Two-Sided Filters for Frame-Based Prediction

Sumam David and Bhaskar Ramamurthi, *Member, IEEE*

***Abstract*—A new linear prediction model, based on a two-sided predictor which predicts on the basis of past and future samples within a frame is presented. The new linear prediction model may be applied wherever frame-based prediction is employed. A stable synthesis procedure is derived by casting the prediction equation as a cyclic convolution in the time domain. When the filter order is the maximum possible, the synthesis filter is shown to have a frequency response proportional to the squared magnitude of the DFT of the frame. A symmetric two-sided predictor is described which has only half the number of coefficients to be coded as compared to a one-sided predictor of the same order. Two-sided prediction showed at least 5 dB improvement in prediction gain over one-sided prediction in our simulations on speech data. Whether this translates to coding gain will be known only after further studies with CELP-type coders.**

I. Introduction

A NEW linear prediction model, called the two-side prediction model (TSP), which predicts on the basis of past and future samples within a frame, is proposed. Since the linear prediction (LP) model [1] is often determined on a frame-by-frame basis, as in speech, all the data samples in an entire frame are available for analysis. Thus, a better estimate of a sample is obtained if we predict based on both the past and future samples.

However, synthesis in the case of TSP is not as straightforward as in the case of one-sided prediction (OSP). For example, in the autocorrelation method of determining the one-sided predictor, the predictor is guaranteed to be minimum phase [2], and hence the inverse filter for synthesis is stable. In the case of TSP, the predictor can never be minimum phase as we will see later. A synthesis procedure is derived in this paper where the problem is overcome by casting the prediction equation as a cyclic convolution rather than a linear one. This leads to a poor prediction of the samples at the two ends of the frame. Modifications which improve the end-sample predictions are also discussed.

II. Two-Sided Prediction

Throughout this paper, a frame is of length N and the order of prediction is $2p$. An estimate of s_n based on p past and p future samples is given by

$$s'_n = -\sum_{i=1}^{p} g_i s_{(n-i)\bmod N} - \sum_{i=1}^{p} g_{N-i} s_{(n+i)\bmod N}. \tag{1}$$

This is equivalent to a cyclic convolution of s_n with g_n:

$$s'_n = -\sum_{\substack{1 \le i \le p \\ N-p \le i \le N-1}} g_i s_{(n-i)\bmod N}, \qquad 0 \le n \le N-1. \tag{2}$$

Thus, the p samples at the beginning of the frame are predicted not only on the basis of samples immediately succeeding them but also those at the end of the frame. For example, the first sample is predicted based on the p samples that follow it as well as the p samples at the end of the frame. Thus the first sample is predicted only using future samples. A similar situation obtains while predicting the last p samples of the frame.

The prediction error e_n is

$$e_n = s_n - s'_n = s_n + \sum_{\substack{1 \le i \le p \\ N-p \le i \le N-1}} g_i s_{(n-i)\bmod N}. \tag{3}$$

Defining $g_0 = 1$, we can express e_n as a cyclic convolution:

$$e_n = s_n \otimes g_n, \qquad 0 \le n \le N-1. \tag{4}$$

The maximum order of prediction possible is $N - 1$. In this case, each sample is predicted based on all the other samples in the frame.

The optimum filter coefficients $\{g_i\}$ are found by minimizing the total squared error $\boldsymbol{E}_{2p}$ with respect to the parameters. Here $\boldsymbol{E}_{2p} = \Sigma_n (e_n)^2$. The optimum predictor that minimizes $\boldsymbol{E}_{2p}$ is obtained by solving the set of $2p$ symmetric normal equations

$$\Gamma_{2p} \underline{g}_{2p} = -\gamma_{2p} \tag{5}$$

where the (i, k)th element of the $2p \times 2p$ matrix Γ_{2p} is given by

$$\gamma(i, k) = \sum_n s_{(n-i)\bmod N} s_{(n-k)\bmod N}, \qquad 1 \le i, k \le p,$$
$$N - p \le i, k \le N - 1 \tag{6}$$

and $\underline{g}'_{2p} = (g_1, \cdots, g_p, g_{N-p}, \cdots g_{N-1})$. The column vector γ_{2p} is defined by $\underline{\gamma}'_{2p} = (\gamma(0, 1), \cdots, \gamma(0, p), \gamma(0, N - p), \cdots, \gamma(0, N - 1))$. The minimum mean-squared error thus obtained is

$$\boldsymbol{E}^o_{2p} = \underline{s}'\underline{s} + \underline{g}'_{2p}\underline{\gamma}_{2p}. \tag{7}$$

