

AD-A096 091

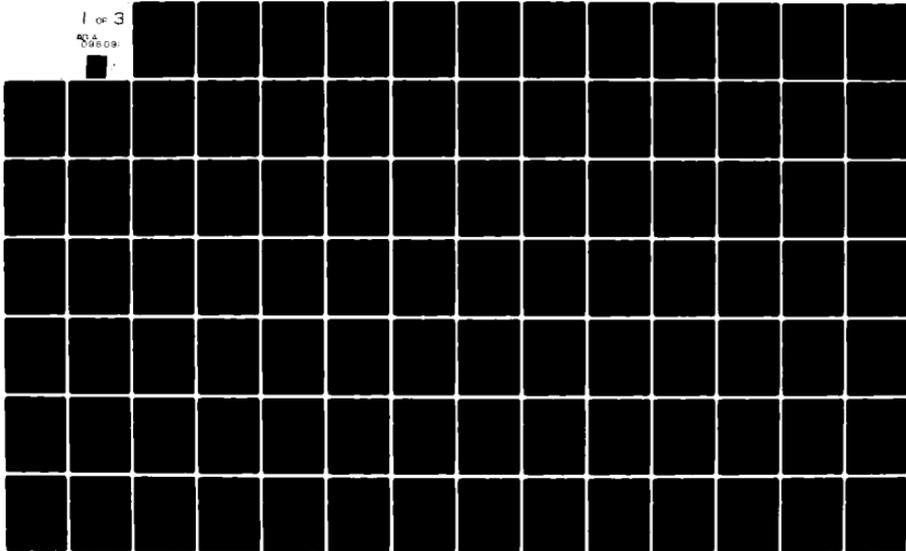
BOLT BERANEK AND NEWMAN INC CAMBRIDGE MA
DESIGN AND REAL-TIME IMPLEMENTATION OF A ROBUST APC CODER FOR S--ETC(U)
DEC 80 R VISWANATHAN, W RUSSELL, A HIGGINS

F/8 17/2.1
DCA100-79-C-0037
NL

UNCLASSIFIED

1 of 3

AD-A096 091



Bolt Beranek and Newman Inc.



AD A 096091

12

Report No. 4565

10/11/80

Design and Real-Time Implementation of a Robust APC Coder for Speech Transmission Over 16 Kb/s Noisy Channels.

Volume 1 - Algorithm Design and Simulation.

Final Report. Jul 1980

December 1980

DTIC
MAR 9 1981

Prepared for:
Defense Communications Agency

DISTRIBUTION STATEMENT A
Approved for public release; distribution is unlimited.

DDC FILE COPY

81 3 9 171

Report No. 4565

DESIGN AND REAL-TIME IMPLEMENTATION OF A ROBUST APC CODER
FOR SPEECH TRANSMISSION OVER 16 Kb/s NOISY CHANNELS

Final Report

Volume I: Algorithm Design and Simulation

Authors: R. Viswanathan, W. Russell, A. Higgins

December 1980

Prepared for:
Defense Communications Agency

Unclassified.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 4565, Vol. I	2. GOVT ACCESSION NO. AD-A09609	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) DESIGN AND REAL-TIME IMPLEMENTATION OF A ROBUST APC CODER FOR SPEECH TRANSMISSION OVER 16 Kb/s NOISY CHANNELS, Vol. I: Algorithm Design and Simulation	5. TYPE OF REPORT & PERIOD COVERED Final Report July 1979 - December 1980	
	6. PERFORMING ORG. REPORT NUMBER BBN Report No. 4565	
7. AUTHOR(s) R. Viswanathan, W. Russell, and A. Higgins	8. CONTRACT OR GRANT NUMBER(s) DCA100-79-C-0037	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, MA 02238	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Agency Contract Management Division, Code 260 Washington, D.C. 20305	12. REPORT DATE December 1980	
	13. NUMBER OF PAGES 285	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report) Unclassified	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce, for sale to the general public.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) speech coding, adaptive predictive coding, mediuiband speech coders, 16 kb/s speech transmission, noise spectral shaping, limit-cycle behavior, robust speech coder, digital voice terminal, real-time speech coder.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes the design and development of a real-time robust APC speech coder that transmits high-quality speech over a 16 kb/s channel with bit-error rates of up to 1%. This volume (Vol. I) describes the development, experimental investigation, and speech-quality optimization of 16 kb/s adaptive predictive coders. Important among the various topics treated in this report are: extensive investigation of several existing and new methods for coding the APC residual, adaptive spectral shaping of the		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

cont'd.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

error-free

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. Abstract cont'd.

cont → quantization noise, causes for and effective remedy of the limit-cycle behavior of the quantizer output, different APC configurations, optimization of APC coders for error-free channels and for (errorful) channels, tradeoff between voice data rate and error-protection data rate, tandeming of the optimized 16 kb/s APC coder with the 2.4 kb/s LPC-10 coder, and coder performance in noisy backgrounds typical in offices and in air-borne command post environments. Most impressive of the numerous interesting and useful results described in this report is that the final optimized, robust coder design yields speech quality that degrades only slightly as the bit-error rate is increased from 0% to 1%.

↖

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

ACKNOWLEDGMENTS

The work by J. Makhoul and M. Berouti on entropy-coded APC system with noise shaping served as the starting point for this project. The authors would like to thank G. Moran, the COTR of this project, for his interest and interaction on technical issues during the course of this project. The author (RV) would also like to thank B. Atal for valuable discussions on some aspects of his work on adaptive predictive coding of speech.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
Ev.	
Distribution/	
Availability Codes	
Dist	Avail and/or Special

TABLE OF CONTENTS

	Page
1. INTRODUCTION	1
1.1 Coder Design Requirements	1
1.2 Summary of the Optimized Algorithm	3
1.3 Overview of the Report	4
2. CHOICE OF 16 KB/S CODER	7
2.1 Rationale for Choosing APC	7
2.2 Sampling Rate of Input Speech	8
2.3 Data Bases	10
3. REVIEW OF THE APC SYSTEM	11
3.1 Basic APC system.	11
3.2 Forward-Adaptive Residual Quantization	15
3.3 Noise Spectral Shaping	17
3.4 Signal-to-Noise Ratio Considerations	17
4. DIFFERENT APC CONFIGURATIONS	19

4.1	Sequencing of Spectral and Pitch Predictors	19
4.2	Prediction-Feedback Configuration	21
4.3	Noise-Feedback Configuration	23
4.4	Hybrid-Feedback Configuration	25
4.5	Comparison of the APC Configurations	27
5.	FEEDBACK GAIN OF APC	28
6.	QUANTIZATION OF APC PARAMETERS	31
6.1	Spectral Parameters	31
6.2	Pitch Parameters	33
6.3	Residual Quantizer Gain	34
7.	ADAPTIVE NOISE SHAPING METHODS	35
7.1	All-Zero Noise Shaping	35
7.2	All-Pole Noise Shaping	38
7.3	Pole-Zero Noise Shaping	40
7.4	Effect of Preemphasis on Noise Shaping	41
7.5	Comparative Evaluation and Experimental Results	43
8.	METHODS FOR CODING THE APC RESIDUAL	46
8.1	Chapter Overview	46
8.2	Pitch Prediction	48
8.2.1	The Method	48
8.2.2	Stability of Multi-tap Pitch Predictor	50
8.2.3	Experimental Results	52
8.3	Segmented Quantization	54
8.3.1	The Method	54

8.3.2	Pitch Prediction with Segmented Quantization	55
8.3.3	Experimental Results	57
8.4	Entropy Coding	60
8.4.1	The Method	60
8.4.2	Variable-to-Fixed Rate Conversion	61
8.4.3	Experimental Results	63
8.5	Pitch-Adaptive Coding	65
8.5.1	Itakura's Method	66
8.5.2	Basic Pitch-Adaptive Method	67
8.5.3	Pitch Prediction and Pitch-Synchronous Segmented Quantization	74
8.6	Segmented Quantization with Bit Allocation	76
8.6.1	The Method	76
8.6.2	Experimental Results	82
9.	LIMIT-CYCLE BEHAVIOR OF THE APC SYSTEM	85
9.1	Discussion of the Problem	85
9.2	S-P versus P-S Prediction Sequence	87
9.3	Saturating Limiter	88
9.4	High-Frequency Correction	91
9.5	Preemphasis	93
10.	OPTIMIZATION OF CODERS FOR ERROR-FREE CHANNELS	96
10.1	Entropy Coding with Pitch Prediction	98
10.2	Pitch Prediction and Segmented Quantization	100
10.3	Pitch-Adaptive Coding With Pitch Prediction and Segmented Quantization	100
10.4	Segmented Quantization With Bit Allocation and Pitch Prediction	101
10.5	The S-P Prediction Sequence	102
10.6	Comparative Evaluation and Recommendations	103
11.	OPTIMIZATION OF ERROR-PROTECTED CODERS	105

11.1	Entropy Coding with Pitch Prediction	106
11.2	Pitch Prediction and Segmented Quantization	107
11.3	Segmented Quantization With Bit Allocation and Pitch Prediction	110
11.4	Comparative Evaluation	110
12.	EVALUATION OF ERROR-PROTECTED CODERS IN 1 $\frac{1}{2}$ CHANNEL ERROR	112
12.1	Channel-Error Simulation	113
12.2	Entropy Coding with Pitch Prediction	114
12.3	Pitch Prediction and Segmented Quantization	117
12.4	Segmented Quantization with Bit Allocation and Pitch Prediction	119
12.5	Effect of Prediction Sequence on Channel-Error Performance	120
12.5.1	The S-P Prediction Sequence	120
12.5.2	Receiver Smoothing of the Decoded Residual	121
12.6	Comparative Evaluation and Recommendations	123
12.7	Folded Binary Code	125
12.8	Performance Over Higher Error-Rate Channels	126
13.	ACOUSTIC BACKGROUND NOISE	129
14.	TANDEMING WITH LPC-10	130
14.1	Simulation of LPC-10	130
14.2	LPC-APC Tandem	131
14.3	APC-LPC Tandem	133
15.	16 KB/S BASEBAND CODERS	135

16.	OPTIMIZED, ROBUST 16 KB/S APC CODER	137
16.1	Simplifications for Real-Time Implementation	137
16.1.1	High-Frequency Correction	137
16.1.2	Pitch-Filter Stability Test	138
16.1.3	Noise Shaping	140
16.2	Refinements to the Coder	141
16.2.1	Laplacian versus Gaussian Quantizer	141
16.2.2	Recomputing Parameter Quantization Tables	143
16.3	Optimized Coder Description	144
16.4	FORTTRAN Simulation of the Optimized Coder	151
16.5	Results of Tests on the Real-Time Coder	152
17.	SUMMARY AND MAJOR CONTRIBUTIONS	156
18.	REFERENCES	160
19.	APPENDIX A: SPECIFICATION OF THE OPTIMIZED 16 KB/S APC ALGORITHM.	163
20.	APPENDIX B: USER'S GUIDE FOR THE FORTTRAN SIMULATION OF THE BBN 16 KB/S APC CODER.	211
21.	APPENDIX C: A LISTING OF THE SOURCE PROGRAMS OF THE FORTTRAN SIMULATION OF THE 16 KB/S APC CODER.	236

Report No. 4565

Bolt Beranek and Newman Inc.

LIST OF FIGURES

FIG. 1.	Basic APC System	12
FIG. 2.	Forward-adaptive quantizer for the APC residual.	16
FIG. 3.	Two methods of sequencing the spectral and the pitch predictors.	20
FIG. 4.	The prediction-feedback configuration of APC.	22
FIG. 5.	The noise-feedback configuration of APC.	24
FIG. 6.	The hybrid-feedback configuration of APC.	26
FIG. 7.	Implementation of the all-pole noise shaping with $B(z) = 1/A(z/\alpha)$	39
FIG. 8.	Implementation of the pole-zero noise shaping with $B(z) = A(z/\alpha)/A(z)$.	42
FIG. 9.	Comparison of noise shaping methods.	44
FIG. 10.	S/Q ratios and relative speech-quality judgments for 6 PP-SQ APC coders.	59
FIG. 11.	Bit assignment in the basic pitch-adaptive coding scheme.	68
FIG. 12.	Segmented quantization with bit allocation.	78
FIG. 13.	A step-by-step description of the bit-allocation algorithm used in the SQ-BA method.	80
FIG. 14.	Illustration of the difference between natural binary code (NBC) and folded binary code (FBC), for a 4-level Laplacian nonuniform quantizer.	127
FIG. 15.	Tandem operation of APC and LPC-10 coders.	132
FIG. 16.	Optimum Laplacian and Gaussian quantizers.	143
FIG. 17a	The transmitter of the optimized 16 kb/s APC system.	146
FIG. 17b	The receiver of the optimized 16 kb/s APC coder.	147

Report No. 4565

Bolt Beranek and Newman Inc.

LIST OF TABLES

TABLE	1.	Quantization of the LARs of an APC system that uses the S-P prediction sequence and no preemphasis.	32
TABLE	2.	Quantization of the LARs of an APC system that uses the P-S prediction sequence and no preemphasis.	32
TABLE	3.	Quantization ranges for the pitch taps.	34
TABLE	4.	Nonuniform quantization of delta gains.	57
TABLE	5.	Bit allocation for various coder parameters.	98
TABLE	6.	S/Q ratios for the 9 entropy-coding APC systems, all using 10zero noise shaping	99
TABLE	7.	16 kb/s error-protected entropy-coding systems.	107
TABLE	8.	16 kb/s error-protected PP-SQ systems.	109
TABLE	9.	Error-protection allocations used for four entropy-coding systems.	116
TABLE	10.	Error-protection allocations used for two PP-SQ systems and one PP-SQ-BA system.	118
TABLE	11.	Revised nonuniform quantization of delta gains.	144
TABLE	12.	Quantization and error protection of parameter data for the optimized 16 kb/s APC system.	148

1. INTRODUCTION

The goal of this project was the design and development of a real-time speech coding system that produces high quality speech at a data rate of 16 kb/s (kilobits/second). The final report of this project is organized in two volumes. Volume I, which is this report, describes our work on the development and optimization of the speech coding algorithm. Volume II deals with the implementation of the final optimized speech coding algorithm as a real-time full duplex system on a CSP Inc. MAP-300 signal processing computer and associated hardware.

In this chapter, we state the design requirements on the speech coder performance (Section 1.1), describe briefly the optimized coder (Section 1.2), and provide an overview of the rest of this report (Section 1.3).

1.1 Coder Design Requirements

The input speech of the coder should have a bandwidth of at least 3.2 kHz. The encoder and decoder of the speech coder should operate independently, with the encoder mapping the analog input signal into an output binary sequence and the decoder mapping the binary sequence into the corresponding analog output

speech. In addition to the requirement that the speech coder in general produce speech of very good quality in the sense that it has a very high degree of user acceptance, there are several specific requirements on the coder performance as given below:

1. Noisy channel: Produce good quality speech under conditions of a transmission bit error rate of up to 1%.
2. Acoustic background noise: Produce toll quality speech under conditions of acoustic background noise such as office noise with a sound pressure level (SPL) of 60 dB re 20 micronewtons per square meter, and good quality speech under 100 dB of acoustic background noise such as in Air-Borne Command Post (ABCP) environment.
3. Tandem operation with LPC-10 coder: Perform satisfactorily in tandem (in both directions) with an LPC-10 speech coder operating at a data rate of 2.4 kb/s. The tandem link should provide speech intelligibility with minimal degradation compared with a single link of the 2.4 kb/s LPC-10 coder alone.

Other objectives of this work have included:

1. Minimize the computational complexity of the speech coding algorithm.

2. Identify and explain the features in the optimized 16 kb/s coder and in the 2.4 kb/s LPC-10 coder that control the quality of the tandem link between the two coders, and indicate how the tandem performance could be further improved.

3. Extend the design of a previously developed 9.6 kb/s baseband LPC coder [3] to the 16 kb/s data rate, and compare the output speech quality of the resulting baseband coder with that of the optimized 16 kb/s coder.

1.2 Summary of the Optimized Algorithm

For the speech coding algorithm, we chose the adaptive predictive coder (APC). The optimized APC algorithm may be summarized as follows. In the transmitter, the analog input speech is lowpass filtered at 3.2 kHz, sampled at 6.621 kHz, and divided into frames of 32.625 ms duration. Each frame of speech is preemphasized using the filter $(1-0.4z^{-1})$ and encoded using the APC encoder, to produce the quantized residual samples for that frame. The APC encoder employs (1) 3-tap pitch prediction and 6-pole spectral prediction to obtain the residual, (2) forward-adaptive quantization of the residual, and (3) pole-zero

spectral shaping of the quantization noise to reduce its perception at the coder output. The parameters of the spectral and pitch predictors and of the adaptive quantizer are quantized, coded, partially error-protected, and transmitted along with the encoded residual samples.

In the receiver, the decoded residual samples are applied to the input of a cascade of the 6-pole spectrum synthesis and the 3-tap pitch synthesis filters. The output of the cascade is deemphasized using the filter $1/(1-0.4z^{-1})$, D/A converted, and lowpass filtered at 3.2 kHz to produce the analog output speech.

1.3 Overview of the Report

In Chapter 2, we provide the rationale for our choice of the APC coder for this work and of the 6.621 kHz sampling rate for its input speech, and we describe three input-speech data bases we employed during this work. Chapter 3 reviews briefly the details of the APC coder. In Chapter 4, we describe three types of APC coder configurations that we investigated in this work. In Chapter 5, we define an important quantity called the feedback gain of the APC transmitter and show how it is related to the various APC parameters. The feedback gain concept is used throughout the later chapters to link the occurrence of

undesirably large amounts of quantization noise with positive values of the feedback gain. In the next four chapters, we describe in detail our work on developing the various aspects of the APC system: quantization of APC parameters (Chapter 6); methods for adaptive noise shaping (Chapter 7); methods for coding the APC residual (Chapter 8); and methods for preventing the "limit-cycle" behavior of the APC system (Chapter 9). The results of our work on APC algorithm optimization at 16 kb/s are reported in Chapter 10 for error-free channels, and in Chapters 11 and 12 for noisy channels. Chapter 11 describes several optimized APC coders, each with bits allocated for error protection of parameters but operating over error-free channels, while Chapter 12 contains the results of evaluation of the performance of these optimized coders in 1% channel error. In Chapter 12, we also report a single APC system design as being the most robust and best overall 16 kb/s coder. The performance of this optimized coder in office noise and in ABCP noise environments is treated in Chapter 13, while its performance in tandem with LPC-10 is described in Chapter 14. In Chapter 15, we present a design of a 16 kb/s baseband coder as well as the comparative results of this coder and the optimized APC coder. In Chapter 16, we describe several modifications to the optimized coder to simplify some aspects of the coder and to make refinements to the coder design. Also in Chapter 16, we

summarize the details of the final optimized 16 kb/s APC coder and present the results of testing the coder's real-time implementation on the MAP-300. Finally, in Chapter 17, we summarize the results of this work, and identify explicitly what we believe are the major contributions of this work.

Contained in the appendices are: Specification of the Optimized 16 kb/s APC Algorithm (Appendix A); User's Guide for the FORTRAN Simulation of the Optimized 16 kb/s APC Coder (Appendix B); and a listing of the source programs of this FORTRAN simulation (Appendix C).

2. CHOICE OF 16 KB/S CODER

2.1 Rationale for Choosing APC

As candidates for our choice of the 16 kb/s speech coding algorithms, we considered a number of coders including adaptive residual coder, APC, delta modulation systems, sub-band coder, adaptive transform coder (ATC), and baseband coder (BBC). We investigated each of these coders to see if it could satisfy the requirements given in Section 1.1. We concluded that some of the coders cannot produce toll quality speech at 16 kb/s. The coders that are capable of transmitting toll quality speech at 16 kb/s (assuming good quality input speech and error-free channel) are APC, ATC, and BBC. Flanagan's recent assessment is in agreement with this conclusion [2].

Previous work at BBN has dealt with both APC [5] and BBC [3,4] systems. The results of this work show that the APC system with an appropriate noise spectral shaping produces output speech at 16 kb/s that is almost indistinguishable from noise-corrupted input speech with a signal-to-noise ratio of 10-30 dB [5]. The BBC coder produces either background roughness or low-level tones, depending on the method of high-frequency regeneration used [3,4]. For the ATC coder [6], proper decoding of the

transmitted signal data (transform coefficients) requires an error-free transmission of the side information. This indicates the strong possibility that the quality of the ATC speech would degrade drastically in a relatively high channel-error environment. In contrast, an error in the side information of the APC coder may change the spectral envelope and cause perceivable distortion, but the degradation may be much more graceful than might happen in ATC. Therefore, we chose APC as the best overall approach to the present application.

2.2 Sampling Rate of Input Speech

One of the requirements on the speech coder is that the bandwidth of the input speech of the coder be greater than or equal to 3.2 kHz. The audio signal interface provided by GTE Sylvania for the MAP-300 array processor provides lowpass filters with -3 dB cutoffs of 3.2 kHz and 3.8 kHz (Appendix A in [3]). Therefore, the input sampling rate F_S may be chosen to have a value around 6.67 kHz or 8 kHz. Since the coding and error-protection of the side information of the APC coder is expected to take up about 3 kb/s, choosing $F_S=8$ kHz leads to a residual quantization accuracy of only about 1.6 bits/sample. The resulting quantization noise may more than offset the advantage that the choice of $F_S=8$ kHz may yield slightly higher

intelligibility. Also, the computational load is greater with the higher sampling rate. Therefore, we chose $FS=6.67$ kHz as the approximate sampling rate of the input speech.

The exact value of the sampling rate has to be selected from the options provided by the real-time clock in the audio signal interface for the MAP-300 [3]. The primitive clock rate provided by the master oscillator within the interface is 384 kHz. A candidate sampling rate is given by $384/D$, where D is the programmable, integer divide-ratio. We chose the sampling rate of $384/58$ (approximately 6.621 kHz), since this choice 1) avoids aliasing and 2) yields a variety of 16 kb/s coder realizations with different frame sizes and having integer numbers of both samples per frame and bits per frame. In all the simulations of the APC coders on our PDP-10 computer, we used a sampling rate of 6.67 kHz (or 150-microsecond sampling period), since it is close to the chosen sampling rate and since all the simulation results can be simply carried over to the real-time system. For example, a frame size of 27 ms contains 180 speech samples at 6.67 kHz; the corresponding frame size for the real-time system is 27.1875 ms, which also has 180 speech samples.

2.3 Data Bases

We employed three data bases of 11-bit linear PCM speech in this project: a high-quality data base, an "office-noise" data base, and an ABCP data base. The high-quality data base has 12 sentences of about 2-3 seconds duration each, with equal numbers of sentences from male and female talkers. This data base is the same as the one used in a previous DCA contract at BBN [3]. The signal-to-noise ratio of the speech in this data base is about 60 dB. The office-noise data base has 10 sentences, which we digitized at 6.67 kHz directly from a sponsor-supplied audio tape recorded in an office-noise environment (with the acoustic background noise at a level of about 60 dB SPL re 20 micronewtons per square meter). For the ABCP data base, we digitized a number of utterances from a sponsor-supplied audio tape containing speech recorded in an ABCP environment. The level of the background acoustic noise in such an environment is typically about 90 dB SPL.

3. REVIEW OF THE APC SYSTEM

3.1 Basic APC system.

The basic APC system is depicted in Fig. 1. The feedback structure, which constitutes the transmitter, encodes the sampled input speech $S(z)$ ¹ in terms of the quantized residual $\hat{W}(z)$ and the spectral (or linear prediction) and pitch inverse filters $A(z)$ and $C(z)$ given by:

$$A(z) = 1 + \sum_{k=1}^P a(k)z^{-k}, \quad (1)$$

and

$$C(z) = 1 + \sum_{k=M-m}^{M+m} c(k)z^{-k}, \quad (2)$$

where $a(k)$, $1 \leq k \leq p$, are the spectral predictor coefficients; $c(k)$, $M-m \leq k \leq M+m$, are the pitch predictor coefficients; and M is the pitch period in number of samples. We refer to the order of the spectral predictor p as the LPC order and the order of the pitch predictor $2m+1$ as the number of pitch-filter taps. The spectral predictor $A(z)$ is designed to remove the redundancy due to

¹ $S(z)$ denotes the z -transform of the time signal $s(n)$.

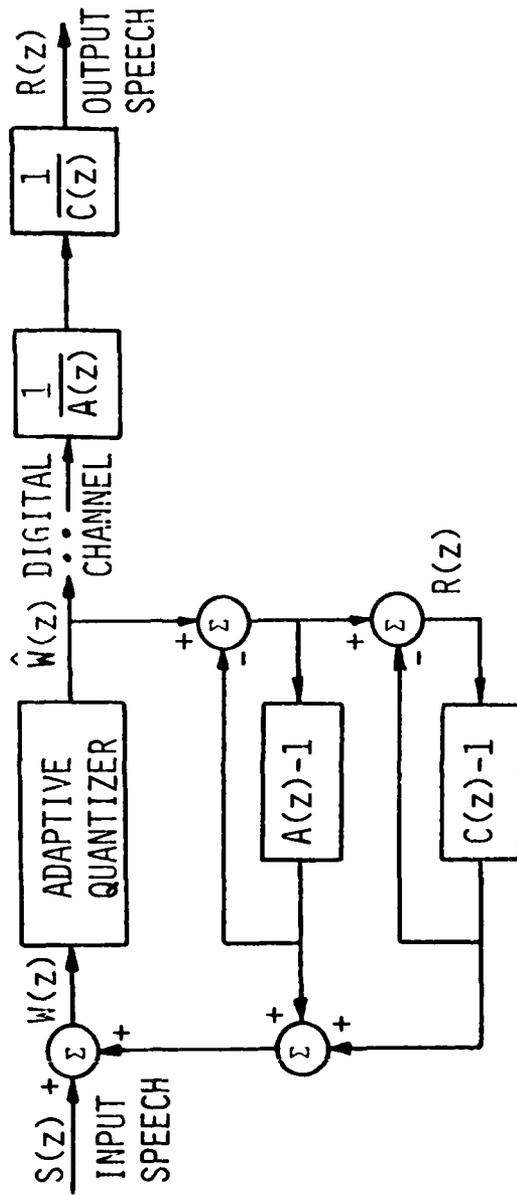


FIG. 1. Basic APC system.

spectral or short-term correlations, while the pitch predictor $C(z)$ removes the redundancy due to long-term correlations produced by pitch periodicity. The early implementations of APC [7-10] have either used no pitch predictor ($c(k)=0$, for all k) or used a 1-tap predictor ($m=0$). Recently, a 3-tap predictor ($m=1$) has been used [11]. The APC residual $W(z)$ is quantized into $\hat{W}(z)$, which is transmitted across the digital channel. At the receiver, the decoded signal is filtered by the all-pole spectral filter $1/A(z)$ and then by the pitch-synthesis filter $1/C(z)$, to produce the output speech $R(z)$. In the absence of channel bit-errors, it can be shown that

$$R(z) = S(z) + Q(z), \quad (3)$$

where $Q(z)$ is the quantization noise:

$$Q(z) = \hat{W}(z) - W(z). \quad (4)$$

The spectral and pitch predictor coefficients and the parameters of the quantizer are varied adaptively in time to track the changing properties of the input speech signal. There are two types of adaptive schemes: (1) forward-adaptive schemes, which transmit the parameters being adapted, once every frame [8,9]; and (2) backward-adaptive schemes, which do not transmit any parameters and estimate them at the receiver from the decoded residual samples [8,12]. Since the performance of the backward-

adaptive schemes would degrade significantly at channel-error rates as high as 1%, we chose to consider only forward-adaptive schemes.

The parameters of the predictors $A(z)$ and $C(z)$ are computed in such a way that the mean-square value of the quantization noise $Q(z)$ is minimized. In our simulation of the APC system in Fig. 1, the pitch and the taps are computed from the speech signal using either the autocorrelation method or the covariance method of linear prediction [13] (see Section 8.2.1); the speech signal is then inverse-filtered with $C(z)$ to produce the "first residual" $E_1(z)$; the spectral coefficients $a(k)$ are computed from $E_1(z)$ using the autocorrelation linear prediction method; the residual signal $E_1(z)$ is inverse-filtered with $A(z)$ to produce the "second residual" $E_2(z)$; and finally, the parameter(s) of the adaptive quantizer are computed from this second residual (see Section 3.2 and Chapter 8 for more details on the adaptive quantizer). The two inverse-filtering operations just mentioned and the two prediction operations $A(z)^{-1}$ and $C(z)^{-1}$ within the feedback structure in Fig. 1 are performed using the corresponding quantized parameters, to correspond to what the receiver does; this also yields a smaller mean-square value for the quantization noise than if we had used unquantized parameters. Similarly, the parameter(s) of the quantizer are also quantized before being used in the APC loop.

3.2 Forward-Adaptive Residual Quantization

The adaptive quantizer that we have used in this work is shown in Fig. 2. This quantizer has a gain normalization $1/G$ followed by an optimum (uniform or nonuniform) unit-variance quantizer. Ideally, the value of G is chosen such that the normalized APC residual $\hat{U}(z)$ has unit variance. Since the APC residual $W(z)$ becomes available only as the APC loop is in operation, G is computed approximately as the rms value of the second residual $E_2(z)$. G is transmitted to the receiver along with the encoded $\hat{U}(z)$. $\hat{W}(z)$ is computed from $\hat{U}(z)$, both at the transmitter and at the receiver, by multiplying it with G . The unit-variance quantizer is usually designed by minimizing the mean-square error between $U(z)$ and $\hat{U}(z)$, assuming a certain probability distribution for $U(z)$ (e.g., Gaussian, Laplacian, gamma, etc.) [14,15].

The residual quantizer produces two kinds of quantization error: 1) clipping error, which is produced whenever the value of signal $u(n)$ lies outside the extreme ranges of the quantizer; and 2) roundoff error or granular noise, which is produced whenever $u(n)$ lies within the extreme ranges of the quantizer. Granular noise, which causes a degradation in the output speech in the form of broad-band background noise, usually constitutes the

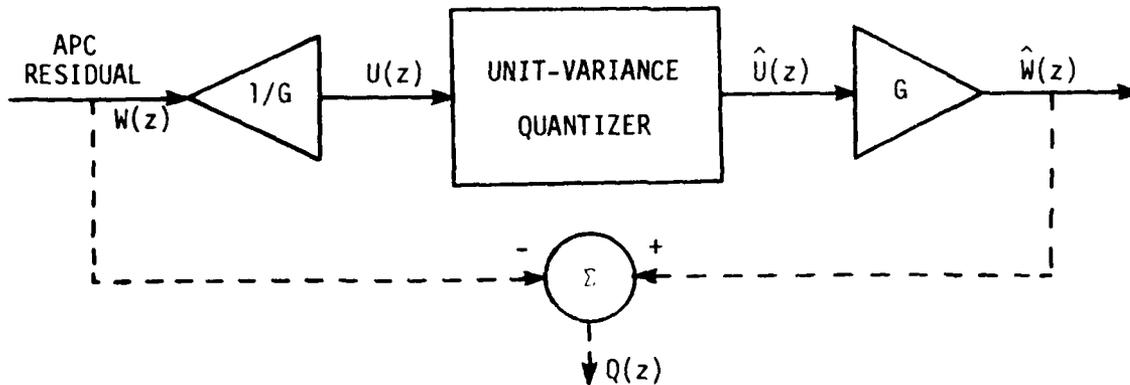


FIG. 2. Forward-adaptive quantizer for the APC residual.

dominant form of degradation. Clipping errors, on the other hand, cause undesirable degradation in the form of "pops" or "clicks"; such effects can be perceived even when the incidence of clipping errors is as low as 0.1% [5]. In Chapter 8, we present several methods for reducing the clipping errors. In the next subsection, we review an approach to reduce the perception of the granular noise.

3.3 Noise Spectral Shaping

If the adaptive quantizer in Figs. 1 and 2 is designed such that the quantization noise $Q(z)$ is mostly granular noise, then $Q(z)$ has a flat spectral envelope, which can mask the speech spectrum at high frequencies; this causes the perception of a hissing background noise in the output speech $R(z)$ [5,11]. To minimize the perception of such noise, proposals have been made recently for proper shaping of the noise spectrum [5,11], resulting in the following revised expression for $R(z)$:

$$R(z) = S(z) + B(z)Q(z); \quad (5)$$

the filter $B(z)$ is designed to shape the noise spectrum in a way that yields a perceptually more pleasing output speech. Let us denote the new output noise as $Q'(z)$

$$Q'(z) = B(z)Q(z). \quad (6)$$

The specific noise shaping methods that we investigated and their implementation issues are treated in Chapter 7.

3.4 Signal-to-Noise Ratio Considerations

The signal-to-output noise ratio in dB is denoted by S/Q' . This may be computed in one of two ways: long-term method or segmental (short-term) method. The long-term S/Q' is computed

from the energies of $S(z)$ and $Q'(z)$ calculated over a long duration (e.g., over individual sentences). The segmental S/Q' is computed as the average over frames (typically 25 ms long) of the frame-based ratio in dB [16]. The segmental S/Q' ratio has been found to correlate better with subjective perceptual judgments than the long-term S/Q' ratio [17,18]. Although we computed both ratios in our simulations, we shall give only the segmental S/Q' values, and we shall also drop the qualifier "segmental" and the prime in S/Q' , for convenience.

We make two observations based on the results of our experimental work. First, the S/Q ratio overestimates the effect of clipping errors introduced by the quantizer. That is, a coder with clipping errors will have a significantly lower S/Q ratio than another coder with only granular noise, notwithstanding that the first coder may in fact have similar or better speech quality. Second, the spectral shaping of the quantization noise $Q(z)$ reduces the S/Q ratio but enhances the perceived speech quality. Both observations should caution the reader not to take S/Q ratios, given in this report or elsewhere, as strictly indicative of perceived speech quality.

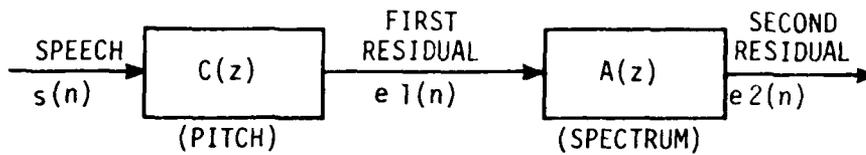
4. DIFFERENT APC CONFIGURATIONS

In this chapter, we present two methods of sequencing the spectral and pitch predictors and three types of configurations for the APC system.

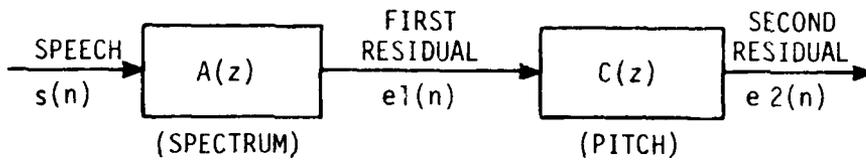
4.1 Sequencing of Spectral and Pitch Predictors

For the APC coder in Fig. 1, we have assumed the prediction sequence of pitch-followed-by-spectrum (denoted by P-S) in the sense that this sequence indicates the order, as discussed in Section 3.1 and as shown in Fig. 3(a), in which the parameters of the pitch and the spectral predictors are estimated from the speech signal. The order of predictors, C followed by A in this case, must be employed in a consistent manner in estimation, encoding (within the APC loop), and synthesis. Inconsistency in the ordering has been found to cause perceivable distortions ([5]; also see Section 12.5.2).

The first APC work employed the P-S sequence [7]; most APC implementations thus far have used this sequence as well. Only recently has the S-P (spectrum-followed-by-pitch) prediction sequence been considered [11]. The APC system using the S-P prediction sequence is obtained from Fig. 1 by simply



(a) Pitch-followed-by-spectrum prediction sequence



(b) Spectrum-followed-by-pitch prediction sequence

FIG. 3. Two methods of sequencing the spectral and pitch predictors.

interchanging the symbols $A(z)$ and $C(z)$. The order of estimation of the predictor parameters in this case is shown in Fig. 3(b).

From a mathematical point of view, since the minimum-mean-square-error estimation of the spectral predictor parameters results in large residual amplitudes at the pitch pulses, it may be argued that removing the pitch redundancy first should improve the subsequent spectral prediction. This viewpoint supports the use of the P-S prediction sequence. On the other hand, since the synthesis for the S-P sequence performs the pitch reconstruction first and then the spectral shaping function, the S-P sequence corresponds to the way the human speech production physically happens. These considerations by themselves do not suggest as to which prediction sequence should be used. Nonetheless, the prediction sequence, as will be seen in the later chapters, plays an important role in the design of the APC coder.

4.2 Prediction-Feedback Configuration

Figure 4 shows the APC configuration that we call the prediction-feedback (PF) configuration, or APC-PF. This configuration is the same as the one in Fig. 1, except that Fig. 4 includes noise shaping as discussed in Section 3.3. Notice that the feedback structure in Fig. 4 employs both $A(z)$ and $C(z)$

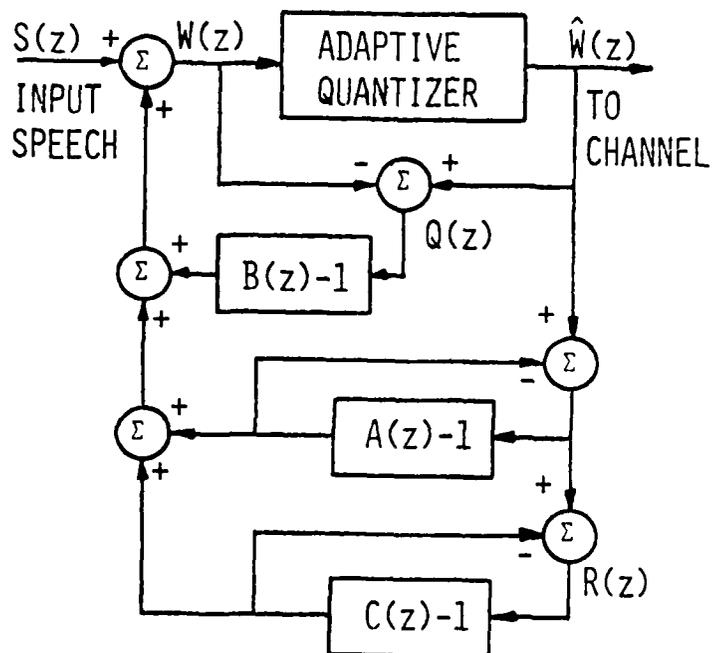


FIG. 4. The prediction-feedback configuration of APC. The figure shows the transmitter only. The receiver contains the filter cascade $[1/A(z)] [1/C(z)]$ as shown in FIG. 1.

in a predictive manner, as follows. The output of the filter $[C(z)-1]$ is the predicted value of the output $R(z)$ using the pitch prediction. Since the input to the filter $[A(z)-1]$ is the receiver's first residual $E1'(z)=C(z)R(z)$, its output is the predicted value of $E1'(z)$ using the spectral prediction. The APC-PF configuration was originally proposed in [5].

4.3 Noise-Feedback Configuration

The APC-NF (NF for noise feedback) configuration is shown in Fig. 5. This configuration was originally proposed in [19] and was used later in [20-22]. Below, we make three observations. First, we observe that in the APC-NF configuration, the input to the feedback structure is the second residual $E2(z)$:

$$E2(z) = A(z)E1(z) = A(z)C(z)S(z). \quad (7)$$

Second, only the quantization noise $Q(z)$ is fed back to the input. That is, there is no feedback path from the quantized residual $\hat{W}(z)$ as in the APC-PF configuration. The third observation that follows is quite important. The feedback transfer function $F1(z)$, which is given by

$$F1(z) = A(z)C(z)B(z) - 1, \quad (8)$$

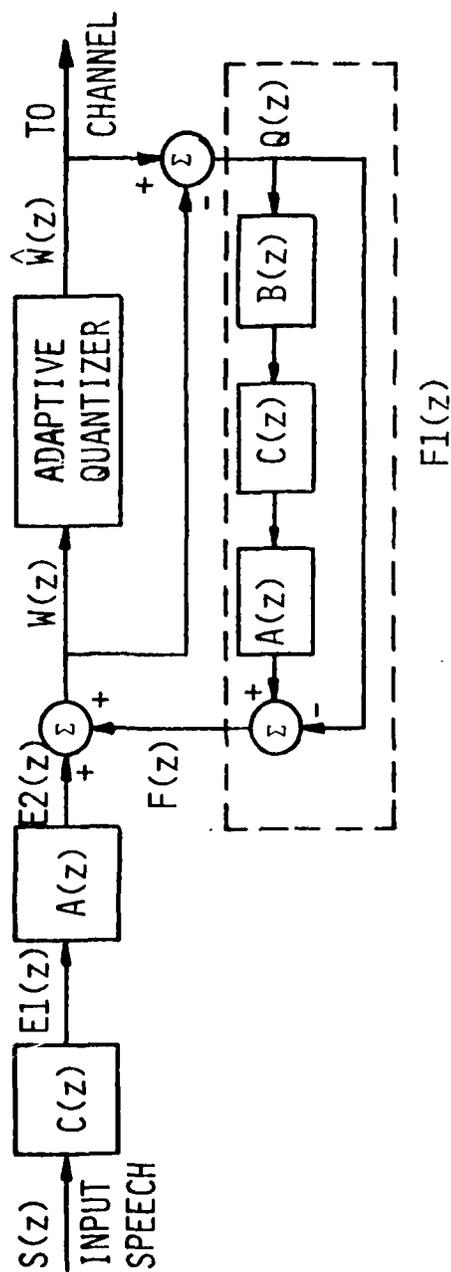


FIG. 5. The noise-feedback configuration of APC. The figure shows the transmitter only. The receiver contains the filter cascade $[1/A(z)] [1/C(z)]$ as shown in FIG. 1.

can be implemented as a single filter by multiplying out the individual filter functions. But, as simple calculations will show, this procedure requires additional storage and/or computations. A simpler procedure is to implement them as a sequence of 3 filters along with a straight feedforward branch with a transfer gain of -1 , as shown in Fig. 5. But, in what order should the 3 filters occur? Does it matter? Yes, it matters. Any order other than the one shown in Fig. 5 will produce, as we have experimentally found, a significant drop in S/Q ratio and may produce occasional "squeals" in the output speech. In one experiment, the S/Q ratio dropped from 18 dB to 16 dB when we switched the filter sequence from BCA to BAC. This non-commutativity is the result of the frame-by-frame time variation of the filters involved. The correct filter sequence can be obtained starting with the APC-PF case and deriving from it the APC-NF case, either by carefully keeping track of the orders of z -transformed quantities or through a series of straightforward block-diagram manipulations.

4.4 Hybrid-Feedback Configuration

An example of the APC-HF (HF for hybrid feedback) configuration is shown in Fig. 6. For this example, which was proposed in [11], the pitch predictor $C(z)$ is placed in a

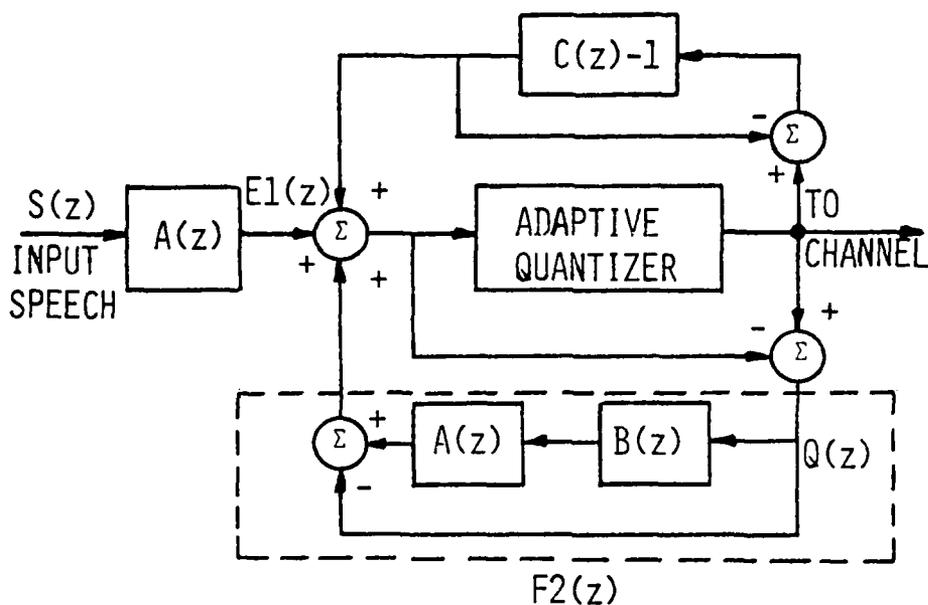


FIG. 6. The hybrid-feedback configuration of APC. The figure shows the transmitter only. The receiver contains the filter cascade $[1/C(z)] [1/A(z)]$.

predictive manner while the spectral predictor $A(z)$ is placed along with $B(z)$ in a noise-feedback manner. One can verify that the correct filter sequence in the NF path is $B(z)A(z)$, as given in the figure. It may also be seen that the example shown in Fig. 6 employs the S-P prediction sequence.

4.5 Comparison of the APC Configurations

Let us summarize the two conditions under which any two configurations become equivalent to each other: 1) same prediction sequence for the two configurations, which is used consistently in estimation, encoding, and synthesis; and 2) correct order of the filters $A(z)$, $B(z)$, and $C(z)$ within the APC encoder. For example, the APC-PF configuration that is equivalent to the APC-HF configuration given in Fig. 6 is obtained by interchanging the symbols $A(z)$ and $C(z)$ in Fig. 4.

The relative properties of the three configurations (and the two prediction sequences) are discussed in later chapters with respect to 1) ease of implementation of noise shaping (Chapter 7) and 2) insertion of a limiter in the path of the quantization noise $Q(z)$ in an attempt to prevent the build-up of excessive quantization noise (Chapter 9). The concept of feedback gain of APC is introduced in the next section, using the noise-feedback configuration.

5. FEEDBACK GAIN OF APC

Referring to the APC-NF configuration in Fig. 5, the APC residual $W(z)$ is given by

$$W(z) = E2(z) + F(z). \quad (9)$$

Thus, the APC residual is the sum of two components: the second residual $E2(z)$ and the quantization noise $Q(z)$ filtered by $F1(z) = A(z)C(z)B(z) - 1$. If the energy of $F(z)$ exceeds the energy of $E2(z)$ for any frame, then the APC residual contains less speech information than quantization noise. Carrying this argument further, if the energy of $W(z)$ is equal to the energy of the filtered noise, then $W(z)$ becomes totally dominated by noise, and the output $R(z)$ becomes non-speech and is usually perceived as "glitches" or "beeps."

To formalize these observations, we define the feedback gain of APC, G_F as the F/W ratio in dB. This definition has the flavor of the definition of the loop gain found in classical control theory texts, where the signal $F(z) = F1(z)Q(z)$ is referred to as the return signal. Our experiments have shown that when the feedback gain G_F is positive for a frame, the quantization noise builds up to excessive values due to frequent clipping errors, leading to a "limit-cycle" behavior of the

quantized signal $\hat{W}(z)$ and producing glitches or beeps in the output speech. In one experiment, we found that the supposedly unit-variance signal $\hat{U}(z)$ took on large values ranging between -20 and 40, and that the quantizer output levels exhibited a limit-cycle behavior, banging between the two extreme levels.

Below, we derive an expression for the feedback gain G_F in terms of the contributions from the quantizer and from the filter $F_1(z)$. Clearly,

$$G_F = F/W = F/Q - W/Q, \quad (10)$$

where the W/Q ratio is the signal-to-quantization-noise ratio of the adaptive quantizer in dB. If we assumed that the quantization noise is uncorrelated, then it can be shown that the F/Q ratio is the power gain of the function $F_1(z)$; this power gain is denoted by $G_P(F_1)$ and given by the sum expressed in dB of the squares of the coefficients of powers of z^{-1} in the function $F_1(z)$. Therefore,

$$G_F = G_P(F_1) - W/Q. \quad (11)$$

From the expression (11) and the results stated above, a power gain $G_P(F_1)$ larger than the W/Q ratio will lead to a limit-cycle behavior of the quantizer output $\hat{W}(z)$ and produce non-speech

output $R(z)$. Notice that the APC system is always stable in the "bounded-input-bounded-output" sense. This is because of the clipping or saturation nonlinearity of the quantizer.

Finally, we make two remarks. First, although we introduced the feedback gain and derived the expression (11) for it using the APC-NF configuration, the conclusions given above are valid for the other two configurations as well. Second, for the purpose of investigating ways of reducing the power gain $G_p(F1)$ or $G_p(ABC-1)$, one can consider the power gain $G_p(ABC)$.

6. QUANTIZATION OF APC PARAMETERS

6.1 Spectral Parameters

After the spectral parameters $a(k)$, $1 \leq k \leq p$, are computed, they are to be quantized for transmission to the receiver as well as for use in inverse filtering and APC encoding, as mentioned in Section 3.1. Our previous work has shown that optimal quantization of the spectral parameters can be accomplished by uniformly quantizing log area ratios (LARs), which are obtained by first converting predictor coefficients $a(i)$ to reflection coefficients $K(i)$ and then using the following logarithmic transformation [23,24]:

$$g(i) = 10 \log [1+K(i)]/[1-K(i)], \quad 1 \leq i \leq p. \quad (12)$$

In most of our investigations, we used $p \leq 8$. The ranges in dB of the 8 LARs obtained for the high-quality data base are given in Table 1 for the case when the spectral parameters are extracted directly from the unpreemphasized speech signal. We used the ranges in Table 1 whenever the S-P prediction sequence was used, or whenever pitch prediction was not used. Table 2 gives the LAR ranges for the P-S sequence. Again, no preemphasis was applied to the input speech. (The LAR ranges for the preemphasized case are given in Table 12, Section 16.3.)

