and express coefficients b_n of (3) as

$$b_n = c_n B_n, \quad 0 \le n \le N-1$$
$$B_n = 2^m, \quad m = 0, \pm 1, \pm 2, \cdots$$
$$\tfrac{1}{2} < |c_n| \le 1. \tag{6}$$

If c_n^* is the version of c_n rounded to $t+1$ bits, the associated rounding error is

$$\epsilon_n' \triangleq c_n^* - c_n \tag{7}$$

and the quantization step is $Q = 2^{-t}$, with $|\epsilon_n'| \le Q/2$. As is customary, the error sequence ϵ_n' with $n = 0, 1, \cdots, N-1$ is considered a sequence of independent random variables uniformly distributed in $[-(Q/2), (Q/2)]$. Error (7) will be very useful for comparisons with the direct form structure.

The permuted coefficients quantization error is

$$E_n \triangleq \bar{a}_{p_n} - a_{p_n} = a_{p_n}\left(\frac{b_n^*}{b_n} - 1\right) = a_{p_n}\frac{\epsilon_n'}{c_n} \quad 0 \le n \le N-1. \tag{8}$$

If $H^*(e^{j\omega})$ denotes the frequency response of the nested structure when quantization is taken into account, i.e., when coefficients a_{p_n} of (2) are replaced with $\bar{a}_{p_n}$, the frequency response error due to quantization is

$$E(e^{j\omega}) \triangleq H^*(e^{j\omega}) - H(e^{j\omega}) = \sum_{n=0}^{N-1} E_n e^{-j\omega p_n}. \tag{9}$$

An upper bound for (9) is given by

$$|E(e^{j\omega})| \le \sum_{n=0}^{N-1}|E_n| = \sum_{n=0}^{N-1}|a_{p_n}|\left|\frac{\epsilon_n'}{c_n}\right| < Q\sum_{n=0}^{N-1}|a_{p_n}| = Q\sum_{n=0}^{N-1}|a_n|. \tag{10}$$

The variance of (9) is

$$\overline{|E(e^{j\omega})|^2} = \overline{\left(\sum_{n=0}^{N-1}E_n\cos(\omega p_n)\right)^2 + \left(\sum_{n=0}^{N-1}E_n\sin(\omega p_n)\right)^2}$$
$$= \sum_{n=0}^{N-1}\left(\frac{a_{p_n}}{c_n}\right)^2\overline{\epsilon_n'^2}\cos^2(\omega p_n) + \sum_{n=0}^{N-1}\left(\frac{a_{p_n}}{c_n}\right)^2\overline{\epsilon_n'^2}\sin^2(\omega p_n)$$
$$= \frac{Q^2}{12}\sum_{n=0}^{N-1}\left(\frac{a_{p_n}}{c_n}\right)^2 < \frac{Q^2}{3}\sum_{n=0}^{N-1}a_{p_n}^2 = \frac{Q^2}{3}\sum_{n=0}^{N-1}a_n^2 \tag{11}$$

where the overbar indicates the expectation operator.

III. Direct Form Frequency Response Error Due to Coefficient Quantization

Similar results are obtainable for direct form structures if the coefficients a_n of (1) are represented in a way similar to (6), i.e.,

$$a_n = A_n c_n, \quad n = 0, 1, \cdots, N-1$$
$$A_n = 2^m, \quad m = 0, \pm 1, \pm 2, \cdots$$
$$\tfrac{1}{2} < |c_n| \le 1. \tag{12}$$

Define $a_n^* = A_n c_n^*$ and

$$E_n^{DF} \triangleq a_n^* - a_n = A_n\epsilon_n' = a_n\frac{\epsilon_n'}{c_n} \tag{13}$$

and denote with $H_{DF}^*(e^{j\omega})$ the frequency response of the direct form structure after coefficient quantization.

