Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing (Toronto, ON), pp. 657-660, May 1991

METHODS OF IMPROVING VECTOR-SCALAR QUANTIZATION OF LPC COEFFICIENTS

John Grass [1] and Peter Kabal [1,2]

[1] Electrical Engineering
McGill University
Montreal, Quebec, H3A 2A7

[2] INRS-Télécommunications
Université du Québec
Verdun, Quebec, H3E 1H6

ABSTRACT

Methods of improving vector-scalar quantization of Linear Predictive Coding (LPC) coefficients with 20 to 30 bits per 20 ms frame are studied in this paper. The approach in this work is to couple the vector and scalar quantization stages. The second innovation is the incorporation of a small adaptive codebook to the larger fixed codebook. Frame-to-frame correlation of the LPC coefficients is exploited at no extra cost in bits. The results of this paper show that the performance of the vector-scalar quantization with the use of the two new techniques introduced is better than that of scalar coding techniques currently used in LPC coders.

1. INTRODUCTION

LPC coefficients are used for the coding of speech because they provide an accurate and economical representation of relevant speech parameters. For low bit rate speech coders in particular, using LPC coefficients has proven to be a popular technique.

Only the coding of LPC coefficients is investigated in this work. The residual signal is passed directly to the receiver without any degradation. In this way, the effects of quantizing the LPC coefficients can be isolated from the effects of the coding of the residual signal. The diagram of the simulation model for studying the coding of LPC coefficients is shown in Fig. 1.

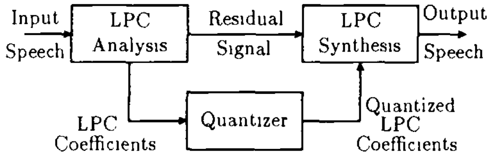

Fig. 1 Quantization of LPC coefficients

There are two basic approaches to quantizing the LPC coefficients. The first, *scalar quantization*, quantizes the LPC coefficients individually The second approach, *vector quantization*, the set of LPC coefficients are considered together as a vector [1], [2]. The disadvantage of vector quantization is the memory required to store the codebook and the number of computations used in comparing the input vector to each codebook vector. Both memory and number of computations increases as the size of the codebook increases. Hence there is a practical limit to the size of the codebook that can be employed.

One method to exploit the advantages offered by vector quantization while mitigating the practical problems is to use *vector-scalar quantization*. First, the input speech frame is vector quantized using a codebook with a moderate number of entries. From this vector quantization stage an error vector results. In the second stage of quantization, the components of the error vector are individually quantized.

The work reported here is only part of a larger study comparing scalar quantization with vector-scalar methods and examining difference error criteria. The emphasis will be on the vector-scalar quantizers which showed better performance than the best scalar quantizers (see for example [3], [4], [5]).

Two new techniques in vector-scalar quantization are introduced and evaluated in this paper. The first approach is to couple the vector and scalar quantization stages. The input LPC coefficient vector is compared to every codebook vector. From these comparisons, error vectors are determined. The components of these error vectors are scalar quantized. The resulting vectors from the overall vector-scalar quantization are all compared to the input vector to determine the closest one. In addition, methods to reduce the computational complexity are suggested.

The second innovation investigated is the incorporation of a small adaptive codebook to the larger fixed codebook. The self-training part of the codebook is based on the previous quantized input vectors. In one approach the adaptive codebook consists of a simple buffer of previously quantized input vectors. In another approach, several methods of constructing the prediction of the next input vector are made based on the previous quantized input vectors. Both approaches exploit the frame-to-frame correlation of the LPC coefficients. In this manner, increased performance is achieved with the vector-scalar quantization at no extra cost in bits. The non-differential portion of the codebook can handle abrupt changes in vector values resulting from abrupt change in vocal tract shape. Simple methods of limiting the propagation of errors in this partially differential scheme are suggested.

2. VECTOR-SCALAR QUANTIZATION

The vector-scalar quantization (VQ-SQ) technique takes advantage of the interparameter correlation between the LPC coefficients. In comparison to conventional vector quantization for a given number of bits, this hybrid can dramatically reduce the amount of memory required and the number of calculations performed.

