

Old Dominion University

ODU Digital Commons

Electrical & Computer Engineering Theses & Dissertations

Electrical & Computer Engineering

Spring 1995

Fetal Heart Rate Detection with a Passive Acoustic Sensor System

Zongyao Zhou
Old Dominion University

Follow this and additional works at: https://digitalcommons.odu.edu/ece_etds



Part of the [Biomedical Commons](#), [Biomedical Devices and Instrumentation Commons](#), [Electrical and Electronics Commons](#), and the [Signal Processing Commons](#)

Recommended Citation

Zhou, Zongyao. "Fetal Heart Rate Detection with a Passive Acoustic Sensor System" (1995). Master of Science (MS), Thesis, Electrical & Computer Engineering, Old Dominion University, DOI: 10.25777/6bb2-h805
https://digitalcommons.odu.edu/ece_etds/578

This Thesis is brought to you for free and open access by the Electrical & Computer Engineering at ODU Digital Commons. It has been accepted for inclusion in Electrical & Computer Engineering Theses & Dissertations by an authorized administrator of ODU Digital Commons. For more information, please contact digitalcommons@odu.edu.

FETAL HEART RATE DETECTION
WITH
A PASSIVE ACOUSTIC SENSOR SYSTEM

Zongyao Zhou

A thesis submitted to the Faculty of
Old Dominion University in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE
ELECTRICAL ENGINEERING

OLD DOMINION UNIVERSITY
May 1995

Approved by:

Stephen A. Zahorian (Director)

John W. Stoughton

Allen Zuckerwar

ABSTRACT

Fetal Heart Rate Detection With A Passive Acoustic Sensor System

Zongyao Zhou

Old Dominion University, 1995

Director: Dr. Stephen A. Zahorian

Research and development is presented of real time signal processing methodologies for the detection of fetal heart tones from noise contaminated signals obtained from passive acoustic sensors. A nonlinear Teager energy operator is utilized for detection of the fetal heart tone event. Autocorrelation and a parallel redundancy correction algorithm derives fetal heart rates. A real time monitoring system is described which records and plots the time history of both the fetal heart rate and the acoustic fetal heart signal. The system is validated in the context of the fetal nonstress test. Comparisons are made with ultrasonic nonstress tests on a series of patients. Comparative data show the clinical reliability of the system for real time detection of fetal heart rates.

ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Stephen A. Zahorian, for his invaluable guidance and advice. His continuous encouragement, availability and support are very much appreciated.

I would also like to thank Dr. David Livingston for his technical support and advice. I would also like to thank the additional members of my thesis advisory committee, Dr. John W. Stoughton and Dr. Allan Zuckerwar, for their generous assistance and time. I also wish to thank NASA for funding this project and EVMS for allowing pilot testing.

With special appreciation to my family for all their love and thoughtfulness.

TABLE OF CONTENTS

LIST OF TABLES	vii
LIST OF FIGURES	viii
INTRODUCTION	1
1.1 Fetal Heart Rate Monitoring	1
1.2 Background of the NASA/ODU Project	2
1.3 Thesis Research Objective	3
1.4 Thesis Organization	5
BACKGROUND AND THEORY	7
2.1 Introduction	7
2.2 Review of Previous Research Work	7
2.2.1 FPCG Signal Characteristics	8
2.2.2 Previous System Performance	12
2.3 Previous Heart Beat Event Detection Methodologies	14
2.3.1 LMS Linear Prediction Signature Matching	16
2.3.2 Average Magnitude Cross Difference	17
2.3.3 Digital Matched Filtering	19

2.3.4 Autocorrelation	20
2.4 Heart Beat Event Detection Methodologies for This Research ...	21
2.4.1 Summary of Detection Methodologies	22
2.4.2 Teager Energy Operator	24
2.4.3 Neural Network	27
FETAL HEART RATE MONITORING SYSTEM	33
3.1 Introduction	33
3.2 Hardware	33
3.2.1 Overview	34
3.2.2 Acoustic Sensor Belt	36
3.2.3 Electronics Support System	39
3.3 Digital Signal Processing	46
3.3.1 Overview	46
3.3.2 Sampling Data	47
3.3.3 Teager Energy Operator	55
3.3.4 Autocorrelation	56
3.3.5 Intelligent Figure of Merit	59
REAL TIME IMPLEMENTATION	62
4.1 Introduction	62
4.2 Software Overview	62

4.3 A View From The ELF-C31	63
4.3.1 SPOX DSP Operating System	65
4.3.2 Main Routine	72
4.3.3 Initialization Routine	74
4.3.4 Sampling Routine	77
4.3.5 Teager Energy Routine	79
4.3.6 Neural Network Routine	79
4.3.7 Autocorrelation Routine	82
4.4 A View From The PC	85
4.4.1 Initialization Routine	87
4.4.2 Main Loop Routine	88
4.4.3 Serial Communication Routine	89
4.4.4 Push-Button Routine	91
4.5 PC-ELF Handshaking and Data Communication	91
4.6 Memory And Timing	93
EXPERIMENTS AND CONCLUSIONS	99
5.1 Introduction	99
5.2 Technical Testing	99
5.2.1 Software Verification	100
5.2.2 Balloon Testing	103

5.3 Clinical Testing	109
5.3.1 Procedure	109
5.3.2 Ultrasound Comparative Study	121
5.4 Conclusions	131
BIBLIOGRAPHY	134
APPENDIX	136

LIST OF TABLES

Table 2.1 Fetal Heart Frequencies per Available Literature	11
Table 3.1 124th Order Equal Ripple Digital Bandpass Filter with Sampling Rate at 500 Hz	52
Table 4.1 RS232 Transmitting Format of Fetal Heart Rates In The Update Mode of Corometrics 116 Unit	90
Table 4.2 Parallel Printer Port Status Line Settings	92
Table 4.3 System Buffers and Memory	95
Table 4.4 Timing Results For Each Routine within The Main Loop	97

LIST OF FIGURES

Figure 1.1 Previous Hardware System on A Cart	4
Figure 2.1 Fetal Heart Tone Localization	9
Figure 2.2 FPCG Signal	10
Figure 2.3 Fetal Heart Tone Spectrum From Pretlow [1]	13
Figure 2.4 Acoustic vs. Ultrasonic NST : Patient 23 of Pretlow[1]	15
Figure 3.1 System Hardware Block Diagram	35
Figure 3.2 Cross-Sectional and Cutaway Views of The Sensor Belt	38
Figure 3.3 Electronics Support System Hardware Block Diagram	40
Figure 3.4 Multiplexer Switch Control Signals And Resultant TDM Signal	44
Figure 3.5 Frequency Response of Digital Bandpass Filter Two	53
Figure 3.6 Impulse Response of Digital Bandpass Filter Two	54
Figure 4.1 SPOX Software Architecture	66
Figure 4.2 Multiple Vector Views	69
Figure 4.3 Matrix View of an Array	71
Figure 4.4 Main Routine	73
Figure 4.5 Initialization Routine	75

Figure 4.6 Sampling Routine	78
Figure 4.7 Teager Energy Routine	80
Figure 4.8 Three Seconds of Acoustic Fetal Heart Signal (Top) and Corresponding Teager Energy (Bottom)	81
Figure 4.9 Three Seconds of Acoustic Fetal Heart Signal (Top) and Corresponding Neural Network Output (Bottom)	83
Figure 4.10 Autocorrelation Routine	84
Figure 4.11 PC System Software Block Diagram	86
Figure 4.12 Structure of Keyboard Command Register	94
Figure 5.1 Fetal Heart Rate Chart For : Patient 23 (Teager Energy Operator)	101
Figure 5.2 Fetal Heart Rate Chart For : Patient 23 (Neural Network) ...	102
Figure 5.3 Simulation Experimental Set Up for The Balloon Test	104
Figure 5.4 Spectrum of A Simulation Fetal Heart Signal in Balloon Test	105
Figure 5.5 Balloon Test Results	107
Figure 5.6 Balloon Test Results (Continued)	108
Figure 5.7 Six Second Acoustic Fetal Heart Beat Signal of Pt1_615	111
Figure 5.8 Six Second Acoustic Fetal Heart Beat Signal of Pt2_615	112
Figure 5.9 Six Second Acoustic Fetal Heart Beat Signal of Pt3_615	113
Figure 5.10 Six Second Acoustic Fetal Heart Beat Signal of Pt2_909	114
Figure 5.11 Six Second Acoustic Fetal Heart Beat Signal of Pt5_1021	115

Figure 5.12 Average Spectrum of Fetal Heart Signal From Pt1_615	116
Figure 5.13 Average Spectrum of Fetal Heart Signal From Pt2_615	117
Figure 5.14 Average Spectrum of Fetal Heart Signal From Pt3_615	118
Figure 5.15 Average Spectrum of Fetal Heart Signal From Pt2_909	119
Figure 5.16 Average Spectrum of Fetal Heart Signal From Pt5_1021	120
Figure 5.17 Ultrasonic NST of Pt1_615	122
Figure 5.18 Acoustic NST of Pt1_615	123
Figure 5.19 Ultrasonic NST of Pt2_615	124
Figure 5.20 Acoustic NST of Pt2_615	125
Figure 5.21 Ultrasonic NST of Pt3_615	126
Figure 5.22 Acoustic NST of Pt3_615	127
Figure 5.23 Ultrasonic NST of Pt2_909	128
Figure 5.24 Acoustic NST of Pt2_909	129
Figure 5.25 Acoustic NST of Pt5_909	130
Figure I System Hardware Connection	136

CHAPTER ONE

INTRODUCTION

1.1 Fetal Heart Rate Monitoring

In 1818 a physician detected the fetal heartbeat by listening to a mother's abdomen. In 1833 a textbook on "Obstetric Auscultation" noted the possible relationship between fetal heart rate (FHR) patterns and fetal health [1]. In 1906 Cremer first measured the fetal electrocardiogram (FECG) by using external abdominal electrodes [2]. Since then FHR monitoring has been clinically utilized as a means of assessing fetal health. FHR monitoring attempts to capture the rhythm of the fetal heart valves opening and closing. If the fetal heart beats at a healthy rate during and following the mother's contractions, it is an indicator that the fetus is coping well with the stress of labor. Furthermore, fetal cardiograms can predict fetal distress so that doctors can prevent them before irreversible damage has occurred.

