

Copyright Warning & Restrictions

The copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyrighted material.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be “used for any purpose other than private study, scholarship, or research.” If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of “fair use” that user may be liable for copyright infringement,

This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve violation of copyright law.

Please Note: The author retains the copyright while the New Jersey Institute of Technology reserves the right to distribute this thesis or dissertation

Printing note: If you do not wish to print this page, then select “Pages from: first page # to: last page #” on the print dialog screen

The Van Houten library has removed some of the personal information and all signatures from the approval page and biographical sketches of theses and dissertations in order to protect the identity of NJIT graduates and faculty.

S I G N A L P R O C E S S I N G T E C H N I Q U E S
F O R A N A L Y S I S O F H E A R T S O U N D S
A N D E L E C T R O C A R D I O G R A M S

by

Alan D. Bernstein

A DISSERTATION
PRESENTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF ENGINEERING SCIENCE IN ELECTRICAL ENGINEERING
AT NEW JERSEY INSTITUTE OF TECHNOLOGY

This dissertation is to be used only with due regard to the rights of the author. Bibliographical references may be noted, but passages must not be copied without permission of the Institute and of the author, and without credit being given in subsequent written or published work.

Newark, New Jersey
1979

A B S T R A C T

Audible heart sounds represent less than 5% of the vibrational energy associated with the cardiac cycle. In this study, experiments have been conducted to explore the feasibility of examining cardiac vibration by means of a single display encompassing the entire bandwidth of the oscillations and relating components at different frequencies. Zero-phase-shift digital filtering is shown to be required in producing such displays, which extend from a recognizable phonocardiogram at one frequency extreme to a recognizable apexcardiogram at the other. Certain features in mid-systole and early diastole, observed by means of this technique, appear not to have been previously described.

Frequency modulation of an audio-frequency sinusoid by a complex signal is shown to be effective in generating sounds analogous to that signal and containing the same information, but occupying a bandwidth suitable to optimum human auditory perception. The generation of such sounds using an exponential-response voltage-controlled oscillator is found to be most appropriate for converting amplitude as well as frequency changes in the original signal into pitch changes in the new sounds,

utilizing the human auditory system's more acute discrimination of pitch changes than amplitude changes. Pseudologarithmic compression of the input signal is shown to facilitate emphasis in the converted sounds upon changes at high or low amplitudes in the original signal. A noise-control circuit has been implemented for amplitude modulation of the converted signal to de-emphasize sounds arising from portions of the input signal below a chosen amplitude threshold. This method is shown to facilitate the transmission of analogs of audible and normally inaudible sounds over standard telephone channels, and to permit the "slowing down" of the converted sounds with no loss of information due to decreased frequencies.

The approximation of an arbitrary waveform by a piecewise-linear (PL) function is shown to permit economical digital storage in parametric form. Fourier series and Fourier transforms may be readily calculated directly from the PL breakpoint parameters without further approximation, and the number of breakpoints needed to define the PL approximation is significantly lower than the number of uniformly-spaced samples required to satisfy the Nyquist sampling criterion; aliasing problems are shown not to arise. Thus data compression is fea-

sible by this means without recourse to a parametric model defined for the signal (e.g., speech) being processed. Methods of automatic adaptive PL sampling and waveform reconstruction are discussed, and microcomputer algorithms implemented for this purpose are described in detail. Examples are given of the application of PL techniques to electrocardiography, phonocardiography, and the digitization of speech.

APPROVAL OF DISSERTATION

SIGNAL PROCESSING TECHNIQUES

FOR ANALYSIS OF HEART SOUNDS AND ELECTROCARDIOGRAMS

by

ALAN D. BERNSTEIN

for

DEPARTMENT OF ELECTRICAL ENGINEERING

NEW JERSEY INSTITUTE OF TECHNOLOGY

BY FACULTY COMMITTEE

APPROVED:

Chairman

Newark, New Jersey

May, 1979

P R E F A C E

This is essentially a feasibility study -- not so much a report of results as an exploration of possible ways to get them. Although other applications are discussed, my primary goal has been the development of techniques to be used in non-invasive cardiac diagnosis and patient monitoring.

Since the techniques we shall be discussing are new and are intended for application in a field other than my own, this work seems to have raised more questions than it has answered and must be considered open-ended. Since I have tried to clarify some of the limitations and implications inherent both in the new techniques and in more conventional methods, the dissertation is a bit long. This is very much an experimental project, and since most of the library research connected with it has perforce consisted of looking for things which were not there in order to confirm their probable nonexistence, there are relatively few references. A few medical terms are used in the text at points where they seemed important; each is briefly defined or described when it first appears.

This work, which began in the spring of 1975, has been supported in part by a research assistantship granted by the Foundation at NJIT during the 1977-1978 and 1978-1979 academic years. It was conducted, for the most part, on the Newark campus of Rutgers University. I am very much indebted to Professor James R. Freeman, Chairman of the Department of Physics at the Newark College of Arts and Sciences of Rutgers University for allowing me unrestricted use of the experimental and computational resources of the department, thus furnishing, at one stroke, most of the instrumentation needed for my experiments as well as a place to conduct them.

It is a pleasure to express my gratitude to my dissertation advisor, Professor Mauro Zambuto, for trusting me enough to permit me to work independently on ideas of my own (and to learn from my own mistakes), for his willingness to explore areas as new to him as they were to me, and for his readiness to place the benefit of his remarkable background at my disposal and to offer sound advice whenever it was needed.

I am indebted to Professor Jacob Klapper for his help in acting as interim advisor during Professor

Zambuto's absence; to Emanuel Goldberg, MD, of the Beth Israel Medical Center in Manhattan, for his patience during hours of enlightening conversation and his kindness in allowing me to perform experiments in the cardiopulmonary laboratory; to Thomas Edery, who assisted in some of the experiments and helped me to learn something about the technical aspects of conventional phonocardiography; to Morton Epstein, MD, for permitting me to make apexcardiogram recordings in the coronary care center of the United Hospitals of Newark, and to Larry Kondra, MD, who made the recordings with me. I have been fortunate in receiving much good advice, from Professor Charles Pine of Rutgers, whose encouragement has been unflagging and whose colleague I feel proud to be; from Professor Benjamin Carroll of Rutgers, with whose guidance I first became involved in science and the teaching of science, and his former graduate student, Harold Tillel; from Ellis D. Cooper, who is doing fascinating work with PL techniques in the generation of electronic music; and from Ardath Day, Robert Gerdes, and Charles Marshall of the Center for Computer Information and Services at Rutgers in Newark, who have helped me out of innumerable computer-related difficulties. To all I offer my most sincere thanks.

It would be foolish to pretend that this work represents the culmination of anything. I shall be very much gratified if it can serve as the beginning of something.

Alan D. Bernstein

Newark, May 1979

Leave me, O Love, which reachest but to dust;
And thou, my mind, aspire to higher things....

-- Sir Philip Sidney

Humbug!

-- Charles Dickens

One of the satisfactions of independent research seems to be its very independence: the chance to decide for oneself on the basis of the evidence -- to conclude, for example, that there is nothing strange, after all, in offering a work of scholarship as an act of love.

-- Alan D. Bernstein

THIS IS FOR RAMA...

WITH LOVE.

T A B L E O F C O N T E N T S

Chapter	Page
ABSTRACT	ii
PREFACE	vi
LIST OF FIGURES	xiv
LIST OF TABLES	xxiv
1 -- INTRODUCTION	1
2 -- WIDEBAND APEXCARDIOGRAPHY	6
2.1 The Conventional Apexcardiogram	9
2.2 Filtering of Apexcardiograms	9
2.3 Zero-Phase-Shift Digital Filtering of Apexcardiograms	19
2.4 Processing of Supine Apexcardio- grams	27
3 -- CONVERSION OF HEART SOUNDS BY FREQUENCY MODULATION	36
3.1 Introduction	36
3.2 Frequency Bandwidth of Heart Sounds	37
3.3 Some Aspects of the Physics of Auscultation	42
3.4 Bandwidth Criteria for Heart-Sound Conversion	53
3.5 Possible Means of Heart-Sound Conversion	55
3.6 Conversion of Heart Sounds by Linear Frequency Modulation	57
3.7 Conversion by Exponential Frequency Modulation	69
3.8 Pseudologarithmic Compression in Exponential FM Conversion	78
3.9 Effect of Pseudologarithmic Conversion on Exponential FM-Conversion Spectra	87
3.10 Pseudologarithmic Compression in System Output Amplitude Control	91
3.11 Special Applications of FM Conversion	101
3.12 Low-Frequency Considerations in FM Conversion	107
4 -- REPRESENTATION OF ANALOG SIGNALS BY PIECEWISE-LINEAR APPROXIMATIONS	110

4.1	Piecewise-Linear Waveforms	110
4.2	The Fourier Transform of a Time-Limited PL Function	111
4.3	The Fourier Series for a Periodic PL Function	117
4.4	Sampling-Rate Considerations in Adaptive PL Sampling	127
5	-- ADAPTIVE PIECEWISE-LINEAR SAMPLING OF ANALOG SIGNALS	149
5.1	Hybrid Implementation of Adaptive PL Sampling	149
5.2	Analog Differentiation	154
5.3	Phase and Gain Error for the Analog Differentiator	157
5.4	Modification of the Analog Circuit	161
5.5	A Microcomputer Algorithm for Adaptive PL Sampling	164
5.6	Slope Detection in PL Sampling	175
5.7	Modifications to the Algorithm for Adaptive PL Sampling	186
6	-- APPLICATIONS OF PL SAMPLING	199
6.1	Reconstruction of PL Signals	199
6.2	PL Sampling of Continuous Speech	217
6.3	PL Sampling of Normal Speech	222
6.4	PL Sampling of Electrocardiograms	225
6.5	PL Sampling of Heart Sounds	260
7	-- CONCLUSIONS AND RECOMMENDATIONS FOR FURTHER INVESTIGATION	269
7.1	Wideband Apexcardiography	269
7.2	Conversion of Heart Sounds by Frequency Modulation	272
7.3	Applications of PL Sampling	275
APPENDIX A	-- COMPUTER PROGRAMS DEVELOPED FOR USE IN THE RESEARCH	281
APPENDIX B	-- FM CONVERSION OF HEART SOUNDS: CIRCUIT DESCRIPTION OF THE EXPERIMENTAL SYSTEM	299
B.1	Introduction	299
B.2	Input Amplifier	301
B.3	Overdrive Indicator	302

B.4	Pseudologarithmic Compression Amplifier	302
B.5	Center Frequency and Frequency Deviation Control Circuit	303
B.6	Artificial Frequency-Foldover Circuit	305
B.7	VCO Control-Voltage Output Circuit	305
B.8	Amplitude-Modulation Control Circuit	306
B.9	Noise Threshold Control Circuit	307
B.10	Adequate-Gain Indicator	309
B.11	Voltage-Controlled Amplifier	310
B.12	The Portable FM Conversion System	311
APPENDIX C	-- OPERATING MANUAL FOR THE FM CONVERSION SYSTEM	325
APPENDIX D	-- MEASUREMENT OF THE ACTIVITY COEFFICIENT FOR NORMAL SPEECH	338
APPENDIX E	-- DESCRIPTION OF THE RECORDED EXAMPLES	342
E.1	Conversion of Heart Sounds by Frequency Modulation	342
E.2	Adaptive PL Sampling of Speech	343
REFERENCES		347
CURRICULUM VITAE		352

L I S T O F F I G U R E S

Figure		Page
1.1	Vibratory spectrum of the heart in relation to the threshold of audibility (after Butterworth et al 1960, p. 25).	2
2.1.1	Apexcardiogram (a) derived by zero-phase-shift digital filtering with a 3-dB corner frequency of 20 Hz from a wideband apexcardiogram (b) obtained from a normal subject in the left lateral decubitus position.	8
2.2.1	Phonocardiogram (a) and apexcardiogram (b) derived by zero-phase-shift digital filtering from a wideband apexcardiogram obtained from a normal subject in the left lateral decubitus position.	10
2.2.2	Analog high-pass filtering of a wideband apexcardiogram obtained from a normal subject in the left lateral decubitus position. Each major division represents 0.1 sec.	12
2.3.1	Zero-phase-shift digital filtering of a wideband apexcardiogram obtained from a normal subject in the left lateral decubitus position.	25
2.4.1	Apexcardiograms obtained from a normal subject (a) in the left lateral decubitus position, and (b) supine.	29
2.4.2	Zero-phase-shift digital high-pass filtering of a wideband apexcardiogram obtained from a normal subject in the supine position.	30
2.4.3	Apexcardiogram (a) derived by zero-phase-shift digital low-pass filtering with a 3-dB corner frequency of 20 Hz from a wideband apexcardiogram (b) obtained from a normal subject in the supine position.	32

2.4.4	Zero-phase-shift digital low-pass filtering, at successively lower frequencies, of a wideband apexcardiogram obtained from a normal subject in the supine position.	33
3.2.1	Spectrum of normal heart sounds, averaged over one cardiac cycle chosen at random from a commercial recording.	39
3.2.2	Spectrum of the author's heart sounds, recorded at the cardiac apex using a stethoscope head with diaphragm attachment and averaged over one cardiac cycle.	40
3.2.3	Arrangement used in making the recording from which the spectrum of Fig. 3.2.2 was plotted.	41
3.2.4	Spectrum of a musical "seagull" murmur, averaged over one cardiac cycle chosen at random from a commercial recording.	43
3.3.1	Arrangement used in making the recording from which the spectrum of Fig. 3.3.2 was plotted.	48
3.3.2	Spectrum of the author's heart sounds, recorded at the cardiac apex using an entire stethoscope with diaphragm attachment and averaged over one cardiac cycle.	49
3.3.3	Frequency response of a stethoscope head measured under free-field conditions.	50
3.3.4	Frequency response of an entire stethoscope measured under free-field conditions.	51
3.3.5	Arrangement used in plotting the frequency response of a stethoscope head and of an entire stethoscope, as seen in Figs. 3.3.3 and 3.3.4.	52
3.6.1	500-Hz carrier used in generating all of the FM spectra shown in this chapter.	62
3.6.2	100-Hz sinusoid used as a modulating signal to obtain the linear-FM spectrum shown in Fig. 3.6.3.	63

3.6.3	Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.2, using a linear VCO (500 Hz/V).	64
3.6.4	100-Hz sinusoid used as a modulating signal to obtain the linear-FM spectrum shown in Fig. 3.6.5.	65
3.6.5	Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.4, using a linear VCO (500 Hz/V).	66
3.6.6	Modulating signal derived from that of Fig. 3.6.4 by means of artificial frequency foldover (in this case, full-wave rectification).	68
3.6.7	Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.6, using a linear VCO (500 Hz/V).	70
3.7.1	Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.	77
3.8.1	Approximate pitch-variation sensitivity of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.	79
3.8.2	Transfer function of the pseudologarithmic compression amplifier.	80
3.8.3	Plot of an analytic approximation of the transfer function of Fig. 3.8.2 for positive voltages as calculated using Eq. (3.8.1).	82
3.8.4	Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) using pseudologarithmic compression, plotted as a function of the amplitude of a sinusoidal input voltage.	83
3.8.5	Approximate pitch-variation sensitivity of an exponential VCO (6V/decade) using pseudologarithmic compression, plotted as a function of the amplitude of a sinusoidal input voltage.	84

3.8.6	Means of continuous selection of the degree of pseudologarithmic compression used in generating the control signal for the VCO.	86
3.9.1	100-Hz sinusoid of amplitude 1V, used as an input to obtain the modulating signal shown in Fig. 3.9.2.	88
3.9.2	100-Hz signal obtained from the sinusoid of Fig. 3.9.1 by pseudologarithmic compression and scaled to the same amplitude (1V).	89
3.9.3	Spectrum of the 100-Hz VCO modulating signal shown in Fig. 3.9.2.	90
3.9.4	An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 1V.	92
3.9.5	An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 2V.	93
3.9.6	An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 3V.	94
3.9.7	An exponential FM spectrum. The input is a 1-V signal derived by pseudologarithmic compression from a 100-Hz sinusoid.	95
3.9.8	An exponential FM spectrum. The input is a 2-V signal derived by pseudologarithmic compression from a 100-Hz sinusoid.	96
3.9.9	An exponential FM spectrum. The input is a 3-V signal derived by pseudologarithmic compression from a 100-Hz sinusoid.	97
3.10.1	Spectrum of a 500-Hz sinusoidal carrier modulated with a 100-Hz sinusoid of 6V amplitude, using an exponential VCO (6V/decade), after pseudologarithmic compression of the modulating signal.	99
3.10.2	Spectrum of the same modulated carrier used to produce Fig. 3.10.1, except that the frequency-modulated carrier has in this instance been subsequently amplitude-modulated by the output of the pseudologarithmic amplifier.	100

3.10.3	Transfer function of the noise-threshold control circuit.	102
3.11.1	Third-octave-band spectrum showing the bandwidth of normal heart sounds processed by pseudologarithmic compression and converted to fit into a telephone bandwidth of 500-5000 Hz.	105
4.1.1	Notation used in discussing piecewise-linear (PL) waveforms.	112
4.2.1	Successive differentiation of a time-limited PL waveform, showing notation used in deriving expressions for the Fourier transform in terms of the breakpoint parameters.	114
4.3.1	Successive differentiation of a periodic PL waveform, showing notation used in deriving expressions for the Fourier-series coefficients in terms of the breakpoint parameters.	118
4.3.2	(a) Three cycles of a normal electrocardiogram; (b) PL approximation of the waveform shown in (a).	123
4.3.3	Computer sketch of one cycle of a periodic PL waveform.	125
4.3.4	Amplitude spectrum of a periodic PL waveform.	128
4.3.5	Amplitude spectrum of a periodic PL waveform, scaled linearly.	129
4.3.6	Phase spectrum of a periodic PL waveform.	130
4.4.1	A triangular pulse defined by three PL breakpoints.	138
4.4.2	(a) Normal electrocardiogram, original waveform. (b) PL approximation, 105 breakpoints/cycle (111.8 breakpoints/sec). (c) Uniform sampling, $f_s = 111.8$ Hz. (d) Uniform sampling, $f_s = 111.8$ Hz, filtered at 55.9 Hz.	140

4.4.3	(a) Normal electrocardiogram, original waveform. (b) PL approximation, 46 breakpoints/cycle (48.9 breakpoints/sec). (c) Uniform sampling, $f_s = 48.9$ Hz. (d) Uniform sampling, $f_s = 48.9$ Hz, filtered at 24.5 Hz.	141
4.4.4	(a) Normal electrocardiogram, original waveform. (b) PL approximation, 31 breakpoints/cycle (33.0 breakpoints/sec). (c) Uniform sampling, $f_s = 33.0$ Hz. (d) Uniform sampling, $f_s = 33.0$ Hz, filtered at 16.5 Hz.	142
4.4.5	Computer sketch of the linear ramp waveform defined by the breakpoints shown in Table 4.4.1.	146
4.4.6	Amplitude spectrum of the linear ramp waveform defined by the breakpoints shown in Table 4.4.1.	147
4.4.7	PL and uniform sampling of the ramp waveform shown in Fig. 4.4.5.	148
5.1.1	A hybrid system for extracting PL-waveform breakpoint parameters by adaptive sampling.	151
5.1.2	Four steps in a process for controlling variable-rate adaptive sampling of an analog signal in order to obtain appropriate breakpoint parameters for a piecewise-linear approximation.	153
5.2.1	An analog circuit for computation of the second derivative of an input waveform.	155
5.2.2	Bode amplitude plot for the circuit of Fig. 5.2.1, whose transfer function is given in Eq. (5.2.5).	158
5.4.1	A three-channel system using the circuit of Fig. 5.2.1 to generate a control voltage proportional to the magnitude of the second derivative of the input signal.	162
5.5.1	Flowchart for adaptive PL sampling algorithm.	176

5.5.2	Flowgram for adaptive PL sampling algorithm.	177
5.6.1	Notation used in discussing slope resolution.	178
5.6.2	An example of the minimum detectable slope as a function of sampling rate for an 8-bit system.	182
5.6.3	A scheme for achieving 16-bit ordinate resolution by means of three conversions.	184
5.7.1	(a) Test waveform: uniform sampling, 264 samples/cycle. (b) Same, PL sampling, 46 samples/cycle. (c) Same, PL sampling, 15 samples/cycle. (d) Same, PL sampling, 11 samples/cycle.	198
5.7.2	Circuit used to generate the test waveform shown in Fig. 5.7.1.	200
6.1.1	A portion of a PL waveform showing the notation used in discussing hybrid interpolation between breakpoints.	202
6.1.2	Combination of high- and low-order bytes of τ_{i+1} using two 8-bit D/A channels.	205
6.1.3	A microcomputer-controlled interpolator for synthesis of PL waveforms.	209
6.1.4	A modification allowing the circuit of Fig. 6.1.3 to be used for PL-waveform reconstruction at a continuously variable rate without recalibration.	212
6.1.5	A scheme for parallel transfer of breakpoint-parameter voltages to an analog interpolator circuit.	216
6.2.1	(a) Spectrum for speech, PL approximation. (b) Equivalent frequency probability density.	219
6.4.1	Normal and abnormal electrocardiograms.	230
6.4.2	Amplitude spectrum for the first cycle of the abnormal waveform shown in Fig. 6.4.1.	233

6.4.3	Amplitude spectrum for the second cycle of the abnormal waveform shown in Fig. 6.4.1.	234
6.4.4	Phase spectrum for the first cycle of the abnormal waveform shown in Fig. 6.4.1.	235
6.4.5	Phase spectrum for the second cycle of the abnormal waveform shown in Fig. 6.4.1.	236
6.4.6	Differential amplitude spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.	237
6.4.7	Differential amplitude spectrum comparing the second cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.	238
6.4.8	Differential amplitude spectrum comparing two consecutive cycles of the abnormal waveform shown in Fig. 6.4.1.	239
6.4.9	Differential phase spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.	240
6.4.10	Differential phase spectrum comparing the second cycle of the abnormal electrocardiogram shown in Fig. 6.4.1 with a normal electrocardiogram.	241
6.4.11	Differential phase spectrum comparing two consecutive cycles of the abnormal electrocardiogram shown in Fig. 6.4.1.	242
6.4.12	Amplitude spectrum of the first cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	243
6.4.13	Amplitude spectrum of the second cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	244
6.4.14	Phase spectrum of the first cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	245

6.4.15	Phase spectrum of the second cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	246
6.4.16	Differential amplitude spectrum comparing two consecutive cycles of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	247
6.4.17	Differential phase spectrum comparing two consecutive cycles of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.	248
6.4.18	Examples of sinus arrhythmia.	258
6.5.1	Normal heart sounds obtained by PL sampling. The first (S1) and second (S2) heart sounds are labelled.	261
6.5.2	Second heart sound from the beginning of the trace shown in Fig. 6.5.1, hanning-windowed for computation of the Fourier transform.	263
6.5.3	Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. Both frequency and amplitude axes are scaled linearly.	264
6.5.4	Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. The frequency axis has been scaled logarithmically to facilitate identification of low-frequency components.	265
6.5.5	Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. Both axes have been scaled logarithmically.	266
B.1.1	Circuit diagram of the experimental FM conversion system. All diodes and transistors are silicon, and all IC's are 1/2 747, except as indicated.	300
3.12.1	Front panel of the portable FM conversion system.	312

B.12.2	Input amplifier for the portable FM conversion system. The maximum voltage gain is 5000 (74 dB).	314
B.12.3	FM center frequency and modulation depth control circuit for the portable FM conversion system.	315
B.12.4	Exponential-response amplifier for the portable FM conversion system.	317
B.12.5	Voltage-controlled oscillator circuit for the portable FM conversion system.	318
B.12.6	Graphic display output circuit for the portable FM conversion system.	319
B.12.7	Operation of the noise threshold circuit.	321
B.12.8	Operation of the pseudologarithmic compressor.	322
B.12.9	Transformer-coupled power amplifier for the portable FM conversion system.	324
D.1	Block diagram of a circuit for measurement of the activity coefficient for speech.	339
E.1	Schematic representation of the abnormal heart sounds recorded on Band 3.	344

L I S T O F T A B L E S

Table		Page
2.2.1	Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole high-pass analog filtering with a 3-dB corner frequency of 40 Hz.	16
2.2.2	Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole high-pass analog filtering with a 3-dB corner frequency of 25 Hz.	17
2.2.3	Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole low-pass analog filtering with a 3-dB corner frequency of 20 Hz.	18
3.7.1	Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.	76
4.3.1	Breakpoint parameters for one cycle of a periodic PL waveform.	124
4.3.2	Harmonic components of a periodic PL waveform.	126
4.4.1	Breakpoint parameters completely defining one cycle of a piecewise-linear ramp waveform having two breakpoints per cycle.	145
5.7.1	Complete Z-80 assembly-language listing of a microcomputer program for adaptive PL sampling, using the algorithm described in Section 5.5 with the modifications given in Section 5.7.	191
6.3.1	Measurements of the activity coefficient for normal speech.	224
6.4.1	Spectrum of the first cycle of the abnormal waveform shown in Fig. 6.4.1.	249

- | | | |
|-------|--|-----|
| 6.4.2 | Spectrum of the second cycle of the abnormal waveform shown in Fig. 6.4.1. | 250 |
| 6.4.3 | Differential spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram. | 251 |
| 6.4.4 | Differential spectrum comparing the second cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram. | 252 |
| 6.4.5 | Differential spectrum comparing two consecutive cycles of the abnormal waveform shown in Fig. 6.4.1. | 253 |
| 6.4.6 | Spectrum of the first cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1. | 254 |
| 6.4.7 | Spectrum of the second cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1. | 255 |
| 6.4.8 | Differential spectrum comparing two consecutive cycles of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1. | 256 |

C H A P T E R 1

INTRODUCTION

Most of the techniques discussed in the following pages were developed primarily for use in connection with the detection, storage, and interpretation of mechanical oscillations arising from the action of the heart. Although some of these techniques may be applicable in other areas (we shall offer examples), we feel that it is appropriate to begin by pointing out some of the characteristics of cardiac-cycle vibration which have motivated our experiments.

Fig. 1.1 shows the vibration spectrum of the heart in relation to the average threshold of human auditory perception. In their caption to the original diagram, contained in a textbook on auscultation of the heart, the authors point out that since the threshold of audibility may vary considerably from one individual to another, differences of opinion occur frequently in the auscultation of heart sounds near the threshold. Certainly one of our objectives is to explore the possibility of implementing techniques for obtaining reproducible quantitative information which will be unaffected by variations in human perception. There is nothing new about




FIGURE 1.1

Vibratory spectrum of the heart in relation to the threshold of audibility (after Butterworth et al 1960, p. 25).

this; however, we find it somewhat disturbing that manufacturers of instrumentation for producing phonocardiograms -- graphic plots of (mainly audible) cardiac vibration as a function of time -- appear to have been quite so assiduous in building into their instruments the limitations of a human observer with respect to frequency response and amplitude discrimination. We have been told that their motivation in so doing is to provide a permanent graphic record of what the physician can hear, an objective whose value is not difficult to grasp. What we find alarming is that we have also been assured that only the audible portion of the heart sounds is regarded as being of clinical significance.

The ordinate axis of the diagram in Fig. 1.1 has dimensions of pressure; but

$$\frac{\text{dynes}}{\text{cm}^2} = \frac{\text{dynes} \cdot \text{cm}}{\text{cm}^3} = \frac{\text{erg}}{\text{cm}^3},$$

so that it may also be regarded as a representation of energy density. Then area on this diagram has units of

$$\left(\frac{\text{erg}}{\text{cm}^3} \right) (\text{Hz}) = \frac{\text{erg} / \text{sec}}{\text{cm}^3},$$

which represent volume power density. If we assume that the shaded portion of the diagram may be extended to the

left by an additional octave to 4 Hz (an assumption whose validity is borne out by our observations; see Fig. 2.3.1, for example), then the portion of the shaded area above the threshold of audibility represents less than 5% of the total shaded area. In other words, of the total energy associated with cardiac vibration, on the order of 95% is inaudible (this is a rough estimate, owing to the variability of the threshold and the fact that the abscissa is scaled logarithmically).

We find it difficult to accept the coincidence that the audible 5% should contain all of the vibratory information of clinical significance, in spite of the fact that many hospitals, even hospitals containing well-equipped coronary care units, lack facilities for studying the remaining 95%. We cannot help suspecting that it is more accurate to say that the clinical significance of the audible 5% is better understood, since the means of detecting it have been available for a considerably longer period of time, and its relationship to the function and condition of the cardiovascular system has been more extensively studied.

In conventional phonocardiography, a distinction is made between the range of audio and subaudio vibrations

(usually about 25 Hz and above) and the lowest frequencies (about 1 - 20 Hz) associated with motion of the chest wall above the cardiac apex. In Chapter 2 we will consider some techniques related to the measurement of heart sounds as a whole to produce graphic displays in the time domain. In Chapter 3 we will show how any audio or subaudio signal may be converted in real time to analogous sounds better suited to the amplitude- and frequency-response characteristics of human hearing.

In Chapters 4 and 5 we will describe a technique for obtaining a time-domain approximation to an arbitrary analog signal in a form which is well-suited to digital storage and particularly convenient for conversion to the frequency domain. In Chapter 6 we will discuss the application of this technique to speech, electrocardiography, and heart sounds.

C H A P T E R 2
WIDEBAND APEXCARDIOGRAPHY

2.1 -- The Conventional Apexcardiogram

The term "apexcardiography" refers to the measurement of precordial (chest wall) displacement as a function of time. It is something of a misnomer since the technique is not restricted to observation at the cardiac apex (Craigie 1974). Precordial motion has been of clinical interest for well over a century, but the popularity of apexcardiography has increased in recent years with the greater availability of transducers having adequate low-frequency response.

Considerable importance is attached to the time constant of the transducer, partly because a limited low-frequency response has the effect of differentiating the waveform and producing potentially misleading time shifts in the resulting tracing of components at different frequencies. This matter has been investigated in some detail (Kesteloot et al 1969). A time constant of 3 sec or more is considered most satisfactory, corresponding to a 3-dB corner frequency of 0.053 Hz. Accordingly, our measurements have been made using a piezoelec-

tric transducer with a (verified) low-frequency response to below 0.02 Hz, which corresponds to a time constant of about 8 sec; the high-frequency response extends to about 2 kHz.

In principle, the frequency content of an apex-cardiographic measurement should be expected to extend from near dc (reflecting motion associated with respiration) through the range of heart sounds audible at the location of the transducer. In practice, the transducer is often attached to the subject so as to limit its sensitivity to respiratory motion. Our transducer, for example, measures displacement with respect to a surrounding retaining ring of about 5 cm diameter. Even so, a wide spectrum is involved in the measurements.

We were therefore surprised to learn that low-pass filtering at 20 Hz was commonly employed in recording apexcardiograms. This fact is less surprising when one observes that most of the energy involved is associated with the lowest frequencies. Apexcardiogram spectra (Kesteloot et al 1969) show amplitudes reduced by about 24 dB in the vicinity of 10 Hz from those near the frequency of the cardiac cycle itself. Fig. 2.1.1 shows




FIGURE 2.1.1

Apexcardiogram (a) derived by zero-phase-shift digital filtering with a 3-dB corner frequency of 20 Hz from a wideband apexcardiogram (b) obtained from a normal subject in the left lateral decubitus position.

that the effect of such filtering is almost imperceptible because of this wide disparity in amplitudes.

2.2 -- Filtering of Apexcardiograms

A single transducer may be used, together with appropriate filters, to obtain simultaneous apex- and phonocardiograms; an example is shown in Fig. 2.2.1. Yet when apexcardiograms are discussed in terms of the mechanisms by which the low-frequency deflections are produced, reference to the high frequencies seems to be restricted largely to the temporal relationships between the low-frequency features and the audible sounds. It seems plausible, however, that something might be gained by examining cardiac vibration as a whole, using a single display encompassing all of the frequencies involved.

One possible way of doing this is to use a spectrograph (such as the Kay "Sonograph") which displays frequency on one axis, time on a second, and amplitude by means of varying degrees of blackness. Equivalently, one might build up a spectral contour by combining separate traces obtained by filtering a signal into many contiguous passbands.




FIGURE 2.2.1

Phonocardiogram (a) and apexcardiogram (b) derived by zero-phase-shift digital filtering from a wideband apexcardiogram obtained from a normal subject in the left lateral decubitus position. The upper trace was obtained by high-pass filtering with a corner frequency of 25 Hz, while the lower was obtained by low-pass filtering with a 3-dB corner frequency of 20 Hz. the first (S1) and second (S2) heart sounds are labelled.

In order to try to clarify the relationship (if any exists) between the audible heart sounds and the apexcardiogram, as exemplified in Fig. 2.2.1, we have chosen an approach involving a display in which the apexcardiogram and audible sounds appear in easily recognizable form. This may be achieved, for example, by using analog high-pass filtering to observe the changes in the waveform as successively more low-frequency components are removed. An example is shown in Fig. 2.2.2. The lowest frequency at which we were able to filter in this instance was 2 Hz, yet these measurements seem to suggest a kind of "evolution" into low-frequency displacements of events beginning earlier at high frequencies. Of particular interest is the interval between the first (S1) and second (S2) heart sounds, during the main contractile portion of the cardiac cycle (note the dashed lines in Fig. 2.2.2). Unfortunately, the hardware limitations which made it impossible to select corner frequencies below 2 Hz prevent us from seeing the complete "evolution" into the apexcardiogram itself, which is shown in the second trace from the bottom. The bottom trace is an electrocardiogram, which was included for timing purposes.




FIGURE 2.2.2

Analog high-pass filtering of a wideband apexcardiogram obtained from a normal subject in the left lateral decubitus position. Each major division represents 0.1 sec. Dashed lines have been added to emphasize "migration" of peaks as discussed in the text.

Analog filtering of apexcardiograms must be approached with a certain amount of caution for the same reason that makes the transducer time constant important. For example, it is easy to show that a single-pole low-pass filter with a 3-dB corner frequency of f_0 produces a phase shift at a frequency f of (Van Valkenburg 1964, 344-345)

$$\varphi = \tan^{-1} \frac{f}{f_0} = \tan^{-1} \frac{\omega}{\omega_0},$$

where ω_0 and ω represent the corresponding angular frequencies. Then

$$T = \frac{\varphi}{\omega} = - \frac{1}{2\pi f} \tan^{-1} \frac{f}{f_0},$$

where T is the corresponding shift along the time axis at this frequency. For instance, if $f = 1$ Hz, we find that the time shift due to a low-pass filter of corner frequency 20 Hz will be -7.95 msec per pole; this corresponds to less than 1% of the cardiac cycle at a pulse rate of 70 beats per minute and may presumably be neglected, even if a two-pole filter is used as in some standard apexcardiography equipment, in which case the value is 1.86%. It should be emphasized that φ and T represent, respectively, relative phase and relative

temporal relationships between signal components at different frequencies. Since the analog filter is a causal system, there is no question of negative time delay, which would imply prediction of the future.

The phase shift for a single-pole high-pass filter is

$$\varphi = \frac{\pi}{2} - \tan^{-1} \frac{f}{f_0} ,$$

whence the time shift per pole is given by

$$T = \frac{1}{2\pi f} \left(\frac{\pi}{2} - \tan^{-1} \frac{f}{f_0} \right) .$$

The worst case for high-pass filtering at a corner frequency of 40 Hz is at the corner frequency itself; the time shift is found to be 3.13 msec per pole. This corresponds to less than 0.4% of the cardiac cycle at a pulse rate of 70 beats per minute and is certainly negligible, even if a two-pole filter is used as in some standard phonocardiography equipment, in which case the value is 0.73%.