The range of summation in (6) is of importance. We specify two possible ranges of summation for n and obtain two methods which are akin to the well-known autocorrelation and covariance methods in one-sided prediction. Following tradition, we dub the first as the autocorrelation method (TSP-A), wherein the range $0 \le n \le N - 1$ is used. In the covariance method (TSP-C), the range used is $p \le n \le N - p - 1$.

In the case of TSP-A, $\gamma(i, k) = r(i, k)$, the cyclic autocorrelation function, defined as

$$r(i, k) = \sum_{n=0}^{N-1} s_{(n-i)\bmod N} s_{(n-k)\bmod N}. \tag{8}$$

Correspondingly, we denote Γ_{2p} in this case by $\boldsymbol{R}_{2p}$ and $\underline{\gamma}_{2p}$ by $\underline{r}_{2p}$. It can be seen that $r(i, k) = r(k, i) = r(i - k)$. Further, the cyclic autocorrelation function is an even function, i.e., $r(i)$

Manuscript received January 12, 1989; revised May 7, 1990.
The authors are with the Department of Electrical Engineering, Indian Institute of Technology, Madras 600036 India.
IEEE Log Number 9042259.

1053-587X/91/0400-0789$01.00 © 1991 IEEE

$= r(-i) = r(N - i)$. Using this property, it is easily seen that the existence of a unique solution to the normal equations implies that $g_i = g_{N-i}$, i.e., the predictor coefficients are symmetric. The $2p$ normal equations then simplify to a set of p symmetric normal equations which are given by

$$(\boldsymbol{R}_p^{(1)} + \boldsymbol{R}_p^{(2)})\underline{g}_p = -\underline{r}_p \tag{9}$$

where $\boldsymbol{R}_p^{(1)}$ and $\boldsymbol{R}_p^{(2)}$ are $p \times p$ matrices whose elements are given, respectively, by $r_1(i, k) = r(i - k)$ and $r_2(i, k) = r(i + k)$, $1 \leq i, k \leq p$, $\underline{g}_p^t = (g_1, \cdots, g_p)$, and $r_p^t = (r(0, 1), \cdots, r(0, p))$. The minimum mean-squared error obtained is

$$E_{2p}^o = \underline{s}^t\underline{s} + 2\underline{g}_p^t\underline{r}_p. \tag{10}$$

Thus, for an analysis filter of length $2p$ we need solve only p linear symmetric equations. A fast algorithm for the solution of this system of equations is given in [3].

In the covariance method, TSP-C, $\gamma(i, k)$ is replaced by $\phi(i, k)$ where $\phi(i, k)$ has the range of summation $p \leq n \leq N - p - 1$. The element $\phi(i, k)$ is symmetric but not periodic even, i.e., $\phi(i, k) \neq \phi(N - i, N - k)$. Further, $\phi(i, k) \neq \phi(i - k)$. Therefore in TSP-C, we must solve $2p$ linear symmetric equations to obtain the predictor coefficients, which are asymmetric.

A. Discussion

The prediction of the first and the last p samples in the frame involve samples which are at the opposite ends of the frame, due to the end-around nature of the cyclic convolution. These predictions are reflected in the normal equations by the presence of the cyclic autocorrelation function $r(i)$, in contrast to the linear autocorrelation function found in the one-sided case. The expression for $r(i)$ is given by $\Sigma_{n=0}^{N-1} s_n s_{(n-i)\bmod N}$ where we see that, apart from products of the type $s_n s_{n-i}$, $i \leq n \leq N - 1$, products such as $s_n s_{N+n-i}$ involving the end samples are also involved. If N is large and i is small, these "spurious" correlations between end samples do not affect the estimate of $r(i)$ significantly. Thus the predictor is essentially optimized for predicting all except the first and last p samples.