Worldwide, there are two methods which are primarily used for fetal monitoring. These are fetal electrocardiogram (FECG) techniques and fetal phonocardiogram (FPCG) techniques. FECG techniques detect the electrical activity of the fetal heart internally, typically using electrodes attached to the scalp of the fetus, or externally, using Doppler shift ultrasound methods. The

internal FECG technique is available only when mothers are committed to delivery. Even the external ultrasound method is considered as an "invasive" technique, since the fetus is exposed to doses of ultrasound waves. However, studies to date have failed to show any risks from this technique. The FPCG techniques detect the fetal heartbeat sounds or fetal heart "tones" by means of a passive microphone attached to the maternal abdomen. The FPCG techniques were used in early fetal monitors but eventually were abandoned in favor of the FECG techniques because of the greater sensitivity of the latter [1]. Compared to FECG techniques, the FPCG techniques have a much lower signal to noise ratio (SNR) due to low acoustic energy in the fetal heart tones and relatively high energy sources of other acoustic signals. However, the disadvantages of the FECG techniques are the cost and complexity of the necessary equipment in addition to the invasive nature of the monitoring. Therefore, fetal monitoring has to be performed in a clinic or physician's office, and only limited monitoring is possible. The main advantages of the acoustic technique are its passivity (non-invasiveness) and its simplicity, thus opening the door to long-term home monitoring.

1.2 Background of the NASA/ODU Project

In 1986 Donald Baker (M.D.) of Spokane, Washington made contact with the Technology Utilization Division of NASA and suggested that long term fetal monitoring of the mother at home might prevent a significant number of

fetal problems such as hypoxic brain damage. He further suggested that a passive acoustic device utilizing the FPCG signal would obviate the invasive risks inherent with long term ultrasonic monitoring. NASA agreed to fund a project to develop Baker's idea into a working monitor. NASA would fabricate the sensors for the monitor utilizing a polyvinylidene fluoride acoustic film used to detect vibrations on wind tunnel models. The Department of Electrical and Computer Engineering at Old Dominion University was enlisted to develop the signal processing hardware and software. The Eastern Virginia Medical School agreed to facilitate the clinical data collection and testing of the acoustic monitor.

In the first phase of the project, Pretlow and Stoughton ([1] and [3]) developed a prototype monitor and demonstrated the feasibility of Dr. Baker's idea. The second phase of the project is to improve the monitor and develop a working portable medical research instrument.

1.3 Thesis Research Objective

The previous fetal heart rate monitor implemented by Pretlow is shown in figure 1.1 [1]. The large size of the required computer make it necessary to use a cart to transport the system. Convenient fetal monitoring at home would not be feasible with such a system. In addition, the signal processing microprocessor utilized in this monitor was a fixed-point digital signal processing (DSP) microprocessor (TMS320C25). Due to the operating speed



Figure 1.1 Previous Hardware System on A Cart

limitation of the TMS320C25, the DSP algorithm for the previous monitor was restricted in performance.

As technologies grow rapidly, a new much more powerful floating point DSP microprocessor (TMS320C31) is presently available commercially at low cost. In addition, compact portable PCs are readily available. The objective of the present research is to develop and implement a new real time signal processing algorithm, and to implement this algorithm with a floating point DSP processor housed in a portable PC. Thus the goal is to improve the accuracy of the algorithm for fetal heart rate detection and also to take advantage of advances to technology to make the monitor truly portable.

The system hardware and software would first be tested at a NASA laboratory. Eventually, the system would be clinically tested with a sensor belt at Norfolk General Hospital. The output of the acoustic fetal heart rate monitor would be compared with that of a Corometrics FIECG monitor on a number of patients.

1.3 Thesis Organization

A new FHR detection algorithm is presented in Chapter Two. Chapter Two also includes a discussion of several existing algorithms. Chapter Three introduces the system hardware and the signal processing procedure in detail. Chapter Four presents the real time software implementation of the new algorithm. The technical tests conducted at the NASA laboratory and the

clinical tests held at EVMS are presented in Chapter five. A discussion of these test results and the problems encountered are also presented in this chapter. Chapter five also contains experimental conclusions and points out potential for improvement for the current version of the acoustic monitor.

CHAPTER TWO

BACKGROUND AND THEORY

2.1 Introduction

The background and theory for the real-time implementation algorithm for the detection of fetal heart rate is presented in this chapter. Section 2.2 reviews the previous research work by Pretlow and Stoughton for the fetal heart rate monitor. First the important characteristics of the acoustic fetal signal are reviewed and secondly the algorithm and the performance of the heart rate monitor are reviewed. In section 2.3, a discussion of several existing (i.e., from the literature) heart beat event detection methodologies for the monitor is presented. It is important to note, however, that, except for the work by Pretlow and Stoughton, none of these methodologies were tested with acoustic fetal signals. Section 2.4 describes two new fetal heart tone detection algorithms which were developed in the present study.

2.2 Review of Previous Research Work

During the period from 1989 to 1991, Pretlow and Stoughton studied the FPCG signal, developed a fetal heart rate detection algorithm and implemented a working real-time system for fetal heart rate detection.

However, due to time, money, and technology limitations, the required system portability was never achieved. In this section, the FPCG signal characteristics are first described, and the previously developed system is then discussed in terms of its advantages and disadvantages.

2.2.1 FPCG Signal Characteristics

The FPCG signal is a relatively low energy signal. Generally, the fetal heart beat can be heard in only a small area of the mother's abdomen. The area is usually no more than three centimeters (cm) diameter, although the range of this local area can encompass up to a 12 cm diameter as illustrated in figure 2.1 [1]. The heart beat signal is thought to result from vibrations produced by the opening and closing of the four valves controlling blood flow through the fetal heart and from vibrations of the heart muscle. The heart beats are classified as the "first heart sound" and the "second heart sound" components. The "first heart sound" is thought to be due to closure of the mitral and tricuspid valves (MT), whereas the "second heart sound" is thought to be due to closure of the aortic and pulmonary valves (AT). Figure 2.2 [5], reproduced from the literature, shows a typical fetal acoustic recording, depicting the MT and AP components.

The frequency spectrum of the fetal heart beat signals is not definitive from the literature. Table 2,1 summarizes several literature findings concerning the frequency spectrum of the fetal heart beat signals and shows