Coeff.#	# of Bits	Minimum (dB)	Maximum (dB)	Step size (dB)
1	6	-28.039	8.961	0.5781
2	5	-7.570	20.930	0.8906
3	4	-10.359	9.141	1.2187
4	4	-5.062	12.937	1.1250
5	4	-7.437	6.563	0.8750
6	4	-5.484	8.016	0.8438
7	3	-8.250	3.750	1.5000
8	3	-4.812	6.187	1.3750

TABLE 1. Quantization of the LARs of an APC system that uses the S-P prediction sequence and no preemphasis.

Coeff.#	# of Bits	Minimum (dB)	Maximum (dB)	Step size (dB)
1	6	-23.816	9.684	0.5234
2	5	-7.547	15.453	0.7188
3	4	-9.500	6.500	1.0000
4	4	-4.219	10.781	0.9375
5	4	-6.563	7.437	0.8750
6	4	-3.516	8.984	0.7813
7	3	-7.031	5.469	1.5625
8	3	-4.812	6.187	1.3750

TABLE 2. Quantization of the LARs of an APC system that uses the P-S prediction sequence and no preemphasis.

We employed a total of 33 bits and optimally allocated them among the 8 LARs using a method reported in [24]. The optimal bit allocation and step sizes for the two cases described above are also given in Tables 1 and 2.

6.2 Pitch Parameters

The pitch parameters are the pitch period M and the taps $c(k)$, $M-m \leq k \leq M+m$. For the pitch frequency range 50-450 Hz that we assumed and the 6.67 kHz sampling rate, the pitch period M takes values in the range 14-133 samples, a total of 120 values. We used 7 bits to represent the pitch; therefore, M was "quantized" without error.

Considering the quantization of pitch taps, we investigated in this work the three cases: 1-tap ($m=0$), 3-tap ($m=1$), and 5-tap ($m=2$). We quantized the taps linearly using 4 bits for the center tap $c(M)$ and 3 bits each for all the other taps. The ranges for the taps that we used in the quantization are given in Table 3 for the two methods of computing the taps: the autocorrelation method and the covariance method.

Tap ID	Autocorrelation Method		Covariance Method	
	Minimum	Maximum	Minimum	Maximum
Center tap	-0.93	0.10	-0.99	0.50
Other taps	-0.66	0.30	-0.90	0.35

TABLE 3. Quantization ranges for the pitch taps.

6.3 Residual Quantizer Gain

The gain parameter G , defined in Section 3.2, of the residual quantizer is quantized logarithmically. Without pitch prediction, G was found to take values from -5 dB to 45 dB. With pitch prediction, we used the range -10 to 46 dB. With the exception of the entropy coding system (see Section 8.4.2), we used 6 bits for quantizing the gain in all our experimental work; for the entropy coding system, we used 10 bits.

7. ADAPTIVE NOISE SHAPING METHODS

In this chapter, we describe the methods that we used for adaptive shaping of the quantization-noise spectrum. As a useful notation, we define

$$A(z/\alpha) = 1 + \sum_{k=1}^p a(k)\alpha^k z^{-k}, \quad (13)$$

where $0 < \alpha < 1$; α may be expressed in terms of a bandwidth parameter w :

$$\alpha = \exp(-\pi w T), \quad (14)$$

where T is the sampling interval. We note that the zeros of $A(z/\alpha)$ have the same frequencies as the zeros of $A(z)$ but have bandwidths larger by w Hz. We observe that each of the noise-shaping methods described below has the property that it simultaneously reduces the S/Q ratio and the perception of the granular noise in the output speech.

7.1 All-Zero Noise Shaping

This method was proposed in [5] for the APC-PF system without pitch prediction. In this method, $B(z)$ is an all-zero or FIR filter:

$$B(z) = 1 + \sum_{k=1}^q b(k)z^{-k}, \quad (15)$$

where the leading term is unity so that the filter $B(z)-1$ is realizable within the APC loop (see Fig. 4). For the APC system without pitch prediction (any of the configurations in Figs. 4-6), it can be shown that

$$W(z) = A(z)S(z) + (A(z)B(z)-1)Q(z). \quad (16)$$

Reference [5] suggests that $B(z)$ be computed as the optimal inverse filter to $A(z)$. While this noise shaping method reduces the perception of the granular noise in the output speech, we note that the particular criterion used for choosing $B(z)$ produces also a second beneficial effect. This effect may be explained in two ways. First, from the second term on the right hand side of (16), we see that the above-mentioned criterion minimizes the noise contribution to the APC residual. Second, computing $B(z)$ as the optimal inverse to $A(z)$ is the same as minimizing the power gain of the filter-product $A(z)B(z)$, for a given order q of $B(z)$. From the discussions given in Chapter 5, the minimum-power-gain choice of $B(z)$ will reduce the extent of the limit-cycle problem.

From our experiments, we found the choice $q=1$ (1-zero

shaping) to produce the best perceptual results. For this choice, the coefficient $b(1)$ is computed as follows:

$$b(1) = -\rho(1)/\rho(0), \quad (17)$$

where

$$\rho(0) = 1 + \sum_{k=1}^P a^2(k), \quad (18)$$

and

$$\rho(1) = \sum_{k=0}^{p-1} a(k)a(k+1), \quad (19)$$

with $a(0)=1$.

For the APC system with pitch prediction, the equation (16) becomes

$$W(z) = A(z)C(z)S(z) + [A(z)C(z)B(z)-1]Q(z). \quad (20)$$

This last equation suggests that $B(z)$ may be chosen by minimizing the power gain of $A(z)B(z)C(z)$ (or as the optimal inverse to $A(z)C(z)$). For the 1-tap pitch prediction and for the 1-zero noise shaping, the minimization procedure yields a different coefficient $\hat{b}(1)$:

$$\hat{b}(1) = b(1)(1+c)/(1+c^2), \quad (21)$$

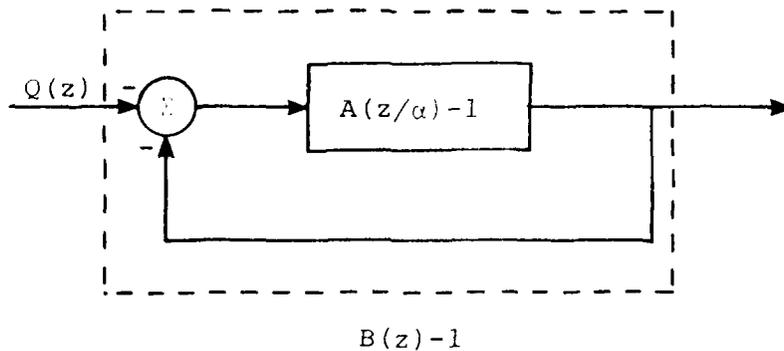
where $b(l)$ is given by (17), and c is the pitch tap $c(M)$. For c close to zero or one, $\hat{b}(l)$ is approximately equal to $b(l)$. A brief experimental comparison of the two ways of computing $b(l)$ showed that using \hat{b} caused slightly additional roughness in the output speech. In all our subsequent work, we used equations (17)-(19) to compute the coefficient of the 1-zero filter.

7.2 All-Pole Noise Shaping

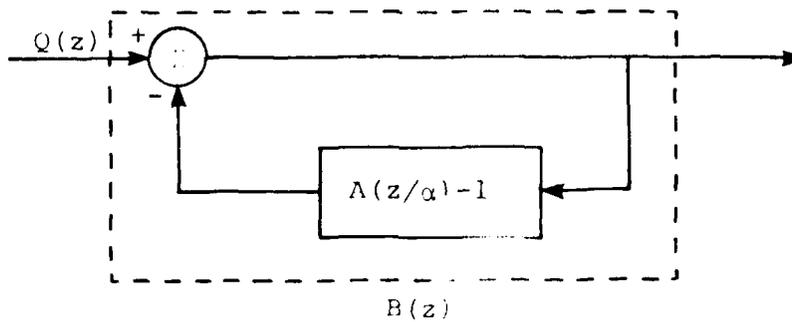
In this method, we set

$$B(z) = 1/A(z/\alpha), \quad (22)$$

where $A(z/\alpha)$ is given by (13). For this choice of $B(z)$, we note that the smaller the value of the bandwidth parameter w , the larger the extent of the noise shaping and the smaller the resulting S/Q ratio. For this method, implementation of $B(z)$ -1 for the APC-PF system is shown in Fig. 7(a), and implementation of $B(z)$ for the APC-NF and the APC-HF systems is shown in Fig. 7(b). It can be shown that this all-pole noise shaping method also reduces the power gain of $A(z)B(z)$ relative to that of $A(z)$.



(a) Implementation of $B(z)-1$, for the APC-PF system.



(b) Implementation of $B(z)$, for the APC-NF and APC-HF systems.

FIG. 7. Implementation of the all-pole noise-shaping with $B(z) = 1/A(z/\alpha)$.

7.3 Pole-Zero Noise Shaping

In this method, $B(z)$ is a pole-zero filter. Following reference [11], we considered the special pole-zero filter

$$B(z) = A(z/\alpha)/A(z). \quad (23)$$

This choice of $B(z)$ yields the following identity:

$$A(z)B(z) = A(z/\alpha), \quad (24)$$

which has two important consequences. First, defining $g_p(A)$ as the power gain in amplitude (rather than in dB) of the filter $A(z)$, we have:

$$g_p(A) = 1 + \sum_{k=1}^p [a(k)]^2, \quad (25)$$

$$g_p(AB) = g_p(A(z/\alpha)) = 1 + \sum_{k=1}^p [a(k)\alpha^k]^2. \quad (26)$$

Since $0 < \alpha < 1$, equations (25) and (26) clearly show that the power gain of AB is less than the power gain of A . Second, a configuration in which $A(z)$ and $B(z)$ occur next to each other in a cascade is best suited for implementing the above pole-zero method, since the cascade reduces to the single filter $A(z/\alpha)$. We have two such configurations: APC-HF in Fig. 6, and APC-NF in

Fig. 5 but with the S-P prediction sequence. For configurations other than these two, pole-zero noise shaping can be implemented as shown in Fig. 8. Finally, for the choice of $B(z)$ in (23), we note that the bigger the value of the bandwidth parameter w , the larger the extent of noise shaping and the smaller the resulting S/Q ratio.

7.4 Effect of Preemphasis on Noise Shaping

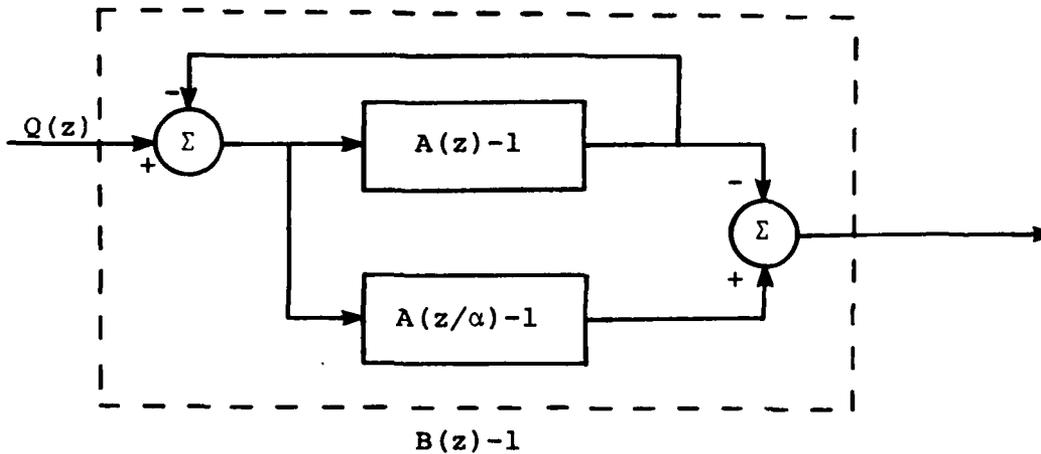
In Section 9.5, we discuss the use of preemphasis of the input speech with a filter

$$P(z) = 1 - \beta z^{-1}, \quad 0 < \beta < 1, \quad (27)$$

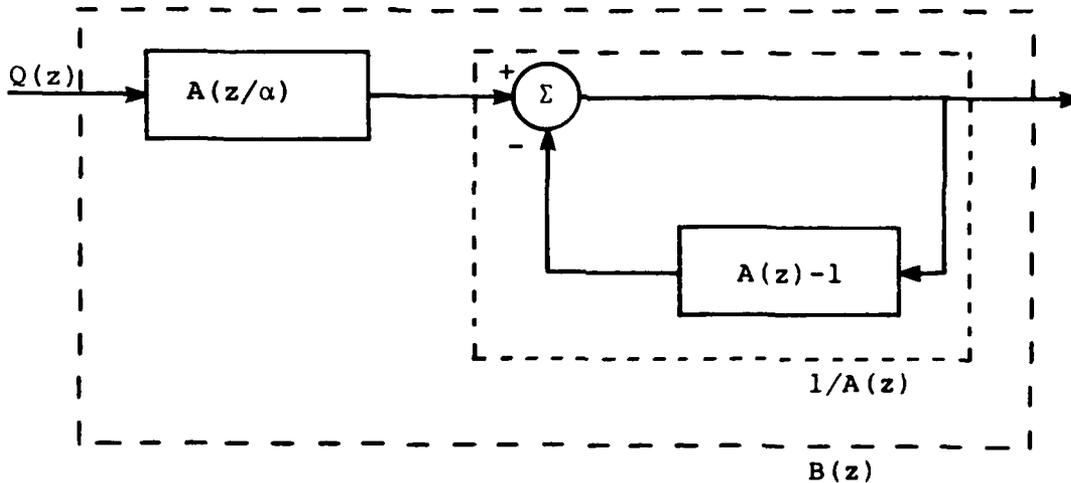
and the associated deemphasis of the coder output with $1/P(z)$. Here, we consider the effect of using preemphasis on a particular noise shaping method. For the APC coder using preemphasis, we can show that the output speech $R(z)$ is given by

$$R(z) = S(z) + [B(z)/P(z)]Q(z). \quad (28)$$

Thus, the quantization noise is further shaped by the all-pole filter $1/P(z)$. This additional shaping decreases the quantization noise at high frequencies at the expense of increasing it at low frequencies. For the case where there is no noise shaping ($B(z)=1$), preemphasis reduced the perception of the



(a) Implementation of $B(z)-1$, for the APC-PF system.



(b) Implementation of $B(z)$, for the APC-NF and APC-HF systems using the P-S prediction sequence.

FIG. 8. Implementation of the pole-zero noise shaping with $B(z) = A(z/\alpha)/A(z)$.

background noise but caused low-frequency roughness in the output speech. When we combined preemphasis with noise shaping, we found that the 1-zero method suffered the most, in that the output speech had perceivable roughness. The all-pole method was affected only slightly, but the pole-zero method was not affected in any perceivable manner.

7.5 Comparative Evaluation and Experimental Results

Figure 9 compares the different quantization noise spectra for a typical voiced sound. In this figure, plot (a) is the spectral envelope of the input speech; (b) is the unshaped (i.e., $B(z)=1$) noise spectrum; and (c) and (d) correspond to, respectively, 1-zero shaping and pole-zero shaping with $w=800$ Hz. An inspection of the noise spectra reveals that the S/Q ratio should be the least for the pole-zero method and the highest for the case without any noise shaping. Notice from Fig. 9 that the pole-zero method redistributes the quantization noise so that it is high at the places the speech spectrum has high amplitudes.

Based on the results of our experiments, we make the following conclusions:

1. Use of noise shaping reduces the perception of the granular noise in the output speech as well as the extent of the limit-cycle problem.

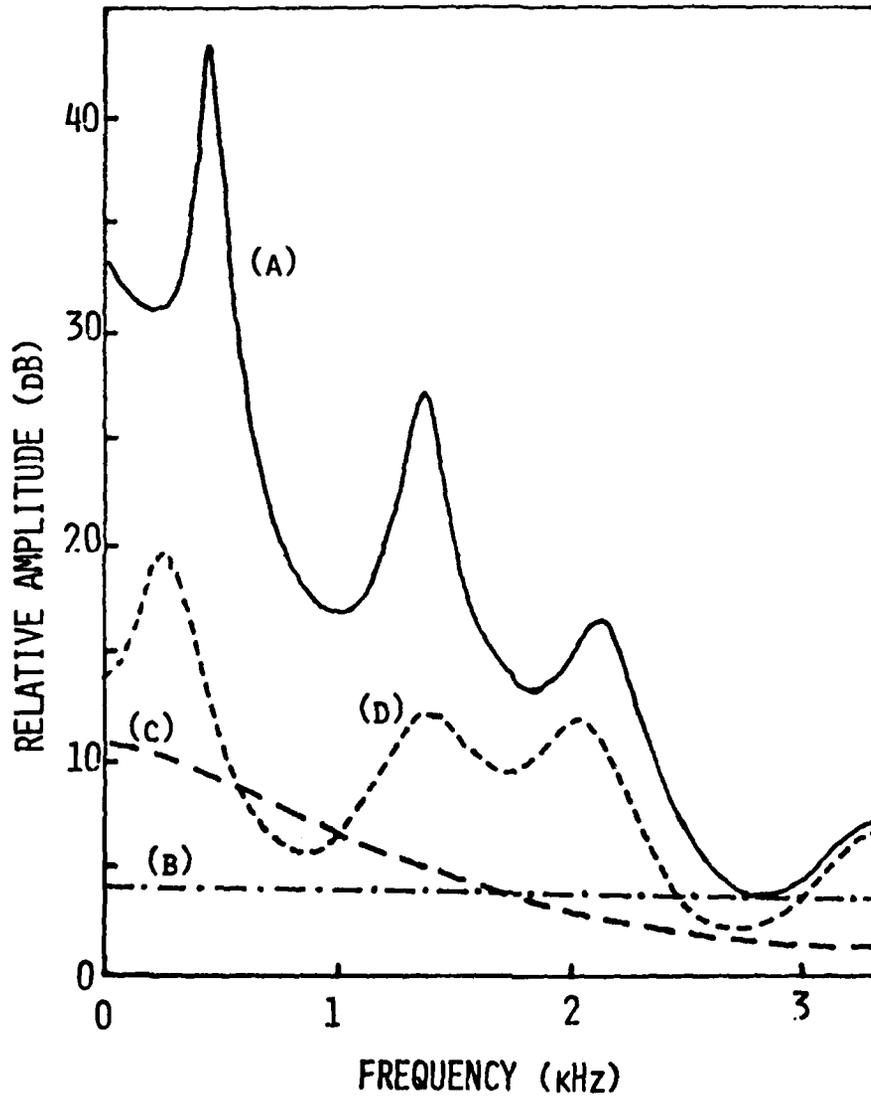


FIG. 9. Comparison of noise shaping methods.

2. The perceptual effect of noise shaping is the most for an APC coder that does not produce any clipping errors; for this case, the pole-zero method produces the best results (see Section 8.4).
3. For APC coders involving clipping errors, the three noise shaping methods produce similar speech-quality improvements, provided no preemphasis is used; the best method for a given configuration may be decided based on the ease of implementation. When preemphasis is used, the pole-zero method produces better speech quality than either of the other two methods.

8. METHODS FOR CODING THE APC RESIDUAL

Recall from Chapter 3 that the output quantization noise is wholly the result of residual coding via the APC feedback loop. Also, the residual constitutes typically about 75-85% of the total bit rate, with the remainder being used for APC parameter transmission. Therefore, proper residual coding is very important. We devoted a substantial part of our effort towards investigating the existing residual coding methods, developing new ones, optimizing them individually, and comparing their relative speech-quality performance. In this chapter, we describe the various coding methods we investigated.

8.1 Chapter Overview

For the discussions given in this chapter, unless stated otherwise, we assume that pitch prediction is not used. For all the coding methods given below, we used the forward-adaptive quantizer described in Section 3.2, with a uniform or an optimal Laplacian nonuniform quantizer.

For the chosen sampling rate of 6.67 kHz, the average number of available bits per residual sample is about 2 bits, with the remaining bits used for the transmission of all other parameters

and their error protection. With only 2 bits available, any method used for coding the residual must have provisions to deal with regions of high amplitude residual samples (e.g., during pitch pulses); otherwise, the resulting clipping errors would significantly degrade the speech quality. Below, we describe five methods of residual coding: Pitch prediction (PP), entropy coding (EC), segmented quantization (SQ), pitch adaptive (PA) coding, and segmented quantization with bit allocation (SQ-BA). Two of these methods (PP and SQ) use a fixed-length code for the residual samples, and the other three use a variable-length code. Given the limited bit resource, each of the five methods attempts to limit the extent of clipping errors in a different manner. The two fixed-length coding methods (PP and SQ) have explicit provisions for reducing clipping errors. The other three methods combat the clipping problem by varying the length of the codeword used for individual samples. Of these three variable-length coding methods, one method (EC) uses a large number of quantizer levels, and the other two (PA and SQ-BA) use only a few possible codeword lengths.

The three variable-length coding methods require variable-to-fixed rate conversion, to be useful in the present fixed-rate transmission application. To avoid involved frame synchronization problems at the receiver [25] and to ensure

reliable decoding in channel error, we chose to accomplish the variable-to-fixed rate conversion over individual frames such that every frame of transmitted data has a fixed number of bits.

Also described below are several composite coding methods obtained by combining two or more of the five basic methods; the composite methods yield significantly better speech quality and higher S/Q ratios. We denote, for ease of reference, the composite methods in terms of the above-introduced abbreviations. For example, the method PP-SQ uses both pitch prediction and segmented quantization.

8.2 Pitch Prediction

8.2.1 The Method

In this method, we use pitch prediction within the APC feedback loop as discussed in Section 3.1, a nonuniform quantizer, and a fixed-length coder for residual samples. The fixed length may be 3 or 4 levels, for example. With 3 levels, we block-code 5 residual samples (a total of $3 \times 3 \times 3 \times 3 \times 3 = 243$ levels) in 8 bits, producing an average of 1.6 bits/sample.

With the use of pitch prediction, we have the choice of employing the S-P or the P-S prediction sequence (see Section

4.1). In our work, we investigated three cases of pitch prediction (1-tap, 3-tap, and 5-tap) and two ways of computing the predictor coefficients $c(k)$: the autocorrelation method and the covariance method [13]. We computed the pitch period M as the nonzero lag corresponding to the peak of the autocorrelation function of the input speech signal in the P-S sequence and of the first residual (see Fig. 3(b)) in the S-P sequence.

Next, we make two observations on the pitch prediction method. First, the pitch-synthesis filter $1/C(z)$ has a long impulse response (on the order of several pitch periods) because of the large-delay terms (z^{-M} , etc.) of $C(z)$. As a consequence, the speech-quality effect of a channel bit-error or of any other anomaly (see Section 9.3) is propagated for a relatively long duration. However, as will be seen in Chapter 12, the propagation problem in channel error is more-than compensated by the reinsertion of the pitch or harmonic structure into the residual, which would otherwise be significantly distorted.

Second, the use of pitch prediction increases the power gain of the feedback transfer function $F(z)$ in Fig. 5. However, pitch prediction decreases clipping errors and hence increases the W/Q ratio (see Chapter 5); this compensates for the power gain increase and hence limits the net feedback gain.

8.2.2 Stability of Multi-tap Pitch Predictor

The autocorrelation method of computing the predictor parameters guarantees the stability of the filter $1/C(z)$ for the 1-tap case only, and the covariance method does not guarantee stability at all. The instances of pitch-filter instability were found to cause pops and beeps in the output speech, especially when the quantizer used a small number (e.g., 3) of levels per sample. Our experiments showed that the computed pitch filter was almost always stable when we used the S-P prediction sequence. For the P-S sequence, we investigated two methods of "stabilizing" the multi-tap pitch filter. The 1-tap filter in the covariance method is stabilized by forcing the tap coefficient to be less than 1 in magnitude. These two methods are described below.

8.2.2.1 Switched Prediction

In this method, we check the stability of the pitch filter each frame, and if the filter is unstable, we switch to 1-tap prediction for that frame. Notice that the additional computation involved in computing the 1-tap coefficient is trivial. However, the stability check of the pitch filter involves a significant amount of computation. A straightforward method for the stability check is to obtain the $M+m$ reflection

coefficients of the pitch filter and to check if their magnitudes are all less than 1. The general recursive procedure for computing the reflection coefficients from the taps or prediction coefficients requires a number of multiplies proportional to M^2 . However, since there are only $2m+1$ nonzero taps, we can reduce the number of multiplies to about $6M$.

8.2.2.2 Stable Lattice Prediction

Using the lattice method of linear prediction [26], we have developed a new pitch prediction method. Below, we consider the lattice with only 3 nonzero reflection coefficients $K(M-1)$, $K(M)$, and $K(M+1)$. It is straightforward to derive expressions for these reflection coefficients in terms of the autocorrelations of the speech signal. The lattice method guarantees the stability of the pitch filter. But, the equivalent pitch prediction filter $C(z)$ has five nonzero taps:

$$C(z) = 1 + c(1)z^{-1} + c(2)z^{-2} + c(M-1)z^{-(M-1)} + c(M)z^{-M} + c(M+1)z^{-(M+1)}. \quad (29)$$

The expressions for the five taps in terms of the reflection coefficients are given below:

$$\left. \begin{aligned}
 c(1) &= K(M) [K(M-1) + K(M+1)], \\
 c(2) &= K(M-1)K(M+1), \\
 c(M-1) &= K(M-1), \\
 c(M) &= K(M) [1 + K(M-1)K(M+1)], \\
 c(M+1) &= K(M+1).
 \end{aligned} \right\} \quad (30)$$

8.2.3 Experimental Results

For the most part in our work, we experimented with the 1-tap and the 3-tap predictors. We investigated the 5-tap case during the parameter optimization study only (see Chapter 11). Using the P-S prediction sequence and a 4-level nonuniform quantizer, we obtained a S/Q ratio of 15.7 dB for 1-tap pitch prediction and 16.5 dB for 3-tap pitch prediction. The autocorrelation method was used in computing the tap(s) in both cases. Perceptually, the 3-tap case produced more clarity of speech than the 1-tap case.

On the matter of pitch filter stability, we found in one experiment that the 3-tap filter was unstable for about 8% of the frames. The instabilities caused beeps in the output speech when the quantizer used 3 levels per sample. The use of the switched prediction method described above proved to be a successful remedy to the instability problem: The output speech in this case did not contain beeps. Even with a 4-level quantizer, the switched 3-tap prediction method reduced discrete distortions such as pops and clicks. Using again a 4-level quantizer, we

found that the stable 3-coefficient lattice method produced a S/Q ratio of about 15.1 dB, as compared to 15.5 dB in the 3-tap case and 15.7 dB in the 1-tap case reported in the previous paragraph. Listening tests indicated that the lattice method produced "squeaky sounds" and perceivably more quantization noise than the 3-tap case. To understand why the lattice method gave a low S/Q ratio, we examined the values of the 5 taps given by (30). We found that the taps $c(1)$ and $c(2)$ tended to take on values close to ± 1 , and that the first part of the filter $(1+c(1)z^{-1}+c(2)z^{-2})$ acted as a preemphasis filter, significantly reducing the spectral dynamic range of the filtered signal. This, in turn, caused the spectral predictor to produce a significantly higher normalized prediction error value [13] than observed in the 3-tap case. This explains the observed drop in the S/Q ratio.

For the P-S prediction sequence, the covariance method of computing the taps produced the same or a slightly lower S/Q ratio than the autocorrelation method. Also, the output speech for the covariance method had low-level "scratchy noises" and clicks. However, using the S-P prediction sequence, we obtained different results. The covariance method produced about 1-2 dB higher S/Q ratios than the autocorrelation method. (We indicate here that with the S-P prediction sequence, we had to use several of the methods given in Chapter 9 to prevent the limit-cycle problem discussed in Chapter 5.)

In conclusion, we make the following observations and recommendations, based on our experimental results:

1. 3-tap prediction produces better speech quality than 1-tap prediction.
2. The pitch filter is almost always stable for the S-P prediction sequence, but this is not so for the P-S sequence. 3-tap switched prediction is a good solution for the instability problem in the latter case.
3. The autocorrelation method is recommended for computing the tap coefficients for the P-S sequence, while the covariance method is recommended for the S-P sequence. (For noisy-channel applications, use of the autocorrelation method is recommended with either prediction sequence; see Chapter 12.)
4. Pitch prediction alone does not provide satisfactory speech quality, as the output speech contains discrete distortions such as pops and clicks.

As will be seen later on in this and several subsequent chapters, pitch prediction, when added to any of the other coding methods, produces better speech quality.

8.3 Segmented Quantization

8.3.1 The Method

The SQ method proposed in [27] employs a nonuniform quantizer and a fixed-length code. The analysis frame is divided into several equal-length segments, and one value of quantizer gain G is computed for each segment. This makes the quantizer

step size adapt rapidly to local variations in the energy of the residual samples. The gain over the whole frame is computed, coded, and transmitted; in addition, the segment gain parameters are computed and their deviations from the (quantized) whole-frame gain (or "delta gain" values) are also transmitted. In our experiments, we used as many as 10 segments.

Ideally, the gain and the delta gain values should be computed for the APC residual, rather than from the first (or the second, with pitch prediction) residual. Since computing the APC residual requires the use of the quantizer, we have a "chicken-or-the-egg" problem. As a suboptimal solution to this problem, the so-called two-spin method is proposed in [27], in which the APC loop is run once with the quantizer gains obtained from the first residual and the resulting APC residual is used in computing an improved estimate for the gains. Except on one occasion (see Section 8.3.3), we used the simple method of computing the gains from the first residual.

8.3.2 Pitch Prediction with Segmented Quantization

The development and extensive investigation of the composite scheme PP-SQ have produced several important results in this work. First, the PP-SQ scheme is computationally the simplest of all the composite schemes we have considered. Second, it is the

only composite scheme that produces a fixed-length code for the residual samples, albeit it offers only two useful codelengths (3 and 4 levels) for a reasonable 16 kb/s coder design. Third, we chose the PP-SQ method for the final design of a robust 16 kb/s APC coder (see Chapter 12).

The SQ method generally provides better results for male speakers than for female speakers, since the dynamic range of the amplitudes of the APC residual within a pitch period is larger for males. The PP-SQ method yields good performance for all speakers, since pitch prediction removes the large peaks of the APC residual at the pitch pulses, and normalization over individual segments rather than over the whole frame provides a better tracking of the short-term amplitude variations of the APC residual.

For the PP-SQ method, the values of gain and delta gains are computed from the second residual $E_2(z)$. We found that the delta gains took values over a wide range from about -25 dB to about 5 dB. But, our experiments have shown that satisfactory quantization of the delta gains can be achieved using only 2 bits per delta gain. The nonuniform quantization of the delta gains that we chose from among several other methods is given in Table 4.

Quantizer Input	Level	Quantizer Output
$-\infty$	0	-8.0
-5.0	1	-3.0
-1.0	2	2.0
3.5	3	5.5
∞		

TABLE 4. Nonuniform quantization of delta gains

8.3.3 Experimental Results

In one test involving the SQ method with 10 segments and a 4-level nonuniform residual quantizer, we found that the output speech contained several beeps over the six sentences (from the high-quality data base) we processed, because of low values of the W/Q ratio and the resulting excessive quantization-noise feedback. When we used the two-spin method for computing the delta gains, the beeps were replaced by loud scratchy noises. We point out that this excessive noise-feedback problem in this case can be effectively solved by the use of a limiter as discussed in

Section 9.3. Although the beeps were not produced at the output when we used the limiter, the granular noise remained at an objectionably high level. We conclude from this and other tests that the SQ method alone is inadequate.

Before we present the experimental results for the PP-SQ method, we introduce a convenient notation. In this notation, PP3-SQ10, for example, denotes the PP-SQ method with 3-tap pitch prediction and 10-segment segmented quantization. In our initial tests of the PP-SQ method using a 3-level quantizer, we encountered severe problems of limit cycles at the quantizer output. In the next chapter, we present effective ways of preventing the limit cycles. For the rest of this subsection, we consider the use only of a 4-level quantizer.

To investigate the performance effect of varying the number of taps and the number of segments in the PP-SQ method, we tested 6 APC coders obtained by considering two values for the number of taps (1 and 3) and three values for the number of segments (1, 5 and 10). We used a frame size of 25.5 ms and quantized all the parameters except the delta gains. The S/Q ratios computed over six utterances are given in Fig. 10 for these six coders. Relative speech-quality judgments obtained via informal listening are also shown in Fig. 10. Adding segmented quantization (SQ5 or SQ10) to pitch prediction (PP1 and PP3) removed certain "tinkling

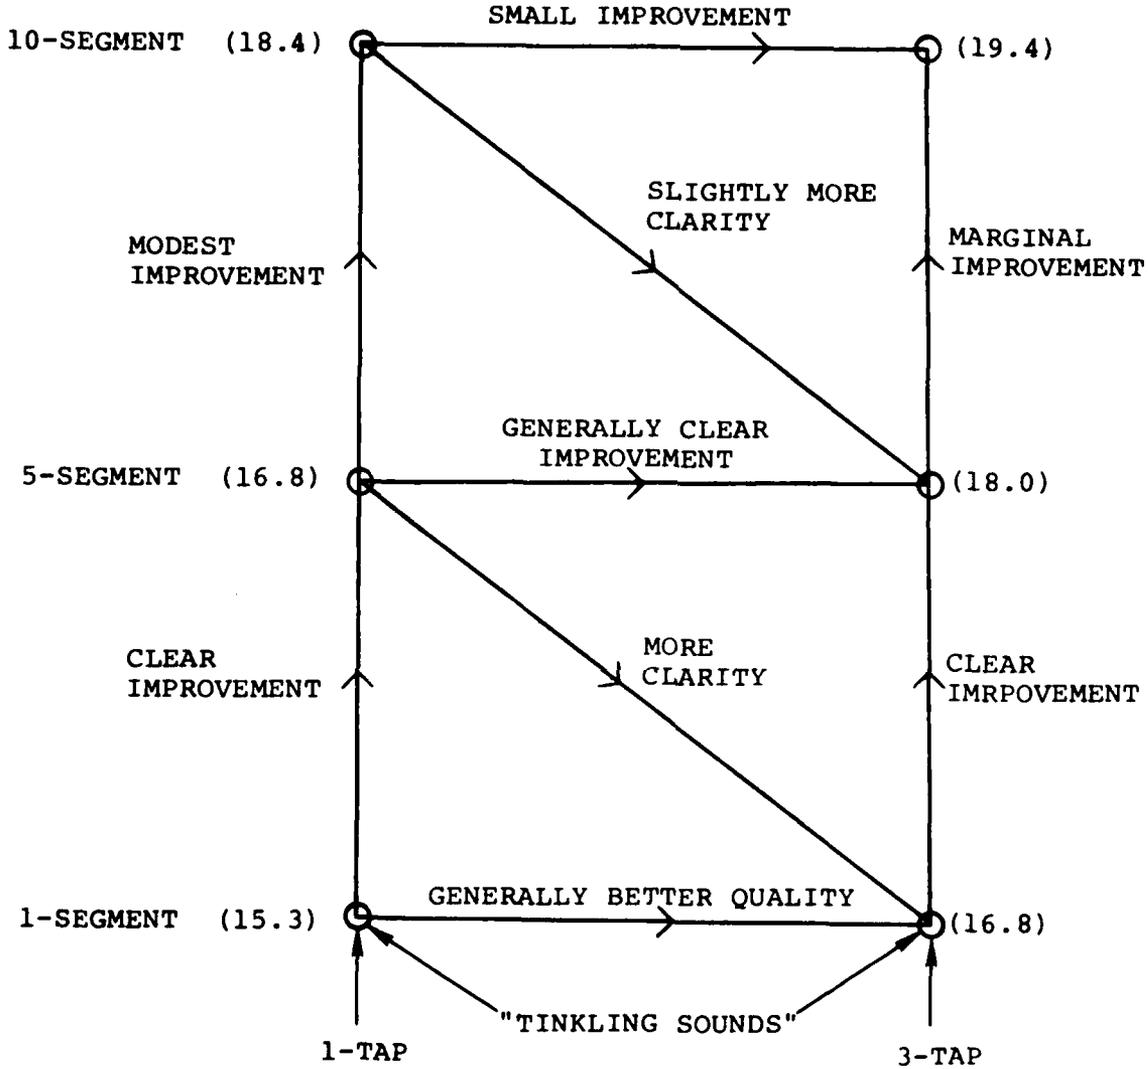


FIG. 10. S/Q ratios and relative speech-quality judgments for 6 PP-SQ APC coders. The number given within parentheses next to each node is the S/Q ratio for the corresponding APC coder. The arrow shown between each pair of APC coders points to the coder judged to produce better speech quality.

sounds" from the output speech and also provided a clear improvement in perceived speech quality. Increasing the number of taps from 1 to 3 generally provided clear improvement only for the 5-segment case. Increasing the number of segments from 5 to 10 yielded a modest improvement for the 1-tap case and only a marginal improvement for the 3-tap case. Of the 6 APC systems, PP3-SQ5 and PP3-SQ10 produced the best overall speech quality.

8.4 Entropy Coding

8.4.1 The Method

This method, described in detail in [5], uses a uniform quantizer with a large (fixed) number of levels to avoid clipping completely. To obtain an average data rate of about 2 bits/sample, the method uses variable-length Huffman or entropy coding [28]. In entropy coding of the residual, frequently occurring residual values (those close to zero) are coded with a small number of bits, and infrequently occurring values (large amplitudes, usually in the clipping range of fixed-length quantizers) are coded with a large number of bits in such a way that the average bit rate is minimized. Following reference [5], we used a (suboptimal) self-synchronizing entropy code, with codewords given by: 0, 10, 110, 1110, etc., because of our noisy-

channel application. Notice that in case of a channel bit-error, the decoding error is in the form of either merging two residual samples or splitting one into two.

The fixed step size of the uniform quantizer is a parameter, whose value is experimentally computed to obtain the desired average number of bits/sample. Increasing the step size decreases the entropy of the code by forcing more samples into the level coded with one bit; this decreases the required average number of bits per sample. Similarly, decreasing the step size increases the average number of bits per sample. With the adaptive quantizer in Fig. 2, changing the quantizer step size can be accomplished by changing the quantizer gain G . To interface the variable-rate coder to a fixed-rate channel, we require variable-to-fixed rate conversion, which is described next.

8.4.2 Variable-to-Fixed Rate Conversion

The scheme we used for achieving this conversion was developed as part of another government contract at BBN [29]. In applying this scheme, we made some modifications to improve its performance. Described below is an outline of the rate conversion scheme.

The scheme forces the number of bits to be a constant at each frame by changing the gain factor in front of the quantizer and repeating the analysis of the APC loop for the whole frame. The new residual thus obtained is coded and the number of bits used is computed. The process is repeated iteratively until the desired convergence is achieved. As a practical matter, we limited the number of iterations to 5. At the end of the 5th iteration, the method chooses the gain value that yielded a number of bits that is closest to, but not exceeding, the desired number; the difference is made up by inserting one-bit codes or "filler bits." The analysis of the APC loop is then repeated with this final gain value; this step avoids having to store all the intermediate residual codewords obtained in the five iterations. In case all five iterations provided total bits for the frame above the desired value (which we encountered a few times in our experiments), the method chooses the one that came closest to the desired number and discards as many of the residual samples from the end of the frame as required. The discarded samples are assumed to be zero at the receiver.

The gain adjustment mentioned above is performed using a procedure outlined next. Denote by δ the difference between the actual number of bits provided in an iteration and the desired number. During the initial iterations, until a "zero-crossing"

is obtained in the sense that δB changes sign, the method adjusts the gain using a modified version of the so-called 6 dB/bit rule, i.e., the gain adjustment in dB is given by $(6 + |\delta B|/20)\delta B/N$, where N is the number of samples in the frame and the term $|\delta B|/20$ is the modification we found to yield good convergence. Once the "zero-crossing" is obtained, the method adjusts the gain using the modified false position method of computing the zero of a function [30].

Since the number of bits used per frame is quite sensitive to even small changes in the gain G (e.g., 0.1 dB), we use the unquantized G at the input of the quantizer and the quantized G at its output (see Fig. 2); to minimize the mismatch between the two gain values, we use 10 bits for quantizing G .

8.4.3 Experimental Results

In testing the variable-to-fixed rate conversion scheme, we found that the number of levels used by the quantizer had to be increased from the previously used value of 19 to 43, to produce about the same perceived speech quality as from the free-running variable-rate system. The S/Q ratio, however, dropped about 1 dB, because of the gain or step size adjustment required by the rate conversion scheme. The transmission rate of the filler bits was found to be about 5 bits per frame or about 200 b/s.

We combined pitch prediction with entropy coding to produce the composite scheme EC-PP. Our investigation has clearly shown that for the same bit rate, EC-PP1 produces better speech quality than EC, and EC-PP3 produces further improvement over EC-PP1. In one test, we used a frame size of 25.5 ms and 1-zero noise shaping. The three 16 kb/s coders, EC, EC-PP1, and EC-PP3, yielded the S/Q ratio values: 19.9 dB, 20.8 dB and 21.2 dB.

Use of noise shaping in the entropy coding method produced significantly better speech quality in the form of reduced granular noise relative to the case without noise shaping. We report an interesting experiment that demonstrates the importance of the type of noise shaping used in the EC method. In the initial stages of this work, we mostly used the 1-zero noise shaping method. With this noise shaping, the 4-level PP1-SQ10 method with a S/Q ratio of 18.2 dB produced a perceivable speech quality improvement (in particular, reduced quantization noise) over the EC-PP3 method with a S/Q ratio of 21.2 dB, notwithstanding the S/Q ratio difference of 3 dB. However, when we later used the pole-zero noise shaping method we observed a reversal in the speech quality ordering of the two coders. The EE-PP3 system with pole-zero noise shaping ($w=800$ Hz) yielded a reduced S/Q ratio of about 19.0 dB, but produced speech that was significantly better than from the same system with 1-zero noise shaping.

8.5 Pitch-Adaptive Coding

The original idea (see Section 8.5.1) that led to our development of the PA method described below was recently proposed in [31], as a way of efficiently allocating the available bit-rate resource for time-domain residual quantization. The basic PA method we describe in Section 8.5.2 is far simpler in both side-information transmission requirements and computational complexity than the method proposed in [31]. Also included in Section 8.5.2 is a simple method of obtaining a fixed number of bits/frame. We have made several modifications to the basic PA method, producing a viable and effective pitch-adaptive coding method that includes pitch prediction and pitch-synchronous segmented quantization (rather than the time-synchronous SQ discussed above in Section 8.3). This novel PA-PP-SQ coding method represents a significant contribution of this work. This notation, although long, allows us to specify conveniently the number of segments in a pitch period, the number of pitch taps, and the number of transmitted delta gains (e.g., PA4-PP3-SQ3). Since the development of the PA-PP-SQ method was based on the experimental results from our earlier versions of the PA method, we present below the experimental results as part of the initial subsections (unlike in the previous Sections 8.2-8.4).

8.5.1 Itakura's Method

This method divides each pitch period of the first (LPC) residual into a small number, e.g., 4, of equal-length segments. The energy per sample E_i is computed for each segment. Let B_i be the number of bits used to quantize the residual samples in the i th segment, where $i = 1, 2, 3, 4$. For convenience, the indexing is chosen such that $E_1 \geq E_2 \geq E_3 \geq E_4$. Let B_0 be the average of the four B_i values. If the pitch period is M samples, the problem is then to allocate the MB_0 bits so that the mean-square quantization error is minimized. It can be shown [6] that for the optimal case, the mean-square errors in individual segments should be equal, and B_i is given by

$$B_i = B_0 + 1/2 \log_2 [E_i / (\prod_{k=1}^4 E_k)^{1/4}]. \quad (31)$$

This means that the segment containing the pitch pulses (segment numbered 1) is assigned the maximum number of bits, and the segment numbered 4 is usually assigned the fewest bits.

The side information transmitted to the receiver consists of the following quantities: pitch period value(s) and as many sets of E_i and of locations of segments as there are pitch periods in the frame. The method requires the calculation of one bit assignment per pitch period at both the transmitter and the

receiver. The details of locating pitch periods and locating segments within a pitch period are not given in [31]. Also, the reference does not consider the issue of fixed-rate transmission.

We make two remarks on Itakura's method. The first remark deals with the apparent similarity between this method and ATC. There is an important difference between the two methods: ATC operates in the frequency (transform) domain, while the above method functions in the time domain. In ATC, some formant peaks may fade in and out of a frequency band, which causes time-varying effects usually perceived as clicks. Such fade-in and fade-out events can also occur in the above method, but they happen in the time domain and thus may not produce a perceptually degrading effect. Second, although Itakura's method computes pitch, it does not use pitch prediction.

8.5.2 Basic Pitch-Adaptive Method

The basic PA method employs a forward-adaptive nonuniform quantizer that uses a variable number of bits/sample. Unlike the EC method, this method uses only a small number (e.g., 4) of code-lengths for the residual samples.

Figure 11 illustrates the basic PA method by means of a 4-segment example. As shown in the figure, one value M of the

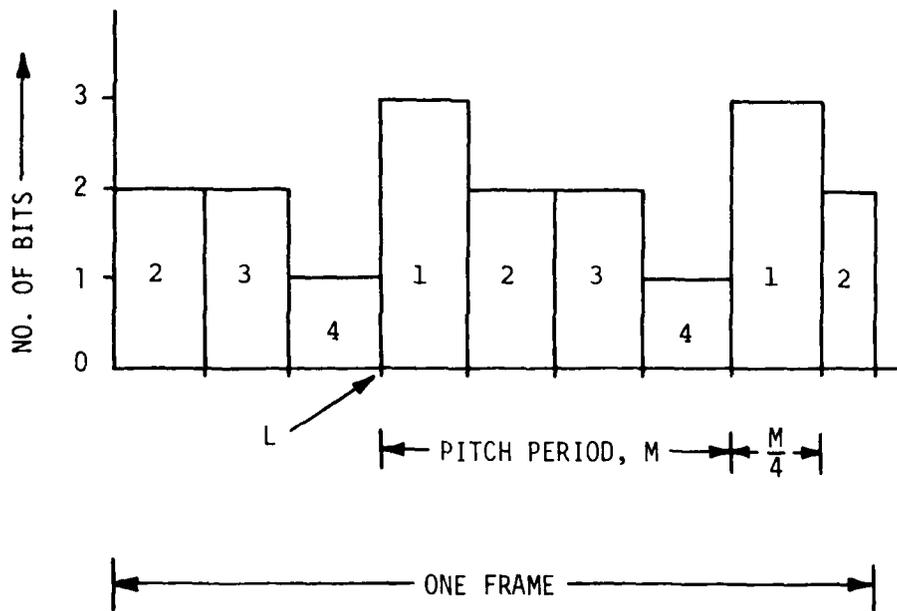


FIG. 11. Bit assignment in the basic pitch-adaptive coding scheme.

pitch period is used in dividing the whole frame into segments of $M/4$ samples each. The segment at the end of the frame has fewer than $M/4$ samples. We explain later how to accommodate cases where $M/4$ is not an integer. For the example in the figure, segments numbered 1 are quantized using 3 bits, segments numbered 2 and 3 are quantized using 2 bits, and segments numbered 4 are quantized using 1 bit. The average number of bits per sample is approximately 2. This allocation of bits among the segments, denoted by $\{3,2,2,1\}$ for this example, is fixed in time, so that this information is not transmitted to the receiver. The side information to be transmitted to the receiver consists of three quantities: the pitch period M , the location L of the beginning of the very first segment numbered 1, and a 1-bit code to be defined in the next subsection. Notice that L can take only 4 values: 0 , $M/4$, $M/2$, and $3M/4$. For the frame shown in the figure, $L=3M/4$. For the 4-segment example, L is transmitted using 2 bits.

Let us consider the case when $M/4$ is not an integer. If $(M-J)/4$ is an integer, where J can be 1, 2, or 3, then 3 of the 4 segments over each pitch period are chosen with $(M-J)/4$ samples. The segment with the assigned bits per sample closest to the desired average bits per sample is made to contain $J+(M-J)/4$ samples.

The above discussion has considered only voiced frames. The case of unvoiced frames is treated as part of the next subsection dealing with the procedure used to compute L.

8.5.2.1 Computation of the Location Parameter

Using optimal nonuniform unit-variance Laplacian quantization, we have precomputed and stored in memory the quantization tables corresponding to the various numbers of bits/sample and the corresponding mean-square quantization errors. For a given frame, we determine the optimal value of the location parameter L as follows. For each allowable value of L, which uniquely defines the segmentation procedure as discussed above, we compute the average over different segments of the quantity which is the product of the sum of the squares of the residual samples in a segment and the stored mean-square quantization error for that segment. This average error measure is computed for two cases: nonuniform bit allocation among segments (e.g., {3,2,2,1}), and uniform bit allocation (e.g., {2,2,2,2}). The two cases are coded using a 1-bit code U/NU (uniform/nonuniform). The values of L and U/NU that yield the least average quantization error are used in the PA scheme for that frame. The inclusion of the uniform bit allocation case allows handling of the unvoiced frames as well. The example

uniform bit allocation given above, namely $\{2,2,2,2\}$, is clearly for the case when the desired average over the frame is 2 bits/sample. If this average is not an integer, say, 1.75 bits/sample, then a bit allocation such as $\{2,2,1,2\}$ may be used for the "uniform" option. In this last case, choosing an optimal value of L is meaningful even for unvoiced frames.