In the covariance method, the filter coefficients $\{g_i\}$ are such as to minimize the error in the range $p \leq n \leq N - p - 1$ and hence no end-around effects are present. In this method, overlapped data frames with p samples overlap at each end are formed. Since the error is minimized only for the samples of interest, namely, $p \leq n \leq N - p - 1$, this method gives the best prediction gain. However, its computational complexity is the highest, as we have to solve twice the number of equations. Further, in a coding application, only half the number of predictor coefficients need be coded when the autocorrelation method is used. Due to these inherent advantages in the autocorrelation method, it is of interest to modify the autocorrelation method to reduce the effects of poor prediction of the end samples.

One approach is to taper the frames at the ends and overlap them (TSP-AT). For a predictor order of $2p$, the p data samples at both the frame ends are tapered and the frames overlapped. The tapered sequence is given by

$$s_n^t = s_n w_m, \qquad 0 \leq n \leq N - 1 \tag{11}$$

where

$$w_n = \begin{cases} n/(p-1), & \text{for } 0 \leq n \leq p - 1 \\ 1, & \text{for } p \leq n \leq N - p - 1 \\ (N - n - 1)/(p - 1), & \text{for } N - p \leq n \leq N - 1. \end{cases}$$

This is equivalent to segmenting the signal using an overlapping trapezoidal window. Tapering decreases the effect of the spurious end-around correlations on the estimates of $r(i)$. The prediction of the samples improves as one moves inwards from the ends of the frame. Further, when the tapered portions of the predicted frames are overlapped and added, the poorly predicted samples are weighted less than the better-predicted ones.

III. Synthesis Model

If the analysis and the synthesis are in the form of a linear convolution, and the analysis filter is symmetric, the causal inverse filter is always unstable. This is because a symmetric two-sided sequence will always have zeros outside the unit circle. However, by casting the analysis in the form of a cyclic convolution, we are able to obtain an "inverse filter" for synthesis in the two-sided case, also in the form of cyclic convolution.

The synthesis model is described as the cyclic convolution between the synthesis filter h_n, and the prediction residual e_n, to obtain the original frame s_n

$$s_n = e_n \otimes h_n. \tag{12}$$

In the frequency domain the synthesis can be expressed as

$$S(k) = E(k)H(k), \qquad 0 \leq k \leq N - 1 \tag{13}$$

where $S(k)$, $E(k)$, and $H(k)$ are the N-point DFT's of s_n, e_n, and h_n, respectively. The analysis model of (4) can be expressed in the frequency domain as

$$E(k) = S(k)G(k), \qquad 0 \leq k \leq N - 1 \tag{14}$$

where $G(k)$ is the N-point DFT of g_n. From (13) and (14), assuming that $G(k) \neq 0$ for any k, we get $H(k) = 1/G(k)$. In the autocorrelation method of solution, g_n is symmetric and real which implies that $G(k)$, $H(k)$, and h_n are all symmetric and real.

A. Existence of Synthesis Filter

In the Appendix, we obtain the following results for TSP-A and TSP-AT that have a bearing on when $G(k)$ becomes zero for some k. If any $S(k)$ is zero, an optimum two-sided predictor can be found which predicts s_n exactly. Further, if $2p = N - 1$, (for this case N must be odd) and $S(k) \neq 0$ for all k, then $G(k)$ for the optimum predictor is nonzero for all k. When $2p < N - 1$, it is possible for $G(k)$ to be equal to zero for some k. In such a case, one solution to the synthesis problem is to vary the order of prediction until $G(k) \neq 0$ for all k. Alternatively, if $G(k_0) = G(N - k_0) = 0$, set $G(k_0) = G(N - k_0) = \epsilon$, where ϵ is a small nonzero value. Each such replacement changes the mean-squared error by a small amount $2/N\, \epsilon^2 |S(k_0)|^2$. In all the frames analysed in our simulations using the different methods, for $2p$ varying between 2 and 16, we never encountered a situation where $G(k)$ was zero. The lowest value of $G(k)$ observed was of the order of 10^{-8}.