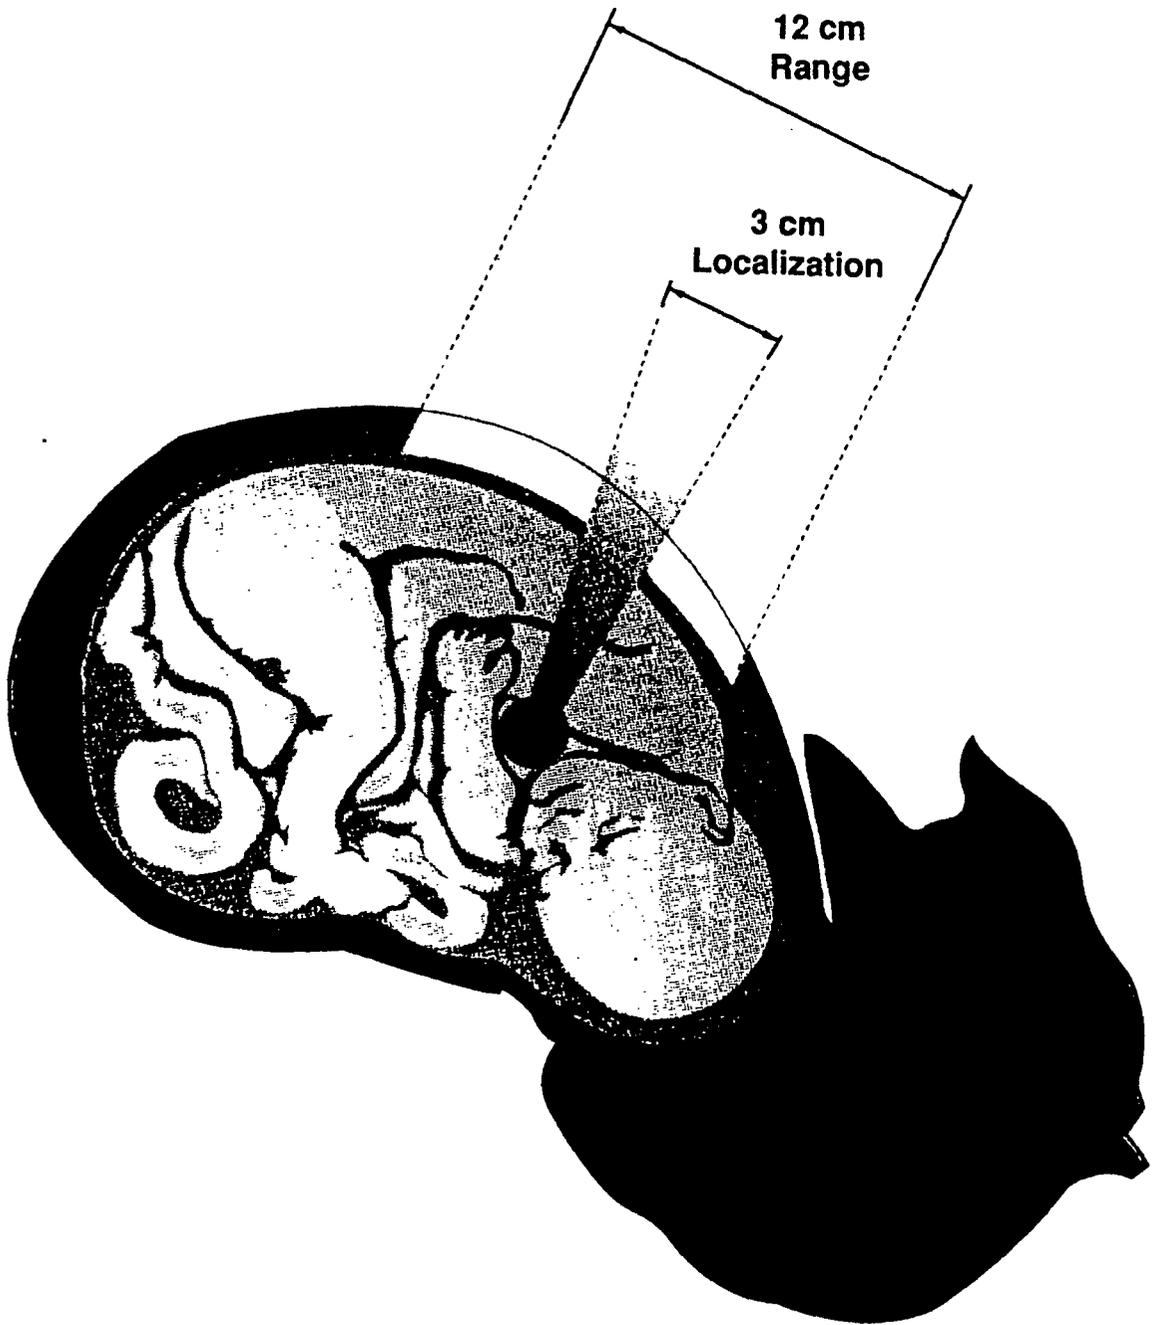


Figure 2.1 Fetal Heart Tone Localization

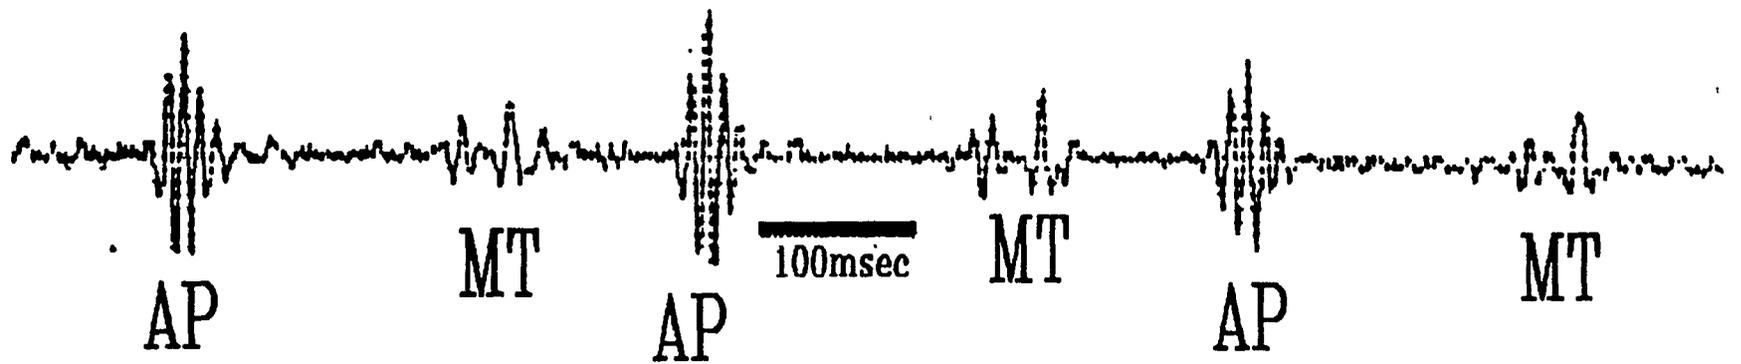


Figure 2.2 FPCG Signal

Table 2.1 Fetal Heart Frequencies per Available Literature

Source	First Sound	Second Sound
Hewlett-Packard[6]	80-110 Hz	——
Jenssen [7]	40 Hz	40-50 Hz
Talbert et al.[5]	60-80 Hz(40 Hz)	60 Hz
Kobayshi et al.[8]	50 Hz	100 Hz
Nagel [9]	20 Hz(45 Hz)	40-50 Hz
Talbert et al.[10]	30 Hz	75-100 Hz

that most of the energy of the fetal heart beat signals is concentrated in the frequency range below 110 Hz. However, Pretlow's data collection pilot study shows that the spectrum of the fetal heart beat signals is mainly within a range from 20 Hz to 50 Hz, and the maternal heart beats have a spectrum from 8 Hz to 15 Hz. In addition, the maternal heart beat signals can be 10 to 100 times stronger than the fetal heart beat signals. Thus, without analog filtering to at least partially separate the two signals, the maternal heart beat signals could saturate the front end amplifier at the gain levels required to detect the fetal heart beat. The fetal heart beat signals would then be dissolved in the quantization noise. Figure 2.3 [1] shows the spectrum of the fetal heart beat signals, as determined by Pretlow.

2.2.2 Previous System Performance

The previous fetal heart rate monitor demonstrated the feasibility of the real-time detection of fetal heart rate from a noise contaminated acoustic signal by comparing favorably to a commercial ultrasonic monitor for heart rate generation for a fetal nonstress test. However, the previous monitor did have unsolved problems, as discussed below. The detection algorithm for the fetal heart rate was based on a Least-Mean-Square (LMS) linear predictor algorithm. Digital signal processing was accomplished with a fixed-point DSP processor (TMS320C25) platform.