Two remarks are in order. First, since L can take only a small set of values (e.g., 4), an exhaustive minimization over this set to compute the optimum L is quite reasonable. Second, computation of segment energies for each L can be simplified by computing once and storing the squares of all the samples over the frame. (Segment energies by themselves cannot be stored since segment widths and locations change with L .)

8.5.2.2 Variable-to-Fixed Rate Conversion

This procedure yields a fixed, prespecified number of bits over each frame. Let B_0 denote the average number of bits per sample corresponding to this desired total number of bits. For a given value of L and the associated segment assignment, the "ideal" periodic bit allocation is assigned for the same, and the total number of bits used for the frame is computed. If this total exceeds the desired number, the following action is taken. Starting from the beginning of the frame, a search is made for a

segment boundary where a segment with its assigned bits/sample less than the average B_0 precedes a segment with its assigned bits/sample equal to or greater than B_0 . This segment boundary is then shifted to the right (more precisely, towards the end of the frame), to increase the size of the first segment by one sample. If the resulting total is still too high, the segment boundary is shifted again at a place one pitch period later. At the end of the frame, if the total is still too high, the above process is repeated. On the other hand, when the ideal bit allocation mentioned at the beginning of this paragraph leads to a total number of bits that is less than the desired total, then the shifting of segment boundaries is done in the opposite direction.

The above iterative process typically converges within a few iterations. Unlike the procedure used in the entropy coding method (see Section 8.4.2), this procedure does not actually quantize the residual samples until the proper segment and bit assignment have been decided; also, the individual iterations in this method are relatively simple to perform.

8.5.2.3 Comparison with Segmented Quantization

For the PA method described above, each segment is made up of similarly located samples from the successive pitch periods of

a frame. In this sense, the PA method employs pitch-synchronous segments. In contradistinction to this, the SQ method uses time-synchronous segments. More important, on the one hand, the SQ method, by using different quantizer gains over individual segments, tracks the residual amplitudes by expanding or contracting the quantizer step size. On the other hand, the PA method uses the same gain and hence the same step size over the whole frame, but tracks the residual amplitudes by increasing or decreasing the number of bits. One can show that the PA method gives a higher S/Q ratio than does SQ. The above interpretation of pitch-synchronous segments was used in developing the PA-PP-SQ scheme described in Section 8.5.3.

8.5.2.4 Experimental Results

In our experimental investigation of the basic PA method, we found that the output speech contained occasional beeps. When we added pitch prediction to the PA method, this limit-cycle problem was eliminated. However, in our tests, a 4-segment, 1-tap PA4-PP1 method produced a S/Q ratio of only 17.1 dB. One reason for this may be that the residual samples within a segment are quantized using a unit-variance quantizer, but they do not have unit variance. A solution to this problem, discussed in the next section, is to employ segmented quantization, but using the pitch synchronous segments just defined above in Section 8.5.2.3.

8.5.3 Pitch Prediction and Pitch-Synchronous Segmented Quantization

In this method, the quantizer of the APC residual uses for each pitch-synchronous segment: (1) a different gain G , which is computed from the second residual for the same pitch-synchronous segment; and (2) in general, a different number of bits/sample. For those frames for which the uniform bit allocation is chosen (see Section 8.5.2.1), the time-synchronous segmented quantization method is used.

The resulting pitch-adaptive method, denoted by PA-PP-SQ, was found to provide a significant increase in perceived speech quality and S/Q ratio over the PA-PP method. For the 1-tap, 4-segment case with the fixed bit allocation $\{3,2,1,2\}$, the PA4-PPI-SQ4 method produced a S/Q ratio of 19 dB, which is 1.9 dB higher than the S/Q ratio of the PA4-PPI method. As another interesting comparison, the 1-tap pitch prediction alone produced a S/Q ratio of only 15.3 dB.

We performed several experiments to investigate various aspects of the PA-PP-SQ method. The results of all but two experiments are stated below briefly, followed by a discussion of the other two experiments.

1. Transmission of $M/4$ (for the 4-segment case) instead of M , to save bits required for pitch transmission, produced perceivable speech-quality distortions.

2. Transmission of the location parameter L using 3 bits instead of 2 bits, for the 4-segment case, did not yield any perceivable improvement in the speech quality. (Note that by increasing the accuracy of L , the segmentation procedure allows the frame to begin with a less-than-full segment.)
3. For an average of about 2 bits/sample, out of the several 3-, 4-, and 5-segment cases we tested, we found that the 4-segment case with the nonuniform bit allocation $\{3,2,1,2\}$ produced the highest S/Q ratio.
4. No perceivable speech-quality degradation resulted when we combined, for the purpose of segmented quantization, segments within a pitch period with the same number of bits per sample. This means, for the example with the bit allocation $\{3,2,1,2\}$, only 3 delta gains need be transmitted. We denote this case explicitly with the notation PA4-PPI-SQ3, for example.
5. Delta gains can be coded using 1 bit for the 1-bit and 2-bit segments, and using 2 bits for segments with larger number of bits/sample.

Comparison with Time-Synchronous Segmented Quantization: In this experiment, we used time-synchronous segmented quantization for all frames rather than only for frames using uniform bit allocation. The segments that this method considers for amplitude normalization are different from those that the above PA method considers, for nonuniform bit allocation. The results of our experiments showed that for the 4-segment case, pitch-synchronous SQ was better than time-synchronous SQ. The former method produced, as noted above, an S/Q ratio of 19 dB. To produce the same value, the number of segments for the time-synchronous method had to be increased to 10.

Adaptive Bit Allocation: In this experiment, we used an adaptive, rather than a fixed, bit allocation, to track frame-by-frame variations in the residual amplitude envelope. A two-bit code was transmitted to indicate which of the following four bit-allocations the current frame employed: $\{3,2,2,1\}$, $\{3,2,1,2\}$, $\{3,1,2,2\}$, and $\{2,2,2,2\}$. The adaptive scheme produced only a slight improvement in perceived speech quality (and about the same S/Q ratio) over the fixed case. We feel that the added complexity is not justified by the small improvement.

8.6 Segmented Quantization with Bit Allocation

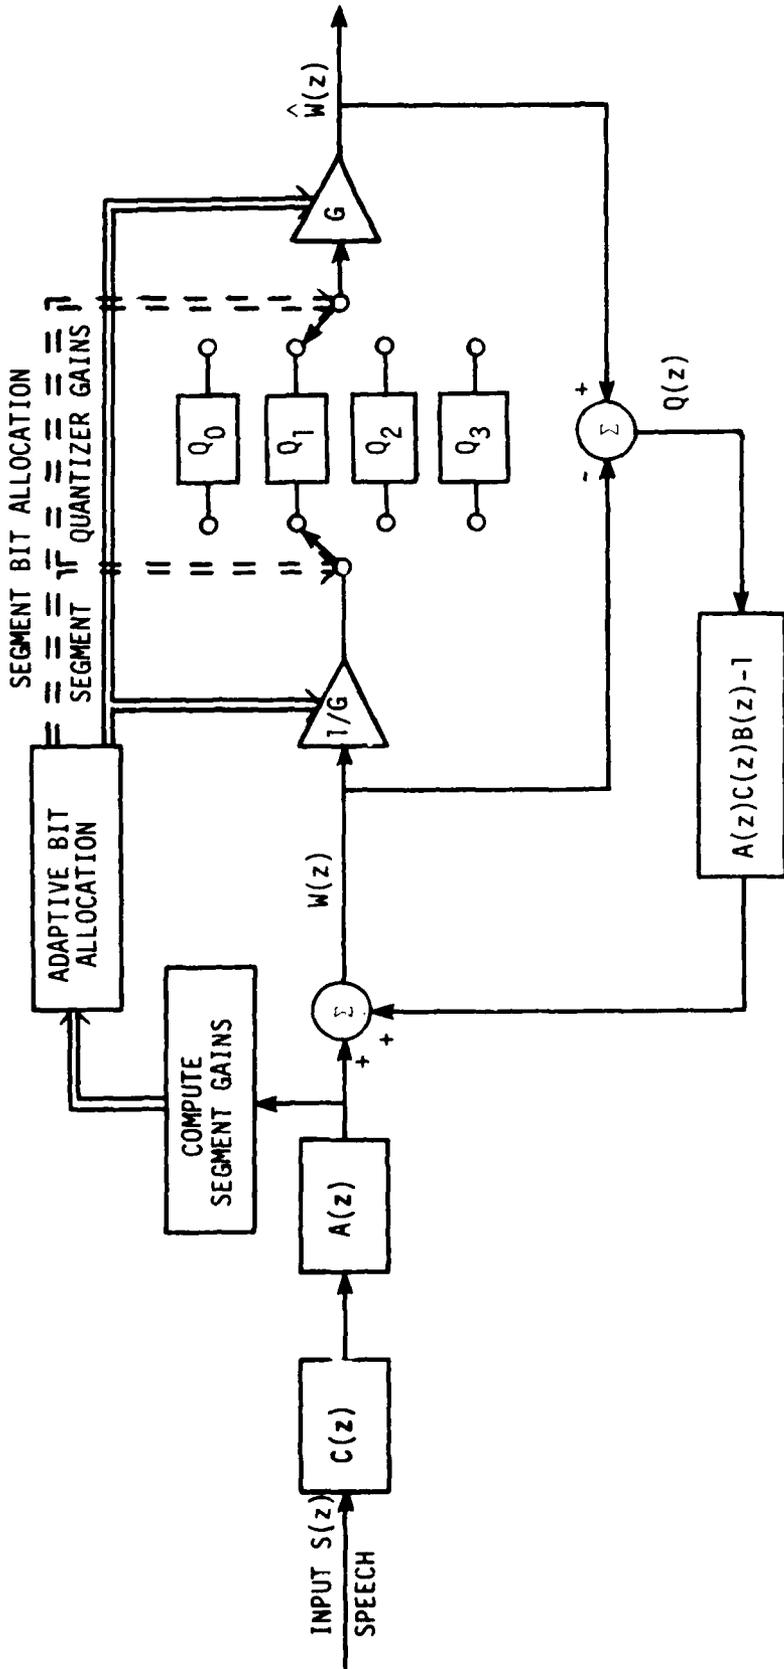
8.6.1 The Method

The SQ-BA method employs a nonuniform quantizer that uses a variable number of bits/sample. This method divides each frame into a fixed number, N , of equal-length segments. The samples in segment i are quantized using B_i (an integer) bits. The bit allocation to be used $\{B_i, 1 \leq i \leq N\}$ is coded and transmitted each frame as side information. A set of optimal n nonuniform unit-variance (Laplacian) quantization tables are stored in memory, where n is the number of distinct values of B_i . In the coding and decoding of the residual samples in the i -th segment, the quantizer gain corresponding to that segment and the quantization

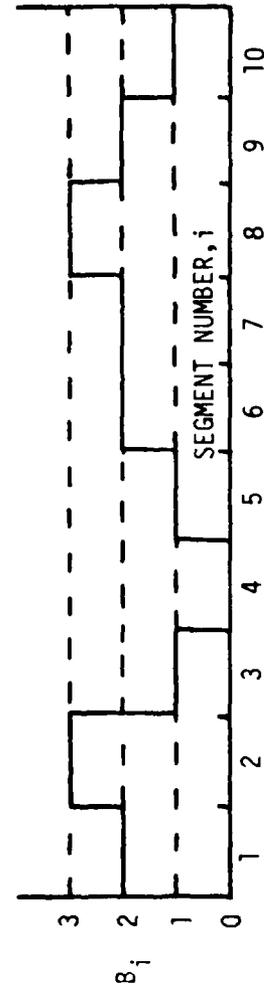
table corresponding to B_i are used. For the case $B_i=0$, all samples in that segment are decoded as zero. Thus, we have combined adaptive bit allocation with segmented quantization, using the same segments for both operations. As discussed below, the quantizer gain of each segment can be obtained from the bit allocation $\{B_i\}$ and the overall gain of the frame. Therefore, a single set of codewords is used to transmit both the bit allocation and the segment gains.

Figure 12(a) depicts an APC-NF system using the SQ-BA method (with pitch prediction). In the figure, Q_i indicates an i -bit quantizer, and the segment bit allocation indicated by the dashed double lines controls the choice of the quantizer via the switch arrangement shown. Figure 12(b) shows the bit allocation for the 10 segments for a frame.

The bit allocation in each frame is chosen to minimize the mean-square quantization error under the constraint that the total number of bits per frame be equal to a given value. The method we have used for determining such an optimal bit allocation is similar to the one used in ATC [6]. Briefly, the optimal allocation is $B_i = B_0 + (P_i - P)/S$, where B_0 is the average number of bits per sample, P_i is the gain G_i in dB of segment i , P is the geometric mean of the segment gains, (i.e., the average of P_i , $1 \leq i \leq N$), and S is a constant equal to 6 (dB/bit) for fine



(a) Block diagram of the APC-NF system using the SQ-BA coding scheme.



(b) Example bit allocation for a frame, for the 10-segment case.

FIG. 12. Segmented quantization with bit allocation.

uniform quantization. For coarse or nonuniform quantizers, the value of S must be chosen experimentally. The set of B_i 's must be rounded to integers and still satisfy the constraint on the total number of bits per frame. To do this, a simple iterative algorithm is used, which typically requires only a small number of iterations. A step-by-step description of the bit-allocation algorithm is given in Fig. 13. Notice that Steps 9 and 10 in this figure give the expressions for the segment gains to be used in the amplitude normalization of the segment residual samples. Since the bit allocation $\{B_i, 1 \leq i \leq N\}$ and the quantized frame gain \hat{P} are transmitted to the receiver, the segment gains \hat{G}_i are also computed at the receiver from the same expressions.

Below, we make several remarks comparing the SQ-BA method just described and the other coding methods. First, the SQ-BA method is different from the SQ method in at least two ways: (1) a variable number of bits/segment rather than a fixed number and (2) transmission of the segment gains via the bit allocation rather than via the delta gains.

Second, although the SQ-BA method uses a bit-allocation procedure similar to the one used in ATC, there are significant differences between the two approaches. The question of time-domain versus frequency-domain coding has already been discussed in Section 8.5.1. With the algorithm used in ATC, the segment

1. $P_i = 20 \log_{10}(G_i), 1 \leq i \leq N.$
2.
$$P = \frac{1}{N} \sum_{k=1}^N P_k$$

$$\Delta B = 0$$
3. $P = P + \Delta B \cdot S$
4. $\hat{B}_i = B_0 + (P_i - P)/S$
5. $\hat{B}_i = [B_i]$ (truncation)
6.
$$\Delta B = \frac{1}{N} \sum_{k=1}^N \hat{B}_k - B_0$$
7. If $|\Delta B| > \epsilon$ (tolerance), go to Step 3
8. Quantize P to \hat{P}
9. $\hat{P}_i = \hat{P} + S(\hat{B}_i - B_0)$
10. $\hat{G}_i = 10^{\hat{P}_i/20}$

FIG. 13. A step-by-step description of the bit-allocation algorithm used in the SQ-BA method.

gains would be coded and transmitted directly; the bit allocation would then be determined from the set of decoded segment gains by means of the algorithm, which would be used in both the transmitter and the receiver. In the SQ-BA method described above, the bit allocation is computed only at the transmitter from the unquantized segment gains and explicitly transmitted to the receiver. Segment gains to be used in both the transmitter and the receiver are computed from the bit allocation as explained above. The difference between these two approaches is apparent in the presence of channel bit-errors. In the ATC approach, a single bit-error on a segment gain may cause errors in any or all of the B_i 's computed at the receiver. Thus, the single bit-error can lead to erroneous decoding of all the samples in that frame. In the approach we have used, a single bit-error causes erroneous decoding of only the samples in the corresponding segment and in the segments that follow.

Finally, we have developed the SQ-BA method (1) as an alternative to the entropy coding method in terms of producing both less computational complexity and potentially better channel-error performance and (2) as an alternative to the pitch-adaptive method in terms of providing an easier implementation on the MAP-300 array processor. The data-dependent segment sizes and locations, among other things, make the PA method extremely difficult to implement on the MAP.

8.6.2 Experimental Results

The SQ-BA method without pitch prediction produced noticeable roughness in the output speech; the roughness was eliminated when we added pitch prediction. The effect of the number of pitch filter taps for the PP-SQ-BA system on speech quality was found to be the same as for the PP-SQ system. Under the 16 kb/s constraint, 3-tap pitch prediction produced better speech quality than 1-tap prediction. We briefly experimented with the choice of the possible values for the number of bits/sample B_i . Considering 4 values, we tested the two sets: $\{0,1,2,3\}$ and $\{1,2,3,4\}$. The first set produced significantly higher S/Q ratios than the second set, under all conditions we tested. We also investigated the use of only 2 values $\{1,2\}$ for B_i . This case produced lower S/Q ratios than the above two cases.

If 10 segments are used, the average number of bits/sample B_0 can be any multiple of 0.1 (e.g., 1.8, 1.9, etc). This flexibility gives a wide choice of systems for investigating the tradeoff involving frame size, number of taps, and B_0 (see Section 10.5 for more results). In one test, we employed 5-tap pitch prediction and found that at 16 kb/s it produced similar quality to that from the 3-tap case. In another test, we observed that using a frame size of 27 ms or larger caused

AD-A096 091

BOLT BERANEK AND NEWMAN INC CAMBRIDGE MA
DESIGN AND REAL-TIME IMPLEMENTATION OF A ROBUST APC CODER FOR S--ETC(U)
DEC 80 R VISWANATHAN, W RUSSELL, A HIGGINS

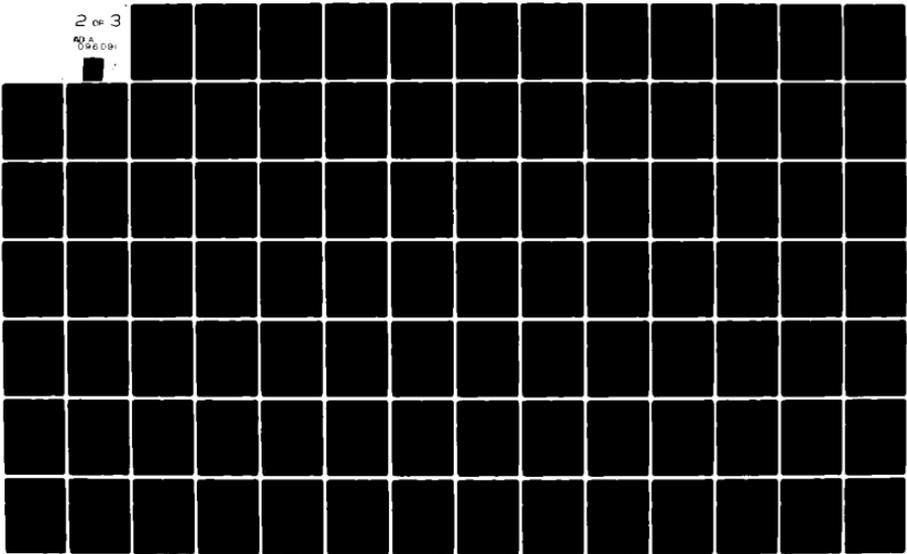
F/S 17/2.1
DCA100-79-C-0037

ML

UNCLASSIFIED

2 of 3

NO. 091



"clicks" and "dropouts" in the perceived speech. These degradations were mitigated but not eliminated by increasing the number of segments per frame.

We experimented with different types of nonuniform quantizers. We found that using a Gaussian quantizer produced speech of similar quality but a slightly lower S/Q ratio than using a Laplacian quantizer. In another experiment, we used a Gaussian 1-bit quantizer, a Laplacian 2-bit quantizer, and a gamma 3-bit quantizer. This resulted in lower speech quality but a slightly higher S/Q ratio, as compared to the case of all Laplacian quantizers. We used the original scheme using Laplacian quantizers in our subsequent work. For Laplacian quantizers, the value $S=4.0$ dB/bit was found to give the best results. (The parameter S is used in Steps 4 and 9 shown in Fig. 13.)

In all our tests of the PP-SQ-BA method, we found that it produced S/Q ratios approximately equal to or less than those from the PP-SQ method under similar conditions, although the speech quality produced by the former method was similar to or sometimes better than that given by the latter method. While we have not thoroughly investigated this issue of unexpectedly low S/Q ratios, we feel that a possible explanation for its cause is that the bit allocation computed from the second residual does

Report No. 4565

Bolt Beranek and Newman Inc.

not adequately match the APC residual, since the two residuals can have different segmental amplitude variations.

9. LIMIT-CYCLE BEHAVIOR OF THE APC SYSTEM

9.1 Discussion of the Problem

As reported in Chapter 5, when the feedback gain G_F of the APC loop is positive, the quantization noise $Q(z)$ builds up to excessive values, causing the quantizer output $\hat{W}(z)$ to exhibit a limit-cycle behavior. Depending on the severity and duration of this behavior, the coder output is perceived as the non-speech-like sounds beeps and glitches or as speech containing discrete distortions e.g., clicks, pops, etc. We noted in Chapter 5 that positive values of G_F are caused by excessive values of the power gain of the filter $F_1 = ABC^{-1}$, inadequate quantization accuracy (which results in low values of the W/Q ratio), or both. In the last chapter, we reported that the basic versions of several of the residual coding methods provide inadequate quantization accuracy and hence cause varying extents of the limit-cycle problem. The various composite coding methods, on the other hand, were found to nearly eliminate the limit-cycle problem, provided the coder uses (1) the P-S prediction sequence and (2) an average number of bits/sample close to 2. Considering the issue of power gain, we reported in Chapter 7 that noise shaping helps to reduce the power gain. In fact, it was found that increasing the amount of pole-zero noise shaping by increasing

the bandwidth parameter w from 200 Hz to 800 Hz eliminated the limit-cycle problem in an entropy-coding APC system, in which, because of its high W/Q ratio, the limit-cycle problem occurs only very infrequently. For all other coding methods we investigated, noise shaping by itself does not provide sufficient reduction in the power gain. In this chapter, we discuss two other methods for reducing the power gain: high-frequency correction (Section 9.4) and preemphasis (Section 9.5).

The limit-cycle behavior can also be triggered if the pitch prediction filter used within the APC loop is unstable. (The spectral filter is always stable since we use the autocorrelation method of linear prediction.) As we reported in Section 8.2, the switched prediction method provides pitch filter stability without perceivably degrading the coder performance.

Another approach towards solving the limit-cycle problem, discussed below in Section 9.3, is to limit or clamp the filtered quantization noise at some value when it is building up. Before we discuss this approach and others mentioned above, we present in the next section our experimental observations regarding the extent of the limit-cycle problem for the two prediction sequences. Throughout our experimental investigation of the limit-cycle behavior of APC, we used several test cases to examine if and how a given approach to cure the problem

controlled the extent of the problem. Some of these test cases are mentioned below as part of our presentation of the experimental results.

9.2 S-P versus P-S Prediction Sequence

At the beginning of this project, we used only the S-P prediction sequence. While the entropy-coding method with a large number of quantizer levels produced no limit-cycle problems, the fixed-length coders that we subsequently implemented caused severe distortions to be perceived in the output speech. For example, the PPl system using 2 bits/sample and 25.5 ms frame size produced loud beeps in the output speech at the rate of one or two per sentence, with each beep lasting about 50 ms. Even with 4 bits/sample, we encountered one beep over 6 sentences of a total duration of about 15 sec. Limit-cycle problems were also encountered with the PP-SQ system.

When we switched to the P-S prediction sequence, we found that for the case of 2 bits/sample, all non-speech-like sounds were eliminated, and the discrete distortions were significantly reduced. The primary reason for this improvement is that using the P-S sequence produced power gain values that were about 1-2 dB lower than those obtained using the S-P sequence. However,

the PP3-SQ10 system with the P-S sequence, a 3-level quantizer, and a frame size of 18 ms still produced beeps at the output.

In summary, the S-P prediction sequence is more likely to encounter the limit-cycle problem than the P-S sequence, since the former yields higher power gain values.

9.3 Saturating Limiter

In this method, the magnitude of the filtered quantization noise is limited to some reasonable value [11]. The APC-NF configuration shown in Fig. 5 serves best to explain the idea of this method and to suggest a way of computing the value of the limit. Since the APC residual $W(z)$ is given by Eq. (9), by limiting the magnitude of the filtered noise $F(z)$, we can ensure that $W(z)$ is not dominated by the quantization noise. Eq. (9) also suggests that the filtered noise samples $f(n)$ may be limited in magnitude to θ times the rms value of the second residual $E_2(z)$. Notice that this rms value is already computed for setting the quantizer gain. As in [11], we used a value of $\theta=2$. We interpret the saturating limiter as serving the role of a "safety valve."

We implemented the limiter first in the APC-NF configuration and investigated its effectiveness in several test cases. The

first case we considered is an SQ10 coder (with no pitch prediction) using 2 bits/sample and 25.5 ms frame size. Without the limiter, this coder produced about 1-2 beeps per sentence. Use of the limiter in this coder was found to entirely eliminate all beeps and other discrete distortions. The S/Q ratio dropped slightly from 16 dB (no limiter used) to 15.8 dB (limiter used). Another interesting case we tested is a PP3-SQ4 coder using the P-S sequence, with 3-level quantization and 19.5 ms frame size. In this case, we found that the limiter was activated for about 0.3% of the filtered noise samples. The S/Q ratio dropped from 14.9 dB to 13.2 dB because of the limiter. Using the limiter in this case eliminated all the beeps and other discrete distortions, but some of the processed sentences sounded objectionably reverberant. The reverberant quality is due to the periodic propagation by the filter $1/C(z)$ of the clipping "errors" introduced by the limiter. This fact can be easily shown as follows. The limiter's output can be written as

$$F'(z) = F(z) + D(z), \quad (32)$$

where $D(z)$ is the limiter-caused clipping noise. From Fig. 5, we obtain

$$\begin{aligned} W(z) &= E_2(z) + F'(z) \\ &= A(z)C(z)S(z) + [A(z)C(z)B(z) - 1]Q(z) + D(z). \end{aligned} \quad (33)$$

From the receiver shown in Fig. 1, the output speech $R(z)$ is given by

$$R(z) = [1/C(z)][1/A(z)][W(z)+Q(z)]. \quad (34)$$

Using (33) in (34), we obtain

$$R(z) = S(z) + B(z)Q(z) + [1/C(z)][1/A(z)]D(z). \quad (35)$$

Thus, the reconstructed speech has a component that is the clipping noise $D(z)$ filtered by $1/A(z)$ and $1/C(z)$. This explains the cause of the reverberant quality mentioned above. It can be easily seen that the only way to avoid this problem is to avoid placing the limiter in the path of the pitch predictor within the APC loop. Such a limiter placement is possible for the APC-PF system with either of the two prediction sequences and for the APC-HF system with the S-P sequence. The output speech expressions in the two cases are given below:

$$R(z) = S(z) + B(z)Q(z) + D(z), \quad (\text{APC-PF}), \quad (36)$$

and

$$R(z) = S(z) + B(z)Q(z) + [1/A(z)]D(z), \quad (\text{APC-HF}). \quad (37)$$

Notice from (37) that the clipping noise is shaped by the

spectral filter for the APC-HF case. Our experiments showed that the limiter used in the APC-PF configuration did not produce any appreciable effect in terms of alleviating the limit-cycle problem. As expected, using the limiter in the APC-HF system did not produce any reverberant quality, and it reduced the severity of (and in some cases eliminated) the limit-cycle problem.

9.4 High-Frequency Correction

This technique, proposed in [11], is a way of reducing the power gain of the predictor $A(z)$, $G_p(A)$. Notice that $G_p(A)$ is simply the integral of the power spectrum of $A(z)$. This inverse spectrum has large amplitudes at the high frequencies near the cutoff frequency of the anti-aliasing (A/D) lowpass filter; the large amplitudes are primarily due to the nonideal lowpass filters. Reference [11] suggests a simple and effective method of reducing $G_p(A)$. In this method, the high-frequency amplitudes of the power spectrum of the signal used for LPC analysis are increased ("corrected") by digitally adding to that signal a highpass-filtered white noise. Since we use the autocorrelation method of LPC analysis, the high-frequency correction (HFC) method reduces to a simple procedure of modifying the autocorrelation coefficients $R(k)$, $0 \leq k \leq p$, used for LPC analysis as follows:

$$\begin{aligned}\hat{R}(k) &= R(k) + \lambda V_p R(0) \mu(k), \quad 0 \leq k \leq 2, \\ \hat{R}(k) &= R(k), \quad 3 \leq k \leq p,\end{aligned}\tag{38}$$

where

$$\mu(0) = 3/8, \quad \mu(1) = -1/4, \quad \mu(2) = 1/16,$$

and V_p is the normalized error [13] of the linear predictor for the unmodified case, which is computed by solving the autocorrelation normal equations. Thus, the above HFC method requires solving the normal equations twice, once for the unmodified case and once for the modified case. (See Section 16.1.1 for a simplified HFC procedure that we have recommended for the real-time implementation.) Notice from (38) that the parameter λ controls the amount of high-frequency correction. From our experiments, we found that the choice $\lambda=0.05$ suggested in [11] was quite reasonable.

When we used the HFC procedure in the PP and PP-SQ systems using 2 bits/sample and the S-P sequence, the intensity of the beeps heard at the output was reduced, but the beeps were not eliminated. For the 4-level PP-SQ systems using the P-S sequence, adding HFC eliminated several of the discrete distortions. However, for the 3-level PP3-SQ4 system mentioned above in Section 9.3, the use of HFC was found only to reduce the number and intensity of beeps in the output speech. While the

HFC technique by itself does not solve the limit-cycle problem, it is quite effective in reducing the power gain, and yet it gives only a small drop in the S/Q ratio for frames not having excessive power gain.

9.5 Preemphasis

Preemphasizing the input speech with a filter $(1-\beta z^{-1})$ reduces its spectral dynamic range and hence reduces the power gain of $A(z)$. The decoded speech at the receiver is deemphasized using the filter $1/(1-\gamma z^{-1})$. There are several ways of choosing β and γ : β fixed at a chosen value, β adaptively chosen as $-R(1)/R(0)$, where the R 's are the autocorrelation coefficients of the input speech signal, $\gamma = \beta$, $\gamma = 0$ (no deemphasis), $\gamma < \beta$. The adaptive preemphasis method has the effect of maximally reducing the power gain, but, as it employs values of β close to 1 for most voiced sounds, the deemphasized output speech was found to have a significant amount of low-frequency roughness and rumble. We found that a fixed value of $\beta=0.4$ is a good compromise in that when used in conjunction with HFC, preemphasis eliminated all the discrete distortions in the output speech and introduced only a small amount of roughness. The choice of $\gamma < \beta$ was found to reduce the roughness slightly relative to what we perceived for $\gamma=\beta$, but the S/Q ratio was significantly lower for

$\gamma < \beta$. We used $\beta = \gamma = 0.4$ in all our subsequent experiments involving preemphasis.

Below, we summarize our experimental results on preemphasis:

1. For systems using 3-level quantization (e.g., the PP3-SQ4 system mentioned in Section 9.3), the use of HFC and preemphasis effectively eliminated the beeps and other discrete distortions. Even for 4-level systems, the use of preemphasis reduced certain discrete distortions.
2. With preemphasized input speech, the number of instabilities of a multi-tap pitch filter (used in the P-S sequence) was reduced.
3. The S/Q ratio was reduced by as much as 1 dB.
4. The output speech was perceived to be slightly rough.
5. When using preemphasis, as mentioned in Section 7.4, pole-zero noise shaping produced better speech quality than the all-pole method, which produced better speech quality than the 1-zero method.

6. Use of preemphasis in the PP-SQ-BA method (see Section 8.6) produced noticeable degradations at the output.

10. OPTIMIZATION OF CODERS FOR ERROR-FREE CHANNELS

Although the final goal of this work has been to develop a robust APC coder for use over noisy channels, initially we conducted the speech-quality optimization study for error-free channels to investigate the speech quality that the various coding methods are capable of producing at 16 kb/s, without the burden of the error protection bits. Also, we felt that parameter tradeoff relations obtained in this study could be used in narrowing the range of parameter values to investigate in the subsequent optimization study for noisy channels. The results of the study reported in this chapter and the recommendations given in Section 10.6 should be useful in the design of 16 kb/s systems for speech communication over error-free channels.

As we explained in the preceding chapters, there are several residual coding methods and noise shaping methods to choose from, and there are several parameters that affect the performance of the APC system. Important among these parameters are: input sampling rate F_S , frame size (or transmission frame rate of coder parameters), number of quantization levels per residual sample, LPC order p , number of pitch-predictor taps, and parameters required by individual residual coding and noise shaping methods (e.g., number of segments). Parameters such as bandwidth w used

in noise shaping are not transmitted and, therefore, are not considered in the bit-rate tradeoff study. We reported our choice of $FS=6.67$ kHz in Section 2.2. When we investigated the choices $p=6,8$ and 10 , we did not find any perceivable difference between the three cases. We decided to continue the use of $p=8$. The bit allocation for the various coder parameters is another dimension that affects transmission rate. As reported in earlier chapters, we chose a bit allocation that produced good results. Table 5 summarizes the bit allocation for the different parameters. Notice that individual APC systems use different subsets of parameters given in Table 5. Below, we report the results of our optimization study involving the remaining parameters, separately for each of the four coding methods. Since this study was conducted before we successfully resolved the limit-cycle problem, we used the P-S prediction sequence, without preemphasis, HFC, and limiter. Subsequent to the work reported in the last chapter, we ran the optimized P-S coder designs but with the S-P sequence. The results of this experiment are reported in Section 10.5. Informal listening tests were used to make all quality judgments reported in this chapter.

<u>Parameter Name</u>	<u>Bit Allocation</u>	<u>Remarks</u>
8 LARs	33	see Tables 1 and 2
Pitch	7	no quantization needed
Quantizer gain	10 6	for entropy coding for other methods
Pitch taps	4 10 16	for 1-tap case for 3-tap case for 5-tap case
Delta gains	2 each	see Table 4
Location parameter	2	see Section 8.5.2
Segment bit allocation	2 each	see Section 8.6

TABLE 5. Bit allocation for various coder parameters

10.1 Entropy Coding with Pitch Prediction

For entropy-coded systems, we conducted a tradeoff study involving a total of 9 coders, obtained from three values of frame size (19.5, 25.5 and 30 ms) and three conditions of pitch prediction (no pitch prediction or 0-tap, 1-tap, and 3-tap). For each of the 9 coders, we used 1-zero noise shaping and the variable-to-fixed rate conversion algorithm to adjust the quantizer step size to produce a 16 kb/s data rate. The S/Q ratios obtained for these 9 systems are given in Table 6. For each frame size, we preferred the 3-tap system over the 0-tap and

1-tap systems. For the three 3-tap systems, the average number of bits/sample used was found to be 1.9, 2.01 and 2.07, respectively. Two 3-tap systems, one with 25.5 ms frame size and the other with 30 ms frame size, were found to produce the best overall speech quality. We investigated these two systems further, using pole-zero noise shaping with various values of the bandwidth parameter w . We found that the choice $w=800$ Hz produced the best perceptual result, which was much better than what we obtained with 1-zero noise shaping. The two 3-tap systems with 800 Hz pole-zero noise shaping produced S/Q ratios of about 19.0 dB (25.5 ms) and 19.3 dB (30 ms). We found the second 3-tap system (30 ms) to produce marginally better speech quality than the first (25.5 ms). For comparisons with other optimized systems, therefore, we chose the EC-PP3 system with a frame size of 30 ms.

No. of Pitch Taps	Frame Size (ms)		
	<u>19.5</u>	<u>25.5</u>	<u>30</u>
0	19.5	19.9	20.0
1	19.7	20.8	21.2
3	20.0	21.2	21.6

TABLE 6. S/Q ratios for the 9 entropy-coding APC systems, all using 1-zero noise shaping

10.2 Pitch Prediction and Segmented Quantization

In Section 8.3.3, we reported the results of an experiment involving 9 PP-SQ systems (three values each of the number of pitch taps and the number of segments), all using a frame size of 25.5 ms. Although these systems did not have the same bit rate, the individual bit rates were close to 16 kb/s. The two systems PP3-SQ5 and PP3-SQ10 produced the best overall speech quality. We then conducted a tradeoff study involving the three 16 kb/s systems: PP3-SQ5 (25.5-ms frame), PP3-SQ10 (29.25-ms frame), and PP1-SQ4 (22.5-ms frame). Comparative speech quality evaluations indicated that the PP3-SQ10 system was the best. This system produced an average S/Q ratio of 18.2 dB with the use of the 1-zero noise shaping. The other methods of noise shaping did not produce any perceivable speech quality improvement.

10.3 Pitch-Adaptive Coding With Pitch Prediction and Segmented Quantization

For pitch-adaptive coding, we considered three 16 kb/s systems trading off frame size, average number of bits/sample, and number of pitch taps: (1) PA4-PP3-SQ3, bit allocation = {3,2,1,2} with the average being 2 bits, and 25.5-ms frame; (2) PA10-PP1-SQ3, bit allocation {3,3,3,2,2,2,2,1,1,2} with the

average being 2.1 bits, and 30-ms frame; and (3) PA5-PP3-SQ3, bit allocation = {3,2,1,1,2} with the average being 1.8 bits, and 19.5-ms frame. Since we combined segments with the same bit allocation for the purpose of segmented quantization, we transmitted only 3 segment gains in each of these three coders (see Section 8.5.3). Informal listening tests indicated that the first coder produced the best overall speech quality. This coder produced an average S/Q ratio of 18.8 dB with the use of 1-zero noise shaping.

10.4 Segmented Quantization With Bit Allocation and Pitch Prediction

For the PP-SQ-BA system, we investigated frame sizes less than 27 ms (see Section 8.6.2), number of pitch taps equal to 1, 3, or 5, number of segments up to 10, and various segment bit allocations. From this investigation, we found that the PP5-SQ10-BA system with a frame size of 25.5 ms and bits/sample from the set {0,1,2,3} produced the best overall speech quality. This system used an average number of bits/sample $B_0=1.9$ and produced an average S/Q ratio of about 14.8 dB with 1-zero noise shaping.

10.5 The S-P Prediction Sequence

We re-investigated the above-described optimized coders using the S-P prediction sequence and the APC-HF configuration. For the EC-PP3 and the PP5-SQ10-BA systems, we used HFC, and a limiter. For the PP3-SQ10 and the PA4-PP3-SQ3 systems, we used preemphasis, HFC, and a limiter. In all the cases, we used pole-zero noise shaping.

By and large we obtained about the same speech-quality performance from the S-P sequence as was observed using the P-S sequence. However, there are three noteworthy differences between the systems resulting from the two prediction sequences. These differences, reported in the preceding chapters, are summarized below:

1. Since the multi-tap pitch filter was found to be stable for the S-P sequence, the switched prediction method, which requires checking the stability of the pitch filter and which is necessary for the P-S sequence (see Section 8.2), is not needed for the S-P sequence.
2. With pole-zero noise shaping, the S-P sequence leads to a simple implementation, since $A(z)B(z)$ reduces to $A(z/\alpha)$ (see Section 7.3).

3. With the use of the S-P sequence and the APC-HF configuration, a limiter can be effectively used in the noise-feedback path as a worthwhile precaution (see Section 9.3).

10.6 Comparative Evaluation and Recommendations

We conducted informal listening tests to compare the four optimized coders: (1) EC-PP3, (2) PP3-SQ10, (3) PA4-PP3-SQ3, and (4) PP5-SQ10-BA. The output speech from the systems (2)-(4) was slightly better than from the system (1), but occasionally it contained low-level discrete distortions. In contrast, the entropy-coded system produced relatively smooth speech. While the perceived speech-quality differences among the four systems were small, three subjects, upon careful listening, rated them in the following order from best to worst: EC-PP3, PA4-PP3-SQ3, PP5-SQ10-BA, PP3-SQ10. From the viewpoint of minimizing computational complexity, the ordering of these systems is just the reverse of the above order, with EC-PP3 being the most complex system (because of its variable-to-fixed rate conversion).

Comparing the output of each of the four systems against the input speech using high-quality headphones with good low- and

high-frequency response, we found that the EC-PP3 system produced speech closest, but not identical (or transparent), to the natural speech.

For 16 kb/s speech communication over error-free channels, if computational complexity is not an issue, we recommend the use of the EC-PP3 system. As for the specific configuration to use, from the observations given in Section 10.5, we recommend the implementation of this system using the S-P prediction sequence, APC-HF configuration, high-frequency correction, and a limiter. For noisy-channel applications, as will be seen in Chapter 12, both the recommendations have to be modified.

11. OPTIMIZATION OF ERROR-PROTECTED CODERS

In this chapter, we present the results of our parameter optimization study, performed in the absence of channel errors, of 16 kb/s APC systems in which bits have been allocated for the error protection of coder parameters. The performance of the error-protected APC systems in 1% channel error is the topic of the next chapter. The objective of our work reported in this and the next chapter was to develop a robust 16 kb/s APC coder to operate over channels with bit-error rates of up to 1%. To meet this objective, we experimentally optimized (1) the tradeoff between the voice data rate and the error-protection data rate, and (2) the amount of error protection for individual transmission parameters. In this chapter, we present several error-protected APC systems for investigating the tradeoff (1) above. We did not protect the coded APC residual. To partially protect important parameters of the coder, we used the Hamming (7,4) code, which protects 4 data bits by adding 3 parity bits; this code detects and corrects all single bit-errors in the resulting 7-bit codeword. In the APC systems reported below, the number of protected bits per frame varies from 28 to 68. Our choice of this moderate to substantial protection of the side-information data was based on our previous experience [3]. We conducted informal listening tests to compare these error-

protected APC systems in the absence of channel errors. The results are presented below separately for the three residual coding methods discussed in Chapter 8. We did not include the pitch-adaptive coding method in the channel-error study, since we felt that its implementation on the MAP would be extremely difficult. For the systems described in this chapter, we used the P-S prediction sequence.

11.1 Entropy Coding with Pitch Prediction

To study the tradeoff between frame size and the number of pitch-filter taps, we considered four systems: EC, EC-PP1, EC-PP3, and EC-PP5. The details of these systems are given in Table 7. Notice that Items 3 and 4 given in the table are both EC-PP3 systems using different noise shaping methods. We used high-frequency correction with $\lambda=0.05$ for all five systems in Table 7. Notice that each of these systems protects a relatively large number of parameter bits. To obtain a 16 kb/s data rate, we used the variable-to-fixed rate conversion algorithm described in Section 8.4.2. Table 7 also gives the average bits/sample used by the individual coders. Systems 1-3 in Table 7 use 1-zero noise shaping. Among these three systems, we found that the EC-PP3 system (System 3) produced the best overall speech quality. When we used 800 Hz pole-zero noise shaping for the EC-PP3

system, the S/Q ratio dropped from 20.1 dB (for System 3) to 17.7 dB (for System 4), but the perceived background noise was significantly reduced relative to the 1-zero case. Comparing the 3-tap system (System 4) with the 5-tap system (System 5), we found that the former system produced slightly better speech quality.

No.	System	Frame Size (ms)	Parameter Bits		Avg. Bits Per Sample	Noise Shaping	S/Q Ratio (dB)
			Total	Protected			
1	EC	25.5	43	32	1.97	1-zero	18.6
2	EC-PP1	27.0	54	40	1.89	1-zero	19.5
3	EC-PP3	30.0	60	44	1.90	1-zero	20.1
4	EC-PP3	30.0	60	44	1.90	pole-zero (800 Hz)	17.7
5	EC-PP5	30.0	66	48	1.86	pole-zero (800 Hz)	17.6

TABLE 7. 16 kb/s error-protected entropy-coding systems

11.2 Pitch Prediction and Segmented Quantization

For the tradeoff study involving LPC order, number of pitch-filter taps, frame size, number of quantizer levels per residual sample, number of segments, and number of bits protected, we considered eight 16 kb/s PP-SQ systems given in Table 8. For all eight systems, we used preemphasis ($\beta=0.4$), high-frequency

correction ($\lambda=0.05$), and pole-zero noise shaping ($w=800$ Hz). The first two systems use 3-level quantization and provide substantial protection of parameters as in the entropy-coding systems considered above. Of these two systems, the PP5-SQ10 system was found to produce speech with more clarity. Notice that these two 3-level systems require block-coding of residual samples; five residual samples are coded using 8 bits. Therefore, a single bit-error causes wrong decoding of five samples. From our experience with the design of the 9.6 kb/s baseband coder [3], we anticipated that such a block-coding method would result in poor channel-error performance. Therefore, we considered six systems (Systems 3-8 in Table 8) using 4-level quantization. In choosing these six systems, we varied the different coder parameters and provided, for the ratio of the number of protected bits to the total number of parameter bits per frame, a range of values from 28/58 (System 3) to 44/56 (System 8). Of these six systems, we found that the PP1-SQ4 system had the highest level of background noise. Comparing the four systems with 33.75 ms frame size, we noted that the speech from the PP3-SQ2 system (System 4) was somewhat rougher than from the other three systems, and that these latter three systems (Systems 5,6, and 8) produced about the same speech quality. The speech from System 7 lacked the clarity produced by the other systems with smaller frame sizes, which indicated that the update

rate for the parameters in that system was not adequate. Therefore, Systems 5, 6, and 8 are our preferred choices of the 4-level systems. Of these three, System 8, which is the 6-pole PP3-SQ3 system, provides the largest protection of parameter bits. When we compared this 4-level system with the 3-level system PP5-SQ10 (System 2), we found that the 4-level system yielded slightly more natural-sounding speech than the 3-level system. We note that the low S/Q ratio values for the systems in Table 8 are primarily because of the noise shaping used. (Without any noise shaping, the S/Q ratio of System 8 was found to be 17.3 dB.)

No.	System	LPC Order (p)	Frame Size (ms)	No. of Levels/ Sample	Parameter Bits		S/Q Ratio (dB)
					Total	Protected	
1	PP3-SQ4	8	19.5	3	64	52	13.9
2	PP5-SQ10	8	25.5	3	82	68	12.9
3	PP1-SQ4	8	30.0	4	58	28	13.2
4	PP3-SQ2	8	33.75	4	60	36	13.4
5	PP3-SQ3	8	33.75	4	62	36	13.6
6	PP3-SQ4	8	33.75	4	64	32	13.8
7	PP3-SQ3	8	36.6	4	62	44	12.9
8	PP3-SQ3	6	33.75	4	56	44	13.8

TABLE 8. 16 kb/s error-protected PP-SQ systems

11.3 Segmented Quantization With Bit Allocation and Pitch Prediction

The results reported above for the PP-SQ system simplified the problem of choosing the PP-SQ-BA system(s) appropriate for the channel-error study. Recall from Section 8.6.2 that frame sizes larger than 27 ms lead to perceivable distortions in the output speech of the PP-SQ-BA system. Based on these considerations, we chose for the channel-error study the PP3-SQ10-BA system with frame size = 25.5 ms, LPC order $p = 6$, segment bit allocation $\{0,1,2,3\}$, and B_0 =average bits/sample=1.7. This system protects 64 bits out of a total of 70 bits/frame of side-information data. Using high-frequency correction ($\lambda=0.05$), no preemphasis, and 400 Hz pole-zero noise shaping, we obtained an average S/Q ratio of 13.8 dB.

11.4 Comparative Evaluation

We compared the three error-protected systems in the absence of channel bit-errors: EC-PP3, PP3-SQ3 (System 8 in Table 8), and PP3-SQ10-BA. We found different types of distortions in the output speech from the three systems. The EC-PP system had the highest level of background noise, but it produced a more-pleasing smooth speech. The PP-SQ and PP-SQ-BA systems had

"choppy" background noise, but their output speech was perceived as more natural than that of the EC-PP system. The PP-SQ-BA system had a "scratchy" quality. Despite these differences, we felt that the three systems were of roughly equivalent overall quality. The final choice of a robust APC system should be determined only after comparing the channel-error performance of these three systems and others described above in this chapter.

12. EVALUATION OF ERROR-PROTECTED CODERS IN 1% CHANNEL ERROR

One of the requirements of this project has been to design a robust 16 kb/s APC system that tolerates adequately channel bit-error rates of up to 1%. We used the following engineering criterion suggested by the COTR as a specific measure of the extent of robustness required of the final APC system design: The speech quality of the error-protected 16 kb/s coder at 1% channel error should be about the same or better than the speech quality of the same coder when it is operated without error protection in 0.1% channel error. Notice that the unprotected, engineering-criterion system will have a bit rate less than 16 kb/s, since it is obtained by discarding the error-protection bits of a 16 kb/s coder.

In the last chapter, we considered the tradeoff between the voice data rate and the error-protection data rate. In this chapter, we present the results of our work on the following issues: (1) distribution of the allocated error-protection bits among individual transmission parameters (Sections 12.2-12.4); (2) selection of the coder having the best channel-error performance for each of the three coding methods considered (Sections 12.2-12.4); (3) effect of the prediction sequence on the channel-error performance of a coder (Section 12.5); (4)

comparative evaluation of the optimized coders and choice of the most robust 16 kb/s coder (Section 12.6); (5) effect of using the so-called folded binary code for encoding the residual (Section 12.7); and (6) investigation of the performance of the robust system over higher-error-rate channels (Section 12.8). Before we proceed to present the results on these topics, we provide in Section 12.1 a brief description of our channel-error simulation.