A diagram of the vector-scalar quantization method is shown in Fig. 2. The LPC coefficients are first quantized using the vector quantization codebook. The error vector resulting from this stage is then scalar quantized using scalar quantization for each component of the vector. The index from the codebook (I_{VQ}) as well as the index vector from the SQ stage (I_{SQ}) are transmitted to the decoder.

Through transformations, the LPC coefficients can be represented in several domains. Two representations of the LPC coefficients that are frequently used for quantization are the *reflection* coefficients and the *Line Spectral Frequencies (LSF's)* [3],[6]. In this work, LSF's will be used.

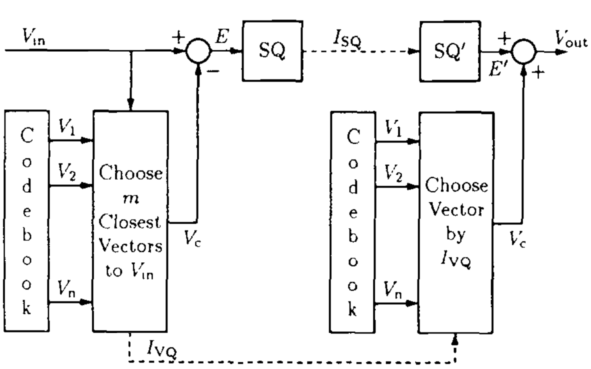

Fig. 2 LPC coefficient coder and decoder using vector quantization followed by scalar quantization

Several VQ-SQ coders were simulated. The codebooks for vector quantization were all trained using the Linde Buzo Gray (LBG) algorithm [7]. LSF vectors of length 10 for a speech frame of 20 ms are used for the LPC representation. The LSF errors were weighted according to the frequency (decreasing weight with increasing frequency) and according to the inter-LSF spacing (more weight to close LSF's). The scalar quantizers consider the LSF error vector. The scalar quantizer is non-uniform and more bits are allocated to the more perceptually important lower LSF's. The SQ levels were chosen based on histograms of the error vectors.

Eight quantizers, VQi-SQj, using vector quantization followed by scalar quantization were studied (i is the number of bits used for vector quantization and j is the number of bits used for scalar quantization).

These quantizers were evaluated using five test sentences. The average performance over the test sentences was evaluated using the average spectral distortion (SD) measure, SNR (dB) and segmental SNR (segSNR in dB). The average value of the spectral distortion measure is given in addition to its percentage of occurrence of spectral distortion values over 2 dB and 4 dB (see Table 1).

Quantizer	ave-SD	% > 2dB	% > 4dB	SNR	segSNR
VQ5-SQ25	1.52	23.59	3 75	11.03	12.34
VQ8-SQ22	1.66	31 08	2 50	11.78	12 92
VQ9-SQ21	1 65	29 49	2 50	12.23	12 50
VQ10-SQ20	1.60	25.95	1 25	10.66	11.63
VQ5-SQ16	2 08	51.77	5.56	7.05	9 99
VQ8-SQ13	2.06	47.78	3.75	5.95	7.70
VQ9-SQ12	2.04	48.71	3 75	6 48	7.86
VQ10-SQ11	1 95	41.59	1.85	6.01	7.98

Table 1 Results for VQ-SQ quantizers

The results show similar performance of the VQ-SQ quantizers to scalar quantizers studied. Perceptually experiments confirm this observation. The allocation of bits between the VQ stage and the SQ stage did not have a strong impact on the performance. We observed that the VQ-SQ technique gives poor frequency resolution in the SQ which can result in some large errors. The problem results from the relatively wide distribution of errors coming from the VQ stage. Histograms of the LSF errors indicate error is only marginally easier to quantize then the LSF's themselves, yet fewer bits are available after the VQ stage.