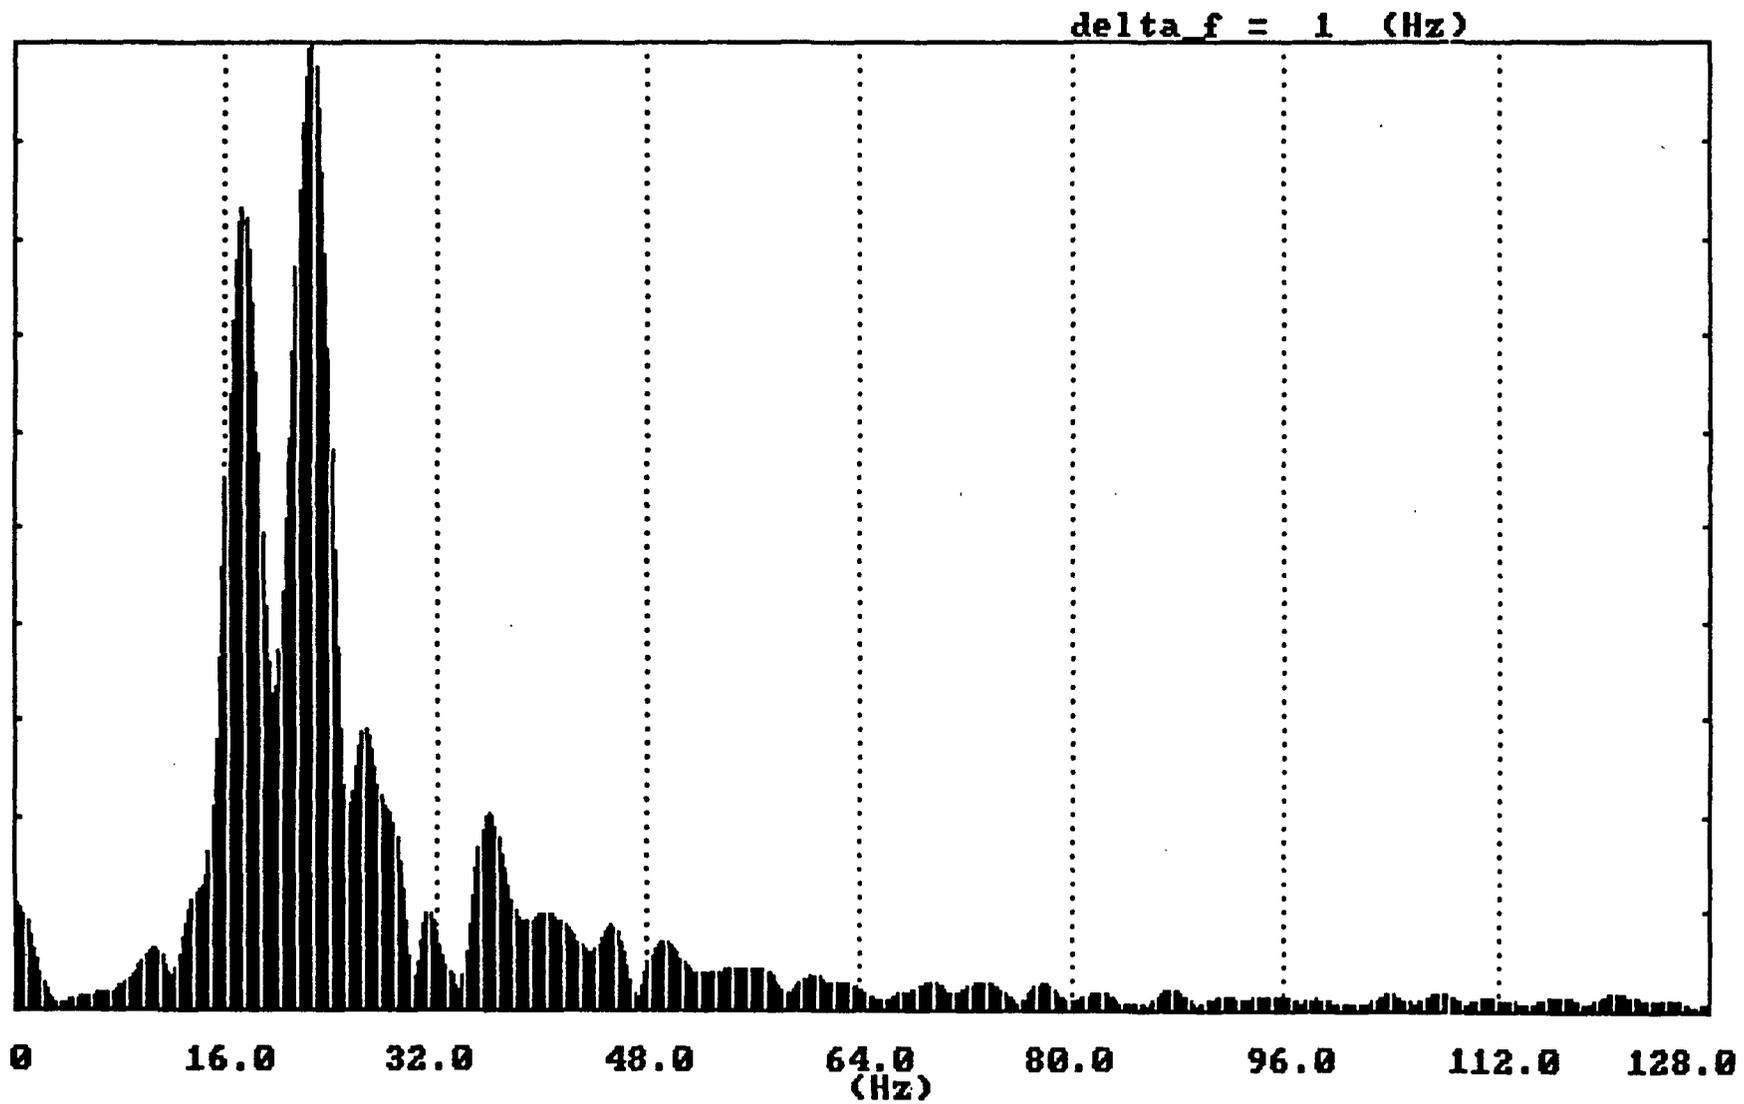


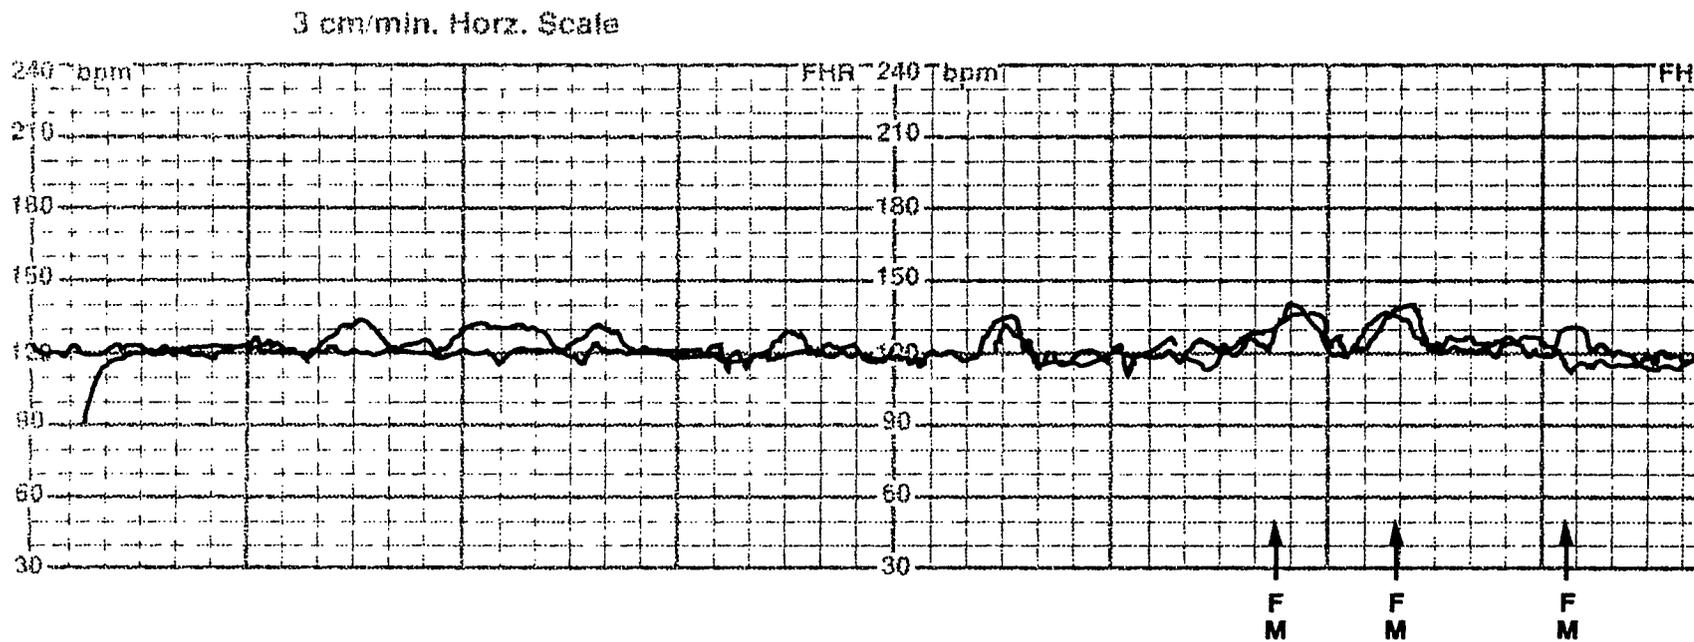
Figure 2.3 Fetal Heart Tone Spectrum From Pretlow[1]

The LMS linear predictor algorithm shows a high degree of efficacy for detection of the fetal heart rate in terms of a good anecdotal correlation in Pretlow's comparative studies. Figure 2.4 [1] depicts this correlation between the LMS linear predictor monitor and the commercial ultrasonic monitor for patient 23.

However, the linear predictor algorithm had a problem of sensitivity to noise, particularly maternal movement artifact. It is also likely that the fetal heart tone sound was erroneously identified in the presence of conditions such as bodily "thumps" like the mother moving her hand against her leg. Other problems encountered include sensitivity of the predictor to maternal heart tones. This problem was significant in the mothers with large overlap in the fetal and maternal spectrum.

2.3 Previous Heart Beat Event Detection Methodologies

Heart beat event (wavelet) detection is necessary in order to determine heart rate. Heart rate is the reciprocal of the time interval between two consecutive heart beat wavelets of the same type. Heart beat wavelet detection methodologies reported in the available literature consist mainly of four approaches: LMS linear prediction signature matching, average magnitude cross-difference, digital matching filtering and autocorrelation techniques. The following four sections describe these four methodologies in detail respectively.



PATIENT 23

ULTRASONIC ———

ACOUSTICAL ———

Figure 2.4 Acoustic vs. Ultrasonic NST: Patient 23 of Pretlow[1]

2.3.1 LMS Linear Prediction Signature Matching [1]

LMS linear prediction theory states that if a signal signature represents the response of a linear system, each point in the signal signature vector is a linear combination of all the preceding points, or:

$$S_{LP}(n) = \sum_{k=1}^N S(n-k)P(k) ,$$

where $S_{LP}(n)$ are the signal signature vector samples, $S(n)$ are the signal vector samples, N is the length of the predictor and $P(k)$ are the predictor coefficients. In the implementation of this method by Pretlow and Stoughton, these coefficients were trained once offline before real time operation. They computed the residual prediction error as:

$$E(i) = S_{LP}(i) - S(i) .$$

The error was used in a feedback loop to adjust the values of the coefficients to more accurately predict the desired signal vector values. Until the mean square error was minimized to a local minimum, the coefficients were iteratively adjusted by a gradient method,

$$P(k+1) = P(k) + 2\delta E(k)S(k) ,$$

where δ is a constant regulating the stability and speed of convergence of the adaption.

Once the coefficients were well trained, the linear predictor was used to signature the acoustic fetal heart signal point by point in real time. The signatured signal is compared to the original acoustic signal to generate a normalized mean square error (MSE). A nominal error threshold was set to determine the occurrence of the fetal heart tone in the acoustic signal. An incoming fetal heart tone in the acoustic signal caused a leading edge in the MSE, while a leaving fetal heart tone led to a trailing edge. The average of the time intervals of the leading edges and the trailing edges was considered to be a fetal heart tone period.

The LMS linear prediction algorithm has advantages of high speed and reasonably accurate detection of signal features if the required training space for the coefficients is large enough. However, the algorithm is intended to predict signals representing the response of a linear system. In the FHR application, the signal is the measured acoustic response to the activity of the fetal heart. This response includes several nonlinear effects. Therefore, the algorithm is limited in its ability to detect the fetal heart tone effectively.

2.3.2 Average Magnitude Cross Difference [2]

Average magnitude cross-difference involves comparing an input signal vector to a known desired template vector point by point. The difference

between corresponding samples in the input signal vector and the template vector are computed using waveform subtraction. The sum, $y(i)$, of the absolute values of the differences is then computed:

$$y(i) = \sum_{k=0}^{N-1} |S_t(k) - S_t - [S(k,i) - S_t]| ,$$

where $S_t(k)$ are samples of the template vector (that is, samples of a template for a fetal heart tone, for this application), $S(k + i)$ are samples of the input signal vector, N is the length of the template vector, and i is the time shift parameter. S_t is the mean value of the template vector given by

$$S_t = \frac{1}{N} \sum_{k=0}^{N-1} S_t(k) ,$$

and S_i the mean value of the i th input signal vector given by

$$S_i = \frac{1}{N} \sum_{k=0}^{N-1} S(k,i) .$$

When the input signal vector and the fetal heart tone template vector are very similar in shape, independently of overall amplitude, the AMCD value, $y(i)$, becomes a minimum (theoretically zero).

However, this method is based on two main assumptions. One assumption is that the shape of the fetal template vector is known *a priori*, and the other one is that fetal heart tones are strongly present in the input signal vectors (i.e., a high signal to noise ratio). In practice, these two assumptions may not be valid. Another important point is that the template should allow, within limits, time expansions and contractions..

2.3.3 Digital Matched Filtering [2]

Matched filtering is commonly used to detect time recurring signals buried in noise. The main underlying assumptions in this method are that the input signal is time limited and has a known wave shape. The problem then is to determine its time of occurrence. The impulse response of a digital matched filter, $h(k)$, is the time-reversed replica of the signal to be detected. Thus in our case, if $S_i(k)$ is the selected template vector then the coefficients of the matched filter are given by

$$h(k) = S_i(N-k-1), \quad k=0,1,\dots,N-1 .$$

The digital matched filter can be represented as an FIR filter with the usual transversal structure, with the output and the input of the filter related as

$$y(i) = \sum_{k=0}^{N-1} h(k)S(i-k) ,$$

where $S(i)$ are the samples of the input signal vector, $S_t(k)$ are the samples of the template vector, N is the filter length, $h(k)$ are matched filter coefficients, and i is the time shift index. It is evident that when the template vector and the fetal heart tone coincide, the output of the matched filter will be maximum. Thus by searching the output of the matched filter for a value above a threshold the occurrence of the fetal heart tone can be tested.