To confirm the relative safety of the conventional process of high-pass filtering to obtain audible-

frequency phonocardiograms, Table 2.2.1 shows the time shifts corresponding to 2-, 4-, and 6-pole filters, together with the equivalent percent of a 70-beat-per-minute cardiac cycle. Similar results for the case of high-pass filtering with a 25-Hz corner frequency are shown in Table 2.2.2. The corner frequency is again the worst case, but the shift at this frequency is fairly small and will be mitigated somewhat if low-pass filtering (at a higher corner frequency) is employed at the same time.

To verify the relative safety of the conventional process of low-pass filtering to obtain apexcardiograms, Table 2.2.3 shows the time shifts corresponding to 2-, 4-, and 6-pole filters, together with the equivalent percent of a 70-beat-per-minute cardiac cycle.

However, for a one-pole high-pass filter of corner frequency 2 Hz, we find

$$f = 2 \text{ Hz} \Rightarrow T = 62.5 \text{ msec} ;$$

$$f = 34 \text{ Hz} \Rightarrow T = 0.275 \text{ msec}.$$

TABLE 2.2.1

Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole high-pass analog filtering with a 3-dB corner frequency of 40 Hz. Frequencies are in Hz and time-shift magnitudes in seconds.

```

*****
FREQUENCY      T(N=2)      T(N=4)      T(N=6)      Z(N=2)      Z(N=4)      Z(N=6)
*****
  40.00        0.00625    0.01250    0.01875          0.729      1.458      2.187
  60.00        0.00312    0.00624    0.00936          0.364      0.728      1.092
  80.00        0.00184    0.00369    0.00553          0.215      0.430      0.646
 100.00        0.00121    0.00242    0.00363          0.141      0.283      0.424
 120.00        0.00085    0.00171    0.00256          0.100      0.199      0.299
 140.00        0.00063    0.00127    0.00190          0.074      0.148      0.221
 160.00        0.00049    0.00097    0.00146          0.057      0.114      0.171
 180.00        0.00039    0.00077    0.00116          0.045      0.090      0.135
 200.00        0.00031    0.00063    0.00094          0.037      0.073      0.110
 220.00        0.00026    0.00052    0.00078          0.030      0.061      0.091
 240.00        0.00022    0.00044    0.00066          0.026      0.051      0.077
 260.00        0.00019    0.00037    0.00056          0.022      0.044      0.065
 280.00        0.00016    0.00032    0.00048          0.019      0.038      0.056
 300.00        0.00014    0.00028    0.00042          0.016      0.033      0.049
 320.00        0.00012    0.00025    0.00037          0.014      0.029      0.043
 340.00        0.00011    0.00022    0.00033          0.013      0.026      0.038
 360.00        0.00010    0.00020    0.00029          0.011      0.023      0.034
 380.00        0.00009    0.00018    0.00026          0.010      0.020      0.031
 400.00        0.00008    0.00016    0.00024          0.009      0.019      0.028
*****

```

TABLE 2.2.2

Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole high-pass analog filtering with a 3-dB corner frequency of 25 Hz. Frequencies are in Hz and time-shift magnitudes in seconds.

```

*****
FREQUENCY      T(N=2)      T(N=4)      T(N=6)      X(N=2)      X(N=4)      X(N=6)
*****
25.00         0.01000    0.02000    0.03000      1.167      2.333      3.500
50.00         0.00295    0.00590    0.00886      0.344      0.689      1.033
75.00         0.00137    0.00273    0.00410      0.159      0.319      0.478
100.00        0.00078    0.00156    0.00234      0.091      0.182      0.273
125.00        0.00050    0.00101    0.00151      0.059      0.117      0.176
150.00        0.00035    0.00070    0.00105      0.041      0.082      0.123
175.00        0.00026    0.00052    0.00077      0.030      0.060      0.090
200.00        0.00020    0.00040    0.00059      0.023      0.046      0.069
225.00        0.00016    0.00031    0.00047      0.018      0.037      0.055
250.00        0.00013    0.00025    0.00038      0.015      0.030      0.044
275.00        0.00010    0.00021    0.00031      0.012      0.024      0.037
300.00        0.00009    0.00018    0.00026      0.010      0.021      0.031
325.00        0.00008    0.00015    0.00023      0.009      0.018      0.026
350.00        0.00006    0.00013    0.00019      0.008      0.015      0.023
375.00        0.00006    0.00011    0.00017      0.007      0.013      0.020
400.00        0.00005    0.00010    0.00015      0.006      0.012      0.017
*****

```


TABLE 2.2.3

Time shifts and equivalent percent of a 70-beat-per-minute cardiac cycle resulting from 2-, 4-, and 6-pole low-pass analog filtering with a 3-dB corner frequency of 20 Hz. Frequencies are in Hz and time-shift magnitudes in seconds.

```

*****
FREQUENCY      T(N=2)    T(N=4)    T(N=6)      X(N=2)    X(N=4)    X(N=6)
*****
  1.00          0.01590  0.03180  0.04771     1.855     3.711     5.566
  2.00          0.01586  0.03173  0.04759     1.851     3.701     5.552
  3.00          0.01580  0.03160  0.04739     1.843     3.686     5.529
  4.00          0.01571  0.03142  0.04712     1.833     3.665     5.498
  5.00          0.01560  0.03119  0.04679     1.820     3.639     5.459
  6.00          0.01546  0.03092  0.04639     1.804     3.608     5.412
  7.00          0.01531  0.03062  0.04593     1.786     3.572     5.358
  8.00          0.01514  0.03028  0.04542     1.766     3.533     5.299
  9.00          0.01496  0.02991  0.04487     1.745     3.490     5.234
 10.00         0.01476  0.02952  0.04428     1.722     3.444     5.165
 11.00         0.01455  0.02910  0.04365     1.698     3.395     5.093
 12.00         0.01434  0.02867  0.04301     1.672     3.345     5.017
 13.00         0.01411  0.02823  0.04234     1.646     3.293     4.939
 14.00         0.01389  0.02777  0.04166     1.620     3.240     4.860
 15.00         0.01366  0.02731  0.04097     1.593     3.186     4.779
 16.00         0.01342  0.02685  0.04027     1.566     3.132     4.698
 17.00         0.01319  0.02638  0.03957     1.539     3.078     4.617
 18.00         0.01296  0.02592  0.03888     1.512     3.024     4.536
 19.00         0.01273  0.02546  0.03819     1.485     2.970     4.455
 20.00         0.01250  0.02500  0.03750     1.458     2.917     4.375
*****

```

Thus for a 5-pole filter such as that used in our experiments, we expect a maximum time shift over this frequency range of about 0.31 sec. According to Fig. 2.2.2, the shift in the midsystolic peak which moves from S1 toward S2 as the frequency decreases is about 0.55 sec over the same range. This analysis suggests that analog filtering is responsible for some of this apparent migration, and that Fig. 2.2.2 is therefore somewhat misleading.

2.3 -- Zero-Phase-Shift Digital Filtering of Apexcardiograms

Digital filtering provides a solution to this problem; however, the filtering cannot be done in real time, since the key to the method is the cancelling out of phase shifts by processing the time record of the signal backwards as well as forwards. Since we shall rely heavily on this extremely useful zero-phase-shift property, we include a discussion developed from a more succinct explanation given by Stearns (Stearns 1975, 138-140).

Following Stearns' notation conventions, we begin with

$$\tilde{H}(z) = \sum_{m=-\infty}^{\infty} h_m z^{-m} \quad (Z\text{-transform});$$

$$\bar{H}(j\omega) = \text{discrete Fourier transform (DFT)}.$$

For linear phase shift, a transfer function in the DFT domain must be of the form

$$\bar{H}(j\omega) = \bar{R}(j\omega) e^{-j\omega n T}, \quad (2.3.1)$$

where $\bar{R}(j\omega)$ is real, T is the interval between samples, and n is an integer. The Z-transform is related to the DFT by

$$\bar{F}(j\omega) = \tilde{F}(e^{-j\omega T}),$$

whence for the case of a transfer function with linear phase shift,

$$\tilde{H}(z) = z^{-n} \tilde{R}(z),$$

where

$$\tilde{R}(z) = \tilde{R}(e^{j\omega T})$$

is real. But in that case, we must have

$$\tilde{R}(e^{-j\omega T}) = \tilde{R}(e^{j\omega T})$$

or

$$\tilde{R}(z^{-1}) = \tilde{R}(z). \quad (2.3.2)$$

From Eq. (2.3.1) it may be seen that the phase and time shifts are $n\omega T$ and nT , respectively. For zero phase shift, we require that n vanish, so that

$$\bar{H}(j\omega) = \bar{R}(j\omega).$$

Stearns points out that Eq. (2.3.2) may be satisfied by any recursive $\tilde{H}(z)$ if time reversal is employed. That is, if there are $2N+1$ samples, the time series is

$$[f_m] = [f_{-N}, f_{-N+1}, \dots, f_0, \dots, f_N],$$

and the reversed series is

$$\begin{aligned} [f_m]^r &= [f_N, f_{N-1}, \dots, f_0, \dots, f_{-N}] \\ &= [f_{-m}]. \end{aligned} \quad (2.3.3)$$

In consequence, the Z-transform of the reversed series satisfies

$$\begin{aligned}\tilde{F}^r &= \sum [f_m]^r = \sum_{-N}^N f_{-m} z^{-m} \\ &= \sum_{-N}^N f_m z^m = \sum_{-N}^N f_m (z^{-1})^{-m},\end{aligned}$$


or

$$\tilde{F}^r(z) = \tilde{F}(z^{-1}), \quad (2.3.4)$$

and the sum, product and reversal operations can be seen to satisfy

$$\begin{aligned}[\tilde{A}(z) + \tilde{B}(z)]^r &= \tilde{A}^r(z) + \tilde{B}^r(z) \\ [\tilde{A}(z) \tilde{B}(z)]^r &= \tilde{A}^r(z) \tilde{B}^r(z) \\ [\tilde{A}^r(z)]^r &= \tilde{A}(z).\end{aligned} \quad (2.3.5)$$

We now consider the system



Here

$$\begin{aligned}\tilde{G}(z) &= \left\{ \tilde{H}(z) \left[\tilde{H}(z) \tilde{F}(z) \right]^r \right\}^r \\ &= \tilde{H}^r(z) \tilde{H}(z) \tilde{F}(z) \\ &= \tilde{H}(z^{-1}) \tilde{H}(z) \tilde{F}(z),\end{aligned}$$

which we obtain by using (2.3.5) and then (2.3.4). The net transfer function is then

$$\tilde{H}_{\text{NET}}(z) = \tilde{H}(z^{-1}) \tilde{H}(z) ,$$


whence, in terms of the DFT,

$$\bar{H}_{\text{NET}}(j\omega) = \bar{H}(-j\omega) \bar{H}(j\omega) = \bar{H}^*(j\omega) \bar{H}(j\omega) ,$$

so that

$$\bar{H}_{\text{NET}}(j\omega) = |\bar{H}(j\omega)|^2 . \quad (2.3.6)$$

Equivalently, if the system is



then

$$\begin{aligned} \tilde{G}(z) &= \tilde{H}(z) [\tilde{H}(z) \tilde{F}^r(z)]^r \\ &= \tilde{H}(z) \tilde{H}^r(z) \tilde{F}(z) , \end{aligned}$$

resulting in Eq. (2.3.6) as before. Thus either of these two systems, with the addition of a positive-square-root operation, can be used to obtain a zero-phase-shift equivalent of a given transfer function, even a transfer function containing a nonlinear phase shift.

We have implemented this approach by adapting digital filter designs to obtain zero-phase-shift high-, low-, and bandpass digital Butterworth filters by means of FORTRAN programs combining standard digital filtering routines with record reversal.

Fig. 2.3.1 shows the result of using zero-phase-shift digital filtering to obtain a wideband apex-cardiogram plot similar to that of Fig. 2.2.2. As before, the filter cutoff frequencies are scaled logarithmically so that the frequency ratio from one trace to the next remains constant. However, in this case filtering was achieved with the digital equivalent of a 12-pole Butterworth filter without phase shift. The migration of the peaks following the first and second sounds is much less pronounced, but not entirely absent (note the dashed lines in Fig. 2.3.1).

It will be noted that the traces corresponding to corner frequencies below the pulse rate of about 1.2 Hz are virtually indistinguishable, which is not surprising. Of particular interest is the midsystolic peak (between S1 and S2), especially pronounced in the 8.36-Hz trace, since it corresponds clearly neither to the apex-




FIGURE 2.3.1

zero-phase-shift digital filtering of a wideband apex-cardiogram obtained from a normal subject in the left lateral decubitus position. The audible first (S1) and second (S2) sounds are labelled, and subaudio peaks of particular interest as described in the text are marked with curved arrows. Peak "migration" (see Fig. 2.2.2) is most absent.

cardiogram nor to the recognized audible or subaudio heart sounds; it is considerably lower in frequency than the lower limits of conventional phonocardiography. We have been unable to find any reference to this feature, nor could it be identified for us by a cardiologist to whom it was made audible by means of the technique described in Chapter 3. At present writing, the most we can say is that we have found this peak to coincide roughly with the peak carotid-artery pressure.

A similar peak, especially pronounced in the 6.12-Hz trace, occurs in the vicinity of the third heart sound, which is normally inaudible in healthy adults.

We mention parenthetically that Figs. 2.2.2 and 2.3.1 were obtained from different recordings, as may be seen by comparing the unfiltered apexcardiogram traces, but from the same subject using the same transducer. Small variations in the location of the transducer, the position of the subject, and the amount of air in the lungs appear to have profound effects upon the apexcardiogram. Apexcardiography is recognized as technically the most difficult of the procedures for measurement of cardiac vibration, and cannot always be

employed successfully (Tavel 1972, 30). One of the problems related to the practical measurement of apexcardiograms is discussed in the next section.

2.4 -- Processing of Supine Apexcardiograms

Apexcardiograms are recorded, by preference, with the subject in the left lateral decubitus position; that is, lying with the left side turned downward and the upper body somewhat elevated. Upon visiting a hospital coronary care unit to obtain apexcardiogram recordings reflecting abnormalities in cardiac function, we learned that even in cases where satisfactory apexcardiograms could presumably be obtained with the subject in the left lateral decubitus position, this position may be sufficiently uncomfortable or even dangerous for some patients to preclude its use.

As a result, the recordings obtained on this occasion were taken with the subjects supine, resulting in apexcardiograms quite different in appearance from those we had seen in literature. This fact raises an obvious question: if the optimal position is inconvenient or actually contraindicated in certain cases where the information obtained from the apexcardiogram might be help-

ful, can the same information somehow be obtained from a supine apexcardiogram, or can such a tracing be somehow processed to obtain the equivalent left-lateral-decubitus apexcardiogram?

Fig. 2.4.1 shows two apexcardiograms obtained from a normal subject in each of these two positions. The top trace is unfiltered; the bottom trace had to be pre-processed with a high-pass analog filter having a corner frequency of 0.1 Hz (the lower limit for the available apparatus) to remove enough low-frequency baseline shift to keep the signal in range long enough to be sampled. This fact is unfortunate from the standpoint of phase shift as discussed in Section 2.3, but system modification would have been required to make the analog filtering unnecessary.

Fig. 2.4.2 shows an apexcardiogram "contour" plot similar to that of Fig. 2.3.1 except that the original signal is the supine apexcardiogram of Fig. 2.4.1(b). The variations from trace to trace are imperceptible below 6.12 Hz, and the high-frequency noise level is higher; in fact the transducer signal was of considerably lower amplitude in this instance.




FIGURE 2.4.1

Apexcardiograms obtained from a normal subject (a) in the left lateral decubitus position, and (b) supine. The lower trace has undergone high-pass analog filtering with a 3-dB corner frequency of 0.1 Hz.




FIGURE 2.4.2

Zero-phase-shift digital high-pass filtering of a wideband apexcardiogram obtained from a normal subject in the supine position. The audible first (S1) and second (S2) sounds are labelled.

Fig. 2.4.3 shows that low-pass filtering of the supine apexcardiogram at 20 Hz is not sufficient to recover from it a waveform at all similar to the left-lateral-decubitus waveform. Fig. 2.4.4 represents an unsuccessful attempt to to recover a normal-looking apexcardiogram waveform from a wideband supine apexcardiogram by digital low-pass filtering at successively lower frequencies. Even the lowest-frequency traces are of doubtful similarity to the apexcardiograms of Figs. 2.1.1 through 2.3.1 obtained from the same subject in the left lateral decubitus postition.

Thus the problem of obtaining a "standard" apexcardiogram from a supine subject appears to be more complex than a matter of simple filtering. One might consider trying to model the result of the change in position in terms of a transfer function similar to that representing a musical-instrument formant structure. In spite of the pulse-rate periodicity in the spectra shown by Kesteloot and his colleagues (Kesteloot et al 1969), it seems plausible that the relationship between the supine apexcardiogram $x(t)$ and the equivalent left-lateral-decubitus apexcardiogram $y(t)$ might be described approximately by the convolution relationship




FIGURE 2.4.3

Apexcardiogram (a) derived by zero-phase-shift digital low-pass filtering with a 3-dB corner frequency of 20 Hz from a wideband apexcardiogram (b) obtained from a normal subject in the supine position.




FIGURE 2.4.4

zero-phase-shift digital low-pass filtering, at successively lower frequencies, of a wideband apexcardiogram obtained from a normal subject in the position.

$$y(t) = h(t) * x(t) ,$$

where $h(t)$ is the impulse response of a linear time-invariant transfer function which may be represented in the frequency domain by

$$H(f) = \frac{Y(f)}{X(f)} .$$

Part of our motivation for this assumption is the suspicion that the periodicity in the aforementioned spectra may be a spurious result arising from windowing to an interval equivalent to the duration of the cardiac cycle; we will return to this subject in Chapter 7.

If the transfer function can be determined from measurements on the same subject of both types of apex-cardiogram -- conceivably by means of deconvolution (Stockham et al 1975) -- there is still no assurance that it will be invariant from one individual to another. It might, for example, be desirable to normalize in the frequency domain according to the size of the subject, or (better yet) according to spectral measurements; double-bass formant frequencies would differ from those of the violin even if the internal relationships were identical.

Another complication arises from the fact that differences in the two types of apexcardiograms may result, not only from changes in the configuration of the vibrating system as a function of position, but from changes in cardiac function itself as a function of the change in position. Thus the transfer function we have defined, even if otherwise invariant or easily adjustable, might conceivably vary according to the nature of the very abnormality one wishes to detect.

We do not pretend to have a solution to this problem; we merely wish to bring its existence to the reader's attention, and will suggest later on that it be studied further.

CONVERSION OF HEART SOUNDS
BY FREQUENCY MODULATION3.1 -- Introduction

An inherent difficulty in cardiac auscultation is the fact that a substantial amount of the energy in heart sounds as heard through a stethoscope is distributed among frequencies near the lower limit of human aural acuity, where the auditory system is comparatively insensitive. Ironically, "Our ears undoubtedly evolved this low-frequency insensitivity so as not to be bothered by internal sounds produced by heartbeats, blood moving through veins, and so on." (Backus 1969, 86) It is not surprising, therefore, that it is recommended that the weak sounds produced by cardiac function be listened to in especially quiet circumstances, preferably in well-insulated locations set aside for the purpose (Ravin 1967, 23).

We shall be discussing instrumentation for use in facilitating the extraction of useful diagnostic information from heart sounds by listening in real time. To establish the necessary basis for that discussion it is

important that we consider the frequency bandwidth actually occupied by heart sounds as well as some pertinent aspects of traditional stethoscopic auscultation.

3.2 -- Frequency Bandwidth of Heart Sounds

The frequency bandwidth actually occupied by heart sounds is given as 15-2000 Hz by Bukstein (Bukstein 1973, 18). Ravin reports that 80% of the energy of the first and second sounds is associated with frequencies below 70 Hz, that most sounds and murmurs are composed of frequencies below 500 Hz, and that frequency components above 650 Hz are of little importance in auscultation (Ravin 1967, 15).

Of course, ascribing little importance to frequencies above 650 Hz provides no assurance that those frequencies are absent from the sounds. However, our own measurements on commercially recorded heart sounds seem to bear out Ravin's assertion, as do spectra plotted from recordings of the author's heart sounds.

A measurement of the spectrum of sounds recorded at the cardiac apex of a normal subject and distributed on

phonograph records for training purposes may be seen in Fig. 3.2.1. The spectrum is averaged over one cardiac cycle. The bandwidth involved appears to be roughly 20-150 Hz; the higher-frequency peaks in the spectrogram are probably a result of 60-Hz hum associated with the recording or the copying process by which the tape loop was prepared so that the spectrogram could be plotted.

Fig. 3.2.2 shows the author's cardiac spectrum as recorded at the cardiac apex using the arrangement sketched in Fig. 3.2.3. The frequency bandwidth involved appears to extend from below 20 Hz to about 170 Hz; 60-Hz hum is clearly evident, and some of its concomitant overtones (180, 240, and 300 Hz) may be seen. The overall shape of this spectrum, representing as it does the frequency content of one entire cardiac cycle, is quite similar to that obtained from one cycle of the commercial recording.

In both of these spectra, the first two (and by far the most prominent) peaks occur at about 40 and 60 Hz, although these large peaks occur in a region of frequencies where aural acuity decreases with frequency at a rate of about 40 dB per decade for weak sounds (Fletcher

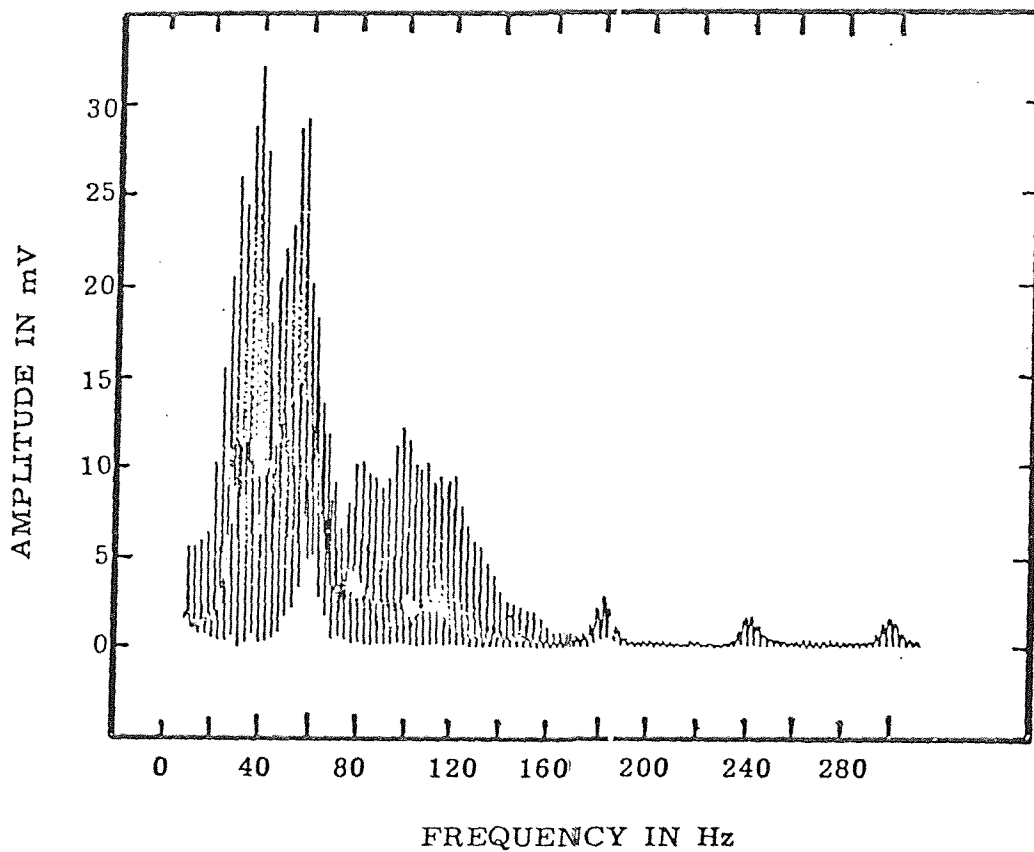


FIGURE 3.2.1

Spectrum of normal heart sounds, averaged over one cardiac cycle chosen at random from a commercial recording.

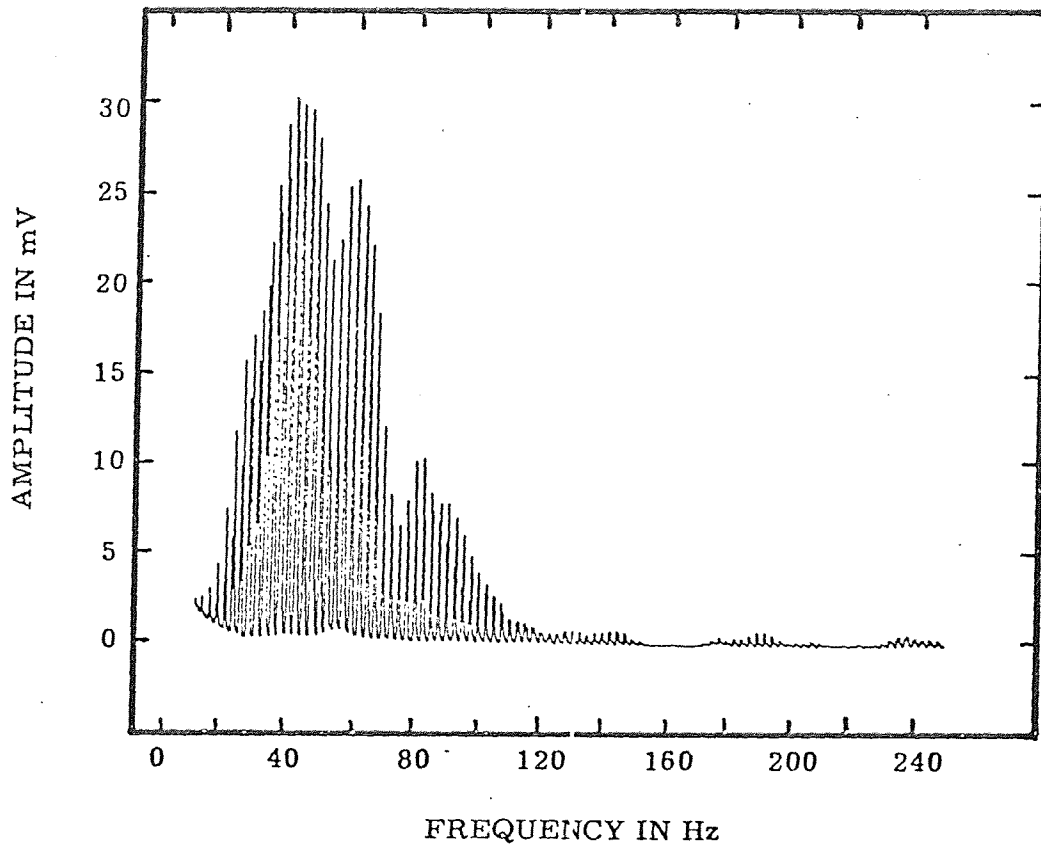


FIGURE 3.2.2

Spectrum of the author's heart sounds, recorded at the cardiac apex using a stethoscope head with diaphragm attachment and averaged over one cardiac cycle.




FIGURE 3.2.3

Arrangement used in making the recording from which the spectrum of Fig. 3.2.2 was plotted. A length of rubber tubing was used to couple the stethoscope head to the microphone; the two elements were sufficiently close together to ensure that no acoustic resonances at audible frequencies would be generated in the tubing.

and Munson 1933), their absence is discernable if they are filtered out by means of a high-pass filter when playing the recordings on reproduction equipment of reasonably good quality.

The highest-frequency components in the recorded heart sounds at the author's disposal were found in the case of the musical "seagull" murmur of early diastole in an eight-year-old girl with rheumatic heart disease and aortic and mitral insufficiency; the spectrum (Fig. 3.2.4) shows a substantial peak extending to 500 Hz.

3.3 -- Some Aspects of the Physics of Auscultation

In proposing instrumentation for aid in real-time auscultation, we will be faced with design choices which can be made realistically only if we consider the bandwidth occupied by the sounds of interest and also some of the characteristics of the human auditory system and the previous conditioning of the medical personnel whom we are trying to assist. We will consider these matters in turn.

Our measurements have verified that most of the energy in heart sounds observed at the surface of the




FIGURE 3.2.4

Spectrum of a musical "seagull" murmur, averaged over one cardiac cycle chosen at random from a commercial recording.

body is associated with frequencies at or below the lower frequency limit of human hearing. In addition, other limitations inherent in the human auditory system are of importance in this context. In cardiac auscultation one is concerned with the presence or absence of various sounds, the timing and relative loudness of the sounds, and to some degree the pitch and tone quality of the sounds. A pencil-and-paper method for recording auscultation findings, described in 1933 and highly recommended with minor modifications in Ravin's auscultation textbook and in the descriptive material accompanying the Merck, Sharp, and Dohme heart sound recordings, is more successful in displaying details of timing and relative loudness than in specifying pitch and tone quality (Segall 1933; Ravin 1968, passim). Clearly, loudness variations are important in detecting and classifying sounds and murmurs and in evaluating the splitting of the first and second sounds. Unfortunately, the human auditory system is better equipped to detect changes in pitch than changes in loudness. Ravin himself points out that "If a faint sound follows immediately upon a loud sound, the ear is not adjusted to hear it. Because of this phenomenon faint murmurs that follow loud sounds are heard with difficulty, or not at all. The same is true of faint sounds that follow loud murmurs." (Ravin 1967, 16)

Furthermore, subjective loudness is a function of duration while pitch is not, provided the duration exceeds the time threshold, or minimum duration required for a pure tone to produce a sensation of pitch. This minimum duration is about 10-15 msec, or 2-3 periods, whichever is longer (Roederer 1973, 81). But the subjective loudness of a tone pulse increases with duration for the first second or so, for frequencies of 200 Hz or less; for low-frequency tones whose duration is on the order of 0.1 sec (consistent with individual heart sounds) the ratio of perceived loudness to actual (steady-state) loudness is roughly doubled if the duration is doubled (Plomp and Bouman 1959).

There exists a slight dependence of pitch upon loudness for tones whose frequencies exceed about 1 kHz; however, considerable changes in loudness are needed to produce this effect (Stevens 1934).

The conditioning of medical personnel to the special peculiarities of the ordinary stethoscope cannot be overemphasized, mainly since that instrument serves as their primary means of access to heart sounds. We note, for example, that the written material accompanying each

of the three sets of commercial heart-sound recordings to which the author has had access contains the suggestion that the recordings be listened to through a stethoscope. Moreover, at least one commercial phonocardiograph is equipped with a "stethoscope filter" for listening purposes (Bukstein 1973, 155).

A stethoscope in normal use may be regarded essentially as a tube closed at one end by the patient's chest wall and at the other by the observer's tympanic membrane. Ravin stresses the importance of a good air seal at each end (Ravin 1967, 19). It may reasonably be expected, therefore, that the stethoscope will exhibit standing-wave resonances for sounds whose wavelengths in air are on the order of twice the length of the tube. The author's stethoscope (which is of the very common Sprague-Bowles type) is about 78 cm in length, taking into account the length of the meatus, or external ear canal, which is about 2.7 cm (Flanagan 1972, 88). For an ambient temperature of 23 °C, the speed of sound in air is (Beranek 1954, 10)

$$v = 33140 + (60.1)(23) = 34536 \frac{\text{cm}}{\text{sec}},$$

whence the lowest expected resonance frequency is

$$f = \frac{v}{2L} = \frac{34536}{(2)(18)} = 221 \text{ Hz}.$$

A second overall spectrum of the author's heart sounds was plotted, this time using the entire stethoscope as shown in Fig. 3.3.1, with the result shown in Fig. 3.3.2. This plot shows no characteristics which may clearly be attributed to the resonance described above, although there is a slight peak at 220 Hz (the second harmonic) which was not present in the earlier plot obtained using the stethoscope head alone.

Stethoscope resonance is of particular interest when one considers the frequent injunction to listen to recorded heart sounds through a stethoscope. In those circumstances the tube is open at one end and closed at the other, whether the bell or diaphragm attachment is used, since the soft diaphragm is virtually transparent to low-frequency sounds. Figs. 3.3.3 and 3.3.4 show the frequency response of the stethoscope head and the entire stethoscope respectively, as measured under these conditions using the arrangement illustrated schematically in Fig. 3.3.5. In this case, we expect resonances at frequencies such that the wavelength λ is related to the length L by




FIGURE 3.3.1

Arrangement used in making the recording from which the spectrum of Fig. 3.3.2 was plotted.




FIGURE 3.3.2

Spectrum of the author's heart sounds, recorded at the cardiac apex using an entire stethoscope with diaphragm attachment and averaged over one cardiac cycle.




FIGURE 3.3.3

Frequency response of a stethoscope head measured under free-field conditions.




FIGURE 3.3.4

ncy response of an entire stethoscope measured free-field conditions.

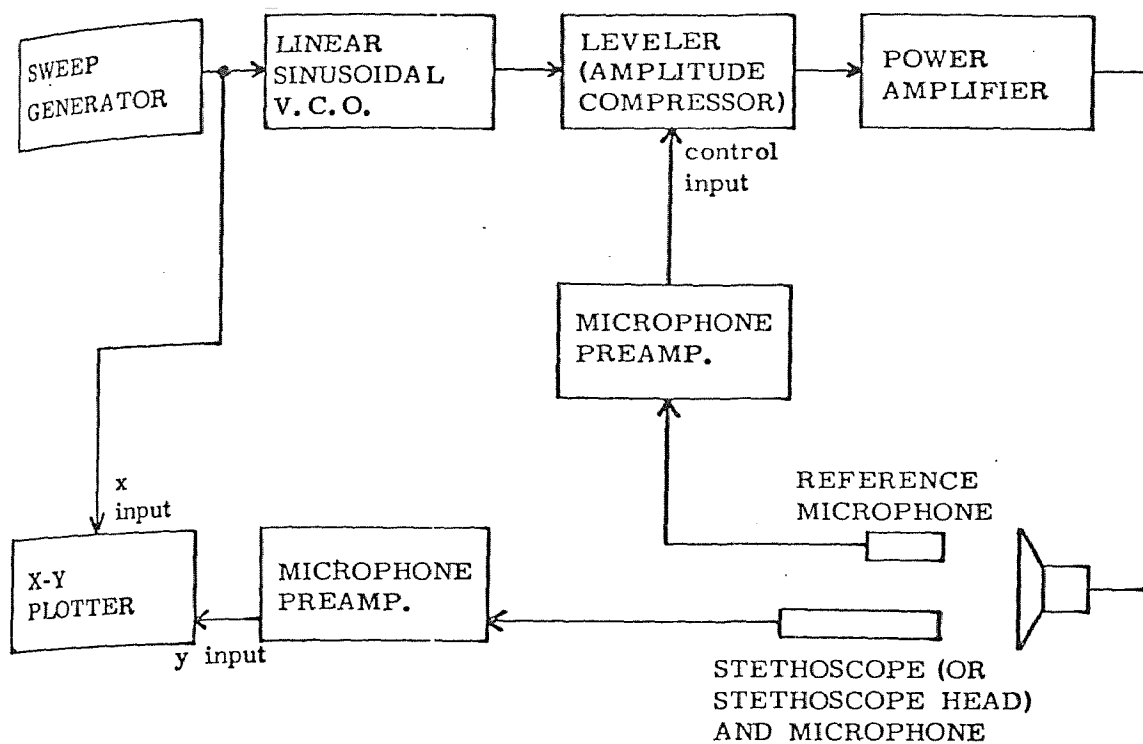


FIGURE 3.3.5

Arrangement used in plotting the frequency response of a stethoscope head and of an entire stethoscope, as seen in Figs. 3.3.3 and 3.3.4.

$$\lambda = \frac{4L}{n}, \quad n = 1, 3, 5, \dots$$

Thus the lowest resonance frequency should be

$$f = \frac{v}{4L} = \frac{34563}{(4)(78)} = 111 \text{ Hz}.$$

This resonance is indeed present in Fig. 3.3.4. It shows that listening to heart-sound recordings through a stethoscope offers certain tangible advantages, and not only that of the comfortable familiarity of heart sounds as heard through a stethoscope: the frequency response of the open-ended stethoscope provides a strong pre-emphasis in a region of important low frequencies where the response of the reproducing equipment is likely to be deficient.