B. Relationship to DFT of the Frame

For the full-order predictor ($2p = N - 1$), from (A2) we get

$$H(k) = 1/G(k) = \text{const.}\ |S(k)|^2. \tag{15}$$

This implies that the synthesis filter has a frequency response proportional to the squared-magnitude of the DFT of the frame. This has been verified in our simulations on speech data. If this synthesis filter is excited by a frame consisting of a unit sample (hence $U(k) = 1$, $\forall k$), the synthesized output is proportional to the cyclic autocorrelation of the input process.

The infinite-order predictor in OSP, $A(z)$, is proportional to the whitening filter for the frame. The squared magnitude of the frequency response of the corresponding synthesis filter $1/A(z)$ is proportional to the power spectral density (psd) of the input process. Thus when this synthesis filter is excited by white noise (psd = 1), the psd of the output is proportional to the psd of the original process. The results in TSP are similar, except for the fact that in TSP the filter is of finite order $N - 1$ while in OSP it is of infinite order. Moreover, in TSP we are dealing with a deterministic frame, whereas in OSP, the theory is developed for a random process.

IV. Simulation Results

The proposed two-sided predictor was tested on speech data. Speech sampled at 8 kHz, segmented into frames of 128 samples ($N = 128$), was used in both TSP and OSP, with order of prediction varying in the range 1 to 16. Four sentences each of male and female voices constituted the data set.

A. Order of Predictor

The order of prediction required was decided by its value at which the average normalized mean-squared prediction error in decibels, $V_{2p}^{T} = 1/M \sum_{j=1}^{M} 10 \log (E_{2p}^{o}/r(0))_j$, (averaged over the entire data set of M frames), flattens off. Fig. 1 shows the variation of V_{2p}^{T} with order of prediction $2p$, for two-sided prediction using the TSP-A, TSP-AT, and TSP-C. It can be seen from Fig. 1 that the normalized mean-squared prediction error flattens off for $2p > 8$ in all the methods. We emphasize again that an eighth-order symmetric TSP has only 4 arbitrary coefficients compared to 8 in OSP.

In OSP, an order of prediction between 8 and 10 is considered to be sufficient for speech sampled at 8 kHz [1]. It was found that in over 98% of the frames, TSP of order 8 performed better than OSP of order 8 in all the above methods.

It can be observed from Fig. 1 that the covariance method achieves the highest prediction gain irrespective of predictor order. However, more coefficients have to be coded in this case and the computational complexity is higher than with the autocorrelation method. TSP-AT performs better than TSP-A as the effect of the end-around predictions decreases with tapering. The difference in prediction gain between TSP-AT and TSP-C was around 0.25–2 dB (1 dB in the case of the eighth-order predictor). Thus we can see that TSP-AT is a good choice as a compromise between prediction gain and computational complexity.

B. Comparison of Mean-Squared Prediction Error in TSP and OSP

Two variants were compared for the autocorrelation method: a) TSP-A, OSP-A where the speech is segmented into frames with 128 data samples and b) TSP-AT, OSP-AT where the data samples are windowed by the 6.25% overlapping trapezoidal window. In the covariance method the mean-squared prediction error is minimized over 112 data samples in both TSP-C and OSP-C, for predictor orders between 1 and 16. The frames overlap to the extent of 8 samples at either end, giving a total frame size of 128.

Let us define α_{2p} as the difference in decibels between the average normalized mean-squared prediction errors in TSP and OSP, i.e., $\alpha_{2p} = V_{2p}^{T} - V_{2p}^{O}$ (dB), ($V_{2p}^{O} = 1/M \sum_{j=1}^{M} 10 \log (D_{2p}^{o}/r(0))_j$, D_{2p}^{o} is the minimum prediction error in OSP), α_{2p} is a measure of the relative performance of the two prediction models. In Fig. 2, the variation of α_{2p} with predictor order is plotted for the two autocorrelation methods and the covariance method. It can be seen from Fig. 2 that an eighth-order TSP gave more than 5 dB improvement compared to OSP in all the methods. Both TSP-AT and TSP-C showed almost the same degree of improvement in prediction gain over OSP-AT and OSP-C, respectively. The worst relative performance is seen in the case of TSP-A and this is due to the end-around predictor which are more pronounced here than in TSP-AT.