The frequency response error is

$$E_{DF}(e^{j\omega}) = H_{DF}^*(e^{j\omega}) - H_{DF}(e^{j\omega}) = \sum_{n=0}^{N-1}E_n^{DF}e^{-j\omega n}. \tag{14}$$

An upper bound for (14) is

$$|E_{DF}(e^{j\omega})| \le \sum_{n=0}^{N-1}|E_n^{DF}| = \sum_{n=0}^{N-1}\left|a_n\frac{\epsilon_n'}{c_n}\right| < Q\sum_{n=0}^{N-1}|a_n|. \tag{15}$$

The variance of the frequency response error is

$$\overline{|E_{DF}(e^{j\omega})|^2} = \overline{\left(\sum_{n=0}^{N-1}E_n^{DF}\cos(\omega n)\right)^2 + \left(\sum_{n=0}^{N-1}E_n^{DF}\sin(\omega n)\right)^2}$$
$$= \sum_{n=0}^{N-1}\left(\frac{a_n}{c_n}\right)^2\overline{\epsilon_n'^2}\cos^2(\omega n) + \sum_{n=0}^{N-1}\left(\frac{a_n}{c_n}\right)^2\overline{\epsilon_n'^2}\sin^2(\omega n)$$
$$= \frac{Q^2}{12}\sum_{n=0}^{N-1}\left(\frac{a_n}{c_n}\right)^2 < \frac{Q^2}{3}\sum_{n=0}^{N-1}a_n^2. \tag{16}$$

The final expression and the upper bound at the right-hand side of (16) are identical to those of the right-hand side of (11).

IV. Conclusions

The above analysis shows that identical expressions are obtainable for the frequency response errors of both direct form structures and nested structures. The key to having identical coefficient sensitivity is to use the same scaled representation for the coefficients.

Acknowledgment

The authors would like to acknowledge the reviewers for their comments and suggestions, which greatly improved the quality of this presentation.

References

[1] A. Mahanta, R. C. Agarwal, and S. C. Dutta Roy, "FIR filter structures having low sensitivity and roundoff noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-30, no. 6, pp. 913-920, Dec. 1982.

Objective and Subjective Optimization of APC System Performance

JERRY D. GIBSON AND WEN-WHEI CHANG

***Abstract*—System configurations and relative performance results are presented for adaptive predictive coding (APC) systems at 9.6 kb/s optimized based upon three different criteria, subjective listening tests,**

Manuscript received September 28, 1988; revised August 7, 1989.

J. D. Gibson is with the Department of Electrical Engineering, Texas A&M University, College Station, TX 77843.

W.-W. Chang is with the Department of Communication Engineering, National Chiao Tung University, Taiwan.

IEEE Log Number 9034982.

0096-3518/90/0600-1053$01.00

signal-to-quantization noise ratio (SNR), and segmental SNR (SNRSEG). The subjectively optimized coder has little granular noise but some spectral distortion and SNR and SNRSEG values less than half those of the other two systems. The objectively optimized systems have significant granular noise but a better reproduction of the signal spectrum. Optimization of SNRSEG does not lead to the subjectively preferred design.

I. Introduction

Adaptive predictive coding has been an important technique for speech digitization at 9.6 kb/s and below since it was first introduced by Atal and Schroeder in 1970 [1]. Here we use adaptive predictive coding (APC) to denote a system that removes both long-term and short-term redundancy prior to quantization and coding of the residual, as is common practice [1]-[3]. Since its introduction, there have been numerous improvements to APC systems, including noise spectral shaping (NSS), multitap long-term predictors, more efficient quantization of side information, and center clipping of the prediction residual, and although we investigate these various modifications here, the detailed descriptions are left to [2] and [3].

It is by now well established that objective performance measures such as signal-to-quantization noise ratio (SNR) are not reliable indicators of the quality and intelligibility of a speech coder's output [3], [4], although segmental SNR has been shown to have a better correlation with coder subjective performance than the standard SNR. These latter comments concerning objective performance measures are particularly true at the data rates of interest here (below 10 kb/s), even though papers on multipulse linear predictive coding (LPC) and code-excited LPC (CELP) often present SNR and segmental SNR values. Of course, there are other objective performance measures, such as the articulation index, and each has proven useful in particular situations [4].

The use of subjective performance indicators presents its own set of problems. Given sufficient time and resources, it is possible to adopt formal testing procedures such as the diagnostic rhyme test (DRT) [5] and the diagnostic acceptability measure (DAM) [6], or with somewhat less of a commitment, establish a mean opinion score (MOS) [4] for a coder. Of course, when one is actually trying to design a coder, informal listening tests are the dominant subjective evaluation procedure because of their relative ease of implementation.

We report here on the adjustment of various parameters of an APC system to optimize, respectively, (informal) subjective performance, the SNR, and the segmental SNR. The data rate is constrained to be 9.6 kb/s or less. The results substantiate some of the existing notions already discussed while at the same time raising a new issue or two. It is hoped that results such as these will be useful to engineers who are considering, or just beginning, the development of an APC system.