12.1 Channel-Error Simulation

In our simulation, we used the binary symmetric channel in which independent, identically distributed random errors are introduced into the transmitted bit stream. A bit error simply changes the state of the bit from 0 to 1 or from 1 to 0. Our simulation system permits the user to vary both the bit-error rate and the amount of error protection separately for each parameter. We used this feature to (1) examine how each transmitted parameter, when subjected to 1% channel error, would independently affect the output speech and (2) investigate, as a diagnostic tool, the cause of the perceived distortions in the output speech. In general, we found that channel bit-errors on the (unprotected) APC residual samples caused the output speech to have a continuously rough or "raspy" character and a reverberant quality. In contrast, uncorrected bit-errors on the

side-information data caused discrete distortions such as pops and clicks in the speech. Specific results are given in the following sections.

12.2 Entropy Coding with Pitch Prediction

For the entropy-coding systems, as mentioned in Section 8.4.1, we used the self-synchronizing code with codewords 0,10,110,etc. Clearly, one bit-error in a codeword can cause one of two decoding errors: splitting a sample into two samples, or merging two samples into one. This will cause the total number of decoded samples in a frame to be larger or smaller than the desired, fixed number. In view of a requirement stemming from the real-time implementation on the MAP, we chose to discard samples at the end of the frame if a larger number of samples were decoded and to append zero samples to the end of the frame if a smaller number of samples were decoded.

Initially, we conducted our experiments using the EC-PP3 system reported in Section 11.1, to determine the amount of protection required by the individual parameters and to understand the source of each of the different distortions perceived in the output speech. We found that the amount and the specific distribution of error protection of parameter data of

the EC-PP3 system given in Table 9 was quite effective in coping with 1% channel bit-errors; when the residual samples were not subjected to channel error, the performance of this coder was found to be approximately invariant as the bit-error rate on the parameters was varied from 0% to 1%. However, with the introduction of 1% bit-errors on the residual samples, we found that the output speech had a "ringing" or reverberant quality.

In our subsequent experiments, we compared the 1% channel-error performance of the entropy-coding systems reported in Section 11.1. The specific error-protection allocations used for these systems are given in Table 9. Listening tests showed that the ringing or reverberant quality was substantially worse for the EC-PP1 system than for the EC-PP3 system. The EC-PP5 system produced slightly less reverberant quality than the EC-PP3 system did but caused perceivable distortions such as pops. The EC system (which does not use pitch prediction) did not produce a reverberant quality but exhibited a continuously rough or raspy character, which degraded the speech quality substantially. The output speech of the EC system sounded almost like whispered speech, without proper pitch periodicity. From the results of these tests, we make the following two conclusions: (1) Although using pitch prediction causes the output speech to sound reverberant, it yields significantly better overall speech

Parameter		EC	EC-PP1	EC-PP3	EC-PP5
Quantizer gain		10(9)	10(9)	10(9)	10(9)
Pitch			7(6)	7(6)	7(6)
Taps	c(M-2)	--	--	--	3(2)
	c(M-1)	--	--	3(2)	3(2)
	c(M)	--	4(3)	4(3)	4(3)
	c(M+1)	--	--	3(2)	3(2)
	c(M+2)	--	--	--	3(2)
Log Area Ratios	1	6(5)	6(5)	6(5)	6(5)
	2	5(4)	5(4)	5(4)	5(4)
	3	4(3)	4(3)	4(3)	4(3)
	4	4(3)	4(3)	4(3)	4(3)
	5	4(3)	4(3)	4(3)	4(3)
	6	4(3)	4(2)	4(2)	4(2)
	7	3(2)	3(2)	3(2)	3(2)
	8	3	3	3	3
Error Protection: Total protected Cost		(32) 24	(40) 30	(44) 33	(48) 36
Sync		1	1	1	1
Total bits/frame		68	85	94	103

TABLE 9. Error-protection allocations used for four entropy-coding systems. Numbers given within parentheses indicate the number of the most significant bits protected.

quality than the scheme without pitch prediction; and (2) 3-tap pitch prediction produces the best overall speech quality. Combining these results with the results reported in Section 11.1 for the 0% error case, we chose the EC-PP3 system as the most robust 16 kb/s entropy-coding APC system.

12.3 Pitch Prediction and Segmented Quantization

Although we tested in channel error several of the PP-SQ coders described in Section 11.2, we present below the results for the two interesting systems: 3-level PP5-SQ10 (System 2 in Table 8) and 4-level PP3-SQ3 (System 8 in Table 8). Both systems provide for substantial error protection of side-information data. Table 10 gives the amount of error protection we used for individual parameters in each of the two systems. We obtained this error protection allocation among parameters through several experiments. We found that full protection of the 2-bit segment gains was necessary to reduce unpleasant pops and clicks.

Output speech from the PP5-SQ10 system operating in 1% channel error contained discrete distortions and had a continuously rough and reverberant quality. To examine the extent to which these quality degradations were caused by the block-coding of the residual samples used in this 3-level system,

Parameter		PP5-SQ10	PP3-SQ3	PP3-SQ10-BA
Quantizer gain		6(6)	6(6)	6(6)
Pitch		7(7)	7(7)	7(7)
Taps	c(M-2)	3(2)	--	--
	c(M-1)	3(2)	3(2)	3(2)
	c(M)	4(3)	4(3)	4(3)
	c(M+1)	3(2)	3(2)	3(2)
	c(M+2)	3(2)	--	--
Segment	Delta gains Bit allocation	10 x 2(20) --	3 x 2(6) --	-- 10 x 2(20)
Log Area Ratios	1	6(5)	6(5)	6(6)
	2	5(4)	5(4)	5(5)
	3	4(3)	4(3)	4(4)
	4	4(3)	4(2)	4(3)
	5	4(3)	4(2)	4(3)
	6	4(3)	4(2)	4(3)
	7	3(2)	--	--
	8	3(1)	--	--
Error protection: Total protected Cost		(68) 51	(44) 33	(64) 48
Sync		1	1	1
Total bits/frame		134	90	119

TABLE 10. Error-protection allocation used for two PP-SQ systems and one PP-SQ-BA system. Numbers given within parentheses indicate the number of the most significant bits protected.

we simulated the same system without block coding (i.e., we used 2 bits to encode the output of the 3-level quantizer). It was noted that the output speech from this latter system was less rough and reverberant. Also, several discrete distortions, observed in the block-coded case, were removed.

The 4-level PP3-SQ3 system, on the other hand, produced speech that was less reverberant and substantially less rough than the PP5-SQ10 system. Any significant reduction in the side-information protection for the PP3-SQ3 system was found to increase the number and the intensity of the perceived discrete distortions in the output speech. Also, recalling from Section 11.2, this PP3-SQ3 system performed at least as well as any other PP-SQ system that we tested in 0% channel error. Therefore, we chose the 4-level, 6-pole PP3-SQ3 system as the most robust PP-SQ system.

12.4 Segmented Quantization with Bit Allocation and Pitch Prediction

For the PP3-SQ10-BA system described in Section 11.3, we protected fully all ten 2-bit codes representing the segment bit allocations, since errors in these codes cause wrong decoding of some or all of the residual samples of a frame. The error protection we chose for other parameters is given in Table 10.

The output speech from this coder in 1% channel error was found to have several discrete distortions and be somewhat reverberant.

12.5 Effect of Prediction Sequence on Channel-Error Performance

The APC systems considered in the above-described channel-error study used the P-S prediction sequence. We investigated the channel-error performance of some of these systems using the S-P prediction sequence. The results of this investigation are given in Section 12.5.1. In an attempt to improve the inferior channel-error performance caused by the S-P sequence, we incorporated at the receiver means of smoothing the decoded residual samples. The results of this study are given in Section 12.5.2.

12.5.1 The S-P Prediction Sequence

For the S-P prediction sequence, we used the APC-HF configuration with a limiter in the noise-feedback path. Preemphasis and high-frequency correction were used in the same way as with the P-S prediction sequence (see Chapter 11). Since we found that using the autocorrelation method of pitch-tap computation yielded fewer discrete distortions in 1% channel error than using the covariance method, we used the

autocorrelation method in the subsequent experiments. Listening tests comparing the 1% channel-error performance produced by the two prediction sequences for otherwise identical APC systems showed that the speech from the S-P system was slightly less reverberant than from the P-S system, but it contained objectionable discrete distortions. The overall speech quality from the S-P sequence was found to be inferior to that from the P-S sequence.

12.5.2 Receiver Smoothing of the Decoded Residual

In an attempt to improve the channel-error performance produced by the S-P sequence, we investigated two modifications to the APC system. Both modifications were motivated by the following considerations. The channel bit-errors may be thought of as an additive random noise corrupting the transmitted bit stream. For the P-S prediction sequence, which leads to a good channel-error performance as we noted above, this additive noise is "smoothed" by the spectral filter $1/A(z)$ before it is processed by the pitch filter $1/C(z)$. Although the pitch filter propagates a bit-error in a periodic manner and with a relatively long time constant, the effect of this smoothing provided by the spectral filter may be responsible for the observed good channel-error performance of the coder using the P-S sequence. With the

S-P sequence, the additive noise mentioned above is processed directly by the pitch filter. This may be responsible for the resulting poor channel-error performance.

First, we reversed the order of the pitch and spectral filters in the receiver of the S-P system so that the receiver processed the channel errors on the residual samples the same way as in the P-S system. Although this change should introduce some additional distortion in the output speech for error-free channels, we hoped that the possible improvement in the coder's channel-error performance might outweigh that bad effect. However, our tests showed that the output speech contained severe distortions and frequently had reverberant quality both in the presence and in the absence of channel bit-errors.

Second, we investigated the effect of inserting a smoothing operation in the receiver, to smooth the decoded residual samples [32]. We investigated two types of smoothing: linear (average) and non-linear (median). We used a 3-point average smoother and both 3-point and 5-point median smoothers. In the presence of channel bit-errors, smoothing reduced some of the discrete distortions in the speech. In this regard, average smoothing was preferred over median smoothing. However, both types of smoothing introduced substantial smearing and muffling of the speech, thus lowering the overall speech quality significantly.

Since neither of the above two modifications improved the channel-error performance of the coder using the S-P prediction sequence, we chose to use the P-S prediction sequence (without smoothing) in our optimized design of the robust APC coder.

12.6 Comparative Evaluation and Recommendations

We compared the three optimized systems, EC-PP3, 4-level PP3-SQ3, and PP3-SQ10-BA, in 1% channel error. A general comment should be made regarding comparisons of different systems in the presence of channel bit-errors. The perceived quality of speech transmitted over an errorful channel depends on the particular realization of the random process causing the channel bit-errors. In comparing two systems, therefore, the speech-quality judgments should be made over a large amount of speech, instead of on a sentence-by-sentence basis. Following this method and using the 12 sentences from the high-quality data base, we found the PP-SQ and PP-SQ-BA systems to be very similar in overall quality. The PP-SQ-BA system produced more discrete distortions and less reverberant quality than the PP-SQ system did. The EC-PP system produced much more reverberant speech and was clearly inferior to the other two systems.

We then tested each of the three systems to check if it

satisfied the engineering channel-error criterion mentioned at the beginning of this chapter. For the EC-PP system, the unprotected system operating in 0.1% channel error produced less reverberant and better overall speech quality than the error-protected system did in 1% channel error. For both the PP-SQ and PP-SQ-BA systems, the protected and unprotected systems yielded roughly the same overall speech quality.

In view of the performance equivalence of the PP-SQ and PP-SQ-BA systems in both error-free and errorful channels, we have chosen the PP-SQ system for real-time implementation. We feel that this is the safer choice, since the PP-SQ system was found to be more robust in the sense that it performed in a more consistent and uniform manner over individual sentences than the PP-SQ-BA system did. The reason for this difference is that while channel bit-errors on the transmitted bit-allocation parameters of the BA scheme can cause erroneous decoding of a part of the residual samples, proper decoding is always ensured in the PP-SQ system. It is therefore reasonable to expect that the PP-SQ system will continue to perform well for speech utterances other than those we used in our study. The PP-SQ system is also less complex and uses a larger frame size, simplifying the requirements on the real-time computation speed.

Finally, we compared the output speech from the chosen PP-SQ

system obtained for two cases: 0% and 1% channel error. Quite impressively, the degradations caused by channel error were found to be perceivable but small; the degradations were in the form of roughness and a slightly reverberant quality.

12.7 Folded Binary Code

In all our channel-error simulations reported thus far, we used the natural binary code (NBC) to encode the quantized residual samples. To improve the robustness of the PP-SQ system further, we investigated the use of an alternate encoding method called the folded binary code (FBC). For FBC, the most significant bit gives polarity information; the remaining bits represent the sample magnitude in natural binary code. Figure 14 illustrates the difference between the two encoding methods, for the 4-level quantizer. For NBC, going from the most negative to the most positive level, the 4 codes are: 00,01,10, and 11. For FBC, the 4 codes are: 01,00,10, and 11. We note that the word "folded" comes from the fact that the codes for the first two levels of NBC have been reversed in FBC. To show when and how FBC yields better performance than NBC, let us assume that only single bit-errors occur within the 2-bit residual code. Referring to Fig. 14, a bit-error causing 00 to be received as 10 results in a decoding error of 2.238 ($=0.419+1.829$; see Fig. 14)

for NBC, but results in a decoding error of only 0.838 ($=0.419+0.419$) for FBC, both decoding errors computed for the Laplacian quantizer shown in Fig. 14. It can be shown that for single bit-errors in the residual code, the mean-square decoding error is the same for all the four levels and equal to 3.52 for NBC. Using FBC reduces the mean-square decoding error of each of the two inner levels to 1.35 at the expense of increasing the decoding error of each of the two outer levels to 7.68. Therefore, if the combined probability of occurrence of the two inner levels is greater than 0.5, which is the case in the PP3-SQ3 system, then using FBC produces a smaller overall mean-square decoding error than using NBC.

When we used FBC in the PP-SQ system for encoding the output levels of a 4-level Laplacian quantizer, we obtained perceivable improvements in the output speech quality in the form of a reduction in both the reverberant quality and the raspy character.

12.8 Performance Over Higher Error-Rate Channels

Since certain Department of Defense applications may have the need to operate 16 kb/s APC coders over channels having error rates higher than 1%, the COTR suggested the testing of the

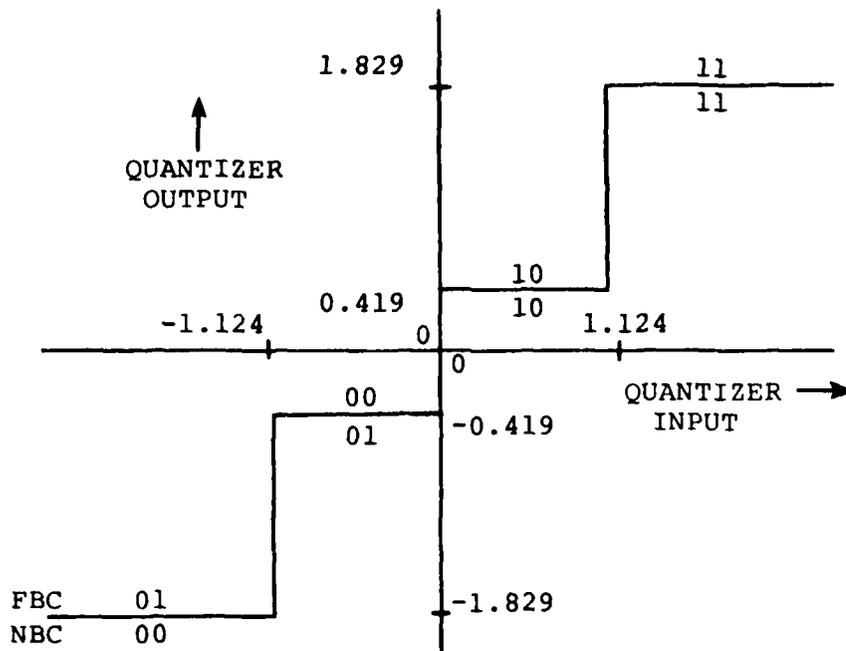


FIG. 14. Illustration of the difference between natural binary code (NBC) and folded binary code (FBC), for a 4-level Laplacian nonuniform quantizer.

chosen robust coder in channel error rates of 2% to 5%. Of course, the design requirement of this project was to achieve a robust performance only for channel error rates of up to 1%. For the higher error rate channels, we obtained the following two results: (1) the output speech intelligibility is satisfactory for 2% bit-errors; and (2) the output speech is not acceptable for 3% and higher bit-error rates, with loud pops and frequent drop-outs of entire words. The reason for this poor performance may be that the effectiveness of the Hamming (7,4) code used in

the APC coder breaks down at these high error rates. Mathematically, it can be shown that the average error rate of the decoded bits is about 0.2% for 1% channel error and about 3% for 4% channel error. We conducted another experiment to determine if the poor performance of the coder was caused by the breakdown of the effectiveness of the Hamming code or by the effect of the channel errors on the unprotected residual signal. Using the FORTRAN simulation of the PP-SQ coder (see Appendix B), we simulated a coder in which only the residual signal was exposed to channel errors, and we processed six sentences at error rates of 3%, 4%, and 5%. The processed speech was quite intelligible even at the 5% error rate. With respect to speech quality, the processed speech sounded generally more raspy as the error rate was increased from 3% to 5%, and the speech from female talkers exhibited more reverberant quality. The perceptual effect of channel errors on only the residual signal seemed similar to that of the quantization noise. From the results of this experiment, we conclude that for an application involving high error-rate channels, more powerful codes than the Hamming (7,4) code and a larger amount of error protection of coder parameters than we have used in the above PP-SQ system are required to yield satisfactory speech intelligibility.

13. ACOUSTIC BACKGROUND NOISE

We tested the performance of the optimized, robust PP3-SQ3 coder for input speech corrupted by one of two acoustic background noise types: office noise (about 60 dB SPL) or ABCP noise (about 90 dB SPL). For these tests, we used sentences from the office-noise data base and from the ABCP data base described in Section 2.3. For the office-noise case, the coder produced output speech with very good quality. For the ABCP noise, the output speech of the coder was found to have good quality and intelligibility. We noted that the output speech quality in the first case was found to be closer to the input speech quality than we observed for the case using the high-quality speech input, and the input-output quality comparison was even closer for the ABCP data base.

14. TANDEMING WITH LPC-10

14.1 Simulation of LPC-10

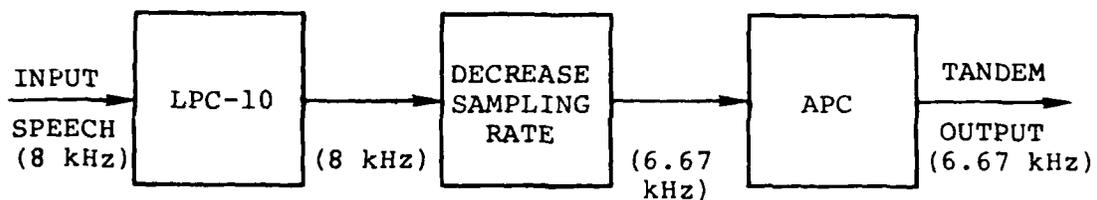
We installed on our DECSYSTEM-20 the PDP-11 FORTRAN implementation of 2.4 kb/s the LPC-10 (version 42) coder [27]. The process of bringing up the LPC-10 coder simulation on our computer involved, among other things, the following two tasks: modification of the input/output sections to accept our formatted speech waveform files and modification to the subroutine "CHANL", which assumes a 16-bit wordlength, to operate properly on our 36-bit computer. The output speech from this implementation of LPC-10 was compared against the audio tape recording of the output from the original PDP-11 implementation of LPC-10. This comparative evaluation and subsequent consultations with the DoD agency that supplied the LPC-10 program clearly indicated that our implementation of LPC-10 was functioning correctly.

Before we present the detailed results of our investigation of the tandem link between APC and LPC-10, we point out that this tandem produced satisfactory performance in either direction unlike the tandem connection between 16 kb/s CVSD coder and LPC-10 [34].

14.2 LPC-APC Tandem

Since the LPC-10 coder employs a 8 kHz sampling rate and our FORTRAN simulation of the APC coder employs a 6.67 kHz sampling rate, the digital tandem interface between the two coders must have provisions for changing the sampling rate. For the LPC-APC tandem, shown in Fig. 15(a), the interface must reduce the sampling rate from 8 kHz to 6.67 kHz; we performed this sampling rate reduction by 5:1 interpolation followed by 6:1 decimation. We used high-order FIR lowpass filters for the operations of interpolation and decimation.

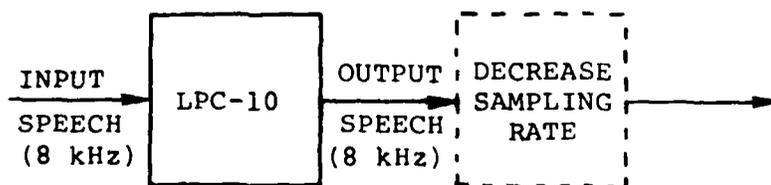
Since the LPC-10 coder uses, as excitation for voiced sounds, a stored version of the impulse response of an allpass filter, its output is not expected to have a high peak factor (peak-to-rms ratio). Signals with a high peak factor may in general increase the clipping errors in an APC system and hence introduce additional distortions in the APC speech. However, the optimized coder design obtained in this work employs 3-tap pitch prediction and 3-segment segmented quantization to track the varying signal amplitudes. In fact, the output speech from the LPC-APC tandem had about the same perceived speech quality as the LPC-10 speech band-limited to 3.33 kHz (shown in dashed lines in Fig. 15(c)). The single-link LPC-10 output was slightly more



(a) LPC-APC tandem



(b) APC-LPC tandem



(c) Single LPC-10 link

FIG. 15. Tandem operation of APC and LPC-10 coders.

"crisp" than the tandem output. Since the LPC-10 speech has energy in the frequency range 3.33-4.0 kHz, potentially it may have a slightly higher speech intelligibility than the speech from the tandem link. We did not conduct any formal speech intelligibility tests.

14.3 APC-LPC Tandem

Because of the difference in the sampling rates of the two coders, the tandem interface must increase the sampling rate from 6.67 kHz to 8 kHz, as shown in Fig. 15(b). We achieved this sampling rate increase by 6:1 interpolation followed by 5:1 decimation.

The output speech from the PP3-SQ3 APC coder has granular noise and some clipping noise. One effect of noise in speech is to reduce its short-term spectral dynamic range. The linear prediction analysis of LPC-10 on APC speech would, therefore, overestimate formant bandwidths, and the resulting speech would in general sound buzzy and be of lower quality than speech from a single LPC-10 link.

In our testing of the APC-LPC tandem, we obtained results very similar to the ones reported above for the LPC-APC tandem, with one difference. The APC-LPC tandem produced output that was

slightly inferior to the output from the LPC-APC tandem. Let us consider ways of improving the performance of the APC-LPC tandem. One method that we suggested in our proposal [1] is to enhance the APC speech through a spectral subtraction method given in [33], before processing through the LPC-10 coder. This method should reduce the distortions caused by the noise in the APC speech. However, we feel that a more serious source of quality problem is at the digital tandem interface. The speech coming out of the interface has a spectrum with a sharp amplitude change (discontinuity) at about 3.33 kHz and with very small amplitudes in a region just below 4 kHz. The subsequent LPC analysis would unduly "spend" some of its resources in attempting to model the spectral discontinuity. Said another way, the LPC analysis makes effective use of fewer than 10 coefficients (which is the number of poles used in LPC-10, for voiced frames). A reasonable solution to this problem would be to "fill in" the spectral gap between 3.33 kHz and 4 kHz using, for example, the high-frequency correction method [11] that we have described in Section 9.4. This spectral correction can be done as part of the interface or as a user-selectable option within the LPC-10 coder. In the latter case, the HFC method can be implemented by simply modifying the elements of the covariance matrix as suggested in [11].

15. 16 KB/S BASEBAND CODERS

As we mentioned in Chapter 2, the baseband coder is one of the three types of coders capable of producing toll quality speech at 16 kb/s. Since we recently designed and implemented on the MAP-300 a 9.6 kb/s BBC system as part of a DCA contract [3], we proposed to extend this design to the 16 kb/s case and compare the performance of the resulting BBC system with that of the optimized, PP3-SQ3 APC system. The results of this work are reported below.

Based on our previous experience [3], we chose to investigate two 16 kb/s baseband coder designs. Both coders use a baseband width of 1.67 kHz, encode the baseband residual using an APC coder with pitch prediction and no spectral prediction, and perform high-frequency regeneration at the receiver using the perturbed spectral folding method [3]. The two 2-band BBC systems are defined in terms of their parameter values, as follows:

System 1

21 ms frame size, 11 quantization levels per baseband residual sample, 3-tap pitch prediction for the baseband APC coder, and 44 of 55 parameter bits protected against channel error.

System 2

27 ms frame size, 16 quantization levels per baseband residual sample, 1-tap pitch prediction for the baseband APC coder, and 28 of 49 parameter bits protected against channel error.

In the error-free channel, System 2 had less background noise and was judged to be of higher quality than System 1. We compared System 2 with the optimized fullband APC coder (PP3-SQ3). The fullband APC coder had noticeably more background noise, particularly at low frequencies, but had better overall speech quality because of the unnatural high-frequency distortions produced by the baseband coder. We repeated the same comparison in the presence of 1% channel error. In this test, the fullband APC coder speech quality was clearly superior to that of the baseband coder.

16. OPTIMIZED, ROBUST 16 KB/S APC CODER

In this chapter, we report the results of further work on the PP3-SQ3 coder that we chose as the most robust coder. This work consisted of 1) making some performance-preserving simplifications to the coder design for facilitating real-time implementation on the MAP-300 and 2) providing some refinements to the coder design, to improve the coder performance further. Then, we summarize the details of the final design of the coder and introduce Appendices A-C, which contain a detailed specification and FORTRAN simulation of the coder. Finally, we present and discuss the results of our tests on the real-time implementation of the coder.

16.1 Simplifications for Real-Time Implementation

16.1.1 High-Frequency Correction

Recall from Section 9.4 and Eq. (38) that HFC requires solving the linear prediction normal equations twice, once using the computed autocorrelation coefficients, and once using the modified autocorrelation coefficients. The first of the two solutions is required only to compute the normalized error V_p of the p -th order linear predictor ($p=6$ in our case). To reduce

the complexity of the first step, we used V_2 as an estimate of V_6 , where V_2 can be computed explicitly in terms of $R(0)$, $R(1)$ and $R(2)$ as given below:

$$V_2 = \left(1 - [R(1)/R(0)]^2\right) \left(\left[1 - \frac{R^2(1) - R(0)R(2)}{R^2(0) - R^2(1)}\right]^2\right) \quad (39)$$

Using this second-order estimate, we reoptimized the parameter λ in Eq. (38) to be $\lambda = 0.035$, the previous choice being $\lambda = 0.05$. The original and the simplified HFC procedures were found to yield the same S/Q ratio and speech quality for the PP3-SQ3 coder. Therefore, we recommend the use of this simplified procedure in the real-time coder.

16.1.2 Pitch-Filter Stability Test

Since the optimized PP3-SQ3 coder uses 3-tap pitch prediction, it requires checking the stability of the pitch filter every frame and switching to 1-tap pitch prediction for frames for which instability is detected (see Section 8.2.2.1). The (exact) method of testing the pitch-filter stability requires $6M$ multiplies/frame, where M is the pitch period in number of samples: a nontrivial computation, especially for male speakers. To simplify the stability testing procedure, we considered an orthogonal linear transformation [35] of the three tap coefficients given below:

$$\left. \begin{aligned} T_1 &= C_1 + C_2 + C_3, \\ T_2 &= C_1 - 2C_2 + C_3, \\ T_3 &= C_1 - C_3, \end{aligned} \right\} \quad (40)$$

where we have denoted the tap coefficients $C(M-1)$, $C(M)$, and $C(M+1)$ as C_1 , C_2 , and C_3 , for convenience. (This notation is also used in Appendix A.) Initially, we investigated a procedure that declared the pitch filter stable if the transformed coefficients satisfy the relations:

$$|T_1| < 1, |T_2| < 1, |T_3| < 1; \quad (41)$$

1-tap pitch prediction was used when the magnitude conditions (41) were not satisfied. Mathematically, the conditions (41) are neither necessary nor sufficient for pitch-filter stability. Experimentally, we found that as a detector of pitch filter instability, the above procedure yielded very high probability of detection (one error in 1200 frames) at the expense of a high false-alarm rate (declared instability for 20% of frames, while only 8% of frames had an unstable filter). Relative to the exact method, this simplified procedure yielded about 0.3 dB decrease in S/Q ratio and slight but audible speech quality degradation for 0% channel error, and it produced slightly more reverberant speech for 1% channel error because of the increased use of 1-tap prediction (see Section 12.2). Upon closer examination of the

false-alarm cases of the above procedure, we found that in each case T_2 was in the range $-2 < T_2 < 1$. Therefore, we modified the "stability conditions" to be:

$$|T_1| < 1, -2 < T_2 < 1, |T_3| < 1. \quad (42)$$

This modified procedure yielded only 2 errors in the detection of instability out of the 1200 frames we considered. More important, this modified procedure yielded the same coder performance as the exact method both in the absence and in the presence of channel bit-errors. Therefore, the stability testing procedure involving Eqs. (40) and (42) is recommended for the real-time implementation.

16.1.3 Noise Shaping

The optimized PP3-SQ3 coder employs the pole-zero noise shaping method with a bandwidth parameter w of 800 Hz. Pole-zero noise shaping requires more computation and more coefficient and data memory than other types of noise shaping. The memory requirement is quite important for implementation on the MAP. The computational complexity of pole-zero noise shaping is roughly twice that of all-pole noise shaping and 2^p ($p=6$ in our case) times that of 1-zero noise shaping. In an attempt to simplify the implementation, the all-pole and the one-zero noise

shaping methods were re-examined in systems otherwise identical to the optimized coder. For the all-pole method, $w=200$ Hz produced the best output speech quality. We then compared the speech produced by each of these two noise shaping methods against the speech produced by the pole-zero method. The speech for the all-pole method sounded more raspy and rough, and the 1-zero method produced noticeably more roughness (because of the use of preemphasis) and more background noise. The output speech obtained without noise shaping contained discrete distortions (e.g., clicks) and an increased level of roughness and background noise. Therefore, the complexity of the pole-zero method is worthwhile to keep in the real-time coder.

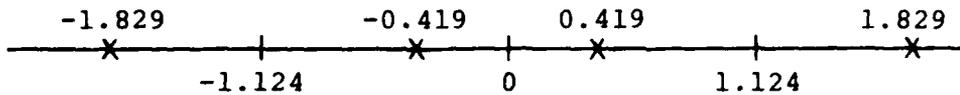
16.2 Refinements to the Coder

16.2.1 Laplacian versus Gaussian Quantizer

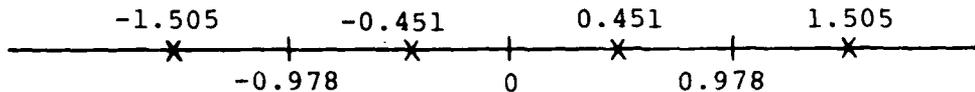
For the error-free channels, we had previously found that both Laplacian and Gaussian optimal nonuniform residual quantizers produced essentially the same perceived speech quality. We used the Laplacian quantizer in most of our simulations, because it produced about 0.5 dB higher S/Q ratio than did the Gaussian quantizer. However, when we repeated the same comparison for 1% channel error, we found that the Gaussian

quantizer produced a noticeable improvement in the speech quality over the Laplacian quantizer. The extent of reverberant quality and the loudness of discrete noises in the output speech were reduced with the use of the Gaussian quantizer. The observed difference in channel-error performance may be explained as follows. Using the decoded values for the two quantizers shown in Fig. 16 and assuming single bit-errors in the residual code, we can show that the mean-square decoding errors for the four levels in FBC (01,00,10,11) are 7.68, 1.35, 1.35, and 7.68 for the Laplacian case and 5.39, 0.96, 0.96, and 5.39 for the Gaussian case. Therefore, for each level, the Gaussian quantizer produces a lower mean-square decoding error than the Laplacian quantizer. This, therefore, explains the observed improvement produced by the Gaussian quantizer over the Laplacian quantizer.

Also, we observed that the benefit provided by the folded binary code was less in the Gaussian case than in the Laplacian case. This result can be explained by the larger width and hence the larger probability of occurrence of the inner levels for the Laplacian case than for the Gaussian case (see Fig. 16) and by the result presented in Section 12.7. The Gaussian quantizer was still judged to be better than the Laplacian quantizer when the two cases were compared, both using FBC. Therefore, we recommend the use of the Gaussian quantizer in the final coder design.



(a) Laplacian quantizer



(b) Gaussian quantizer

FIG. 16. Optimum Laplacian and Gaussian quantizers. The tick-marks are used to indicate quantizer input boundaries, and the symbols x are used to indicate quantizer output values.

16.2.2 Recomputing Parameter Quantization Tables

Having completed the coder design, we recomputed the statistics of each transmission parameter for the purpose of checking the ranges and step sizes used for the quantization. Using the 12-sentence high-quality data base, histograms were prepared for each parameter. Maximum and minimum parameter values (to be used in quantization) were then estimated by visual inspection of the histograms. As a result of the recomputed

statistics, we revised the quantization of log area ratios and delta gains. The revised quantization of delta gains is given below in Table 11, while the LAR quantization is included in the overall description of the optimized coder given in the next section.

Quantizer Input	Level	Quantizer Output
$-\infty$		
-3.6	1	-6.2
-0.5	2	-2.0
2.2	3	1.0
∞	4	3.5

TABLE 11. Revised nonuniform quantization of delta gains.

16.3 Optimized Coder Description

Before we describe the optimized coder, we discuss the choice of the frame size to be used in the real-time coder. Recall that the sampling rate of the real-time system is 6.621 kHz, while that used in the simulations is 6.67 kHz (Section 2.2). We found 32.625 ms to be the best choice of the frame size for the real-time system corresponding to the value of 33.75 ms

that we used in our simulations. Below, we summarize the details of the optimized, robust 16 kb/s coder. A detailed specification of the coder design is given in Appendix A.

A block diagram of the optimized coder is shown in Fig. 17. Table 12 provides information regarding the quantization and error protection of parameter data of the APC system. At the transmitter, the analog input speech is lowpass filtered at 3.2 kHz and sampled at 384/58 (or about 6.621) kHz. Referring to Fig. 17(a), the sampled speech $s(t)$ is divided into frames of 216 samples (32.625 ms duration). Each frame of speech is preemphasized using the filter $(1 - 0.4z^{-1})$. The preemphasized speech $s'(t)$, before being encoded by the APC encoder, is processed, as explained below, to extract in order pitch predictor parameters, spectral predictor parameters, and segment gains of the quantizer. Extraction of the pitch predictor parameters consists of the following steps: computing the autocorrelation function of $s'(t)$ for lags 0-134, from an interval of 265 samples (216 from the current frame and 49 from the previous frame); determining the pitch value M as the peak of this function over lags 14-133; solving for the 3 pitch taps for the 3-tap filter and for the single tap for the 1-tap filter from the corresponding autocorrelation normal equations; checking for the stability of the 3-tap filter using Eqs.(40) and (42) and, if

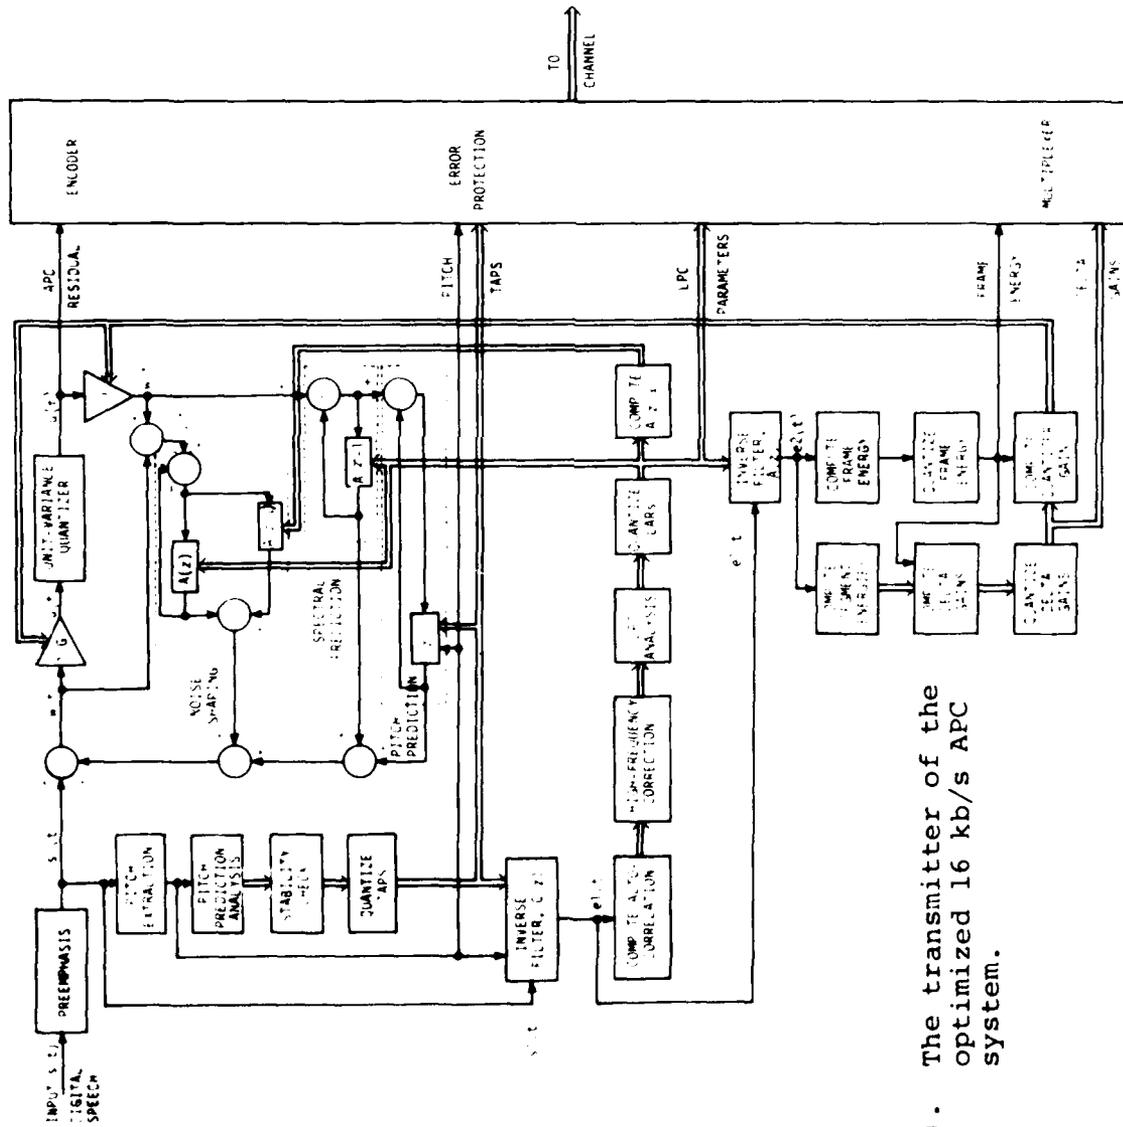


FIG. 17(a). The transmitter of the optimized 16 kb/s APC system.

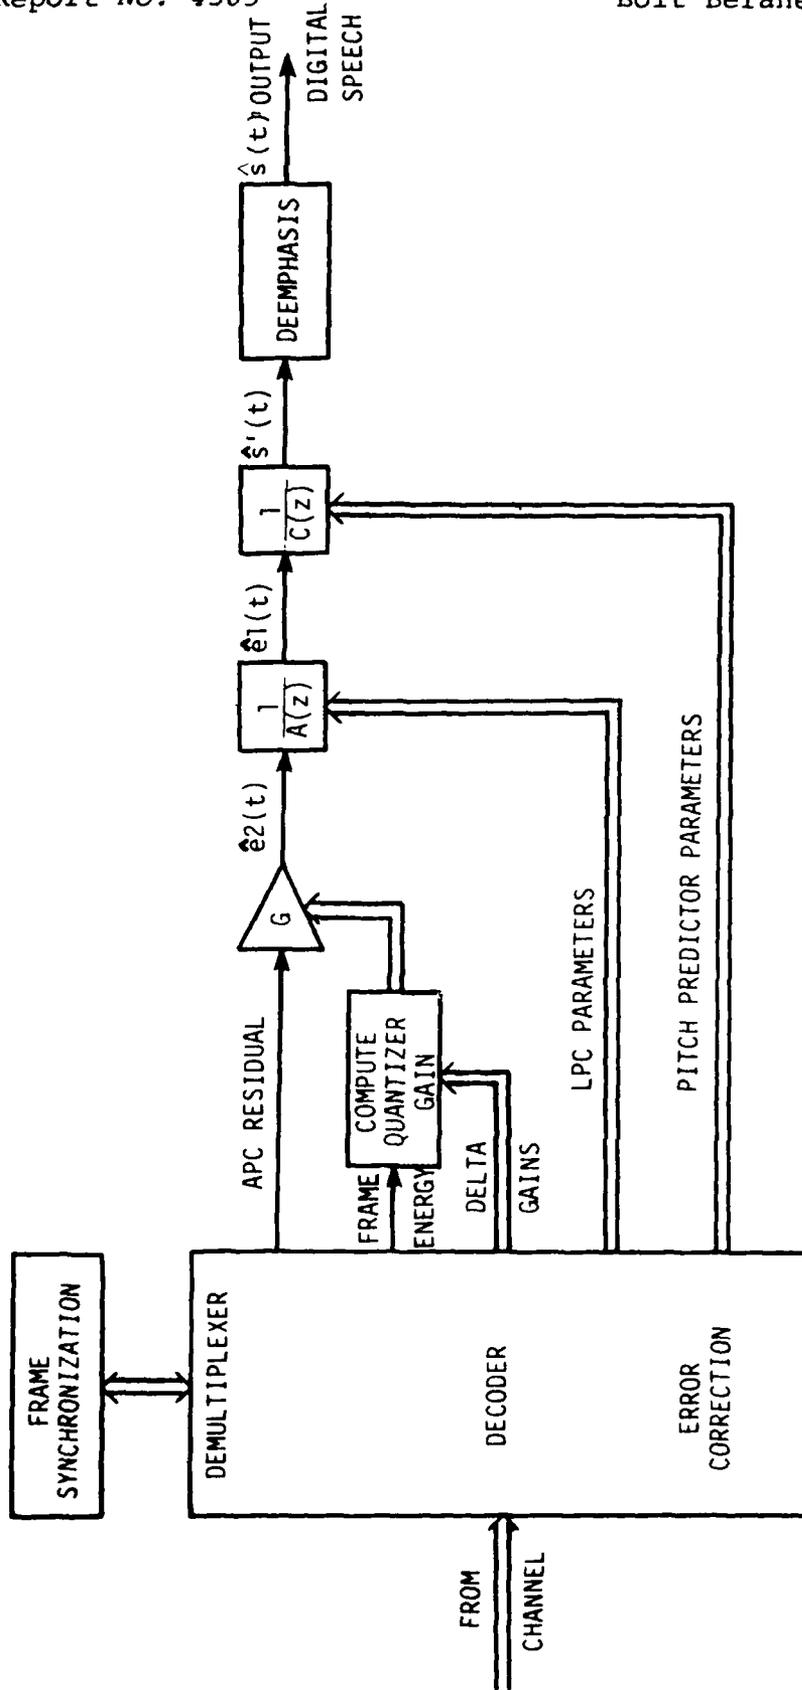


FIG. 17(b). The receiver of the optimized 16 kb/s APC coder.

Parameter		Min	Max	Step Size	# of Bits	# of Bits Protected
Pitch		14	133	No Quantization	7	7
Pitch Taps	C1	-0.549	0.427	0.122	3	2
	C2	-0.95	0.12	0.067	4	3
	C3	-0.549	0.427	0.122	3	2
Second residual energy (dB)		-10.0	46.0	0.875	6	6
Delta gains	1	See	Table	11	2	2
	2				2	
	3				2	
Log Area Ratios (dB)	1	-21.849	11.053	0.514	6	5
	2	- 8.789	13.711	0.703	5	4
	3	- 9.031	7.969	1.063	4	3
	4	- 6.094	8.906	0.938	4	2
	5	- 5.281	7.719	0.813	4	2
	6	- 3.741	9.559	0.831	4	2
Total bits per frame					56	44

TABLE 12. Quantization and error protection of parameter data for the optimized 16 kb/s APC system.

found unstable, replacing the computed 3-tap filter with another 3-tap filter with zero side taps and its center tap equal to the computed 1-tap coefficient; quantizing the 3 pitch taps; and inverse filtering the signal $s'(t)$ to produce the first residual $e_1(t)$, using the pitch-inverse filter $C(z)$ with quantized coefficients. Extraction of spectral parameters consists of the following steps: computing the autocorrelation function of $e_1(t)$ for lags 0-6; modifying the computed values of this autocorrelation function using Eqs.(38) and (39) and with $\lambda=0.035$; obtaining the reflection coefficients via autocorrelation LPC analysis; quantizing the reflection coefficients (via the log area ratio transformation); computing the coefficients of the numerator of the noise shaping filter, $A(z/\alpha)$ from the LPC predictor coefficients; and inverse filtering of the signal $e_1(t)$ to produce the second residual $e_2(t)$, using the spectral inverse filter $A(z)$. Extraction of the segment quantizer gains consists of the following steps: computing the energy of $e_2(t)$ over the frame and over each of the three 72-sample segments in the frame; quantizing the frame energy; computing the three delta gains as the ratio of the segment energy and the quantized frame energy; quantizing the delta gains; and computing the segment quantizer gains from the quantized frame energy and the quantized delta gains. The quantized values of the various extracted parameters are used to

update the corresponding parameters of the APC encoder, which is set up as in the APC-PF configuration (see Fig. 4). The residual quantizer in the APC encoder is the optimal, 4-level, Gaussian, nonuniform quantizer (see Fig. 16(b)). The quantized parameter data, 3 pitch taps, 6 LARs, and frame energy and segment delta gains of the second residual, and the unquantized pitch are all binary encoded, error protected using 11 Hamming (7,4) codewords (see Table 12), multiplexed with one synchronization bit and 432 bits (2 bits/sample x 216 samples) of folded-binary encoded residual data, and transmitted over the channel.

At the receiver, shown in Fig. 17(b), the received data are demultiplexed, decoded, and error-corrected. The three segment quantizer gains are computed from the decoded frame energy and the delta gains. The decoded APC residual samples are multiplied by the corresponding segment quantizer gain and filtered first by the spectrum-synthesis filter $1/A(z)$ and then by the pitch-synthesis filter $1/C(z)$. The filtered output $\hat{s}'(t)$ is deemphasized using the filter $1/(1 - 0.4z^{-1})$ to produce the digital speech output $\hat{s}(t)$ (as an approximation to the original input $s(t)$). This digital output is passed through a D/A converter and an analog lowpass filter with its cutoff at 3.2 kHz to produce the analog output speech.

The COTR was supplied with an audio demonstration tape in

June 1980. The tape contained the recordings of the output speech obtained from the simulation of the above described 16 kb/s APC system. The recorded sections on the tape successfully demonstrated the performance of the robust coder, respectively, for high-quality input speech, in acoustic background noise, over a noisy channel in 1% bit-errors, and in tandem with the 2.4 kb/s LPC-10 coder. In each of these cases, the coder performance met the requirements stated in Chapter 1.

16.4 FORTRAN Simulation of the Optimized Coder

During the project, we developed a general software package to simulate the APC coder. It contained many features that aided us in program debugging and in the coder optimization and evaluation. This general software package was modified to produce a FORTRAN simulation of only the final optimized system.

A user's guide for this FORTRAN simulation is included with this report as Appendix B, and a listing of the FORTRAN source programs is contained in Appendix C. We have tested and verified that the FORTRAN simulation of the optimized coder produced synthesized speech identical to that produced by the general software package with the parameters set as in the optimized coder.

16.5 Results of Tests on the Real-Time Coder

We tested the real-time 16 kb/s coder on the MAP-300, using input from several tapes containing speech from a number of males and females, for informal evaluation of the coder performance for different speech materials and different speakers. Except for two problems mentioned below, the coder was found to produce high quality speech output. First, the output speech for one low-pitched male talker (with an average pitch of 95 Hz) contained audible roughness. Second, the coder produced audible background noise at the output for some female talkers.

To investigate the causes of these problems, we performed several tests on the real-time coder and on the FORTRAN simulation. First, using the RT-11 debugging program (FDT) on our PDP-11, the values of three of the coder parameters, which are specified in DATA statements, were varied about their nominal (previously optimized) values. The three parameters are: preemphasis constant β , noise shaping bandwidth parameter w , and high-frequency correction coefficient λ . After each parameter change, we listened to the output of the real-time coder, with its input speech from a tape. For each of the parameters, we concluded that the nominal value produced the best overall speech quality.

