. A fixed rate lattice vector quantizer

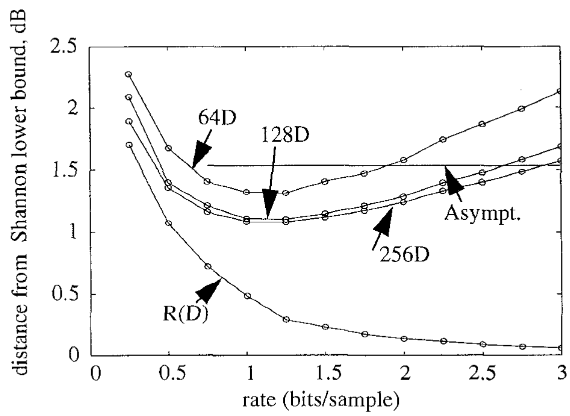

Figure 2: Distortion-rate performance relative to the Shannon lower bound for v = .5

Table 1: Quantization Complexity							
Number of operations per sample: Typical							
			mult	add	round	lookup	compare
Vectors within lattice			1	1	1	1	1/dim
Overload vectors			1	6	1	3	~log(dim)
Average total operations per sample: Examples for source with v = .5							
rate and dimension	prob. of overload	$\frac{MSE_{gran}}{MSE_{ovld}}$	mult	add	round	lookup	compare
1 bpp, dim=16	.63	.19	1	4.1	1	2.3	2.5
1 bpp, dim=64	.90	.07	1	5.9	1	2.9	5.9
3 bpp, dim=16	.26	1.0	1	2.5	1	1.6	1.3
3 bpp, dim=64	.56	.50	1	4.1	1	2.2	3.7

Table 2: Quantization SNR for a unit variance generalized Gaussian source, v = .5								
rate	Theoretical results (dB)			Simulation results (dB)				
	SLB	D(R)	Asympt.	dim = 8	dim = 32	dim = 64	dim = 128	dim = 256
1	9.72	9.23	8.19	6.17	8.05	8.40	8.60	8.64
2	15.74	15.61	14.21	11.99	13.76	14.17	14.39	14.50
3	21.76	21.71	20.23	16.71	19.04	19.63	19.98	20.19