3.4 -- Bandwidth Criteria for Heart-Sound Conversion

Clearly our basic goal is to bring heart sounds, by suitable processing, into a range of frequencies where they can be dealt with more easily by human observers. This frequency range should obviously be one in which the human auditory system is relatively sensitive. In practice we may want to modify our choice of frequencies in order to accommodate certain limitations of the equipment

which will be used. For example, it may be desirable to place as much of the converted information as possible within a radio or telephone communications passband.

The choice of frequency limits for the signals to be converted may be somewhat more subtle. It is to our advantage, from a signal-to-noise ratio standpoint, to exclude from the acceptance bandwidth of our system frequencies at which no useful diagnostic information is to be found. We note that Ravin assigns the following bandwidths to various normal and abnormal sounds for purposes of instructional simulation (Ravin 1968, 2).

First sound	70	-	110 Hz
Diastolic murmur	70	-	110 Hz
Second sound	80	-	135 Hz
Systolic murmur	370	-	380 Hz
Aortic stenosis murmur	180	-	400 Hz

The implication is that frequencies outside a bandwidth of 70-400 Hz may be disregarded in auscultatory diagnosis. The upper limit of 400 Hz differs from Ravin's own specification of 650 Hz, mentioned earlier, by 0.21 decades or a pitch interval of 8.41 semitones,

between a minor and major tempered sixth. In the spectra of Figs. 3.2.1, 3.2.2, and 3.3.2 the first two (and by far the most prominent) peaks occur at about 40 and 60 Hz, below the 70-Hz limit specified for simulation by Ravin. This finding is consistent with his statement, reported earlier, that most of the energy in heart sounds is associated with low frequencies.

It seems plausible, however, that these low frequencies are ignored in synthesis not because they contain no potentially useful information but because they are largely inaudible in conventional auscultation. We therefore choose a working bandwidth of approximately 20-650 Hz for our purposes, realizing that most of the energy in the converted sounds may well be derived from a previously inaudible portion of the input signal.

3.5 -- Possible Means of Heart-Sound Conversion

An extremely attractive approach to the real-time frequency conversion of heart sounds is offered by the phase vocoder (Flanagan 1972, 378). This device, combined with a linear predictor, has been used in the generation of electronic music to achieve independent control over time, pitch, and spectrum (Moorer 1976).

Amplitude modulation could certainly be used to shift the spectrum along the frequency axis to appropriately higher frequencies. However, from the Fourier transform pair

$$f(t) \cos \omega_c t \longleftrightarrow \frac{1}{2} [F(\omega + \omega_c) + F(\omega - \omega_c)]$$

it may be seen that the spectrum of $f(t)$ is shifted in frequency but not scaled, so that the bandwidth is unaltered. This means, among other things, that the shifted version of a harmonic spectrum will, in general, no longer be harmonic. We offer the following illustration to demonstrate that amplitude modulation is unsuitable for our purposes. Let the original bandwidth of 20-650 Hz be shifted to the most sensitive region of aural acuity, say 2020-2650 Hz. The original bandwidth covers about 1.5 decades or 5 octaves, while the transformed bandwidth covers about 0.1 decades, less than 5 semitones (a tempered fourth). While this would provide clearly audible sounds whose overall amplitude varied with that of the original sounds, one could expect little discernable variation in tone quality; moreover little advantage would be taken of the observer's pitch discrimination.

Our experiments have dealt with linear and exponential frequency modulation of an audio-frequency carrier, since this approach seems to offer relatively simple and inexpensive means of dealing with the problem. We shall discuss each of these two methods in turn.

3.6 -- Conversion of Heart Sounds by Linear Frequency Modulation

As a basis for our discussion and in order to establish some convenient notation, we begin by setting down some equations pertinent to frequency modulation. Let a signal $v_m(t)$ be used to modulate a sinusoidal carrier of fixed angular frequency ω_c in such a way that the instantaneous frequency $\omega(t)$ of the modulated carrier varies linearly with the modulating signal, thus:

$$\omega(t) = \alpha v_m(t) , \quad (3.6.1)$$

where α is a constant of proportionality. In the time domain, the modulated carrier $v(t)$ may be represented by the familiar FM equation: for a sinusoidal modulating signal of angular frequency ω_m ,

$$\begin{aligned} v(t) &= A \cos \int_0^t \omega(\tau) d\tau \\ &= A \cos [\omega_c t + m \sin(\omega_m t)] , \end{aligned} \quad (3.6.2)$$

where A is the carrier amplitude and m is the modulation index, the ratio of peak frequency deviation to modulating frequency, so that the peak frequency deviation is

$$\Delta\omega = m \omega_m \text{ (rad/sec)} \quad (3.6.3)$$

or

$$\Delta f = m f_m \text{ (Hz)}, \quad (3.6.4)$$

where f_m is the modulating frequency in Hz. Eq. (3.6.2) may be rewritten in terms of a Fourier series whose coefficients are Bessel functions of the first kind (Panter 1965, 279):

$$v(t) = A \sum_{k=-\infty}^{\infty} J_k(m) \cos [(\omega_c + k\omega_m)t] \quad (3.6.5)$$

Equivalently, the modulated carrier

$$v(t) = A \sin [\omega_c t + m \sin(\omega_m t)] \quad (3.6.6)$$

may be expressed as

$$v(t) = A \sum_{k=-\infty}^{\infty} J_k(m) \sin [(\omega_c + k\omega_m)t] \quad (3.6.7)$$

Eqs. (3.6.5) and (3.7.6) show that modulation of a sinusoid by another sinusoid can produce a complex spectrum; certainly modulation of a sinusoid by a complex waveform, using an adequately large modulation index, can produce a signal rich in harmonics.

Narrowband FM, where the modulation index is small enough so that only the terms corresponding to $k = 0, 1,$ and -1 in Eqs. (3.6.5) and (3.6.7) are significant, is used for transmission of electrocardiograms over ordinary telephones. It appears also to have been used with heart sounds to facilitate monitoring of cardiac function during surgery (Holloway et al 1975).

With large values of the modulation index, linear FM has been applied to the computer generation of electronic music (Chowning 1973), where the dependence of the spectrum upon the modulation index as well as the input signal can be used to advantage to achieve flexible control over tone quality. Our own experiments seem to indicate that this method can be used to advantage for our purposes as well, but with certain limitations which will be discussed below.

It must be borne in mind that in this application of FM the carrier, as well as the modulating signal, is of audio frequency, and that no subsequent demodulation is involved. Linear FM may be realized in real time by means of a voltage-controlled oscillator (VCO) whose output frequency is related to a control voltage by

$$\begin{aligned}
 \omega_{VCO} &= \alpha v_m(t) \\
 &= \alpha (v_{mdc} + v_{mac}) \\
 &= \omega_c + \alpha v_{mac} \geq 0.
 \end{aligned}
 \tag{3.6.8}$$

The FM equation for this case corresponding to Eq. (3.6.2) is

$$v(t) = A \cos \left[\omega_c t + \alpha \int_0^t v_{mac}(\tau) d\tau \right].
 \tag{3.6.9}$$

In the case of a sinusoidal modulating signal, the practical restriction to positive frequencies imposed by Eq. (3.6.8) limits the modulation index m :

$$\begin{aligned}
 \Delta \omega &= \alpha v_{mac} \leq \omega_c \\
 m = \frac{\Delta \omega}{\omega_m} &= \frac{\alpha v_{mac}}{\omega_m} \leq \frac{\omega_c}{\omega_m}.
 \end{aligned}
 \tag{3.6.10}$$

Even with the modulation depth limited by this practical restriction, however, complex spectra can be obtained by modulating one sinusoid by another in this way. When a 500-Hz carrier (Fig. 3.6.1) is modulated by a 100-Hz sinusoid with the maximum possible modulation index of 5 (Fig. 3.6.2), the spectrum of Fig. 3.6.3 is obtained. In this case the carrier frequency is a simple integral multiple of the modulation frequency, and a harmonic spectrum is obtained having significant frequency components up to 1100 Hz.

If the input voltage amplitude is increased (see Fig. 3.6.4), the VCO remains at 0 Hz for the negative-going portion of the input waveform (dotted). This is equivalent to modulating with a half-wave rectified sinusoid, which itself possesses a complex spectrum; not surprisingly, the VCO output spectrum (Fig. 3.6.5) is more complex than that of the last example, although still harmonic since the modulating signal remains periodic and the carrier frequency is still five times the fundamental frequency of the modulating signal.

A clear discontinuity in tone quality is noticed as the input signal amplitude is increased beyond the point




FIGURE 3.6.1

500-Hz carrier used in generating all of the FM spectra shown in this chapter.




FIGURE 3.6.2

100-Hz sinusoid used as a modulating signal to obtain the linear-FM spectrum shown in Fig. 3.6.3.




FIGURE 3.6.3

Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.2, using a linear VCO (500 Hz/V). The peaks at 50 Hz and at dc, corresponding respectively to those at 950 and 1000 Hz, are not shown.




FIGURE 3.6.4

100-Hz sinusoid used as a modulating signal to obtain the linear-FM spectrum shown in Fig. 3.6.5.




FIGURE 3.6.5

Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.4, using a linear VCO (500 Hz/V).

of clipping, placing a severe limitation on the practical generation of rich spectra by this method. Were negative frequencies admissible in practice, it may be seen from Eq. (3.6.5) or (3.6.7) that perfectly acceptable output voltages would be obtained. From Eq. (3.6.5), for example, since the cosine is an even function of its argument and the Bessel function satisfies

$$J_{-k}(m) = (-1)^k J_k(m),$$

negative frequencies will be obtained whenever

$$\left\{ \begin{array}{l} k < 0 \\ |k|m > \omega_c \end{array} \right\},$$

in which case the negative frequency component will be reflected about 0 Hz to reappear as a positive-frequency component with a 180-degree phase shift if k is odd.

This "frequency foldover" may be approximated with the VCO by using a modified full-wave rectifier to reflect input signal voltages about the voltage corresponding to 0 Hz (which may not be 0 V). For example, the sinusoid of Fig. 3.6.4 would be transformed into that of Fig. 3.6.6; the FM output spectrum as measured is




FIGURE 3.6.6

Modulating signal derived from that of Fig. 3.6.4 by means of artificial frequency foldover (in this case, full-wave rectification). The corresponding linear-FM spectrum is shown in Fig. 3.6.7.

shown in Fig. 3.6.7. Because of thermal drift in the foldover circuit, the reflected frequencies are not quite where they belong, but they clearly contribute to the spectrum.

3.7 -- Conversion by Exponential Frequency Modulation

There is another way around the problem of the limitations on modulation depth imposed by Eq. (3.6.10), which our experiments lead us to believe offers significant advantages for our purpose. This approach is motivated in part by the human auditory system's ability to recognize pitch intervals corresponding to frequency ratios. The number n of semitones between the two frequencies f and f' satisfies (Roederer 1973, 152)

$$\frac{f}{f'} = 2^{n/12},$$

so that

$$f = 2^{n/12} f' = \left[\exp\left(\frac{n \ln 2}{12}\right) \right] f'$$

and a VCO will produce the same pitch displacement from an arbitrarily-chosen starting frequency f' for the same increment V in control voltage provided its response satisfies




FIGURE 3.6.7

Spectrum of a 500-Hz sinusoid modulated by the signal shown in Fig. 3.6.6, using a linear VCO (500 Hz/V).

$$f = e^{\alpha V} f' ,$$

where α is a constant.

Suppose the input signal is applied to a voltage-controlled oscillator whose instantaneous angular frequency is given by

$$\omega(t) = \omega_0 e^{\alpha V_m(t)} \tag{3.7.1}$$

where ω_0 is the center (carrier) frequency obtained when $v_m(t) = 0$, and α is a constant. If the modulating signal is a sinusoid given by

$$V_m(t) = V_m \cos(\omega_m t) ,$$

the instantaneous VCO frequency is then

$$\omega(t) = \omega_0 e^{\alpha V_m \cos(\omega_m t)} .$$

The exponential in this expression may be expanded in a Fourier series whose coefficients are modified (hyperbolic) Bessel functions (Everett and Annars 1956, 376, Eq. 9.6.34):

$$\omega(t) = \omega_0 \left\{ I_0(\alpha V_m) + 2 \sum_{k=1}^{\infty} I_k(\alpha V_m) \cos(k\omega_m t) \right\}, \quad (3.7.2)$$

whence the modulated carrier may be expressed as

$$\begin{aligned} v(t) &= A \sin \left[\int_0^t \omega(\tau) d\tau \right] \\ &= A \sin \left\{ \omega_0 I_0(\alpha V_m) t + 2\omega_0 \sum_{k=1}^{\infty} \frac{I_k(\alpha V_m)}{k\omega_m} \sin(k\omega_m t) \right\} \quad (3.7.3) \end{aligned}$$

Again, modulation of a sinusoid by another sinusoid produces a complex spectrum, but with two important differences. One is that there is no restriction on modulation depth except that imposed by the upper frequency limit of the VCO itself. No negative frequencies are involved, and no hardware is required for simulation of frequency foldover.

The second major difference is reflected in the carrier term of the output signal. In the linear case, this term is seen from Eq. (3.6.7) to be

$$A J_0(m) \sin(\omega_c t) . \quad (3.7.4)$$

The amplitude of the carrier component varies with the modulation index, but the frequency is constant. In the exponential case, the carrier term is seen from Eq. (3.7.3) to be approximated for low V_m by

$$A \sin \omega_0 I_0(\alpha V_m) t . \quad (3.7.5)$$

Here the carrier term is of constant amplitude A , but its frequency is proportional to the amplitude V_m of the modulating signal and independent of the modulating frequency. Only the amplitudes of the higher-order terms in the series vary with the input signal frequency as well as its amplitude.

This is extremely important, because it affords a simple way of converting changes in heart-sound intensity into corresponding changes in frequency which the observer is equipped to discern far more easily.

We consider a practical example. The exponential VCO used in our experiments is calibrated to yield a frequency shift of one decade for a 6-V change in input voltage. Thus the instantaneous frequency in Hz is

$$f(t) = f_0 e^{\alpha V_m} = f_0 10^{\frac{V_m}{6}} , \quad (3.7.6)$$

whence

$$e^{\alpha V_m} = 10^{\frac{V_m}{6}} . \quad (3.7.7)$$

Taking the natural logarithm of both sides, we obtain

$$\alpha V_m = \frac{V_m}{6} \ln 10$$

$$\alpha = \frac{\ln 10}{6} = 0.3838 , \quad (3.7.8)$$

so that

$$f(t) = f_0 e^{0.3838 V_m} .$$

Let us calculate the change in pitch to be expected as a function of modulating signal amplitude for our apparatus. From Eq. (3.7.3) we see that if the frequency when $V_m = 0$ is f_0 , the shifted frequency f' is

$$f' = f_0 I_0(\alpha V_m) = f_0 I_0(x) ,$$

where the input-signal amplitude is low and

$$X = \alpha V_m = \frac{1}{6} V_m$$

for our apparatus. The pitch variation in semitones may be found from

$$\frac{f'}{f_0} = 2^{n/12},$$

whence, for low values of x ,

$$n = \frac{12 \log \left(\frac{f'}{f_0} \right)}{\log 2} = \frac{12 \log [I_0(x)]}{\log 2}, \quad (3.7.9)$$

where n is the number of semitones of shift. The zeroth-order modified Bessel functions may be calculated from (Spiegel 1968, 138, Eq.24.35)

$$I_0(x) = 1 + \frac{x^2}{2^2} + \frac{x^4}{2^2 \cdot 4^2} + \frac{x^6}{2^2 \cdot 4^2 \cdot 6^2} + \dots \quad (3.7.10)$$

From Eqs. (3.7.9) and (3.7.10) we are pleased to note that the shift in pitch is symmetric with respect to the input signal amplitude. Table 3.7.1 shows the variation of x , $I_0(x)$, f' , and semitone shift with the amplitude of a sinusoidal input voltage; these results are plotted in Fig. 3.7.1.

TABLE 3.7.1

Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.

V_m = input signal amplitude
 $I_0(x)$ = zeroth-order modified Bessel function
 f' = shifted frequency for f_0 of 500 Hz
 S = pitch variation in semitones

V_m (V)	x	$I_0(x)$	f' (Hz)	S (semitones)
0.0	0.0	1.0000	500.0	0.0
0.25	0.0959	1.0023	501.2	0.04
0.50	0.1919	1.0092	504.6	0.16
0.75	0.2878	1.0208	510.4	0.36
1.00	0.3838	1.0372	518.6	0.63
1.25	0.4797	1.0584	529.2	0.98
1.50	0.5756	1.0846	542.3	1.41
1.75	0.6716	1.1160	558.0	1.90
2.00	0.7675	1.1528	576.4	2.46
2.25	0.8635	1.1953	597.6	3.09
2.50	0.9594	1.2437	621.8	3.78
2.75	1.0554	1.2984	649.2	4.52
3.00	1.1513	1.3598	679.9	5.32
3.25	1.2472	1.4284	714.2	6.17
3.50	1.3432	1.5045	752.3	7.07
3.75	1.4391	1.5888	794.4	8.01
4.00	1.5351	1.6817	840.9	9.00
4.25	1.6310	1.7841	892.1	10.02
4.50	1.7269	1.8966	948.3	11.08
4.75	1.8229	2.0200	1010.0	12.17
5.00	1.9188	2.1552	1077.6	13.29
5.25	2.0148	2.3032	1151.6	14.44
5.50	2.1107	2.4651	1232.5	15.62
5.75	2.2066	2.6419	1320.9	16.82
6.00	2.3026	2.8350	1417.5	18.04




FIGURE 3.7.1

Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.

3.8 -- Pseudologarithmic Compression in Exponential FM Conversion

It is clear from Fig. 3.7.1 that the sensitivity of pitch change as a function of input sinusoid amplitude is low at low amplitudes and fairly uniform at high amplitudes. This may be seen in another way by plotting sensitivity in cents per volt as a function of the input amplitude, as we have done in Fig. 3.8.1.

The decrease in sensitivity at low output amplitudes may be a blessing in disguise, since it means that background noise will have little effect compared even with very small amplitude variations in the signal when the amplitude is reasonably large. On the other hand, the detection of low-amplitude sounds and murmurs might be facilitated if we had some way of making the sensitivity greater at low amplitude levels. Accordingly, we wish to provide this option if possible.

We have found that this capability can be incorporated into our system by means of a simple circuit which (as we will show) serves another useful purpose elsewhere in the system at the same time. It is the pseudologarithmic compression amplifier whose transfer function is plotted in Fig. 3.8.2.




FIGURE 3.8.1

Approximate pitch-variation sensitivity of an exponential VCO (6V/decade) as a function of the amplitude of a sinusoidal input voltage.




FIGURE 3.8.2

Transfer function of the pseudologarithmic compression amplifier. The dashed curves represent the true logarithmic response for input voltages of low magnitude.

In order to evaluate the results obtained when this circuit is incorporated into our system for modification of exponential FM sensitivity as discussed earlier, it is convenient to have an analytic expres-

sion for the transfer function of the pseudologarithmic amplifier. By successive approximation, we find that the transfer function shown in Fig. 3.8.2 may be closely approximated by

$$V_{OUT} = 1.655 \frac{V_{IN}}{|V_{IN}|} \left(1 - e^{-0.55331 |V_{IN}|} \right). \quad (3.8.1)$$

This approximation is plotted for positive input voltages in Fig. 3.8.3.

With the aid of Eqs. (3.7.9) and (3.8.1) we can calculate the pitch variation for exponential modulation when pseudologarithmic compression is used. The results, analogous to the earlier case shown in Fig. 3.7.1, are plotted in Fig. 3.8.4. As in the earlier case, the sensitivity of pitch change as a function of input sinusoid amplitude may be seen in another way by plotting sensitivity in semitones per volt as a function of the input amplitude; this has been done in Fig. 3.8.5.




FIGURE 3.8.3

Plot of an analytic approximation of the transfer function of Fig. 3.8.2 for positive voltages as calculated using Eq. (3.8.1). The actual transfer function, shown dashed, is superimposed.




FIGURE 3.8.4

Approximate pitch variation from the center frequency of an exponential VCO (6V/decade) using pseudologarithmic compression, plotted as a function of the amplitude of a sinusoidal input voltage.




FIGURE 3.8.5

Approximate pitch-variation sensitivity of an exponential VCO (6V/decade) using pseudologarithmic compression, plotted as a function of the amplitude of a sinusoidal input voltage.

In this instance it is clearly seen that the system will respond strongly to amplitude changes at low amplitudes, while the sensitivity to amplitude change is markedly decreased at high amplitudes. This is exactly what we hoped to achieve; our experiments have shown that low-level murmurs are much more readily detected in the converted sounds when pseudologarithmic compression is in use. Of course, increased sensitivity at low amplitudes means that low-amplitude background noise will be more evident in the converted sounds. We have therefore felt it desirable to incorporate into our system a noise threshold control circuit, whose operation will be described below, which seems to be quite helpful in dealing with this problem.

In order to facilitate alteration of system amplitude sensitivity so that the user may concentrate more easily on heart-sound components of high or low amplitude, a simple means of shifting continuously between the two extremes of sensitivity reflected in Figs. 3.8.4 and 3.8.5 is provided, as shown in Fig. 3.8.6. The 10-K potentiometer is used to balance the system so that a 5-V input signal will produce the same output amplitude regardless of the setting of the 100-K potentiometer.




FIGURE 3.8.6

Means of continuous selection of the degree of pseudologarithmic compression used in generating the control signal for the VCO.

The voltage follower (IC6) is included to prevent loading of the 100-K voltage divider output.

3.9 -- Effect of Pseudologarithmic Compression on Exponential FM-Conversion Spectra

Since pseudologarithmic compression is a nonlinear operation, it is to be expected that the spectra obtained with input signals so treated will be affected in some way. In practice, the variation in tone quality is not as pronounced as the difference in response to low-level changes in the input signal. At the same time, however, pseudologarithmic processing of a periodic input signal generates a certain amount of harmonic distortion. The signal shown in Fig. 3.9.2, for example, derived by pseudologarithmic compression from the sinusoid of Fig. 3.9.1, has the spectrum plotted in Fig. 3.9.3. Because of the symmetry of the pseudologarithmic amplifier transfer function, the processed waveform displays half-wave symmetry; as a result only odd-numbered harmonics are present in the spectrum.

To illustrate the differences in the exponential FM spectra obtained with and without the nonlinear processing for this example, appropriate spectra are shown in




FIGURE 3.9.1

100-Hz sinusoid of amplitude 1V, used as an input to obtain the modulating signal shown in Fig. 3.9.2.




FIGURE 3.9.2

100-Hz signal obtained from the sinusoid of Fig. 3.9.1 by pseudologarithmic compression and scaled to the same amplitude (1V).




FIGURE 3.9.3

Spectrum of the 100-Hz VCO modulating signal shown in Fig. 3.9.2.

Figs 3.9.4 through 3.9.9 for three different values of VCO control-voltage amplitude.

3.10 -- Pseudologarithmic Compression in System Output Amplitude Control

Since the system under discussion converts both instantaneous changes in the input voltage and alterations in input signal amplitude into changes in frequency, no information need be associated with system output amplitude unless it is so desired. To be sure, some sort of control over output amplitude is desirable, if only to avoid listening to the unmodulated carrier while the system is quiescent. It would not, however, be helpful to impress the input signal envelope upon the system output, since this would mean attenuation of information associated with low-level input signals which we have been at pains to extract.

We have experimented with "soft-switch" (click-free) gating of the output controlled by an amplitude-threshold circuit triggered by the input signal, but we feel that this approach is dangerous since changes in the output signal associated with input voltage changes below the switching threshold are totally absent and valuable information may be inadvertently overlooked.




FIGURE 3.9.4

An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 1V.




FIGURE 3.9.5

An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 2V.




FIGURE 3.9.6

An exponential FM spectrum. The input is a 100-Hz sinusoid of amplitude 3V.




FIGURE 3.9.7

An exponential FM spectrum. The input is a 1-V signal derived by pseudologarithmic compression from a 100-Hz sinusoid.




FIGURE 3.9.8

An exponential FM spectrum. The input is a 2-V signal derived by pseudologarithmic compression from a 100-Hz sinusoid.




FIGURE 3.9.9

An exponential FM spectrum. The input is a 3-V signal derived by pseudologarithmic compression from a 100-H sinusoid.

A more satisfactory approach is to control the output amplitude using the output of the pseudologarithmic compressor, so that the system output amplitude decreases in relation to the input amplitude, but is not cut off entirely, and approaches its maximum value more rapidly than does the input amplitude in order to facilitate observation of low-level changes in the original signal.

Because of the close relationship between input and output signals, it was felt desirable to include the option of controlling the output amplitude in either of two ways: either by means of an envelope follower (a full-wave rectifier and peak detector) operating upon the output of the pseudologarithmic compressor, or by means of a full-wave rectifier output without the peak detector. In the latter case the output sounds somewhat "cleaner;" amplitude-modulating the frequency-modulated carrier in this way results in a relative attenuation of higher-frequency components, as may be seen by comparing Figs 3.10.1 and 3.10.2.

A means of controlling the input signal amplitude at which the output amplitude begins to rise is useful, particularly when the frequency-shift sensitivity is




FIGURE 3.10.1

Spectrum of a 500-Hz sinusoidal carrier modulated with a 100-Hz sinusoid of 6V amplitude, using an exponential VCO (6V/decade), after pseudologarithmic compression of the modulating signal.




FIGURE 3.10.2

Spectrum of the same modulated carrier used to produce Fig. 3.10.1, except that the frequency-modulated carrier has in this instance been subsequently amplitude-modulated by the output of the pseudologarithmic amplifier.

chosen so as to favor low-level changes, which also results in greater sensitivity to noise. Accordingly, we have incorporated into the system a noise threshold control circuit whose transfer function is shown in Fig. 3.10.3. The threshold voltage ϵ is adjustable from zero to -5 V. With this circuit the maximum and minimum output amplitudes are constant; all that changes is the rate of transition from minimum to maximum as the input amplitude increases.

3.11 -- Special Applications of FM Conversion

We have touched on the possible need for placing as much of the converted heart-sound information as possible within a particular bandwidth such as that of a telephone communications channel. This turns out to be feasible by means of the calculations given below. We consider only the exponential-VCO case since we feel that it is unquestionably the method of choice.

Let the output bandwidth desired extend from f_1 to f_2 . We then require a center frequency f_0 satisfying

$$\frac{f_2}{f_0} = \frac{f_0}{f_1} = 10^{\delta V_m} \quad , \quad (3.11.1)$$




FIGURE 3.10.3

Transfer function of the noise-threshold control circuit. The threshold voltage ϵ may be varied from zero to $-5V$. Since this portion of the circuit deals with inverted signals, increasing instantaneous signal amplitudes are reflected in increasingly large negative values of the input and output voltages.

where γ is the VCO response in decades per Volt, and V_m is the amplitude of the modulating signal, which is assumed to have an average value of 0 V. Solving this equation for f_o , we obtain

$$f_b = (f_1 f_2)^{1/2} . \quad (3.11.2)$$

The required bandwidth will be "filled" if V_m is chosen such that

$$\gamma V_m = \log \left(\frac{f_2}{f_o} \right) ,$$

whence

$$V_m = \frac{1}{\gamma} \log \left[\frac{f_2}{\sqrt{f_1 f_2}} \right] = \frac{1}{2\alpha} \log \left(\frac{f_2}{f_1} \right) . \quad (3.11.3)$$

To test the efficacy of this approach, let us try to fit FM-converted heart sounds into a telephone bandwidth of 500 to 5000 Hz. From Eq. (3.11.2) we find that the required center frequency is

$$f_o = \sqrt{(500)(5000)} = 1581 \text{ Hz} ,$$

while the peak amplitude is given by Eq. (3.11.3) as

$$V_m = \frac{1}{(2)(1/6)} \log \left(\frac{5000}{500} \right) = 3 \text{ V}$$

for our apparatus ($\gamma = \frac{1}{6}$ decades/V). A third-octave-band spectrum for normal heart sounds converted according to these specifications is shown in Fig. 3.11.1.

FM conversion by means of an exponential VCO offers convenient independent control over center frequency and pitch interval variation. Accordingly, a second special application suggests itself.

Conventional auscultation is often made especially difficult as a result of several aural phenomena (murmurs, clicks, snaps, etc.) occurring in close-temporal proximity to normal sounds or to each other. If, in addition, tachycardia is present, distinguishing the various individual sounds during auscultation may be next to impossible.

Although our technique is intended primarily as a "real-time" aid to cardiac diagnosis and monitoring, we find that it offers a simple and inexpensive means of slowing down the heart sounds in conjunction with the




FIGURE 3.11.1

Third-octave-band spectrum showing the bandwidth of normal heart sounds processed by pseudologarithmic compression and converted to fit into a telephone bandwidth of 500-5000 Hz. The spectrum is averaged over 8 sec.

frequency conversion if an ordinary audio-frequency tape recorder is available. For example, to slow the apparent pulse rate by a factor of two, the center frequency is set to twice its desired value; the converted sounds are recorded on tape, and subsequently played back at half speed. This brings the center frequency back to normal while the apparent pulse rate is half of the original value.

Similarly, the apparent pulse rate can be decreased by a factor of n if the center frequency is chosen n times its normal value, and the recording of the converted sounds is played back at $1/n$ times the speed at which the recording was made, subject only to the limitations imposed by the high-frequency response of the recorder and the speeds available. Alternatively, two recorders may be used to reduce the apparent pulse rate by reducing the playback speed in steps. The author, being inexperienced in auscultation, can testify that this method is very helpful in learning to separate sounds occurring in close temporal proximity and to estimate splitting of the first and second heart sounds.

Among the recordings accompanying this dissertation are examples of speed reduction by factors of 2 (Band 4) and 4 (Band 7).

3.12 -- Low-Frequency Considerations in FM Conversion

In Chapter 1 we pointed out that the audible frequencies account for less than 5% of the vibrational energy associated with the cardiac cycle. The measurements made in connection with the material of this chapter were limited by the frequency response of the microphone used in some cases, and by the frequency range of the commercial recordings at our disposal in others. Later experiments performed using a piezoelectric contact transducer with a frequency response reasonably flat from 0.02 Hz to 2 kHz have shown that the technique we have described can be used just as well at the lowest frequencies.

However, two practical problems are found in connection with the conversion of very low frequencies into audible sounds. One is that the far greater amplitudes at these frequencies tends to obscure the weaker "landmarks" in the converted sounds associated with the normally audible components of the original vibrations; but

this problem could, of course, be dealt with by filtering or equalization.

A somewhat more complex problem, touched upon in Chapters 1 and 2, is related to the fact that in clinical practice the wide spectrum of cardiac-cycle vibration seems to be treated, not as an organic whole, but piecemeal in terms of apexcardiograms, phonocardiograms, and audible sounds; moreover, facilities for apex- and phonocardiography are evidently unavailable in many hospitals.

It seems reasonable that the relative availability of a particular class of data should be expected to influence the development of techniques and skills for evaluation and interpretation of those data; but the fact remains that the element of unfamiliarity in converted heart sounds, especially those obtained from normally inaudible frequencies, presents a challenge to the medical observer who has had to expend a considerable amount of effort in learning to interpret those components of cardiac-cycle vibrations audible by conventional auscultation.

The reader is urged to listen with some care to the examples of unaltered and converted heart sounds contained in the tape cassette accompanying this dissertation (these are described in Appendix E). It is felt that this is much the best way, short of actual experimentation, to judge the potential usefulness of the technique and the practical difficulties involved in its application to clinical diagnostic practice.

C H A P T E R 4

REPRESENTATION OF ANALOG SIGNALS BY PIECEWISE-LINEAR APPROXIMATIONS

4.1 -- Piecewise-Linear Waveforms

In this chapter we will discuss the representation of an arbitrary time-domain signal by an approximation consisting entirely of conjunct linear segments; we will also explore some of the properties of such piecewise-linear (PL) waveforms, particularly in relation to the frequency domain. The only restriction we will place upon the signals to be approximated in this manner is that they and their first derivatives be finite.

A PL waveform, by its very nature, is completely defined by its breakpoints, the junctions between adjacent linear segments. Each breakpoint, in turn, may be completely defined by its ordinate value relative to an arbitrary reference level and some kind of time parameter. The actual choice of variables to be used depends to some degree upon the circumstances, as we shall see.

Fig. 4.1.1 shows a portion of a PL waveform, with each breakpoint circled for emphasis and identified by a pair of breakpoint parameters consisting of the breakpoint time and the corresponding ordinate value. An equivalent set of parameters, more suitable for computer storage because of the more limited range of the time variable, is the pair (τ_i, ν_i) , where the breakpoint time is replaced by the time elapsed since the previous breakpoint occurred.


We will show that once a PL approximation has been defined, i.e. once the breakpoint parameters are known, the Fourier transform or Fourier series may be calculated directly from the breakpoint parameters by summation of series of finite length without further approximation. Two additional quantities useful in this connection are shown in Fig. 4.1.1: the segment slope and the difference in the slopes of adjacent segments.