This method is similar in concept to the AMCD described above, with a slightly different implementation. It is essentially a template comparison, as is the AMCD method, except it will be sensitive to variations in the signal amplitude. It has some of the same limitations as the AMCD method.

2.3.4 Autocorrelation

Autocorrelation detection is generally used to determine fetal heart rate from Doppler-shift ultrasonic signals.[11] The fetal heart beat is manifested as reflected tone burst wavelets analogous to those present in the fetal phonocardiogram. The tone bursts correspond to Doppler frequency shifts of the reflected continuous outgoing ultrasound signal produced by moving heart valves. Autocorrelation processing detects the tone bursts among the noise, using

$$R_a(k) = \frac{1}{N-k} \sum_{n=0}^{N-k-1} S(n)S(n+k) ,$$

where $0 \leq k \leq N - 1$ and N is frame length of an autocorrelation buffer. The frame length of the autocorrelation buffer determines whether the detecting method emphasizes sensitivity to a long term average of the tone burst or a short term sudden change of the tone burst. Generally, this frame length is from one to three seconds long for Doppler-shift ultrasonic fetal signals [11].

Autocorrelation emphasizes periodically occurring correlated wavelets by eliminating non-periodic uncorrelated noise such as movement artifacts, baseline drift, and random environmental noise. However, autocorrelation techniques also require fairly high signal noise ratio (SNR) of detected signals, and will also be sensitive to any periodic noise sources.

2.4 Heart Beat Event Detection Methodologies For This Research

The detection of the fetal heart tones is a complex process. The ambient background noise and the large-amplitude maternal heart tones prevent the desired high SNR for the fetal heart tones. Other noise sources, including aortic pulse waves, placental flow sounds, and bodily movement artifacts, further degrade the SNR. As yet another problem, acoustic fetal heart tones are much weaker than the ultrasonically-derived fetal heart signals. All these

conditions add to the difficulty of the detection of the fetal heart tones using passive acoustic methods.

2.4.1 Summary of Detection Methodologies

In the previous section, four detection methodologies have been introduced for detecting fetal heart rate. They all have their advantages and disadvantages for acoustic fetal monitoring, due to underlying assumptions. As mentioned previously, except for the LMS linear predictor method, these techniques have not been tested with acoustic signals. These advantages and disadvantages of these methods for acoustic fetal monitoring can be summarized as follows:

LMS linear prediction has the advantage of working in a low SNR situation as this work requires. The previously developed monitor used the linear prediction as the algorithm for the detection of the fetal heart tone, and was found to perform adequately under some conditions. However, under other conditions, performance was degraded. As noted before, due to fundamental invalid assumptions regarding linearity, this approach was not further developed in this research.

Both average magnitude cross difference and digital matched filtering algorithms would theoretically detect the fetal heart tones. However, both algorithms require high-SNR signals. An accurate template is also required for both detection methodologies. However, the shape of the fetal heart tones

change from patient to patient and even over time within the same patient. The lack of generality is one of disadvantages for these algorithms. Major modification would be necessary in order for these methods to work well for the present application.

Autocorrelation, by itself, would not be an optimum technique for detection because of its tendency to lock on other periodic noise such as 60 Hz EMI or maternal heart signals. However, autocorrelation does have the advantage of eliminating non-periodic noises even though the SNR is low.

As mentioned above, the previous detection methodologies are not directly suitable for the present application due to reasons mentioned above. However, certain aspects of these methodologies can be combined with new methodologies to improve overall performance. Two new methodologies, the Teager energy operator and a neural network, are introduced in the following two sections. Both methodologies use non-linear processing and can work with a low SNR signal.

Another very important consideration is the speed of the TMS320C31 is more than double that of the older TMS320C25 processor used in the previous implementation. The newer processor also is floating point, has more memory, and much better software support tools. Therefore, much more complex algorithms are possible, including hybrid methods which include parallel redundant processing of signals from multiple sensors. For this research, multiple detection methodologies are simultaneously executed in real time on

signals from multiple sensors, thus computing multiple estimates of the heart rate at each time instant. An evaluation of the performance of these estimates is also made in real time, and the final instantaneous fetal heart rate is then derived from the comparison.

2.4.2 Teager Energy Operator

The fundamental goal of the signal processing is simply to determine the temporal locations of the fetal heart tones. These should be regions of local high energy. One operator which can be used to trace the energy of the acoustic fetal heart signal, is a non-linear operator, called the "Teager energy operator." The operator was first proposed by James F. Kaiser [4]. In this section we summarize the basic theory of this operator.

Applying Newton's law of motion to a mass m suspended by a spring of force constant k , yields a well-known energy expression of the system:

$$E = \frac{1}{2} m \omega^2 A^2$$

or

$$E \propto A^2 \omega^2 ,$$

where A is the amplitude of the system, and ω is the frequency of the system. Thus the energy of the system is proportional not only to the square of

amplitude but also to the square of the frequency of the system. Now, let us look for simple means for calculating a running estimate of the energy E produced by the vibration of the opening and closing of the valves of a fetal heart tone.

Let S_n be the samples of the signal representing one of the frequency components of the vibration of the opening and closing valves. Therefore

$$S_n = A \cos(\Omega n + \phi) ,$$

where Ω is the digital frequency in radians/sample and is given by $\Omega = 2\pi f / f_s$ where f is the analog frequency and f_s is the sampling frequency. ϕ is the arbitrary initial phase in radians.

Using the same notation, the product of S_{n-1} and S_{n+1} is computed and simplified by a trigonometric identity to form:

$$S_{n-1}S_{n+1} = A^2 \cos^2(\Omega n + \phi) - A^2 \sin^2(\Omega) .$$

Notice that the first term on the right-hand side of the equation is simply the square of S_n . Thus, substituting for S_n , the following equation results:

$$A^2 \sin^2(\Omega) = S_n^2 - S_{n-1} S_{n+1} .$$

It is important to note that the equation is exact and unique provided that the value Ω is restricted to value less than $\pi/2$, the equivalent of one-fourth of the sampling frequency. In addition, for small values of Ω , $\sin(\Omega) \approx \Omega$. Now if the value of Ω is limited to $\Omega < \pi/4 = 0.7854$, i.e. $f / f_s, 1/8$, then the relative error is always below 11%. There results

$$E_n = S_n^2 - S_{n-1} S_{n+1} = A^2 \sin^2(\Omega) \approx A^2 \Omega^2 ,$$

where E_n is the energy output of the algorithm.

The Teager energy operator as summarized above is nonlinear, independent of the initial phase ϕ , robust even if the signal passes through zero, and also responds rapidly to changes in both A and Ω (only two sampling instants). In addition, the Teager energy operator is very general in the sense that a fetal heart tone template is not required for the algorithm. However, the Teager energy operator traces energy not only of the fetal heart tones but also of other noise sources. In our implementation, additional signal processing procedures (described in Chapter Three), such as autocorrelation and intelligent merit computations, were combined with the Teager energy operator to reliably detect the fetal heart tones.

2.4.3 Neural Network

Neural network models have been studied for many years in the hope of achieving human-like performance in the field of pattern recognition. These models are composed of many nonlinear computational elements operating in parallel and arranged in patterns reminiscent of biological neural nets. Computational elements or nodes are connected via weights that are typically adapted during use to improve performance. This section introduces one of the static neural nets -- Multilayer Perceptron (MLP) to identify fetal heart tone patterns from a noisy acoustic signal. It also summarizes the MLP learning algorithm -- Backpropagation.

The Multilayer Perceptron (MLP)

The multilayer perceptron can be considered as a series of products between input vectors and weight matrices, followed by a non-linear function. In this research, a sigmoid function is chosen as the non-linear function for the non-linearity of the multilayer perceptron. A sigmoid function of a variable X is computed as:

$$F_{sig}(X) = \frac{1}{1 + e^{-X}}$$

The dimension of the input vectors and the weight matrixes can be determined based on the number of nodes in each layer. An output vector of each layer

can be considered as an input vector of the next layer. Except for the input layer, an input vector of each layer can be added with a bias vector and then passed through a non-linear function (such as a sigmoid function) to become a output vector of that layer. A output vector of each layer times a weight matrix produces an input vector for the next layer. In this way, the weight matrix connects each layer to the next layer. The number of the products needed for a whole procedure of a feed-forward neural network depends on the number of layers in the neural net. In order to illustrate the procedure more clearly, an example with a one-hidden-layer neural net (including one input layer and one output layer) is given in the next paragraph.

The neural net has N nodes in the input layer, M nodes in the hidden layer and P nodes in the output layer. Therefore, three vectors associated with each layer, two bias vectors associated with the hidden layer and the output layer respectively, and two weight matrixes are needed for the neural network. The input vector V_i has N dimensions. The hidden vector V_h has M dimensions, and the output vector V_o has P dimensions. The weight matrix M_{ih} connecting the input layer to the hidden layer is an N by M matrix, and the other matrix M_{ho} connecting the hidden layer to the output layer is an M by P matrix. The two bias vectors have the same dimensions as their corresponding I/O vectors. The input of the hidden layer can be computed as

$$V_h = V_i \cdot M_{ih} ,$$

and the output of the hidden layer is

$$V_h = \frac{1}{1 + e^{-(V_h \cdot B_h)}} ,$$

where B_h is the bias vector associated with the hidden layer.

In the same way, the input of the output layer can be obtained as

$$V_o = V_h \cdot M_{ho} ,$$

and the output of the output layer is

$$V_o = \frac{1}{1 + e^{-(V_o \cdot B_o)}} .$$

In this research, the neural net used had a 20-30-1 structure. The neural net was trained to identify the time-domain signature of the fetal heart tones. At each sample time the neural network is presented the current input plus the previous 19 acoustic samples as its input vector. For each such sample time, the neural network computes one output value. The output indicates whether the corresponding acoustic signal sample is part of a fetal heart tone.

The higher value the output is, the more probable the sample is part of a fetal tone. The weight matrices were trained, offline, using labeled training data, using the backpropagation algorithm, as described in the next section.