3. VECTOR QUANTIZATION COUPLED WITH SCALAR QUANTIZATION

An alternate approach to having the SQ stage following the VQ stage is to have the SQ stage nested within the VQ stage as shown in Fig. 3. The input vector is compared to each codebook vector and the error vector calculated. Each component of the error vectors is scalar quantized. The resulting quantized error vectors are then added back to the corresponding codebook vectors. The result is that for every codebook vector a new vector is formed From these new vectors, the smallest overall error is determined. The index of the codebook vector and the set of scalar quantization indices (determined for the error vector) associated with the new vector selected are passed to the decoder.

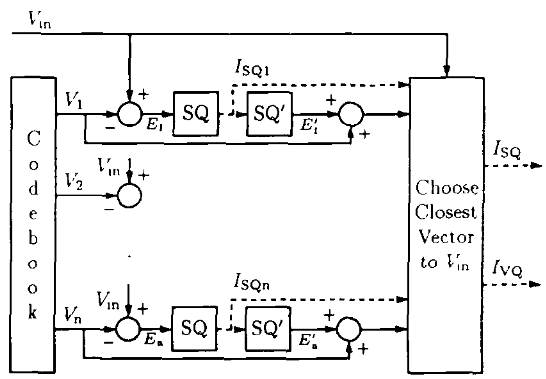

Fig. 3 LPC coefficient coder using vector quantization coupled with scalar quantization

In the nested VQ-SQ approach, the true closest vector resulting from the combination of VQ-SQ is chosen. A codebook vector that is not the closest one to the input vector may become the closest vector when combined with the scalar quantization. The coupled VQ-SQ scheme will always perform as well or better than when the two stages are performed sequentially. The disadvantage lies in the increased number of computations since the scalar quantization must be performed for all the vectors in the codebook for each frame of speech input. A modification that would reduce the amount of calculations is to have the VQ stage select the m best matches from the codebook. SQ is performed on these m matches and the final decision is made from these m combinations.

The results are shown in Table 2 for the coupled VQ-SQ quantizers for m equal to 512. Coupling the SQ and VQ stages results in better performance than when the stages were separate. The average spectral distortion as well as the number of outliers are reduced for the coupled VQ-SQ coders. Small improvements were seen in the segmental SNR value.

To examine the effects of lowering the value of m, the quantizer VQ9-SQ21-C was implemented with varying values of m. Table 3 shows that for values of m as low as 10 there is not a significant reduction in the performance of the coupled VQ-SQ quantizer. Hence the computational complexity of the coupled VQ-SQ quantization can easily be reduced without diminishing performance.

4. PARTIALLY ADAPTIVE VECTOR CODEBOOK

LPC parameters have frame-to-frame correlation that is not exploited in the VQ scheme previously examined. A method to incorporate frame-to-frame correlation in the VQ coders as developed in the previous section is shown in Fig. 4. The codebook is comprised of two sections; one which is fixed and one which is variable. The fixed section is from the regular trained codebook. The variable codebook is based on the last outputs of the coder, creating a buffer of previously quantized vectors. In this fashion,

Quantizer	ave-SD	% > 2dB	% > 4dB	SNR	segSNR
VQ5-SQ25-C	1 26	13 03	2.50	10.50	12.53
VQ8-SQ22-C	1.12	10.93	0 0	10.17	13.82
VQ9-SQ21-C	1 05	8.69	0 0	10.23	12.93
VQ10-SQ20-C	1 00	7.51	0.0	10.03	12.94
VQ5-SQ16-C	1.84	33 56	5 00	7.42	10.04
VQ8-SQ13-C	1 51	20 27	2 50	7.59	9 27
VQ9-SQ12-C	1 51	20.57	0 0	6.35	8.99
VQ10-SQ11-C	1 42	16.49	0.0	6.51	9 00

Table 2 Results for coupled VQ-SQ quantizers

m	ave-SD	% > 2dB	% > 4dB	SNR	segSNR
1	1 65	29 49	2 5	12 23	12 50
5	1 24	12 57	0 0	11 61	13 51
10	1 15	10 60	0 0	10 72	13 18
25	1 08	9 81	0 0	10.05	12 96
50	1 05	8 88	0 0	10 56	12.94
512	1.05	8 69	0.0	10 23	12 93

Table 3 Results for quantizer VQ9-SQ21-C with varying values of m.

the VQ-SQ coder can use frame-to-frame correlation if it exists or it can rely on the fixed codebook if a large correlation is not present. A coder with only a fixed codebook cannot benefit from the correlation. Further, a codebook based completely on previous frame vectors can perform poorly if it so happens that there are suddenly large differences in the LPC vectors frame-to-frame such as during an abrupt change in vocal tract shape.