4.2 -- The Fourier Transform of a Time-Limited PL Function

Let $v(t)$ be a periodic piecewise-linear function of time having n finite breakpoints

$$(t_1, \nu_1) \dots (t_n, \nu_n)$$

where



SEGMENT DURATION: $\tau_k = t_{k+1} - t_k, \quad k = 1, 2, \dots$

SEGMENT SLOPE: $\alpha_k = \frac{v_{k+1} - v_k}{t_{k+1} - t_k} = \frac{v_{k+1} - v_k}{\tau_k}, \quad k = 1, 2, \dots$

SLOPE DIFFERENCE: $\beta_k = \alpha_k - \alpha_{k-1}, \quad k = 1, 2, \dots$

FIGURE 4.1.1

Notation used in discussing piecewise-linear (PL) waveforms.

$$v(t) = 0 \quad \text{for} \quad \left\{ \begin{array}{l} t \leq t_1 \\ t \geq t_n \end{array} \right\}$$

and

$$t_k > t_{k-1}, \quad k = 2, 3, \dots, n.$$

The segment slopes are given by (see Fig. 4.2.1)

$$\alpha_k = \left\{ \begin{array}{ll} \frac{v_{k+1} - v_k}{t_{k+1} - t_k}, & k = 1, 2, \dots, n-1 \\ 0, & k = 0, n \end{array} \right\}$$

The first derivative of $v(t)$ is piecewise constant, and has the property

$$\dot{v}(t) = \left\{ \begin{array}{ll} \alpha_k, & \left\{ \begin{array}{l} t_k < t < t_{k+1} \\ (k = 1, 2, \dots, n) \end{array} \right\} \\ 0, & \left\{ \begin{array}{l} t \leq t_1 \\ t \geq t_n \end{array} \right\} \end{array} \right\}.$$

The second derivative is a string of n impulses:

$$\ddot{v}(t) = \sum_{k=1}^n \beta_k \delta(t - t_k),$$

where




FIGURE 4.2.1

Successive differentiation of a time-limited PL waveform, showing notation used in deriving expressions for the Fourier transform in terms of the breakpoint parameters.

$$\beta_k = \alpha_k - \alpha_{k-1}, \quad k = 1, 2, \dots, n.$$

Beginning with the Fourier transform of a unit impulse,

$$\mathcal{F}[\delta(t-t_k)] = e^{-j\omega t_k},$$

we may write

$$\mathcal{F}[\ddot{v}(t)] = (j\omega)^2 \mathcal{F}[v(t)] = \sum_{k=1}^n \beta_k e^{-j\omega t_k}.$$

Using the integration property of the Fourier transform, we obtain

$$V(j\omega) = \mathcal{F}[v(t)] = -\frac{1}{\omega^2} \sum_{k=1}^n \beta_k e^{-j\omega t_k};$$

letting

$$X(\omega) = \operatorname{Re} V(j\omega) = -\frac{1}{\omega^2} \sum_{k=1}^n \beta_k \cos \omega t_k$$

and

$$Y(\omega) = \operatorname{Im} V(j\omega) = \frac{1}{\omega^2} \sum_{k=1}^n \beta_k \sin \omega t_k,$$

the amplitude spectrum is given by

$$\Psi(\omega) = |V(j\omega)| = \sqrt{[X(\omega)]^2 + [Y(\omega)]^2}$$

while the phase spectrum is given by

$$\varphi(\omega) = \arg V(j\omega) = \tan^{-1} \left(\frac{Y(\omega)}{X(\omega)} \right).$$

At zero frequency (dc), the Fourier transform may be found from the definition

$$F(j\omega) = \int_{-\infty}^{\infty} f(t) e^{-j\omega t} dt,$$

which implies that

$$F(0) = \int_{-\infty}^{\infty} f(t) dt.$$

For the case of a PL function this reduces to

$$F(0) = \frac{1}{2} \sum_{k=1}^{n-1} (V_k + V_{k+1}) (t_{k+1} - t_k).$$

4.3 -- The Fourier Series for a Periodic PL Function

Let $v(t)$ be a periodic piecewise-linear function of time described over one period by n finite breakpoints (see Fig. 4.3.1)

$$(t_1, v_1) \cdots (t_n, v_n)$$

where

$$v_n = v_1$$

and

$$v_{k+(n-1)j} = v_k, \quad \left\{ \begin{array}{l} k = 1, 2, \dots, n \\ j = 0, \pm 1, \pm 2, \dots \end{array} \right\}.$$

The segment slopes are

$$\alpha_k = \frac{v_{k+1} - v_k}{t_{k+1} - t_k}$$

for all integer values of k . Because of the periodicity of $v(t)$,

$$\alpha_{k+(n-1)j} = \alpha_k;$$

in particular,

$$\alpha_0 = \alpha_{n-1}$$

$$\alpha_n = \alpha_1.$$




FIGURE 4.3.1

Successive differentiation of a periodic PL waveform, showing notation used in deriving expressions for the Fourier-series coefficients in terms of the breakpoint parameters.

The first derivative of $v(t)$ is piecewise-constant with ordinates equal to the slopes of $v(t)$ (see Fig. 4.3.1):

$$\dot{v}(t) = \alpha_k, \quad t_k < t < t_{k+1}$$

while the second derivative is a periodic string of impulses:

$$\ddot{v}(t) = \sum_{k=-\infty}^{\infty} \beta_k \delta(t - t_k), \quad k = \text{integer}$$

where

$$\beta_k = \alpha_k - \alpha_{k-1}.$$

Expanding $v(t)$ in a Fourier series, we obtain

$$\ddot{v}(t) = \sum_{m=-\infty}^{\infty} c_m e^{jm\omega_0 t},$$

where

$$\omega_0 = \frac{2\pi}{T} = \frac{2\pi}{t_2 - t_1}$$

and

$$c_m = \frac{1}{T} \int_{t_1}^{t_1+T} \ddot{v}(t) e^{-jm\omega_0 t} dt$$

$$\begin{aligned}
 c_m &= \frac{1}{T} \int_{t_1}^{t_1+T} \left[\sum_{k=1}^n \beta_k \delta(t-t_k) \right] e^{-jm\omega_0 t} dt \\
 &= \frac{1}{T} \sum_{k=1}^n \beta_k e^{-jm\omega_0 t_k}
 \end{aligned}$$

Integrating $v(t)$ twice, term by term, we obtain

$$v(t) = -\frac{1}{m\omega_0^2} \sum_{m=-\infty}^{\infty} c_m e^{jm\omega_0 t}$$

i.e.,

$$v(t) = \sum_{m=-\infty}^{\infty} \gamma_m e^{jm\omega_0 t}$$

where

$$\gamma_m = \frac{-c_m}{m^2\omega_0^2} = \frac{-1}{m^2\omega_0^2} \frac{1}{T} \sum_{k=1}^n \beta_k e^{-jm\omega_0 t_k}$$

or

$$\begin{aligned}
 \gamma_m &= \frac{-T}{4\pi^2 m^2} \sum_{k=1}^n \beta_k e^{-jm\omega_0 t_k} \\
 &= \frac{t_1 - t_n}{4\pi^2 m^2} \sum_{k=1}^n \beta_k e^{-jm\omega_0 t_k}, \quad m \neq 0. \quad (4.3.1)
 \end{aligned}$$

The dc term is obtained by averaging $v(t)$ over one period:

$$\begin{aligned} \bar{v}_0 &= \langle v(t) \rangle = \frac{1}{T} \int_0^T v(t) dt \\ &= \frac{1}{T} \sum_{k=1}^{n-1} \left(\frac{v_k + v_{k+1}}{2} \right) (t_{k+1} - t_k) \end{aligned}$$

i.e.,

$$\bar{v}_0 = \frac{1}{2T} \sum_{k=1}^{n-1} (v_k + v_{k+1}) (t_{k+1} - t_k). \quad (4.3.2)$$

The cross-correlation function for two periodic signals having the same period (or regarded as having been normalized to the same period) may be calculated from their respective Fourier coefficients. If n harmonics have been calculated for each function, then

$$R_{12}(\tau) \approx A_{01} A_{02} + \sum_{k=1}^n \frac{A_{k1} A_{k2}}{2} \cos(k\omega_0\tau + P_{k2} - P_{k1}),$$

where the A's and P's represent, respectively, the amplitudes and phases of the Fourier coefficients.

To illustrate the principal results of this section, we will consider an example of a PL approximation obtained from a normal electrocardiogram by means of adaptive PL sampling as described in Chapter 5. Fig. 4.3.2

shows three cycles of the original waveform together with the corresponding PL approximation. In this example, the approximation, although reasonably good, has been sufficiently limited in terms of slope resolution (discussed in Section 5.3) to be recognizable as an approximation; note in particular the first complete cycle shown.

The breakpoint times and ordinates defining one cycle of the PL waveform are shown in Table 4.3.1, together with the slope of each linear segment. A rough sketch of one cycle, obtained using an interactive computer terminal, is shown in Fig. 4.3.3; the sketch is produced by linear interpolation between breakpoints followed by 21-level quantization of the times. The horizontal axis is drawn at zero ordinate value.

The results of Fourier-series analysis of one cycle of the PL waveform, performed as we have outlined above, are shown in Table 4.3.2. For each harmonic, the amplitude and phase are obtained by calculating the real and imaginary parts, respectively, of the complex Fourier coefficients, each of which has been determined directly from the breakpoint parameters by means of Eqs. (4.3.1) and (4.3.2). The amplitude spectrum, plotted by the




FIGURE 4.3.2

- (a) Three cycles of a normal electrocardiogram;
- (b) PL approximation of the waveform shown in (a).

TABLE 4.3.1

Breakpoint parameters for one cycle of a periodic PL waveform.

NORMAL ECG, PL APPROXIMATION

```

*****
BREAKPOINT   BREAKPOINT   BREAKPOINT   SEGMENT
  NUMBER      TIME           ORDINATE     SLOPE
*****
  1          8.600E-01    8.200E+01   -6.750E+03
  2          8.667E-01    3.700E+01   -1.005E+04
  3          8.733E-01   -3.000E+01   -9.000E+03
  4          8.800E-01   -9.000E+01   -4.350E+03
  5          8.867E-01   -1.190E+02    3.000E+02
  6          8.933E-01   -1.170E+02    3.000E+03
  7          9.000E-01   -9.700E+01    3.750E+03
  8          9.067E-01   -7.200E+01    2.550E+03
  9          9.133E-01   -5.500E+01    1.200E+03
 10         9.200E-01   -4.700E+01   -3.375E+02
 11         9.467E-01   -5.600E+01    1.895E+02
 12         1.073E+00   -3.200E+01    4.500E+02
 13         1.100E+00   -2.000E+01    7.500E+02
 14         1.113E+00   -1.000E+01    4.500E+02
 15         1.120E+00   -7.000E+00    7.501E+02
 16         1.127E+00   -2.000E+00    3.000E+02
 17         1.147E+00    4.000E+00   -7.500E+01
 18         1.160E+00    3.000E+00   -6.000E+02
 19         1.167E+00   -1.000E+00   -3.000E+02
 20         1.173E+00   -3.000E+00   -8.727E+02
 21         1.247E+00   -6.700E+01   -4.500E+02
 22         1.260E+00   -7.300E+01   -7.500E+01
 23         1.273E+00   -7.400E+01   -3.000E+02
 24         1.280E+00   -7.600E+01    1.500E+02
 25         1.293E+00   -7.400E+01    2.786E+02
 26         1.340E+00   -6.100E+01   -1.227E+02
 27         1.413E+00   -7.000E+01    5.357E+01
 28         1.507E+00   -6.500E+01    3.000E+01
 29         1.573E+00   -6.300E+01    3.000E+02
 30         1.580E+00   -6.100E+01    2.400E+02
 31         1.647E+00   -4.500E+01    0.0
 32         1.653E+00   -4.500E+01    3.000E+02
 33         1.660E+00   -4.300E+01    5.250E+02
 34         1.673E+00   -3.600E+01    1.500E+02
 35         1.687E+00   -3.400E+01   -4.500E+02
 36         1.693E+00   -3.700E+01   -9.000E+02
 37         1.713E+00   -5.500E+01   -6.000E+02
 38         1.720E+00   -5.900E+01   -7.500E+01
 39         1.733E+00   -6.000E+01    1.125E+02
 40         1.760E+00   -5.700E+01    4.500E+02
 41         1.767E+00   -5.400E+01    1.350E+03
 42         1.773E+00   -4.500E+01    3.750E+03
 43         1.780E+00   -2.000E+01    6.450E+03
 44         1.787E+00    2.300E+01    7.649E+03
 45         1.793E+00    7.400E+01    1.200E+03
 46         1.800E+00    0.200E+01

```




FIGURE 4.3.3

Computer sketch of one cycle of a periodic PL waveform.

TABLE 4.3.2

Harmonic components of a periodic PL waveform.

NORMAL ECG, PL APPROXIMATION

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	4.825E+01	26.08	-11.674	180.00
1	1.064E+00	1.850E+02	100.00	0.0	-177.65
2	2.128E+00	5.211E+01	28.16	-11.006	174.11
3	3.191E+00	1.725E+01	9.33	-20.607	163.48
4	4.255E+00	1.019E+01	5.50	-25.185	-163.85
5	5.319E+00	7.605E+00	4.11	-27.722	163.51
6	6.383E+00	1.394E+00	0.75	-42.462	95.46
7	7.447E+00	1.413E+00	0.76	-42.340	111.42
8	8.511E+00	2.193E+00	1.19	-38.524	74.44
9	9.574E+00	2.878E+00	1.56	-36.162	67.74
10	1.064E+01	3.017E+00	1.63	-35.754	52.85
11	1.170E+01	2.727E+00	1.47	-36.630	51.11
12	1.277E+01	2.572E+00	1.39	-37.139	43.22
13	1.383E+01	2.513E+00	1.36	-37.340	29.16
14	1.489E+01	2.542E+00	1.37	-37.242	17.58
15	1.596E+01	2.597E+00	1.40	-37.053	12.99
16	1.702E+01	2.543E+00	1.37	-37.236	6.27
17	1.809E+01	2.189E+00	1.18	-38.538	6.62
18	1.915E+01	2.012E+00	1.09	-39.270	3.22
19	2.021E+01	1.546E+00	0.84	-41.562	0.56
20	2.128E+01	1.196E+00	0.65	-43.788	-3.53

```
*****
```

computer in Fig. 4.3.4, shows that most of the energy in this signal is associated with the first few harmonics. A somewhat more detailed picture of the relative harmonic amplitudes may be obtained by logarithmic (dB) scaling of the amplitudes prior to plotting, as seen in Fig. 4.3.5. The phase spectrum for this example is shown in Fig. 4.3.6.

4.4 -- Sampling-Rate Considerations in Adaptive PL Sampling

In sampling a time-domain signal at a fixed rate, one of the major practical considerations is the relationship between the sampling rate and the range of frequencies present in the signal being sampled. In general, the sampling frequency must be made greater than twice the highest frequency contained in the original signal in order to ensure that the signal may be recovered from the samples; i.e., the sampling rate must exceed the Nyquist frequency.

This restriction does not apply to adaptive PL sampling. In order to show this, we feel that it is appropriate to review the problem as it applies to fixed-rate sampling, although it is by no means a new result.

NORMAL ECG, PL APPROXIMATION




FIGURE 4.3.6

Phase spectrum of a periodic PL waveform.

Let us assume first that a continuous, time-limited, aperiodic signal $v(t)$ is sampled by multiplication in the time domain by a periodic sampling function $s(t)$ consisting of a string of unit impulses. This multiplication yields a discrete signal $\bar{v}(t)$:

$$\bar{v}(t) = v(t) s(t) , \quad (4.4.1)$$

where

$$s(t) = \sum_{k=-\infty}^{\infty} \delta(t - kT)$$

and T is the sampling period, the reciprocal of the sampling frequency f_s . The sampling signal $s(t)$ is periodic; hence it may be represented by a Fourier series

$$s(t) = \sum_{n=-\infty}^{\infty} c_n e^{jn\omega_s t}$$

$$(\omega_s = 2\pi f_s = \frac{2\pi}{T})$$

where the Fourier coefficients c_n are given by

$$c_n = \frac{1}{T} \int_{-T/2}^{T/2} \delta(t - kT) e^{-jn\omega_s t} dt .$$

The integrand in this expression is zero over the entire range of integration except at $t = 0$, so that

$$c_n = \frac{1}{T} e^0 = \frac{1}{T} ,$$

whence

$$s(t) = \frac{1}{T} \sum_{n=-\infty}^{\infty} e^{jn\omega_s t} .$$

We now take the Fourier transform of both sides of this equation to obtain

$$S(f) = \frac{1}{T} \sum_{n=-\infty}^{\infty} \delta(f - nf_s) . \quad (4.4.2)$$

Evidently $\bar{v}(t)$ has the Fourier transform

$$\begin{aligned} \bar{V}(f) &= V(f) * S(f) \\ &= \frac{1}{T} V(f) * \sum_{n=-\infty}^{\infty} \delta(f - nf_s) , \end{aligned} \quad (4.4.3)$$

as may be seen from Eqs. (4.4.1) and (4.4.2). But the convolution of an arbitrary function $f(\xi)$ with an impulse function $\delta(\xi - \eta)$ has the property

$$f(\xi) * \delta(\xi - \eta) = f(\xi - \eta) , \quad (4.4.4)$$

which may be applied to Eq. (4.4.3) to yield

$$\bar{V}(f) = \frac{1}{T} \sum_{n=-\infty}^{\infty} V(f - n f_s). \quad (4.4.5)$$

Eq. (4.4.5) shows that the sampled signal $v(t)$ is periodic in frequency with "period" f_s . The only stipulation with which we began is that $\bar{v}(t)$ is a discrete periodic function of time, defined only at certain instants spaced uniformly along the time axis; evidently this stipulation is sufficient to account for the frequency-domain periodicity of the sampled signal. It is, of course, this periodicity that makes it necessary that f_s be set high enough to keep the spectral "cycles" from overlapping, with concomitant distortion in the time domain. It remains, however, to show that the problem arises, not because $\bar{v}(t)$ is discrete, but because $s(t)$ is periodic.

More generally, fixed-rate sampling is accomplished by multiplying a continuous signal $v(t)$ in the time domain by a periodic (not necessarily discrete) sampling signal $s(t)$ of arbitrary waveshape and period $T = 1/f$. We can generalize the argument presented above as follows.

Since $s(t)$ is periodic, it may be represented by a Fourier series:

$$s(t) = \sum_{n=-\infty}^{\infty} c_n e^{jn\omega_s t}$$

where the Fourier coefficients c_n depend upon the waveshape; they may be written

$$c_n = \frac{1}{T} \int_{-T/2}^{T/2} s(t) e^{-jn\omega_s t} dt$$

Taking the Fourier transform of our expression for $s(t)$, we obtain

$$S(f) = \sum_{n=-\infty}^{\infty} c_n \delta(f - nf_s), \quad (4.4.6)$$

and multiplication of $v(t)$ by $s(t)$ in the time domain to yield $\bar{v}(t)$ is equivalent in the frequency domain to

$$\begin{aligned} \bar{V}(f) &= V(f) * S(f) \\ &= \sum_{n=-\infty}^{\infty} c_n V(f) \delta(f - f_s), \end{aligned} \quad (4.4.7)$$

where we have used Eq. (4.4.4) in Eq. (4.4.6) to obtain the last expression. But Eq. (4.4.7) may be rewritten as

$$\bar{V}(f) = \sum_{n=-\infty}^{\infty} c_n V(f - nf_s). \quad (4.4.8)$$

Here $\bar{V}(f)$ is a periodic function of frequency with period f_s , "modulated" by the Fourier coefficients c_n , which will in general be different for different values of n . Although $\bar{V}(f)$ is no longer strictly periodic in frequency, there remains a pattern of periodic recurrences of the spectrum of the original waveform, so that the sampling frequency f_s must still be made high enough to prevent these patterns from overlapping. This is evidently necessitated by the fact that $s(t)$ is a periodic function of time.

A PL approximation to a time-limited aperiodic signal is not obtained merely by multiplying that signal in the time domain by a periodic sampling signal. Although breakpoint parameters are defined for discrete instants, those instants are in general not uniformly spaced along the time axis, and aliasing does not occur even when they are.

We have shown that for a time-limited PL function the Fourier transform may be written

$$V(j\omega) = \left\{ \begin{array}{l} -\frac{1}{\omega^2} \sum_{k=1}^n \beta_k e^{-j\omega t_k}, \quad \omega \neq 0 \\ \frac{1}{2} \sum_{k=1}^n (\nu_k + \nu_{k+1})(t_{k+1} - t_k), \quad \omega = 0 \end{array} \right\} \quad (4.4.9)$$

where the notation is defined in Section 2 of this chapter. But Eq. (4.4.9) shows that $V(j\omega)$ is a continuous function of frequency, without periodic recurrences of a spectral pattern. Hence there is no problem of potential spectral overlap to be dealt with. This means that if we wish to choose PL breakpoints from among samples of the original waveform taken at a uniform rate, that rate must be high enough to ensure adequate resolution of rapid changes in the waveform -- but the average number of breakpoints per unit of time required to define an acceptable PL approximation and the uniform sampling frequency itself may be considerably lower than the Nyquist frequency. Moreover, there is no need to ensure strict bandlimiting of the sampled signal by preliminary filtering, nor to filter the PL waveform when reconstructing the approximation from stored breakpoint parameters. Finally, in computing Fourier transforms of aperiodic functions, time-domain multiplication by a "window" function is desirable just as in the case of fixed-rate (uniform) sampling; but here it is needed only to avoid unnecessary spectral distortion when a portion of a PL record is to be isolated. It is not required in order to avoid spectral overlap resulting from strict time-limiting (the "leakage" effect).

A special problem arises when a PL record containing relatively few breakpoints is to be windowed. For example, a triangular pulse defined by three breakpoints, such as that shown in Fig. 4.4.1, will emerge unaltered when multiplied by a window function centered about its peak. It may therefore be desirable to generate some superfluous breakpoints by linear interpolation between the essential breakpoints prior to point-by-point multiplication by the window function.

In the case of a periodic PL function, we have found the complex Fourier series to be

$$v(t) = \sum_{m=-\infty}^{\infty} \gamma_m e^{jm\omega_0 t}, \quad (4.4.10)$$

where the Fourier coefficients are given by

$$\gamma_m = \left\{ \begin{array}{l} \frac{1}{2T} \sum_{k=1}^{n-1} (v_k + v_{k+1}) (t_{k+1} - t_k), \quad m=0 \\ \frac{-T}{4\pi^2 m^2} \sum_{k=1}^n \beta_k e^{-jm\omega_0 t_k}, \quad m \neq 0 \end{array} \right\}$$

and the notation is defined in Section 3 of this chapter. Taking the Fourier transform of the PL waveform described by Eq. (4.4.10), we obtain

$$V(f) = \sum_{m=-\infty}^{\infty} \gamma_m \delta(f - mf_0)$$

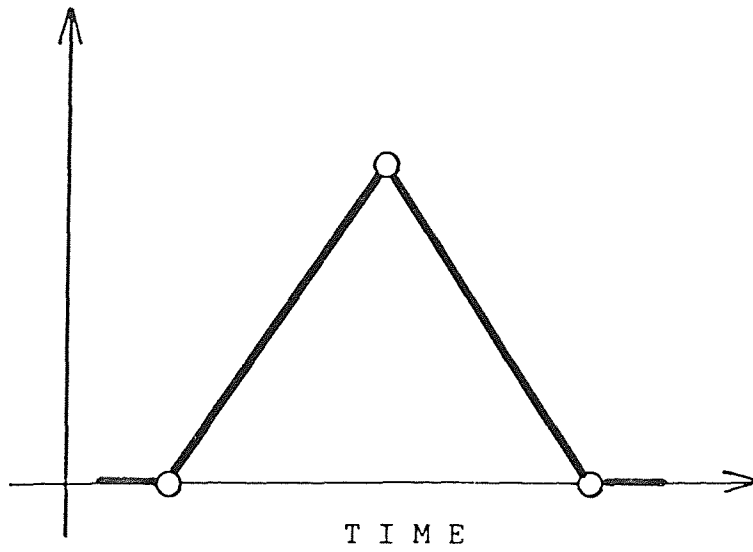


FIGURE 4.4.1

A triangular pulse defined by three PL breakpoints.

This equation shows that $V(f)$ is nonzero only for discrete frequencies. For $f_0 > 0$, there is no repetitive spectral pattern and no possibility of spectral overlap. Again, the average number of breakpoints per unit of time needed to define an adequate PL approximation may be considerably lower than the Nyquist frequency; and similarly, preliminary anti-aliasing filtering and low-pass filtering during reconstruction of the PL waveform are unnecessary.

To illustrate the results of this analysis, we offer several examples of PL representations of a normal electrocardiogram, obtained by adaptive PL sampling as described in Chapter 5 (see Figs. 4.4.2-4). These examples have successively fewer breakpoints per unit of time and consequently successively poorer resolution of the original signal. Each PL example is shown together with the original signal and two additional signals obtained from the original by fixed-rate sampling (see Fig. 4.1.2) at a frequency equal to the average breakpoint-per-second rate for the PL approximation. In each case, the first of these waveforms was obtained without anti-aliasing precautions, while the second was obtained by low-pass filtering of the original signal at half the sampling




FIGURE 4.4.2

- (a) Normal electrocardiogram, original waveform.
- (b) PL approximation, 105 breakpoints/cycle (111.8 breakpoints/sec).
- (c) Uniform sampling, $f_s = 111.8$ Hz.
- (d) Uniform sampling, $f_s = 111.8$ Hz, filtered at 55.9 Hz.




FIGURE 4.4.3

- (a) Normal electrocardiogram, original waveform.
- (b) PL approximation, 46 breakpoints/cycle (48.9 breakpoints/sec).
- (c) Uniform sampling, $f_s = 48.9$ Hz.
- (d) Uniform sampling, $f_s = 48.9$ Hz, filtered at 24.5 Hz.




FIGURE 4.4.4

- (a) Normal electrocardiogram, original waveform.
- (b) PL approximation, 31 breakpoints/cycle (33.0 breakpoints/sec).
- (c) Uniform sampling, $f_s = 33.0$ Hz.
- (d) Uniform sampling, $f_s = 33.0$ Hz, filtered at 16.5 Hz.

frequency (to satisfy the Nyquist condition) and low-pass filtering of the resultant piecewise-constant sampled waveform during reconstruction at the same frequency. The preliminary filtering was performed using a 4-pole analog Butterworth filter, while the reconstruction filtering was implemented by means of a 12-pole zero-phase-shift digital Butterworth filter.

The points to be emphasized in considering these examples are (1) that distortion in the fixed-rate waveforms is avoided only at the expense of excluding potentially useful information by filtering prior to sampling, and (2) that a useful PL approximation may be defined by fewer data than a fixed-rate waveform of comparable quality.

It is, perhaps, appropriate to point out that the examples shown in Figs. 4.4.1-3 were chosen in large measure because of our special interest in heart sounds and electrocardiograms; we feel that they offer useful insight into some of the possible applications of our technique. However, most of the energy in the normal electrocardiogram is contained in the first few harmonics, as we have seen in Section 3 of this chapter.

Therefore, spectral overlap produces less waveform distortion than would be the case if the high-frequency components were more prominent.

If a more dramatic example is desired, one might consider a 1-Hz linear ramp. This waveform may be completely defined by two PL breakpoints per cycle, or two breakpoints per second in the present case. If this signal is to be sampled uniformly at 2 Hz, anti-aliasing filtering at half the sampling frequency will necessarily remove a great deal of the fundamental and all of the overtones of this signal. For reference, a breakpoint table, a one-cycle sketch, and an amplitude-spectrum plot are shown in Table 4.4.1 and Figs. 4.4.5-6. Fig. 4.4.7 shows how anti-aliasing precautions used in uniform sampling of this waveform at the same average rate limit the reconstructed waveform to the fundamental-frequency component.

TABLE 4.4.1

Breakpoint parameters completely defining one cycle of a piecewise-linear ramp waveform having two breakpoints per cycle.

1-HZ RAMP

BREAKPOINT NUMBER	BREAKPOINT TIME	BREAKPOINT ORDINATE	SEGMENT SLOPE
----------------------	--------------------	------------------------	------------------

1	0.0	0.0	1.333E+00
2	7.500E-01	1.000E+00	-4.000E+00
3	1.000E+00	0.0	




FIGURE 4.4.5

Computer sketch of the linear ramp waveform defined by the breakpoints shown in Table 4.4.1.




FIGURE 4.4.7

PL and uniform sampling of the ramp waveform shown in Fig. 4.4.5. (a) Original waveform; (b) PL approximation; (c) original waveform after anti-aliasing filtering at 1 Hz; (d) filtered waveform sampled at the Nyquist rate (2 Hz); (e) sampled waveform filtered at 1 Hz. All waveforms shown have been normalized to the same amplitude.

C H A P T E R 5

ADAPTIVE PIECEWISE-LINEAR SAMPLING OF ANALOG SIGNALS

5.1 -- Hybrid Implementation of Adaptive PL Sampling

In this chapter we will explore techniques and algorithms for automatic real-time determination of acceptable PL approximations of analog waveforms. Whatever the method used, the result of the PL sampling must be the storage in digital form of the breakpoint parameters which define the PL waveform. We should like to limit the memory requirements for breakpoint-parameter storage as sharply as possible, consistent with adequate resolution; moreover, the fewer breakpoints needed the more quickly we may calculate Fourier series and Fourier transforms from the breakpoint parameters.

What makes PL sampling adaptive is the fact that a satisfactory PL approximation may be fitted to an analog waveform with few breakpoints wherever that waveform is changing slowly, while more breakpoints per unit of time will be needed in regions where the waveform is changing rapidly. In other words, where the curvature of the waveform is high (in the region of a sharp peak, for

example), more breakpoints will be required than in regions where the waveform is not changing as rapidly and the curvature is lower. Where the curvature is zero, i.e. where the slope is constant, the waveform and its PL approximation are identical and no breakpoints at all are needed.

In more formal terms, no breakpoints will be required in regions where the slope is constant and the curvature, which is reflected in the second derivative of the waveform, is zero. The greater the curvature, the greater the rate at which the slope changes, and the greater the magnitude of the second derivative. Evidently the absolute value of the second derivative may be taken as an index of the appropriate sampling rate at any given time. Although there is some question as to whether a nonlinear relationship might not be preferable, we have chosen to make the sampling rate directly proportional to the magnitude of the second derivative in each of the methods discussed below.

Fig. 5.1.1 shows a schematic representation of a hybrid method for obtaining breakpoint parameters by adaptive PL sampling according to the approach we have




FIGURE 5.1.1.

A hybrid system for extracting PL-waveform breakpoint parameters by adaptive sampling.

just described. The magnitude of the second derivative is calculated by analog computation; it controls the frequency of a voltage-controlled oscillator (VCO) whose function is to generate the trigger pulses which activate the analog-to-digital converters which perform the actual sampling of ordinate value and time.

This diagram also shows a threshold comparator (a Schmitt trigger) whose function is to multiply the VCO control voltage by zero whenever the curvature is so small that nonzero values of the second derivative may be attributed to noise. This part of the system is not unlike the noise threshold control discussed in Chapter 3 and Appendices B and C in connection with FM conversion of heart sounds. (This feature was dispensed with in later work, since we find that it places an undesirable restriction on curvature resolution.)

The operation of this hybrid system may be visualized with reference to Fig. 5.1.2. In the example shown, part of the input waveform is already piecewise-linear, so that the system chooses breakpoints at the junctions between the linear segments. In the last portion of the input waveform, the nonzero curvature is




FIGURE 5.1.2

Four steps in a process for controlling variable-rate adaptive sampling of an analog signal in order to obtain appropriate breakpoint parameters for a piecewise-linear approximation.

reflected in the second derivative, whose magnitude decreases as the input waveform flattens out, decreasing the rate at which samples are taken as it approaches zero.

5.2 -- Analog Differentiation

If the second derivative is to be calculated by analog circuitry, considerable attention must be paid to the maintenance of acceptable signal-to-noise ratios. We have experimented with modifications of a differentiator circuit described by Graeme (Graeme 1973, 78) which approximates the derivative by summing an input signal and its time integral; this approach has been used in designing circuits for analog computers (Computer Handbook 1962, 2-28).

Fig. 5.2.1 shows a second-derivative circuit built from two integrators and a summing inverter. For simplicity in analysis we will consider all of the resistances equal, although it will be seen that in practice the circuit gain is quite low if this is done.

In order to evaluate the potential usefulness of this circuit, we first obtain its transfer function by




FIGURE 5.2.1

An analog circuit for computation of the second derivative of an input waveform.

analysis in the Laplace transform domain, using the notation

$$G = \frac{1}{R}$$

$$\tau = RC = \frac{C}{G} .$$

Assuming the current I_1 to be negligibly small, Kirchhoff's current rule may be applied at the inverting input to the first-stage amplifier to yield

$$V_1(s) = \frac{-V_{in}(s)}{1 + \tau s} . \quad (5.2.1)$$

Similarly, for the second stage we find that

$$V_2(s) = -\frac{V_{in}(s)}{1 + \tau s} - \frac{V_1(s)}{1 + \tau s} \quad (5.2.2)$$

while for the third stage, analysis yields

$$V_{out}(s) = -V_{in}(s) - V_1(s) - V_2(s) \quad (5.2.3)$$

We now combine Eqs. (5.2.1), (5.2.2), and (5.2.3) to obtain the transfer function

$$H(s) = \frac{V_{out}(s)}{V_{in}(s)} = \frac{-\tau^2 s^2}{(1 + \tau s)^2} . \quad (5.2.4)$$

This is indeed the transfer function of a double-differentiator provided that τs is small compared to unity. For examination of the steady-state behavior of the circuit at a single frequency, we replace the Laplace-transform variable s by $j\omega$ and define

$$\omega_0 = \frac{1}{\tau} = \frac{1}{RC}$$

whence Eq. (5.2.4) may be rewritten as

$$H(j\omega) = -\frac{1}{\omega_0^2} \left[\frac{\omega^2}{1 + j\left(\frac{\omega}{\omega_0}\right)} \right], \quad (5.2.5)$$

The Bode plot corresponding to this transfer function is shown in Fig. 5.2.2. It may be seen that for two decades of operation (from 10 Hz to 1 KHz, for example) a dynamic range of 80 dB must be available in the useful part of the curve, i.e., where ω is small compared to ω_0 .

5.3 -- Phase and Gain Error for the Analog Differentiator

If phase error is excessive, the pulses which trigger the A/D converter will be delayed. The phase of $H(j\omega)$ is given by

$$\varphi = 180^\circ - 2 \tan^{-1} \left(\frac{\omega}{\omega_0} \right),$$

so that the phase error for the circuit is of magnitude




FIGURE 5.2.2

Bode amplitude plot for the circuit of Fig. 5.2.1, whose transfer function is given in Eq. (5.2.5).

$$\epsilon = 2 \tan^{-1} \left(\frac{\omega}{\omega_0} \right) = 2 \tan^{-1} \left(\frac{f}{f_0} \right) . \quad (5.3.1)$$

Clearly, the greatest phase error occurs at the highest operating frequency. Solving Eq. (5.3.1) for the break frequency f_0 , we obtain

$$f_0 = \frac{f}{\tan (\epsilon/2)} . \quad (5.3.2)$$

For a maximum phase error of 5% at a maximum operating frequency of 1 KHz, for example, we may calculate the required break frequency as follows:

$$(0.05)(180^\circ) = 9.0^\circ$$

$$f_0 = \frac{1 \text{ KHz}}{\tan 4.5^\circ} = 12.7 \text{ KHz} .$$

From Eq. (5.2.5), the gain error associated with the circuit is seen to be

$$\eta = \frac{\omega^2}{\omega_0^2} - \frac{\omega^2}{\omega_0^2} \left[\frac{1}{1 + \frac{\omega^2}{\omega_0^2}} \right] .$$

Defining

$$X = \frac{\omega^2}{\omega_0^2} ,$$

this may be written as

$$\eta = \frac{x^2}{1+x} \quad (5.3.3)$$

Evidently η is zero at dc and increases monotonically with frequency. At the maximum operating frequency of 1 KHz specified for our example,

$$x = \left[\frac{1 \text{ KHz}}{12.7 \text{ KHz}} \right]^2 = 0.0062 \quad .$$

When this value is substituted in Eq. (5.3.3), η is found to be 3.8×10^{-5} . The gain at this frequency is found from Eq. (5.2.5) to be 6.2×10^{-3} , so that the gain error is about 0.61%. Evidently the critical factor in this instance is the phase error.

Phase and gain error are unaffected if the feedback resistor in the final stage is increased with respect to R in order to obtain added gain. Amplifier noise in the first two stages is integrated and thereby made to decrease with increasing frequency (Graeme 1973, 78). Signal-to-noise ratio problems remain, however, at low operating frequencies mainly because of the large dynamic range required. These problems may be alleviated by decreasing ω_0 , but only at the expense of increased phase error.

5.4 -- Modification of the Analog Circuit

Our experiments have shown that one way of overcoming the signal-to-noise problems inherent in the circuit described above is to replace it with a group of similar circuits, each having a different time constant, and arranged for parallel operation. At any given time, only one of the channels is actually connected to the circuit controlling the output to the sampling VCO. Fig. 5.4.1 shows a three-channel system built along these lines. The priority logic controlling the choice of channel is implemented by means of three comparators and three analog switches controlled by an appropriate combination of digital gates. Each comparator goes high as its input voltage passes a preselected threshold beyond which the noise level is considered acceptable. As the operating frequency increases, the channel having the lowest ω_0 approaches the point where its phase error becomes significant, but by that time the logic circuitry has disconnected it from the output, replacing it by the second channel. Similarly, the third channel (having the highest ω_0) replaces the second channel as soon as its output rises above the threshold level as the frequency of operation continues to increase.