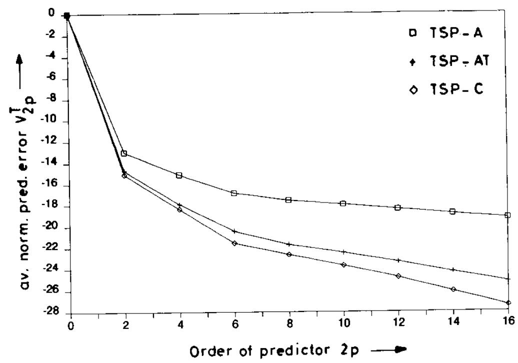

Fig. 1. Variation of average normalized mean-squared prediction error V_{2p}^{TSP} with order of prediction $2p$, for two-sided prediction in autocorrelation and covariance methods.

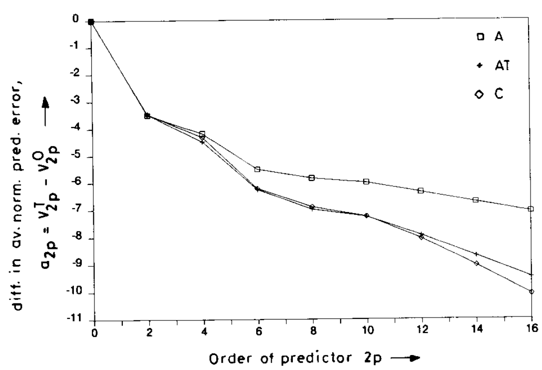

Fig. 2. Variation of the difference (decibels) in average normalized mean-squared prediction error between TSP and OSP, α_{2p}, in autocorrelation and covariance methods.

In the above performance comparisons, the errors in OSP in the range N to $N + 2p - 1$ were also included in the evaluation. When the same comparisons were made taking into account only the errors in the data samples a similar performance was seen.

The percentage of frames in which TSP performed better than OSP as indicated by α_{2p} was always above 87%, and the percentage of such frames was 100% in our simulations when the predictor order was greater than 12.

The proposed TSP model is seen to predict much better than the conventional OSP model as can be seen from Fig. 2. An eighth-order TSP is sufficient for speech modelling, providing a prediction gain of around 5 dB over OSP and it requires only half the number of filter coefficients to be coded when compared to eighth-order OSP.

It is instructive to compare the prediction residuals in the case of TSP and OSP filters. The tapered autocorrelation method was

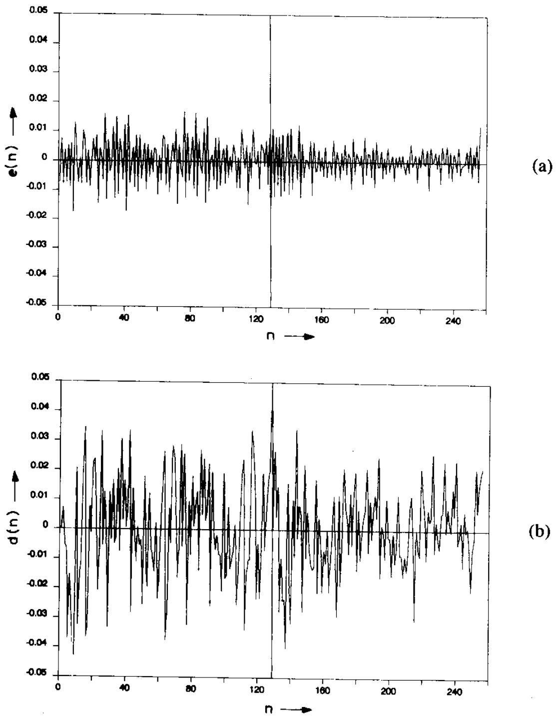

Fig. 3. Prediction residual of the eighth-order filter (a) TSP and (b) OSP, for the speech frames of Fig. 5(a).