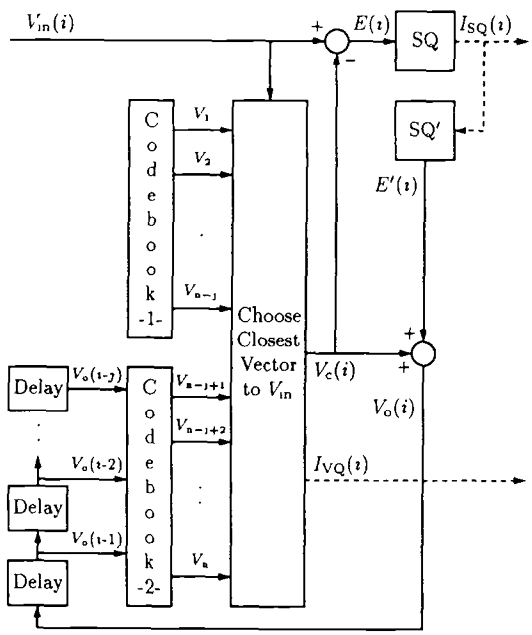

Fig. 4 LPC coefficient coder using vector quantization followed by scalar quantization In the vector codebook, a buffer of the past quantized vectors of length j is used

The use of previous quantized vectors to be included in part of the codebook can be termed as optional differential time domain coding. The cost for this differential coding is quite low. Later results show that only a small number of vectors in a codebook of 512 are required for differential coding.

An improvement to the partially adaptive quantizers is to include the technique previously examined of coupling the vector and scalar quantization stages together. The results of the quantizers using this scheme are shown in Table 4.

Quantizer	ave-SD	% > 2dB	% > 4dB	SNR	segSNR
VQ5-SQ25-AC	1.19	11.02	0.93	12.02	13.37
VQ8-SQ22-AC	1.06	8.88	0.0	11.28	14.48
VQ9-SQ21-AC	1.01	8.19	0.0	10.96	13.43
VQ10-SQ20-AC	0.97	6.33	0.0	10.36	13 19
VQ5-SQ16-AC	1.87	37.81	1.85	8.38	10.32
VQ8-SQ13-AC	1.47	18.59	2.50	7.53	9.68
VQ9-SQ12-AC	1 49	19.89	0 0	6 19	9 12
VQ10-SQ11-AC	1.39	15 44	0.0	6.32	9.29

Table 4 Results for coupled VQ-SQ quantizers using an partially adaptive codebook

All eight quantizers improved with the addition of the variable component to the codebook. The previous quantized vectors were selected a significant fraction of the time (around 35 % of frames).

It is of interest to determine the ideal length of the buffer of previously quantized vectors. Experiments showed that the previous vector is chosen very frequently (one third of the time) with the next five being selected occasionally. The rest of the vectors are chosen as frequently as any other vector in the fixed codebook (around 0.2 % of the time). Hence a good choice would be to store the previous six quantized input vectors in the variable codebook.

The amount of correlation that exists between the LPC parameters of a speech file and its previous frame can be roughly divided into three categories; little, some and considerable. The case of little correlation is handled by the fixed codebook section while the variable codebook section handles the cases of some and considerable correlation. An improvement would be to have two scalar quantizers available for the variable section of the codebook. One scalar quantizer for the case of some correlation and scalar quantizer with smaller quantization levels for the case of considerable correlation. A diagram showing the implementation of these two scalar quantizers is shown in Fig. 5. The buffer of previously quantized vectors is stored in two separate codebooks. One of these codebooks uses the regular scalar quantizer while the other uses the quantizer with small quantization levels. The decoder uses the scalar quantizer as determined by the vector quantization index. The cost of this scheme is the small reduction of the fixed codebook size as the variable codebook size codebook is doubled. The overall bit-rate of the coder, however, is not increased. Results of the quantizers using this scheme are shown in Table 5.