FIGURE 5.4.1

A three-channel system using the circuit of Fig. 5.2.1 to generate a control voltage proportional to the magnitude of the second derivative of the input signal. The full-wave rectifier output is negative-going, the threshold voltage is negative, and the output is negative-going. The feedback resistance R controls the proportionality relating the sampling rate to the curvature of the input signal.

In order to ensure smooth channel transitions, the resistors R_1 , R_2 , and R_3 must be reasonably stable and two of them must be carefully balanced; these resistances satisfy the relationship

$$R_1 > R_2 > R_3 .$$

The logic equations describing the switching from channel to channel are

$$\begin{aligned} S_1 &= D\bar{E}\bar{F} && \text{(low frequency)} \\ S_2 &= E\bar{F} && \text{(medium frequency)} \\ S_3 &= F && \text{(high frequency)} \end{aligned}$$

Small switching transients may be present in the output of this three-channel circuit even when high-speed comparators or Schmitt triggers are used. It must be borne in mind, however, that the output of this circuit serves only to control the frequency of sampling, so that the transients are normally imperceptible in the sampling pattern itself and therefore need not be filtered out.

5.5 -- A Microcomputer Algorithm for Adaptive piecewise-Linear Sampling

Let us assume that an arbitrary input signal $v(t)$ is sampled at a fixed rate f , yielding successive values v_i , $i = 1, 2, 3, \dots$. For simplicity, let us further assume for the purpose of this discussion that our A/D converter has 201 quantization levels, -100 through +100 (including zero). We define

$$v_i = i^{\text{th}} \text{ sample of } v(t), \quad i = 1, 2, 3, \dots$$

$$(d1)_i = \frac{v_i}{2} - \frac{v_{i-1}}{2}, \quad i = 2, 3, 4, \dots$$


$$(d2)_i = (d1)_i - (d1)_{i-1}, \quad i = 3, 4, 5, \dots$$

The quantities $d1$ and $d2$ are not quite first and second differences because of the division by 2 inherent in our definition of $(d1)_i$; the necessity for this division will emerge presently.

The quantities $(d2)_i$ may be used in determining an appropriate piecewise-linear (PL) approximation to $v(t)$. How the choice of appropriate breakpoints may be made by the system can perhaps be seen most clearly by considering a series of examples.

CASE 1: dc


i	v	(d1)	(d2)
1	10	--	--
2	10	0	--
3	10	0	0
4	10	0	0
5	10	0	0



In this instance the total lack of activity in d2 will cause the system to do nothing but count the number of samples taken since the last breakpoint. When the counter overflows, the previous sample is taken as a breakpoint whose ordinate value is 10 and whose sample count since the preceding breakpoint is the maximum value which the counter register can contain; in our case, this value is FFFF hexadecimal, or 65536. The counter is now reset to zero, and the process begins again.

CASE 2: Finite Impulse

i	v	$(d1)$	$(d2)$
1	-60	--	--
2	-60	0	--
3	-60	0	0
4	-60	0	0
5	80	70	70
6	-60	-70	-140
7	-60	0	70
8	-60	0	0



Each of the nonzero values of $d2$ may be used as an indication that the previous sample must be taken as a breakpoint, resulting in the approximation shown dashed in the sketch. Clearly the impulse would be ignored by the system if it did not coincide with a sampling instant, and the input signal would be seen by the system as dc (Case 1). The implication, of course, is that if spikes of finite amplitude and duration are anticipated in $v(t)$ we will risk missing them or seeing them as having substantially reduced amplitude if the sampling rate is too low.


This case also serves to illustrate the necessity for division of the ordinate values by 2 in calculating $d1$. Since $d2$ will be calculated from two values of $d1$, each of these values must be valid; $d1$ cannot be permitted to go out of range of our registers. In this example, we require

$$-100 \leq (d1)_i \leq 100, \quad i = 1, 2, 3, \dots$$

On the other hand, $d2$ has been used so far only as an index of abrupt changes in curvature reflected by changes in $d1$. It may be permitted to overflow, provided that the system is aware of its having done so and is able to interpret that event as an indication that the previous sample is to be taken as a breakpoint.

CASE 3: Step Function


i	v	(d1)	(d2)
1	10	--	--
2	10	0	--
3	10	0	0
4	10	0	0
5	90	40	40
6	90	0	-40
7	90	0	0
8	90	0	0




Here again, each of the nonzero values of $d2$ may be used as an indication that the previous sample is to be taken as a breakpoint, resulting in the approximation shown dashed in the sketch.

CASE 4: Triangle

i	v	$(d1)$	$(d2)$
1	20	--	--
2	30	5	--
3	40	5	0
4	50	5	0
5	40	-5	-10
6	30	-5	0
7	20	-5	0



i	v	$(d1)$	$(d2)$
1	20	--	--
2	40	10	--
3	60	10	0
4	52	-4	-14
5	32	-10	-6
6	12	-10	0



In the first example shown above, the fortuitous coincidence of the waveform peak with a sampling instant


results in a single nonzero value of d_2 (for $i = 5$), indicating that v is to be taken as a breakpoint. When the waveform peak does not correspond with a sampling instant, as in the second example, two nonzero values of d_2 are obtained; these are equal if the peak is midway between two adjacent samples and the slope is of equal amplitude on either side of it. The approximation obtained when each sample preceding a nonzero value of d_2 is taken as a breakpoint is shown for this example as a dashed line connecting the points at $i = 3$ and $i = 4$.

It may be argued that, provided the sample rate is high enough, the breakpoint at $i = 4$ is unnecessary in this example, and that an appropriate approximation would be that obtained by choosing only the higher of the two adjacent breakpoint candidates, and connecting that point (at $i = 3$) with the point at $i = 6$ as shown in the sketch. We have in fact implemented a routine for performing editing of this kind, relying for its operation on the fact that this is the only situation in which breakpoint candidates occur as adjacent samples whose corresponding values of d_2 are of the same sign (see Cases 2 and 3, above). If that sign is negative we choose the point having the higher value; if it is positive we choose the one whose value is lower.

Unfortunately, we find that for our experimental system the saving of breakpoints achieved in this manner is offset by considerably increased length and complexity of the sampling program and concomitantly lower speed; this routine has therefore not been incorporated in the program used in obtaining the PL approximations to be considered below.

CASE 5: Smooth Curve

i	v	$(d1)$	$(d2)$
1	-76	--	--
2	-36	20	--
3	-2	17	-3
4	26	14	-3
5	48	11	-3
6	64	8	-3
7	76	6	-2
8	84	4	-1
9	88	2	-1
10	90	1	0
11	90	0	0
12	90	0	0



We would like to have greater breakpoint density in our PL approximation where the curvature of $v(t)$ is greater, or more precisely, in those regions where the slope is changing more rapidly as reflected in larger magnitudes of $d2$. One way of accomplishing this might be to sum the absolute values of the $(d2)_i$ as they are ob-

tained, taking a breakpoint (and beginning a new summation) whenever the sum exceeds a suitably-chosen curvature-resolution constant. In practice this approach is found to be unsatisfactory, for reasons which we will explore below.

In none of the examples given thus far has consideration been paid to the quantization error inherent in the sampling of $v(t)$ or to the error resulting from the finite-precision integer arithmetic by means of which our system calculates d_1 and d_2 . In the first three cases considered, this error creates no problems in choosing breakpoints because d_2 is zero except immediately following finite discontinuities in an otherwise constant signal.

The worst case is found to be the triangle wave (Case 4), since in practice we obtain a random "bobble" of small positive and negative values of d_1 and d_2 in regions where the slope of $v(t)$ is in fact constant, resulting from quantization error and finite-precision arithmetic as we have described. Consequently, summation of the absolute values of d_2 eventually causes unnecessary breakpoints to be taken in the interior of a region

where $v(t)$ has constant slope, even if the resolution constant described above is set so high that curvature resolution is degraded to the point where the system does not recognize the discontinuity in the slope at the peak of the waveform.

The solution to this problem lies in the fact that the troublesome "bobble" is indeed random. In practice, we sum the $d2$ values themselves, instead of their absolute values. In this way the "bobble" tends to cancel out in the long run, as we might expect, and the summation goes out of a suitably-chosen range of values symmetric about zero only when systematic curvature is present in the input waveform, or when it is driven abruptly out of range by a sample following a discontinuity in the slope. Our experimental program, which operates on this basis, actually performs two sets of threshold comparisons: one for $d2$ itself, and the other for the summation of $d2$ values. In this way, resolution constants can be chosen separately for slope-discontinuity (corner) detection and curvature resolution.

An algorithm by which our system can choose breakpoints according to the approach we have described is

shown below in simplified flowchart form, and again as a flowgram (Karp 1978). In the flowgram, A and B have been used to represent the corner-detection and curvature-resolution constants respectively.

5.6 -- Slope Detection in PL Sampling

Experiment has shown that the use of an 8-bit microprocessor such as the Z-80 on which our experimental system is based places a potentially serious restriction on the usefulness of that system for PL sampling according to the algorithm we have just described. This restriction may be explained with reference to Fig. 5.6.1.

For the portion of the waveform shown on a trace of duration t_c , the (constant) slope may be expressed as

$$k = \frac{m}{t_c} = \frac{n}{t_s} \quad (\text{levels/sec}) .$$

At a given sampling rate f_s , the minimum slope detectable from adjacent samples will be denoted by

$$k_{\min} = \frac{1}{t_s} = f_s .$$




FIGURE 5.5.1
Flowchart for
adaptive PL sampling
algorithm.





FIGURE 5.5.2

Flowgram for
adaptive PL sampling
algorithm.



M = total quantization levels
 m = integer, $0 \leq m \leq M$
 n = integer, $0 \leq n \leq M$
 t_s = sampling period
 t_c = trace period
 f_s = sampling frequency = $1/t_s$

FIGURE 5.6.1

Notation used in discussing slope resolution.

For detection of a minimal slope of P percent of our system's maximum of M quantization levels over a trace of duration t_c seconds, the sampling rate may not exceed

$$(f_s)_{\max} = \frac{PM}{100 t_c}$$

For example, if

$$P = 10$$

$$t_c = 1 \text{ sec}$$

$$M = 2^8 = 256 \text{ levels,}$$

the maximum sampling rate allowed is 25.6 Hz, which is clearly inadequate for satisfactory resolution of high-frequency events. With a 12-bit system (A/D converter and microcomputer), however, we have

$$P = 10$$

$$t_c = 1 \text{ sec}$$

$$M = 2^{12} = 4096 \text{ levels,}$$

and the maximum allowable sampling rate is 409.6 Hz, which will do nicely for sampling electrocardiograms in real time, for example; a sampling rate often used for

uniform sampling of electrocardiograms in Holter-tape systems is 400 hz.

An alternative way of improving slope resolution would be to modify the algorithm so that when two adjacent samples fall within the same quantization interval, the next sample is considered to see whether there is in fact a nonzero slope too small to be detected from adjacent samples. Such a modification would involve several decision-making steps not included in our present algorithm; for example, in considering groups of three samples it would be important to distinguish between slopes of low magnitude, where the second and third samples might differ by one quantization level, and changes in curvature which would cause the second and third samples to differ by more than one level. Moreover, keeping parameters for three points at a time available to the decision-making routines would involve operations with data in RAM rather than in CPU registers, which would impose significant speed limitations in addition to those already inherent in our sampling program. Evidently, with an M-level system operating on the basis of our present algorithm, the best resolution we can hope for, as a function of the sampling rate, is

$$P = 100 \left(\frac{f_s t_c}{M} \right).$$

For our system (8 bits, 256 levels),

$$P = \frac{100 f_s t_c}{256} = 0.391 f_s t_c$$

For a one-second trace, the order of magnitude of the trace duration necessary to display one cardiac cycle,

f (Hz)	P
100	39.1
200	78.1
300	117.2
400	156.3
500	195.3

This may be visualized in terms of a one-second example of constant slope, as shown in Fig. 5.6.2.




FIGURE 5.6.2

An example of the minimum detectable slope as a function of sampling rate for an 8-bit system.

These considerations are of especially great concern to us when we recall that the algorithm we have described involves division of the slopes by 2, using integer arithmetic. This means that our slope-detection sensitivity is immediately degraded by a factor of 2, which we are not in a position to afford. Clearly some method is required of improving the slope-detection resolution of our system, even if it is only to be used to demonstrate the feasibility of obtaining reasonable PL approximations by means of the algorithm we have developed.

It is tempting to consider using two of the seven A/D channels in our experimental system to achieve 16-bit resolution, since

$$2^{16} = 65,536$$

and 65,536 levels would allow us to sample at over 6.5 khz for 10-percent slope resolution as described above, assuming that we can find a way around the need for division of the slopes by 2. A block diagram of a circuit which might permit this is shown in Fig. 5.6.3.




FIGURE 5.6.3

A scheme for achieving 16-bit ordinate resolution by means of three conversions. The low-order byte is obtained by scaling and converting the difference between the input signal and its quantized output.

Unfortunately, several considerations make this approach impractical. First, the three conversion operations required (A/D or D/A) must be performed serially; each takes about $11 \mu\text{s}$, which is not serious. Second, the two bytes (high-order and low order) must be combined by 16-bit addition the preparation for which depends on the sign of each byte; a routine to achieve this would take on the order of 200 machine cycles, or $100 \mu\text{s}$, for each sample, which is more serious. Third, all of the processing of the 16-bit signed ordinate values -- the calculation of d_1 and d_2 , the summation of the d_2 values, and the comparisons with the four threshold constants -- would require 16-bit arithmetic where we now use 8-bit arithmetic; some of these operations would require special subroutines since they are not already part of the system's repertoire of 16-bit operations. The result would be extremely unwieldy and result in a long (and slow) program.

We shall demonstrate, using the system at our disposal, that PL sampling is indeed practical and can result in substantial data compression even with the limitations inherent in an 8-bit system such as ours; but we must also conclude that practical implementation of

our techniques for sampling of input signals of any kind should really be carried out using a 12- or 16-bit A/D converter and microprocessor.

5.7 -- Modifications to the Algorithm for Adaptive PL Sampling

We have pointed out that the second difference, or a fraction thereof, can be used in estimating waveform curvature only if it determined from valid first differences or fractions of first differences. We have also shown that halving the sampled ordinate values prior to subtracting in order to obtain d_1 prevents d_1 from overflowing and insures the validity of the quantities from which we will obtain d_2 .

We have also pointed out that a significant and unfortunate drawback in this "adjustment" is the fact that the division by two doubles the minimum slope our system is able to detect, and therefore restricts the rate at which the input waveform may be sampled. Since that rate is already restricted by the limitations on slope resolution inherent in 8-bit A/D conversion and 8-bit arithmetic, we must seek a way around this problem if we are to use our 8-bit system effectively.

We find that a satisfactory alternative to division by two is to limit the first difference instead of halving it. That is, we calculate the first difference by subtracting each sample from its successor, and replace the resulting d_1 value by the appropriate maximum or minimum value consistent with 8-bit twos-complement representation each time the subtraction causes register overflow. Naturally, such "correction" really means that the d_1 values obtained in the event of an overflow are incorrect; however, (1) overflow will occur very seldom because the range of slopes which can cause it is narrow, and (2) the worst that can happen when overflow does occur is that we will take one additional valid breakpoint which we could have done without.

The program modifications necessary to achieve the adjustments we have described may be represented in flow-gram form as follows:

```

:
:
SAMPLE WAVEFORM; STORE vi
|
CALCULATE d1 = vi - vi-1
|
NO OVERFLOW? d1 IS VALID
|
d1 < 0? SET d1 = MAX
|
d1 > 0? SET d1 = MIN
|
←
STORE d1
:
:

```

The backwards-looking corrections in the event of an overflow are easily explained in terms of the method by which overflow is detected by the system, which uses twos-complement representation in performing binary arithmetic. For example, if in adding two numbers of the same sign the sign of the result is found to change, the register in which the sum is formed has cycled past its maximum binary value (7F hexadecimal in our case) and has entered the range of negative numbers (80 through FF hexadecimal) in which the most significant binary digit is a one. Overflow can occur in the hypothetical (-100 through 100) system used in the examples given earlier only when two adjacent samples are of different sign and sufficiently remote from each other (i.e., more than 100

units apart). Since subtraction of these values is equivalent to negation of the second followed by addition, the overflow-detection scheme described above will indicate whether the result is out of range, in which case it must be replaced by the value of greatest available magnitude and opposite sign.

With this modification, all of the values for $d1$ and $d2$ given in the tables for the examples of the unmodified algorithm must be doubled. The only case in which overflow occurs in $d1$ is that of the finite impulse; with adjustments as required, the table becomes:

i	v_i	$(d1)_i$		$(d2)_i$
1	-60	--		--
2	-60	0		--
3	-60	0		0
4	-60	0		0
5	80	140	--> 100	100
6	-60	-140	--> -100	-200
7	-60	0		100
8	-60	0		0

In this instance, the system "corrects" the overflow values of d_1 for $i = 5$ and $i = 6$. The overflow in d_2 for $i = 6$ need not be altered; as before, it is sufficient that the system be aware that overflow has occurred in order that it recognize the breakpoint at $i = 5$.

Table 5.7.1 is a complete assembly-language listing of the program actually used in our work to obtain PL waveforms by adaptive sampling. This program utilizes the algorithm described in Section 5.5, with the modifications we have just discussed. Fig. 5.7.1 illustrates the results of PL sampling by means of this program. The test waveform shown was generated by summation of an asymmetric rectangular waveform with its derivative, using the circuit shown in Fig. 5.7.2.

TABLE 5.7.1

Complete Z-80 assembly-language listing of a microcomputer program for adaptive PL sampling, using the algorithm described in Section 5.5 with the modifications given in Section 5.7.

```

E000 ALAN D. BERNSTEIN
E001 DEPARTMENT OF PHYSICS, RUTGERS NCAS
E002
E003
E004 THIS PROGRAM PERFORMS PIECEWISE-LINEAR
E005 SAMPLING IN REAL TIME. INPUT IS VIA
E006 A/D CHANNEL 1, AND A STARTING-TRIGGER
E007 PULSE MAY BE PROVIDED VIA A/D CHANNEL
E008 2. IF THE INPUT SIGNAL IS CONNECTED TO
E009 BOTH CHANNELS, SAMPLING WILL BEGIN AT
E010 THE ONSET OF THE INPUT SIGNAL.
E011
E012 REGISTER INITIALIZATION:
E013 ENTER WITH
E014 IX <--- RAM START ADDRESS
E015 IY <--- RAM BYTE COUNT
E016 B <--- CORNER THRESHOLD CONSTANT
E017 C <--- CURVATURE THRESHOLD CONSTANT
E018
E019
E020 LD SP,STACK #POSITION STACK
E021 PUSH BC #SAVE STARTING ENVIRONMENT
E022 PUSH IY #ON STACK
E023 PUSH IX
E024 POP HL #HL <--- BYTE COUNT
E025 LD (IY),H #WRITE BYTE COUNT
E026 LD (IY+1),L #AT START OF RAM
E027 INC IY #POINT TO START OF
E028 INC IY #DATA-STORAGE AREA
E029 LD DE,-5
E030 ADD HL,DE #HL <--- HL-5
E031 LD IX,0 #CLEAR CYCLE COUNT
E032 PUSH IY #SAVE RAM POINTER
E033 LD IY,TABLE #POINT TO TABLE
E034 LD A,B #A <--- CORNER THRESHOLD
E035 LD (IY),A #STORE IN TABLE

```

```

0010 #
0020 #
0030 #
0040 #
0050 #
0060 #
0070 #
0080 #
0090 #
0100 #
0110 #
0120 #
0130 #
0140 #
0150 #
0160 #
0170 #
0180 #
0190 #
0200
0210
0220
0230
0240
0250
0260
0270
0280
0290
0300
0310
0320
0330
0340
0350

```

```

31 FF E2
C5
FD E5
DD E5
E1
FD 74 00
FD 75 01
FD 23
FD 23
11 FR FF
19
DD 21 00 00
FD E5
FD 21 00 E3
78
FD 77 00

```

E025	ED	44	0360	NEG				!A <---- -A
E027	3C		0370	INC				!ADJUST
E028	FD	77 01	0380	LD				!STORE LOWER THRESHOLD
E02B	79		0390	LD				!A <---- CURVATURE THRESHOLD
E02C	FD	77 02	0400	LD				!STORE IN TABLE
E02F	ED	44	0410	NEG				!A <---- -A
E031	3C		0420	INC				!ADJUST
E032	FD	77 03	0430	LD				!STORE LOWER THRESHOLD
E035	FD	E1	0440	POP				!RESTORE RAM POINTER
E037	11	00 00	0450	LD				!CLEAR OLD (Y,D1)
E03A	D9		0460	EXX				!RESTORE BC, DE, HL
E03B	3E	03	0470	LD				!INITIALIZE STROBE
E03D	D3	6E	0480	OUT				!AND TIMER GATE
E03F	DB	1A	0490	IN				!READ TRIGGER CHANNEL ----
E041	DB	1A	0500	IN				!FIRST SAMPLE INVALID
E043	FE	00	0510	CP				!V >= 0?
E045	F2	4A E0	0520	JF				!IF SO SKIP NEXT STEP
E048	ED	44	0530	NEG				!A <---- -A
E04A	FE	16	0540	CP				!V >= 0.5 VOLTS?
E04C	F2	60 E0	0550	JF				!IF SO BEGIN SAMPLING
E04F	DB	18	0560	IN				!ELSE EXAMINE KEYBOARD
E051	CB	7F	0570	BIT				!KEY PRESSED?
E053	CA	3F E0	0580	JF				!IF NOT LOOP BACK
E056	CD	61 84	0590	CALL				!IF SO GET CHARACTER
E059	CB	BF	0600	RES				!RESET SIGN BIT
E05B	FE	53	0610	CP				!"S" FOR "START"?
E05D	C2	3F E0	0620	JF				!IF NOT LOOP BACK
E060	3E	02	0630	LD				!READY TO START ----
E062	D3	6E	0640	OUT				!OPEN TIMER GATE
E064			0650					
E064			0660					!BREAKPOINT--STORAGE MODULE
E064			0670					
E064	97		0680	BREAK:	SUB	A		!A <---- 0
E065	D3	6E	0690	OUT		TARBC,A		!INITIATE STROBE PULSE
E067	DD	E5	0700	PUSH		IX		!TRANSFER CYCLE COUNT

E069	C1	0710	POP	BC	‡ TO BC
E06A	FD 70 00	0720	LD	(IY)‡B	‡HIGH-ORDER BYTE TO RAM
E06D	FD 71 01	0730	LD	(IY+1)‡C	‡LOW-ORDER BYTE TO RAM
E070	D9	0740	EXX		‡GET BC‡ DE‡ HL‡
E071	FD 72 02	0750	LD	(IY+2)‡D	‡PREVIOUS SAMPLE TO RAM
E074	0E 00	0760	LD	C‡0	‡INITIALIZE D2 SUMMATION
E076	D9	0770	EXX		‡GET BC‡ DE‡ HL
E077	11 03 00	0780	LD	DE‡3	‡DE ‹--- 3
E07A	FD 19	0790	ADD	IY‡DE	‡INCREMENT RAM POINTER BY 3
E07C	A7	0800	AND	A	‡RESET CARRY FLAG
E07D	ED 52	0810	SBC	HL‡DE	‡DECREMENT BYTE COUNT BY 3
E07F	DA F9 E0	0820	JP	C‡FINISH	‡DONE? GO TO RESTORE MODULE
E082	DD 21 00 00	0830	LD	IX‡0	‡INITIALIZE CYCLE COUNT
E086	C3 E8 E0	0840	JP	WAIT1	‡AND AWAIT SAMPLE TRIGGER
E089		0850	‡		
E089		0860	‡SAMPLE--TEST MODULE		
E089	3E 02	0870	‡		
E089	D3 6E	0880	TEST‡	A‡2	‡A ‹--- 2
E08B	D9	0890	LD	TARBC‡A	‡TERMINATE STROBE PULSE
E08D	EB	0900	OUT		‡GET BC‡ DE‡ HL‡
E08E	DB 19	0910	EXX	DE‡HL	‡DE‡ ‹--- OLD (Y‡D1)
E08F	67	0920	EX	A‡AD1	‡TAKE INPUT SAMPLE
E091	92	0930	IN	H‡A	‡H‡ ‹--- NEW Y
E092	E2 A0 E0	0940	LD	D	‡A ‹--- NEW D1
E093	FA 9E E0	0950	SUB	FO‡GOOD	‡D1 VALID‡ BRANCH
E096	3E 80	0960	JP	M‡OVER	‡OVERFLOW‡ NEGATIVE? BRANCH
E099	C3 A0 E0	0970	JP	A‡80H	‡CORRECT A TO MINIMUM
E09B	3E 7F	0980	LD	GOOD	‡D1 NOW VALID‡ BRANCH
E09E	6F	0990	LD	A‡7FH	‡CORRECT A TO MINIMUM
E0A0	93	1000	LD	L‡A	‡L ‹--- NEW D1
E0A1	FD E5	1010	SUB	E	‡A ‹--- D2
E0A2	FD 21 00 E3	1020	PUSH	IY	‡SAVE RAM POINTER
E0A4	EA D3 E0	1030	LD	IY‡TABLE	‡POINT TO CONSTANT TABLE
E0A8	FD BE 00	1040	JP	PE‡EXIT1	‡D2 OV? HAVE BREAKPOINT
E0AB		1050	CP	(IY)	‡COMPARE W. UPPER THRESH.

```

EOAE F2 D3 E0
EOB1 FD RE 01
EOB4 FA D3 E0
EOB7 81
EOB8 FD RE 02
EOBB F2 D3 E0
EOBE FD RE 03
EOC1 FA D3 E0
EOC4 4F
EOC5 FD E1
EOC7 D9
EOC8 11 01 00
EOCB DD 19
EOCD DA E1 E0
EOD0 C3 E8 E0
EOD3 D9
EOD4 FD E1
EOD6 11 01 00
EOD9 DD 19
EODB DA E1 E0
EODE C3 64 E0
EOE1 DD 21 FF FF
EOE5 C3 64 E0
EOES
EOEB
EOEB
EOEB DB 6E
EOEA CB 47
EOEC CA E8 E0
EOEF DB 6E
EOF1 CB 47
EOF3 C2 EF E0
EOF6 C3 89 E0
EOF9
EOF9

1060 $SIGN +? HAVE BREAKPOINT
1070 $COMPARE W. LOWER THRESH.
1080 $SIGN -? HAVE BREAKPOINT
1090 $A <---- UPDATED D2 SUMMATION
1100 $COMPARE W. UPPER THRESH.
1110 $SIGN +? HAVE BREAKPOINT
1120 $COMPARE W. LOWER THRESH.
1130 $SIGN -? HAVE BREAKPOINT
1140 $STORE D2 SUMMATION
1150 $RESTORE RAM POINTER
1160 $GET BC, DE, HL
1170 $DE <---- 1
1180 $INCREMENT CYCLE COUNT BY 1
1190 $OVERFLOW? HAVE BREAKPOINT
1200 $NO BREAKPOINT? WAIT
1210 EXIT1:
1220 $GET BC, DE, HL
1230 $RESTORE RAM POINTER
1240 $DE <---- 1
1250 $INCREMENT CYCLE COUNT BY 1
1260 $OVERFLOW? HAVE BREAKPOINT
1270 $GO TO STORAGE MODULE
1280 $IX <---- FFFF HEXADECIMAL
1290 $GO TO STORAGE MODULE
1300 $EXTERNAL TRIGGER MODULE
1310 $
1320 WAIT1:
1330 $A <---- STATUS WORD
1340 $EXAMINE SAMPLING CLOCK
1350 $LOW? SAMPLE IT AGAIN
1360 $A <---- STATUS WORD
1370 $EXAMINE SAMPLING CLOCK
1380 $HIGH? AWAIT TRAILING EDGE
1390 $GO TO SAMPLE--TEST MODULE
1400 $RESTORATION OF STARTING ENVIRONMENT

```

```

E0F9      3E 03      1410  †
E0FB      D3 6E      1420  †FINISH:
E0FD      DD E1      1430  †LD
E0FF      FD E1      1440  †OUT
E101      C1        1450  †POP
E102      C3 48 1F  1460  †POP
E105                1470  †POP
                        1480  †JF
                        1490  †MONIT
E2FF      0E2FFH   1490  †EQU
E300      0E300H   1500  †EQU
006E      6EH     1510  †EQU
001A      1AH     1520  †EQU
0018      18H     1530  †EQU
8461      8461H   1540  †EQU
0019      19H     1550  †EQU
1F48      8008    1560  †EQU
E105                1570  †MONITOR (WARM--START ENTRY)
E105                †TOP OF STACK
E105                †TABLE OF THRESHOLDS
E105                †CONTROL/STATUS PORT
E105                †TRIGGER PORT
E105                †KEYBOARD PORT
E105                †KEYBOARD--READ SUBROUTINE
E105                †SIGNAL PORT
E105                †MONITOR (WARM--START ENTRY)

E105                †REGISTER ASSIGNMENTS DURING PROCESSING
E105                †
E105                †HL' -- NEW (Y,D1)
E105                †DE' -- OLD (Y,D1)
E105                †C' -- D2 SUMMATION
E105                †HL -- BYTE COUNT
E105                †DE -- INCREMENTS/DECREMENTS
E105                †IX -- SAMPLING--CYCLE COUNTER
E105                †IY -- RAM POINTER

```

CROSS REFERENCE

AD1	0019	0920	
AD2	001A	0490	0500
BREAK	E064	1260	1280
EXIT1	E0D3	1040	1060 1080 1110 1130
EXIT2	E0E1	1190	1250
FINISH	E0F9	0820	
GCHAR	8461	0590	
GOOD	E0A0	0950	0980
INIT1	E03F	0580	0620
INIT2	E04A	0520	
INIT3	E060	0550	
KB	0018	0560	
MONIT	1F48	1470	
OVER	E09E	0960	
STACK	E2FF	0200	
TABLE	E300	0330	1030
TARBC	006E	0480	0640 0690 0890 1320 1350 1430
TEST	E089	1380	
WAIT1	E0E8	0840	1200 1340
WAIT2	E0EF	1370	




FIGURE 5.7.1

- (a) Test waveform: uniform sampling, 264 samples/cycle.
- (b) Same, PL sampling, 46 samples/cycle.
- (c) Same, PL sampling, 15 samples/cycle.
- (d) Same, PL sampling, 11 samples/cycle.

C H A P T E R 6

APPLICATIONS OF PL SAMPLING

6.1 -- Reconstruction of PL Signals

We have suggested that adaptive PL sampling techniques may permit adequate representation of analog signals by approximations which may be described by fewer data than might be necessary in conventional uniform sampling. We have also shown that anti-aliasing precautions are unnecessary in PL sampling. We now address ourselves to the question of reconstructing PL waveforms from stored breakpoint parameters in order that the PL approximation may be used in the time domain.

Essentially, we require a waveform-synthesis technique which will permit the real-time regeneration of PL waveforms by interpolating between adjacent breakpoints. We should also like to be able to perform time-axis scaling as desired, to facilitate the display or plotting of the waveforms. Two approaches suggest themselves.

If a digital computer of sufficient speed and floating-point arithmetic capability is available, the PL




FIGURE 5.7.2

Circuit used to generate the test waveform shown in Fig. 5.7.1. Summation of the rectangular waveform with its derivative takes place in the first stage; the feedback capacitors are included for stability.

waveform may be (in effect) sampled at a high uniform rate by the computer, and then converted into an analog staircase approximation of the PL waveform which can be low-pass filtered to remove the spurious high-frequency components generated by the uniform sampling. The result, of course, is an approximation of an approximation. We have in fact used a variation of this technique, although not in real time, to obtain computer sketches of filed PL waveforms and to convert PL waveforms by uniform sampling so that digital filtering could be performed.

A second approach requires somewhat more thought but less sophisticated hardware. It involves a hybrid reconstruction scheme in which interpolation between breakpoints is achieved by analog computation under the control of a digital microcomputer. This method has the advantage that it can be incorporated fairly easily into other systems, even portable ones.

The PL breakpoint parameters stored in digital memory are (see Fig. 6.1.1)




FIGURE 6.1.1

A portion of a PL waveform showing the notation used in discussing hybrid interpolation between breakpoints.

- $v_i = i^{\text{th}}$ breakpoint voltage
 $\tau_i =$ number of sampling cycles elapsed
 between $(i-1)^{\text{th}}$ and i^{th} breakpoints
 $i = 2, 3, 4, \dots$

For use in our experiments, we have designed a hybrid system which performs analog interpolation under microcomputer control. It is based on the the Analog Devices AD531JD, which (together with auxiliary active and passive components) can be made to yield the voltage-transfer function

$$V_{\text{out}} = \frac{(x_1 - x_2) y}{z}, \quad (6.1.1)$$

where x_1 , x_2 , y , and z are input voltages. The interpolation equation to be synthesized is evidently (see Fig. 6.1.1)

$$v(t_i \leq t \leq t_{i+1}) = v_i + \left(\frac{v_{i+1} - v_i}{\tau_{i+1}} \right) (f_s)(t - t_i), \quad (6.1.2)$$

where f_s is the sampling rate used to obtain the approximation and $(t - t_i)$ is a unit ramp beginning at $t = t_i$. This may also be expressed as

$$v(t_i \leq t \leq t_{i+1}) = v_i + \left(\frac{v_{i+1} - v_i}{\tau_{i+1}} \right) \mathcal{R}(t), \quad (6.1.3)$$

where $R(t)$ is the scaled ramp function

$$R(t_i \leq t \leq t_{i+1}) = - (V_c)(t - t_i)$$

and V_c is a negative-going voltage equal in magnitude to f_s .

One of the restrictions associated with the AD531JD is that the input z in Eq. (6.1.1) must be non-negative. This restriction is met in our system by having the computer supply the quantity

$$- \tau_{i+1}$$

which is then inverted and shifted by an operational amplifier prior to conversion to a current by the auxiliary circuitry associated with the AD531JD.

The adaptive PL sampling program discussed in Section 5.7 stores time intervals between breakpoints as unsigned 16-bit sampling-cycle counts. Having only 8-bit D/A converters, we might convert these values to corresponding analog voltages by using two D/A channels as shown in Fig. 6.1.2. If τ''_{i+1} and τ'_{i+1} represent, respectively, the high- and low-order bytes of τ_{i+1} , then

$$\tau_{i+1} = 256 \tau''_{i+1} + \tau'_{i+1} \quad (6.1.4)$$




FIGURE 6.1.2

Combination of high- and low-order bytes of v_{i+1} using two 8-bit D/A channels. The scaling constant k is $(22K\Omega/1.024M\Omega) = 2.15 \times 10^{-3}$, as may be seen from the diagram and from Eq. (6.1.3).

In practice we must choose R so that the amplifier's output remains within the active range for all values of τ_{i+1} . The extrema to be anticipated are 65536 (FFFF hexadecimal), corresponding to approximately 2.54 V, and zero, corresponding to about -2.54 V. Thus R must be on the order of 19 k Ω for an output swing of ± 12 V.

An alternative (which we have implemented) is to resolve the time-interval information into 8 bits before conversion to analog form. This is done by shifting each of the interval counts stored in memory to the right, one bit at a time, until the high-order byte is zero in every case. Then only the low-order byte is converted to an analog voltage.

To illustrate the practical considerations involved in deciding how to deal with the time-interval data, we offer two examples. In sampling electrocardiograms with a sampling-clock frequency of 300 hz, the largest cycle count (FF hexadecimal) which will fit into a single 8-bit register or memory location is

$$\frac{256 \text{ cycles}}{300 \text{ Hz}} = 853 \text{ ms .}$$

since it is very unlikely that such a large interval would occur between electrocardiogram breakpoints obtained from a living subject, 8-bit resolution is apparently adequate for this application.