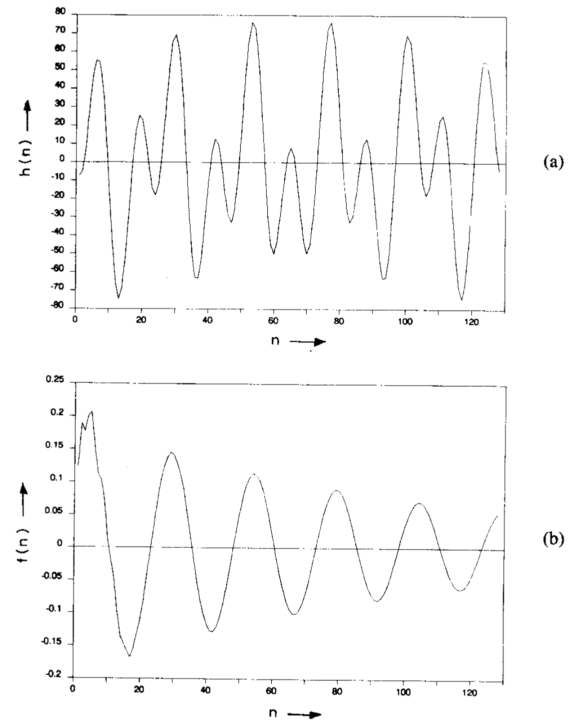

Fig. 4. Impulse response of the eighth-order synthesis filter (a) TSP and (b) OSP, for the first frame of Fig. 5(a).

used to find the prediction filter for the two adjacent frames of speech shown in Fig. 5(a). The residuals obtained for eighth-order TSP (Fig. 3(a)) and OSP (Fig. 3(b)) filters for these two frames are shown in Fig. 3. The TSP residual does not show any characteristics significantly different from that of OSP residual except that it has lower energy.

The synthesis filter obtained in the autocorrelation method is symmetric. The synthesis-filter's impulse response in the one-sided case typically decays rapidly. In contrast, the synthesis filter response in TSP typically has significant values over the whole frame. The unit-sample responses of the eighth-order TSP and OSP synthesis filters for the first frame in Fig. 5(a) are shown in Figs. 4(a) and (b), respectively. The magnitude-spectrum of these synthesis filters are also shown in Fig. 5(b) along with the magnitude spectrum of the first speech frame. A noticeable feature of the TSP spectrum is that it is not accurate for the higher frequencies where the speech-spectral level is very low.

V. Symmetric Filters in Speech Synthesis

The prediction residual, when used as the excitation with the appropriate synthesis filter, gives as output the original speech itself without any distortion in both TSP and OSP. To achieve low bit rates, the excitation is coded, as in CELP [4] or multipulse excited LPC (MPELPC) [5]. *It is not evident whether in such schemes, if OSP is replaced with TSP and the coder is suitably optimized, if the remarkably high improvement in prediction gain with TSP will directly translate to a high coding gain.*

A conclusive test of the applicability of TSP in speech coding would be its performance in CELP coders. This is, however,

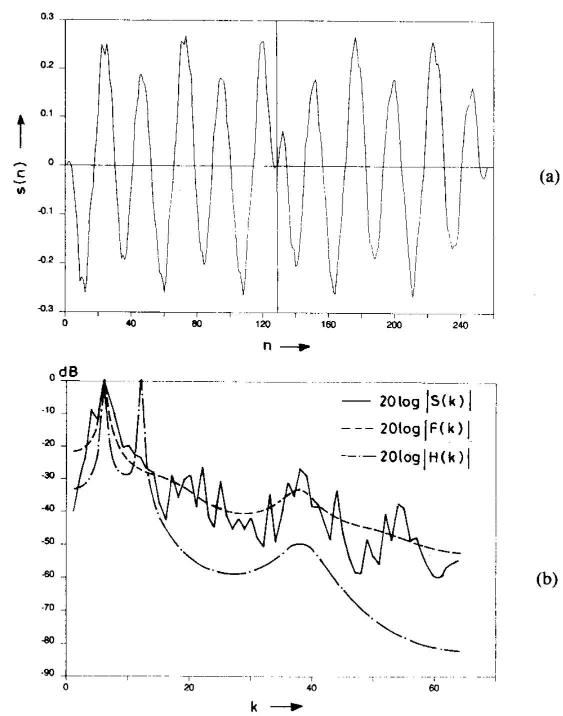

Fig. 5. (a) Original speech waveforms for two adjacent frames of 16 ms duration each. (b) Magnitude spectrum of the eighth-order TSP and OSP synthesis filters shown in Fig. 4 and magnitude spectrum of the first speech frame of Fig. 5(a).