Second, using the CSPI-supplied program MPLOOK, we made changes to the coding and decoding tables of the 4-level residual quantizer. Two types of changes were investigated: 1) Each element of the coding and decoding tables was multiplied by a constant (called quantizer load factor) to investigate the tradeoff between clipping and granular quantization errors; and 2) different types of unit-variance quantizers were employed. For the first item, we used values of 0.8, 1.0 (nominal value), and 1.2 as load factors. For the second item, we compared Laplacian and gamma quantizers with the nominal Gaussian quantizer. This investigation of changes to the quantizer also resulted in no perceivable improvement in the overall speech quality of the coder.

Since the two types of testing, described above, on the real-time coder did not uncover the observed speech-quality problems, we decided to pursue the subsequent work using our earlier versions of the FORTRAN simulation of the APC coder. We chose two specific sentences that suffered the greatest quality degradation and digitized them at the 6.67 kHz sampling rate. We processed one of these sentences (spoken by a low-pitched male talker) using the simulation program, without quantization of any parameters (i.e., with only the residual being quantized). The output speech was found to be nearly identical to the output of

the real-time coder. We then investigated the following changes to the coder (one change at a time), but we observed no significant improvement in the coder performance: 1) Pitch filter stability check was not used; 2) analysis frame size used for pitch computation was varied between 35 and 45 ms; and 3) LPC order used in spectral prediction was increased from 6 to 10 poles.

In a subsequent set of tests, we found that each of the following changes did produce a significant increase in speech quality: (1) variable-rate entropy coding, with an average entropy of 2 bits/sample; (2) increase from 3 (optimized value) to 10 in the number of segments used for segmented quantization; and (3) use of pitch-adaptive quantization. We did not use variable-to-fixed rate conversion in (1), and we did not readjust the bit allocation to limit the data rate to 16 kb/s in (2). For the low-pitched male speaker, the entropy coding and 10-segment schemes each produced slightly higher speech quality than the pitch-adaptive scheme. For the second sentence from a female speaker, increasing the number of segments to 10 did not improve the speech quality. Also, we found that the 5-segment scheme produced about the same overall speech quality as the 3-segment scheme, even for the low-pitched male speaker.

Based on these experimental investigations, we offer the

following conclusions. The observed speech quality degradations were caused by the relatively large dynamic range of the input to the residual quantizer. Both entropy coding and pitch-adaptive quantization methods represent effective ways of dealing with the problem. However, the performance of the entropy coding method under channel errors is substantially worse than the performance produced by the optimized coder (see Section 12.6). As for the pitch-adaptive method, its implementation on the MAP is extremely difficult, as we reported in Chapter 11. Increasing the number of segments from 3 to 10 prevents only the roughness problem observed for low-pitched males. Further, such a change would involve a reoptimization of the coder and may perhaps lead to a less robust channel-error performance than our original optimized coder. All things considered, we believe that the coder design implemented on the MAP is still the most robust coder meeting the design requirements given in Chapter 1. The test results reported in this section have shown that for some speakers, the coder output speech degrades perceptibly relative to the input speech.

17. SUMMARY AND MAJOR CONTRIBUTIONS

In summary, we have investigated and compared several methods, some already existing ones and some new ones developed in this work, for coding the residual signal and for shaping the spectrum of the quantization noise, in the course of optimizing the APC system to meet the specific needs of this project. As part of this work, we have also optimized the values of various parameters as well as the bit allocation for those parameters that are transmitted to the receiver, to produce the best output speech quality at a synchronous data rate of 16 kb/s and for an input-speech sampling rate of 6.67 kHz. For the noisy channel application, we have considered in the optimization study the tradeoff between the voice data rate and the error-protection rate and the allocation of the error protection bits among individual transmission parameters.

As a result of this work, we have developed two best 16 kb/s APC systems, one for use over perfect or noiseless channels and the other for noisy channel applications involving as much as 1% bit-errors. For an error-free transmission, the best system uses 8-pole spectral prediction, 3-tap pitch prediction, entropy coding with a large number of quantizer levels (43 levels used in our tests), and pole-zero noise shaping. For operation over

noisy channels, the most robust system uses 6-pole spectral prediction, 3-tap pitch prediction, 3-segment segmented quantization with a 4-level nonuniform Gaussian quantizer, and pole-zero noise shaping; allocates to error protection of parameters slightly over 6% of the total transmission bit rate (or about 37% of the bit rate used for parameter transmission); and encodes the quantized residual samples with the folded binary code. Informal listening tests have shown that the second system satisfies all the design requirements of this project: speech-quality requirements for high-quality speech inputs and for acoustic background noise environments, robustness requirement in channel bit-errors of 1%, and speech-intelligibility requirement for tandem operation with a 2.4 kb/s LPC-10 coder. We have made specific suggestions for improving the speech quality of the APC-LPC tandem. Quite impressively, the robust coder produces only a slight speech quality degradation as the channel bit-error rate is increased from 0% to 1%.

In this work, in addition to designing a robust APC coder that meets the requirements of this project, as mentioned above, we have made a number of significant contributions, which when put together represent, in our view, an advance in the state of the art in adaptive predictive coding of speech. The specific contributions of this work are stated below:

1. Demonstration of the important role played by the sequencing of spectral and pitch predictors.
2. Establishment of performance equivalence conditions for the several configurations of the APC system. (Any violation of these conditions has been found to yield a significant performance degradation.)
3. Demonstration of the effects of, and development of a successful remedy for, the instability problem of multi-tap pitch prediction.
4. Identification of excessive quantization-noise problems as the limit-cycle behavior of the quantizer output, interpretation of the causes of the limit cycles in terms of the feedback gain of the APC loop, and comprehensive solution of the limit-cycle problem by reducing the feedback gain.
5. Demonstration of the dual benefits of noise shaping: suppression of quantization-noise perception and reduction of feedback gain, and of how the role of noise shaping is affected by other system components (e.g., preemphasis).

6. Development of new methods for the coding of APC residual: multi-tap pitch prediction and segmented quantization; pitch-adaptive coding with multi-tap pitch prediction and pitch-synchronous segmented quantization and using variable number of bits/sample over segments; and segmented quantization with multi-tap pitch prediction and using a variable number of bits/sample over segments.

7. Demonstration of the importance of (multi-tap) pitch prediction for significantly improving the coder performance both over noiseless channels and over noisy channels. That a robust APC coder design must include pitch prediction has been vividly demonstrated in one of our experiments comparing three 16 kb/s entropy-coded (0-tap, 1-tap and 3-tap) systems operating in 1% channel error.

REFERENCES

1. Proposal for Speech Algorithm Optimization at 16 kb/s, BBN Proposal No. P79-ISD-42, Bolt Beranek and Newman Inc., February 1979, Submitted to Defense Communications Agency.
2. J.L. Flanagan, "Opportunities and Issues in Digitized Voice," Proceedings EASCON '78, EASCON, Washington, D.C., September 1978, pp. 709-712.
3. R. Viswanathan, J. Wolf, L. Cosell, K. Field, A. Higgins, and W. Russell, "Design and Real-Time Implementation of a Baseband LPC Coder for Speech Transmission Over 9600 Bps Noisy Channels - Final Report," BBN Report No. 4327, Bolt Beranek and Newman Inc., February 1980.
4. R. Viswanathan, A. Higgins, W. Russell and J. Makhoul, "Baseband LPC Coders for Speech Transmission over 9.6 kb/s Noisy Channels," IEEE International Conf. Acoustics, Speech and Signal Processing, IEEE, Denver, CO, April 1980, pp. 348-351.
5. J. Makhoul and M. Berouti, "Adaptive Noise Spectral Shaping and Entropy Coding in Predictive Coding of Speech," IEEE Trans. Acoustics, Speech and Signal Processing, Vol. ASSP-27, February 1979, pp. 63-73.
6. R. Zelinski and P. Noll, "Adaptive Transform Coding of Speech Signals," IEEE Trans. Acoustics, Speech and Signal Processing, Vol. ASSP-25, August 1977, pp. 299-309.
7. B.S. Atal and M.R. Schroeder, "Adaptive Predictive Coding of Speech Signals," Bell Syst. Tech. J., Vol. 49, October 1970, pp. 1973-1986.
8. N.S. Jayant, "Digital Coding of Speech Waveforms: PCM, DPCM, and DM Quantizers," Proc. IEEE, Vol. 62, May 1974, pp. 611-632.
9. P. Noll, "A Comparative Study of Various Quantization Schemes for Speech Encoding," Bell Syst. Tech. J., Vol. 54, November 1975, pp. 1597-1614.
10. A.J. Goldberg, R. L. Freudberg and R. S. Cheung, "High Quality 16 kb/s Voice Transmission," IEEE International Conf. Acoustics, Speech and Signal Processing, IEEE, Philadelphia, PA, April 1976, pp. 244-246.
11. B.S. Atal and M.R. Schroeder, "Predictive Coding of Speech Signals and Subjective Error Criteria," IEEE Trans. Acoustics, Speech and Signal Processing, Vol. ASSP-27, June 1979, pp. 247-254.

12. D.L. Cohn and J.L. Melsa, "The Residual Encoder - An Improved ADPCM System for Speech Digitization," IEEE Trans. Communications, Vol. COM-23, September 1975, pp. 935-941.
13. J.Makhoul, "Linear Prediction: A Tutorial Review," Proc. IEEE, Vol. 63, April 1975, pp. 561-580.
14. J. Max, "Quantizing for Minimum Distortion," IRE Trans. Info. Theory, Vol. IT-6, March 1960, pp. 7-12.
15. M.D. Paez and T.H. Glisson, "Minimum Mean-Squared-Error Quantization in Speech PCM and DPCM Systems," IEEE Trans. Commun. Tech., Vol. COM-20, April 1972, pp. 224-230.
16. P. Noll, "Adaptive Quantizing in Speech Coding Systems," Proc. 1974 IEEE Zurich Seminar on Digital Commun., IEEE, March 1974, pp. B3(1)-B3(6). (Reprinted in the IEEE Press Book Waveform Quantization and Coding, N.S.Jayant, (Ed.), 1976, pp.204-209.)
17. B.J. McDermott, C. Scagliola and D. Goodman, "Perceptual and Objective Evaluation of Speech Processed by Adaptive Differential PCM," Bell Syst. Tech. J., Vol. 57, May-June 1978, pp. 1597-1618.
18. J.M. Tribolet, P. Noll, B.J. McDermott and R.E. Crochiere, "A Study of Complexity and Quality of Speech Waveform Coders," Proc. 1978 IEEE International Conf. Acoustics, Speech and Signal Processing, IEEE, Tulsa, OK, April 1978, pp. 586-590.
19. C.C. Cutler, "Transmission Systems Employing Quantization," U.S. Patent 2,927,962, 1960.
20. E.G. Kimme and F.F Kuo, "Synthesis of Optimal Filters for a Feedback Quantization System," IEEE Trans. Circuit Theory, Vol. CT-10, 1963, pp. 403-413.
21. R.C. Brainard and J. C. Candy, "Direct-Feedback Coders: Design and Performance with Television Signals," Proc. IEEE, Vol. 57, May 1969, pp. 776-786.
22. P. Noll, "On Predictive Quantizing Schemes," Bell Syst. Tech. J., Vol. 57, May-June 1978, pp. 1499-1532.
23. R. Viswanathan and J. Makhoul, "Quantization Properties of Transmission Parameters in Linear Predictive Systems," IEEE Trans. Acoustics, Speech and Signal Processing, Vol. ASSP-23, June 1975, pp. 309-321.
24. R. Viswanathan, J. Makhoul and A.W.F. Huggins, "Speech Compression and Evaluation," BBN Report No. 3794, Bolt Beranek and Newman Inc., April 1978.

25. E. Blackman, R. Viswanathan, W. Russell and J. Makhoul, "Narrowband LPC Speech Transmission over Noisy Channels," IEEE International Conf. Acoustics, Speech and Signal Processing, IEEE, Washington, D.C., April 1979, pp. 60-63.
26. J. Makhoul, "Stable and Efficient Lattice Methods for Linear Prediction," IEEE Trans. Acoustics, Speech and Signal Processing, Vol. ASSP-25, October 1977, pp. 423-428.
27. T.E. Tremain, J.W. Fussell, R.A. Dean, B.M. Abzug, M.D. Cowing and P.W. Boudra, Jr., "Implementation of Two Real-Time Narrowband Speech Algorithms," Proc. EASCON '78, EASCON, Washington, D.C., September 1978, pp. 698-708.
28. D.A. Huffman, "A Method for the Construction of Minimum-Redundancy Codes," Proc. IRE, Vol. 40, September 1952, pp. 1098-1101.
29. M. Berouti, J. Makhoul, and M. Krasner, "Efficient Encoding and Decoding of Speech," BBN Report No. 4567, Bolt Beranek and Newman Inc., Dec. 1980.
30. R. Hamming, Introduction to Applied Numerical Analysis, McGraw-Hill Book Company, 1971.
31. F. Itakura, "Research on Low Bit Rate Speech Coding at the Electrical Communication Laboratory, NTT," Presented at the 96th Meeting of the Acoust. Soc. Amer., Honolulu, Hawaii, November 27-December 1, 1978.
32. N.S. Jayant, "Average- and Median-Based Smoothing Techniques for Improving Digital Speech Quality in the Presence of Transmission Errors," IEEE Trans. Communications, Vol. COM-24, September 1976, pp. 1043-1045.
33. M. Berouti, R. Schwartz and J. Makhoul, "Enhancement of Speech Corrupted by Acoustic Noise," 1979 IEEE International Conf. Acoustics, Speech and Signal Processing, IEEE, Washington, D.C., April 1979, pp. 208-211.
34. T.P. Barnwell, R.W. Schafer, and A.M. Bush, "Evaluation of LPC/CVSD Tandem Connections," IEEE International Conf. Acoustics, Speech and Signal Processing, Tulsa, OK, Apr. 1978, pp. 326-329.
35. B.S. Atal, Personal Communication with R. Viswanathan.

APPENDIX A

SPECIFICATION OF THE OPTIMIZED
16 KB/S APC ALGORITHM

TABLE OF CONTENTS

1.	SOME GENERAL ITEMS.	167
2.	TRANSMITTER	169
2.1	Preemphasis.	171
2.2	Pitch Analysis	171
2.2.1	Pitch Extraction.	172
2.2.1.1	Remove DC.	172
2.2.1.2	Hamming Window	173
2.2.1.3	Compute Autocorrelation Coefficients	173
2.2.1.4	Compute Pitch.	174
2.2.2	Pitch Prediction Computation.	175
2.2.3	Stability Check	175
2.2.4	Code-Decode Taps.	177
2.2.5	Code-Decode Pitch	177
2.3	Inverse Filter to Obtain First Residual.	178
2.4	Spectral and Noise-Shaping Analysis.	178
2.4.1	Compute Autocorrelation Coefficients.	179
2.4.2	High-Frequency Correction (HFC)	179
2.4.3	Spectral Predictor Computation.	180
2.4.4	Code-Decode Reflection Coefficients	180
2.4.5	K-to-A Conversion	181
2.4.6	Noise-Shaping Analysis.	181
2.5	Inverse Filter to Obtain the Second Residual	182
2.6	Gain Computation	182
2.6.1	Compute Energy.	183
2.6.2	Code-Decode Energy.	184
2.6.3	Compute Segment Energies.	184
2.6.4	Compute Delta Gains	184
2.6.5	Code-Decode Delta Gains	185
2.6.6	Compute Quantizer Scale Factors	185
2.7	APC Loop	186
2.7.1	Compute Predictions for Sample <i>i</i>	186
2.7.2	Compute APC Residual, WO (Step 5)	188
2.7.3	Normalize APC Residual (Step 6)	188
2.7.4	Code-Decode Normalized Residual (Steps 7 and 8).	189
2.7.5	Scale the Quantized Residual (Step 9)	189
2.7.6	Update Arrays (Steps 10-13)	189
2.8	Folded Binary Code (FBC) for Encoding the Residual Samples.	190
3.	RECEIVER	191
3.1	Decode Parameters.	193

Table of Contents (Cont.)

3.2	Compute scale factors.	193
3.3	Scale the quantized residual	193
3.4	K-to-A conversion.	193
3.5	LPC synthesis.	193
3.6	Pitch Synthesis.	194
3.7	Deemphasis	195
REFERENCE	196

1. SOME GENERAL ITEMS

- 1.1 Speech signal sampling rate = 384/58 KHz (~6.621 KHz)
- 1.2 Frame size = 32.625 ms or 216 samples
- 1.3 Spectral predictor order = 6
- 1.4 Pitch predictor order = 3
- 1.5 Parameter Coding

Coding and decoding tables (Tables 1-5) are given at the end of this appendix. Each of the tables has three columns, $X(J), J, R(J)$, where

$X(J)$ = quantization boundary
 J = code or level
 $R(J)$ = decoded or quantized parameter value.

When a parameter has a value A , which satisfies $X(J) \leq A < X(J+1)$, it is coded as J and decoded as $R(J)$.

- 1.6 Data rate = 16 kb/s or 522 bits/frame

<u>Item</u>	<u>Bits/frame</u>
Parameter data	56
Protection	33
Residual samples (216 X 2)	432
Sync	1
Total	<u>522</u>

1.7 Bit Allocation for Quantization and Protection

Parameter		Bits (Total)	Most Significant Bits Protected
Reflection Coefficients	IK(1)	6	5
	IK(2)	5	4
	IK(3)	4	3
	IK(4)	4	2
	IK(5)	4	2
	IK(6)	4	2
Gain	IG	6	6
Pitch	IM	7	7
Delta Gains	IDG(1)	2	2
	IDG(2)	2	2
	IDG(3)	2	2
Pitch Taps	IC(1)	3	2
	IC(2)	4	3
	IC(3)	<u>3</u>	<u>2</u>
Total		56	44

The 44 bits are protected using 11 Hamming (7,4) codewords. Error protection and correction are done as in our 9.6 kb/s BBC coder [1], and therefore these items are not discussed below.

Bolt Beranek and Newman Inc.

2. TRANSMITTER

A block diagram of the transmitter is given in Figure 1. In this section, we specify each of the various transmitter components.

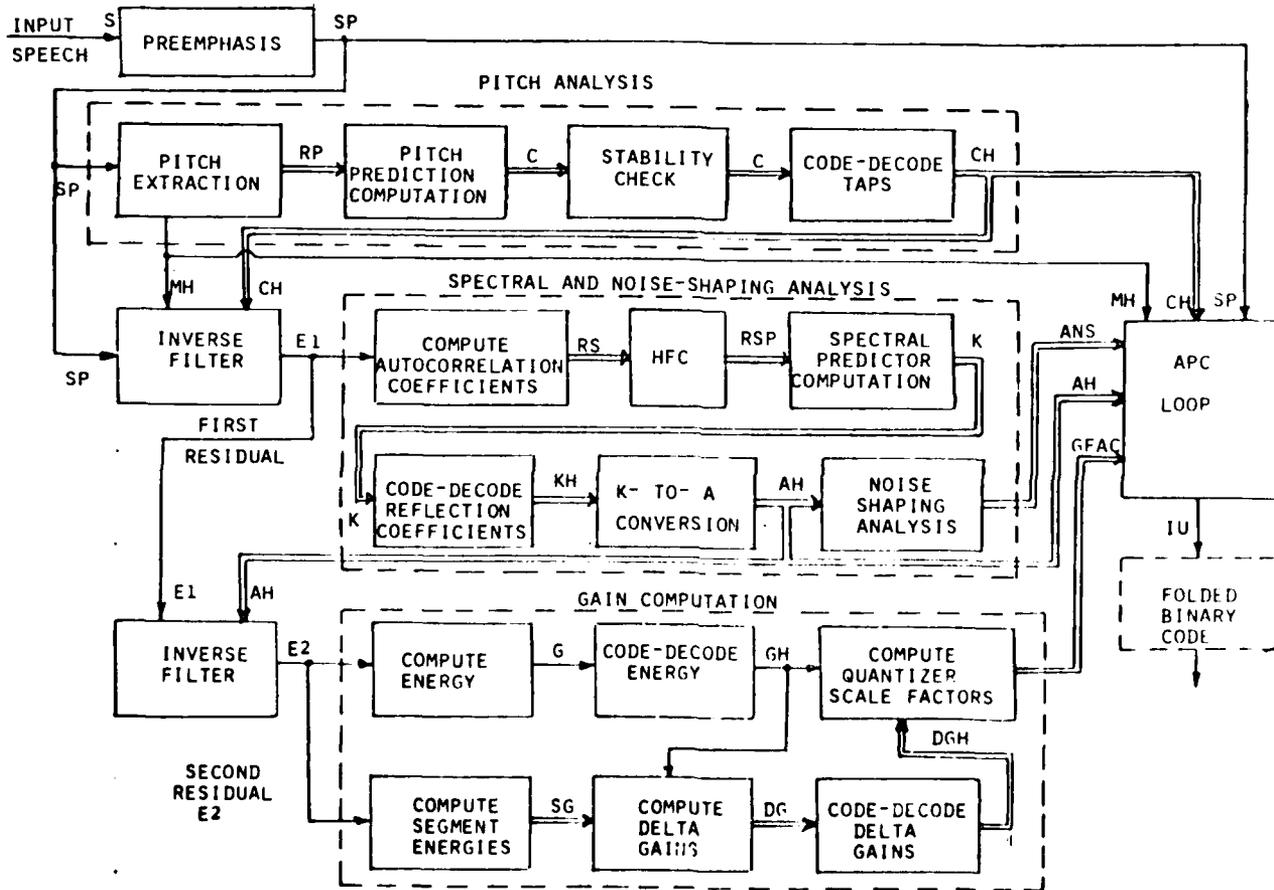


Figure 1. Block diagram of the transmitter of the APC coder

2.1 Preemphasis

$$SP(n) = S(n) - \text{ALPHA} * S(n-1), \quad 1 \leq n \leq 216,$$

where

S = input speech samples (216 samples of the present frame;
last one sample of the past frame)
ALPHA = constant = 0.4
SP = output samples (216 total)

Save the last input sample of the present frame as initial condition of the next frame.

2.2 Pitch Analysis

Pitch analysis consists of the following steps, as shown in Figure 2: pitch extraction, pitch prediction computation, stability check, and coding and decoding of pitch and pitch taps. The symbols given in Figure 2 denote the various quantities as listed in the next page.

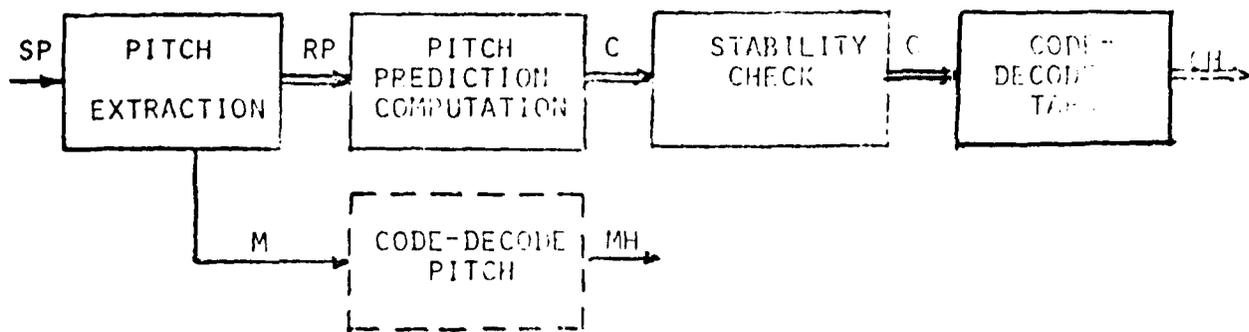


Figure 2. Block diagram for the pitch analysis.

SP = preemphasized speech
RP = autocorrelation coefficients of SP
M = pitch period in number of samples
C = pitch predictor taps: C1,C2,C3
CH = quantized taps
MH = quantized pitch

2.2.1 Pitch Extraction

Pitch extraction uses a frame of 265 preemphasized speech samples (~40 ms): 216 samples of the present frame and 49 samples from the past frame. Pitch extraction consists of the following sequence of operations: remove DC, hamming window, compute autocorrelation coefficients, RP, and compute pitch, M.

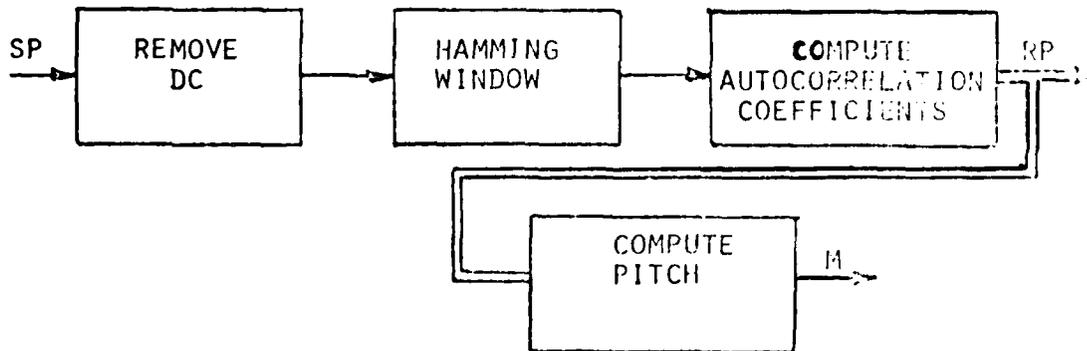


Figure 3. Block diagram for the pitch extraction

2.2.1.1 Remove DC

$$y(n) = x(n) - DC, 1 \leq n \leq N,$$

where

x = input preemphasized samples SP (265 total)
 y = output samples (265 total)
 $DC = 1/N * \sum_{1 \leq i \leq N} x(i), N=265$

2.2.1.2 Hamming Window

$y(n) = x(n) \{ \text{ALPHA} - \text{BETA} * \cos [2\pi(n-1)/(N-1)] \}, 1 \leq n \leq N, N=265,$

where

x = input preemphasized and DC-removed samples, (265 total)
 y = output samples (265 total)
 ALPHA = 0.54
 BETA = 1.0 - ALPHA = 0.46

2.2.1.3 Compute Autocorrelation Coefficients

Direct Method

$RP(m) = \sum_{1 \leq n \leq N-m} x(n) * x(n+m), 0 \leq m \leq MX, N=265,$

where

x = Hamming-windowed input samples (265 total)
 RP(m) = autocorrelation coefficient of lag m
 MX = maximum lag = 134

FFT Method

(a) pad with zeros

$x(n) = 0.0, 266 \leq n \leq 512$

(b) Compute 512-point FFT of x(n)

$X(k) = \text{FFT}(x(n)), 1 \leq k \leq 512$

(c) Compute power spectrum of x(n)

$|X(k)|^2 = [X_R(k)]^2 + [X_I(k)]^2$

where

$X_R(k)$ = real part of $X(k)$
 $X_I(k)$ = imaginary part of $X(k)$

(d) Compute 512-point inverse FFT of $|X(k)|^2$

$V(m) = \text{FFT}^{-1}[|X(k)|^2], 1 \leq m \leq 512$

(e) Autocorrelations are defined as:

$RP(m), = V(m+1), 0 \leq m \leq MX$

NOTE: It is possible to reduce the computation, as follows:

Since input sequence $x(n)$ is real,

$X_R(k)$ is even (i.e. $X_R(k) = X_R(512-k)$), and
 $X_I(k)$ is odd (i.e. $X_I(k) = -X_I(512-k)$)

Therefore, in step (b) compute the lower half of FFT

$X(k), 1 \leq k \leq 257,$

and in step (c) compute the lower half of the power spectrum

$|X(k)|^2, 1 \leq k \leq 257.$

Then, fill $|X(k)|^2$ array from $k = 258$ to 512 as:

$|X(k)|^2 = |X(512-k+2)|^2, 258 \leq k \leq 512.$

Compute steps (d) and (e) as above.

2.2.1.4 Compute Pitch

Search the autocorrelation function $RP(m)$ for a maximum between the range of $m=14$ to $m=133$. Pitch M is computed as the lag, m , at which the autocorrelation coefficient, $RP(m)$, is maximum.

2.2.2 Pitch Prediction Computation

Compute the 3-tap and the 1-tap pitch predictor coefficients. (The 1-tap coefficient is used if the stability check of the 3-tap filter fails).

(a) Compute the 3-tap coefficients from the normal equations:

$$\begin{bmatrix} \text{RP}(0) & \text{RP}(1) & \text{RP}(2) \\ \text{RP}(1) & \text{RP}(0) & \text{RP}(1) \\ \text{RP}(2) & \text{RP}(1) & \text{RP}(0) \end{bmatrix} * \begin{bmatrix} \text{C1} \\ \text{C2} \\ \text{C3} \end{bmatrix} = - \begin{bmatrix} \text{RP}(\text{MH}-1) \\ \text{RP}(\text{MH}) \\ \text{RP}(\text{MH}+1) \end{bmatrix}$$

where

RP = autocorrelation coefficients (6 total)

C1,C2,C3 = pitch predictor coefficients

MH = quantized pitch period (See Section 2.2.5 for pitch quantization)

The solution for the above normal equations may be obtained by the Levinson recursion or from expressions derived by solving the 3 equations. Note that the right-hand-side vector in the above normal equations does not have the elements RP(1), RP(2) and RP(3); this means that the recursive solution used in the standard autocorrelation method cannot be employed here.

(b) Compute the 1-tap coefficient, C2P, as:

$$\text{C2P} = -\text{RP}(\text{MH})/\text{RP}(0)$$

2.2.3 Stability Check

Transform pitch predictor coefficients (C1,C2,C3) as:

$$\text{T1} = \text{C1} + \text{C2} + \text{C3}$$

$$\text{T2} = \text{C1} - 2*\text{C2} + \text{C3}$$

$$\text{T3} = \text{C1} - \text{C3}$$

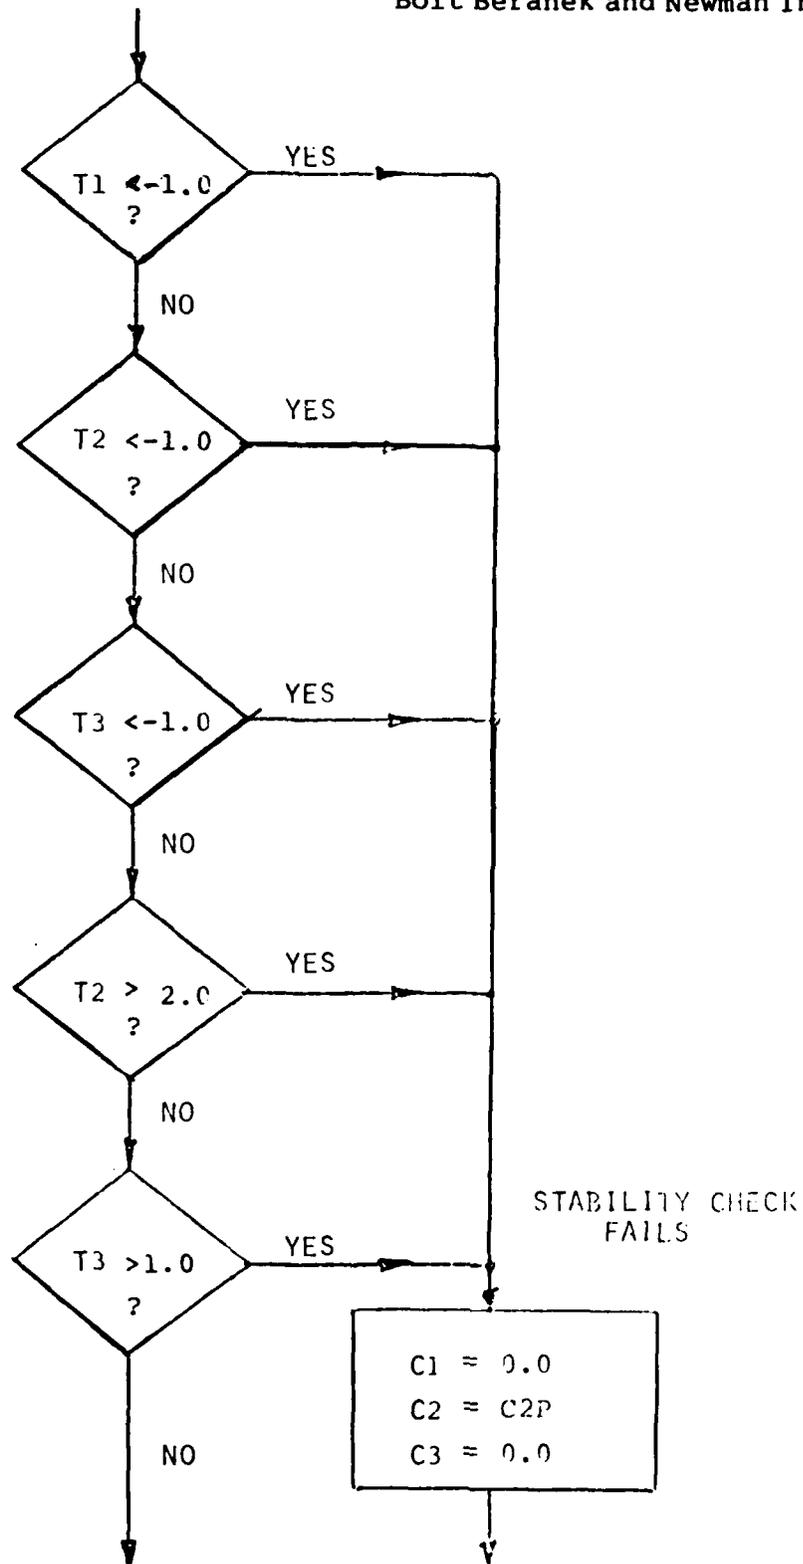


Figure 4. Flowchart for stability check

Check the range of transformed parameters (T_1, T_2, T_3) as shown in Figure 4. Whenever the stability check fails, use the 3-tap predictor $C_1=0, C_2=C_2P$, and $C_3=0$, which is actually the 1-tap predictor.

2.2.4 Code-Decode Taps

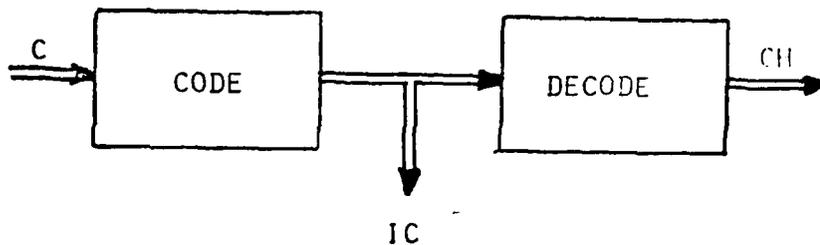


Figure 5. Encoding and decoding of the pitch prediction taps

where

C = pitch predictor coefficients: C_1, C_2, C_3

IC = transmitted codes

CH = quantized values, C_{1H}, C_{2H}, C_{3H}

The coefficients C_1, C_2 , and C_3 are coded using 3, 4, and 3 bits, respectively. The coding and decoding tables for taps are given in Table 1.

2.2.5 Code-Decode Pitch

Since the pitch period M takes integer values in the range 14-133 (a total of 120 values), it is coded directly in 7 bits, as:

$$IM = M - 14$$

Decoded value MH is given by:

$$MH = IM + 14$$

Thus, pitch is quantized without error i.e., $MH = M$

2.3 Inverse Filter to Obtain First Residual

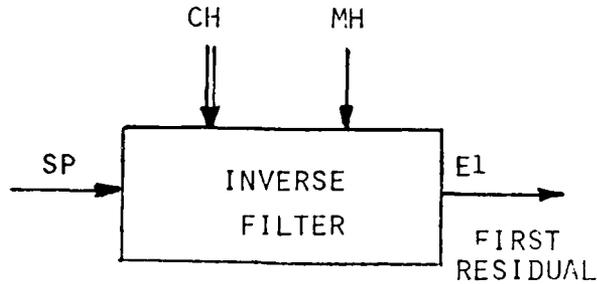


Figure 6. Inverse filter for first residual

$E1(n) = SP(n) + C1H*SP(n-MH+1) + C2H*SP(n-MH)+C3H*SP(n-MH-1)$,
 where

SP = input preemphasized speech samples (216 samples of the present frame and up to 134 samples of the past frame)

CH = quantized pitch predictor coefficients: C1H,C2H,C3H

MH = quantized pitch

E1 = output samples of the first residual (216 total)

2.4 Spectral and Noise-Shaping Analysis

Spectral or linear prediction analysis consists of the following sequential steps: compute autocorrelation coefficients RS, high-frequency correction (HFC), spectral predictor computation, code and decode reflection coefficients, and convert quantized reflection coefficients to predictor coefficients (KH to AH). Noise-shaping analysis involves computing the predictor coefficients ANS from the coefficients AH. In Figure 7, we have used the following terminology:

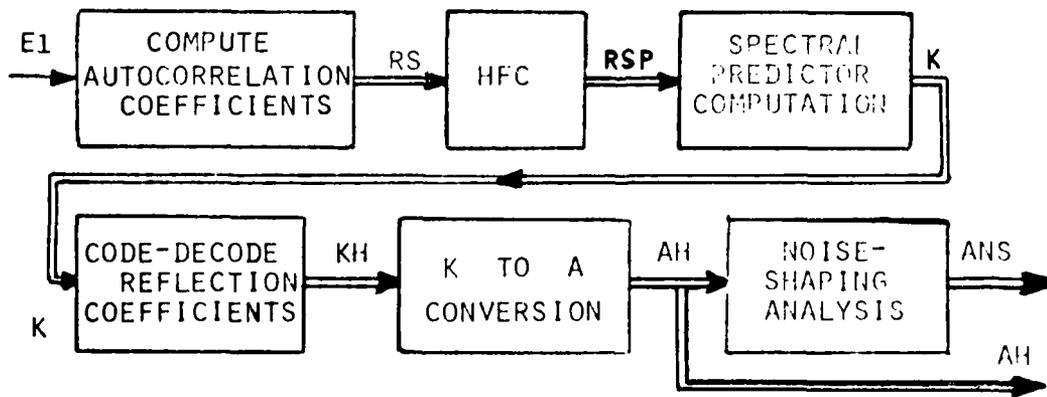


Figure 7. Block diagram of spectral and noise shaping analysis

E1 = first residual samples
RS = autocorrelation coefficients
RSP = high-frequency corrected autocorellation coefficients
K = reflection coefficients
KH = quantized reflection coefficients
AH = predictor coefficients
ANS = noise shaping filter coefficients

2.4.1 Compute Autocorrelation Coefficients

First, Hamming-window 216 samples of the first residual, E1 (see Section 2.2.1.2); then, compute the 7 autocorrelation coefficients $RS(0), RS(1), \dots, RS(6)$, using the direct method, as in 9.6 kb/s BBC coder [1].

2.4.2 High-Frequency Correction (HFC)

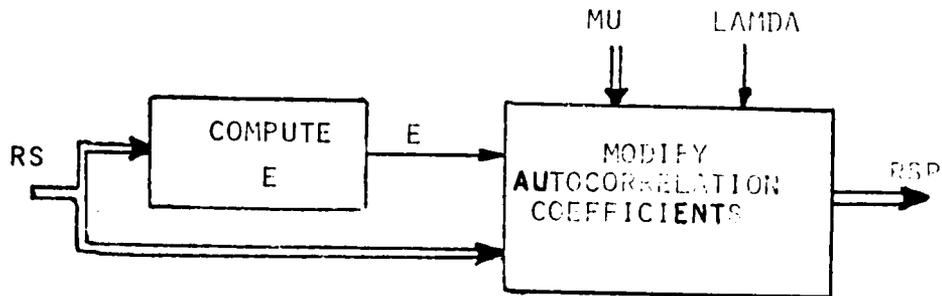


Figure 8. Block diagram of high-frequency correction.

2.4.2.1 Compute minimum mean-squared prediction error, E

$$E = RS(0) * \left[1.0 - \left(\frac{RS(1)}{RS(0)} \right)^2 \right] * \left[1.0 - \left(\frac{RS(1)^2 - RS(0) * RS(2)}{RS(0)^2 - RS(1)^2} \right)^2 \right]$$

where

RS = autocorrelation coefficients

2.4.2.2 Modify autocorrelation coefficients

$$RSP(n) = RS(n) + LAMBDA * E * MU(n), \quad 0 \leq n \leq 2,$$

$$RSP(n) = RS(n), \quad 3 \leq n \leq 6,$$

where

RS = autocorrelation coefficients (7 total)

LAMBDA = 0.035

MU(0) = +0.375

MU(1) = -0.25

MU(2) = +0.0625

RSP = output autocorrelation coefficients (7 total)

2.4.3 Spectral Predictor Computation

Use the standard routine employed in the 9.6 kb/s BBC coder [1].
The input and output quantities are:

RSP = input autocorrelation coefficients (7 total)

K = output reflection coefficients: K(1), K(2), ..., K(6)

2.4.4 Code-Decode Reflection Coefficients

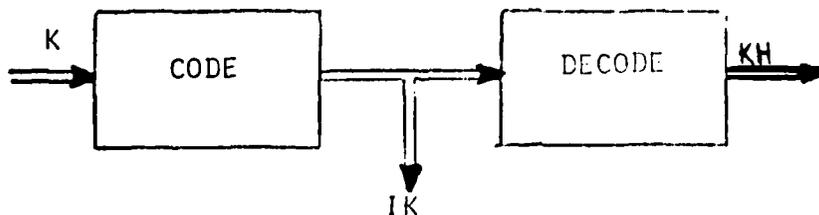


Figure 9. Encoding and decoding of the reflection coefficients

The symbols K, IK and KH used in Figure 9 are defined as follows:

K = reflection coefficients
IK = transmitted codes
KH = quantized reflection coefficients

The reflection coefficients K(1), K(2), K(3), K(4), K(5) and K(6) are coded using 6,5,4,4,4, and 4 bits, respectively. The coding and decoding tables are given in Table 2.

2.4.5 K-to-A conversion

The routine required for K-to-A conversion is the same as the one used in the 9.6 kb/s BBC coder [1], with the following input and output quantities:

KH = input quantized reflection coefficients (6 total)
AH = predictor coefficients for the quantized case (6 total)

2.4.6 Noise-Shaping Analysis

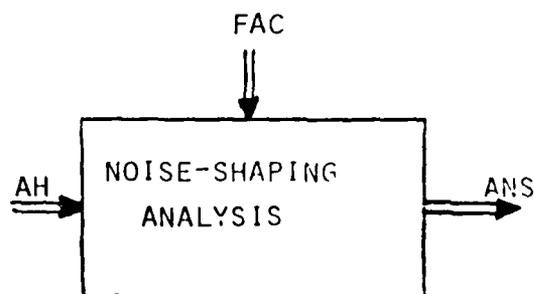


Figure 10. Noise shaping analysis

$$ANS(k) = (FAC)^k * AH(k), \quad 1 \leq k \leq 6$$

where

AH = predictor coefficients for the quantized case (6 total)
ANS = output coefficients (6 total)

The values of $(FAC)^k$ are:

FAC¹ = 0.684128
FAC² = 0.468032
FAC³ = 0.320194
FAC⁴ = 0.219054
FAC⁵ = 0.149861
FAC⁶ = 0.102524

2.5 Inverse Filter to Obtain the Second Residual

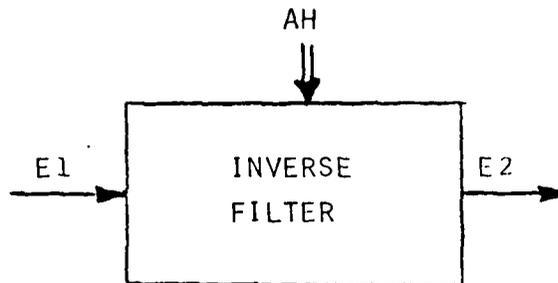


Figure 11. Inverse filter for the second residual

$$E2(n) = E1(n) + \text{SUM}[AH(i) * E1(n-i), 1 \leq i \leq 6]$$

where

E1 = input samples of the first residual (216 samples of present frame; 6 samples of past frame)
AH = predictor coefficients for the quantized case (6 total)
E2 = output samples of the second residual (216 total)

2.6 Gain Computation

The gain computation consists of the following steps: compute and code-decode energy (or mean-squared value) of the second residual, compute segment energies, compute the delta gains and code-decode them, and compute quantizer scale factors.

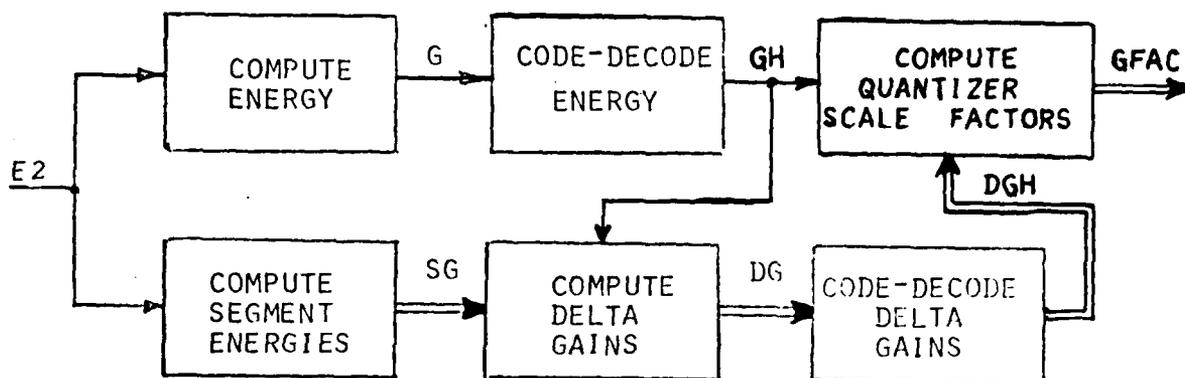


Figure 12. Block diagram of Gain Computation

The symbols used in Figure 12 are defined below:

E2 = second residual samples
 G = energy of E2
 GH = square root of the quantized energy
 SG = segment energies
 DG = delta gains
 DGH = square root of the quantized delta gains
 GFAC = quantizer scale factors

2.6.1 Compute Energy

$$G = 1/N * \text{SUM}[E2(i)^2, 1 \leq i \leq N], N=216$$

where

E2 = input samples of the second residual (216 total)
 G = output energy

2.6.2 Code-Decode Energy

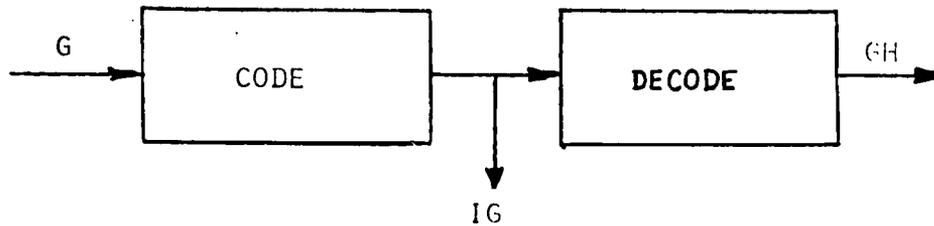


Figure 13. Encoding and decoding of energy

In Figure 13, we have

G = energy
 IG = transmitted code
 GH = square root of the quantized value

The energy G is coded using 6 bits. The coding and decoding tables are given in Table 3.

2.6.3 Compute Segment Energies

$$SG(j) = 1/72 * \text{SUM}[E2(i)^2, (j-1)*72+1 \leq i \leq j*72], j=1,2,3$$

where

E2 = input samples of the second residual (216 total)
 SG = segment energies (3 total)

2.6.4 Compute Delta Gains

$$DG(j) = SG(j) / (GH * GH), \quad j=1,2,3$$

where

DG = delta gains (3 total)
SG = segment energies (3 total)
GH = square root of the quantized energy

2.6.5 Code-Decode Delta Gains

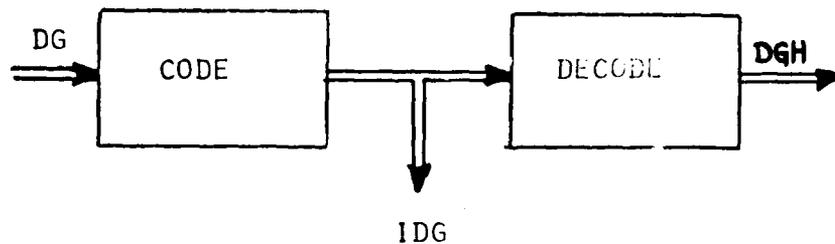


Figure 14. Encode and decode of the delta gains

In Figure 14, we have

DG = delta gains
IDG = transmitted codes
DGH = quantized values

The delta gains are coded using 2 bits each. Table 4 contains the coding and decoding tables for the delta gains.

2.6.6 Compute Quantizer Scale Factors

$$GFAC(j) = [GH * DGH(j)] \quad j=1,2,3$$

where

DGH = quantized delta gains (3 total)
GH = square root of the quantized energy
GFAC = quantizer scale factors (3 total)

2.7 APC loop

For each input preemphasized speech sample $SP(i)$, the following steps are performed.

- Steps 1-4. Compute predictions (4 total)
- Step 5. Compute APC residual, WO
- Step 6. Normalize APC residual
- Steps 7,8. Code-decode residual
- Step 9. Scale quantized residual
- Steps 10-13. Compute Q and update arrays; $Q1, VH, RH$

The output from each of these steps is marked in the block diagram of the APC loop, given in Figure 16, by a circled number; these numbers indicate the order in which the outputs are computed.