Such is not the case if we are to sample normal speech, or recorded speech slowed to quarter-speed and sampled with a sampling-clock frequency of 3 khz (which would correspond to a generous 12 khz in real time). In these circumstances, 8-bit resolution limits us to an interval of

$$\frac{256 \text{ cycles}}{3000 \text{ Hz}} = 85.3 \text{ ms,}$$

i.e., less than a tenth of a second between breakpoints. Since the pauses between words in normal speech may last several seconds, 8-bit resolution is inadequate unless we shift the data as described above until so much low-order-byte information has been lost that the active portions of the speech record are seriously distorted as a result. In this instance, it is clearly advisable to retain all 16 bits, allowing an interval of as much as

$$\frac{65536 \text{ cycles}}{3000 \text{ Hz}} = 21.8 \text{ sec}$$

between breakpoints describing recorded speech at quarter-speed, corresponding to 5.46 seconds of real time.

If the active portions of a speech record may be handled with 8-bit resolution of time-interval data, we may effect a compromise that will allow us to get by with a single 8-bit D/A conversion for each interval. The time between breakpoints will be large enough to make an 8-bit register overflow only during pauses between words; but during every such pause, the numerator $v_{i+1} - v_i$ of the factor in Eq. (6.1.2) containing the interval count will be zero or approximately zero, so that the term containing this factor will vanish as it should in spite of the overflow, and the output will still correctly reflect the zero slope which gave rise to the cycle-count overflow in the first place.

The circuit shown in Fig. 6.1.3 is controlled by the microcomputer, which provides four analog signals (v_i , $-v_i$, v_{i+1} , and $-v_{i+1}$) and one digital control signal (V_m) for mode control of the ramp generator, IC4.




FIGURE 6.1.3 A microcomputer-controlled interpolator for synthesis of PL waveforms.

Since v_i is connected to a noninverting input of the AD531JD and v_{i+1} to an inverting input, an inverting adder (IC5) can provide $v(t)$ as described by Eq. (6.1.3) if the AD531JD output is connected to one of its inputs and $-v_i$ is connected to the other. IC1 performs the necessary scaling and shifting of $-v_{i+1}$ prior to conversion to a current by IC2 as prescribed in the Analog Devices specification sheet for the AD531JD. IC3 is a voltage follower whose function is to buffer the slope-control voltage V' .

The ramp generator (IC5) may be reset by raising the computer-supplied control voltage V_m sufficiently to turn on transistor Q1; the ramp is enabled when V_m is brought to zero after the computer has updated the four analog input voltages at each breakpoint. The diode in series with the mode-control voltage source prevents leakage of the timing current I into that source when the ramp is enabled (Jung 1974, 388-389). The timing current for this circuit is given by

$$I_t = \frac{V'}{R} ,$$

and the slope of the output is

$$\text{slope} = \frac{I_t}{C} = \frac{V'}{RC} . \quad (6.1.4)$$

Manual control of the slope is provided by means of the ten-turn potentiometer which determines V' . This feature facilitates relatively convenient calibration of the interpolation circuit for a wide range of time-interval scaling, so that the system may be used for plotting as well as reconstruction of signals at their original speed.

Primarily in order to protect the AD531JD, power for this circuit is provided by the microcomputer, ensuring that all input voltages will be zero whenever the power to the interpolator is turned off.

A drawback inherent in the circuit of Fig. 6.1.3 is that whenever the software delay controlling the rate of waveform generation is changed, the circuit must be recalibrated by readjustment of V_c (or by alteration of the ramp-generator time constant) to restore proper interpolation. This restriction may be removed by modifying the circuit as shown in Fig. 6.1.4, although the control software must be modified as well.

With this modification, the rate of reconstruction of the PL waveform is controlled by a hardware clock




FIGURE 6.1.4

A modification allowing the circuit of Fig. 6.1.3 to be used for PL-waveform reconstruction at a continuously variable rate without recalibration. Two series resistances provide coarse and fine adjustment of the ramp-generator time constant.

rather than by software delay. This clock consists, in our system, of a VCO whose output frequency f is related to its control voltage V_c by

$$f = (-0.947 V_c + 0.526) n \quad (6.1.5)$$

where n is the frequency-decade number and V ranges from -0.5 V to -10 V as the frequency increases by one decade.

It may be seen from Eq. (6.1.2) or (6.1.3) that if the interval count τ_{i+1} and the ramp are scaled by the same constant, the output of the interpolator will be unaltered except for time-axis scaling. To achieve this, we need a voltage V' (see Eq. (6.1.4)) which is directly proportional to the clock frequency; we obtain it from V_c by shifting to remove the constant term in Eq. (6.1.5). The reconstruction rate may now be varied continuously over an entire decade without recalibration of the interpolator.

The system we have described has been implemented and tested with various PL waveforms. In its present form, however, the speed limitations of the D/A converters and of the microcomputer controlling system opera-

tion precludes its use for the reconstruction of PL speech approximations, even at quarter-speed.

At each breakpoint, the computer must provide four analog and two digital outputs to permit the analog interpolation circuit to generate the next linear segment of the approximation. In our system these outputs are issued sequentially, a process requiring $117 \mu\text{s}$ for each breakpoint. In experimenting with PL sampling of speech limited to a bandwidth of 76-7600 Hz, this interval represents over 22% of the basic sampling period t_s (see Fig. 5.6.1) required for processing the speech samples at quarter-speed, assuming the basic sampling frequency is chosen equal to the highest frequency of interest, i.e., 1900 Hz. In PL sampling of speech, adequate approximation of consonant waveforms -- particularly sibilants -- involves numerous breakpoints separated by a single sampling period. These waveforms are seriously disfigured by the breakpoint transients, which are moreover of the same order of magnitude as the signal itself and cannot be removed by baseband filtering.

Evidently the limiting factor here is the speed of transfer of the breakpoint-parameter information from the

digital microcomputer to the analog interpolation circuit, rather than the slew rates of the operational amplifiers used in segment generation. Clearly, the problem could be alleviated by using a microcomputer and D/A conversion system of sufficiently high speed; we propose an alternative solution, which we hope to implement with our present system at a future date.

First, since our system has only seven D/A channels, the inversion of v_i will have to be done by means of an operational amplifier, thus leaving two groups of three channels each as shown in Fig. 6.1.5. Each group of three channels is sufficient to provide the three breakpoint-parameter voltages (corresponding to v_i , v_{i+1} , and $-v'_{i+1}$) needed for the synthesis of any given segment. While a segment is being generated, the microcomputer is free to update the three currently inactive channels with the parameters describing the next segment; when the next breakpoint arrives, the computer issues a single digital output, changing the state of the channel-group-select line, thus connecting the other group of channels to the analog circuit all at once, while at the same time resetting the ramp generator via another of the eight digital lines which make up the same digital port.

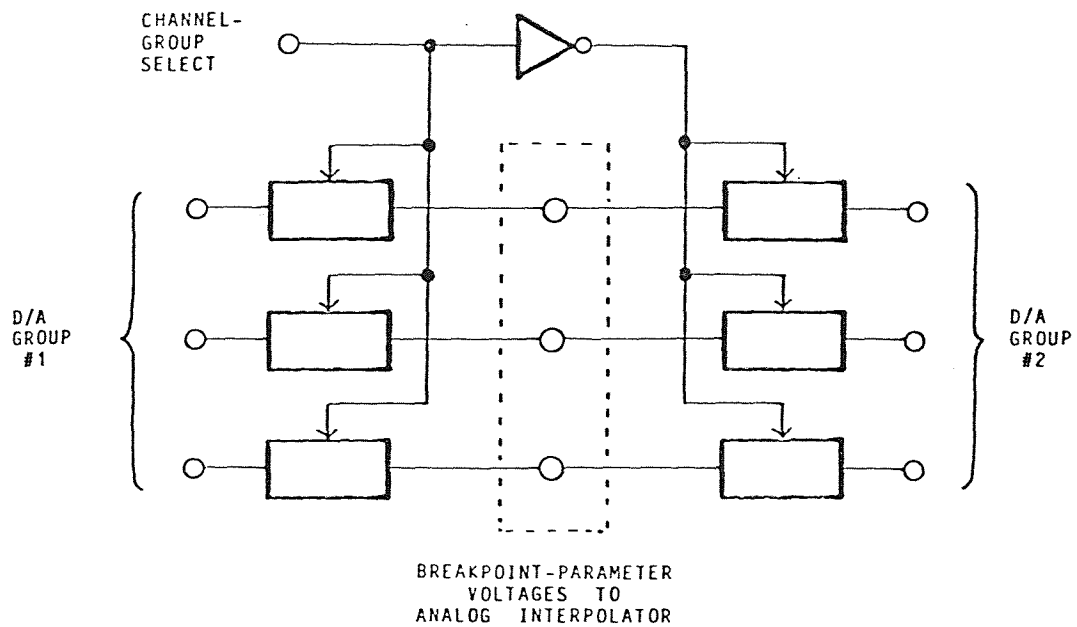


FIGURE 6.1.5

A scheme for parallel transfer of breakpoint-parameter voltages to an analog interpolator circuit. Six analog switches are used for alternate connection of two groups of D/A channels to the analog circuit.

With this arrangement, the limiting factor becomes the speed of the analog circuitry, particularly the speed with which the ramp generator can be reset; but it appears probable that the breakpoint transients could be effectively eliminated using this approach.

6.2 -- PL Sampling of Continuous Speech

The PL sampling technique we have described was developed primarily for use with biomedical signals such as heart sounds and electrocardiograms. Among its most important features is the ease with which frequency-domain information may be obtained directly from the PL breakpoint parameters. In the time domain, a significant characteristic of PL sampling is that it is a technique for obtaining waveform approximations without regard for the nature of the system producing the waveform, but with respect only to a geometrical feature of the waveform itself -- the curvature.

In a sense, we are trying to achieve a certain amount of data compression in the most general way possible, by considering only the shape of the data while ignoring their source. There exist various schemes for compression of speech data, each based upon some kind of

model by which speech may be described in terms of a few well-defined parameters; it is the values of these parameters which constitute the speech data to be transmitted or stored (Schafer and Rabiner, 1975).

Although our technique may be found useful primarily where frequency-domain information is desired, we feel that it is of interest to see whether any substantial data compression is possible, without a model for speech, by means of the same technique. We will therefore perform an order-of-magnitude calculation, starting with published information about the frequency distribution of the energy associated with speech, in order to estimate the degree of data compression which might be achieved by means of adaptive PL sampling.

A PL approximation to a published speech spectrum (Flanagan 1972, 163) is shown in Fig. 6.2.1; the bandwidth extends from 50 to 8000 Hz. With uniform sampling, satisfaction of the Nyquist sampling criterion requires a sampling frequency of more than 16 kHz. In order to find an estimate of the average sampling rate needed if PL sampling is used, we will convert this spectrum into a frequency probability density $P_f(f)$. The instantaneous PL sampling rate is given by




FIGURE 6.2.1

- (a) Spectrum for speech, PL approximation.
 (b) Equivalent frequency probability density.

$$N = \frac{\text{samples}}{\text{sec}} = \left(\frac{\text{samples}}{\text{cycle}} \right) \left(\frac{\text{cycles}}{\text{sec}} \right) ,$$

and its expected value is

$$\langle N \rangle = \langle nf \rangle = \int_{\text{all freq.}} nf P_f(f) df , \quad (6.2.1)$$

where n represents the number of samples needed per cycle at a given signal frequency f . We will allot 7 break-points per cycle for low frequencies and 3 for high frequencies for the purpose of our estimate; these figures have been found to be reasonable in using PL methods for waveform synthesis in electronic music (Bernstein and Cooper, 1976).

To find the probability density, we note first that the drop of 15 dB from 500 to 2000 Hz means that on a linear ordinate scale such as that of the lower sketch in Fig. 6.2.1,

$$20 \log \frac{A}{B} = 15 ,$$

or

$$\frac{A}{B} = 10^{15/20} ,$$

or

$$B = 0.1778 A . \quad (6.2.2)$$

But the curved portion of the lower sketch must be a decaying exponential function of frequency, i.e.

$$P_f(500 \leq f \leq 8000 \text{ Hz}) = A e^{-k(f-500)} . \quad (6.2.3)$$

At 2000 Hz, Eqs. (6.2.2) and (6.2.3) imply that

$$B = 0.1178A = A e^{-1500k} ,$$

from which k may be found to have the value 1.1513 ms . A is found from the normalization requirement

$$\int_{\text{all freq.}} P_f(f) df = A \left[450 + \int_{500}^{8000} e^{-k(f-500)} df \right] = 1 ,$$

where the first term in the square brackets arises from the area under the constant portion of the curve; upon evaluation of the integral and division by the quantity in the square brackets, we find that A must be $758.5 \mu\text{s}$. Thus

$$P_f(f) = \left\{ \begin{array}{ll} 758.5 \mu\text{s} , & 50 \leq f \leq 500 \text{ Hz} \\ 758 \times 10^{-6} e^{-(f-500)/1.1513 \mu\text{s}} , & 500 \leq f \leq 8000 \text{ Hz} \\ 0 , & \text{elsewhere} \end{array} \right\} \quad (6.2.4)$$

The expected value of the PL sampling rate for continuous speech is then found from Eqs. (6.2.1) and (6.2.4) to be

$$\begin{aligned}
 \langle N \rangle &= (7)(758.5 \times 10^{-6}) \int_{50}^{500} f \, df \\
 &\quad + (3)(758.5 \times 10^{-6}) \int_{500}^{8000} e^{-1.153 \times 10^{-3}(f-500)} \, df \\
 &= 1649 \text{ Hz}
 \end{aligned}
 \tag{6.2.5}$$

Eq. (6.2.5) says that for continuous speech we should expect the PL sampling rate to be on the order of one-tenth that required for uniform sampling. Each breakpoint, however, is defined by two numbers (time and ordinate value); hence the number of numbers required for PL sampling over a given period is roughly one-fifth that required for uniform sampling of continuous speech.

6.3 -- PL Sampling of Normal Speech

In this section we shall see that it is very much to our advantage that normal speech is far from continuous. During the pauses between words our system will take no samples at all, a circumstance which implies further reduction in the number of breakpoints needed to describe a given speech record.

In order to make the necessary adjustment in our estimated data-compression figure for continuous speech, we must relate the effective average sampling rate for normal speech to the corresponding sampling rate estimated for continuous speech. This relationship may be expressed as

$$r_{\text{effective}} = a r_{\text{continuous}} \quad (6.3.1)$$

where a is an "activity coefficient" defined by

$$a = \frac{T_1}{T_2} = \frac{\text{"active" time}}{\text{total record duration}} \quad (6.3.2)$$

In other words, for ordinary speech the data compression afforded by PL sampling is improved over that for continuous speech by a factor of $1/a$.

Using the method described in Appendix D, we have measured the activity coefficient defined by Eq. (6.3.2) for several examples of ordinary speech. The results of these measurements are shown in Table 6.3.1.

It is not particularly surprising that the activity coefficient for radio commercials was found to be greater

TABLE 6.3.1

Measurements of the activity coefficient for normal speech.

SOURCE	T_1 (sec)	T_2 (sec)	a	1/a
Radio commercial, popular music station	90.6	135.6	0.67	1.5
Radio commercial, WQXR, New York	75.8	127.9	0.59	1.7
Radio commercial, WQXR, New York	63.1	102.7	0.61	1.6
Radio news, WQXR, New York	67.7	168.1	0.40	2.5
Radio news, WQXR, New York	159.3	244.7	0.65	1.5
Radio news, WQXR, New York	185.7	289.6	0.64	1.6
Poetry reading	42.8	84.8	0.50	2.0
Prose reading	81.1	148.3	0.55	1.8
MEAN, COMMERCIALS			0.62	1.6
MEAN, NEWS			0.56	1.8
MEAN, GLOBAL			0.58	1.7

than the average of 0.58. We have the impression that the figure for news would be higher were it not for the pauses between news items. In any event, it appears that for ordinary speech the effective sampling rate expressed by Eq. (6.3.1) can be expected to be on the order of 60% of the corresponding estimate for continuous speech, or on the order of 1000 breakpoints per second -- about one-eighth the number of measurements required for uniform sampling at 16 kHz.

This matter will be explored further in Appendix E, which describes the examples of digitized speech to be found on the tape cassette accompanying the dissertation.

6.4 -- PL Sampling of Electrocardiograms

We have already given several examples of PL sampling of electrocardiograms (see Fig. 4.3.2) and of Fourier-series calculations from the PL breakpoint parameters (see Figs. 4.3.2 through 4.3.6). In this section we will offer some tentative observations regarding the application of such calculations.

The possible applications of our technique to electrocardiography seem, at this stage, to fall into three

categories. First, the results of our experiments suggest that the PL technique might be used to advantage for time-domain storage and retrieval of electrocardiogram waveforms as an alternative to digitization by uniform sampling. We have considered this possibility to some degree in Chapter 5, within the limitations of our 8-bit A/D converter and 8-bit microcomputer system.

A second category is related to the ease with which the Fourier series may be determined from the PL breakpoint parameters, providing two frequency-domain relationships -- amplitude and phase as functions of frequency -- as alternatives to the representation of voltage as a function of time. It must be remembered that the skin potentials measured in electrocardiography result from potential differences occurring in a three-dimensional structure. A complete set of electrocardiograms measured on a single subject often consists of twelve separate plots, differing only in the locations of the points on the surface of the body between which the potential differences are measured. It is tempting to consider the possibility that the separation of amplitude and phase characteristics in the

frequency domain might provide as complete a picture of the electrical activity of the heart in fewer displays (and with fewer data). Moreover, one cannot help but wonder if the alternative representation might not prove easier to interpret in some circumstances, either by human observers or by machine. If useful information can be obtained, for example, from the first 20 harmonics of an electrocardiographic measurement, it is encouraging to recall that the amplitude spectrum is completely defined in that case by only 21 numbers (twenty harmonic amplitudes plus the pulse rate, or fundamental frequency; the dc level appears to be universally regarded as unimportant), and the phase spectrum by another 20. Clearly, the storage of electrocardiographic information in this form is not uneconomical.

The third category is that of machine-aided cardiac diagnosis and monitoring of cardiac function. Here it seems plausible that the convenient conversion to the frequency domain might prove useful. To begin with, in the frequency domain the pulse rate is isolated, in a way, from the remaining characteristics of the signal. Examination of the harmonic amplitudes, for example, may be regarded from this point of view as examination of an

electrocardiogram which has been normalized with respect to the pulse rate.

The Fourier series represents a quantitative measurement, and suggests a means of quantitative comparison of anomalous electrocardiograms with stored "standards" following normalization with respect to fundamental frequency (a trivial operation with Fourier series), or with previous measurements made on the same subject, or with successive cycles of an ongoing electrocardiographic measurement intended to monitor changes in cardiac function.

The use of electrocardiogram spectra for purposes of diagnosis and monitoring can be approached in at least two ways. We can try to identify spectral characteristics or changes corresponding to characteristics or changes in the electrocardiogram waveform which have come to be recognized as indicative of various disorders or conditions of the heart. Alternatively, we can try to establish relationships between the spectra themselves and cardiac conditions or disorders, without the intermediate consideration of the electrocardiographic waveform. We will return to this matter in Chapter 7.

To illustrate this point (and, perhaps, to hint at its apparent complexity) we offer a series of examples drawn from two electrocardiogram tracings, the normal waveform shown in Fig. 4.3.2, and the abnormal waveform shown in Fig. 6.4.1. The latter tracing is that of an electrocardiogram characteristic of supraventricular rhythm, a condition in which the electrical impulse which triggers contraction of the ventricles does not come from the usual location.

The amplitude and phase spectra shown in Figs. 6.4.2 through 6.4.11 are related to the first two cycles shown in Fig. 6.4.1. It will be seen that the major difference between the first two cycles occurs just prior to the R-waves marked in the figure. Differential spectra, obtained by term-by-term subtraction of complex Fourier series in the frequency domain, have been used to compare these two cycles with a normal electrocardiogram and with each other. For this purpose the fundamental frequencies have been normalized to unity.

The most striking difference between the amplitude spectra of Figs. 6.4.2 and 6.4.3 appears to be in the second harmonic; the same is true for the corresponding




FIGURE 6.4.1

Normal and abnormal electrocardiograms. (a) Normal waveform as shown in Fig. 4.3.2; (b) abnormal waveform, obtained by PL sampling, corresponding to a supraventricular rhythm in which ventricular contraction is not triggered from its usual location. The first two R waves are labelled.

phase spectra of Figs. 6.4.4 and 6.4.5. At present writing, this is all we can say; whether one can monitor changes in this particular condition merely by measuring the second harmonic remains very much to be seen.

The differential amplitude spectra of Figs. 6.4.6. and 6.4.7, comparing the first and second cycles of Fig. 6.4.1 with a normal electrocardiogram, look very much like the normal spectrum itself (see, for example, Fig. 6.4.18). But this may be misleading, partly because of the difference in the overall amplitudes of the two signals; the normal electrocardiogram is the stronger signal in this case. Moreover, all of our spectrum plots are scaled to fit in the same page area. Again, the best we can do is to speculate; should two signals be normalized before comparisons of this sort? And if so, should they be normalized to a common energy, or in some other way? This, too, remains to be seen.

The differential phase spectra in Figs. 6.4.9 and 6.4.10 may be seen to be similar to each other, but quite different from phase spectra obtained from a normal electrocardiogram (see Figs. 6.4.14 and 6.4.15). Does this mean that the phase spectrum could be useful in recognizing a supraventricular rhythm?

Figs. 6.4.12 through 6.4.16, showing spectra obtained from two consecutive cycles of a normal electrocardiogram, are included to emphasize the fact that one must be able to distinguish (and disregard) changes from cycle to cycle which do not arise from abnormal or undesirable changes in cardiac function. It is interesting to note that the differential phase spectrum of Fig. 6.4.17 shows angles most of which are close to 180 degrees in magnitude, in sharp contrast with the spectra of Figs. 6.4.9 and 6.4.10, in which the normal and abnormal electrocardiograms are compared.

Clearly the "right" way to go about using electrocardiographic spectra in diagnosis and monitoring is not immediately obvious from our results. What we wish to emphasize is the ease with which Fourier series may be obtained for these purposes through adaptive PL sampling.

Partly to compensate for the potentially misleading results of amplitude normalization in our plots (some of which, however, would have been unclear without it) we include the "raw" amplitude information for each of the spectra shown. These data may be found in Tables 6.4.1 through 6.4.8.

THE FOLLOWING DISPLAY SHOWS HARMONIC NUMBER,
 FREQUENCY IN HERTZ, DB WITH RESPECT TO THE FUNDAMENTAL,
 AND LOGARITHMICALLY-SCALED LINE SPECTRUM.

SUPRAVENTRICULAR RHYTHM, 1 CYCLE




FIGURE 6.4.3

Amplitude spectrum for the second cycle of the abnormal waveform shown in Fig. 6.4.1.

SUPRAVENTRICULAR RHYTHM, 1 CYCLE




FIGURE 6.4.4

Phase spectrum for the first cycle of the abnormal waveform shown in Fig. 6.4.1.

SUPRAVENTRICULAR RHYTHM, 1 CYCLE




FIGURE 6.4.5

Phase spectrum for the second cycle of the abnormal waveform shown in Fig. 6.4.1.

THE FOLLOWING DISPLAY SHOWS HARMONIC NUMBER,
FREQUENCY IN HERTZ, DB WITH RESPECT TO THE FUNDAMENTAL,
AND LOGARITHMICALLY-SCALED LINE SPECTRUM.

DIFFERENTIAL SPECTRUM, SUPRAVENTRICULAR RHYTHM




FIGURE 6.4.6

Differential amplitude spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.




FIGURE 6.4.9

Differential phase spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.




FIGURE 6.4.10

Differential phase spectrum comparing the second cycle of the abnormal electrocardiogram shown in Fig. 6.4.1 with a normal electrocardiogram.




FIGURE 6.4.11

Differential phase spectrum comparing two consecutive cycles of the abnormal electrocardiogram shown in Fig. 6.4.1.

THE FOLLOWING DISPLAY SHOWS HARMONIC NUMBER,
FREQUENCY IN HERTZ, DB WITH RESPECT TO THE FUNDAMENTAL,
AND LOGARITHMICALLY-SCALED LINE SPECTRUM.

SECOND CYCLE, NORMAL ECG




FIGURE 6.4.13

Amplitude spectrum of the second cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.

TABLE 6.4.1

Spectrum of the first cycle of the abnormal waveform shown in Fig. 6.4.1.

SUPRAVENTRICULAR RHYTHM, 1 CYCLE

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	3.882E+00	8.00	-21.937	0.0
1	1.146E+00	4.852E+01	100.00	0.0	-9.34
2	2.293E+00	4.443E-01	0.92	-40.766	-89.04
3	3.439E+00	7.528E+00	15.52	-16.184	104.19
4	4.585E+00	5.191E+00	10.70	-19.413	78.73
5	5.732E+00	4.866E+00	10.03	-19.975	126.85
6	6.878E+00	6.970E+00	14.37	-16.853	125.97
7	8.024E+00	4.687E+00	9.66	-20.300	112.25
8	9.171E+00	2.977E+00	6.14	-24.243	115.45
9	1.032E+01	3.731E+00	7.69	-22.281	110.14
10	1.146E+01	2.222E+00	4.58	-26.784	143.52
11	1.261E+01	1.472E+00	3.03	-30.357	-157.68
12	1.376E+01	2.050E+00	4.23	-27.482	-141.78
13	1.490E+01	2.244E+00	4.63	-26.696	-130.84
14	1.605E+01	2.529E+00	5.21	-25.659	-128.99
15	1.720E+01	2.199E+00	4.53	-26.875	-126.27
16	1.834E+01	2.333E+00	4.81	-26.360	-135.41
17	1.949E+01	1.841E+00	3.80	-28.415	-118.62
18	2.063E+01	1.256E+00	2.59	-31.737	-100.32
19	2.178E+01	8.388E-01	1.73	-35.245	-92.91
20	2.293E+01	7.113E-01	1.47	-36.677	-80.94

```
*****
```


TABLE 6.4.2

Spectrum of the second cycle of the abnormal waveform shown in Fig. 6.4.1.

SUPRAVENTRICULAR RHYTHM, 1 CYCLE

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	9.284E+00	10.10	-19.918	0.0
1	1.294E+00	9.197E+01	100.00	0.0	-5.39
2	2.587E+00	1.154E+01	12.54	-18.031	15.86
3	3.881E+00	1.054E+01	11.46	-18.818	53.73
4	5.174E+00	4.873E+00	5.30	-25.517	55.63
5	6.468E+00	6.502E+00	7.07	-23.012	104.51
6	7.761E+00	5.740E+00	6.24	-24.094	84.80
7	9.055E+00	4.729E+00	5.14	-25.778	90.86
8	1.035E+01	2.947E+00	3.20	-29.885	97.63
9	1.164E+01	1.499E+00	1.63	-35.755	99.14
10	1.294E+01	1.202E+00	1.31	-37.675	162.94
11	1.423E+01	1.626E+00	1.77	-35.053	-162.14
12	1.552E+01	1.842E+00	2.00	-33.966	-142.12
13	1.682E+01	1.836E+00	2.00	-33.994	-138.37
14	1.811E+01	2.302E+00	2.50	-32.032	-140.63
15	1.940E+01	1.472E+00	1.60	-35.917	-133.17
16	2.070E+01	1.158E+00	1.26	-37.998	-113.68
17	2.199E+01	9.488E-01	1.03	-39.729	-113.92
18	2.328E+01	7.272E-01	0.79	-42.040	-97.74
19	2.458E+01	6.447E-01	0.70	-43.085	-71.27
20	2.587E+01	4.387E-01	0.48	-46.430	-48.99

```
*****
```

TABLE 6.4.3

Differential spectrum comparing the first cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.

DIFFERENTIAL SPECTRUM, SUPRAVENTRICULAR RHYTHM

```

*****
HARMONIC   FREQUENCY   AMPLITUDE   PERCENT OF   DR W.R.T.   PHASE
NUMBER     IN HERTZ      FUNDAMENTAL FUNDAMENTAL FUNDAMENTAL IN DEG.
*****
  0         0.0         5.194E+01   23.75        -12.486     0.0
  1         1.000E+00   2.1E9E+02   100.00        0.0        -0.19
  2         2.000E+00   4.822E+01   22.02        -13.142    -7.17
  3         3.000E+00   1.253E+01    5.72        -24.845    10.79
  4         4.000E+00   1.254E+01    5.73        -24.840    39.38
  5         5.000E+00   4.500E+00    2.06        -33.742    28.43
  6         6.000E+00   6.211E+00    2.84        -30.942   136.16
  7         7.000E+00   3.427E+00    1.57        -36.109   112.21
  8         8.000E+00   2.477E+00    1.13        -38.926   171.70
  9         9.000E+00   2.655E+00    1.21        -38.323   162.41
 10        1.000E+01   3.949E+00    1.80        -34.877    -165.25
 11        1.100E+01   4.345E+00    1.98        -34.047    -143.03
 12        1.200E+01   4.889E+00    2.23        -33.021   -142.40
 13        1.300E+01   4.726E+00    2.16        -33.317   -145.15
 14        1.400E+01   4.951E+00    2.26        -32.913   -149.10
 15        1.500E+01   4.389E+00    2.00        -33.959   -153.49
 16        1.600E+01   4.658E+00    2.13        -33.441   -155.56
 17        1.700E+01   3.551E+00    1.62        -35.800   -154.71
 18        1.800E+01   2.595E+00    1.19        -38.522   -152.80
 19        1.900E+01   1.667E+00    0.76        -42.369   -155.53
 20        2.000E+01   1.253E+00    0.57        -44.848   -154.63
*****

```

TABLE 6.4.4

Differential spectrum comparing the second cycle of the abnormal waveform shown in Fig. 6.4.1 with a normal electrocardiogram.

DIFFERENTIAL SPECTRUM, SUPRAVENTRICULAR RHYTHM

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DEG W.R.T. FUNDAMENTAL	PHASE IN DEG.
-----------------	--------------------	-----------	------------------------	------------------------	---------------

0	0.0	5.739E+01	21.85	-13.210	0.0
1	1.000E+00	2.626E+02	100.00	0.0	-0.32
2	2.000E+00	5.898E+01	22.46	-12.972	-2.35
3	3.000E+00	2.070E+01	7.88	-22.069	9.86
4	4.000E+00	1.334E+01	5.08	-25.881	31.07
5	5.000E+00	6.940E+00	2.64	-31.559	40.89
6	6.000E+00	4.380E+00	1.67	-35.557	88.24
7	7.000E+00	3.586E+00	1.37	-37.295	83.47
8	8.000E+00	1.671E+00	0.64	-43.924	159.31
9	9.000E+00	1.922E+00	0.73	-42.790	-140.58
10	1.000E+01	3.743E+00	1.43	-36.921	-148.19
11	1.100E+01	4.459E+00	1.70	-35.402	-145.08
12	1.200E+01	4.681E+00	1.78	-34.979	-142.56
13	1.300E+01	4.387E+00	1.67	-35.543	-149.56
14	1.400E+01	4.881E+00	1.86	-34.616	-155.34
15	1.500E+01	3.846E+00	1.46	-36.687	-160.88
16	1.600E+01	3.330E+00	1.27	-37.937	-156.09
17	1.700E+01	2.819E+00	1.07	-39.383	-164.20
18	1.800E+01	2.282E+00	0.87	-41.219	-162.90
19	1.900E+01	1.349E+00	0.51	-45.784	-159.93
20	2.000E+01	9.702E-01	0.37	-48.649	-170.19

TABLE 6.4.5

Differential spectrum comparing two consecutive cycles of the abnormal waveform shown in Fig. 6.4.1.

DIFFERENTIAL SPECTRUM, SUPRAVENTRICULAR RHYTHM

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	5.402E+00	12.36	-18.157	180.00
1	1.000E+00	4.369E+01	100.00	0.0	179.00
2	2.000E+00	1.166E+01	26.69	-11.475	-162.03
3	3.000E+00	8.167E+00	18.69	-14.566	-171.57
4	4.000E+00	2.039E+00	4.67	-26.618	148.38
5	5.000E+00	2.725E+00	6.24	-24.102	-118.23
6	6.000E+00	4.615E+00	10.56	-19.525	-179.06
7	7.000E+00	1.747E+00	4.00	-27.960	-167.11
8	8.000E+00	9.179E-01	2.10	-33.552	-165.50
9	9.000E+00	2.278E+00	5.21	-25.658	117.35
10	1.000E+01	1.159E+00	2.65	-31.524	123.36
11	1.100E+01	1.947E-01	0.45	-47.020	-18.14
12	1.200E+01	2.084E-01	0.48	-46.430	-138.76
13	1.300E+01	4.875E-01	1.12	-39.048	-101.27
14	1.400E+01	5.396E-01	1.24	-38.166	-69.61
15	1.500E+01	7.587E-01	1.74	-35.207	-112.79
16	1.600E+01	1.329E+00	3.04	-30.341	-154.24
17	1.700E+01	8.993E-01	2.06	-33.730	-123.58
18	1.800E+01	5.308E-01	1.21	-38.309	-103.85
19	1.900E+01	3.375E-01	0.77	-42.241	-137.76
20	2.000E+01	4.110E-01	0.94	-40.532	-115.34

```
*****
```

TABLE 6.4.6

Spectrum of the first cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.

FIRST CYCLE, NORMAL ECG

```

*****
HARMONIC  FREQUENCY  AMPLITUDE  PERCENT OF  DB W.R.T.  PHASE
NUMBER    IN HERTZ    FUNDAMENTAL  FUNDAMENTAL  IN DEG.
*****
  0         0.0       4.811E+01  28.10       -11.026    180.00
  1       1.064E+00  1.712E+02  100.00         0.0      -177.61
  2       2.128E+00  4.816E+01  28.13       -11.017    173.36
  3       3.191E+00  1.500E+01  8.76        -21.150    160.72
  4       4.255E+00  9.138E+00  5.34        -25.452   -161.73
  5       5.319E+00  7.095E+00  4.14        -27.651    165.71
  6       6.383E+00  1.393E+00  0.81        -41.792     73.92
  7       7.447E+00  1.261E+00  0.74        -42.658    112.34
  8       8.511E+00  2.609E+00  1.52        -36.342     63.30
  9       9.574E+00  2.975E+00  1.74        -35.201     65.22
 10      1.064E+01  3.088E+00  1.80        -34.875     48.86
 11      1.170E+01  2.944E+00  1.72        -35.292     44.24
 12      1.277E+01  2.839E+00  1.66        -35.606     37.15
 13      1.383E+01  2.610E+00  1.52        -36.336     22.59
 14      1.489E+01  2.719E+00  1.59        -35.983     12.23
 15      1.596E+01  2.633E+00  1.54        -36.259      4.06
 16      1.702E+01  2.596E+00  1.52        -36.385      6.39
 17      1.809E+01  2.330E+00  1.36        -37.322     -2.45
 18      1.915E+01  2.084E+00  1.22        -38.292     -1.36
 19      2.021E+01  1.482E+00  0.87        -41.255     -5.71
 20      2.128E+01  1.255E+00  0.73        -42.698     -7.59
*****

```

TABLE 6.4.7

Spectrum of the second cycle of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.