Backpropagation Algorithm

The backpropagation algorithm is a generalization of the LMS algorithm. It uses a gradient search technique to minimize a cost function equal to the mean square difference between the desired and the actual net outputs. The desired output of all nodes is typically "low"(0 or < 0.1) unless that node corresponds to the pattern which is the fetal heart tone is "high"(1.0 or > 0.9). The algorithm requires a continuous differentiable non-linearity. This is the reason why the sigmoid function is chosen for the multilayer perceptron. The backpropagation training algorithm consists of four main steps, including initializing weight matrices and bias vectors, presenting input vectors and desired outputs, calculating actual output vectors and adapting weight matrices.

The first initialization step is to set all weights and node biases to small random values. The absolute values of these weights and biases are usually between zero and one at the beginning. In this way, the neural network training time will not be impeded because of beginning at a place far away from the minimum in the error surface.

The second step is to properly present input vectors and desired outputs. The desired outputs are selected as "high"(1.0 or > 0.9) when the input vectors are in the middle of fetal heart tones, and as "low"(0 or < 0.1) when apparently not fetal heart tones. When the input vectors are in the transition period from non-fetal-heart-tone period to fetal-heart-tone period, the corresponding desired outputs are set "linearly " from 0.1 to 0.9, and vice versa.

The third step is to calculate actual outputs. The actual outputs are calculated as a full process of a feed-forward perceptron as mentioned above.

The final step is to adapt weight matrices. The backpropagation algorithm first adapts weights in the output layer, and adjusts weights and biases in the hidden layer.

The weight matrix M_{ho} is adjusted using

$$M_{hoi}(k+1) = M_{hoi}(k) + \eta \delta_{ho} \odot V_h .$$

In this equation, $M_{hoi}(k)$ is i th row of $M_{hoi}(k)$, \odot is a dot product sign for vectors and η is a gain term which controls the training speed of the neural net. The larger value η is, the larger step the neural net jumps from the place before adapting weight matrices to the place after adapting weight matrices in the error surface. In addition, δ_{ho} is an error term for the output nodes and computed as

$$\delta_{ho} = V_o \odot (1 - V_o) \odot (V_d - V_o) .$$

where V_d is a desired output vector.

The weight matrix M_{ih} is adjusted a little differently than M_{ho} :

$$M_{ihj}(k+1) = M_{ihj}(k) + \eta \delta_{ih} \odot V_i \odot \alpha (M_{ihj}(k) - M_{ihj}(k-1)) ,$$

where j is row index of matrix M_{ih} and α is usually a value between zero and one. The δ_{ih} error term for weight matrix M_{ih} is also computed differently than that for M_{ho} :

$$\delta_{ih} = V_i \odot (1 - V_i) \odot (\delta_{ho} * M'_{ho}) ,$$

where " ' " is a transpose sign for a matrix.

The algorithm repeats steps 2-4 until the mean square error reaches an acceptable minimum. The MLP neural net is then well trained.

The MLP neural network has the advantage of more accurate detection of fetal heart tones than other methodologies. However, the MLP neural network is difficult to train well due to limited data. In addition, one full feed-forward pass of the MLP neural network requires a very large number of floating-point calculations. This is one of the main reasons why Pretlow didn't adopt this methodology for his version of the fetal heart rate monitor.

CHAPTER THREE

FETAL HEART RATE MONITORING SYSTEM

3.1 Introduction

The methodologies are presented in this chapter for the real time signal processing procedure for the detection of fetal heart tones. The technique used for heart rate derivation is explained, as well as a parallel redundancy technique which was used for rate correction in case spurious signals were initially identified as heart tones. In section 3.2 the hardware system is described. A detailed description of the digital signal processing (DSP) procedure is given in section 3.3. Note that a description of the software used to implement the DSP algorithms is given in Chapter four.

3.2 Hardware

Hardware for a portable acoustic fetal heart monitor has been developed and tested. A detailed description of the hardware system is given in this section, including an overview of the hardware system, a introduction of an acoustic pressure sensor belt developed by NASA, and a description of the electronics support system for the fetal heart monitor.

3.2.1 Overview

The portable acoustic fetal heart rate monitor consists of four main components: a pressure sensor array mounted on a belt worn by the mother, an electronics support system, a 33.3 MFLOPS DSP platform (TMS320C31 based) and a portable IBM PC. Figure 3.1 is a block diagram of the monitor system hardware. Both the DSP platform and the portable PC are commercial products, from Atlanta Signal Processing Inc. and Modgraph Inc., respectively. The pressure sensor belt was fabricated at NASA Langley. The electronics support system was designed and tested in the Department of Electrical and Computer Engineering at Old Dominion University.

The acoustic sensor belt contains a seven-element sensor array. Each element of the sensor array utilizes piezoelectric polymer film polyvinylidene fluoride (PVF2 or PVDF) to detect the fetal heart tone and thereby permit the derivation of heart rate. The choice of this material is based on several particularly fitting properties: high sensitivity, low surface mass density, mechanical flexibility, ruggedness, chemical stability, and low cost.

The electronics support system contains five main parts, including instrumentation amplifiers, band pass filters, two time-division multiplex channels, two final-stage amplifiers, one push button detector and one audio amplifier. The instrumentation amplifiers suppress 60 Hz noise and amplify the fetal heart signals. The band pass filters not only further suppress 60 Hz noise but also suppress maternal heart signals. Two time-division multiplex

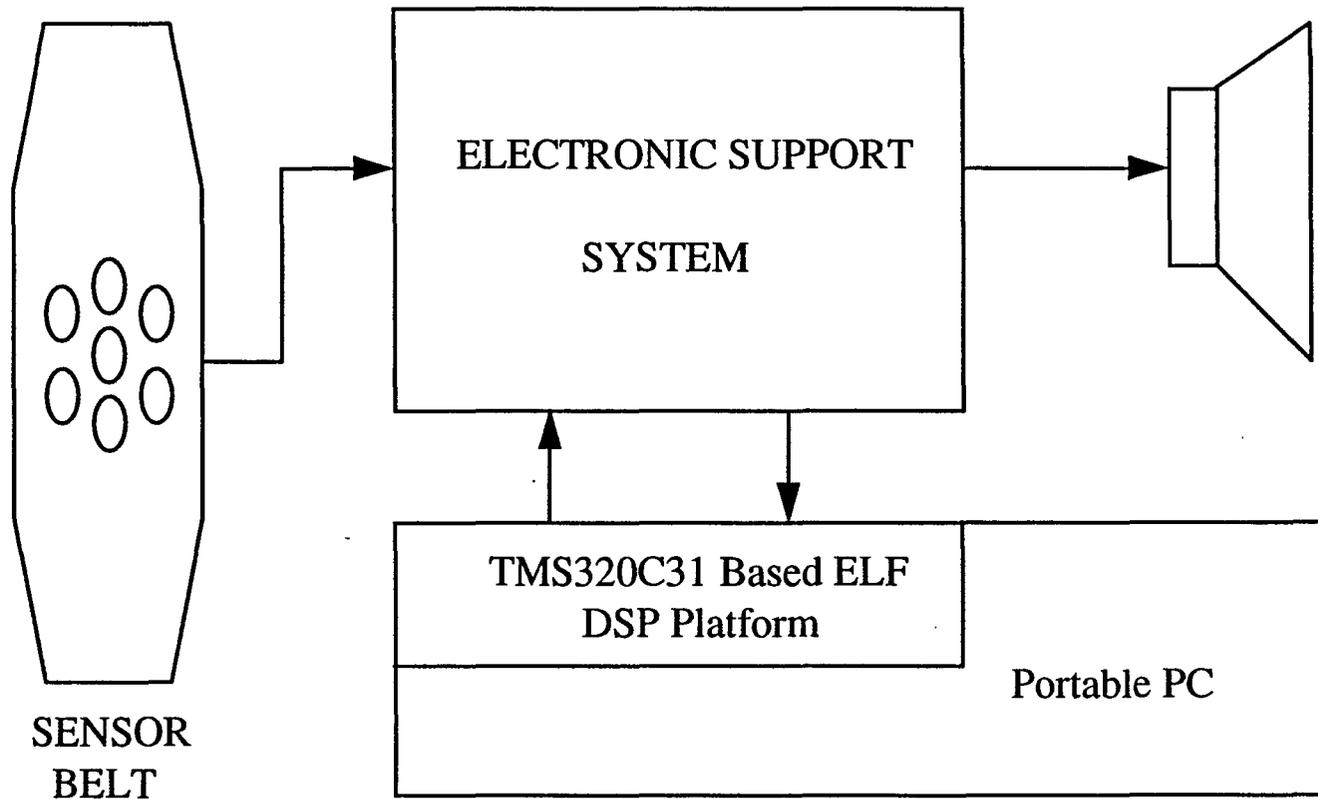


Figure 3.1 System Hardware Block Diagram

(TDM) channels combine seven acoustic fetal heart signals to two time-division multiplexed signals (since the DSP board has only two AD channels). Two final-stage amplifiers have variable gain from 1 to 10, allowing an overall signal gain of 400 to 4000. The push button detector is designed to allow mothers to indicate periods of fetal movement and contractions, both of which would be expected to be correlated to periods of elevated fetal heart rate. The audio amplifier gives an option to listen to the fetal heart tones in real time.

3.2.1 Acoustic Sensor Belt

The fetal heart tone is highly localized on the maternal abdominal surface. The localization is attributed to preferred acoustic propagation to sites where the fetal back is in contact with the maternal abdomen. This contact site favors an ideal acoustic impedance match, and propagation to other sites is greatly attenuated. While the fetus is curled up in the classic itol position, acoustic transmission from the fetal chest is dampened by amniotic fluid and a longer path to the maternal surface, while lateral transmission is suppressed by the fetal arms. As mentioned previously in chapter two, the contact site has a typical diameter of three cm. As the fetus moves about, this site also moves, but generally remains within a range of approximately twelve cm in diameter. The present version of the acoustic sensor belt, representing the third generation of development, contains a seven-element sensor array which covers the twelve cm diameter site expected for

possible pickup of the fetal heart tone. Consequently, the sensor belt should always have at least one sensor which is in good contact with the acoustic fetal heart tone on the maternal abdominal surface, even if the fetus moves during a fetal heart monitoring test. However, the best sensor element may change at different time instants during monitoring, since individual sensors only cover approximately a 4-cm-diameter area. Figure 3.2 [3] shows cross section and cutaway views of the sensors installed on the acoustic sensor belt.