2.7.1 Compute Predictions for Sample i

(a) For noise shaping and spectral predictors (Steps 1,2,3)

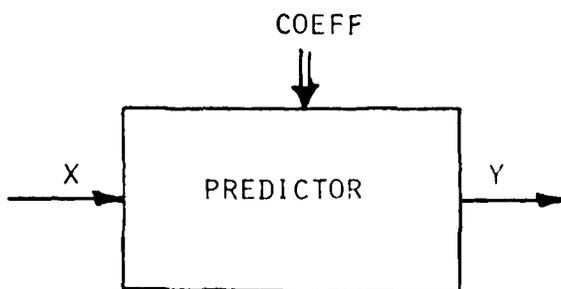


Figure 15. Noise shaping and spectral predictor in the APC loop

$$Y = \text{SUM}[\text{COEFF}(j) * x(i-j), 1 \leq j \leq 6], \text{ for sample } i$$

PREDICTOR TYPE	INPUT X	RANGE OF X	COEFF VECTOR	# OF COEFF	OUTPUT Y
Noise-shaping (zeros)	$Q1(j)$	$i-6 \leq j \leq i-1$	ZNS	6	QP
Noise-Shaping (poles)	$Q1(j)$	$i-6 \leq j \leq i-1$	AP	6	QP1
Spectral Predictor	$VH(j)$	$i-6 \leq j \leq i-1$	AH	6	VHP

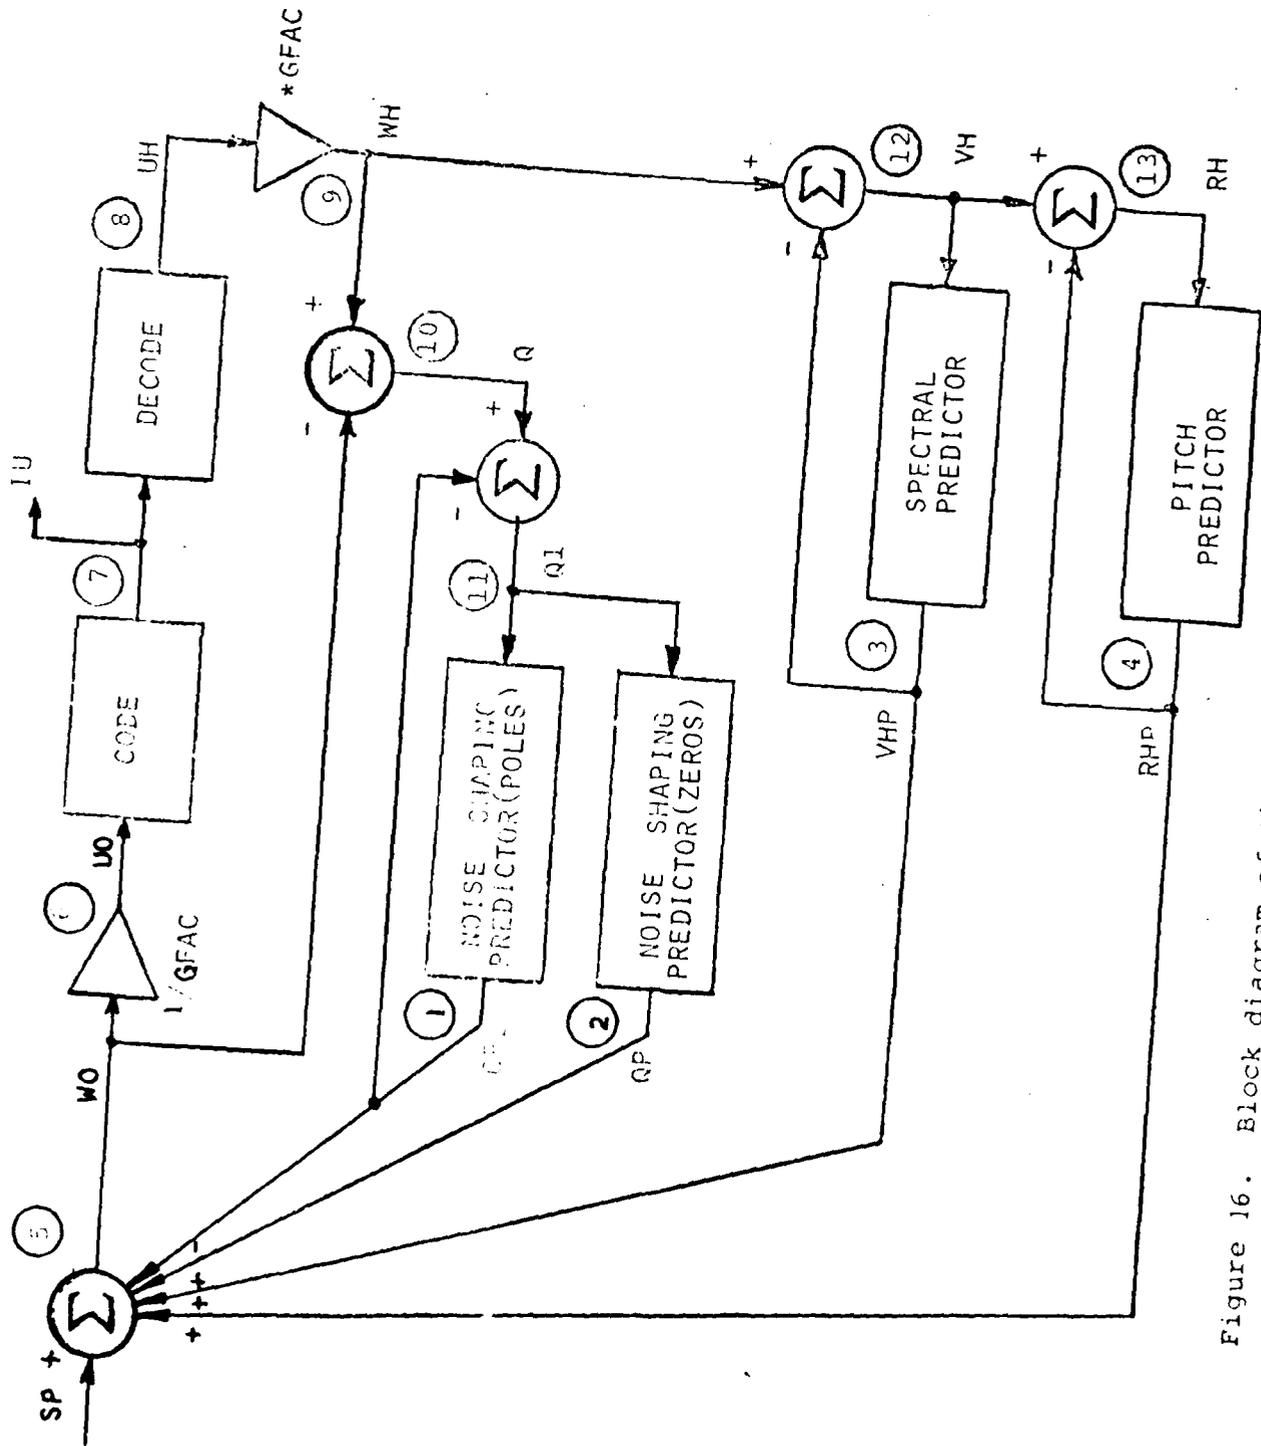


Figure 16. Block diagram of the APC loop

(b) For pitch predictor (Step 4)

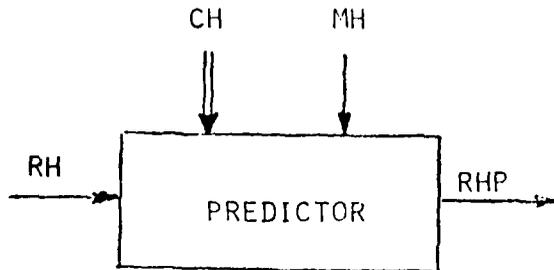


Figure 17. Pitch predictor in the APC loop

$$RHP = C1H \cdot RH(i-MH+1) + C2H \cdot RH(i-MH) + C3H \cdot RH(i-MH-1), \text{ for sample } i$$

where

RH = input samples (range from $i-MH-1$ to $i-MH+1$ for sample i)
 CH = quantized pitch predictor coefficients: $C1H, C2H, C3H$
 MH = quantized pitch
 RHP = output prediction

2.7.2 Compute APC Residual, WO (Step 5)

$$WO = SP(i) + QP - QP1 + RHP + VHP$$

2.7.3 Normalize APC Residual (Step 6)

$$UO = WO / GFAC(j)$$

where, for the i -th residual sample,

$j=1$ for $1 < i < 72$
 $j=2$ for $73 < i < 144$
 $j=3$ for $145 < i < 216$

2.7.4 Code-Decode Normalized Residual (Steps 7 and 8)

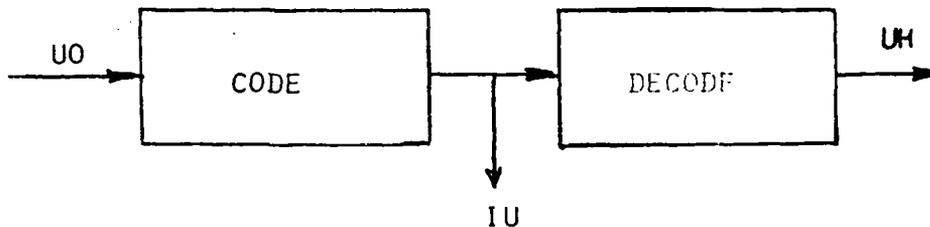


Figure 18. Encoding and decoding of the normalized residual in the APC loop

In Figure 18, we have

UO = normalized residual sample
 IU = transmitted code
 UH = quantized value

The normalized residual is coded using 2 bits. Table 5 gives the coding and decoding tables.

2.7.5 Scale the Quantized Residual (Step 9)

$$WH = UH * GFAC(j)$$

where, for the i -th sample,

$$\begin{aligned}
 j=1 & \text{ for } 1 < i < 72 \\
 j=2 & \text{ for } 73 < i < 144 \\
 j=3 & \text{ for } 145 \leq i < 216
 \end{aligned}$$

2.7.6 Update Arrays (Steps 10-13)

Arrays are updated in the following sequence:

- (a) $Q = WH - WO$
- (b) $Q1(i) = Q - QP1$
- (c) $VH(i) = WH - VHP$
- (d) $RH(i) = VH(i) - RHP$

Repeat the Steps 1 to 13 for each of the 216 input samples.

2.8 Folded Binary Code (FBC) for Encoding the Residual Samples

The FBC encoding may be performed as a separate operation or may be included as part of the coding table.

- (a) Coding table approach: interchange the "J" values of 0 and 1 in Table 5.
- (b) Separate operation: The residual is coded as in Table 5; then the codes are interchanged as in Figure 19. At the receiver the codes are again interchanged as in Figure 19 and then decoded as in Table 5.

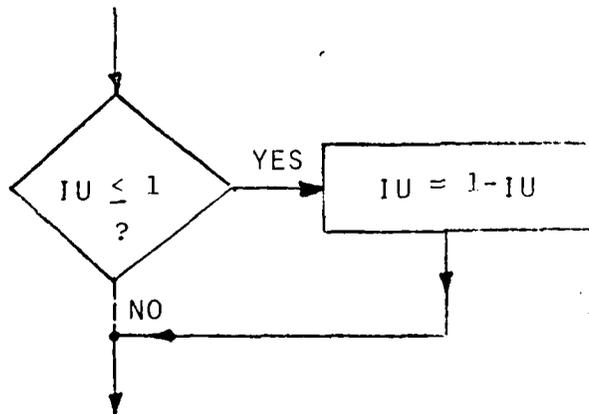


Figure 19. Flowchart for folded binary code (FBC); separate operation

3. RECEIVER

A block diagram of the receiver is given in Figure 20. In this section, we give specification for each of the receiver components.

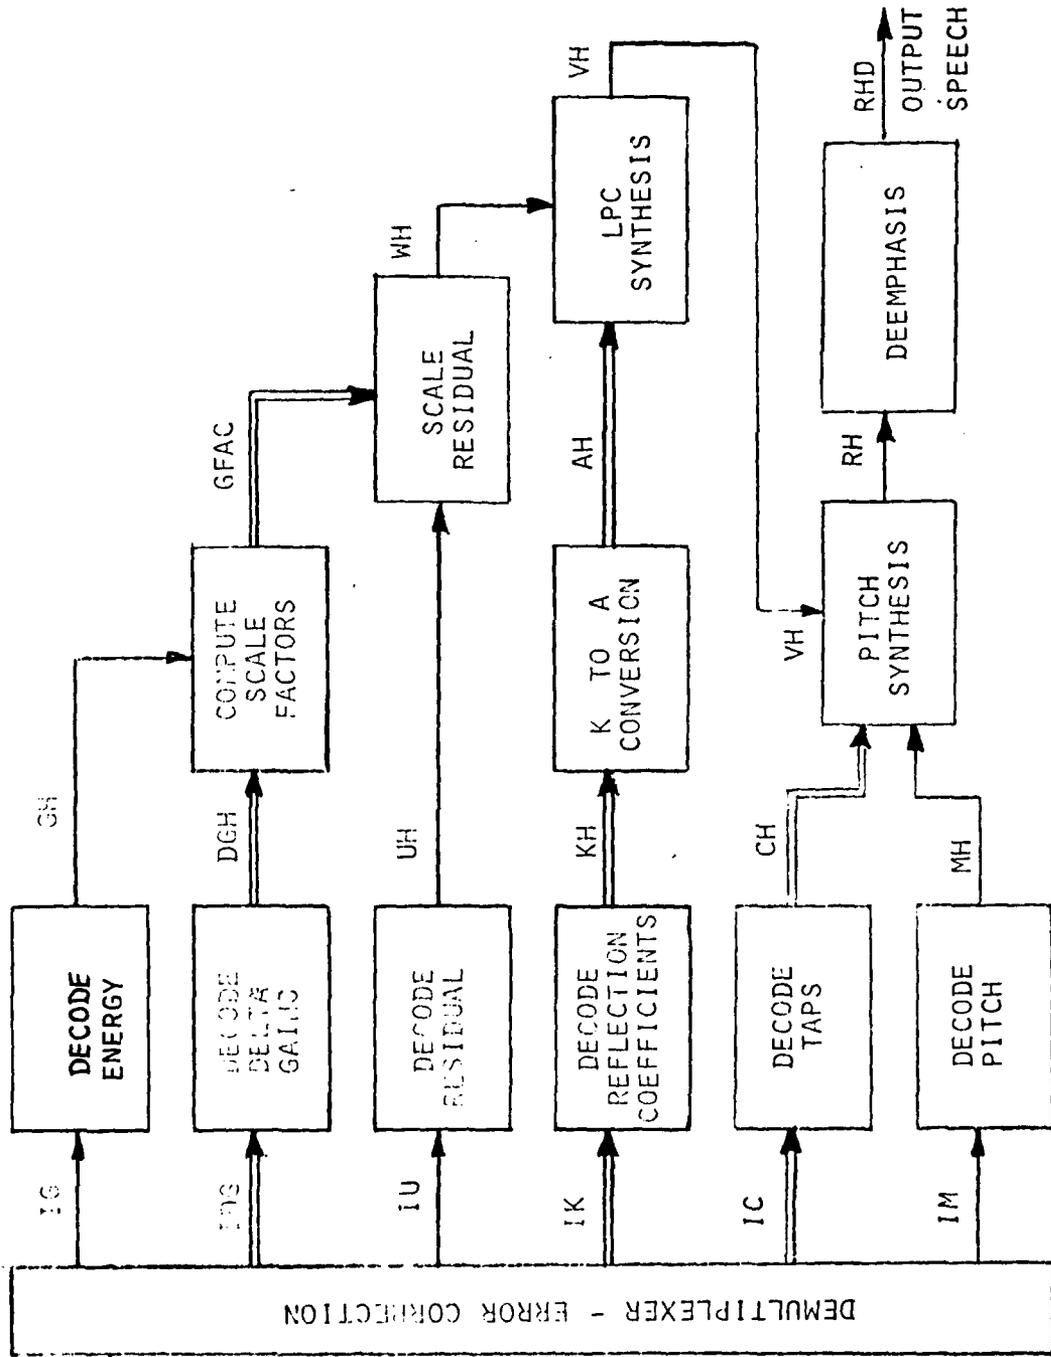


Figure 20. Block diagram of the receiver of the APC coder

3.1 Decode Parameters

Decode the received parameters using appropriate tables listed below.

<u>Parameter</u>	<u>Tables</u>
Energy	3
Delta Gains	4
Residual *	5
Reflection Coefficients	2
Pitch taps	1
Pitch	see Section 2.2.5

*Decode folded binary code (FBC) as shown in Fig. 19, if FBC encoding is done as a separate operation rather than as part of the coding table. (See Section 2.8).

3.2 Compute scale factors

See Section 2.6.6

3.3 Scale the quantized residual

See Section 2.7.5

3.4 K-to-A conversion

See Section 2.4.5

3.5 LPC synthesis

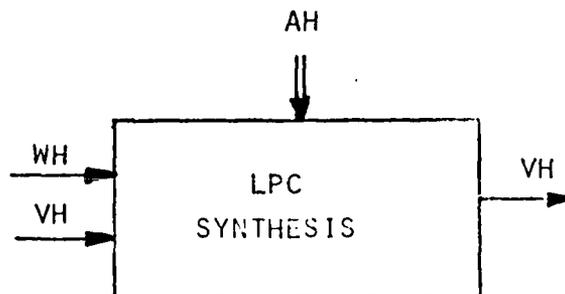


Figure 21. LPC synthesis

$$VH(i) = WH(i) - \text{SUM}[AH(j)*VH(i-j), 1 \leq j \leq 6], 1 \leq i \leq 216$$

where

WH = input residual samples (216 total)
 AH = predictor coefficients (6 total)
 VH = for each output sample i, six prior values (i-1 to i-6) of VH are required as input.

Save the last 6 samples of VH as initial condition of next frame. Also, the functions in Sections 3.4 and 3.5 can be combined into one, if lattice-form synthesis is used.

3.6 Pitch Synthesis

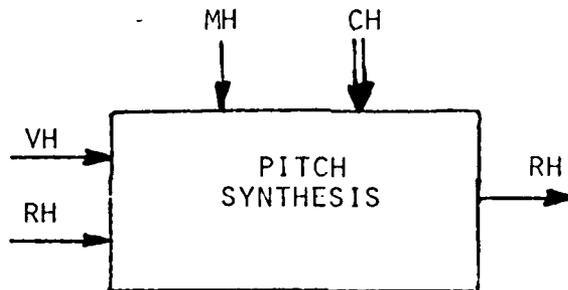


Figure 22. Pitch synthesis

$$RH(i) = VH(i) - C1H * RH(i - MH + 1) - C2H * RH(i - MH) - C3H * RH(i - MH - 1)$$

where

VH = input samples (216 total)
 MH = pitch
 CH = pitch predictor taps: C1H, C2H, C3H
 RH = for each output sample i, three prior values: i-MH+1, i-MH, i-MH-1 of RH are required as input.

Save the last 134 samples of RH as initial condition of next frame.

3.7 Deemphasis

$$\text{RHD}(i) = \text{RH}(i) + \text{ALPHA} * \text{RHD}(i-1), \quad 1 \leq i \leq 216$$

where

RH = input samples (216 total)

ALPHA = constant = 0.4

RHD = output speech samples (216 total)

Save the last sample of RHD as initial condition of next frame.
Note that RHD is the synthesized output speech.

REFERENCE

1. R. Viswanathan, J. Wolf, L. Cosell, K. Field, A. Higgins and W. Russell, "Design and Real-Time Implementation of a Baseband LPC Coder for Speech Transmission over 9600 BPS Noisy Channels", Final Report, BBN Report No. 4327, Vol. I and II, Bolt Beranek and Newman Inc., Cambridge, MA, February 1980.

Table 1(a)

CODING AND DECODING TABLES FOR PITCH TAPS

C1,C3 (3 Bits)

X(J)	J	R(J)
$-\infty$	0	-.488
-.427	1	-.366
-.305	2	-.244
-.183	3	-.122
-.061	4	.000
.061	5	.122
.183	6	.244
.305	7	.366
∞		

Table 1(b)
 CODING AND DECODING TABLES FOR PITCH TAP
 C2 (4 Bits)

X(J)	J	R(J)
-∞	0	-.91656251
-.88312500	1	-.84963750
-.81625000	2	-.78281250
-.74937499	3	-.71593750
-.68249999	4	-.64906249
-.61562499	5	-.58213749
-.54874998	6	-.51531249
-.48187498	7	-.44843748
-.41499998	8	-.38156248
-.34812498	9	-.31463748
-.28124999	10	-.24781249
-.21437499	11	-.18093749
-.14749999	12	-.11406249
-.08062499	13	-.04718749
-.01374999	14	.01966752
.05312502	15	.08655252
∞		

Table 2(a)

CODING AND DECODING TABLES FOR THE REFLECTION COEFFICIENT

K(1) (6 BITS)

X(J)	J	R(J)
- ∞	0	-.98623391
-.98540065	1	-.98451736
-.98358107	2	-.98258864
-.98153679	3	-.98042203
-.97924066	4	-.97798879
-.97666232	5	-.97525692
-.97376801	6	-.97219076
-.97052009	7	-.96875066
-.96687680	8	-.96489257
-.96279173	9	-.96056771
-.95821358	10	-.95572206
-.95308557	11	-.95029607
-.94734517	12	-.94422411
-.94092366	13	-.93743417
-.93374563	14	-.92984749
-.92572883	15	-.92137819
-.91678373	16	-.91193306
-.90681341	17	-.90141150
-.89571356	18	-.88970540
-.88337240	19	-.87669948
-.86967121		

Table 2(a)
K(1) (cont.)

X(J)	J	R(J)
-.85448478	20	-.86227170
-.83768241	21	-.84629394
-.81912912	22	-.82863319
-.79868747	23	-.80915297
-.77621973	24	-.78771539
-.75159076	25	-.76418365
-.72467132	26	-.73842508
-.69534210	27	-.71031488
-.66349818	28	-.67974033
-.62905408	29	-.64660559
-.59194915	30	-.61083690
-.55215304	31	-.57238807
-.50967117	32	-.53124595
-.46454968	33	-.48743577
-.41637981	34	-.44102574
-.36680007	35	-.39213083
-.31450189	36	-.34091512
-.26022125	37	-.28759249
-.20424442	38	-.23242530
-.14609926	39	-.17572091
-.08854935	40	-.11782601
-.02958530	41	-.05911889
.02958537	42	.00000004
.08854942	43	.05911896

Table 2(a)
K(1) (cont.)

X(J)	J	R(J)
.14689934	44	.11782609
.20424450	45	.17572098
.26022132	46	.23242537
.31450195	47	.28759256
.36680095	48	.34091519
.41687988	49	.39213091
.46454975	50	.44102581
.50967123	51	.48743584
.55215311	52	.53124601
.59194921	53	.57230812
.62905414	54	.61083695
.66349822	55	.64660564
.69534215	56	.67974038
.72467136	57	.71031493
.75159079	58	.73842513
.77621976	59	.76410369
.79868749	60	.78771544
.81912915	61	.80915300
.83763244	62	.82863321
	63	.84629396

Table 2(b)

CODING AND DECODING TABLES FOR THE REFLECTION COEFFICIENT

K(2) (5 BITS)

X(J)	J	R(J)
- ∞	0	-.74930506
-.73101083	1	-.71160168
-.69104202	2	-.66929988
-.64634765	3	-.62216299
-.59672958	4	-.57003794
-.54200627	5	-.51288126
-.48243376	6	-.45078440
-.41795417	7	-.38399474
-.34896366	8	-.31292941
-.27597119	9	-.23817849
-.19965046	10	-.16049504
-.12032709	11	-.08077108
-.04045165	12	.00000000
.04045156	13	.08077099
.12032701	14	.16049496
.19965037	15	.23817841
.27597111	16	.31292932
.34896358	17	.38399466
.41795409	18	.45078433
.48243363	19	.51288119
.54200621	20	.57003786
.59672952		

Table 2(b)
K(2) (cont.)

X(J)	J	R(J)
.64634760	21	.62216294
.69104198	22	.66929983
.73101079	23	.71160164
.76652286	24	.74930502
.79789401	25	.78270507
.82546644	26	.81213313
.84959197	27	.83793808
.87061901	28	.86047144
.88888308	29	.88007615
.90470025	30	.89707866
•	31	.91178370

Table 2(c)

CODING AND DECODING TABLES FOR REFLECTION COEFFICIENT

K(3) (4 BITS)

X(J)	J	R(J)
- ∞	0	-.75245864
-.72468231	1	-.69433287
-.66129490	2	-.62547535
-.58680991	3	-.54526938
-.50086620	4	-.45366014
-.40376312	5	-.35134251
-.29662241	6	-.23988272
-.18145542	7	-.12171834
-.06108626	8	.00000000
.06108626	9	.12171833
.18145541	10	.23988271
.29662241	11	.35134250
.40376312	12	.45366013
.50086619	13	.54526937
.58680990	14	.62547535
.66129489	15	.69433286
∞		

Table 2(d)

CODING AND DECODING TABLES FOR REFLECTION COEFFICIENT

K(4) (4 BITS)

X(J)	J	R(J)
- ∞	0	-.57005347
-.53250506	1	-.49273689
-.45079815	2	-.40677000
-.36076742	3	-.31293977
-.26347055	4	-.21257565
-.16050062	5	-.10751649
-.05391451	6	.00000000
.05391451	7	.10751648
.16050061	8	.21257564
.26347053	9	.31293977
.36076741	10	.40677000
.45079814	11	.49273689
.53250505	12	.57005347
.60530267	13	.63843988
.66931576	14	.69804088
.72468230	15	.74932021
∞		

Table 2(e)

CODING AND DECODING TABLES FOR REFLECTION COEFFICIENT

K(5) (4 BITS)

X(J)	J	R(J)
- ∞	0	-.50890809
-.47343143	1	-.43634928
-.39772316	2	-.35763384
-.31618157	3	-.27348579
-.22968443	4	-.18493246
-.13940015	5	-.09327063
-.04673719	6	.00000000
.04673717	7	.09327062
.13940014	8	.18493246
.22968442	9	.27348579
.31618155	10	.35763383
.39772316	11	.43634928
.47343142	12	.50890808
.54273631	13	.57489084
.60536267	14	.63415766
.66129489	15	.68680496
∞		

Table 2(f)

CODING AND DECODING TABLES FOR REFLECTION COEFFICIENT

K(6) (4 BITS)

X(J)	J	R(J)
- ∞		
-.32298331	0	-.36516072
-.23480363	1	-.27948269
-.14258225	2	-.18910990
-.04781699	3	-.09541580
.04781697	4	.00000000
.14253224	5	.09541579
.23430362	6	.18910989
.32298330	7	.27948267
.40589048	8	.36516071
.48261537	9	.44506935
.55258511	10	.51846758
.61554603	11	.58494609
.67152148	12	.64439614
.72075575	13	.69695902
.76365393	14	.74296685
∞	15	.78288344

Table 3
CODING AND DECODING TABLES FOR ENERGY (6 BITS)

X(J)	J	R(J)
-	0	0.33256382
0.12232071	1	0.36781124
0.14962357	2	0.40679443
0.18302761	3	0.44990933
0.22387212	4	0.49759385
0.27384197	5	0.55033230
0.33496544	6	0.50866034
0.40973211	7	0.67317039
0.50118723	8	0.74451765
0.61305579	9	0.82342681
0.74989421	10	0.91069929
0.91727594	11	1.00722152
1.12201344	12	1.11397386
1.37246095	13	1.23204054
1.67300398	14	1.36262074
2.05352500	15	1.50704074
2.51168536	16	1.56676740
3.07255733	17	1.84342298
3.75037393	18	2.03880173
4.59726077	19	2.25488809
5.62341309	20	2.49387682

Table 3
ENERGY (cont.)

X(J)	J	R(J)
5.07859881	21	2.75919513
8.41395115	22	3.05052787
10.27200480	23	3.37384397
12.58925340	24	3.73142737
15.39926480	25	4.12690991
16.83649020	26	4.56430846
23.04092360	27	5.04806560
25.15332310	28	5.58309466
34.47455350	29	6.17482972
42.15954320	30	6.82928115
51.50221390	31	7.55309576
63.09572790	32	8.35362506
77.17914580	33	9.23900044
94.40610220	34	10.21821390
115.47818800	35	11.30121100
141.25374100	35	12.49899180
172.75250900	37	13.32372160
211.34888500	38	15.28885570
255.52346400	39	16.90927510
316.22775300	40	18.70143700
390.51201900	41	20.58354540
473.15122200	42	22.87573150

Table 3
ENERGY (cont.)

X(J)	J	R(J)
578.76192500	43	25.30025960
757.94572500	44	27.93175740
865.95427900	45	30.94745760
1059.25362800	46	34.22748420
1295.63658000	47	37.35515120
1584.89297400	48	41.36730190
1938.65242400	49	46.33468990
2371.37352000	50	51.21238140
2900.68485300	51	56.54022640
3548.13379800	52	62.54335060
4340.10215300	53	69.28272910
5308.04351400	54	76.52579440
6493.11567200	55	84.74712280
7943.28168700	56	93.72921370
9716.27893000	57	103.66328700
11835.32250000	58	114.65023700
14537.84370000	59	126.90167400
17782.79220000	60	140.24098800
21752.03560000	61	155.10471300
26607.24850000	62	171.54377700
32545.17310000	63	189.72517200

Table 4

CODING AND DECODING TABLES FOR DELTA GAINS (2 BITS EACH)

X(J)	J	R(J)
$-\infty$		
0.43651563	0	0.48977882
0.09125794	1	0.79432824
1.05958589	2	1.12201844
∞	3	1.49623564

Table 5

CODING AND DECODING TABLES FOR RESIDUAL SAMPLES (2 BITS)

(NATURAL BINARY CODE)

X(J)	J	R(J)
$-\infty$		
-0.97517319	0	-1.53489251
0.48758659	1	-0.45145388
0.97517319	2	0.45145333
∞	3	1.53489251

APPENDIX B

USER'S GUIDE
FOR THE FORTRAN SIMULATION OF
THE BBN'S 16 kb/s APC CODER

August 1980

Prepared by

W. Russell, A. Higgins, and R. Viswanathan
Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

TABLE OF CONTENTS

1. INTRODUCTION.	215
2. SIMULATION SOFTWARE	216
3. TYPICAL USER SESSION.	224
4. SIMULATION OF SOME VARIATIONS OF THE OPTIMIZED APC CODER	226
4.1 Flags.	226
4.2 Variables.	228
5. CODER OPERATION WITH CHANNEL-ERROR SIMULATION	232

1. INTRODUCTION

This guide provides information necessary to use the FORTRAN simulation of the BBN 16 kb/s APC coder. Installation of this simulation on the user's computer system will require some software modifications. These modifications are specified in detail in Section 2 of this guide. In Section 3, a typical user session is described. Section 4 outlines how the user may alter the operation of the coder by resetting various flags and coder parameters. The simulation of the coder operating in the presence of channel bit-errors is discussed in Section 5.

2. SIMULATION SOFTWARE

The simulation software consists of a main program 16KMN, and the following five subroutine packages:

1. 16KIO - File handling and data I/O routines
2. 16KGEN - General utility routines
3. 16KCOD - Quantization, encoding and decoding routines
4. 16KER - Channel bit-error simulation, error-protection, and error-correction routines
5. 16KPR - All other routines

The FORTRAN listings of the main program and of each of the five subroutine packages are given in Appendix C. The simulation also uses one routine from the IBM scientific subroutine package (Routine NDTR for evaluating the normal distribution function, called by the subroutine OPTQ in the 16KCOD package) and several routines from our BBN speech library package. The FORTRAN code for these latter routines is not included in the supplied software, since they have been designed specifically for the BBN computer system. For the user's reference, a list of these routines from the BBN speech library, their calling sequence and a brief description of their purpose are given at the end of this section.

The user must substitute his own software to perform the tasks of the missing routines. The locations within the main module

16KMN, where substitutions should be made, are specially marked with a string of asterisks and comments. The steps required to perform the substitutions are listed below.

1. Speech I/O

The BBN simulation system employs disk files for speech I/O, in that digitized speech samples are read into a buffer from an input disk file and processed speech samples are written out from a buffer into an output disk file. (The input and output disk files may be compared using a separate D/A playout program, which is not part of the simulation software.) The following parts of the main program 16KMN have to be modified to suit the user's I/O facility.

- a) Specification statements for file handling: The specification statements at the top of the main program labeled "DATA FOR FILE HANDLING" should be replaced with appropriate ones that may be needed for the user's specific speech I/O.
- b) Opening input and output speech files: The user must replace the code labeled "OPEN INPUT AND OUTPUT SPEECH FILE" below statement 100 in the main program and the subroutines OPNIF and OPNOF in the 16KIO package, with his own software to provide access for input and output speech samples. Also, at this place in the main program, the quantity NFRAME (number of samples/frame) must be computed from the sampling frequency in Hz, FREQS and the frame size in ms, TFRAME. (The BBN-specific subroutine OPNIF reads in the value of FREQS from the header of the input speech file, or allows the user to specify it in the case of an unheadered file.)
- c) Reading in speech samples: The code labeled "READ IN NFRAME SAMPLES" after statement 2000 in the main program and the subroutine ISAMP in the 16KIO package should be replaced with the user's own code to read in NFRAME number

of speech samples. These samples should be stored as floating-point numbers in the buffer SPEECH starting from the location N2S. The user's code must also check for the end of the input speech data. When the end is detected, the program control should be transferred to statement 2010.

- d) Writing-out speech samples: The user must replace the code labeled "OUTPUT SAMPLES" below statement 4000 in the main program and the subroutine OSAMP in the 16KIO package, with his own code to write out NFRAME number of output speech samples from the buffer SLAST, starting from location N2S.
- e) Closing input and output speech files: The user must replace the code labeled "CLOSE FILES" around statement 4050 with his own code to close the input and output access.

2. FFT of Real Data

- a) The subroutine PITCH in the package 16KPR calls another subroutine FFTR to perform FFT of real data. A description of FFTR is given at the end of this section. The user must replace FFTR with his own subroutine.
- b) A related subroutine WRWI is called by the main program (after statement 5 at the top of the program) to set up cosine table to be used by FFTR. The user must either remove this call or replace with another depending upon how his own FFTR subroutine is organized.

3. Random Number Generation

- a) The subroutine ERRCHN in the package 16KER calls another subroutine RANDOM to generate pseudo-random numbers. A description of RANDOM may be found at the end of this section. Again, the user must replace RANDOM with his own subroutine.
- b) A related subroutine ZETRAN is called by the main program (after statement 10 at the top of the program) to initialize the random-number generator at a prespecified point. This is necessary if one wants to employ an identical sequence of random numbers in two separate experiments. Again, the user must either remove this call or substitute ZETRAN with his own subroutine.

Software Change to Suit Different Input-Speech Wordlengths

The BBN 16 kb/s APC coder has been designed with the assumption that its input is 11-bit (including the sign bit) linear PCM speech. (To avoid a possible confusion, note that we store input speech samples using 12 bits each after extending or duplicating the sign bit to the left, and that three such 12-bit samples are packed in one 36-bit computer word.) If the user plans to use a different linear PCM speech as coder input, he must change the gain quantization ranges in dB, GMAX and GMIN, in the subroutine QTBLG (16KCOD routine package), to ensure proper gain quantization. This is accomplished by setting the value of the quantity DBCHANG, specified via a DATA statement in the same subroutine, to be equal to 6 times (actual speech sample size in bits -11). The factor 6 is due to the 6 dB/bit rule. For example, if the coder input is 9-bit linear PCM speech, then DBCHANG = -12.0.

A LIST OF SUBROUTINES FROM THE BBN SPEECH LIBRARY

INIT - CLOSSES ALL OPEN FILES AND INITIALIZES SYSTEM
CALL INIT

OPENIF - OPENS FILE AS INPUT FILE

CALL OPENIF(JFN,BYTSIZ)
CALL OPENIF(JFN,BYTSIZ,FILNAM)
CALL OPENIF(JFN,BYTSIZ,FILNAM,IERR)

JFN = JOB FILE NUMBER, RETURNED BY SUBROUTINE

BYTSIZ = FILE BYTE SIZE

FILNAM = POINTER TO FILE NAME. IF THIS ARGUMENT IS 0 OR NOT
GIVEN THEN THE FILE NAME IS TO BE TYPED IN.

IERR = OPTIONAL ERROR STATE ARG. IF NOT GIVEN, FILE OPENING
ERRORS WILL BE HANDLED BY THE IO ERROR HANDLER. IF THIS
ARG IS GIVEN, THEN THIS SUBROUTINE WILL ALWAYS RETURN,
WITH IERR=0 IF THE FILE WAS OPENED. IF THE FILE WASN'T
OPENED, THEN IERR=JSYS ERROR CODE AND RETURNED JFN = -1.

OPENOF - OPEN FILE AS OUTPUT FILE

CALL OPENOF(JFN,BYTSIZ)
CALL OPENOF(JFN,BYTSIZ,FILNAM)
CALL OPENOF(JFN,BYTSIZ,FILNAM,IERR)

ARGUMENTS SAME AS FOR OPENIF

CLOSF - CLOSSES FILE, GIVEN JFN

CALL CLOSF(JFN,NOREL)

NOREL = OPTIONAL ARGUMENT: IF GIVEN AND NONZERO, THE FILE IS
CLOSED WITHOUT RELEASING THE JFN. IF ZERO OR NOT GIVEN,
THE FILE IS CLOSED AND/OR THE JFN RELEASED, AS
APPROPRIATE

FILNAM - GETS FILE NAME, GIVEN JFN

CALL FILNAME (JFN,ARRAY)

JFN = JOB FILE NUMBER OF FILE

ARRAY = POINTER TO ARRAY WHERE FILE NAME IS TO BE STORED

SFBSZ - (RE) SETS FILE BYTE SIZE

CALL SFBSZ (JFN,IBSIZE)

JFN = JOB FILE NUMBER

IBSIZE = NEW BYTE SIZE

SFPTR - SETS FILE POINTER

CALL SFPTR (JFN,NBYTE)

JFN = JOB FILE NUMBER

NBYTE = BYTE NO. TO WHICH POINTER IS TO BE SET

=-1, WILL POINT TO CURRENT END OF FILE

RFPTR - READS FILE POINTER
CALL RFPTR(JFN,NBYTE)
JFN = JOB FILE NUMBER
NBYTE = BYTE NUMBER OF POINTER IN FILE

SINB - STRING INPUT FROM FILE WITH ARBITRARY BYTE SIZE
CALL SINB(EOF,JFN,POINTR,BYTSIZ,NBYTE)
CALL SINB(EOF,JFN,POINTR,BYTSIZ,NBYTE,ENDCHR)
EOF =ASSIGNED STATEMENT NO. FOR END OF FILE TRANSFER
JFN =JOB FILE NUMBER
POINTR = POINTER TO ARRAY WHERE STRING IS TO BE STORED
BYTSIZ = BYTE SIZE IN ADDRESS SPACE; IT CAN BE DIFFERENT FROM
FILE BYTE SIZE. BYTE IS ALWAYS RIGHT-JUSTIFIED WITH
EXTRA ZEROS TO THE LEFT OR TRUNCATION IF NECESSARY
DEPENDING ON THE RELATION BETWEEN THE TWO BYTE SIZES.
NBYTE = NUMBER OF BYTES ACCORDING TO THE FOLLOWING:
=0, ZERO BYTE TERMINATES
>0, EXACT BYTE COUNT
<0, NEGATIVE BYTE COUNT OR A BYTE OF -1, WHICHEVER COMES
FIRST.
ENDCHR = OPTIONAL RIGHT JUSTIFIED BYTE ON WHICH TO TERMINATE INPUT.
OVERRIDES -1 TERMINATION WHEN NBYTE<0

SOUTB - STRING OUTPUT TO FILE, ARBITRARY BYTE SIZE
CALL SOUTB(JFN,POINTR,BYTSIZ,NBYTE)
POINTR = POINTER TO ARRAY FROM WHICH STRING IS OUTPUT
OTHER ARGUMENTS SAME AS IN SINB.

PSOUT - ASCII STRING(S) OUTPUT TO TTY
CALL PSOUT(POINTR1,POINTR2,...)
CALLS ASCZA IF HOLLERITH ARGUMENT

ASCZA - SEARCHES A 7-BIT STRING FOR A ZERO WORD. THEN
TRACES BACK LOOKING FOR A WORD WITH AN '&' AND THE REST
FILLED WITH BLANKS. IT WILL ONLY SKIP BACK OVER WORDS WHICH ARE
ALL SPACES.
IF FOUND, THE '&' AND THE ALL
BLANKS ARE REPLACED WITH NULLS. IS USEFUL FOR FORTRAN LITERALS.
IF (OPT) 2ND ARG IS GIVEN IT IS A LEFT JUSTIFIED TERMINATOR BYTE
TO BE USED INSTEAD OF '&'. IF IT IS SPACE, THEN
THIS MEANS TO DELETE ALL TRAILING SPACES BEFORE THE ZERO WORD.
CALL ASCZA(String,TERM)
STRING = A 7-BIT STRING WHICH MUST BE TERMINATED BY A ZERO WORD

USUALLY A HOLLERITH.

PSOUTR - ASCII STRING(S) OUTPUT TO TTY FOLLOWED BY CR-LF
CALL PSOUTR(POINTR1,POINTR2,...)
PUTS OUT ALL STRINGS, THEN A CR-LF. CALLS ASCZA IF HOLLERITH
ARGUMENT

RALPH - READ ALPHANUMERIC STRING FROM TTY.
ALLOWS CTRL-A OR RUBOUT EDITING.
ALSO ALLOWS CTRL-R VIEWING OF THE STRING
ALSO ALLOWS CTRL-U START OVER.
STRING IS TERMINATED BY CARRIAGE RETURN OR THE 400TH CHARACTER,
NEITHER OF WHICH IS PUT INTO THE ARRAY.

CALL RALPH(ASCI,NCHAR)
ASCI = ARRAY IN WHICH STRING IS STORED WITH A NULL TERMINATOR
NCHAR = NUMBER OF CHARACTERS IN THE STRING

LSH - LOGICAL SHIFT
JFOO=LSH(WORD,NPLACES)
WORD = WORD TO BE SHIFTED
NPLACES = NUMBER OF LEFT SHIFTS (NEGATIVE IF TO BE A RIGHT SHIFT)

EXTFLT - SIGN-EXTENDS, THEN FLOATS, ASSUMING SIZE <= 27 BITS
X=EXTFLT(IX,IEXWD)
IX = WORD TO BE SIGN-EXTENDED
IEXWD = 1 IN THE MOST SIGNIFICANT BIT OF THE BYTE
=LSH(1,BYTESIZE-1)

NRBYTS - FUNCTION TO COUNT BYTES IN A TERMINATED STRING
ICNT=NRBYTS(FROM,IDX,BYTSIZ,TERM)
FROM = STRING ADDRESS (I.E., AN ARRAY ELEMENT)
IDX = OPTIONAL STRING INDEX
IF ABSENT OR <= 0, DEFAULT VALUE OF 1 IS
USED.
BYTSIZ = OPTIONAL BYTE SIZE. IF ABSENT OR <= 0,
DEFAULT VALUE OF 7 IS USED.
TERM = OPTIONAL TERMINATOR BYTE. IF ABSENT,
DEFAULT VALUE OF 0 IS USED.

THE TERMINATOR BYTE IS NOT COUNTED.

CHMOVE - SUBROUTINE TO MOVE A CHARACTER STRING (NCHARS LONG)
CALLING SEQUENCE:
CALL CHMOVE(FROM,IDX1,TO,IDX2,NCHARS)
ARGUMENTS AS IN NRBYTS

ICHAR - FUNCTION WHICH RETURNS THE IDX-TH CHARACTER OF
THE STRING CONTAINED AT "FROM", LEFT JUSTIFIED AND PADDED WITH
SPACES

(SO IT CAN BE COMPARED WITH A FORTRAN SINGLE-CHARACTER LITERAL).
J=ICHAR(FROM,IDX)
WHERE FROM AND IDX ARE ARRAY PTR AND INDEX, AS IN NRBYTS.

WRWI - Subroutine to generate cosine table required for the
subroutine FFTR

CALL WRWI
CALCULATES 513 COSINES EQUALLY SPACED BETWEEN AND INCLUDING
0 AND 90 DEGREES

FFTR -- FFT of a Real Function

CALL FFTR(LOG2N,NSAMP,S,TR,TI)

COMPUTES THE (LOWER HALF + 1) OF THE FFT OF A REAL
FUNCTION

ARGUMENTS:

LOG2N = LOG2(N) WHERE N IS THE ORDER OF THE FFT
= MAXIMUM OF 10

NSAMP = NUMBER OF REAL SAMPLES TO BE TRANSFORMED

S = VECTOR OF LENGTH NSAMP, CONTAINS SAMPLES

TR = VECTOR OF LENGTH N/2+1, REAL PART OF TRANSFORM

TI = VECTOR OF LENGTH N/2+1, IMAG PART OF TRANSFORM

VECTORS S AND TR OR TI MAY BE IDENTICAL

ZETRAN - SETS THE RANDOM NUMBER "INITIAL VALUE" AND IS USED
TO SET THE ORIGIN OF THE RANDOM NUMBER SEQUENCE.

CALL ZETRAN(X,Y)

X=HIGH ORDER PART OF SEED

Y=LOW ORDER PART OF SEED

RANDOM - RANDOM REAL NUMBER GENERATOR

GENERATES A RANDOM REAL

NUMBER UNIFORMLY DISTRIBUTED BETWEEN TWO LIMITS.

X = RANDOM(A,B)

A=LOWER LIMIT

B=UPPER LIMIT

3. TYPICAL USER SESSION

The operation of the FORTRAN simulation requires only two inputs from the user: (1) a source of digitized input speech samples and (2) a location for the storage of the processed speech samples. At BBN, speech waveform samples are stored on disk files, as mentioned above. A typical user session, using disk file I/O, is described below. User input is underlined. In this session the input data file is <DCA16>BV1M.WAV and the output storage file is <DCA16>BV1M.TES. After inserting these two file names, the full coder simulation (transmitter and receiver) is executed without further intervention from the user. When all data has been processed, the program will print out the total number of frames processed and signal-to-quantization-noise (S/Q) ratios. The control of the program is then returned to the user. At that time the user may choose to process another speech utterance or abort the session.

Typical User Session

RUN 16KMN ÷

INPUT SPEECH FILE: <DCA16>BV1M.WAV ÷

10440 12-BIT SAMPLES AT 150 USEC = 1.5660 SECONDS

OUTPUT SPEECH FILE: <DCA16>BV1M.TES.1 ÷

FRAME COUNT = 48

S/Q RATIO in dB: LONG-TERM = 12.363; SEGMENTAL = 13.456

CONTINUE? (YES=-1,NO=0)=0 ÷

CPU TIME: 23.91 ELAPSED TIME: 1:35.95

NO EXECUTION ERRORS DETECTED

(Note that the symbol ÷ at the end of each input from the user denotes carriage return.)

4. SIMULATION OF SOME VARIATIONS OF THE OPTIMIZED APC CODER

The simulation has been designed to give the user the flexibility to modify the operation of the APC coder without software changes. Two methods of modification have been provided:

- (1) The user may set (True=-1) or clear (False=0) various flags that control the sequence of operations in the main program; and
- (2) The user may change the values of variables that specify important coder parameters. All flags and variables that the user may change are given their default values via DATA statements at the top of the main program 16KMN.

4.1 Flags

Flags have been provided so that the user may choose to keep or abort the execution of a specific section of the coder by setting or clearing the appropriate flag. For example, if the user wishes not to quantize the residual samples, he accomplishes this by simply clearing the flag IQ(1) (i.e., IQ(1)=0), prior to the execution of the coder. A list of the names of the flags and a description of the section of the coder each controls are given in Table 1. All flags, with the exception of ICHAN, ICHANE, and ICHANP, have their default value specified as True. The flags ICHAN, ICHANE and ICHANP are specified as False i.e., the coder is defaulted to operate in the absence of channel error.

Table 1. Flags in the FORTRAN simulation of the BBN 16 kb/s coder

<u>Flag</u>	<u>Description</u>
IPREM	Preemphasis - Deemphasis
KPPF	Pitch Prediction
INSTFX	3-Tap Pitch Predictor Stability Check
IHFCR	High Frequency Correction
KNSF	Noise Shaping
ISEGFG	Segmented Quantization
IFBC	Folded Binary Coding
ICHAN	Channel Simulation (Bit streaming)
ICHANE	Channel Error Simulation
ICHANP	Error Protection
IQ(1)-IQ(7)	Parameter Quantization
IQ(1)	Residual Samples
IQ(2)	Energy
IQ(3)	Delta Gains
IQ(4)	Spectral Coefficients
IQ(6)	Pitch
IQ(7)	Pitch Predictor Taps
NOPRNT	Listing of quantization tables

4.2 Variables

A list of the variables defining coder parameters that may be modified by the user, and their default values are given in Table 2. This table also specifies the limits of parameter values within which the user may reset them without any software changes. Changes in the parameter values should be made in a manner that preserves the consistency of interdependent parameters such as TFRAME and NENSEC. Note also that choosing NPOLE > 6 requires specification of additional data, as indicated in Table 2 under NPOLE.