II. APC Systems and Subjective Optimization

The block diagram of a generalized APC system is shown in Fig. 1. Depending upon the choice of the transfer functions in the several blocks, a number of different APC system configurations can be realized [1]-[3], [7]. For example, with $H = I$ (the identity), P_1 a short-term predictor, P_2 a long-term predictor, and F appropriately chosen, we have an APC system with noise spectral shaping where both predictors have the quantizer within the loop. An alternative form is obtained if $H = 1 - P_3$, P_3 a long term predictor, $P_2 = 0$, and all other system quantities as before, where now only the short-term prediction is based upon the quantized speech. There are many other possibilities, and some of these are discussed in the references [3], [7].

Our APC system design routine consisted of starting with configurations and methods presented in the paper by Atal [3], synthesizing speech using the various APC systems, and then selecting that system which gave the best subjective output speech quality as judged by informal listening tests. Our only constraint was a maximum data rate of 9.6 kb/s, and we investigated a wide variety of system design details, including long- and short-term predictor orders, quantization procedures, pitch extractors, frame lengths, and bit allocations among parameters, to name a few. We studied most common APC system configurations, and our speech data base was about 9 minutes of speech, with 5 female and 9 male speakers, which included telephone recorded speech as well as directly digitized speech.

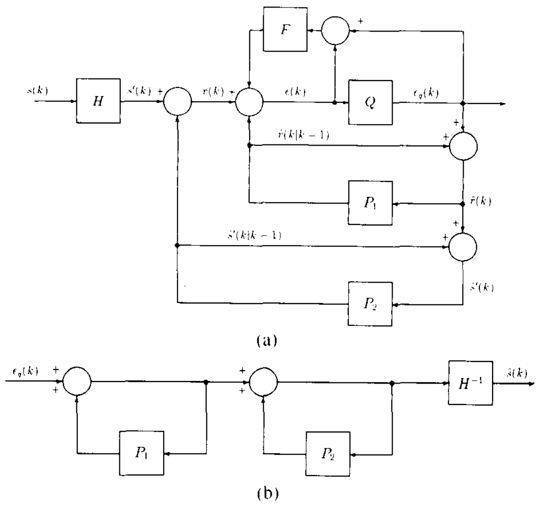

Fig. 1. (a) APC system transmitter. (b) APC system receiver.

TABLE I
Bit Allocations per Frame for Subjectively Optimized APC System (APC1)

Parameter	No. of Bits Assigned
Pitch	7
Step Size	10 (Twice)
β_1	4
β_2	5
β_3	4
k_1	7
k_2	7
k_3	5
k_4	5
k_5	4
k_6	4
k_7	4
k_8	3

The APC system that produced the best subjective output speech quality according to our informal listening tests has the following characteristics. It uses a three-tap long-term (pitch) predictor, an eighth-order short-term predictor, and three-level, center-clipped quantization of the residual. With respect to Fig. 1, $H = 1 - P_3$, P_3 a long-term predictor, $P_2 = 0$, and P_1 a short-term predictor, where the P_1 coefficients are calculated using the autocorrelation method and the P_3 parameters are calculated as described in [3]. Stabilization of the three long-term predictor coefficients is necessary and was accomplished via the procedures developed by Ramachandran and Kabal [8]. The analysis frame length is 160 samples, or 25 ms, at the chosen sampling rate of 6400 samples/s. The step size is updated twice per frame, while all other parameters are transmitted once per frame. The total number of bits allocated per frame is 79, with the breakdown according to parameters as shown in Table I, where the β_i represent the long-term predictor coefficients and the k_i are the short-term predictor reflection coefficients. A variable length code, which maps two ternary digits into 1 to 6 binary digits, is employed to transmit the three-level quantizer output at an average rate of 1 b/sample, so with the rate of 3160 b/s for the side information, the total transmitted data rate is 9560 b/s. Our final design does not include noise spectral shaping because although it improved the quality of some speakers, it seemed to slightly degrade others. For further discussion, we denote this subjectively designed system as APC1.

Note that what we are claiming about the APC1 system is that it provides good quality, intelligible output speech based upon informal subjective listening tests. We are *not* stating that this is *the* subjectively optimal APC system configuration. The experimental design procedure, performed over a period of several months, required numerous tradeoffs to be made, and when optimizing a particular quantity, other system parameters were set at some nominal values. A change in these nominal values could cause a different sequence of decisions. However, we do feel that, given our particular system constraints, the invaluable information in [3], and our experience in speech coding, the final APC1 system output speech quality is comparable to any alternative APC system design.