The acoustic sensor belt is designed to fulfill five functions: signal detection, acceleration cancellation, acoustical isolation, electrical shielding, and electrical isolation of the mother. Each element sensor contains two PVF2 elements arranged in a bimorph structure -- mechanically in series and electrically in parallel. The internal sensor detects pressure pulses on the maternal abdomen situated directly opposite. Electrical contacts to the Ni electrodes are made with conducting epoxy. Although the bimorph structure is usually associated with flexure, the sensor shown here operates in the compressional mode -- a normal stress producing a normal electric field with e_{33} the active piezoelectric modules. The external sensor, identical in construction, is intended to cancel accelerations due to rigid body motion of the mother. It is acoustically isolated with kevlar wool to attenuate signals due to ambient noise. The Ni plate attached to the belt is bonded to both bimorphs and assures that both are subjected to the same accelerations. A foil of Cu-coated kapton completely surrounds the sensor assembly. When the ground

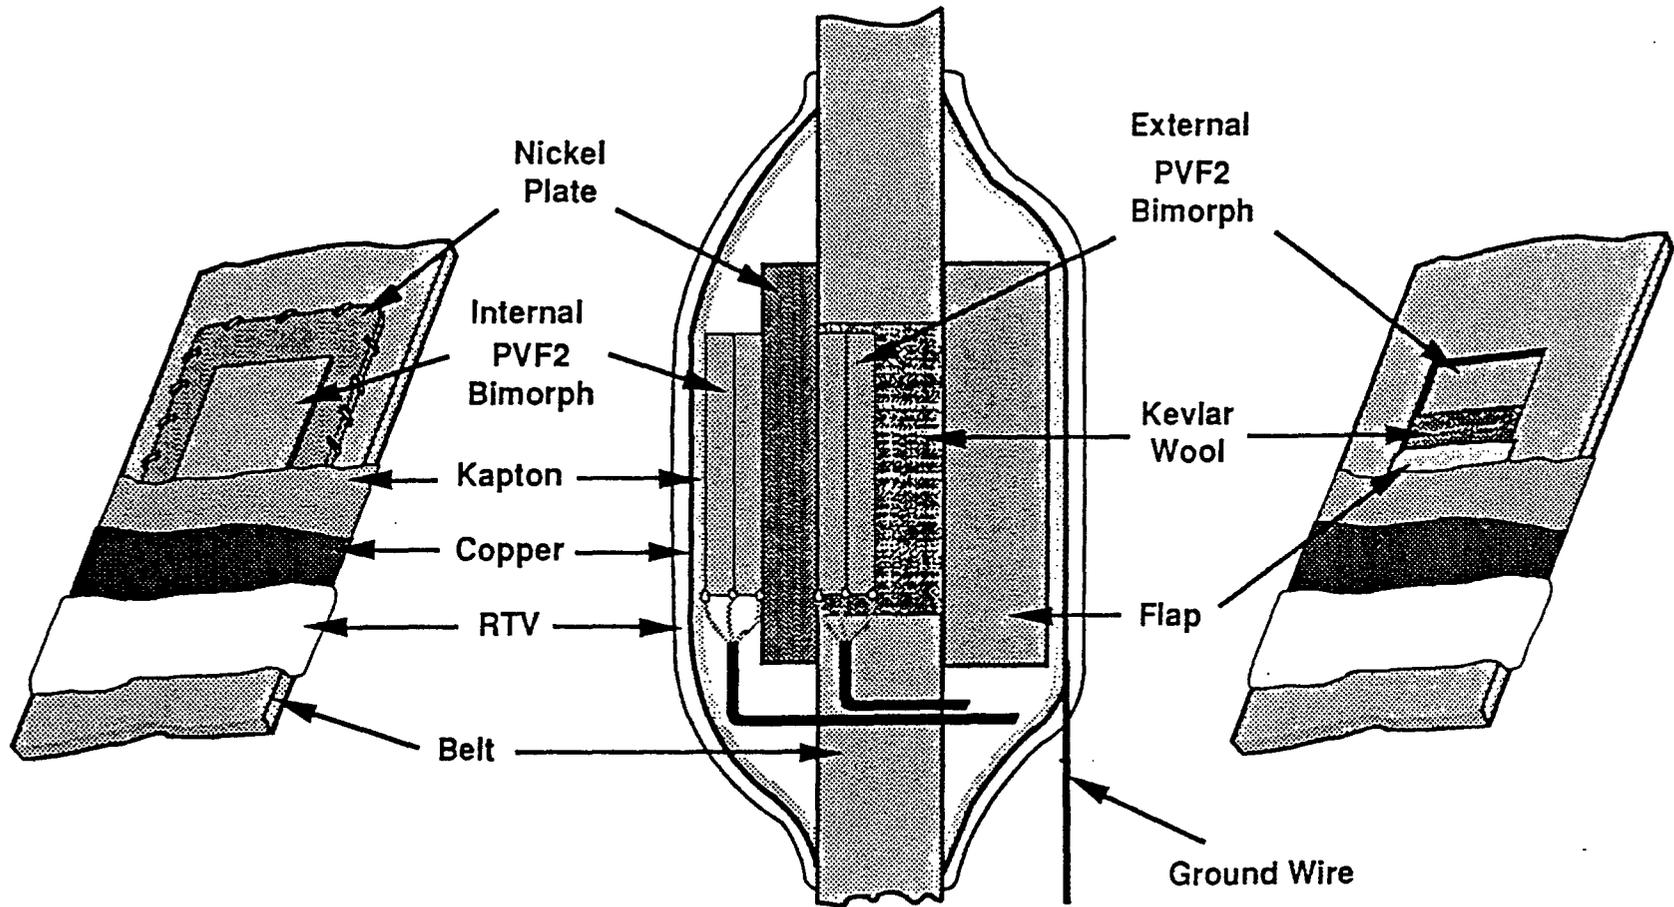


Figure 3.2 Cross-Sectional and Cutaway Views of The Sensor Belt

wire connected to the Cu coating contacts earth ground, the ubiquitous 60 Hz EMI merges into the background. A final layer of RTV silicone rubber, completely covering the Cu shield, isolates the mother from earth ground, but is acoustically transparent. The belt itself is made of nylon parachute webbing. It does not have to be drawn tightly around the mother, rather requiring only minimal acoustic contact, but must have a sufficiently high stiffness to resist displacement by the incident pressure pulses and thus assure adequate compression of the PVF2 foil.

3.2.3 Electronic Support System

The main purpose of the electronic support system for the fetal heart monitor is to amplify the sensor signals to match the dynamic range of the AD convertors, while keeping the fetal heart signal to noise ratio (SNR) as high as possible. To achieve this purpose requires that the electronic support system introduce as little electronic noise as possible, suppress unwanted signals (such as remaining 60 Hz noise, ambient noise and maternal heart tones) as much as possible, and amplify the fetal heart tones enough to avoid AD quantization noise problems.

The electronics support system contains five main parts, including instrumentation amplifiers, band pass filters, two time-division multiplex channels, two final-stage amplifiers, one push button detector and one audio amplifier. Figure 3.3 is a block diagram of the electronic support system. The