Quantizer	ave-SD	% > 2dB	% > 4dB	SNR	segSNR
VQ5-SQ25-A2C	1 25	14 17	1 85	13 20	14 21
VQ8-SQ22-A2C	1 07	9.07	0 0	12 24	15.14
VQ9-SQ21-A2C	1.02	8.88	0 0	11 60	14 28
VQ10-SQ20-A2C	0.96	5 90	0 0	12 08	14.36
VQ5-SQ16-A2C	1.98	37.88	11.11	9.94	11.30
VQ8-SQ13-A2C	1.47	19.82	3.75	8.23	10 76
VQ9-SQ12-A2C	1.48	19.64	0.0	7 49	9.96
VQ10-SQ11-A2C	1.39	15.25	0.0	7 33	9.84

Table 5 Results for coupled VQ-SQ quantizers using an partially adaptive codebook with two quantizers available to the adaptive codebook

The quantizers select the scalar quantizer with finer quantization levels frequently. The result is better performance of around 1 dB segSNR higher than the quantizers that did not use two scalar quantizers.

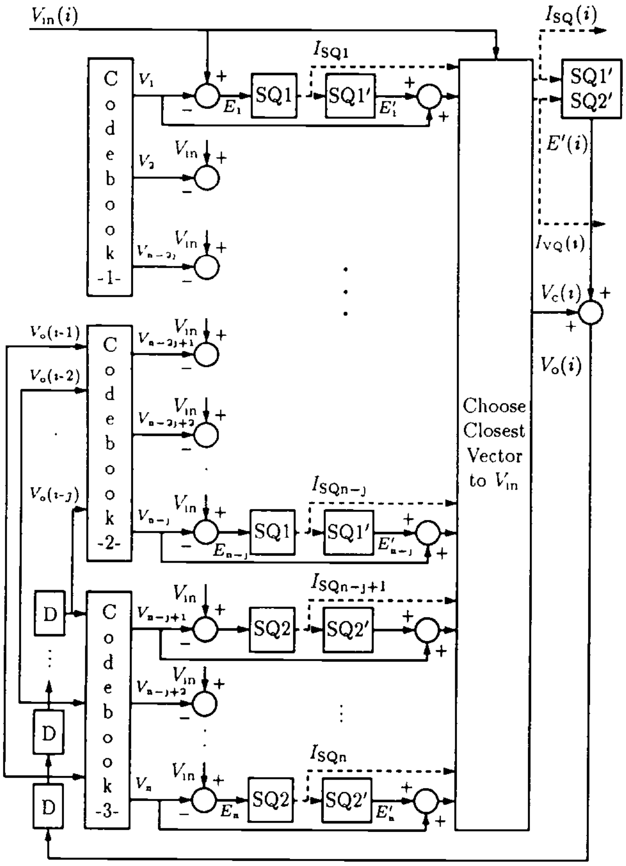

Fig. 5 LPC coefficient coder using vector quantization coupled with scalar quantization. In the vector codebook, a buffer of the past quantized vectors of length j is used with two scalar quantizers available for the buffer vectors.

The idea of predicting the input vector can be incorporated into the VQ coder developed so far in this work. Several simple linear prediction schemes using up to 3 past vectors are used to predict several possibilities of the input vector. These predictions are then placed in a small section of the codebook. Also, two scalar quantizers can be used for the predicted vectors. If a very good prediction is made, a SQ with small quantization steps can be used while if a good prediction is made, a SQ with medium quantization levels can be used. If the predicted vectors are all poor, the coder would pick a vector from another part of the codebook.

The hoped for improvements resulting from using predictive techniques were not realized (see Table 6). The predicted vectors were selected fairly often (15 % of frames) but there was a proportionable decrease in the number of vectors selected from the buffer of previous vectors (20 %, down from 35 %). This indicates that the past input vectors offer as good a prediction of the next input vector as the prediction schemes used in the coder.