III. Objective APC System Designs

We also undertook a series of experiments to optimize APC system designs based upon the two objective criteria, signal-to-quantization noise ratio defined as

$$\text{SNR} = 10 \log_{10} \frac{\langle s^2(k) \rangle}{\langle [s(k) - \hat{s}(k)]^2 \rangle} \tag{1}$$

where $\langle \cdot \rangle$ denotes time averaging over the complete utterance, $\{s(k)\}$ is the input sequence, and $\{\hat{s}(k)\}$ is the output sequence, and segmental SNR given by

$$\text{SNRSEG} = \frac{1}{\text{NB}} \sum_{i=1}^{\text{NB}} \text{SNR}_i \tag{2}$$

where NB is the total number of frames or blocks in an utterance and SNR_i is computed according to (1) for each block.

As before, an extensive set of experiments was performed to optimize the APC system to yield the maximum SNR. The final APC design that gave the highest SNR, denoted APC2, has the same configuration with respect to Fig. 1 as APC1, a three tap long-term predictor, an eighth-order short-term predictor, a three-level quantizer, and a 160 sample frame length, and all parameters (including step size) are transmitted once per frame. The total number of bits per frame is 79 with bits allocated to parameters as listed in Table II.

The design of an APC system to maximize the segmental SNR in (2) resulted in a system with the same configuration, components, and bit allocations as for APC2 but with a center-clipped three-level quantizer design similar to APC1. Specifically, APC1 has a smaller center-clipping threshold and a smaller step size than APC2, and the SNRSEG-optimized APC system, designated APC3, has much the same threshold and step size as APC1. See Appendix B.

We now provide some additional details concerning the relative contributions and importance of the several APC2 and APC3 system parameters. All frame rate, bit allocation, and predictor order decisions are constrained by the overall 9.6 kb/s data rate. With a sampling rate of 6400 samples/s and 1 b/sample allocated to quantize the residual error signal, this leaves 3.2 kb/s for the side information. The three-tap long-term predictor makes an important contribution to the performance of all three APC systems that cannot be accounted for elsewhere. For the APC2 and APC3 coders, the three-tap long-term predictor provides a nominal 2-dB advantage in SNR and SNRSEG, respectively, over that produced by the single-tap long-term predictor. We examined both autocorrelation based and average magnitude difference function (AMDF) [4], [7] type pitch extractors and found no clear difference in overall system

TABLE II
Bit Allocations per Frame for SNR Optimized APC System (APC2)

Parameter	No. of Bits Assigned
Pitch	7
Step Size	5
β_1	7
β_2	7
β_3	7
k_1	8
k_2	8
k_3	6
k_4	6
k_5	5
k_6	5
k_7	4
k_8	4

performance between the two. However, it is imperative that the pitch extraction algorithms operate on every sample and not skip samples as is sometimes recommended to reduce computations.

To maximize SNR and SNRSEG, we found that the maximum number of taps possible in the short-term predictor should be used and that both long- and short-term predictor parameters should be allocated as many bits as possible. It is clear in comparing Tables I and II that more accurate representations of the predictor parameters is a primary difference between APC1 and the objectively optimized coders.

Increasing the frame rate above 40 frames/s entails reducing the number of parameters in the side information or using fewer bits to represent the parameters. Performing these tradeoffs does not produce increases in SNR and SNRSEG and so a faster frame rate was not adopted. A slower frame rate of 32 frames/s would allow a higher order short-term predictor to be used, but at the expense of performance during spectral transitions, and hence, is not attractive.

The APC2 coder has a minimum threshold (step point) for the zero level quantization region that is more than six times larger than that for APC1, the step size is updated only half as often, and fewer bits are used to quantize the step size. The APC3 system uses a quantizer nearer the APC1 coder in order to reduce the granular noise. It is evident, however, that SNR and SNRSEG-based coders put less emphasis on the minimization of granular noise than APC1.

For all three systems, the step size is logarithmically quantized, the pitch and β_i coefficients are uniformly quantized, and the reflection coefficients are uniformly quantized. Other quantization methods were examined, the most notable of which was the quantization of log area ratios rather than reflection coefficients. No significant advantage was found. It is felt that APC is perhaps less sensitive to coefficient quantization methods than LPC. The allocation of bits among the eight reflection coefficients was investigated with the major conclusion being that the first two coefficients should receive the most bits. Taking bits away from the k_5 through k_8 allocations in Table II to give to k_3 and k_4 generally reduces performance.