Table 2. A list of APC coder simulation variables and their description

<u>Variable Name</u>	<u>Description</u>	<u>Default Value</u>
1. FREQS	Sampling frequency (Hz) Sampling frequency must correspond to that of the input data	6666.66
2. TFRAME	Interframe interval (ms) Set TFRAME such that NFRAME \leq 300 where NFRAME = IFIXR (TFRAME * FREQS/1000.) NOTE: Check NENSEC when modifying TFRAME as it also depends on NFRAME Function IFIXR is given in the subroutine package I6KGEN	32.4
3. BWF	Preemphasis bandwidth (Hz)	972.21465
4. T40	Pitch extraction frame size (ms) Set T40 such that I40 \leq 600 where I40 = IFIXR (T40*FREQS/1000.)	34.75
5. FOL	Lower limit of pitch frequency (Hz) This parameter is used to compute the upper limit on pitch period IF0L, defined in samples, where IF0L = IFIXR ((FREQS/FOL)+.5)	50
6. F0H	Upper limit of pitch frequency (Hz) This parameter is used to compute the lower limit on pitch period, IF0H, defined in samples where IF0H = IFIX(FREQS/F0H)	450
7. LOG2P	FFT order (exponent of 2) 2**LOG2P \geq I40+IF0L	9

Table 2
(cont.)

<u>Variable</u>	<u>Description</u>	<u>Default</u>
8. NTAPS	Number of pitch prediction taps 1 and 3 are permissible values, if no pitch prediction is required clear flag KPPF. Do not set NTAPS to zero.	3
9. NPOLE	Number of poles for LPC analysis Buffers accommodate up to 14 coefficients. However, default values for quantization and channel error are specified for up to NPOLE=6 only. To operate the simulation with NPOLE >6, the user must make the following modifications: <ol style="list-style-type: none"> 1. Provide additional values for the coding ranges CMIN and CMAX in subroutine QTBLC (16KCOD Package) 2. Provide additional default values for the number of bits protected and the number of bits transmitted (arrays NPERC and NBITC respectively at the top of the main routine, 16KMN) 	6
10. NPHFC	Order of the computation of the minimum mean-squared prediction error used in the high frequency correction module HFCOR (package 16KPR) $NPHFC \leq NPOLE$	2
11. HFLAM	Scalar constant used in the high frequency correction module	0.035
12. CHF	Autocorrelation coefficients of a high pass filter used in the high-frequency correction module 3 values required	+ .375, - .25, + .0625
13. BWANS	Noise shaping bandwidth (Hz)	800
14. NENSEC	Number of segments used in the segmented quantization scheme. Set NENSEC such that NSMSEC is an integer where $NSMSEC = NFRAME/NENSEC$	3
15. IRUNG	Switch to simulate rungs or stages of the phased real-time implementation. (IRUNG=0 for full simulation; 1, for stage 1; 2, for stage 2; and 3, for stage 3).	0

Table 2
(cont.)

<u>Variable</u>	<u>Description</u>	<u>Default</u>
16. NTYP	Switch to set the type of distribution used in the optimal quantization of the residual samples 1 = Gamma, 2 = Exponential, 3 = Gaussian	3
<p>The following parameters set the number of bits for quantization of the transmitted parameters. If quantization is not required for a specific parameter, clear the appropriate IQ flag -- do not set the parameter below to zero.</p>		
17. NBITR	Number of bits for quantization of the residual samples Range - 1 to 3	2
18. NBITG	Number of bits for quantization of the Energy Range - 1 to 7	6
19. NBITSC	Number of bits for quantization of the delta gains Range - 1 to 3	2
20. NBITC(I)	Number of bits for quantization of the spectral coefficient I Range - 1 to 7	6 5 4 4 4
21. NBITP	Number of bits for quantization of the pitch Range - $2^{**}NBITP \geq IF0L - IF0H$	7
22. NBITT(I)	Number of bits for quantization of the pitch predictor taps Range = 1-5	3 4 3

5. CODER OPERATION WITH CHANNEL-ERROR SIMULATION

The FORTRAN simulation is defaulted to operate in the absence of channel error. To initiate the simulation of channel error on all transmitted parameters the user must set (true=-1) the two flags ICHAN and ICHANE. Also, the flag ICHANP must be set to invoke error protection of parameters. The coder is defaulted to operate at 1% channel error when these flags are set. The user may change the percentage of channel error and the amount of protection for each parameter independently by resetting their default values at the top of the main program. Table 3 gives the names of the variables that specify the number of high-order bits protected for each transmitted parameter. When changing the values of these variables, the user must keep in mind that the total number of protected bits for all transmitted parameters should be an integer multiple of 4. The percentage of channel error for each transmitted parameter is defined in the array ERP. In Table 4, the correspondence between the array entries and the transmitted parameters is defined.

Table 3. A list of transmitted parameters, along with the number of high-order bits of each that are error-protected

Variable	Transmitted Parameter	Default Value
NPERC (I)	Spectral Coefficients $NPERC(I) \leq NBITC(I)$	5,4,3,2,2,2
NPERP	Pitch $NPERP \leq NBITP$	7
NPERG	Energy $NPERP \leq NBITG$	6
NPERSC	Delta gains $NPERSC \leq NBITSC$	2
NPERT (I)	Pitch Predictor Taps $NPERT(I) \leq NBITT(I)$	2,3,2

(No protection is provided for the residual samples.)

Table 4. Correspondence between the elements of the array ERP and the type of transmitted parameters. The value of the array element indicates the percentage of channel error due to which the corresponding transmitted parameter(s) are exposed.

<u>Variable</u>	<u>Transmitted Parameters</u>	<u>Default Value</u>
ERP (1)	Spectral Coefficients	0.01
ERP (2)	Pitch	0.01
ERP (3)	Energy	0.01
ERP (4)	Delta Gains	0.01
ERP (5)	Pitch Predictor Taps	0.01
ERP (6)	Residual Samples	0.01

APPENDIX C

A LISTING OF THE SOURCE PROGRAMS OF THE
FORTRAN SIMULATION OF THE 16 KB/S APC CODER


```

C <UCAL0>166MM.V4.2  Prt 21-Nov-88 9:36AM
C WARNING WORD MESSAGES
  DATA IHAMHD(1)/4/,IHAMHD(2)/16/,IHAMHD(3)/32/
  DATA IHAMHD(4)/64/,IHAMHD(5)/1/,IHAMHD(6)/72/
  DATA IHAMHD(7)/8/,IHAMHD(8)/128/
C AMOUNT OF PROTECTION FOR EACH PARAMETER
C DEFINED IN VARIABLES MPERC - SEE USER'S
C GUIDE FOR CORRESPONDENCE BETWEEN VARIABLE
C AND PARAMETERS
  DATA MPERC(1)/5/,MPERC(2)/4/,MPERC(3)/3/,MPERC(4)/2/
  DATA MPERC(5)/2/,MPERC(6)/2/
  DATA MPERC(6)/MPERC(7),MPERC(8)/
  DATA MPERT(1)/2/,MPERT(2)/3/,MPERT(3)/2/

```

```

C <UCAL0>166MM.V4.2  Prt 21-Nov-88 9:36AM
C CURFF
  DIMENSION MBITC(14),MBLVC(14),XCUP(129,14),KCOF(128,14)
  DIMENSION CSTEP(14)
  DIMENSION ICG(14),ICM(14)
  DATA MBITC(1)/6/,MBITC(2)/5/,MBITC(3)/4/
  DATA MBITC(4)/4/,MBITC(5)/4/,MBITC(6)/4/
  DATA ICG(1)/-1/
C OVERALL GAIN
  DIMENSION KG(129),KG(128)
  DATA MBITC(6)/,IG(2)/-1/
C DELTA GAINS
  DIMENSION XSEC(9),YSEC(8)
  DIMENSION IGMS(3),IGMS(J)
  DATA MBITC(2)/,IG(3)/-1/
C RESIDUAL
  DATA MBITR(2),MBTR(3),IQ(1)/-1/,IPRC/-1/
C COMMUN/RESCUD/MBLVC,MBES(9),YRES(8)
C CHANNEL ERROR SIMULATION AND PROTECTION
  -----
  INTEGER SHANA,SRANY
C
  DIMENSION ERP(6),ISEC(50),MPERC(50),MBERC(50)
  DIMENSION IDAT(50),ER(50),IDAYS(50),IDS(50)
  DIMENSION ERMS(50),IDATE(50),IDATR(50)
  DIMENSION IHW(50),IHW(50),IHW(50),IHW(50)
  DIMENSION ID(50),IDRC(50),HOSPA(600)
  DIMENSION MPERC(14),MPERT(3)
C
  COMMON/INASK/INASK(32)
  COMMON/IHAMHD/IHAMHD(8)
  COMMON/CNN/ICNNBT,IECCUM,ISERR
C
  DATA MBTR/3/,MAREP/500/
  DATA MBWT/50/,RMDNE/1/
  DATA SHANA/28835966/,SRANY/21743579315/
  DATA ICHANE/8/,ICAN/8/,ICAMP/8/
  DATA ICHANE/8/,ICAN/8/,ICAMP/8/
C PERCENTAGE OF CHANNEL ERROR FOR EACH
C TRANSMITTED PARAMETER TYPE DEFINED IN
C ARRAY ERP - SEE USER'S GUIDE FOR CORRESPONDENCE
C BETWEEN ARRAY ENTRY AND PARAMETER TYPE
  DATA ER(1)/8.01/,ERP(2)/8.01/,ERP(3)/8.01/
  DATA ER(4)/8.01/,ERP(5)/8.01/,ERP(6)/8.01/

```



```

C SET INITIAL DATA FOR CIDER OPERATION
C
C PREEMPHASIS
  ALPHA=1.0-ALPHA
  PEN=DP*TMU01/PRES
  PENFAC=EXP(PEN)
  SNT FRAME LOCATION AND ANALYSIS WINDOW
  (ANALYSIS WINDOW NOT USED IN THIS SIMULATION)
  N100=NFRAME
  N25=N25N
  N25=N25*NFRAME-1
  PITCH EXTRACTION
  F40=F40*PRES/1000.
  I40=IPIR(F40)
  I(100.LI.NFRAME) I40=NFRAME
  I40N=I40-NFRAME
  N25P=N25-140N
  I40=IPIR((PRES/POL)*.5)
  I40N=I40-NFRAME
  SECURED QUANTIZATION
  RN5SEC=NFRAME/NESEC
  NOISE SHAPING
  PZPAC=PI*WANS*TSAMP
  PZPAC=EXP(PZPAC)
  PITCH QUANTIZATION
  NLEVP=2**NBITP
  NLEVP=2**NBITP
  GENERATE QUANTIZATION TABLES FOR
  PITCH TAPS
  CALL QTLT(UTAPS,NBITT,NLEVT,RYAP,RYAP,TSTEP)
  REFLECTION CORPUS
  CALL WTMC(WPUL,NBITC,NLEVC,RC,FCUP,CSTEP)
  OVERALL ENERGY
  CALL QTLG(QBETC,NLEVC,RC,FC,CSTEP)
  DELTA GAINS
  CALL WTMS(WMITSC,NLEVSC,ISAC,YSEC)
  RESIDUAL SAMPLES
  NLEVP=2**NBITN
  IPULO=NLEVP/2-1
  CALL WPTC(WTTP,NLEVP,NLGS,YRNS)
C PRINT QUANTIZATION DATA
  I(INDPMT) GO TO 299
  TYPE 297
297 FORMAT(10,10R,34HTYPE QUANTIZATION TABLES - STARTED)
298 CONTINUE
  DO 202 I=1,NTAPS
  PRINT 201
  FORMAT(1M1)
  PRINT 200,1
  FORMAT(1M,10R,6HTAP 0,12,/)
  PRINT 201

```

```

202 FORMAT(1M,12R,4HX(J),12R,1HJ,10R,4HR(J),/)
  PRINT 204,RYAP(1,1)
  FORMAT(1M,10R,P10.0)
  DO 206 J=1,NLEVT(1)
  JJ=J-1
  PRINT 200,JJ,RYAP(J,1)
  FORMAT(1M,27R,13,7R,P10.0)
  PRINT 204,RYAP(J+1,1)
  CONTINUE
  C CUEFF
  DO 224 I=1,NPOLE
  PRINT 201
  PRINT 204,1
  FORMAT(1M,27R,8MC0EFF 0,12,/)
  PRINT 202
  PRINT 204,ICOP(1,1)
  DO 211 J=1,NLEVC(1)
  JJ=J-1
  PRINT 200,JJ,VCOP(J,1)
  PRINT 204,ICOP(J+1,1)
  CONTINUE
  C GAIN
  PRINT 201
  PRINT 206
  FORMAT(1M,30R,4MCAIM,/)
  PRINT 202
  PRINT 222,IC(1)
  FORMAT(1M,10R,P14.0)
  DO 224 J=1,NLEVC
  JJ=J-1
  PRINT 226,JJ,VC(J)
  PRINT 222,IC(J+1)
  FORMAT(1M,31R,13,7R,P14.0)
  CONTINUE
  C DELTA GAINS
  PRINT 201
  PRINT 208
  FORMAT(1M,26R,11DELTA GAINS,/)
  PRINT 202
  PRINT 222,ASEC(1)
  DO 234 J=1,NLEVSC
  JJ=J-1
  PRINT 226,JJ,VSEC(J)
  PRINT 222,ASEC(J+1)
  CONTINUE
  C RESIDUAL
  PRINT 201
  PRINT 209

```

```

C PITCH TAPS
L=30
IF (NTAPS.GT.1) GO TO 322
L=L+1
K=K+1
ISEQ(K)=L
MERR(K)=MERRT(I)
MPER(K)=MPERT(I)
GO TO 330
322 DO 325 I=1,NTAPS
K=K+1
L=L+1
ISEQ(K)=L
MERR(K)=MERRT(I)
MPER(K)=MPERT(I)
CONTINUE
325 CONTINUE
330 MDATA=M
C IF NO ERROR PROTECTION SET MPER=0
IF (LCMWP) GO TO 335
DO 332 I=1,MDATP
MPER(I)=0
CONTINUE
332 CONTINUE
335 CONTINUE
C SET START POSITION IN SEQUENCE FOR RESIDUAL SAMPLES
MSTRES=0
C COMPUTE TOTAL BITS FOR RESIDUAL SAMPLES
MRTT=MRTT+MFRAME
MDATP=MDATP+MRTT
L=L+1
DO 340 I=MDATP+1,MDATP
L=L+1
ISEQ(I)=L
MPER(I)=0
MERR(I)=MNONE
CONTINUE
340 CONTINUE
C SET UP BIT ARRAY FOR RESIDUAL SAMPLES (BIT STREAMING)
DO 343 I=M23,M28
MBSPA(I)=MRTT
CONTINUE
343 CONTINUE
C INITIALIZE CUMULATIVE ERROR COUNTS
ITCUMB=0
ITCUNE=0
MTEERR=0
CONTINUE
350 CONTINUE

```

```

289 FORMAT(1H ,20L,MRESIDUAL,/)
PRINT 202
PRINT 222,MRES(I)
DO 244 J=1,MLEVR
JJ=J-1
PRINT 226,JJ,MRES(J)
PRINT 227,MRES(J+1)
CONTINUE
TYPE 296
FORMAT(1H ,10L,9MTYPE DUNE)
CONTINUE
299
C
C
C SET UP PARAMETERS REQUIRED FOR CHANNEL ERROR SIMULATION
C TOTAL NUMBER OF BITS PER TRANSMITTED PARAMETER STORED IN MLEN
C TOTAL NUMBER OF PROTECTED BITS FOR EACH TRANSMITTED PARAMETER
C STORED IN MPER
C SEQUENCE OF TRANSMITTED PARAMETERS STORED IN ISEQ
C (SEE SAVER FOR CORRESPONDENCE BETWEEN PARAMETER
C NUMBER AND PARAMETER TYPE)
C
C IF (.NOT. LCMWP) GO TO 350
C NUMBER OF SEQUENCE DEFINED BY K
K=0
C SPECTRAL CURVES
DO 315 I=1,MPOLE
K=K+1
MERR(K)=MERRT(I)
ISEQ(K)=K
MPER(K)=MPERT(I)
CONTINUE
315 CONTINUE
C OVERALL ENERGY
K=K+1
ISEQ(K)=16
MERR(K)=MRTT
MPER(K)=MPERC
C PITCH
K=K+1
ISEQ(K)=15
MERR(K)=MRTT
MPER(K)=MPEPP
C DELTA CHIRP
L=L+1
DO 320 I=1,MEWSEC
L=L+1
K=K+1
ISEQ(K)=L
MERR(K)=MRTT
MPER(K)=MPEPC
CONTINUE
320 CONTINUE

```



```

C
2315 CONTINUE
   ENTE-ENERGY(DM,NAER,N25,N2E)
   IF(ENTE-EO-0) ENTE-ENERO
   ENTE-SUMT(ENTE)
C QUANTIZE GAIN
   IF(.NOT.IG(2)) GO TO 2396
   CALL CUB(ENYA,IG,MLTVC,IRNY)
   CALL DCOM(ENQA,IG,ENRE)
2316 CONTINUE
C SECURED QUANTIZATION
   IF(.NOT.ISEC7G) GO TO 2320
   COMPUTE DELTA GAINS - ESTE
   ENQ3-AMR*ENRE
   CALL DELTAC(EN,NAER,ENQ3,N25,ESTE)
   IF(.NOT.IG(3)) GO TO 2322
C QUANTIZE DELTA GAINS
   DO 2321 I=1,NERSEC
   CALL CUB(ENYA(I),ISEC,MLTVC,IRNY(I))
   CALL DCOM(ENRS(I),YSEC,ENRE(I))
2321 CONTINUE
   GO TO 2327
C NO QUANTIZATION OF DELTA GAINS
2322 DO 2326 I=1,NERSEC
   ASRE(I)=SUMT(ESTE(I))
2326 CONTINUE
2327 CONTINUE
C COMPUTE QUANTIZER SCALE FACTOR - ESSEC
   GO TO 2340
   CALL ASCALA(LSAC,ENRE,N25,ESRE)
C NO SECURTED QUANTIZATION
2328 DO 2330 I=N25,N2E
2330 LSAC(I)=ENRE
2340 CONTINUE
C
C . . . . . APC LOOP . . . . .
C
   CALL APC(S,CSEC,ANS,DA,DR,MPYCHR,N25,CO,IR,VR,RH,Q,Q1
   I,BTAS,IQ(I),EPFF,EMSF)
C
C
C PULSED BINARY CUBING
   IF(.NOT.IFBC) GO TO 2289
   CALL PULBC(IR,VR,RH,IFOLD)
2289 CONTINUE
C

```

```

C CHANNEL CUBING
C
2400 IF(.NOT.ICMAB) GO TO 2399
C . . . ERROR PROTECTION . . .
C ZERO DATA ARRAY
   DO 2405 I=1,NDATT
2405 IDAT(I)=0
C CONVERT QUANTIZATION LEVELS OF RESIDUAL SAMPLES
C IN THE ORDER GIVEN BY ISEQ. PERCENT OF CHANNEL ERROR FOR
C EACH PARAMETER,ERP, ALSO LOADED (IN SAME
C SEQUENCE AS PARAMETERS) IN ARRAY EP.
C CALL SEQZ(NDATT,ISEQ,ICQ,IPTCH,IEHQ,IEHS,IPTAP,IE,IDAT,ERP,ERP)
C EXTRACT BITS TO BE PROTECTED FROM THE ARRAY IDAT
C AND FORM 4-BIT BYTES. 4-BIT BYTES ARE PUT IN IHW. BITS NOT
C PROTECTED ARE PUT IN IDATS.
C CALL EITHM(NDATP,IDAT,MPER,NBER,IMW,IMWT,IDATS,IRS,ENH,ER)
C FILL IDATS WITH ALL OTHER PARAMETERS
C (I.E. ALL PARAMETERS IN WHICH NO BITS WERE PROTECTED)
   DO 2407 I=NDATP+1,NDATT
2407 IDATS(I)=IDAT(I)
C HANNING CURVE 4-BIT BYTES (IHW)
C INTO 7,4 HANNING CODEWORDS
   DO 2410 I=1,IMW
   CALL CODR(IHW(I),HOPT,INWC(I))
2410 CONTINUE
C
C CHANNEL
C
   IF(.NOT.ICMAB) GO TO 2435
   ITCUMB=0
   IERCUM=0
   MSLRN=0
C ERROR HANNING CODEWORDS
   DO 2420 I=1,IMW
2420 CALL EARCH(INWC(I),HOPT+4,ERH(I),INWE(I),0)
C ERROR UNPROTECTED DATA BITS
   DO 2430 I=1,NDATT
   KRESFG=0
   IF(I.CE-NSTRG) KRESFG=-1
2430 CALL EARCH(IDATS(I),IRS(I),ER(I),IDATE(I),AMESFL)
C UPDATE COUNT OF TOTAL BITS
   ITCUMB=ITCUMB+ITCUMBT
C UPDATE COUNT OF ERRORED BITS
   ITCUMB=ITCUMB+IERCUM

```

```

C * * * * RECEIVER * * * *
3080 CONTINUE
C * * * *
C PUSH DOWN ARRAYS
CALL PSHN(ARR,MAK,MPFRAME)
CALL PSHN(BRR,MAK,MPFRAME)
CALL PSHN(KMS,MAK,MPFRAME)
CALL PSHN(ASP,MAK,MPFRAME)
CALL PSHN(SLAST,MAK,MPFRAME)
C
C
C RUNC BRANCHING
IF(IHUNG.NE.1) GO TO 3084
DO 3082 I=N25,N28
SLAST(I)=S(I)
CONTINUE
GO TO 4088
3082
C
3084 IF(IHUNG.NE.2) GO TO 3088
DO 3086 I=N25,N28
SP(I)=R(I)
CONTINUE
GO TO 3518
3086
C
3088 IF(IHUNG.NE.3) GO TO 3012
DO 3018 I=N25,N28
MM(I)=M(I)
CONTINUE
GO TO 3418
3018
C
3012 CONTINUE
C * * * * * DECODE ENERGY * * * * *
C NO QUANTIZATION
EMR=SAINT(ENTR)
GO TO 3112
C QUANTIZED ENERGY
CALL DCW(EMR,FC,EMR)
3110
3112 IF(ISECP) GO TO 3116
C NO SEGMENTED QUANTIZATION
DO 3114 I=N25,N28
3114 656(I)=EMR
GO TO 3135
C SEGMENTED QUANTIZATION
3116 IF(.NOT.IQ(3)) GO TO 3125
L QUANTIZED DELTA GAINS
DO 3116 I=L,MEUSEC
CALL DCOD(IEMR(I),YSEC,ESR(I))
3118 CONTINUE
GO TO 3138

```

```

C UPDATE COUNT OF ENCODED RESIDUAL SAMPLE BITS
NTSEB=NTSEB+NSER
GO TO 2339
C NO CHANNEL ERROR
2435 CONTINUE
DO 2436 I=1,MMU
MM(I)=INC(I)
2436
DO 2437 I=1,MBATT
IBATE(I)=IBATS(I)
2437
C
C CHANNEL DECODING
C * * * * * ERROR CORRECTION * * * * *
C DECODE NAMING COMMANDS
DO 2440 I=1,MMU
CALL DCOD(IME(I),MUPT,IMR(I),MOPFC)
2440 CONTINUE
C COMMING PROTECTED (IMR) AND UNPROTECTED (IDATE) DATA BITS
CALL CUMR(IMR,MPE,MGR,MBATT,IDATE,IBATR)
C EXTRACT PARAMETERS FROM IBATR
CALL SEQR(IBATT,ISEQ,ICUR,IPTCR,IEMR,IENSR,IPTAP,IBR,IBATR)
C
C EXTRACT RESIDUAL SAMPLES FROM BIT STREAM
2486 CALL MSTMU(EN,MBUT,MBSPA,N25,IM,MPHANG,MSRCY)
C
C NO CHANNEL SIMULATION
2588 CONTINUE
DO 2518 I=1,MPULE
ICUR(I)=ICQ(I)
2518
DO 2520 I=1,MPFRAME
IM(I)=IN(I)
2520
DO 2530 I=1,MLNSEC
IENSR(I)=IENS(I)
2530
DO 2540 I=1,MTAPS
IPTAP(I)=IPTAP(I)
IEMR=IEMR
IPTCR=IPTCR
C
C

```



```

SUBROUTINE USAMP(JPRO,S,ISCR,NST,NEW)
C SUBROUTINE TO OUTPUT SAMPLES TO FILE
C INPUTS:
C JPRO = JOB OF OUTPUT FILE
C S = VECTOR OF SAMPLES (REAL)
C NST = NUMBER OF SAMPLES IN 3 FROM NST TO NEW ARE TRANSFERRED TO FILE
C ISCR = SCRATCH VECTOR (SAME SIZE AS S)
C DIMENSION S(1),ISCR(1)
COMMON/NAKAMP/NAKAMP,NOVER
C NOVER IS THE NUMBER OF TIMES SAMPLE AMPLITUDES
C EXCEEDS AND HENCE SATURATES AT NAKAMP.
C
K=0
DO 10 I=NST,NEW
K=K+1
C FIX AND ROUND THE SAMPLES
ISCR(K)=IFIX(S(I))
C SATURATE SAMPLE AMPLITUDE AT NAKAMP
IF (ISCR(K).GE.NAKAMP) GO TO 20
GO TO 30
ISCR(K)=NAKAMP-1
NOVER=NOVER+1
30 IF (ISCR(K).LT.-NAKAMP) GO TO 40
GO TO 10
ISCR(K)=-NAKAMP
NOVER=NOVER+1
10 CONTINUE
C TRANSFER K SAMPLES IN ISCR
C TO OUTPUT FILE
CALL SOUTB(JPRO,ISCR,36,K)
RETURN
END

```

```

SUBROUTINE ISAMP(IPLG,JPRO,S,ISCR,NST,NEW)
C SUBROUTINE TO INPUT SAMPLES FROM FILE
C INPUTS:
C IPLG = JOB FILE NUMBER OF INPUT FILE
C NST = VECTOR OF SAMPLES (REAL)
C NEW = SAMPLES STORED IN S AS 36-BIT FLOATING-POINT NUMBERS
C FROM NST TO NEW
C ISCR = SCRATCH VECTOR (SAME SIZE AS S)
C IPLG = IPLG IS SET (IPLG=1) IF END OF FILE
C IS ENCOUNTERED WHEN READING IN SAMPLES
COMMON/NAVFL/NAVFL,IBVSIZ,LEXND,INBYT1
COMMON/FUNSIZ/FPFRAC,NUIND
COMMON/INSEPT/INSEPT
C DIMENSION S(1),ISCR(1)
IPLG=C
ASSIGN 100 TO IEND
SET FILE POINTER TO CORRESPOND TO THE
BEGINNING OF THE DESIRED FRAME.
CALL SPTX(IPLG,INBYT1,INSEPT)
READ BUIND NUMBER OF INTEGER SAMPLES OF SIZE
IBVSIZ BITS EACH, INTO THE VECTOR ISCR(SAMPLES ARE RIGHT JUSTIFIED)
CALL SIM(IEND,FPFRAC,ISCR(NST),36,NUIND)
PUT EACH SAMPLE, EXTEND SIGN BIT INTO
(36-IBVSIZ) MOST SIGNIFICANT BITS, AND CONVERT
TO FLOATING-POINT NUMBER.
DO 10 I=NST,NEW
S(I)=REALT(IISCR(I),ICRND)
CONTINUE
C ADVANCE THE COUNTER TO CORRESPOND
C TO THE BEGINNING OF THE NEXT FRAME.
INSEPT=INSEPT+FPFRAC
RETURN
C SET END OF FILE FLAG
IPLG=1
RETURN
END

```

```

C READ IN 1ST 38 WORDS FROM FILE AND FIND OUT IF HEADER EXISTS
326 ASSICH 321 TO IERF
CALL SIMH(IECF,JFMI,INDR,36,30)
GO TO 322
C IERF CALL
321 CALL PSUMTR('UNEXPECTED IEOP ON HEADER READ - RESTART')
IERF--1
RETURN
322 IF(INDR(1).EQ.HORTES(1))
1.AND.INDR(2).EQ.HORTES(2)
2.AND.INDR(3).EQ.HORTES(3)
4.AND.INDR(6).EQ.HORTES(6)
5.AND.INDR(7).EQ.HORTES(7)) GO TO 332
C FILE HAS NO MESSAGES
C ASK FOR FILE BYTE SIZE AND SAMPLING RATE
YESMOR=0.
CALL PSUMTR(' FILE DOES NOT HAVE HEADER')
CALL PSUMTR(' FILE SAMPLING RATE (M2): L')
ACCEPT 101,FREQS
FORMAT(F)
TSAMP=1./FREQS
CALL PSUMTR(' FILE BYTE SIZE (IN BITS): L')
ACCEPT 102,IMBSIZ
FORMAT(I)
IERM=LSH(1,IMBSIZ-1)
IHWI=0
MARAMP=2**((IMBSIZ-1)
MOVER=9
C FIND FILE LENGTH IN TERMS OF BYTE SIZE JUST GIVEN
CALL SFBSZ(JFMI,IMBSIZ)
CALL SFPR(JFMI,-1)
CALL SFPR(JFMI,MFLPTS)
CALL SFPR(JFMI,0)
DO 338 I=1,38
338 IHDR(I)=HORTES(I)
SMP=((./FREQS)*1.0E+6
IHDR(11)=IFIX(SMP/10.)
IF(SMP.NE.TMP) IHDR(11)=-FIX(SMP)
IERPT=0
GO TO 335
C FILE HAS HEADER - USE HEADER DATA

```

```

SUBROUTINE OPNIF(IECR,JFMI)
C SUBROUTINE TO OPEN INPUT FILE
C OUTPUT: JFMI = JFM OF INPUT FILE
C JFMI IS 0 IF FILE
C WAS SUCCESSFULLY OPENED.
C IF UNSUCCESSFUL,
C JFMI=-1
C ICRM = FLAG TO INDICATE THUMBLE
C WITH FILE OPENING
C
COMMON/MBFIL/IMBSIZ,IECR,IBVTTI
COMMON/NAME/TPNAM(10),IMAME(10),UNAME(10)
COMMON/INDR/INDR(30),TESMOR
COMMON/FREQS/FREQS,TSAMP
COMMON/MARAMP/MARAMP,MOVER
COMMON/IERPT/IERPT
COMMON/MFLPTS/MFLPTS
INTEGER HORTES(30)
DATA (HORTES(I),I=1,30)/36,10,1,0,1,1,23,400000000000/
C
308 CALL PSUMTR
CALL PSUMTR(' INPUT SPEECH FILE: ')
CALL RALPH(IECR,ICR)
IF(ICR.EQ.0) GO TO 302
GO TO 310
302 CALL PSUMTR('? NAME NOT GIVEN - RESTART')
IERF--1
RETURN
C FILE NAME NOT EMPTY - CHECK IF EXTENSION IS SPECIFIED
310 NOTM=HORTES(IECR)
DO 312 I=1,NOTM
IF(IECR(I).EQ.'.') GO TO 314
312 CONTINUE
C NO EXTENSION GIVEN
CALL PSUMTR('NO EXTENSION GIVEN - RESTART')
IERF--1
RETURN
C EXTENSION GIVEN - STORE <IECR>NAME IN TPNAM
314 CALL CHMOVE(IECR,1,TPNAM,1,1)
C IF FILE IS OPEN (JFMI IS NOT NEGATIVE)
C CLOSE THE FILE WITH JFM=JFMI
IF(JFMI.GT.0) CALL CLOS(JFMI)
JFMI--1
C OPEN FILE
CALL OPEN(JFMI,36,NAME,IER)
IF(IECR.EQ.0) GO TO 320
CALL PSUMTR('THUMBLE WITH UPNIF - RESTART')
JFMI--1
IERF--1
RETURN

```

C <DC16>16K10.F4.2 Tue 8-Jul-88 11:58AM

Page 312

C <DC16>16K10.F1.2 Tue 8-Jul-88 11:58AM

```

332 RESMUR=1.
   CALL SFMSZ(JFNU,IMDR(4))
   IAMSIZ=IMDR(4)
   MOUTMD=36/IMBSIZ
   IEMD=ASH(1,IMBSIZ-1)
   IMBVT1=IMDR(1)*MOUTMD
   MAREAMP=1/IMBSIZ-1
   MOVER=0
   MFLPTS=IMDR(5)
   IF (IMDR(11).LT.0) GO TO 333
   FREQS=(1./FLUAT(IMDR(11)))*1E+6
   GO TO 334
333 FREQS=1./FREQS/(FLUAT(1685*(IMDR(11))))
334 FSAWP=1./FREQS
   IREPT=0
C COMPUTE UTTERANCE LENGTH IN SECONDS
C AND TYPE IT OUT AT THE USER'S TERMINAL
335 I1=MFLPTS*100./FREQS*0.5
   JJ=FSAMP*1E+6
   UUR=1/1000.
   TYPE J37,MFLPTS,IMBSIZ,JJ,UUR
   FORMAT(' ',17,13,'-BIT SAMPLES AT ',14,' USEC = ',
17,4,' SECONDS',/)
C
   RETURN
   END
C SUBROUTINE UPHOP(IERR,JFNU)
C SUBROUTINE TO OPEN OUTPUT SPEECH FILE
C OUTPUT
   JFNU = JFNU IN OUTPUT FILE
   (AS IN UPHIF)
   IERR = FLAG TO INDICATE TROUBLE
   WITH FILE OPENING
C
   COMMON/MAYTTL/IMBSIZ,IEMD,IMBVT1
   COMMON/IMDR/IMDR(36),FESHR
   COMMON/NAME/TPNAM(10),IMAME(10),OMAME(10)
C
   IERR=0
   CALL PSOUT('OUTPUT SPEECH FILE:')
   CALL UPENDF(JFNU,36,9,IERR)
   IF(IERR.EQ.0) GO TO 20
   CALL PSOUT('TROUBLE WITH OPENUP - RESTART')
   IERR=-1
   RETURN
C STORE NAME
20 DO 22 1=1,10
21 OMAME(1)=0.
22 CALL FILNAM(JFNU,OMAME)
   CALL SFMSZ(JFNU,36)
C OUTPUT HEADER AND SET BYTE SIZE
   CALL SOUTS(JFNU,IMDR,36,30)
   RETURN
   END

```

```

FUNCTION IPIXR(A)
C FIX AND ROUND A FLOATING
C POINT NUMBER
C
AB=ABS(A)
IF(AB.NE.0.) GO TO 10
IPIXR=0
RETURN
10 ASIGN=A/AN
TEMP=A*ASIGN*(0.5)
IPIXR=IPIXR(TEMP)
RETURN
END

```

```

SUBROUTINE PSHDR(A,ND,N)
C ROUTINE TO PUSH DOWN (PSHDR) DATA IN AN
C ARRAY (A) OF TOTAL LENGTH NU BY AN AMOUNT N
C THAT IS , THE CONTENTS OF A(N+1) THROUGH
C A(ND) ARE COPIED INTO A(1) THROUGH A(ND-N))
C THE CONTENTS OF A(ND-N+1) THROUGH A(ND) ARE
C UNCHANGED.
C
DIMENSION A(ND)
N1=N+1
K=0
DO 10 I=N1,ND
K=K+1
A(K)=A(I)
RETURN
END

```

```

SUBROUTINE SETZK(X,MARK)
C ROUTINE TO ZERO OUT AN ARRAY (X)
C UP LENGTH MARK
C
DIMENSION X(1)
DO 10 I=1,MARK
X(I)=0.0
CONTINUE
RETURN
END

```

```

FUNCTION ALARC(C)
C REFLECTION COEFF TO LOG AREA RATIO
C CONVERSION FUNCTION
C
CRR=(1.-C)/(1.+C)
ALAR=DB(CRR)
RETURN
END

```

```

FUNCTION REFL(A)
C LOG AREA RATIO TO REFLECTION COEFF
C CONVERSION FUNCTION
C
VV=(10.-0)**(A/10.-0)
REFL=(VV-1.)/(1.+VV)
RETURN
END

```

```

FUNCTION DB(X)
C DB FUNCTION: 10*LOG10(X)
C
XX=X
IF(XX.LT.1.-10) XX=1.-0E-10
DB=10.*ALOG10(XX)
RETURN
END

```

```

FUNCTION RUB(A)
C "REVERSE" DB FUNCTION: 10**(A/10)
C
W=A/10.-0
PDB=(10.-0)**W
RETURN
END

```


SUBROUTINE QTRLC(INPULSE,ABITC,MLEV,RCOF,ICOF,CSTEP)

C ROUTINE TO SET CODING TABLES FOR REFLECTION COEFFICIENTS.

C INPUT:

C NPULSE = NUMBER OF COEFFICIENTS

C ABITC = VECTOR OF NUMBER OF BITS PER COEFFICIENT

C MLEV: MLEV = VECTOR OF NUMBER OF LEVELS PER COEFFICIENT

C RCOF: RCOF = TRANSMITTER CODING TABLE

C ICOF: ICOF = RECEIVER DECODING TABLE

C CSTEP: CSTEP = VECTOR OF STEP SIZES

C

C DIMENSION MLEV(1),MLEV(1),RCOF(129,14),ICOF(129,14)

C DIMENSION CSTEP(1),CMIN(14),CMAX(14)

C

C THE FOLLOWING DATA STATEMENTS CONTAIN LAR CODING RANGES IN DBI

C DATA CMIN(1)/-21.849249/,CMIN(2)/-8.7887499/,CMIN(3)/-9.831125/

C DATA CMIN(4)/-6.893175/,CMIN(5)/-5.28125/,CMIN(6)/-3.7485/

C DATA CMAX(1)/11.853149/,CMAX(2)/13.718449/,CMAX(3)/7.46875/

C DATA CMAX(4)/8.96825/,CMAX(5)/7.71875/,CMAX(6)/9.5599499/

C

C DO 110 I=1,MPULSE

C MLEV(I)=2**ABITC(I)

C CONTINUE

C

C DO 111 J=1,MPULSE

C CALL QTRLC(CMAX(I),CMIN(I),MLEV(I),RCOF(I),ICOF(I),CSTEP(I))

C CONTINUE

C

C CONTINUE

C

C CONVERT TABLE VALUES FOR DIRECT CODING

C OF REFLN COEFF

C DO 112 I=1,MPULSE

C MLEV(I)

C ICUF(I,1)=REFLN(ICUF(I,1))

C DO 112 J=1,M

C ICUF(J,1,1)=REFLN(ICUF(J,1,1))

C ICUF(J,1)=REFLN(ICUF(J,1))

C CONTINUE

C RETURN

C END

112

7L

SUBROUTINE QTRLC(MBITC,MLEV,RC,VC,CSTEP)

C ROUTINE TO SET UP CODING TABLES FOR OVERALL ENERGY

C INPUTS:

C MBITC NUMBER OF BITS FOR CODING ENERGY

C MLEV NUMBER OF QUANTIZATION LEVELS

C RC TRANSMITTER CODING TABLE

C VC RECEIVER DECODING TABLE

C CSTEP STEP SIZE

C

C DIMENSION IC(1),IC(1)

C

C CMAX AND CMIN DEFINE GAIN CODING RANGE IN DBI

C NOTE: THE VALUES CMAX AND CMIN GIVEN BELOW HAVE

C BEEN EMPIRICALLY OBTAINED FOR INPUT SPEECH

C SAMPLES THAT HAVE BEEN NORMALIZED(OTHER INDIVIDUAL

C UTTERANCES) TO BE EACH 11 BITS (INCLUDING ONE SIGN BIT).

C FOR OTHER SPEECH SAMPLE BYTE SIZES, FIGURE

C APPROXIMATELY 6 DB/BIT CHANGE ON THE VALUES OF

C CMAX AND CMIN. THIS IS ACCOMPLISHED BY SETTING

C THE VALUE OF ORCHANG TO BE EQUAL TO THE TOTAL CHANGE IN

C DB. THIS CHANGE IS POSITIVE IF THE SPEECH BYTE SIZE IS

C GREATER THAN 11 BITS AND NEGATIVE IF IT IS LESS.

C FOR EXAMPLE, IF THE BYTE SIZE IS 9, ORCHANG=-12.0

C DATA CMAX/46.8/,CMIN/-18.8/

C DATA ORCHANG/8.8/

C

C MLEV=2**MBITC

C CMAX=CMAX+ORCHANG

C CMIN=CMIN+ORCHANG

C

C CALL QTRLC(CMAX,CMIN,MLEV,RC,VC,CSTEP)

C

C CONVERT TABLE VALUES FOR DIRECT CODING

C OF ENERGY

C IC(1)=ROD(IC(1))

C DO 115 I=1,MLEV

C IC(I,1)=ROD(IC(I,1))

C VC(I,1)=SQRT(ROD(VC(I,1)))

C CONTINUE

C RETURN

C END

115

7L

```

SUBROUTINE VTRLS(NBITSC,MLEVSC,ISEC,VSEC)
C ROUTINE TO SET UP CODING TABLES FOR DELTA GAINS.
C INPUT:
C NBITSC NUMBER OF BITS FOR CODING DELTA GAINS
C ( SET UP FOR ONLY NBITSC = 2 BITS)
C OUTPUTS:
C MLEVSC NUMBER OF QUANTIZATION LEVELS
C ISEC TRANSMITTER CODING TABLE
C VSEC RECEIVER DECODING TABLE
C
C DIMENSION ISEC(1),VSEC(1)
C DIMENSION ISEC(5),VSEC(4)
C THE FOLLOWING TABLE VALUES ARE IN DBI
C DATA ISEC(1)/-15.0/,ISEC(2)/-3.6/,ISEC(3)/-0.5/
C DATA ISEC(4)/2.2/,ISEC(5)/30.0/
C DATA VSEC(1)/-6.2/,VSEC(2)/-2.0/,VSEC(3)/1.0/
C DATA VSEC(4)/3.5/
C MLEVSC=2**NBITSC
C
C CURRENT TABLE VALUES FOR DIRECT CODING
C UP ENERGY RATIOS
C ISEC(1)=DBI(ISEC(1))
C DO 100 I=1,MLEVSC
C ISEC(I+1)=MIN(ISEC(I+1))
C VSEC(I)=SQRT(DBI(VSEC(I)))
C CONTINUE
C RETURN
C END
100
~

```

```

SUBROUTINE COD(R,A,MLEV,IR)
C GENERAL CODING ROUTINE (TABLE LOOKUP)
C INPUT:
C R = PARAMETER TO BE CODED
C A = TRANSMITTER CODING TABLE ( QUANTIZER THRESHOLDS)
C MLEV= # OF QUANTIZATION LEVELS
C OUTPUTS:
C IR = TRANSMITTED CODE
C DIMENSION X(MLEV)
C
C DO 10 M=2,MLEV
C IF(R-LT.X(M)) GO TO 100
C CONTINUE
C M=MLEV+1
C IR=M-2
C RETURN
C END
10
100
~

```

```

SUBROUTINE DCOB(IR,Y,R)
C GENERAL DECODING ROUTINE (TABLE LOOKUP)
C INPUT:
C IR = RECEIVED CODE FOR PARAMETER R
C Y = RECEIVER DECODING TABLE
C OUTPUT:
C M = QUANTIZED VALUE OF R
C
C DIMENSION V(1)
C R=Y(IR+1)
C RETURN
C END
~L
C
C SUBROUTINE OPTQ(ITYP,MLEVS,X,Y)
C PROGRAM TO SET UP CODING AND DECODING TABLES FOR OPTIMUM
C BINARY QUANTIZATION. IT IS ASSUMED THAT THE INPUT TO
C THE QUANTIZER IS ZERO MEAN AND UNIT VARIANCE. THE
C ALGORITHM IS GIVEN IN MAX, "QUANTIZING FOR MINIMUM DISTORTION,"
C IRE TRANS. INFORM. THEORY, VOL. IT-6, PP.7-12, MAR 1960.
C
C INPUTS:
C ITYP MODEL P.D.F. 1=CANON 2=LAPLACIAN 3=GAUSSIAN
C MLEVS NUMBER OF QUANTIZATION LEVELS.
C OUTPUTS:
C X CODING TABLE.
C Y DECODING TABLE. X(I) < Y(I) < X(I+1).
C
C DIMENSION P(2000),X(1),Y(1)
C DATA BINWID/0.01, EPS/1.0E-10, TOL/0.002/
C
C EVALUATE MODEL DISTRIBUTION FUNCTION, P(X), AT INCREMENTS
C OF BINWID TO GENERATE A ONE-SIDED HISTOGRAM, P.
C P(I) = PROBABILITY THAT (I-1)*BINWID < ABS(X) < I*BINWID.
C
CUT=0.
SIGMA=B.
P2= 0.5
I=0
I= I+1
Z= FLOAT(I)*BINWID
P1=P2
IF(ITYP.NE.1) GO TO 05
IF(I.NE.1) GO TO 05
Z= SQRT(2.*Z)
CALL WTRZ(Z,P2,B)
IF(ITYP.EQ.2) P2= 1.-.5*EXP(-1.4142*Z)
IF(ITYP.EQ.3) CALL NTRH(Z,P2,B)
DIF= P2 -P1
C UNIT WHEN HISTOGRAM APPROACHES ZERO.
IF(DIF.LT.EPS) GO TO 70
P(I)= DIF
TOT= TOT+DIF
AR= FLOAT(I-1)
SIGMA= SIGMA +P(I)*(AR*AR +.333)
GO TO 75
SIGMA= SQRT(SIGMA/TOT)
MAR= I-1
C
C DETERMINE THE OPTIMAL (M.S.E.) QUANTIZER FOR THE GIVEN
C HISTOGRAM. START WITH GUESS OF INNERMOST DECODING LEVEL.
C UQSS= SIGMA/LOAT(MLEVS)

```

```

C GIVEN UN UNO NUMBER OF LEVELS.
  N= NUN*(NLEVS,2)
  NLEVS= NLEVS/2
  IF(N.EQ.0) X(1)= 0.
  IF(N.EQ.1) V(1)= 0.
C ITERATE UP TO 100 TIMES
  DO 10 I=1,100
    V(1,0)= GUESS
    IF(N.EQ.1) X(1)= V(1)/2.
C COMPOSITE QUANTILAR RECURSIVELY
    DO 50 I=1,NLEVS-1
      CALL ERECT(P,MAX,X(I),V(1,0),X(1,1),IPLAC)
      IF(.NOT.3PLAC) GO TO 30
      GUESS= GUESS*0.9
30   GO TO 40
50   V(1,0)= 2.*I(1,1) -V(1,0)
      CONTINUE
      AMP1= FLOAT(MAX)-EPS
      X(NLEVS,1)= AMP1
C CHECK CENTROID OF LAST BIN AND ADJUST GUESS.
      CALL CENTRD(P,MAX,X(NLEVS),TEST,AMP1)
      CUR= (TEST-V(NLEVS,0))/V(NLEVS,0)
      IF(ABS(CUR).LT.TOL) GO TO 95
      GUESS= V(1,0)+V(1,0)*CUR*0.75
C SCALE BY STD. DEV. AND MAKE SYMMETRIC W.R.T. 0.
95   IF(N.EQ.1) TEMP= V(NLEVS,1)
      DO 99 I=1,NLEVS-1
        X(NLEVS+1,0)= X(I)/SIGMA
        V(NLEVS+1,1)= V(I)/SIGMA
        X(NLEVS,1)= 1000.
        IF(N.EQ.1) V(NLEVS)= TEMP/SIGMA
      DO 100 I=1,NLEVS-1
        X(I)= -X(NLEVS+2-I)
        V(I)= -V(NLEVS+1-I)
100  RETURN
      END
C
C

```

```

C SUBROUTINE CENTRO(P,MAX,XBEG,Q,XEND)
C COMPUTES THE CENTROID OF A HISTOGRAM WITHIN A SPECIFIED
C INTERVAL. DETAILS OF THE NUMERICAL INTEGRATION METHOD
C ARE GIVEN IN SYRON AND PACE, "A COMPARISON OF OPTIMUM AND
C LOGARITHMIC QUANTIZATION FOR SPEECH PCM AND DPCM SYSTEMS,"
C IEEE TRANS. ON COMMUNICATIONS, JUNE 1973.
C INPUTS:
C P HISTOGRAM. P(K)= PROBABILITY OR NUMBER OF
C OBSERVATIONS OF X IN THE KTH 'BIN', K-1 < X < K.
C MAX NUMBER OF HISTOGRAM BINS.
C XBEG,XEND ENDPOINTS OF THE INTERVAL OVER WHICH THE CENTROID
C IS TO BE COMPUTED. B < XBEG < XEND < MAX.
C INPUTS:
C Q CENTROID (XBEG < Q < XEND).
C DIMENSION P(MAX)
C XBEG= IFIX(XBEG)+1
C XEND= IFIX(XEND)+1
C FM= .5*(XEND-XBEG)*P(XBEG)-XBEG*P(XBEG)
C FD= XEND*P(XEND)-XBEG*P(XBEG)
C IF(XEND.EQ.XBEG) GO TO 75
C XBEG AND XEND NOT IN THE SAME BIN.
C FM= FM+.5*(FLOAT(XBEG+2)*P(XBEG)-FLOAT((XEND-1)+2)
C *P(XEND))
C FD= FD+ FLOAT(XBEG)*P(XBEG) -FLOAT(XEND-1)*P(XEND)
C IF(XEND.LE.XBEG+1) GO TO 75
C XBEG AND XEND NOT IN ADJACENT BINS.
C DO 50 I1=XBEG+1,XEND-1
  FB= FM+ (FLOAT(I1)-.5)*P(I1)
  FD= FD+ P(I1)
50   Q= FM/FD
C RETURN
C END
C

```


C <OCAL0>SECUR.74.2 THU 21-NOV-68 9:37AM

3-ROUTINE FUNCTION,WFNAME,IFOLD)

C COUNT FROM NATURAL BINARY TO POLAR
C BINARY COUNT, OR VICA-VERSA. IS IS AN
C EMPTY/EMPTY VECTOR OF LENGTH NFRAME,
C WHICH CONTAINS THE CODE WORDS.