SECOND CYCLE, NORMAL ECG

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	4.628E+01	178.99	5.057	180.00
1	1.064E+00	2.586E+01	100.00	0.0	-164.11
2	2.128E+00	1.204E+01	46.57	-6.638	150.19
3	3.191E+00	6.621E+00	25.60	-11.834	50.50
4	4.255E+00	3.530E+00	13.65	-17.296	-59.05
5	5.319E+00	2.630E+00	10.17	-19.854	50.24
6	6.383E+00	6.647E+00	25.70	-11.800	7.10
7	7.447E+00	5.048E+00	19.52	-14.189	15.76
8	8.511E+00	5.984E+00	23.14	-12.712	24.67
9	9.574E+00	5.053E+00	19.54	-14.181	21.50
10	1.064E+01	4.883E+00	18.88	-14.479	24.07
11	1.170E+01	4.479E+00	17.32	-15.238	19.64
12	1.277E+01	4.101E+00	15.86	-15.994	16.68
13	1.383E+01	4.067E+00	15.73	-16.067	11.16
14	1.489E+01	3.635E+00	14.06	-17.041	3.71
15	1.596E+01	3.213E+00	12.42	-18.114	-0.73
16	1.702E+01	3.176E+00	12.28	-18.216	-2.95
17	1.809E+01	2.272E+00	8.78	-21.125	-0.19
18	1.915E+01	1.618E+00	6.26	-24.070	-11.31
19	2.021E+01	1.292E+00	5.00	-26.024	-13.51
20	2.128E+01	8.713E-01	3.37	-29.449	-17.67

```
*****
```

TABLE 6.4.8

Differential spectrum comparing two consecutive cycles of the normal electrocardiogram shown in Figs. 4.3.2 and 6.4.1.

DIFFERENTIAL SPECTRUM, NORMAL ECG

```
*****
```

HARMONIC NUMBER	FREQUENCY IN HERTZ	AMPLITUDE	PERCENT OF FUNDAMENTAL	DB W.R.T. FUNDAMENTAL	PHASE IN DEG.
0	0.0	1.823E+00	1.25	-38.084	180.00
1	1.064E+00	1.462E+02	100.00	0.0	-179.97
2	2.128E+00	3.739E+01	25.58	-11.843	-179.36
3	3.192E+00	1.837E+01	12.57	-18.016	-179.51
4	4.256E+00	1.049E+01	7.18	-22.878	179.11
5	5.320E+00	8.562E+00	5.86	-24.647	-178.19
6	6.384E+00	6.232E+00	4.26	-27.406	175.25
7	7.448E+00	5.342E+00	3.65	-28.744	-177.80
8	8.512E+00	4.269E+00	2.92	-30.691	-177.75
9	9.576E+00	3.557E+00	2.43	-32.275	166.20
10	1.064E+01	2.449E+00	1.68	-35.518	172.16
11	1.170E+01	2.180E+00	1.49	-36.530	165.43
12	1.277E+01	1.750E+00	1.20	-38.436	162.12
13	1.383E+01	1.595E+00	1.09	-39.245	172.22
14	1.490E+01	1.029E+00	0.70	-43.052	160.66
15	1.596E+01	6.283E-01	0.43	-47.334	158.77
16	1.702E+01	7.451E-01	0.51	-45.853	142.59
17	1.809E+01	1.079E-01	0.07	-60.000	0.0
18	1.915E+01	5.641E-01	0.39	-48.271	28.35
19	2.022E+01	2.669E-01	0.18	-54.772	35.35
20	2.128E+01	4.254E-01	0.29	-50.722	13.44

```
*****
```

To conclude this section, we offer a final observation with respect to spectral analysis of electrocardiograms based on the three waveforms shown in Fig. 6.4.18. In each of these waveforms, the actual configuration of the electrocardiogram is much the same; the main differences lie in the duration of diastole as reflected in the interval between the T- and P-waves. This shows that the rapid pulse of sinus tachycardia and the slow pulse of sinus bradycardia are not merely variations in the fundamental frequency, or repetition rate, of an otherwise invariant waveform. In other words, the Fourier series of each of these waveforms will be different. The fundamental frequencies (and therefore the overtone frequencies) differ, and so will the corresponding harmonic amplitudes; one cannot obtain one waveform from another merely by time-axis scaling.

On the other hand, owing to the similarities of the waveform configurations during systole (the "active" portions of the tracings), two approaches suggest themselves for examination of waveform configuration in the frequency domain isolated from differences in repetition rate arising from differences in the duration of diastole.




FIGURE 6.4.18

Examples of sinus arrhythmia. (a) Normal sinus rhythm; (b) sinus bradycardia; (c) sinus tachycardia.

One possible approach is the use of the Fourier transform rather than the Fourier series, since if a single cardiac cycle is extracted from any of these tracings by time-domain windowing, the periodicity of the cardiac cycle is ignored and the resulting spectrum does not reflect changes in the duration of diastole.

A second possible approach is to retain the use of the Fourier series, but to normalize each waveform, prior to calculation of the Fourier series, by extracting one cycle (from the onset of the P-wave to the end of the T-wave) and scaling it to an arbitrary duration. The resulting Fourier-series spectra no longer reflect variations in the duration of diastole, and may perhaps be used for investigation of other conditions or disorders in the manner we have suggested.

The Fourier-series approach seems to offer a possible advantage in that fewer data are needed to define the Fourier series than the Fourier transform for the same range of frequencies. Even if it should be found that the frequencies and amplitudes of only the few strongest peaks in the Fourier transform are sufficiently useful in diagnosis or monitoring, many values must still

be calculated even though only a few are to be stored, unless it should be found that the frequencies at which the major peaks occur can be calculated from the breakpoint parameters prior to calculation of the Fourier transform itself.

6.5 -- PL Sampling of Heart Sounds

Since the PL sampling technique is oriented to data compression in terms of waveform curvature and not predicated on any particular model of the source of the signal under consideration, it may be applied to heart sounds -- and with the same hardware and software -- just as readily as to electrocardiograms or speech. The same system limitations apply as well; the sampling clock must be of sufficiently high frequency to permit adequate approximation of high-frequency changes in the waveform, but it must also be of low enough frequency to allow the system to recognize small slopes.

Fig 6.5.1 shows a portion of a commercial recording of normal heart sounds which has been sampled by our system using the microcomputer program described in Section 5.7. The sampling-clock frequency was 503 Hz, which may be seen from the figure to be high enough to produce




FIGURE 6.5.1

Normal heart sounds obtained by PL sampling. The first (S1) and second (S2) heart sounds are labelled.

some equivocal nonzero slopes between major components of the heart sounds. On the other hand, when a portion of this waveform is extracted from the PL record by windowing, as described in Section 4.4, the approximation (see Fig. 6.5.2) may be seen to provide adequately-detailed signal peaks.

The Fourier transform of the second heart sound, calculated from the parameters of the 69 breakpoints which make up Fig. 6.5.2 by means of the equations derived in Section 4.2, is shown in Fig. 6.5.3. Fig. 6.5.4 shows a second plot of the Fourier transform, but in this case the frequency axis is scaled logarithmically to make the spectrogram easier to read at low frequencies. It is clearer in this plot, for example, that the two major peaks occur at about 40 and 53 Hz. A third plot of the Fourier transform for the normal second heart sound is shown in Fig. 6.5.5; here both axes are scaled logarithmically, and the "fine structure" at low amplitudes is more clearly visible.

We are not in a position to attempt to interpret this example in terms of cardiac function, and will content ourselves for the time being with illustrating the



FIGURE 6.5.2

Second heart sound from the beginning of the trace shown in Fig. 6.5.1, hanning-windowed for computation of the Fourier transform.




FIGURE 6.5.3

Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. Both frequency and amplitude axes are scaled linearly.




FIGURE 6.5.4

Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. The frequency axis has been scaled logarithmically to facilitate identification of low-frequency components.




FIGURE 6.5.5

Fourier transform of the second heart sound from the beginning of the trace shown in Fig. 6.5.1. Both axes have been scaled logarithmically.

feasibility of obtaining Fourier transforms from PL breakpoints. One reason for this is that the strongest peak seems to be at approximately 53 Hz. Although we are more confident of the stability of frequency calibration in this method than in that used to obtain the long-term spectra shown in Section 3.2, we have less reason to trust the turntable and tape speeds involved in copying the commercial recordings for analysis and cannot be entirely certain that this strong peak is not in fact 60-Hz hum.

A second reason is related to the low amplitudes at frequencies below 40 Hz. As we showed in Chapter 2, most of the vibrational energy associated with the cardiac cycle is of low -- in fact, subaudio -- frequency. The commercial recording is limited to audible frequencies, and we have also filtered out subaudio frequencies in order to exclude noise in this region prior to calculating the Fourier transform. Although it may be common practice to use transducers insensitive to frequencies below 40 Hz or to filter out components of lower frequency, as we have been told, we feel that to exclude these frequencies and attempt to develop machine-aided analysis techniques without them would be premature at

best. Lacking evidence to support the assumption that the low-frequency components are of little or no clinical significance, we would prefer to start by examining Fourier transforms of heart sounds sampled from recordings of whose quality we could be more confident.

It should be clearly understood that the examples of PL processing we have presented in this chapter are intended merely to suggest some of the ways in which our techniques might be applied to cardiology-related problems. We make no attempt to interpret the results of our measurements from a medical standpoint; we wish only to show that the techniques we have developed can in fact be used to perform such measurements.

C H A P T E R 7

CONCLUSIONS AND RECOMMENDATIONS FOR FURTHER INVESTIGATION

7.1 -- Wideband Apexcardiography

We have proposed a technique for consideration of cardiac-cycle vibrations as a whole by means of a single display extending from the apexcardiogram through the audible heart sounds. We have seen that although analog filtering is safe enough for obtaining audible-frequency phonocardiograms and apexcardiograms in the conventional manner, it cannot be used to produce the type of wideband display under consideration because of excessive relative time shifts. However, we have found that this difficulty may be entirely circumvented by means of zero-phase-shift digital filtering.

Examination of cardiac vibration in this fashion shows certain features, which appear not to have been investigated, at frequencies lower than those usually considered in conventional phonocardiography but still sufficiently high to be obscured by stronger components at lower frequencies in conventional apexcardiography.

Having found that the left lateral decubitus position cannot always be used in obtaining apexcardiograms, we have established that low-pass filtering (even zero-phase-shift digital filtering) is not sufficient to recover "normal-looking" apexcardiograms from signals recorded with the subject supine.

We feel that clinical evaluation of displays such as that of Fig. 2.3.1, obtained from wideband measurements of precordial displacement at the cardiac apex and elsewhere, should be conducted in each of two ways. Examination of wideband apexcardiograms representative of various known disorders should establish on an empirical basis what potential diagnostic usefulness may be embodied by this technique. Somewhat less straightforward but perhaps of potentially greater value in the long run would be an attempt to relate, through wideband apexcardiography, the events giving rise to the audible heart sounds and significant features of the conventional audiocardiogram to the low-frequency precordial displacements measured in conventional apexcardiography.

In particular, we suggest that an investigation of the genesis of the midsystolic and early diastolic peaks

discussed in Section 2.3 might prove fruitful, and should be conducted at least to the point where the events giving rise to these peaks are identified as a necessary first step in determining their clinical significance, if any.

We suggest further that the possibility of recovering "normal-looking" apexcardiograms from signals recorded with the subject supine might be worthy of study, since such recovery by deconvolution or other means could extend the usefulness of conventional and wideband apexcardiography to cases where the left lateral decubitus position is contraindicated. Moreover, in cases where recordings may safely be made in each of the two positions, it seems possible that the same technique might help to isolate changes in cardiac function attendant upon changes in position.

Finally, we suggest that the possibility of obtaining useful diagnostic information by examination of apexcardiographic signals in the frequency domain be considered. As a possible starting point, it might be helpful to determine whether the pulse-rate periodicity in some published apexcardiogram spectra is in fact

spurious, as we suspect since the pulse rate is not a modal frequency of the heart, considered as a vibrating mechanical system; we have seen, for example, that the morphology of the electrocardiogram is not necessarily linked to the pulse rate (see Fig. 6.4.18). Such a determination could be made fairly easily using the methods discussed in Chapters 4 and 5 to obtain Fourier transforms such as those shown in Section 6.5. Care must be taken, however, to ensure that the waveform is shifted vertically prior to windowing so that the windowing process itself does not produce misleading deflections from zero near the beginning and end of the record. This preliminary shifting may be performed with impunity since it constitutes no more than an alteration of the dc level of the waveform. If it is done successfully, the beginning and end of the record of interest will already be zero before windowing takes place.

7.2 -- Conversion of Heart Sounds by Frequency Modulation

In Chapter 3 we discussed techniques whereby frequency modulation of an audio-frequency sinusoid could be used to generate sounds containing the same information as an input signal, but occupying a bandwidth in which the human auditory system is more sensitive. We have

shown that although frequency modulation performed using a linear-response voltage-controlled oscillator can produce new sounds having complex spectra even if the carrier and modulating signal are both sinusoidal, this approach is inconvenient because artificial means of "frequency foldover" must be employed to deal with negative frequencies. On the other hand, we have shown that frequency modulation achieved through the use of an exponential-response VCO cannot give rise to foldover problems and imposes no restrictions upon depth of modulation other than that implicit in the finite upper frequency limit of the VCO itself; no negative frequencies are involved. We have found that this approach provides a means of producing a close parallel between the instantaneous amplitude of the input signal and the perceived pitch of the converted sounds.

We have demonstrated that pseudologarithmic compression of the input signal may be used with this technique to emphasize amplitude changes occurring in the input signal at high or low amplitudes. Moreover, since both frequency and amplitude changes in the input signal are reflected in pitch changes in the converted sounds, we are free to deal with output amplitude as we see fit, and

have described a noise threshold control circuit permitting de-emphasis of low-amplitude noise while ensuring rapidly-increasing amplitude of the converted sounds whenever the input signal rises above the chosen threshold.

We have shown that the converted sounds may be "tailored" to occupy a desired bandwidth, for example to facilitate telephone transmission, in which case the sounds received contain all of the information implicit in the input signal, no matter how low the original frequencies involved, and may be heard without further processing or special apparatus. Finally, we have demonstrated that this technique affords a simple and inexpensive means of "slowing down" the converted sounds with no loss of information due to reduced frequencies.

On the basis of our experiments, we feel that this technique should be subjected to clinical tests in order to determine what difficulties are involved in becoming accustomed to the converted sounds, in relating them either to familiar sounds or directly to cardiac function, and in learning to adjust system response so as to isolate desired features of the sounds heard. A portable

conversion system is already available for this purpose (see Appendix C), although there is clearly room for improvement and considerable miniaturization of the instrument.

Finally, we suggest that a study of the potential usefulness of this technique with otherwise inaudible signals might be fruitful. Such signals include subaudio cardiac vibrations, but the system can be used just as easily to convert electrocardiographic, electroencephalographic, or electromyographic signals, to name a few. We feel that if the hurdle of becoming accustomed to unfamiliar sounds can be overcome, this technique could be useful in diagnosis, monitoring, and biofeedback applications. We suggest parenthetically that the basic principle of exponential-response frequency modulation of an audio-frequency sinusoid may be used to advantage in any circumstance where the conversion of instantaneous amplitude to perceived pitch (as opposed to frequency) might be helpful.

7.3 -- Applications of PL Sampling

In Chapter 4 we discussed the representation of an arbitrary analog waveform by a piecewise-linear approx-

imation completely defined by its breakpoint parameters. We have shown that the Fourier transform and Fourier series can be easily calculated directly from the breakpoint parameters, without further approximation.

Our analysis has shown that the number of breakpoints needed to define an adequate PL approximation may be much smaller than the minimum number of uniformly-spaced samples necessary to satisfy the Nyquist sampling criterion for the same record, and that the problem of aliasing does not arise when PL sampling is used.

We have described experiments aimed at finding a method of choosing appropriate breakpoints automatically. We have presented in detail an algorithm which has proven capable of successfully performing adaptive PL sampling with breakpoint density determined by the curvature of the original waveform, and we have examined some of the special system restrictions inherent in this approach.

We have shown on the basis of probability estimation that adaptive PL sampling offers a potential means of data compression in the digitization of speech, although it implies no assumption of a particular "model" for

speech; the method is especially advantageous in connection with normal speech containing pauses during which no breakpoints are needed. We have described a method of reconstructing PL waveforms by computer-controlled analog interpolation between stored breakpoints, and have given examples of Fourier series and Fourier transforms obtained from electrocardiograms and heart sounds digitized by means of our technique.

In no respect is the feasibility-study nature of our work more clearly apparent than in the matter of adaptive PL sampling. Our efforts appear to have raised more questions than they have answered, and to some degree our recommendations for further investigation will perforce refer to entire categories of studies rather than to specific projects.

To begin with, we feel that the possibility of applying PL techniques to the digital synthesis of transfer functions in general and digital filters in particular should be looked into. Similarly, it may well be that computational "shortcuts," analogous to the fast-Fourier-transform methods used in calculating the discrete Fourier transform, could be found for more efficient

computation of Fourier series and Fourier transforms from PL breakpoint parameters.

Another theoretical project of potential usefulness would be an attempt to work out a means of calculating the frequencies at which major Fourier-transform amplitude peaks occur without actually calculating the transform itself, again directly from the breakpoint parameters. In the spirit of formant analysis, it may be that calculation of the amplitudes of only the most significant peaks would provide sufficient information for some types of heart-sound analysis or monitoring, or might be applied to computer recognition of speech.

We suggest that a PL-sampling system be set up using a 12-bit A/D converter and a 16-bit microcomputer, thus circumventing the limitations inherent in an 8-bit system such as that used in our experiments, for use in evaluating the potential usefulness of PL techniques in each of several areas.

It would be of interest to determine whether PL sampling could offer significant data compression of speech signals for economical storage of such signals in

digital form, even though no speech "model" is implicit in the technique. Similarly, it would be of interest to investigate the possibility of using PL sampling to digitize musical signals for recording and analysis. Along these lines, Professor Mauro Zambuto of NJIT has suggested that the technique might afford sufficient data compression in digitizing arbitrary waveforms to permit the recording of sound for motion pictures in digital form using an optical sound track.

We suggest that the possibility of using PL sampling to record electrocardiograms digitally over long periods of time be considered as an alternative to Holter-tape recording, partly in order to record over longer periods of time and partly with a view towards classifying and analyzing the recordings with computer assistance, either by feature recognition or waveform comparison in the time domain or by conversion to the frequency domain using the methods we have described.

In particular, we feel that advantage could be taken of the ease of access to the frequency domain afforded by PL techniques to explore the possibilities of using frequency-domain displays and comparison and recognition

algorithms in cardiac monitoring and diagnosis. For example, the sampling algorithm described in Sections 5.5 and 5.7 can easily be modified to permit multiplexed sampling of several analog channels. We suggest that a full set of electrocardiograms be recorded by simultaneous sampling in this manner for subsequent conversion into the frequency domain in order to determine to what degree examination of amplitude and phase spectra might be helpful in compressing redundant information.

As in the case of wideband apexcardiography, investigations of the possible medical applications of PL techniques may be carried out on an empirical basis, in terms of establishing relationships between observed spectral features and known cardiac disorders, conditions, or events, or on an analytical basis, in the form of attempts to learn better to understand the events which give rise to the spectral features; nor must such investigations be limited to cardiography. Spectral investigations have been conducted before, but it is our hope that the techniques we have begun to develop will be found to simplify the making of spectral measurements in such studies, and to facilitate the practical application of their results.

A P P E N D I X A

COMPUTER PROGRAMS DEVELOPED FOR USE IN THE RESEARCH

In the following pages brief functional descriptions will be given of some of the programs which have been developed during the course of this research. The experimental and computational resources available for this work included a microcomputer system (suitable for A/D and D/A conversion, but limited as to computational ability) and a full-scale computer system based on an IBM 360, with batch-processing and interactive capabilities (well-suited for computation, but without facilities for interfacing with measurement apparatus). Since disc- or tape-storage facilities were not locally available which would have allowed data to be transferred to the IBM system for batch processing, a certain amount of software had to be developed so that the microcomputer could transfer data to and from the IBM system by emulating an interactive terminal. This approach was found to be relatively slow (being limited to 300 baud by the telephone link), but otherwise proved to be adequate for our purposes.

Many of the programming details are, of course, system-dependent, which is one reason that program descriptions were felt to be more appropriate for this appendix than program listings, flowcharts, or flowgrams. We have made only one significant exception to this philosophy, that being the complete assembler listing of the adaptive piecewise-linear sampling program (central to our work) which has been given in Chapter 5.

The programs described in this appendix do not include those which were needed only for "housekeeping," such as file management, indexing, and cataloguing.

P R O G R A M S

Program	Language	Page
ASMBLIST	FORTRAN	284
BAND	FORTRAN	284
CORR	FORTRAN	284
CYCLE	FORTRAN	284
CYCLE3	FORTRAN	285
FIXIN	Z-80 ASSEMBLY LANGUAGE	285
FIXOUT	Z-80 ASSEMBLY LANGUAGE	286
FT	FORTRAN	286
GNIVRI	FORTRAN	286
GNIVRI	Z-80 ASSEMBLY LANGUAGE	287
HIGH	FORTRAN	287
IRVING	Z-80 ASSEMBLY LANGUAGE	288
IRVING	FORTRAN	288
LINAXIS	FORTRAN	288
LOW	FORTRAN	289
ORDCAL	Z-80 ASSEMBLY LANGUAGE	289
PLOT1,PLOT2,...,PLOT13	FORTRAN	289
PLOUT1	Z-80 ASSEMBLY LANGUAGE	290
PLOUT2	Z-80 ASSEMBLY LANGUAGE	290
PLOUT3	Z-80 ASSEMBLY LANGUAGE	290
PLS3	Z-80 ASSEMBLY LANGUAGE	291
PLS3\$	Z-80 ASSEMBLY LANGUAGE	291
SAMPFIX	FORTRAN	292
SAMPLOT	FORTRAN	292
SAMPLOT	Z-80 ASSEMBLY LANGUAGE	293
SCALE	FORTRAN	293
SIMUL	Z-80 ASSEMBLY LANGUAGE	293
SP1	FORTRAN	294
SP2	FORTRAN	294
SP3	FORTRAN	294
SP4	FORTRAN	295
TRANSFIX	Z-80 ASSEMBLY LANGUAGE	295
TRANSFIX	FORTRAN	296
WF1	FORTRAN	296
WF2	FORTRAN	296
WF3	FORTRAN	296
WF4	FORTRAN	297
WINDOW	FORTRAN	297

ASMBLIST (FORTRAN)

Produces a hard-copy assembler listing from a data file on disc. Such files are prepared by means of a terminal-emulation driver by which the microcomputer sends alphanumeric characters via a telephone modem as though it were a terminal. Appropriate indenting and page formatting are provided.

BAND (FORTRAN)

Performs zero-phase-shift digital filtering of a stored waveform which has been sampled at a fixed rate. From 8 to 80 poles may be chosen. The resulting waveform is stored on disc in standard waveform format. Each basic filter section is a 2-pole band-pass Butterworth filter, and zero phase shift is achieved through reversal of the record following each filtering operation.

CORR (FORTRAN)

Calculates the cross-correlation function of two signals from their Fourier-series coefficients as stored on disc. The result is stored in standard waveform format in a third disc file.

CYCLE (FORTRAN)

Extracts one chosen cycle from a periodic waveform

(either PL or fixed-rate) stored on disc, by examination of upper and lower peak values and comparison of those values with the mean value of the entire waveform. The result is stored in standard waveform format. This program is used primarily in preparation for the calculation of Fourier-series coefficients (see SPL) and for determination of the number of breakpoints defining one cycle of a PL waveform.

CYCLE3 (FORTRAN)

Extracts three contiguous cycles from a periodic waveform stored on disc, beginning with a chosen cycle. Similar in operation to CYCLE, this program is used to facilitate visual comparison of two or more waveforms by plotting them one under the other with the same time scale and with coincidence of major features.

FIXIN (Z-80 ASSEMBLY LANGUAGE)

Samples an input signal at a chosen fixed rate (controlled by software delay) via an a/d converter. Samples are stored in RAM in twos-complement representation. Sampling may be initiated manually, by means of a trigger pulse monitored by a second A/D channel, or by the onset of the sampled signal itself.

FIXOUT (Z-80 ASSEMBLY LANGUAGE)

Plots a waveform which has been sampled at a fixed rate and whose ordinate values have been stored in microcomputer RAM. Plotting is accomplished by means of an H/P x-y pen recorder having an internal sweep triggered by the computer at the start of the plotting operation; the y-axis channel is driven by a D/A converter. Time scaling may be accomplished either by software delay or by adjustment of the recorder's sweep rate.

FT (FORTRAN)

Calculates the Fourier transform of a stored PL waveform over a chosen range of frequencies. The output is stored on disc in standard waveform format, with separate files for amplitude spectrum and phase spectrum. Linear or logarithmic scaling may be chosen for the frequency axis.

GNIVRI (FORTRAN)

Transfers a PL waveform from a disc file to microcomputer RAM for subsequent plotting or signal reconstruction. Filed breakpoint times are converted to intervals between breakpoints; time intervals and ordinate values are then quantized (to 256 levels for

ordinates and 65536 levels for time intervals) and converted from decimal to signed hexadecimal representation before being transmitted in ASCII code as though to a terminal.

GNIVRI (Z-80 ASSEMBLY LANGUAGE)

Transfers a PL waveform file from disc to microcomputer RAM for subsequent plotting or signal reconstruction. A terminal-emulation routine reads data via a telephone modem in ASCII code, converting it to signed hexadecimal notation. Ordinate values are converted to twos-complement hexadecimal representation prior to being stored as 8-bit binary words; time intervals are stored as 16-bit binary words.

HIGH (FORTRAN)

Performs zero-phase shift digital filtering of a stored fixed-rate waveform. From 2 to 20 poles may be chosen; the resulting waveform is stored on disc in standard waveform format. Each basic filter section is a 2-pole high-pass Butterworth filter, and zero phase shift is achieved through reversal of the record following each filtering operation.

IRVING (Z-80 ASSEMBLY LANGUAGE)

Transfers PL-waveform breakpoint parameters from microcomputer RAM to an alphanumeric disc file. Data stored in RAM as 8-bit (for ordinates) and 16-bit (for time intervals) binary words are converted to signed hexadecimal representation and transmitted in ASCII code by means of a terminal-emulator routine and a telephone modem.

IRVING (FORTRAN)

Creates a PL waveform file from data which has been transferred from the microcomputer to an alphanumeric file on disc. Data are converted from signed hexadecimal to decimal, and time intervals are converted to breakpoint times, prior to filing. The program reports the total number of breakpoints read as well as the total duration of the record.

LINAXIS (FORTRAN)

Provides a scaling table for use with a computer-generated plot. Input consists of the maximum and minimum values of a plotted variable and the displacement range on the associated axis.

LOW (FORTRAN)

Performs zero-phase-shift digital filtering of a stored waveform which has been sampled at a fixed rate. From 2 to 20 poles may be chosen. The resulting waveform is stored on disc in standard waveform format. Each basic filter section is a 2-pole low-pass Butterworth filter, and zero phase shift is achieved through reversal of the record following each filtering operation.

ORDCAL (Z-80 ASSEMBLY LANGUAGE)

Provides continually refreshed calibration voltages via a D/A channel permitting accurate calibration of the vertical sensitivity of the x-y plotter driven by the microcomputer. After calibration, LINAXIS or LOGAXIS may be used to relate chart displacement to whatever variable is plotted on the vertical axis.

PLOT1,PLOT2,...,PLOT13 (FORTRAN)

Each of these programs calculates values of a different function and formats the result for plotting by means of a computer-driven plotter linked to a terminal. The programs are executed in batch mode, and a batch/timesharing interface system (such as COBI or SUPERWYLBUR) is used in plotting.

PLOUT1 (Z-80 ASSEMBLY LANGUAGE)

Resolves stored 16-bit time-interval data describing a piecewise-linear waveform into 8 bits by a series of logical shifts. An indication of the number of shifts required to reduce all of the high-order time-interval bytes to zero is given on the console when the processing is completed. The data are modified in place; the file format remains unaltered.

PLOUT2 (Z-80 ASSEMBLY LANGUAGE)

Controls the reconstruction of PL waveforms in real time from stored breakpoint parameters. Input to the program consists of the file address and a delay constant which determines the reconstruction rate. This program, which uses four analog and two digital outputs, controls an analog interpolator circuit which may be used to produce audible sounds or to plot waveforms at reduced speed.

PLOUT3 (Z-80 ASSEMBLY LANGUAGE)

Similar to PLOUT2, except that the reconstruction rate is determined by an external clock signal provided by a voltage-controlled oscillator. The control voltage of this oscillator is utilized by the analog interpolator circuit to permit continuous alteration, without recalibration, of the reconstruction rate.

PLS3 (Z-80 ASSEMBLY LANGUAGE)

Performs real-time adaptive PL sampling of an input waveform via an A/D converter. Separate control may be exercised over second-difference (corner) sensitivity and second-difference summation (curvature) sensitivity. Sampling may be initiated manually, by a trigger pulse, or by the onset of the signal itself. An oscilloscope strobe pulse is provided following each breakpoint so that the operation of the program may be monitored as it is being executed. Time intervals are stored in RAM as unsigned 16-bit words; ordinate values are stored in twos-complement representation as 8-bit words. An external sampling clock is used to determine the maximum rate at which the input signal is examined for possible breakpoints.

PLS3\$ (Z-80 ASSEMBLY LANGUAGE)

Essentially the same as PLS3, except that for each sample taken, the sample value, the first and second differences, the second-difference summation, and the absolute value of the second-difference summation are displayed on the console in hexadecimal representation. Moreover, whenever a breakpoint is chosen by the system, a symbol is displayed on the console to identify the

mechanism responsible for the choice: corner threshold exceeded, curvature threshold exceeded, or sampling-cycle count overflow.

SAMPFIX (FORTRAN)

Samples a PL waveform filed on disc at 1000 points equally distributed over the duration of the record. Ordinate values at times not coinciding with PL breakpoints are calculated by linear interpolation. The result is stored in standard waveform format. This program is used mainly in converting PL waveforms to their fixed-rate counterparts prior to windowing (see WINDOW) and digital filtering (see BAND, HIGH, and LOW).

SAMPLOT (FORTRAN)

Transfers fixed-rate waveform ordinates from disc to microcomputer RAM for plotting or signal reconstruction. Decimal values are quantized to 256 levels and converted to signed hexadecimal representation prior to transmission in ASCII code as though to a terminal. This program is similar to GNIVRI except that times are not transmitted (fixed-rate sampling is assumed), and the format is somewhat different.

SAMPLOT (Z-80 ASSEMBLY LANGUAGE)

Transfers fixed-rate waveform ordinates from disc to microcomputer RAM for plotting or signal reconstruction. This program reads waveform ordinates in ASCII code (in groups of three alphanumeric characters) via a terminal-emulation routine and a telephone modem. Data are converted from signed hexadecimal to twos-complement representation prior to storage in RAM as 8-bit words. Time intervals are neither read nor stored (fixed-rate sampling is assumed).

SCALE (FORTRAN)

Prepares a PL or fixed-rate waveform file for plotting on an x-y pen recorder driven by the microcomputer. The ordinate values are scaled and shifted as required so that the resulting ordinates will fall optimally within the range of an 8-bit register containing bipolar values in twos-complement representation.

SIMUL (Z-80 ASSEMBLY LANGUAGE)

Samples two input signals at a chosen fixed rate controlled by software delay, using two A/D channels. This program is similar to FIXIN, except that it is intended for use when two synchronous waveforms are to be

sampled at the same time (for example, an apexcardiogram and an electrocardiogram).

SP1 (FORTRAN)

Calculates Fourier-series coefficients for a periodic signal one cycle of which is stored on disc in a PL waveform file. The output consists of a "spectrum file" containing the amplitude and phase of up to 100 harmonics. The Fourier coefficients are calculated by means of the equations derived in the text, which are based upon the properties of piecewise-linear waveforms.

SP2 (FORTRAN)

Displays filed Fourier-series amplitude spectra in any or all of three formats: (1) a table showing harmonic number, frequency, amplitude, percent or fundamental-frequency amplitude, amplitude in dB relative to the fundamental, and phase; (2) a linearly-scaled histogram plot of the amplitude spectrum; (3) a logarithmically-scaled (dB) plot of the amplitude spectrum. The program will display as many harmonics as desired up to the number actually calculated and filed by SP1 or SP2.

SP3 (FORTRAN)

Plots Fourier-series phase spectra in histogram form

from spectrum files created by SP1 or SP2. Phase is expressed in degrees, and as many harmonics as desired may be displayed up to the number actually calculated by SP1 or SP2.

SP4 (FORTRAN)

Performs superposition of signals in the time domain by means of term-by-term addition of their Fourier series. Each component signal is assumed to have the same fundamental frequency. Each may be scaled by a constant factor or shifted in phase as desired prior to superposition. The output is stored on disc as a spectrum file.

TRANSFIX (Z-80 ASSEMBLY LANGUAGE)

Transfers fixed-rate waveform samples from microcomputer RAM to disc. Ordinate values stored as 8-bit twos-complement binary words are converted first to signed hexadecimal representation, and then transmitted in groups of three ASCII characters by means of a terminal-emulation routine and a telephone modem. Disc storage is in alphanumeric-file format.

TRANSFIX (FORTRAN)

Creates a waveform file from ordinates sampled at a fixed rate and transferred to an alphanumeric disc file by the microcomputer. Breakpoint times are calculated from the sampling rate as entered from the terminal keyboard, and ordinate values are converted from signed hexadecimal representation to decimal prior to filing.

WF1 (FORTRAN)

Creates a PL waveform file from breakpoint parameters entered from the terminal keyboard. If desired, the program provides a table of breakpoint times, ordinate values, and slopes; in addition, it sketches the waveform upon demand by means of linear interpolation between breakpoints.

WF2 (FORTRAN)

Permits editing of filed PL waveforms by terminal-keyboard entry of breakpoint parameters. A breakpoint-parameter table (similar to that generated by WF1) is provided as desired, and the revised waveform is refiled after editing has been completed.

WF3 (FORTRAN)

Produces a sketch of a filed waveform at the ter-

minal by linear interpolation between breakpoints. Scaling is automatic, and either PL or fixed-rate waveforms may be sketched.

WF4 (FORTRAN)

Performs superposition of waveforms in the time domain. Each component waveform consists of 1000 samples, assumed to be uniformly distributed over the total duration of the record. Each may be scaled by a constant factor, or shifted in phase as desired. The output is stored in standard waveform-file format.

WINDOW (FORTRAN)

Extracts a desired portion of a stored waveform prior to calculation of the Fourier transform (see FT). The Hanning (cosine-squared) window function is used for extraction, and the result is stored in standard waveform format. Identification of the time limits of the window is achieved by reference to the start and total duration of the input record. The program may be used either with PL or fixed-rate waveforms; however, care must be exercised when using it with PL waveforms in which the breakpoint density is low. For example, a triangle defined by three breakpoints beginning and ending with zero am-

plitude will emerge exactly as though the window were rectangular. In cases where the breakpoint density is low, therefore, it is preferable to convert the waveform to fixed-rate (see SAMPFIX) prior to windowing.

A P P E N D I X B

FM CONVERSION OF HEART SOUNDS: CIRCUIT DESCRIPTION OF THE EXPERIMENTAL SYSTEM

B.1 -- Introduction

The system described below was used to make the observations discussed in Chapter 3, and to obtain the results reflected in the various plots of waveforms, spectra, transfer functions, and sensitivity to input amplitude changes. It has been used with outboard voltage-controlled oscillators, both linear and exponential, and incorporates capabilities needed for dealing with each.