IV. System Performance Comparisons

The SNR and SNRSEG values for the APC1, APC2, and APC3 systems are listed in Table III for the five sentences described in Appendix A. One thing that is extraordinarily striking in Table III is the extremely low values of SNR and SNRSEG for the APC1 coder compared to the other two. Of course, since APC1 was designed independently of these quantities we should not be surprised, but differences of greater than a factor of two point out how ineffective these two objective measures can be for system design.

To develop some feeling as to why the APC1 system is subjectively preferable to APC2 and APC3 in spite its much lower SNR

TABLE III
OBJECTIVE PERFORMANCE OF APC SYSTEMS

	SNR(dB)			SNRSEG(dB)		
Sentence	APC1	APC2	APC3	APC1	APC2	APC3
1	6.19	16.48	15.95	4.78	12.90	13.42
2	5.87	16.53	15.87	5.31	12.85	12.88
3	4.89	13.76	13.23	4.14	11.36	11.64
4	4.45	13.03	12.77	4.91	12.27	12.90
5	5.59	14.31	13.87	5.48	12.33	13.34

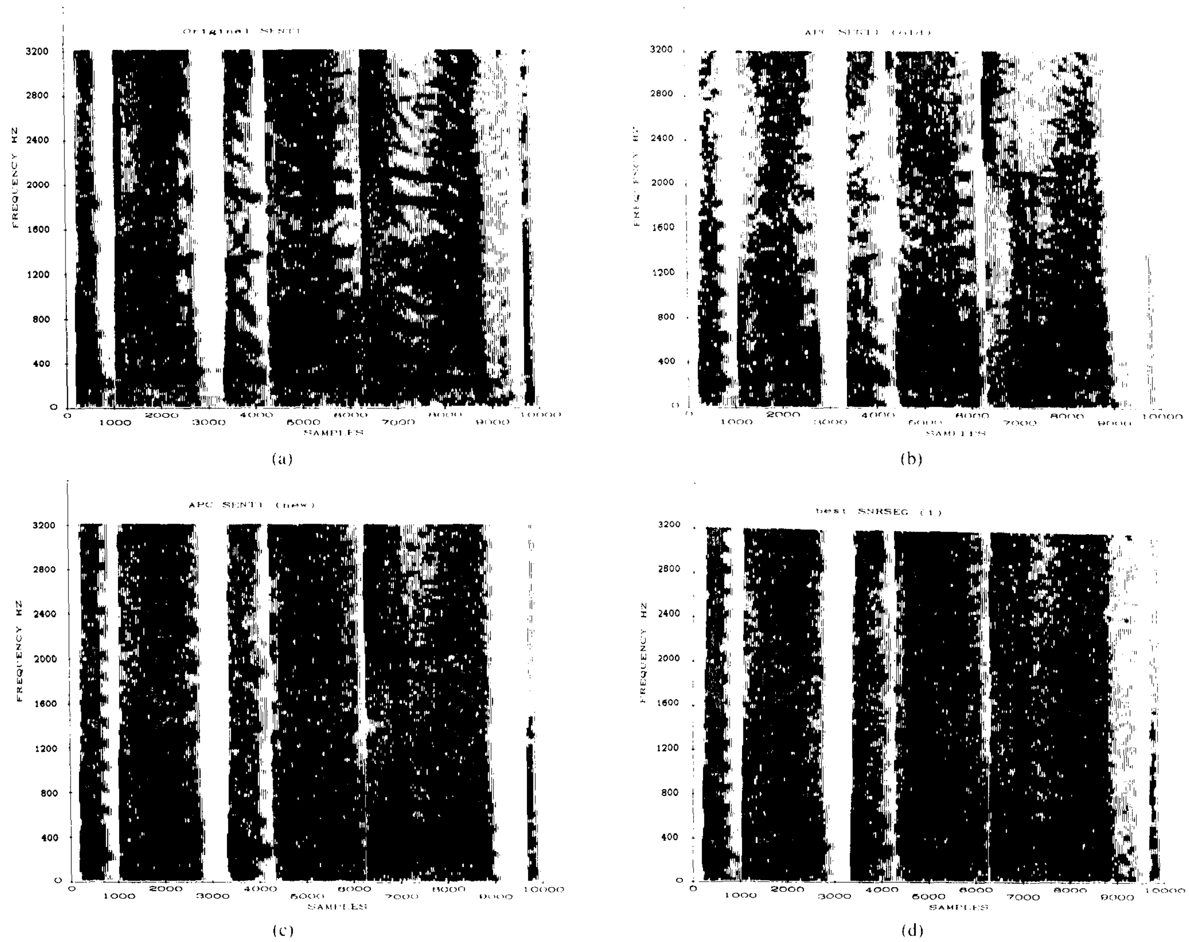

Fig. 2. Spectrogram comparisons—Sentence 1. (a) Original. (b) Subjectively optimized—APC1. (c) SNR optimized—APC2. (d) SNRSEG optimized—APC3.