not straightforward as CELP has to be modeled using cyclic convolutions. It is easy to implement TSP synthesis using DFT. A multiband excited vocoder [6] where the short-term spectrum of speech is modeled as a product of the speech spectral envelope and an excitation spectrum in the frequency domain is currently being studied. It employs a speech production model based on the TSP filter to reconstruct the spectral envelope.

Another approach is to use variable-sized pitch synchronous frames. In such a case, samples one pitch-period away can be used along with the adjacent samples for two-sided prediction. Further, the problem of end-around predictions is eliminated here. However, variable frame-size coders usually imply variable bit-rate and variable coding delay.

A quick, but less informative, alternative to study the performance of TSP in speech coders is the MPELPC scheme (without the long-term predictor). The OSP filter was replaced by the TSP filter, and the analysis-by-synthesis procedure modified slightly to find the excitation. The tapered autocorrelation method was employed to find the prediction filter as it gives good performance with low computational complexity.

From the results obtained, it was clear that no catastrophic problems were associated with TSP synthesis when the excitation differed from the residual. Further, there was a significant improvement of nearly 3 dB in segmental SNR when TSP was used instead of OSP in about 50% of the frames. The cyclic nature of the synthesis in the two-sided case, and the fact that the unit-sample response of the TSP synthesis filter has significant values over the entire frame, make the sequential pulse optimization technique of [5] less effective in the case of TSP. Any one excitation pulse significantly affects all the synthesized samples in the frame. Therefore, it is difficult to interpret whether the poor performance in the remaining frames is due to bad choice of excitation or to the synthesis filter.

TSP can be profitably employed in MPELPC coders in the following manner. Both TSP and OSP synthesis filters can be used for each frame and the one that performs better can be used. This improves the overall performance of the MPELPC scheme while at the same time keeping the bit rate unchanged. The type of the predictor can easily be indicated to the receiver by the addition of an extra bit while encoding.

In the autocorrelation method for TSP, the filter is symmetric. A lattice-form realization, with the inherently low sensitivity of the reflection coefficients, can be found for the causal half of the symmetric response, which is sufficient to recover the two-sided response at the decoder.

Appendix

We present here some results pertaining to the characteristics of the TSP analysis filter when the autocorrelation method is used.

A. Result 1

If any $S(k)$ is zero, then a two-sided predictor can be found that predicts exactly.

The mean-squared prediction error is $E_{2p} = 1/N \, \Sigma_{k=0}^{N-1} |S(k)|^2 \, G^2(k)$. Therefore, if $S(k_0) = S(N - k_0) = 0$, then with $G(k_0) = G(N - k_0) = N/2$, the signal can be completely predicted giving zero prediction error. Note that this choice of $G(k)$ satisfies the constraint $g_0 = 1$.

B. Result 2

When the predictor order is $N - 1$, if $S(k) \neq 0$ for all k, then $G(k)$ of the optimum predictor is nonzero for all k.

In the frequency domain $E(k) = G(k)S(k)$. Since g_n is symmetric, and $G(k)$ is real, the mean-squared prediction error $E_{2p} = 1/N \, \Sigma_{k=0}^{N-1} |S(k)|^2 \, G^2(k)$. With two-sided predictor order equal to $N - 1$, we can obtain the optimum predictor $\{G(k)\}$ by minimizing E_{N-1} with respect to $G(i)$, $0 \leq i \leq N - 1$, subject to the constraint $g_0 = 1/N \, \Sigma_{k=0}^{N-1} G(k) = 1$. The Lagrange auxiliary function to be minimized is therefore

$$f(G) = \frac{1}{N} \sum_{k=0}^{N-1} |S(k)|^2 G^2(k) - \lambda \left(\frac{1}{N} \sum_{k=0}^{N-1} G(k) - 1 \right). \tag{A1}$$