One problem with differential coding is the propagation of error that can result from channel-induced decoding errors. A method of limiting this problem is to prevent a vector from the variable section of the codebook being returned to the codebook a second time. Hence if the receiver incorrectly stores a vector, the vector can only be selected a fixed number of times ($n-i$, where n is the total number of codebook vectors and i is the length of the buffer) before being discarded. A second solution

Quantizer	ave-SD	%>2dB	%>4dB	SNR	segSNR
VQ5-SQ25-AP2C	1 28	13.80	0.0	12.07	13.78
VQ8-SQ22-AP2C	1.05	8 07	0.0	12.00	15.18
VQ9-SQ21-AP2C	0.98	7 77	0 0	12.16	14.55
VQ10-SQ20-AP2C	0.96	5.90	0.0	12.08	14.28
VQ5-SQ16-AP2C	2 19	50.19	15 00	9 24	10 87
VQ8-SQ13-AP2C	1 45	18.22	3 75	7 84	10 63
VQ9-SQ12-AP2C	1 50	20.01	0 93	7 76	9 74
VQ10-SQ11-AP2C	1 38	15 00	0 0	7 28	9.82

Table 6 Results for coupled VQ-SQ quantizers using a partially adaptive codebook with two quantizers available to the adaptive codebook

is to force the coder to choose a vector from the fixed codebook every so often.

5. CONCLUSION

Vector-scalar quantization is a two stage coding scheme that exploits the advantages offered by vector quantization while drastically reducing the memory and computational requirements for a given number of bits per frame. Two new techniques were developed to increase the performance of the vector-scalar quantizers without increasing the bit-rate.

The first implementation of vector-scalar quantization does not perform better than straight scalar quantization. Coupling the vector and scalar stages helped considerably. Note that the scalar stage was the same for all VQ outcomes. Further improvements may be obtained if the SQ is tailored to each VQ outcome. We achieve a similar effect in our partially adaptive codebook by SQ allowing for two different scalar quantizers.

The final results of this paper show that the best of the vector-scalar quantization techniques is a good method of coding LPC coefficients. The performance of the vector-scalar quantization is better than that of the conventional scalar quantization methods examined. For similar bit rates, the vector-scalar quantization with the use of the two new techniques introduced has significantly lower average spectral distortion and fewer outliers. This reduction can be important for low-bit rate coders. A further area of research is the prediction techniques used in the adaptive codebooks.

REFERENCES

1. R. M. Gray, "Vector quantization", *IEEE Acoust. Speech, Signal Processing Mag.*, pp. 4-29, April 1984.
2. J. Makhoul, S. Roucos and H. Gish, "Vector quantization in speech coding", *Proc. of IEEE*, vol. 73, pp. 1551-1588, Nov. 1985.
3. O. Ghitza and J. L. Goldstein, "Scalar LPC quantization based on formant JND's", *IEEE Trans. Acoust. Speech, Signal Processing*, vol. ASSP-34, pp. 697-708, Aug. 1986.
4. G. S. Kang and L. J. Fransen, "Low-bit rate speech encoders based on line-spectrum frequencies (LSFs)", *Naval Research Laboratory Report 8857*, Nov. 1984.
5. J. R. Crosmer and T. P. Barnwell, III, "A low bit rate segment vocoder based on line spectrum pairs", *Proc. Int. Conf. on Acoust. Speech, Signal Processing*, Tampa, Florida, pp. 7.2.1-7.2.4, April 1985.
6. G. S. Kang and L. S. Fransen, "Application of line spectrum pairs to low bit rate speech encoders", *Proc. Int. Conf. on Acoust. Speech, Signal Processing*, Tampa, Florida, pp. 7.3.1-7.3.4, April 1985.
7. Y. Linde, A. Buzo and R. M. Gray, "An algorithm for vector quantization design", *IEEE Trans. on Commun.*, vol. COM-28, pp. 84-95, Jan. 1980.