and SNRSEG values, we point to two factors. First, the APC1 coder has much less granular noise at high frequencies and between formants. This claim is supported by the narrow-band spectrograms in Fig. 2, which show the first 10 000 samples of the original, the APC1 output, the APC2 output, and the APC3 output in parts (a)–(d), respectively. Compare especially the granular noise at 3500–4000 samples, around 6000 samples, at 7000–8000 samples, and at high frequencies throughout. This granular noise is particularly audible in the APC2 and APC3 coder outputs. A second reason for the subjective preference for APC1 is that the reconstruction error in APC1 has a spectral content that is close to the spectrum of the original speech. Fig. 3 shows the spectrum of a 25 ms frame of speech in (a), followed by spectra of the corresponding reconstruction errors for the APC1, APC2, and APC3 systems in (b)–(d), respectively. Clearly, the APC1 reconstruction error spectrum much more closely duplicates that of the original speech than either the APC2 or APC3 error spectrum, and thus, even though the reconstruction error power is greater than APC1, it is sufficiently signal correlated to be less objectionable subjectively. The differences between the APC2 and APC3 outputs are less easy to discern, although a comparison of Fig. 2(c) and (d) around 9000 samples shows a preference for the SNRSEG optimized system.

A comparison of the subjective quality of the three systems is also instructive. The APC1 coder is judged to be perceptually better, almost solely because of the lack of quantization noise and in spite of the presence of spectral distortion. The APC2 output has a greatly increased level of granular noise over APC1, but, if this noise can be ignored, the speech sounds much crisper due to the