C
C

 DIMENSION IR(1)

 ND TO 1-1-NFRAME
 IF(ND(1)-1-N-IFOLD) IR(1)=1-ND-IR(1)
 CONTINUE
 RETURN
 END

10

```

C <COA16>I&ER.F0.2
C INDEX BIT COUNT - 4 BIT WORDS
15 I&ER-I&ER+1
C EXTRACT BITS FROM DATA - FOR 4 BIT WORD
I&ER(I)=0
IF(I&ER(I)-I&ER(I+1)).EQ.0) GO TO 28
C BIT EXTRACTED IS 1 - SET THIS BIT IN I&ER=0
I&ER-I&ER-I&ER(I&ER-I&ER+1)
C TURN ON IB
I&ER(I&ER)=1
28 CONTINUE
C CHECK IF 4 BITS COLLECTED
IF(I&ER(I&ER-I&ER)) GO TO 48
C 4 BITS COLLECTED - MAKE 4 BIT OUTPUT WORD
J=I&ER-I&ER+1
I&ER=I&ER+1
GO TO 15, I&ER
J=J-1
I&ER(I&ER)=I&ER(I&ER)+I&ER(I&ER+1, J)
CONTINUE
38 I&ER(I&ER)=I&ER(I&ER)
C RESET 4 BIT WORD INDEX - READY NEW WORD
I&ER=0
C CHECK IF PARAMETERS FINISHED
IF(I&ER-I&ER) GO TO 50
C WORD NOT FINISHED CONTINUE
IP=IP+1
GO TO 15
C WORD FINISHED - SET UP SUB WORD & SUB WORD BITS
50 I&ER(I&ER)=I&ER
I&ER(I&ER)=I&ER(I&ER)
IF(I&ER(I&ER)) RETURN
C NOT DONE - GET NEW PARAMETER
I&ER=I&ER+1
GO TO 18
RETURN
END

```

```

C <COA16>I&ER.F0.2
C COUNTING EXTRINSIC, I&ER, I&ER, I&ER, I&ER, I&ER, I&ER, I&ER
I&ER(I&ER)
C ROUTINE TO EXTRACT BITS TO BE PROTECTED FROM CODED PARAMETERS
C PUMP-BIT WORDS ARE FORMED PRIOR TO ERROR PROTECTION
C VIA MATH(7,4) ON (0,4) CODE.
C INPUTS:
C I&ER VECTOR OF QUANTIZED LEVELS OF A PRESPECIFIED
C SEQUENCE OF TRANSMITTED PARAMETERS.
C I&ER VECTOR GIVING THE NUMBER OF BITS PROTECTED
C CORRESPONDING TO THE SEQUENCE IN I&ER
C I&ER VECTOR GIVING TOTAL NUMBER OF BITS FOR EACH
C PARAMETER (SEQUENCE AS IN I&ER)
C I&ER VECTOR OF PERCENT CHANNEL BIT ERROR FOR EACH PARAMETER
C (SEQUENCE AS IN I&ER)
C I&ER TOTAL NUMBER OF PARAMETERS
C I&ER TOTAL NUMBER OF 4-BIT WORDS TO BE ERROR PROTECTED
C OUTPUTS:
C I&ER COUNTER OF THE NUMBER OF 4-BIT WORDS FORMED
C I&ER VECTOR OF 4-BIT WORDS TO BE ERROR PROTECTED
C I&ER VECTOR OF PERCENTAGE OF CHANNEL BIT ERROR
C FOR EACH 4-BIT WORD IN I&ER
C I&ER VECTOR OF SUB WORDS I.E. WHAT IS LEFT
C AFTER EXTRACTING PROTECTED BITS FROM
C PARAMETERS CODED
C I&ER VECTOR OF BITS/SUB WORD
C
C DIMENSION I&ER(1), I&ER(1), I&ER(1), I&ER(1), I&ER(1)
C DIMENSION I&ER(4), I&ER(1), I&ER(1)
C DIMENSION I&ER/I&ER(32)
C DATA I&ER/4/
C INITIALIZE COUNTS
I&ER=0
I&ER=0
I&ER=0
C ZERO 4 BIT WORDS
CALL SETEM(I&ER(1), I&ER)
IF(I&ER(I&ER)) RETURN
C SET UP FOR NEW PARAMETER
C PUT DATA, PROTECTED BITS, & TOTAL BITS IN DUMMY VARIABLES
I&ER=I&ER(I&ER)
I&ER=I&ER(I&ER)
I&ER=I&ER(I&ER)
I&ER=I&ER(I&ER)
I&ER=I&ER(I&ER)
I&ER=I&ER(I&ER)
C SET I&ER TO 1 - 0 PROTECTED BITS
I&ER=1

```

```

SUBROUTINE ERCHN(LMD,MB,PCTE,INDOUT,KRESFG)
C CHANGES BITS (0 TO 1, UM 1 TO 8) OF A
C WORD TO SIMULATE THE EFFECT OF CHANNEL BIT ERRORS
C INPUTS:
C LMD INPUT WORD TO BE ERRURED
C MB TOTAL NUMBER OF BITS OF LMD
C PCTE PERCENT ERROR FOR LMD
C KRESFG FLAG TO INITIATE COUNT OF RESIDUAL BIT ERRORS
C OUTPUT:
C INDOUT OUTPUT ERRURED WORD
C COMMON/CUM/ICUMBT,IERCUM,ISERR
C ICUMBT= COUNTER OF TOTAL BITS PROCESSED
C IERCUM= COUNTER OF TOTAL BITS WHICH HAVE BEEN ERRURED
C ISERR= COUNTER OF TOTAL NUMBER OF RESIDUAL BITS ERRURED
CUMBT/IMASK/IMASK(32)
INDOUT=LMD
IF(MB.EQ.0) RETURN
IF(PCTE.GE.0.) GO TO 50
DO 10 I=1,MB
C CALL RANDOM NUMBER GENERATOR FUNCTION
C (NOTE: THIS FUNCTION NOT PART OF SUPPLIED SOFTWARE
C USER MUST SUBSTITUTE APPROPRIATE ROUTINE)
R=RANDOM(0.,0.,1-0)
IF(R.LT.PCTE) GO TO 20
GO TO 30
20 INDOUT=INDOUT.XOR.IMASK(I)
IF(KRESFG)ISERR=ISERR+1
IERCUM=IERCUM+1
ICUMBT=ICUMBT+1
CONTINUE
RETURN
C NO ERROR
50 ICUMBT=ICUMBT+MB
RETURN
END

```

```

SUBROUTINE COON(LMD,NPPT,IMWD)
C CUMBT 4-BIT WORDS INTO (7,4) UM (0,4) WARNING CODEWORDS
C INPUTS:
C LMD 4-BIT DATA WORD TO BE CODED
C NPPT SWITCH TO SPECIFY (0,4) OR (7,4) WARNING CODEWORD
C (7,4) WARNING CODING
C (0,4) WARNING CODING
C OUTPUT:
C IMWD (7,4) UM (0,4) WARNING CODEWORD
C BIT(1) UM(1),IC(1),IV(0)
C EVALUATE((10(1),10(1)),((10(1)),((10(1)),10(5)))
C COMMON/IMWD/IMWD(8)
CUMBT/IMASK/IMASK(32)
DATA ICNST/4/
C INITIALIZE
IMWD=0
IMWD=0
C EXTRACT BITS
DO 10 I=1,ICUMBT
IM(I)=0
IF((IM(I).AND.IMASK((ICNST-1+I)).EQ.0) GO TO 10
I=I+1
CONTINUE
C COMPUTE CHECK BITS
IC(1)=MOD((IM(1)+IM(2)+IM(3)+IM(4)),2)
IC(2)=MOD((IM(1)+IM(3)+IM(4)),2)
IC(3)=MOD((IM(2)+IM(3)+IM(4)),2)
FOR (0,4) CHECK COMPUTE AN ADDITIONAL PARITY CHECK BIT
IF(NPPT.EQ.4)IC(4)=MOD((IM(1)+IC(1))+IC(2))+2)
C MAKE 8-BIT WORD FROM UM=IM+IC
DO 20 I=1,NPPT+4
I=I+1
IMWD=IMWD+UM+IMAND(I)
CONTINUE
RETURN
END

```

```

C EXTRACT 4 DATA BITS FROM N4M WORD
J=ICUMST
DO 50 I=1,ICUMST
  TURN ON IWORD BIT J (IE SET BIT J TO 1)
  IWORD=IWORD.OR.IMASK(J)
  TURN OFF IWORD BIT J IF IWORD I BIT IS ZERO
  IF((IWORD.AND.IMASK(1)).EQ.0)IWORD=IWORD.IUR.IMASK(J)
  J=J-1
50 CONTINUE
RETURN
END
-L

```

```

SUBROUTINE DCUM(IWD,SUPT,IWORD,NB4FC)
  DCUMES 7,4 OR 0,4 MARKING CODE WORD
  C OUTPUT IS A 4-BIT ERROR CORRECTED A4M4
  C INPUTS:
  C IWD 7,4 OR 0,4 MARKING CODEWORD
  C SUPT SWITCH TO SPACE (0,4) OR (7,4) MARKING CODE A4M4
  C 3, (7,4) CODEWORD
  C 5, (0,4) CODEWORD
  C OUTPUTS:
  C I4M4 4-BIT ERROR CORRECTED WORD
  C NB4FC FLAG TO INDICATE EVEN ERRORS(0,4 CODE ONLY)
  C -1, 5BZ ERRORS DETECTED
  C 0, 4BZ ERRORS DETECTED

```

```

  DIMENSION IRI(4),IC(4),IM(0)
  EQUIVALENCE (IRI(1),I4(1)),(IC(1),IIC(5))
  COMMON/I4M4/IMASK(8)
  COMMON/IMASK/IMASK(32)
  DATA ICUMST/4/

```

```

C INITIALIZE
IWORD=IWD
IWORD=0
NB4FC=0
C EXTRACT BITS FROM N4M WORD
DO 10 I=1,NOPT+4
  IM(I)=0
  IF((IWORD.AND.IMASK(I)).NE.0) IM(I)=1
  CONTINUE
C IF 7,4 CODE - CO COMPUTE CHECK BITS
  IF(SUPT.EQ.4) GO TO 40
  C 0,4 CODE - GET PARITY BIT
  ISUM=0
  DO 20 I=1,8
    ISUM=ISUM+I4(I)
  CONTINUE
  IC(4)=IWORD(I5UM,2)
  C COMPUTE CHECK BITS
  IC(1)=IWORD(I4(1))+IC(2)+I4(2)+I4(3)+I4(4)+I4(5)+I4(6)+I4(7)+I4(8)
  IC(2)=IWORD(I4(1))+IC(3)+I4(3)+I4(4)+I4(5)+I4(6)+I4(7)+I4(8)
  IC(3)=IWORD(I4(2))+IC(4)+I4(2)+I4(3)+I4(4)+I4(5)+I4(6)+I4(7)+I4(8)
  C COMPUTE CORRECTIVE BIT LOCATION
  ICUR=LSH(IC(1),2)+LSH(IC(2),1)+IC(1)
  C FLIP BIT IF ERROR = 7,4 OR 0,4
  IF(ICUR.NE.0) IWORD=IWORD.XOR.IMASK(ICUR)
  C IF 0,4 TURN ON FLAG IF TWO ERRORS
  IF((IMPT.EQ.4).AND.(IC(4).EQ.0).AND.(IC(7).NE.0)) NB4FC=-1

```

```

C COMPLETE WORD
50 IDAT(MW)=IDAT(MW)+IDATS(MW)
C IF COMPLETE RETURN
IF(MW.EQ.MDATT) RETURN
C MORE LEFT TO DO
MW=MW+1
J=J+1
L=L-1
IF(MW.EQ.0) GO TO 50
C CHECK IF SINGLE BIT WORDS USED UP
GO TO 30
RETURN
END

```

```

SUBROUTINE CUMWD(LNW,MP,ND,MDATT,IDATS,IDAT)
C CUMBINES 4-BIT ERROR CORRECTED WORDS
AND SUB WORDS TO RECREATE PARAMETER CODES
C INPUTS:
C LNW VECTOR OF ERROR CORRECTED 4-BIT WORDS
C MP VECTOR OF NUMBER OF PROTECTED BITS OF PARAMETER CODES
C ND VECTOR OF NUMBER OF BITS OF PARAMETER CODES
C MDATT NUMBER OF PARAMETER CODES
C IDATS VECTOR OF SUB WORDS
C (NOTE: ALL ARRAYS ARE OF A PRESPECIFIED SEQUENCE)
C OUTPUTS:
IDAT = VECTOR REGENERATED PARAMETER CODES
C
C
C DIMENSION IM(1),MP(1),ND(1),IDATS(1),IDAT(1)
C DIMENSION IB(4)
C COMMON/IMASK/IMASK(32)
C DATA ICNST/4/
C
C SET INDICES
L=0
A=0
J=ND(1)
MW=1
ND=1
ICNST=ICNST+1
C ZERO DATA WORDS
CALL SETM(IDAT(1),MDATT)
C SEPARATE 4-BIT WORDS INTO SINGLE BIT WORDS IB
DO 20 I=1,4
IB(I)=
IF((ND(MW)-AND(IMASK(ICNST-I))).EQ.0) GO TO 20
L=L+1
CONTINUE
IF(MP(MW)-EQ.0) GO TO 50
C LUMP UP FORMATION OF DATA WORDS
J=J-1
E=E+1
L=L+1
IDAT(MW)=IDAT(MW)+LSH(IB(L),J)
C CHECK IF PROTECTED BITS COMPLETE
IF(L.LT.MP(MW)) GO TO 30
C PROTECTED BITS COMPLETE - COMPLETE WORD
GO TO 50
C CHECK IF ANY SINGLE BIT WORDS LEFT
IF(L.LT.ICNST) GO TO 40
C OUT OF SINGLE BIT WORDS - GET NEXT IN'S
MW=MW+1
E=E+1
GO TO 10

```

```

C (UCA16)16KER.F4.2
C COMPLETE WORD
50 IDAT(MW)=IDAT(MW)+IDATS(MW)
C IF COMPLETE RETURN
IF(MW.EQ.MDATT) RETURN
C MORE LEFT TO DO
MW=MW+1
J=J+1
L=L-1
IF(MW.EQ.0) GO TO 50
C CHECK IF SINGLE BIT WORDS USED UP
GO TO 30
RETURN
END

```

SHOULDN'T BE HERE. I, ER, ERP

C PUTS PARAMETER CODES AND CHANNEL ERROR PERCENT
C FOR EACH PARAMETER INTO THE ARRAYS IDAT AND ER
C RESPECTIVELY. SEQUENCE OF CODE AND ERROR STORAGE IN THESE ARRAYS
C IS PREDEFINED BY THE ARRAY ISEQ. ISEQ CONTAINS
C NUMBERS WHICH CORRESPOND TO THE PARAMETER TYPES AS GIVEN BELOW.

C INPUTS:
C ISEQ TOTAL NUMBER OF PARAMETER CODES
C IDAT VECTOR OF NUMBERS DESCRIBING THE SEQUENCE OF
C PARAMETER CODES
C ERP PERCENT OF CHANNEL ERROR FOR EACH PARAMETER TYPE
C
C ERP(1) - REFLECTION COEFFS
C ERP(2) - PITCH
C ERP(3) - OVERALL GAIN
C ERP(4) - DELTA GAINS
C ERP(5) - PITCH PREDICTOR TAPS
C ERP(6) - RESIDUAL SAMPLES
C
C SINGLE PARAMETER CODES VARIABLE ISEQ #
C TYPE -----
C DELTA COEFF I TO 14
C PITCH 15
C GAIN 16 TO 18
C DELTA GAINS IEM5
C PITCH TAPS IPTAP 21 TO 30
C RES SAMPLES IR 31 TO 35
C 36 ----

C UNTPUTS:
C IDAT VECTOR OF PARAMETER CODES
C ISEQ SEQUENCE AS IN ISEQ
C ER VECTOR OF PERCENT ERROR FOR EACH PARAMETER CODE
C
C ISEQ SEQUENCE AS IN ISEQ

C DIMENSION ISEQ(1),ICM(1),IEMS(1),IPTAP(1),IR(1)
C DIMENSION ER(1),ERP(1),IDAT(1)

C DO 1000 I=1,NDATT
C IPNAM=ISEQ(I)
C REFLECTION COEFFS
C IF(IPNAM.LE.14) GO TO 10
C PITCH
C IF(IPNAM.EQ.15) GO TO 20
C GAIN
C IF(IPNAM-15.LE.5) GO TO 30
C DELTA GAINS
C IF(IPNAM-20.LE.10) GO TO 40

C <OCAL0>IGERR.F4.2 Thu 21-Aug-88 9:39AM
C PITCH TAPS
C IF(IPNAM-30.LE.5) GO TO 50
C GO TO 60
C REFLECTION COEFFS
C IDAT(I)=ICM(IPNAM)
C ER(I)=ERP(I)
C GO TO 1000
C PITCH
C IDAT(I)=IPTCH
C ER(I)=ERP(2)
C GO TO 1000
C GAIN
C ICMX=IPNAM-15
C IDAT(I)=IEM5
C ER(I)=ERP(3)
C GO TO 1000
C DELTA GAINS
C ICMX=IPNAM-20
C IDAT(I)=IEMS(IIDX)
C ER(I)=ERP(4)
C GO TO 1000
C PITCH TAPS
C ICMX=IPNAM-30
C IDAT(I)=IPTAP(IIDX)
C ER(I)=ERP(5)
C GO TO 1000
C RESIDUAL SAMPLES
C ICMX=IPNAM-35
C IDAT(I)=IR(IIDX)
C ER(I)=ERP(6)
C CONTINUE
C RETURN
C END

```

SUBROUTINE SEQW(MDATE,ISEW,ICW,IPTCM,IEMW,IENS,IPTAP,IR,IQAT)
C DERIVE INDIVIDUAL PARAMETERS ON RESIDUAL SAMPLES
C FROM THE COMPOSITE ARRAY IQAT.(SEE SEQX FOR DEFINITION
C OF ARGUMENTS)
C
DIMENSION ISEQ(1),ICQ(1),IENS(1),IPTAP(1),IR(1)
DIMENSION IQAT(1)
C
DO 1000 I=1,MDATT
  IPRM=ISEQ(I)
C REFLECTION COEFFS
  IF(IPRM.LE.14) GO TO 10
C PITCH
  IF(IPRM.EQ.15) GO TO 20
C GAIN
  IF(IPRM-15.LE.5) GO TO 30
C DELTA GAINS
  IF(IPRM-20.LE.10) GO TO 40
C PITCH TAPS
  IF(IPRM-30.LE.5) GO TO 50
  GO TO 60
C
C REFLECTION COEFFS
  ICQ(IPRM)=IQAT(I)
  GO TO 1000
C PITCH
  IPTCH=IQAT(I)
  GO TO 1000
C GAIN
  IEMG=IQAT(I)
  GO TO 1000
C DELTA GAINS
  INDE=IPRM-20
  IENS(IEMG)=IQAT(I)
  GO TO 1000
C PITCH TAPS
  INDI=IPRM-30
  IPTAP(INDI)=IQAT(I)
  GO TO 1000
C RESIDUAL SAMPLES
  INDE=IPRM-35
  IN(IINDE)=IQAT(I)
  CONTINUE
  RETURN
  END

```

```

SUBROUTINE BSTWMI(IR,NBPS,NSMP,NST,IB,NBITT)
C BIT STREAMING OF THE RESIDUAL SAMPLES
C
C EXTRACTS BITS FROM THE RESIDUAL SAMPLE PARAMETER CODES
C (IN ARRAY IR) AND FORMS AN ARRAY IB OF RESIDUAL
C SAMPLE BITS (1 DATA BIT /WORD)
C
C INPUTS:
C IR VECTOR OF RESIDUAL SAMPLE CODES
C NBPS VECTOR OF BITS PER RESIDUAL SAMPLE
C (THIS ARRAY ALLOWS A VARIABLE # OF BITS/SAMPLE)
C NSMP TOTAL NUMBER OF RESIDUAL SAMPLE CODES
C NST INDEX OF THE FIRST ENTRY IN NBPS
C
C OUTPUTS:
C IB VECTOR OF RESIDUAL SAMPLE BITS(1 DATA BIT / WORD)
C NBITT TOTAL NUMBER OF 1 DATA BIT WORDS IN IB
C
DIMENSION IR(1),NBPS(1),IB(1)
COMMON/IMASK/IMASK(32)
C
C COUNT ON BITS TOTAL
  NBITT=0
C INDEX FOR BITS/SAMPLE
  II=NST-1
C LOOP FOR FRAME
  DO 100 I=1,NSMP
    C UPDATE INDEX FOR BITS/SAMPLE
      II=II+1
    C CHECK IF ZERO BITS USED - IF SO GO TO NEXT SAMPLE
      IF(NBPS(II).EQ.0) GO TO 100
    C BITS NOT ZERO - PUT SAMPLE IN DUMMY VARIABLE
      IDUM=IR(II)
    C INITIALIZE BITS/SAMPLE COUNTER
      IBITT=1
    C UPDATE TOTAL BIT COUNTER
      NBITT=NBITT+1
    C INITIALIZE BIT ARRAY
      IN(MBITT)=0
      IF((IDUM.AND.IMASK(NBPS(II)-IBITT+1)).EQ.0) GO TO 20
      IDUM=(IDUM-IMASK(NBPS(II)-IBITT+1))
    C TURN ON IB
      IN(MBITT)=1
    CONTINUE
      IF(IBITT.EQ.NBPS(II)) GO TO 100
      IBITT=IBITT+1
      GO TO 15
    CONTINUE
  RETURN
  END

```

```

C IF SO DONT ENTER BIT LOOP - GO TO NEXT SAMPLE
IF(J,C2,0) GO TO 100
CONTINUE
MSMP=MSMP-1
RETURN
END

```

```

SUBROUTINE SYSTEMATION,MBIT,MBPS,NST,IRR,MSBUC,MSMP)
C FORMS RESIDUAL SAMPLE CODES FROM
C THE BIT STREAM FORMED IN STRM1
C INPUTS:
C IRR VECTOR OF RESIDUAL SAMPLE BITS(1 DATA BIT / WORD)
C MBIT TOTAL NUMBER OF 1 BIT DATA WORDS IN IRR
C MBPS VECTOR OF # OF BITS /RESIDUAL SAMPLE CODE
C NST INDEX OF FIRST ENTRY IN MBPS
C MSBUC TOTAL NUMBER OF RESIDUAL SAMPLE CODES
C OUTPUTS:
C MSMP COUNTER OF NUMBER OF RESIDUAL SAMPLE CODES
C IRR VECTOR OF RESIDUAL SAMPLE CODES

```

```

C DIMENSION IRR(1),MBPS(1),IRR(1)
COMMON/INASK/INASK(12)

```

```

C ENTER IN BITS/SAMPLE ARRAY
II=NST
MSMP=1
C INITIALIZE SAMPLE COUNT
C ENTER IN BITS/SAMPLE
J=MBPS(II)-1
C SET SAMPLE TO ZERO
IRR(MSMP)=0
C CHECK IF BITS = 0 - IF SO DONT ENTER BIT STREAM LOOP
C LOOK AT NEXT SAMPLE
IF(J,C2,0) GO TO 20
II=II+1
MSMP=MSMP+1
GO TO 10

```

```

C BIT STREAM LOOP
CONTINUE
J 100 I=1,MBIT
IRR(MSMP)=IRR(MSMP)+LSW(IRR(I),J)
J=J-1
IF(J,E4,-1) GO TO 50
GO TO 100

```

```

C SAMPLE FINISHED - RESET
MSMP=MSMP+1
IRR(MSMP)=0
II=II+1
J=MBPS(II)-1
C IF TOT NUM OF SAMPS - JUMP OUT
IF(MSMP-UT-MSBUC) GO TO 110
C CHECK IF NEXT SAMP HAS ZERO BITS

```

```

C      FUNCTION ENERGY(S,MAXS,NST,MEAN)
C      COMPUTES ENERGY IN MEAN-SQUARED VALUE.
C      ARGUMENTS:
C      S      VECTOR OF SAMPLES.
C      MAXS  MAXIMUM DIMENSION OF S.
C      NST,MEAN INDICES OF THE FIRST AND LAST SAMPLES FOR WHICH
C      THE MEAN-SQUARED VALUE IS TO BE COMPUTED.
C      DIMENSION S(MAXS)
C      ENERGY=0.0
C      NTOT=MEAN-NST+1
C      DO 10 I=NST,MEAN
C      ENERGY=ENERGY+S(I)*S(I)
C      RETURN
C      END

```

10

-L

```

C      SUMMING ZERODC(S,MAXS,NST,MEAN,TEMP)
C      COMPUTES AND SUBTRACTS DC OR MEAN VALUE.
C      INPUTS:
C      S      VECTOR OF INPUT SAMPLES
C      (DC-REMOVED OUTPUT STORED ALSO IN S).
C      MAXS  DIMENSION OF S.
C      NST,MEAN INDICES OF THE FIRST AND LAST SAMPLES FOR WHICH
C      THE MEAN IS TO BE REMOVED.
C      OUTPUT:
C      TEMP  MEAN VALUE.
C      DIMENSION S(MAXS)
C      TEMP=0.0
C      NTOT=MEAN-NST+1
C      DO 10 I=NST,MEAN
C      TEMP=TEMP+S(I)
C      DO 20 I=NST,MEAN
C      S(I)=S(I)-TEMP
C      RETURN
C      END

```

10

20

-L

```

C SUBROUTINE DEEMPH(S,MAIS,NST,NEH,DEMFAC)
C 1ST ORDER DEEMPHASIS.
C INPUTS:
C S VECTOR OF INPUT SAMPLES.
C (DEEMPHASIZED SAMPLES ARE ALSO STORED IN S)
C MAI MAXIMUM DIMENSION OF S
C NST,NEH INDICES OF THE FIRST AND LAST SAMPLES FOR WHICH
C DEEMPHASIS IS TO BE PERFORMED.
C DEMFAC DEEMPHASIS FILTER CONSTANT.
C OUTPUT:
C S DIMENSION S(MAIS)
C DO 10 I=NST,NEH
C S(I)=S(I)*DEMFAC*(I-1)
C RETURN
C END

```

```

C SUBROUTINE PREEM(X,Y,MAI,NST,NEH,PEMFAC)
C 1ST ORDER PREEMPHASIS.
C INPUTS:
C X VECTOR OF INPUT SAMPLES
C MAI MAXIMUM DIMENSION OF X AND Y
C NST,NEH INDICES OF THE FIRST AND LAST SAMPLES FOR WHICH
C PREEMPHASIS IS TO BE PERFORMED.
C PEMFAC PREEMPHASIS FILTER CONSTANT((1-PEMFAC)**-1))
C OUTPUT:
C Y VECTOR OF PREEMPHASIZED SPEECH SAMPLES
C DIMENSION Y(MAI),Y(MAI)
C DO 10 I=NEH,NST,-1
C Y(I)=X(I)-PEMFAC*(I-1)
C RETURN
C END

```

C SUBROUTINE RES(S,MAKS,R,MAKR,A,BSTR,MP,BST,HEM)

C LPC (SPECTRAL) INVERSE FILTERING TO COMPUTE RESIDUAL

C INPUTS:
 C S VECTOR OF INPUT (SPEECH) SAMPLES.
 C MAKS DIMENSION OF S.
 C MAKR DIMENSION OF R.
 C A PREDICTOR COEFFICIENTS.
 C BSTR INDEX OF FIRST OUTPUT SAMPLE.
 C MP NUMBER OF POLES (=LENGTH OF A).
 C BST,HEM INDICES OF THE FIRST AND LAST INPUT SAMPLES
 C TO BE FILTERED.

C OUTPUT: VECTOR OF RESIDUAL SAMPLES.

C DIMENSION S(MAKS),R(MAKR),A(MP)

C M=MAKR-MST
 DO 100 I=MST,HEM
 SHAT=0.0
 DO 10 K=1,MP
 J=I-K
 SHAT=SHAT-A(K)*S(J)
 RE(K)=S(I)-SHAT
 RETURN
 END

10
 100
 -L

C SUBROUTINE SYNTH(S,MAKS,A,R,MAKR,BSTR,MP,BST,HEM)

C LPC (SPECTRAL) SYNTHESIS.

C INPUTS:
 C M VECTOR OF EXCITATION SAMPLES.
 C MAKR MAXIMUM DIMENSION OF R.
 C MAKS MAXIMUM DIMENSION OF S.
 C A PREDICTOR COEFFICIENTS.
 C BSTR INDEX OF FIRST INPUT SAMPLE.
 C MP NUMBER OF POLES (=LENGTH OF A).
 C BST,HEM INDICES OF THE FIRST AND LAST OUTPUT SAMPLES
 C TO BE FILTERED.

C OUTPUT: VECTOR OF SYNTHESIZED (SPEECH) SAMPLES.

C DIMENSION S(MAKS),A(MP),R(MAKR)

C M=MAKR-MST
 DO 100 I=MST,HEM
 J=M+I
 TEMP=R(J)
 JJ=0
 DO 10 KK=1,MP
 JJ=JJ+1
 INJJ=I-JJ
 TEMP=TEMP+A(KK)*S(INJJ)
 S(I)=TEMP
 CONTINUE
 RETURN
 END

10
 100
 -L

```

SUBROUTINE AUTCOM(S,MAES,AR,ARMAX,NST,NEP,NP)
C COMPUTE AUTOCORRELATION FUNCTION.
C
C INPUTS:
C S VECTOR OF (WINDOWED) SAMPLES.
C MAES DIMENSION OF S.
C ARMAX DIMENSION OF AR.
C NST,NEP INDICES OF THE FIRST AND LAST SAMPLES FOR WHICH
C THE AUTOCORRELATION FUNCTION IS COMPUTED.
C NP MAXIMUM LAG OF INTEREST.
C
C OUTPUTS:
C AR AUTOCORRELATION COEFFICIENTS.
C NOTE THAT AR(1) IS THE ZEROth AUTOCORRELATION
C COEFFICIENT AND AR(NP+1) IS THE AUTOCORRELATION
C COEFFICIENT FOR THE LAG NP.
C
C DIMENSION S(MAES),AR(MAXAR)
C
C NP1=NP+1
C DO 100 J=1,NP1
C RLOD=NEP-J+1
C SUM=0.0
C DO 50 N=NST,MLOUP
C R2=N+J-1
C SUM=SUM+S(N)*S(R2)
C CONTINUE
C AR(J)=SUM
C RETURN
C END

```

```

SUBROUTINE HANNING(S,MAES,NST,NEP,SUM)
C HANNING WINDOW.
C
C INPUTS:
C S VECTOR OF INPUT SAMPLES.
C MAES DIMENSION OF S.
C NST,NEP INDICES OF THE FIRST AND LAST INPUT SAMPLES
C TO BE WINDOWED.
C
C OUTPUTS:
C SUM VECTOR OF HANNING WINDOWED SAMPLES.
C
C DIMENSION S(MAES),SUB(MAES)
C
C SUM=PI/DI,TWOPI
C
C PI2=MAES-NST+1
C PI2=PI2*PI2/PLDPT(MPT2)
C PI=PI-1.0
C DO 10 I=NST,NEP
C PI=PI+.0
C SUB(I)=S(I)*(ALPHA-ALPHA*PI+COS(PI*PI2/I))
C CONTINUE
C RETURN
C END

```



```

C SUBROUTINE PREDI,ENAT,IMPR,F,LF,MPITCH)
C LINEAR PREDICTION OF A SAMPLE VALUE FROM PREVIOUS SAMPLES OF
C THE SAME SIGNAL.
C
C INPUTS:
C I VECTOR OF INPUT SAMPLES.
C IMPR INDEX OF THE SAMPLE WHOSE VALUE IS TO BE PREDICTED.
C F COEFFICIENTS OF PREDICTION FILTER.
C LF NUMBER OF PREDICTION FILTER COEFFICIENTS.
C MPITCH IF MPITCH>0, THEN PITCH PREDICTION IS PERFORMED,
C USING MPITCH AS THE PITCH PERIOD. IN THIS CASE,
C F CONTAINS THE PITCH FILTER TAPS AND LF
C SHOULD BE LDF. IF MPITCH=0, THEN SPECTRAL
C LINEAR PREDICTION IS PERFORMED.
C
C OUTPUTS:
C ENAT PREDICTION OF X(IMPR).
C
C DIMENSION X(1),F(1)
C
C IF(MPITCH.EQ.0) M=0
C IF(MPITCH.GT.0) M=MPITCH-(LF+1)/2
C ENAT=0.
C IF(LF.LT.1) GO TO 10
C DO 10 A=1,LF
C I=IMPR-M+1
C ENAT=ENAT+F((I)*X(A))
C CONTINUE
C RETURN
C END
    
```

10 7

```

C SUBROUTINE PITCH(S,NPTS,LOG2N,IF0N,LF0L,MPIT,FR,FL,PI)
C COMPUTE PITCH, USING THE AUTOCORRELATION METHOD.
C
C INPUTS:
C S VECTOR OF INPUT SAMPLES FROM WHICH THE PITCH
C IS ESTIMATED.
C NPTS NUMBER OF SAMPLES OF S.
C LOG2N A (2**LOG2N)-POINT FFT IS USED TO EVALUATE THE
C AUTOCORRELATION FUNCTION.
C IF0N, SMALLEST AND LARGEST LAGS CORRESPONDING TO PERMISSIBLE
C VALUES OF PITCH.
C OUTPUTS:
C MPIT PITCH, MEASURED IN NUMBER OF SAMPLES.
C FR AUTOCORRELATION COEFFICIENTS (TO BE USED LATER FOR
C COMPUTING THE PITCH FILTER TAPS).
C SCRATCH VECTORS:
C FR,FL 2**LOG2N IN THE CALLING PROGRAM.
C
C DIMENSION S(1),FR(1),FL(1),PH(1)
C
C NFFT=2**LOG2N
C NFFT2=NFFT/2
C DO 50 I=1,NPTS
C FR(I)=S(I)
C CALL ZERUOC(FR,NPTS,1,NPTS,PMEAN)
C CALL WINDOW.
C CALL HANNING WINDOW.
C DO 1777 I=NPTS+1,NFFT
C PH(I)=0.
C DO 7777 I=1,NPTS
C PH(I)=FR(I)
C COMPUTE DFT OF REAL SIGNAL.
C (FFTH IS PART OF A NON-SOURCE LIBRARY, AND IS
C NOT INCLUDED IN THIS SUBROUTINE PACKAGE.)
C CALL FFTR(LOG2N,NFFT,PH,FR,PI)
C COMPUTE SPECTRUM.
C DO 7799 I=1,NFFT2+1
C FR(I)=FR(I)**2+PI(I)**2
C MAKE SYMMETRIC D.R.T. FR(NFFT2+1), SO THAT
C INVERSE DFT WILL BE REAL.
C DO 2777 I=2,NFFT2
C N=NFFT2-I
C FR(N)=FR(I)
C ZERO OUT IMAGINARY PART.
C DO 3777 I=1,NFFT
C FL(I)=0.
C INVERSE DFT.
C CALL FFTR(LOG2N,NFFT,FR,FR,PI)
C PHAT=0.
    
```

```

C SEARCH FOR LARGEST AUTOCORRELATION COEFFICIENT.
DO 4777 I=1,IPMS1,1/BL,1
IF (R(I)-LT-PHASE) GO TO 4777
PHASE=PR(I)
MPIT=I-1
LUNTIME=
RETURN
END
4777
~

C SUBROUTINE PTAPS(AUTOC,NTAPS,TAPS,MPIT,NO)
C COMPUTE PITCH FILTER TAPS.
C
C INPUTS:
C AUTOC AUTOCORRELATION COEFFICIENTS, WITH
C AUTOC(I) BEING THE ZEROTH COEFF.
C NTAPS NUMBER OF PITCH FILTER TAPS (NO MORE THAN 5).
C MPIT PITCH, MEASURED IN NUMBER OF SAMPLES.
C
C OUTPUTS:
C TAPS FILTER TAPS, SMALLEST LAG FIRST.
C NO THE TAP FOR A 1-TAP PITCH FILTER.
C (TO BE USED IN CASE THE MULTI-TAP CASE
C IS UNSTABLE)
C
C DIMENSION AUTOC(I),TAPS(I)
C DIMENSION RMS(S),PEU(S)
C
C NO=0.
C CALL SETM(TAPS,NTAPS)
C IF (NTAPS.LE.0) RETURN
C EXPLICIT SOLUTION FOR THE 1-TAP CASE
TAPS(1)=-AUTOC(MPIT+1)/AUTOC(1)
NO=TAPS(1)
IF (NTAPS.LE.1) RETURN
II=1
C GENERATE THE RIGHT-HAND-SIDE VECTOR RMS
C ON THE NORMAL EQUATIONS.
DO 10 I=MPIT-NTAPS/2,MPIT+NTAPS/2
RMS(II)=-AUTOC(II+1)
II=II+1
10
C LAYMSON RECURSION FOR THE MULTITAP CASE.
CALL LEVREC(NTAPS,AUTOC,RMS,TAPS,PEO)
RETURN
END
~

```

```

SUMMATION TRANSFORM (U,ITTP)
C
C ROUTINE TO LINEARLY TRANSFORM (AND INVERSE TRANSFORM)
C THE PITCH PREDICTOR TAP COEFFICIENTS FOR
C THE 3-TAP CASE.
C ITTP= FALSE, FORWARD TRANSFORM
C TRUE, INVERSE TRANSFORM
C
C DIMENSION M(1),PSI(3)
DO I =1,3
  PSI(I)= 0(1)
  IF(ITTP) GO TO 10
C FORWARD TRANSFORM
  M(1)= PSI(1)+PSI(2)+PSI(3)
  M(2)= PSI(1)-2.*PSI(2)+PSI(3)
  M(3)= PSI(1)-PSI(3)
  GO TO 100
C INVERSE TRANSFORM
  M(1)= (2.*PSI(1)+PSI(2)+3.*PSI(3))/6.
  M(2)= (PSI(1)-PSI(2))/3.
  M(3)= (2.*PSI(1)+PSI(2)-3.*PSI(3))/6.
CONTINUE
RETURN
END
100
~L

```

```

SUMMATION TRANSFORM (L,R,G,P,A)
C
C SOLVES AUTOCORRELATION NORMAL EQUATIONS USING LEVINSON RECURSION.
C PROGRAM FROM BUSHNIN'S BOOK, "MULTICHANNEL TIME SERIES ANALYSIS
C WITH DIGITAL COMPUTER PROGRAMS", PG.44.
C
C INPUTS:
C L= LENGTH OF FILTER P
C R= VECTOR OF AUTOCORRELATION CURVES
C G= VECTOR OF RIGHT-HAND SIDE CURVES IN THE NORMAL EQUATIONS
C OUTPUTS:
C P= VECTOR OF FILTER CURVES, WHICH IS THE SOLUTION TO THE NORMAL EQU
C A= PREDICTION ERROR OPERATOR COEFFICIENTS.
C
C DIMENSION R(1),G(1),P(1),A(1)
DO I =1,3
  A(I)= 1.
  F(I)= G(I)/R(I)
  IF(LM-3*P-1) RETURN
  DO J = L-2,LR
    A(J)= -R/J
    IF(LM-2) GO TO 2
    L1= (L-2)/2
    L2= L1+1
    IF(L2-L1-2) GO TO 5
    DO J1 =J2,L2
      MUL0= A(J)
      A(J)= A(J1)+A(L1)*A(K)
      A(K)= A(K1)+A(L1)*MUL0
      IF(L1-LM-1) GO TO 2
      V= 0+R(L1)*B
      F(L1)= (G(L1)-B)/V
      L3= L1-1
      DO J =J1,L3
        A(J)= F(J)+F(L1)*A(K)
        IF(LM-LR) RETURN
      B=0.
      DO J1 =L1,L
        B= B+A(J1)*R(K)
        M= M+F(L1)*R(K)
      RETURN
    END
  END
~L

```

SUBROUTINE DELTAG(E,MAZE,ETOT,NST,EDRY)

C COMPUTES DELTA GAINS

C INPUTS:

C I VECTOR OF INPUT SAMPLES
C MAZE MAXIMUM DIMENSION UP I
C EDRY TOTAL MEAN SQUARED FRAME ENERGY
C NST INDEX OF THE START SAMPLE IN I FROM
C WHICH THE SECTION ENERGIES ARE TO BE COMPUTED.

C OUTPUTS:

C ADVA VECTOR OF DELTA GAINS
C DIMENSION I(MAZE),EDRY(I)
C CUMMUB /NSSEC/MSSEC,MS#SEC

C I1=NST
C DO 10 I=1,NESEC
C N2=N1+MS#SEC-1
C EDRY(I)=ENERGY(I,MAZE,I1,N2)
C EDRY(I)=EDRY(I)/ETOT
C N1=N2+1
C CONTINUE
C RETURN
C END

10

7

SUBROUTINE OSCALE(G,STOT,NST,EDRY)

C COMPUTES THE QUANTIZER SCALE FACTORS

C INPUTS:

C STUT TOTAL MEAN SQUARE FRAME ENERGY
C CDEP VECTOR OF DELTA GAINS
C NST INDEX OF THE START LOCATION IN THE
C ARRAY G WHERE THE SCALE FACTORS ARE STORED.

C OUTPUTS:

C G ARRAY OF QUANTIZER SCALE FACTORS
C NOTE THAT THE ARGUMENT OF G IS THE SAMPLE NUMBER,
C I.E. THE SCALE FACTOR FOR INDIVIDUAL SAMPLES ARE,
C IN GENERAL, BUT THE SAME.

C DIMENSION G(I),EDRY(I)

C CUMMUB /NSSEC/MSSEC,MS#SEC

C I1=NST

C DO 20 I=1,NESEC

C N2=N1+MS#SEC-1

C E=ETOT*EDRY(I)

C DO 10 J=N1,N2

C G(J)=E

C CONTINUE

C N1=N2+1

C CONTINUE

C RETURN

C END

10

20

7

```

SUBROUTINE PZMS(A,AMS,PZPAC)
C COMPUTES COEFFICIENTS OF THE NUMERATOR
C OF THE NOISE SHAPING FILTER.
C INPUTS:
C   A      SPECTRAL PREDICTOR COEFFS (MPOLE IN NUMBER)
C   PZPAC  CONSTANT
C OUTPUT:
C   AMS    NUMERATOR COEFFS OF NOISE SHAPING FILTER
C   DIMENSION A(1),AMS(1)
C   COMMON/MPOLE/MPOLE
C   BWPACT=1.0
C   DO 2154 I=1,MPOLE
C     BWPACT= BWPACT*PZPAC
C     AMS(I)= A(I)*BWPACT
C   RETURN
C   END
2154

```

-L

```

SUBROUTINE STOLZE(N,NTAPS,NO)
C CHECK FOR STABILITY OF THE 3-TAP PITCH FILTER
C AND REPLACE IT BY 1-TAP FILTER WHENEVER INSTABILITY
C IS OCCURRED.
C INPUTS:
C   N      PITCH PREDICTOR TAPS
C   NO     (REPLACED BY 1-TAP FILTER COEFFS
C   IF INSTABILITY OCCURS, OTHERWISE
C   UNCHANGED AND RETURNED TO CALLING PROGRAM)
C   NTAPS  NUMBER OF PITCH PREDICTOR TAPS
C   NO     PITCH PREDICTOR TAP FOR 1-TAP CASE
C   DIMENSION O(1)
C   COMMON/TSCN/TSCN(1)
C   DO 10 I=1,NTAPS
C     TSCN(I)=O(I)
C   CALL LTRANS(TSCN,1)
C THE STABILITY - CHECK ANGLES THAT FOLLOW
C HAVE BEEN EMPIRICALLY OBTAINED
20  IF(TSCN(1)-17.-1.0) GO TO 25
   IF(TSCN(2)-67.7.0) GO TO 25
   IF(TSCN(3)-67.1.0) GO TO 25
C NO INSTABILITY
   RETURN
C INSTABILITY DECLARED ; REPLACE WITH 1-TAP FILTER
25  NTAPS=(NTAPS-1)/2
   CALL SATEN(N,NTAPS)
   RETURN
   END

```

-L

```

SUMMATION APC(S,GSEC,ANS,DA,DR,MPTCMB,N2S,WD,IB,VH,HW,4,QI
,1,NTAPS,IO,EPFP,ANSF)
C APC LOOP WITH POLE-ZERO NOISE SHAPING. USES P-S
C PREDICTION SCHEME AND PREDICTIVE FORWARD. PITCH
C PREDICTION AND/OR NOISE SHAPING CAN BE DELETED FROM
C THE LOOP BY ZEROING THE FLAGS DESCRIBED BELOW.
C INPUTS:
C S = INPUT SAMPLES (START AT N2S)
C GSEC = QUANTIZER SCALE FACTOR ARRAY (START AT N2S)
C ANS = FILTER COEFFICIENTS USED TO IMPLEMENT THE
C NUMERATOR (ZEROS) OF THE NOISE SHAPING FILTER
C (MPLE TOTAL).
C DA = SPECTRAL PREDICTOR COEFFICIENTS (MPOLE TOTAL)
C NR = PITCH PREDICTOR COEFFICIENTS (NTPOLE TOTAL)
C MPTCMB = PITCH LAG IN # OF SAMPLES
C EPFP = PITCH PREDICTION FLAG (TRUE=ON,FALSE=OFF)
C ANSF = NOISE SHAPING FLAG (TRUE=ON,FALSE=OFF)
C IQ = FLAG FOR RESIDUAL SAMPLE QUANTIZATION
C (TRUE(-1)=ON, FALSE(0)= OFF)
C OUTPUTS:
C IR = CIRCLED RESIDUAL SAMPLES
C WD,WH,HW,Q,I-INTERMEDIATE ARRAYS OF APC LOOP
C (ALL SAME SIZE AS S)
C SEE REAL TIME SPECIFICATIONS
C OF 800 1600/S APC CODER FOR
C DESCRIPTION OF ARRAY FUNCTIONS
C DIMENSION S(I),GSEC(I),ANS(I),DA(I),NR(I),IR(I)
C DIMENSION WD(I),WH(I),HW(I),Q(I),I(I),Q(I)
COMMON/FRNSIZ/MFRAME,NSIMO
COMMON/MPOLE/MPOLE
COMMON/RESORB/MLERR,MBRES(9),YRES(8)
II=N2S-1
C TOP OF APC LOOP
DO 2600 N=1,MFRAME
II=II+1
C NOISE SHAPING
UP=0.
WPI=0.
IF(.NOT.ANSF) GO TO 2630
CALL PRED(OL,JP,IT,ANS,MPOLE,0)
CALL PRED(OL,DP,IT,DA,MPOLE,0)
CONTINUE
C SPECTRAL FILTER
CALL PRED(VH,VHP,II,DA,MPOLE,0)
MFP=0.
IF(.NOT.EPFP) GO TO 2650
CALL PRED(RH,RHW,II,DR,NTAPS,MPTCMB)

```

```

2650 CONTINUE
C COMPUTE WD
WD=S(II)+QP-QPI+RMP+VHP
C ADAPTIVE QUANTIZER
WQ(II)=WQ/GSEC(II)
C QUANTIZE
IF(.NOT.IQ) GO TO 2660
CALL CODI(WQ(II),MBRES,MLERR,IR(N))
CALL DCOW(IR(N),YRES,WH)
GO TO 2665
C NO CODING
2660 WH=WQ(II)
2665 CONTINUE
C
C WH=WH*GSEC(II)
C COMPUTE Q AND QI
Q(II)=WQ-WH
QI(II)=Q(II)-QPI
C COMPUTE VH
VH(II)=WH-VHP
C COMPUTE RH
RH(II)=VH(II)-RHP
C END OF LOOP
2600 CONTINUE
RETURN
END

```

```

SUBROUTINE HFCOR(RSCR,MPMFC,CMF,MFLAN)
C APPLIES HIGH FREQUENCY CORRECTION TO AUTOCORRELATION
C COEFFICIENTS.
C ARGUMENTS:
C   RSCR- VECTOR CONTAINING INPUT AUTOCORRELATION
C         COEFFICIENTS. (ALSO CONTAINS HIGH FREQUENCY CORRECTED
C         OUTPUT AUTOCORRELATIONS)
C   MPMFC- NUMBER OF PULSES USED FOR A/MIN COMPUTATION.
C   CMF- AUTOCORRELATION COEFFICIENTS FOR HIGHPASS
C         FILTERED WHITE NOISE.
C   MFLAN- CONSTANT (.01 < MFLAN < .1). THAT DETERMINES
C         THE EXTENT OF CORRECTION.
C
C   DIMENSION RSCR(1),CMF(1)
C   DIMENSION SCR1(10),SCR2(10),SCR3(10)
C
C   CALL ALPCOR(MPMFC,1,RSCR,SCR1,SCR2,V,SCR3)
C   ENR=RSCR(1)*V
C   DO 10 I=1,3
C     RSCR(I)=RSCR(I)+MFLAN*ENR*CMF(I)
C   RETURN
C   END

```