As a result, the complete circuit, which is shown in Fig. B.1.1, is more complex than is necessary if operation is to be restricted to exponential FM conversion, since in that case the artificial frequency foldover circuit is not needed. Other features, which will be described below, were included because it was felt that they might be useful. Although some of these could conceivably be simplified or omitted, it was felt that the decision to do so could be made realistically only after the system had been tried by practicing physicians, whose




FIGURE B.1.1

Circuit diagram of the experimental FM conversion system. All diodes and transistors are silicon, and all IC's are 1/2 747, except as indicated.

reactions would be important in deciding upon the characteristics of a miniaturized system for general use.

B.2 -- Input Amplifier (IC1, IC2)

IC1 is a fixed-gain (20dB) preamplifier intended primarily for microphone input or the low-level output of a tape player. It is capacitively coupled to the second stage, which is a variable-gain (20 dB maximum) inverting amplifier built around IC2.

Capacitive coupling was placed between stages in order to deal simultaneously with the voltage offset of the operational amplifier IC1 and especially that of one of the tape recorders used in our experiments, whose output included a substantial dc component. Considering the output impedance of IC1 to be negligibly small, we find that the low-frequency pole associated with the capacitive coupling occurs at approximately

$$f = \left[(2\pi)(10 \text{ k}\Omega)(2.2 \mu\text{F}) \right]^{-1} = 7.2 \text{ Hz} ,$$

which is low enough for our purposes.

B.3 -- Overdrive Indicator (IC3, Q1)

The proper operation of the pseudologarithmic compression circuit to be described below depends upon the peak amplitude at its input being slightly under 5 V. The output of the voltage comparator IC3 changes from -15 to +15 V as the output of IC2 increases beyond 5 V, whereupon Q1 goes from cutoff to saturation, permitting current to flow through the neon indicator lamp in its collector circuit.

In a miniaturized system it would probably be more appropriate to use a low-voltage LED indicator; the neon lamp and high-voltage power supply were used simply because they were immediately available in the breadboarding frame used in our experiments.

B.4 -- Pseudologarithmic Compression Amplifier (IC4, IC5)

The two diodes in the feedback path of the first operational amplifier (IC4) provide the logarithmic response for positive- or negative-going input voltages, as the case may be; when one conducts the other is reverse-biased and its effective resistance is large compared with that of the conducting diode.

The 18-K resistor in parallel with the two diodes is required to deal with input voltages at or near 0 V, since the log function is not defined for an argument of zero (Operational Amplifiers, 1971, 268). For very small input voltages the diode resistances are both large compared with 18 K; for zero input this amplifier has a voltage gain of

$$\lim_{V_{in} \rightarrow 0} \left(\frac{V_{out}}{V_{in}} \right) = - \frac{18 \text{ K}}{100 \text{ K}} = -0.18 .$$

As the input voltage increases in either direction, the forward resistance of the conducting diode becomes increasingly small with respect to 18 K, and the response becomes more closely logarithmic.

The second operational amplifier (IC5) has a fixed gain of -10. The two amplifiers in conjunction provide an output of -1.5 to +1.5 V for input voltages of -5 to +5 V, as shown earlier in Fig. 3.8.2.

B.5 -- Center Frequency and Frequency Deviation Control Circuit (IC6, IC7, IC8)

The input to this part of the system is derived from the pseudologarithmic compression amplifier or directly from the input amplifier, or from both in the desired

proportion as described in the last chapter. When the input gain is properly adjusted the peak voltage is the same (slightly below 1.5 V) at each end of the 100-K potentiometer, and the makeup of the voltage to be processed for control of the VCO is determined by the setting of this potentiometer.

After buffering by the voltage follower, IC6, the output gain is adjusted as desired by means of the variable-gain (zero to -10) inverting amplifier, IC7. The center frequency may be adjusted independently by means of a ten-turn potentiometer at the noninverting input of IC8, which is used as a unity-gain inverting amplifier. This potentiometer forms part of a voltage-divider circuit whereby the dc component of the output of IC8 may be set to any desired value from -1.36 to +1.36 V. The choice of this range was made in consideration of the characteristics of the outboard VCO's used in our experiments, and does not necessarily represent an appropriate choice for a miniaturized, self-contained, dedicated instrument.

B.6 -- Artificial Frequency-Foldover Circuit (IC9, IC10)

This circuit is a modification of a precision full-wave rectifier (Graeme 191973, 124), consisting essentially of a unity-gain inverting amplifier (IC9) and a voltage follower (IC10). Gating is accomplished by diodes placed in the feedback loops of their associated amplifiers to minimize error.

The voltage at which the foldover occurs is determined by the setting of the ten-turn 10-K potentiometer; whenever the input to the circuit drops below the foldover voltage at the noninverting input of IC9, it is reflected about the foldover voltage as shown in Fig. 3.6.6. When the system is used with an exponential VCO (for which artificial frequency foldover is unnecessary) the circuit is disabled by setting the foldover voltage to -15 V, in which case IC9 remains off and the circuit is reduced to a voltage follower.

B.7 -- VCO Control-Voltage Output Circuit (IC11, IC12)

In order to accommodate the many exponential outboard VCO's used in our experiments, two low-impedance outputs were provided. One is a positive-going control voltage at the output of the voltage follower IC11, and

the other is a negative-going control voltage of the same magnitude at the output of the unity-gain inverter built around IC12.

B.8 -- Amplitude-Modulation Control Circuit (IC13, IC14, IC15)

In normal circumstances the peak voltage at the output of the pseudologarithmic compression amplifier is 1.5 V. This voltage is amplified by an inverter (IC13) with a fixed gain of -6.67. In a miniaturized system a gain of -10 would have been preferred, but the limited output swing of the 709 operational amplifier used in our breadboarding system makes this impractical. The signal is then full-wave rectified by means of a gated unity-gain inverter (IC14) and voltage follower (IC15). The inverter's gain is matched to that of the voltage follower by adjusting the resistance in the feedback loop.

If it is desired to control the amplitude of the frequency-modulated VCO output according to the envelope of the pseudo-logarithmic compressor output, switch S1 is closed. The associated diode, capacitor, and 10-K resistor form an envelope-follower circuit with a time constant of

$$(10\text{ K})(2.2\mu\text{F}) = 22\text{ ms}.$$

If S1 is opened, the diode is never reverse-biased (since the full-wave rectifier yields a negative-going output) and these three components do nothing.

B.9 -- Noise Threshold Control Circuit (IC16, IC17)

The transfer function of this circuit was shown in Fig. 3.10.3. Ignoring the "tweaking" of gain necessitated in our breadboarded system by the output-swing limitations of the 709 operational amplifier, this circuit was designed as follows.

The active portion of the transfer function has the slope

$$m = \frac{30}{-E - 15} \quad (\text{B.9.1})$$

and its equation is found to be

$$V_{\text{out}} = m V_{\text{in}} + 15(m+1), \quad -15\text{ V} \leq V_{\text{in}} \leq E \quad (\text{B.9.2})$$

Arbitrarily selecting the maximum threshold voltage to be one-third of the peak amplitude, we have

$$0 \leq -\epsilon \leq 5V$$

whence from Eq. (B.9.1)

$$-3 \leq m \leq -2 \quad (\text{B.9.3})$$

We now rewrite Eq. (B.9.2) in a form more suitable for realization by analog circuitry:

$$V_{\text{out}} = m(V_{\text{in}} + 15) + 15, \quad -15V \leq V_{\text{in}} \leq \epsilon \quad (\text{B.9.4})$$

If we define

$$k = -\frac{m}{3},$$

then from Eqs. (B.9.3) and (B.9.4)

$$V_{\text{out}} = -3 \left[\frac{k}{3} (V_{\text{in}} + 15) - 5 \right], \quad \frac{2}{3} \leq k \leq 1$$

or

$$V_{\text{out}} = 3 \left[-k (V_{\text{in}} + 15) + 5 \right], \quad \frac{2}{3} \leq k \leq 1. \quad (\text{B.9.5})$$

This equation is tractable from the standpoint of analog-circuitry realization. In our circuit, the term

$$-k(V_{\text{in}} + 15)$$

is calculated by an inverting summing amplifier (IC16), whose gain (k) is adjustable by means of the noise threshold potentiometer; 5 V is added at the noninverting input of IC17, which is connected as a noninverting amplifier with a nominal gain of 3 in order to provide the required output.

The output of this circuit is -15 V when the system is quiescent, and approaches +15 V for the maximum input-voltage peaks when the input gain has been properly adjusted.

B.10 -- Adequate-Gain Indicator (IC18, Q2)

This circuit is quite similar to the overdrive-indicator circuit described above. The only difference is that the reference voltage at the inverting input of the comparator (IC18) is 12.5 V, so that as the peak input voltage is reached the comparator's output will change state, provided the input gain is sufficient to yield a voltage greater than 12.5 V at the output of the noise threshold control circuit, saturating Q2 and illuminating the neon indicator lamp.

B.11 -- Voltage-Controlled Amplifier (IC19, IC20)

This circuit is a two-quadrant multiplier built around an operational transconductance amplifier (IC19). The transconductance of this amplifier is controlled by the bias current supplied through the 27-K resistor; the current in turn depends upon the voltage at the output of the noise threshold control circuit. The system output voltage is developed across the 10-K resistor by the output current of the amplifier, and is buffered by a voltage follower (IC20). The 50-pF capacitor is included for stability; in conjunction with the 10-K resistor it produces a pole at

$$f = [(2\pi)(10\text{K})(50\text{pF})]^{-1} = 318\text{ kHz},$$

which has no effect upon the operation of the system.

This circuit as it stands has a dynamic range of 28.2 dB, determined by measurement of the input and output voltages for the positive extremes of the noise threshold control circuit output. Although a range of over 60 dB could be achieved by the addition of an n-p-n emitter follower (RCA ICAN-6668), it is felt that in this application it would be undesirable for the system output to become completely inaudible as the input voltage drops to zero.

B.12 -- The Portable FM Conversion System

A portable version of the system described above was designed and built in order that experiments could be made in hospitals or clinics. It was desired that the operation of the system be reasonably simple and reliable, so that medical personnel could familiarize themselves with the converted sounds by independent use of the system.

The portable system was built in three separately-packaged modules: a power supply, a control module, and a telephone-interface module containing a small loudspeaker and space for a standard telephone handset. Many comparatively minor design choices were made on the basis of what materials were readily available. In addition, several modifications were felt to be desirable; these will be outlined below.

Fig. B.12.1 shows the front-panel layout of the control module. Controls for input amplification and monitoring, frequency modulation, amplitude modulation, and output gain are grouped separately, as are input, system output, and what we may term "diagnostic" output connectors.




FIGURE B.12.1

Front panel of the portable FM conversion system.

In order that the system be able to stand alone with various types of transducers or signal sources, the input amplifier was modified as shown in Fig. B.12.2. Provision for the connection of an outboard filter was provided as well as a simple unity-gain inverter (not shown in the figure) which can be switched into the circuit so that pitch changes will occur in the desired direction, whatever the phase convention relating the transducer voltage to actual transducer displacement. In general, a rise in pitch corresponding to outward motion of the surface of the skin seems to be preferred.

A provision for preset center frequency and modulation depth was added to the circuit controlling center frequency and frequency deviation described in Section B.5. This modification, shown in Fig. B.12.3, is intended to permit immediate selection of parameters appropriate to optimum transmission of converted sounds via standard telephone channels. The center frequency in this mode is about 1581 Hz, and the modulation depth is such that the output spectrum is mainly confined to frequencies between 500 Hz and 5 kHz (see Section 3.11).




FIGURE B.12.2

Input amplifier for the portable FM conversion system. The maximum voltage gain is 5000 (74 dB).




FIGURE B.12.3

FM center frequency and modulation depth control circuit for the portable FM conversion system. Preset conditions may be chosen for optimum transmission of converted sounds via a standard telephone channel.

Since the system contains its own (linear-response) voltage-controlled oscillator, an exponential-response amplifier was included (see Fig. B.12.4) to drive the VCO, which is shown in Fig. B.12.5 together with its input scaling and shifting circuit and the voltage follower used to buffer its output.

A graphic display output was added to provide a visual analog, with the aid of an oscilloscope or graphic recorder, of the sounds generated by the system. Its purpose is to facilitate the familiarization with converted sounds by association of those sounds with a visual counterpart; it was motivated partly by the fact that some of these sounds contain information previously available, if at all, only in visual form.

The graphic display signal is generated, as shown in Fig. B.12.6, by amplitude-modulation of the signal at the input to the exponential-response amplifier by the output of the noise threshold circuit. Each of these component signals is available at a front-panel connector. The first stage of the circuit is a level translator whose output range (roughly -14 to $+14$ V) is suitable for amplitude modulation of the input to the second stage.




FIGURE B.12.4

Exponential-response amplifier for the portable FM conversion system.




FIGURE B.12.5

Voltage-controlled oscillator circuit for the portable FM conversion system.




FIGURE B.12.6

Graphic display output circuit for the portable FM conversion system. The output consists of the voltage at the input to the exponential amplifier of Fig. B.12.4 amplitude-modulated by the output of the noise threshold circuit.

Fig. B.12.7 shows examples of the AM and graphic display output signals. The input was derived from a recording of normal heart sounds by high-pass filtering at 40 Hz to retain only audible frequencies. The AM signal (second trace) shows greater activity at low amplitudes because of pseudologarithmic compression, while the third trace shows the effect of raising the noise threshold; the graphic display output signal is shown in the fourth trace.

A second example is shown in Fig. B.12.8. Here the input was a recording of audible heart sounds containing a systolic murmur, which appears in the first trace as slightly elevated amplitude between the first (S1) and second (S2) heart sounds. The remaining three traces show the effects on the graphic output signal of pseudologarithmic compression and suppression of low amplitudes by means of the noise threshold control. These features allow the operator to distinguish readily between the murmur and the baseline noise in both the visual analog and the converted sounds themselves.

The telephone interface module contains a small loudspeaker driven by the single-stage transformer-




FIGURE B.12.7

Operation of the noise threshold circuit. (a) Normal heart sounds, audible frequencies; (b) AM output, no noise suppression; (c) AM output with raised noise threshold; (d) graphic display output signal corresponding to noise suppression as shown in (c).




FIGURE B.12.8

Operation of the pseudologarithmic compressor. (a) Heart sounds (audible frequencies) with systolic murmur, shown by elevated amplitude between S1 and S2; (b) graphic display output with moderate pseudologarithmic compression; (c) graphic display output with pronounced pseudologarithmic compression; (d) graphic display output as in (c) but with raised noise threshold.

coupled power amplifier whose circuit is shown in Fig. B.12.9. Although this circuit is unsuitable for high-quality sound production of wideband converted sounds (no matter how good the loudspeaker used), it has been found entirely adequate for sending appropriately "tailored" sounds over standard telephone channels as we have described above.




FIGURE B.12.9

Transformer-coupled power amplifier for the portable FM conversion system. This stage drives a small loudspeaker in the telephone interface module.

A P P E N D I X C

OPERATING MANUAL FOR THE FM CONVERSION SYSTEM

This appendix consists of an operating manual prepared for use by medical and paramedical personnel working with the portable FM conversion system built during the course of the research. Several recorded examples of heart sounds processed by the system may be found on the tape cassette accompanying the dissertation; they are described in Appendix E.

OPERATING NOTES

FREQUENCY-CONVERSION SYSTEM
FOR BIOMEDICAL SIGNALS

Alan D. Bernstein
Department of Physics
Rutgers University
Newark, NJ 07102

May 1978
Revised June 1978

INTRODUCTION

The system discussed below was designed primarily for use in connection with heart sounds, but may also be used to obtain easily audible sounds containing the same information as in other signals of audio or subaudio frequency. It consists of three interconnected modules: a power supply, a control unit, and a telephone interface unit containing a small loudspeaker. Its basic function is the conversion of changes in signal amplitude into changes in pitch.

OVERALL DESCRIPTION

Figure 1 shows a block diagram of the system; each of the basic elements appearing in the diagram will be discussed in turn.

Input Amplifier

The input is ac-coupled so that dc offset in the input signal will not affect system operation. This feature is incorporated in the design mainly so that baseline shifts will not be troublesome when the system is used with electrocardiographic signals. Light-emitting diodes are provided to aid in adjusting the input gain for proper system operation.

Pseudologarithmic Compressor

This circuit serves a dual purpose. In consideration of the nonlinear amplitude-response characteristics of the human auditory system, it is used to control the system output amplitude through amplitude modulation in such a way that the output




FIGURE 1
BLOCK DIAGRAM

amplitude rises synchronously with, but faster than, the input amplitude. In this way audibility is increased during low-amplitude portions of the input signal. In addition, this circuit may be used via the response-control potentiometer to emphasize the changes in output pitch caused by changes in input amplitude occurring at low or high amplitudes as desired, to facilitate concentration on high- or low-level changes in the input signal.

Exponential Voltage-Controlled Oscillator

This circuit performs the conversion of amplitude changes into pitch changes by means of frequency modulation. Its exponential response to input voltages is of paramount importance, since it ensures that equal changes in input amplitude cause equal pitch intervals in the output (not equal frequency shifts), in consideration of the nonlinear relationship between fundamental frequency and perceived pitch characteristic of human hearing.

Noise Threshold Circuit

This circuit allows the operator to choose the lowest input signal amplitude at which the output signal amplitude begins to rise. It may be used to limit system response to noise inherent in the input signal or to facilitate concentration on high-amplitude signal components, such as S1 and S2 in heart sounds or the R wave in an electrocardiographic signal.

Voltage-Controlled Amplifier

This circuit amplitude-modulates the output of the voltage-controlled oscillator according to the instantaneous amplitude of the input signal after pseudologarithmic compression, subject to any "special instructions" from the noise threshold circuit. It is designed so that the output signal is suppressed but not eliminated below threshold, so that below-threshold events of possible significance will not intrude but will be less likely to go unnoticed.

Power Amplifier

This is a modest power amplification stage whose job is to drive the small loudspeaker in the telephone interface module. A wide-band external loudspeaker may be used with a high-quality audio amplifier driven from the "system audio out" jack.

Telephone Interface Module

When the system is in "telephone" mode, the frequency band occupied by the output signal is internally calibrated to match that of standard telephone channels, and the system may be used with the telephone interface to transmit over the telephone all information inherent in the original signal.

No equipment is necessary for listening at the receiving end, although it may be found helpful in analysis to record the received signal for later playback at slower speed. This procedure is impractical with unprocessed heart sounds since low frequencies would be lost. However, the converted sounds may be reduced without losses even to quarter-speed with the most modest recording equipment.

INPUTS AND OUTPUTS


INPUT	System input; gain is adjustable over three amplitude ranges.
INPUT MONITOR	Signal output following input amplifier.
FM OUT	For monitoring signal used for frequency modulation; reflects setting of amplitude-response control.
AM OUT	For monitoring signal used for amplitude modulation; reflects pseudo-logarithmic compression and setting of noise threshold control.
SYSTEM AUDIO OUT	Output to high-impedance headphones, tape recorder, or high-quality audio amplifier. Amplitude adjustable by means of output gain control.
SPEAKER OUT	Output to telephone interface module (system should be in "telephone" mode). Loudness adjustable by means of output gain control. Not for use with wide-band external loudspeaker with system in "variable" mode; for this purpose use a high-quality audio amplifier driven from the system audio output.

CONTROLS AND OPERATION

- INPUT GAIN For proper system operation, adjust so that "adequate gain" lamp flashes frequently, but "reduce gain" lamp does not.
- GAIN RANGE Determines maximum input amplifier gain: low (L), medium (M), or high (H). Set at lowest range in which adequate gain can be obtained by advancing "input gain" control from fully counterclockwise position.
- INVERT May be used to reverse polarity of the input signal to accomodate various transducers. This option ensures that the pitch changes accompanying transducer motion will occur in the desired direction.
- CENTER FREQ. Sets oscillator frequency about which FM takes place; determines central pitch of output signal.
- MODULATION DEPTH Determines degree of deviation from central pitch caused by instantaneous amplitude changes in input signal.
- FM BANDWIDTH In "variable" mode the center frequency and modulation depth controls are active. In "telephone" mode these parameters are fixed internally for transmission of the output signal over a standard telephone channel via the telephone interface module.
- AMPLITUDE SENS. In fully clockwise position, the system responds to equal instantaneous changes in the input signal with equal pitch intervals. As the control is rotated counterclockwise, changes at low amplitudes produce increasingly greater pitch variations.

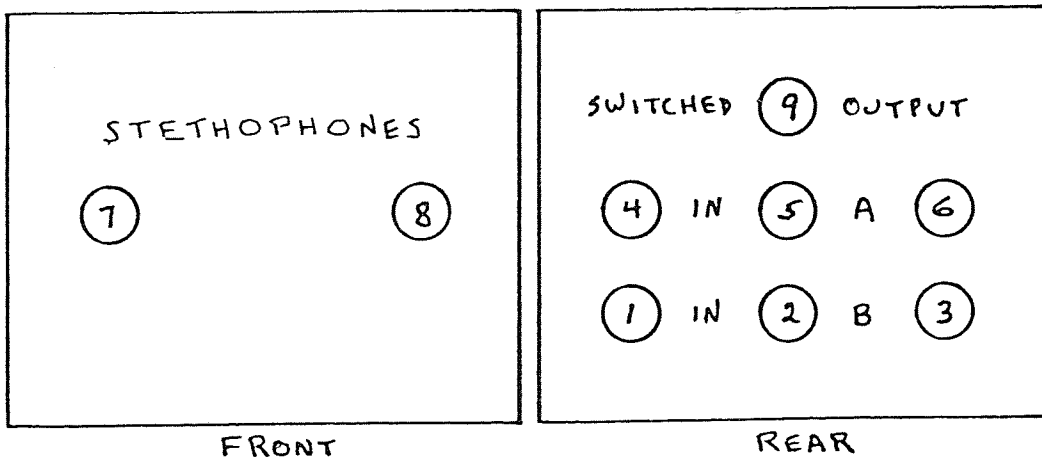

AM NOISE THRESH. Determines minimum instantaneous amplitude of input signal at which output amplitude begins to rise. May be advanced from fully counter-clockwise position to de-emphasize tape noise and/or low-amplitude events in the input signal.

OUTPUT GAIN Determines output signal amplitude at the "system audio out" and "speaker out" jacks.



SUGGESTED PATCH FOR FAMILIARIZATION
WITH SYSTEM OPERATION

STETHOPHONE SWITCHING BOX FOR USE
WITH FREQUENCY-CONVERSION SYSTEM



A P P E N D I X D

MEASUREMENT OF THE ACTIVITY COEFFICIENT FOR NORMAL SPEECH

It was pointed out in Section 6.3 that the average number of breakpoints per second needed to define an acceptable PL approximation to a speech waveform could, in practice, be expected to be lower than the equivalent value for continuous speech, simply because normal speech is not continuous. In the course of that discussion, the "activity coefficient" for speech was defined as the ratio of "active" time (actual speech) to the duration of the entire record, which includes the pauses between utterances. The activity coefficients shown in Table 6.3.1 were obtained by means of energy measurements not unlike those which may be used to distinguish between voiced and unvoiced speech components (Schafer and Rabiner, 1975). A block diagram of the circuit used for our measurements is shown in Fig. D.1.

The energy in a signal $v(t)$ of duration T may be written as

$$E = \int_0^T v^2(t) dt .$$




FIGURE D.1

Block diagram of a circuit for measurement of the activity coefficient for speech.

The arrangement shown in Fig. D.1 is conceived along these lines. The signal is first bandlimited to exclude frequencies outside the speech bandwidth. It is then multiplied by itself, resulting in a new signal containing low-frequency components associated mainly with voiced speech, components which were absent because of the initial bandlimiting.

Integration over a finite interval is simulated by a one-pole low-pass filter whose positive-going output is then compared with a suitably-chosen threshold voltage. Two digital timers are used: one (operated manually) to measure the total record time, and a second to measure the total "active" time during that interval. By means of an AND gate, the second timer is enabled only while the first is running, and counts during that period only when the integrator output exceeds the chosen threshold.

The major difference between our measurements and those appropriate to the separation of voiced and unvoiced speech segments is that in our case the threshold is set just above the ambient noise level so as to include in the "active" time measurement as many of the unvoiced speech components as possible. The integrator

time constant was chosen so that the comparator would respond as quickly as possible to changes in signal energy without oscillating in response to low-frequency components in the squared signal. It is also partly for this reason that some low frequencies are excluded from the input signal.

A P P E N D I X E

DESCRIPTION OF THE RECORDED EXAMPLES

E.1 -- Conversion of Heart Sounds by Frequency Modulation

Several examples have been prepared to illustrate the conversion of heart sounds by means of the technique described in Chapter 3 and Appendices B and C. All have been recorded at approximately the same level, but the unprocessed heart sounds in Bands 1 and 3, because of their low frequency, will probably be largely inaudible if the recordings are played using portable equipment.

Band 1. Normal heart sounds from a commercial recording; audible frequencies.

Band 2. Normal heart sounds as in Band 1, converted using a 500-Hz carrier.

Band 3. Abnormal heart sounds from a commercial recording; audible frequencies. The patient has mitral stenosis; there is a slight crescendo presystolic murmur immediately preceding the normal first sound, and an opening snap following the normal second sound. The

temporal relationships among these components are shown schematically in Fig. E.1.

Band 4. Abnormal heart sounds as in Band 3, converted using a 500-Hz carrier and slowed to half-speed.

Band 5. Normal heart sounds, including all frequencies above 10 Hz, converted using a 500-Hz carrier.

Band 6. Normal heart sounds, including all frequencies above 10 Hz, converted for telephone transmission as described in Section 3.11.

Band 7. Normal heart sounds, including all frequencies above 10 Hz, converted using a 2-kHz carrier and slowed to quarter-speed.

E.2 -- Adaptive PL Sampling of Speech

As explained in Section 6.1, the microcomputer-controlled analog interpolator implemented for use in our experiments is not sufficiently fast at present to be used for reconstruction of PL approximations of speech waveforms, even at quarter-speed. However, the adaptive PL sampling program itself can be used in these circum-




FIGURE E.1

Schematic representation of the abnormal heart sounds recorded on Band 3. A presystolic murmur (M) precedes the normal first sound (S1), and the normal second sound (S2) is followed by an opening snap (S).

stances, and was tried with an extreme example of speech containing a large proportion of sibilants and "noisy" consonants: the phrase "this is an example of recorded speech." The curvature sensitivity was set well beyond that needed to obtain an approximation of the quality specified in Section 6.2, yet the average number of breakpoints chosen per second was found to be about 17% of the Nyquist rate for this case, so that the rate of data storage was about 34% of the Nyquist rate, since two numbers are required to define each breakpoint.

If our system could reconstruct the PL approximation in this case, we would expect that approximation to sound very much like the original signal. We felt that it might be of interest to prepare an example showing the results of uniform sampling at a frequency equivalent to the average number-storage rate needed for high-resolution adaptive PL sampling of this speech signal.

Band 8. Original speech signal, bandlimited to the two decades extending from 76 Hz through 7.6 kHz.

Band 9. Uniform sampling of the signal recorded in Band 8. The sampling frequency was 5104 Hz, about 34% of

the Nyquist rate, and no anti-aliasing precautions were taken.

Band 10. Uniform sampling of the signal recorded in Band 8. The sampling frequency was 5104 Hz as before, but the signal was low-pass filtered at half the sampling rate (i.e., at about 2552 Hz) before and after sampling. As a result, aliasing has been largely eliminated (although no guard band was provided), but at the expense of eliminating a substantial portion of the high-frequency content of the signal and seriously degrading its intelligibility.

R E F E R E N C E S

"Applications of the CA3080 and CA3080A High-Performance Operational Transconductance Amplifiers," Applications Note ICAN-6668, Somerville, NJ (RCA), 1971.

Backus, J., The Acoustical Foundations of Music, New York (W. W. Norton), 1969.

Beranek, L., Acoustics, New York (McGraw-Hill), 1954.

Bernstein, A. D., and E. D. Cooper, "The Piecewise-Linear Technique of Electronic Music Synthesis," J. Audio Eng. Soc. 24:6, 1976.

Bukstein, E., Introduction to Biomedical Electronics, Indianapolis, Indiana (H. W. Sams), 1973.

Butterworth, J. S., M. R. Chassin, R. McGrath, and E. H. Reppert, Cardiac Auscultation, 2nd ed., New York (Grune and Stratton), 1960.

Chowning, J., "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation," J. Audio Eng. Soc. 21:526, 1973.

Computer Handbook, H. D. Huskey and G. A. Korn (eds.),
New York (McGraw-Hill), 1962.

Craige, E., "Apexcardiography," in Noninvasive Cardiology, A. M. Weissler (ed.), New York (Grune and Stratton), 1974.

Everitt, W. L., and G. E. Anners, Communication Engineering, New York (McGraw-Hill), 1956.

Flanagan, J., Speech Analysis, Synthesis and Perception,
2nd ed., New York (Springer-Verlag), 1972.

Fletcher, H., and W. Munson, "Loudness, its Definition,
Measurement and Calculation," J. Acoust. Soc. Am.
5:82, 1938.

Graeme, J., Applications of Operational Amplifiers, New
York (McGraw-Hill), 1973.

Holloway, G. A., W. E. Martin, D. W. Watkins, "A Clinical
Instrument for Translation of Heart Sounds to Mid-
Auditory Frequencies," Proc. 28th Ann. Conf. on
Engineering in Medicine and Biology, Sep. 20-24,

1975, Chevy Chase, Maryland (Alliance for Engineering in Medicine and Biology), 1975.

Hutchins, B., "The Frequency Modulation Spectra of an Exponential Voltage-Controlled Oscillator," J. Audio Eng. Soc. 23:200, 1975.

"IC Programmable Multiplier Divider Computation Circuit" (specification sheet for the AD531 integrated circuit), Norwood, Massachusetts (Analog Devices), undated.

Karp, T., "Develop Software with Flowgrams," Electronic Design 16, Aug. 2, 1978.

Kesteloot, H., J. Willems, and E. van Vollenhoven, "On the Physical Principles and Methodology of Mechanocardiography," Acta Cardiologica 24:147, 1969.

Moorer, J., "The Use of the Phase Vocoder in Computer Music Applications," presented at the 55th Convention, Audio Eng. Soc., Oct.-Nov. 1976.

- Operational Amplifiers, G. Tobey, J. Graeme, and L. Huelsman (eds.), New York (McGraw-Hill), 1971.
- Panter, P., Modulation, Noise, and Spectral Analysis, New York (McGraw-Hill), 1965.
- Plomp, R., and M. Bouman, "Relation between Hearing Threshold and Duration of Tone Pulses," *J. Acoust. Soc. Am.* 31:749, 1959.
- Ravin, A., Auscultation of the Heart, 2nd ed., Chicago (Year Book Medical Publishers, Inc.), 1967.
- Ravin, A., descriptive material accompanying heart sound recordings distributed by Merck, Sharp and Dohme, West Point, PA, 1968.
- Roederer, J., Introduction to the Physics and Psychophysics of Music, New York (Springer-Verlag), 1973.
- Schafer, R., and L. Rabiner, "Digital Representations of Speech Signals," *Proc. IEEE* 63:4, 1975.

Segall, H. N., "A Simple Method for Graphic Description of Cardiac Auscultatory Signs," *Am. Heart J.* 8:533, 1933.

Spiegel, M. R., Mathematical Handbook, New York (McGraw-Hill), 1968.

Stearns, S. D., Digital Signal Analysis, Rochelle Park, NJ (Hayden), 1975.

Stevens, S., "The Attributes of Tones," *Proc. U. S. Natl. Acad. Sci.* 20:457, 1934.

Stockham, T., T. Cannon, and R. Ingebretson, "Blind Deconvolution through Digital Signal Processing," *Proc. IEEE* 63:4, 1975.

Van Valkenburg, M. E., Network Analysis, 2nd ed., Englewood Cliffs, NJ (Prentice-Hall) 1964.

C U R R I C U L U M V I T A E

(TO APRIL 1979)

NAME: Alan D. Bernstein

EDUCATION

New Jersey Institute of Technology, Newark, NJ:
M.Sc. (electrical engineering), May 1975:

Bowdoin College, Brunswick, ME:
B.A. (music history and theory) magna cum laude
in musica, 1960.

Studies Subsequent to B.A.

- 1970-1978: Undergraduate and graduate courses in electrical engineering, graduate courses in mathematics; NJIT.
1968-1969: Graduate courses in physics and mathematics; Rutgers, New Brunswick.
1964-1967: Audited undergraduate courses in physics and mathematics; Hebrew University, Jerusalem.
1961-1963: Undergraduate courses in mathematics, physics, and chemistry; Rutgers, Newark.
1960-1961: Graduate courses in musicology; NYU.

Teaching

- 1970-1979: Instructor in physics; Rutgers, Newark. Course planning and lecturing: musical acoustics; course planning, lecturing, and laboratory instruction: introductory physics (statics, dynamics, waves, optics, electrostatics, passive circuits); some course planning, lecturing, and laboratory instruction: instrumentation, introductory analog and digital circuits.
1970-1977: Director, General Physics Laboratories; Rutgers, Newark.
1969-1970: Assistant instructor in physics; Rutgers, Newark.

- 1968-1969: Teaching assistant in physics, including two summer appointments; Rutgers, Newark.
- 1960-1961: Graduate assistant in musicology; some undergraduate-level teaching; NYU.
- 1958-1960: Student assistant in music; some undergraduate-level teaching; Bowdoin.

Research

- 1975-1979: Doctoral research, signal processing techniques for analysis of heart sounds and electrocardiograms (research assistantship, 1977-1979); NJIT.
- 1974-1975: Master's research, problems in the measurement of perceived fundamental frequency; NJIT.
- 1961-1963: Research assistantship, thermogravimetric telemetry of isolated systems; Rutgers, Newark.

ADDITIONAL TECHNICAL EXPERIENCE

Subsequent to B.A.

- 1970-1978: Informal consultation in electronic music, design of electronic musical instruments, sound distribution systems, and instrumentation for research in visual perception and ethnomusicology.
- 1968-1978: Design and construction of apparatus for demonstrations and undergraduate experimentation in physics; Rutgers, Newark.
- 1964-1967: System design, special recording techniques, electronic-music studio engineering, melographic studies of vocal music; Israel Center for Electronic Music, Hebrew University, Jerusalem.
- 1963-1964: Sound system and radio repair and maintenance; Kibbutz Hazorea, Israel.
- 1960-1961: Sound system repair and maintenance; NYU.

Prior to B.A.

Sound system installation, repair, and maintenance; college radio station studio engineering and maintenance; lighting and stagecraft.

MUSICAL EXPERIENCE

Piano, double bass, guitar, carillons, organ, five-string banjo; modern jazz quintet (Bowdoin), choral arranging and conducting, solo and ensemble singing, musical comedy score composition and musical direction, piano and organ tuning.

COMPUTER PROGRAMMING

FORTRAN, Z-80 Assembler, ATS, APL, CSMP, Superwylbur.

FOREIGN LANGUAGES

Hebrew (fluent), French, German.

MEMBERSHIP

IEEE, AAPT, Sigma Xi, Catgut Acoustical Society.

PUBLICATIONS

Experiments in College Physics, Rutgers NCAS, Department of Physics, 1970; revised 1971, 1974.

"The Piecewise-Linear Technique of Electronic Music Synthesis" (with Ellis D. Cooper), Journal of the Audio Engineering Society, July/August 1976.

"Listening to the Coefficient of Restitution," American Journal of Physics, January 1977.

"Tuning the Ill-Tempered Clavier," American Journal of Physics, August 1978.

"Improvement in the Detection and Analysis of Acoustical Information" (with Mauro Zambuto), presented at the Thirteenth Congress of UNIATEC in Paris on September 25, 1978 (in process of publication); also published as "Miglioramento della Rivelazione e dell'Analisi della Informazione Acustica," Note di Tecnica Cinematografica, Trimesters 1-2, 1978.