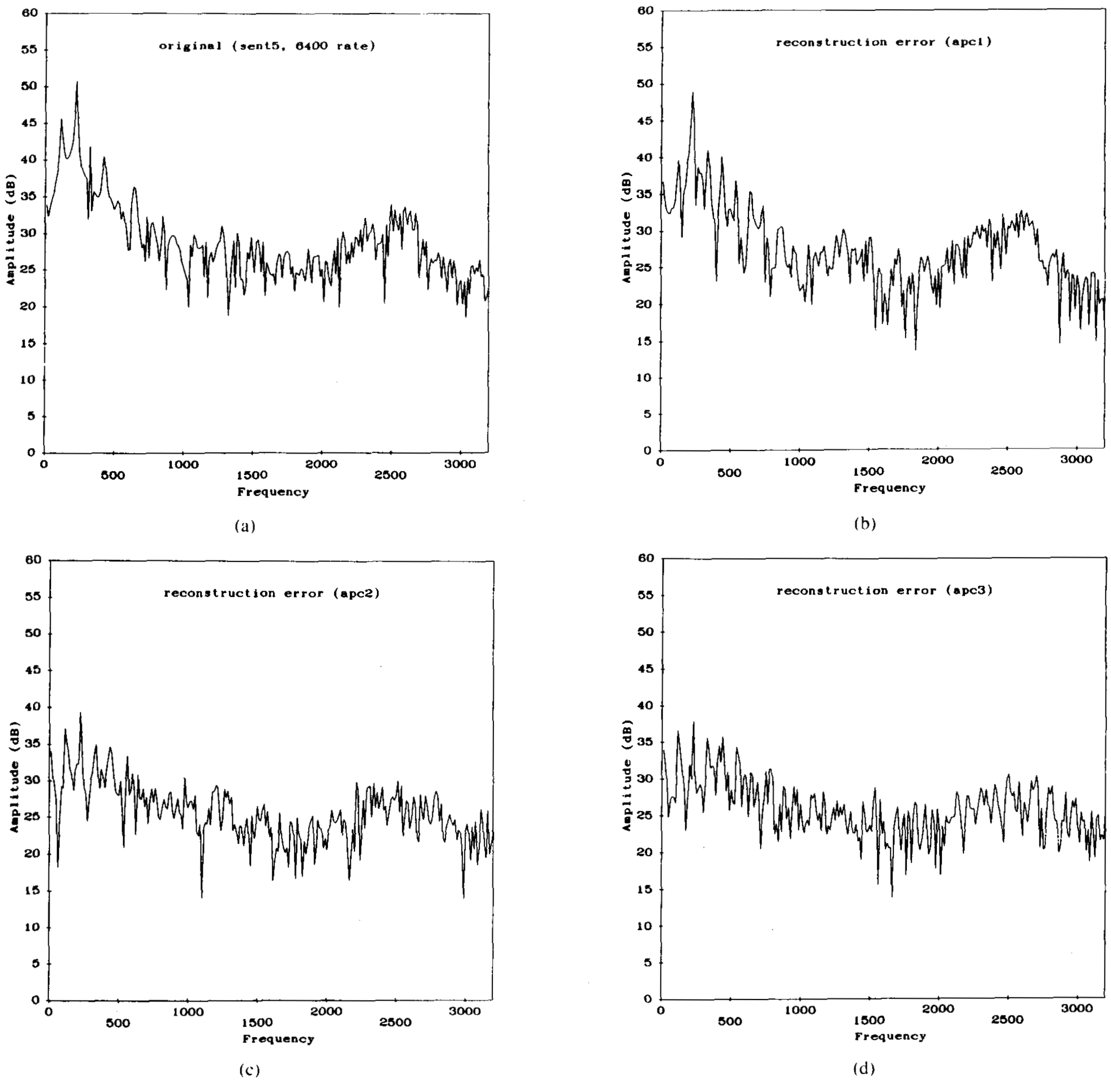

Fig. 3. Spectrum comparisons—Sentence 5. (a) Original speech spectrum. (b) APC1 reconstruction error spectrum. (c) APC2 reconstruction error spectrum. (d) APC3 reconstruction error spectrum.

better spectral representation. The SNRSEG optimized APC3 coder falls between the other two systems in that it has increased granular noise over APC1 but some noticeable spectral errors compared to APC2. This discussion accentuates a difficulty with subjective listening tests, either formal side-by-side preference tests or informal listening tests, in that the choices are subjective and influenced by personal preference. For example, a side-by-side listening test of APC1 and APC2 might result in a preference for APC1 because of much less granular noise or, for those who could focus on the speech quality and intelligibility and ignore the hissing, a preference for APC2. It is our experience that most (but not all) listeners would prefer APC1 but some object to spectral distortions and select APC2 as best.

Another interesting point is that although the APC3 output is subjecively preferable to APC2, the optimization of SNRSEG did not lead to the subjectively optimized APC1 design. We feel that this occurs because maximizing any unweighted SNR criterion (or minimizing unweighted mean squared error) is equivalent to whitening the unpredictable part of the input speech (the residual error). The resulting whitened error sequence, after quantization, is a perceptually less pleasing excitation for the decoder synthesis filter. The APC1 coder allocates a greater portion of the bit rate to accurately quantizing the residual error greater than the center-clipping threshold, and thus retains any pitch-related redundancies that have not been removed. Hence, SNRSEG may be more indicative of perceptual cues than SNR, but it still falls short of a purely subjective criterion.

V. Conclusions

The comparison of the subjectively and objectively optimized APC systems presented here points out the difficulties involved in coder performance optimization during the design phase when the time (or perhaps, financial support) is not usually available for formal testing, such as the DRT or DAM, after each system modification. Certainly, as pointed out by this research, objective measures can vary drastically over different coder designs and subjective judgements depend on whether a listener finds granular or spectral distortion more objectionable.

Appendix A
Sentence Data

Results presented in this paper are for five sentences. Each of the sentences was low-pass filtered to 3200 Hz, sampled 6400 times/s, and digitized to 12 b accuracy. The exact utterances are as follows.

Sentence	
1	"The pipe began to rust while new." (female speaker)
2	"Add the sum to the product of these three." (female speaker)
3	"Thieves who rob friends deserve jail." (male speaker)
4	"Oak is strong and also gives shade." (male speaker)
5	"Cats and dogs each hate the other." (male speaker)

Appendix B
Quantizer Parameters

All three APC systems have a three-level, center-clipped quantizer, however, the center-clipping threshold T, and the output step size Δ, are calculated differently for each coder. To specify these quantities, we note that the number of samples per frame is 160 and that the quantizer input sequence is $\{e(k)\}$.