Thus we obtain

$$G(k) = 1 \Big/ \frac{1}{N} \left[\sum_{i=0}^{N-1} \left(|S(k)|^2 / |S(i)|^2 \right) \right] \tag{A2}$$

and

$$E_{N-1}^o = 1 \Big/ \left(\frac{1}{N} \sum_{k=0}^{N-1} 1/|S(k)|^2 \right). \tag{A3}$$

Thus when the predictor order is $N - 1$, i.e., if the prediction is based on all the remaining samples in the frame, and if $S(k) \neq 0$ for all k, then $G(k)$ of the optimum predictor is nonzero for all k. Further, the optimum value of E_{N-1}^o is also nonzero. This implies that any lower order predictor will also have nonzero prediction error. In the one-sided case, too, the mean-squared prediction error cannot become zero if $S(\omega) \neq 0$ for all ω.

C. Result 3

For a TSP filter of order $2p$, where $2p < N - 1$, $G(k)$ may become zero for some value of k even if $S(k) \neq 0$ for all k (implying that the minimum value of mean-squared prediction error is nonzero).

As an example, consider the 5-sample frame $s(0) = 0.78987$, $s(1) = 0.38987$, $s(2) = 0.18987$, $s(3) = 0.18987$, and $s(4) = 0.38987$. In this case, $S(0) = 1.949$, $S(1) = S(4) = 0.7236$, and $S(2) = S(3) = 0.2763$. The cyclic autocorrelation is given by $r(0) = 1$, $r(1) = r(4) = 0.8$ and $r(2) = r(3) = 0.6$. With a TSP of order 2, $g_1 = r(1)/r(0) + r(2) = -0.5$, and $E_2^o > 0$ but $G(0) = 0$. This is in contrast to the one-sided case, where the zeros of the prediction filter will lie on the unit circle if and only if the mean-squared prediction error goes to zero [1].

References

[1] J. Makhoul, "Linear prediction: A tutorial review," *Proc. IEEE*, vol. 63, no. 4, pp. 561-580, 1975.

[2] S. M. Kay and S. L. Marple, "Spectral analysis: A modern perspective," *Proc. IEEE*, vol. 69, pp. 1380-1419, Nov. 1981.

[3] G. Carayannis, N. Kalouptsidis, and D. G. Manolakis, "Fast recursive algorithms for a class of linear equations," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-30, no. 2, pp. 227-239, Apr. 1982.

[4] B. S. Atal and M. R. Schroeder, "CELP: High quality speech at very low bit rates," in *Proc. ICASSP*, 1985, pp. 937-940.

[5] B. S. Atal and J. R. Remde, "A new model of lpc excitation for producing natural sounding speech at low bit rates," in *Proc. ICASSP*, 1982, pp. 614-617.

[6] D. W. Griffin and J. S. Lim, "Multiband excitation vocoder," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-36, no. 8, pp. 1223-1235, Aug. 1985.

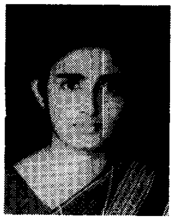

Sumam David was born in Quilon, India, on December 11, 1962. She received the B.Tech. degree in electronics and communications engineering from the University of Kerala, India, in 1985, and the M.Tech. degree in electrical engineering, from the Indian Institute of Technology (I.I.T.), Madras, India, in 1986. She is currently working toward the Ph.D. degree at I.I.T., Madras, India. Her main research interests are in the fields of digital signal processing and speech coding.

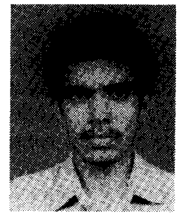

Bhaskar Ramamurthi (S'81–M'84) received the B.Tech. degree in electronics in 1980 from the Indian Institute of Technology (I.I.T.), Madras, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Santa Barbara, in 1982 and 1985, respectively.

He was with AT&T Bell Laboratories, Holmdel, NJ, from January 1985 until April 1986. Since June 1986 he has been with the Department of Electrical Engineering at I.I.T., Madras, India, as Assistant Professor. His research interests are in the areas of source coding, communication theory, signal processing, and packet communications.