APC1

The step size and threshold are updated twice per frame according to

$$\Delta = 0.875\left[\frac{\sum_{k=1}^{80}|e(k)|}{80}\right]$$

and

$$T = \max\left\{4.0,\ 2.5\sqrt{\sum_{k=1}^{80}\frac{e^2(k)}{80}}\right\}$$

where the time index on $e(\cdot)$ is adjusted appropriately.

APC2

The step size and threshold are calculated once per frame as

$$\Delta = 1.897\sqrt{\sum_{k=1}^{160}\frac{e^2(k)}{160}}$$

and

$$T = \max\left\{25.8,\ 1.475\sqrt{\sum_{k=1}^{160}\frac{e^2(k)}{160}}\right\}.$$

APC3

The step size and threshold are computed once per frame from

$$\Delta = 3.33\left(\frac{\sum_{k=1}^{160}|e(k)|}{160}\right)$$

and

$$T = \max\left\{4.0,\ 1.68\sqrt{\sum_{k=1}^{160}\frac{e^2(k)}{160}}\right\}.$$

References

[1] B. S. Atal and M. R. Schroeder, "Adaptive predictive coding of speech signals," *Bell Syst. Tech. J.*, vol. 49, pp. 1973–1986, Oct. 1970.

[2] B. S. Atal and M. R. Schroeder, "Predictive coding of speech signals and subjective error criteria," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-27, pp. 247–254, June 1979.

[3] B. S. Atal, "Predictive coding of speech at low bit rates," *IEEE Trans. Commun.*, vol. COM-30, pp. 600–614, Apr. 1982.

[4] N. S. Jayant and P. Noll, *Digital Coding of Waveforms*. Englewood Cliffs, NJ: Prentice-Hall, 1984.

[5] W. D. Voiers, "Diagnostic evaluation of speech intelligibility," in *Speech Intelligibility and Recognition*, M. E. Hawley, Ed. Stroudsburg, PA: Dowden, Hutchinson, and Ross, 1977.

[6] W. D. Voiers, "Diagnostic acceptability measure for speech communication systems," in *Proc. 1977 IEEE Int. Conf. Acoust., Speech, Signal Processing*, pp. 204–207.

[7] J. D. Gibson, "Adaptive prediction in speech differential encoding systems," *Proc. IEEE*, vol. 68, pp. 488–525, Apr. 1980.

[8] R. P. Ramachandran and P. Kabal, "Stability and performance analysis of pitch filters in speech coders," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-35, pp. 937–946, July 1987.

A Curiosum Concerning Discrete Time Convolution

ERIC B. HALL AND GARY L. WISE

Abstract—**It is shown that the discrete time convolution of two absolutely summable nowhere zero sequences may be identically equal to zero.**

Development

Consider two absolutely summable sequences $a = \{a_n: n \in \mathbf{Z}\}$ and $b = \{b_n: n \in \mathbf{Z}\}$ of real numbers. Further, assume that the sequences a and b are *nowhere* zero. Does it follow that the discrete time convolution $a * b$ is nowhere zero? Does it follow that $a * b$ is nonzero on some nonempty subset of $\mathbf{Z}$? From a linear systems viewpoint, does an absolutely summable nowhere zero input to a discrete time linear time-invariant system described via discrete time convolution with a fixed absolutely summable nowhere zero sequence result in an output which is nonzero somewhere? The following development, inspired by [1, pp. 354–356], addresses these questions.

To begin, we will use the following notation. For an absolutely summable sequence of real numbers $\alpha = \{\alpha_n: n \in \mathbf{Z}\}$, let T_α map absolutely summable sequences of real numbers into absolutely summable sequences of real numbers via

$$[T_\alpha(a)]_k = \sum_{n=-\infty}^{\infty} \alpha_n a_{k-n}$$

where $a = \{a_n: n \in \mathbf{Z}\}$ is any absolutely summable sequence of real numbers. For any two absolutely summable sequences of real

Manuscript received May 25, 1989; revised July 20, 1989. This work was supported in part by the Office of Naval Research under Grant N00014-90-J-1712, and in part by the Air Force Office of Scientific Research under Grant AFOSR-86-0026.

E. B. Hall is with the Department of Electrical Engineering, Southern Methodist University, Dallas, TX 75275.

G. L. Wise is with the Department of Statistics, University of California, Berkeley, CA 94720.

IEEE Log Number 9034983.