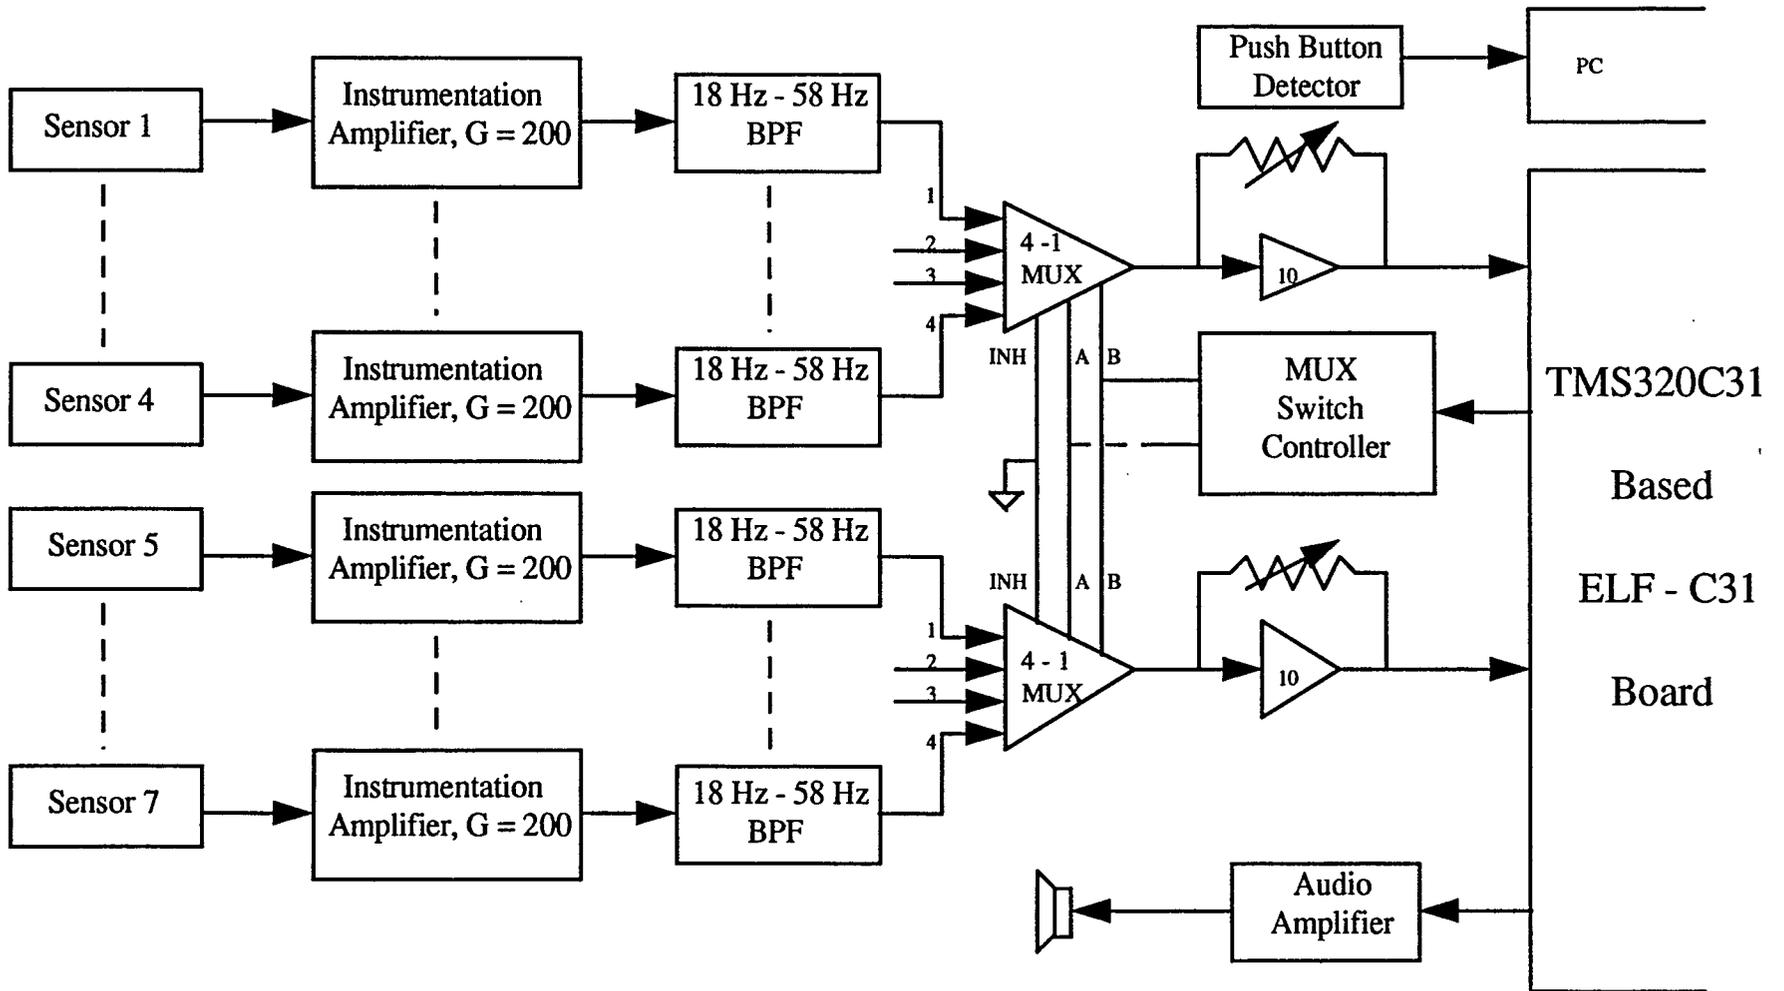


Figure 3.3 Electronics Support System Hardware Block Diagram

I/O configuration of the electronic support system consists of three input connectors and three output connectors. The three input connectors are: (1), one 25-pin male db connector from the acoustic sensor belt; (2), one audio stereo jack containing one TDM switch control signal and one fetal heart playback signal; and (3), one push button. The three output connectors are: (1), one 25-pin female db connector which is used to route the push button signal to the parallel printer port on the PC; (2), one stereo headphone jack for the audio amplifier output; and (3) one audio stereo jack containing the two TDM acoustic signals (i.e., the seven sensor signals).

The front-end of the electronics support system consists of fourteen low noise high gain instrumentation amplifiers (AD524). Although the sensor belt is well shielded, there is still some remaining 60 Hz EMI which is non negligible compared to the acoustic fetal heart tones. Each amplifier, connected to a single sensor element has a gain of 200, yet suppresses 60 Hz noise, described as follows. There are three leads coming out of each sensor element on the sensor belt. One is from the internal sensor, another is from the external sensor and the other is the common lead for the two sensors of a pair. Theoretically, the 60 Hz EMI is in phase for all three leads. Thus the common mode 60 Hz signal is suppressed by the differential inputs for the instrumentation amplifiers. The need for separate amplifiers for the two elements in a sensor pair arises from the configuration of the sensors, and the desire to suppress motion artifact signals. That is, the motion signals are out

of phase for the inner and outer sensors, whereas the fetal signals are in phase for the two sensors. Thus the sensor signals must be added to enlarge fetal signals and reduce motions signals. Experimentally the motion signals on the outer sensor was found to be about one fourth as large as the motion signal for the inner sensor. Therefore the signals are added, using a differential gain factor selected to reduce the effects of motion artifacts. A differential gain adjustment is available in the circuit, providing a gain range from one to four.

After the initial amplifier section, each of the seven resultant signals is band pass filtered. Note that a total of seven band pass filters were required, rather than only one as was used in a previous version of the monitor, since, as described later, our technique samples all seven sensors simultaneously to improve overall performance. The band pass filter consists of a four-pole highpass Butterworth filter with a cutoff frequency at 18 Hz, and a four-pole lowpass Butterworth with a cutoff frequency at 58 Hz. The cutoff frequencies for the band pass filter were chosen to suppress acoustic noise such as maternal heart tones and 60 Hz EMI, while passing as much fetal tone energy as possible. Although there is a large energy component in the fetal heart tone in the frequency range from 16 Hz to 20 Hz, as shown in Figure 2.3, there is also typically an even stronger energy component of the maternal heart tone in this frequency range. However, above 20 Hz, the energy of the fetal heart tone is expected to remain high whereas the energy of the maternal signal is

expected to be lower. The 60 Hz EMI could have been better suppressed if the high-end cutoff frequency of the band pass filter were lower than 58 Hz. However, a bandpass filter with a narrow bandwidth could result in unwanted ringing on the filtered signal. In addition, the 60 Hz EMI was not judged to be enough of a problem, because of the other measures already taken to suppress this, to warrant this additional filtering.

A time-division multiplexed (TDM) circuit was necessary since there are only two A/D channels available for the ELF DSP board while there are a total of seven acoustic fetal heart signals which must be sampled. The two multiplexers are controlled by two logic signals, derived from a single DA control signal generated by the ELF DSP board, in order to combine the seven acoustic fetal heart signals onto the two AD inputs for the DSP board. One TDM acoustic signals contains acoustic signals from the first four sensors, and the other one contains acoustic signals from the last three sensors, plus a grounding signal needed for demultiplexing, as described below. A fundamental point is that the multiplexing operation has to be synchronized with the AD sampling, in order that the signals could be properly demultiplexed. The only option was to use the DSP DA capability for synchronization purposes. The multiplexers have a truth table as shown in figure 3.4, where B and A are the logic controls for the device. Since only one DA DSP channel was available to generate the control signal, it was necessary to generate the second signal by a simple divide by two counter (i.e., a JK flip-flop). In the implementation, the

CD4052B Truth Table

INPUT STATES			"ON" CHAN
INHIBIT	B	A	
0	0	0	1
0	0	1	2
0	1	0	3
0	1	1	4

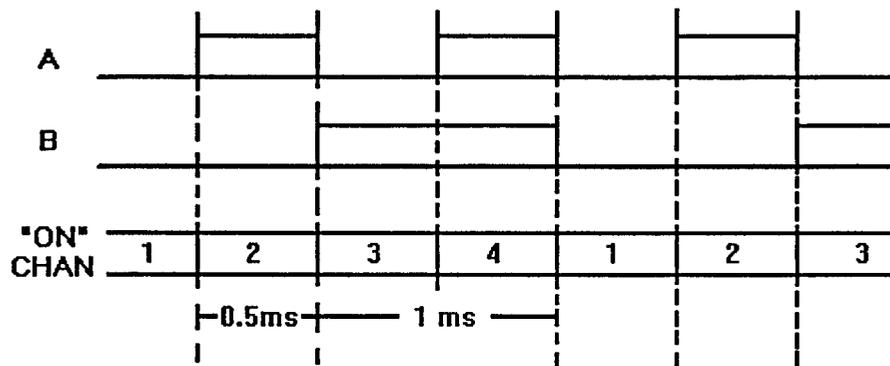


Figure 3.4 Multiplexer Switch Control Signals And Resultant TDM Signal

A logic signal is the DA output, and the B signal is the divide by two signal. The resulting TDM signal, in terms of channel number of the multiplexer, is also shown in figure 3.4 along with the control signals. Note that the multiplexer is thus synchronized to the AD converter, since the DA control signal is synchronized to the AD by the DSP card. The multiplexer was switched every .5 ms, or 2000 times per second. Using an initial sampling rate of 8 kHz, there are 4 consecutive samples of each channel in each multiplexer "time slice." The procedure for demultiplexing is explained later.

Finally, the two resulting TDM acoustic signals are amplified to match the dynamic range of the A/D interface of the ELF DSP board, and sent to the ELF DSP board through the audio-stereo-jack output connectors. The gains for both TDM acoustic signals can be manually controlled by tuning a dual-ganged 100 k Ω pot, configured to insure that the amplification for both channels is the same.

In addition, there are two extra capabilities of the electronic support system. One is an audio amplifier, used to provide sufficient power for a speaker so that one channel of the acoustic signal can be listened to while the monitor is in operation. This provides audio feedback so that nurses and patients can better position the belt over the fetus. The other is the push button detector. The push button detector is basically a simple RC circuit connected to set one of twenty five bits on the parallel printer port while the

button is pushed. This detector is designed to allow the mother to indicate periods of fetal movement and contractions.

3.3 Digital Signal Processing

The two TDM acoustic signals are sampled by the ELF DSP board. A series of digital signal processing procedures are applied to the signals to detect the fetal heart rates. An overview of the digital signal processing is first presented. The next four subsections describe the method for sampling data, fetal heart tone detection, computing the autocorrelation, and computing and updating a figure of merit for the detected rates.

3.3.1 Overview

The two analog TDM acoustic signals containing seven acoustic fetal heart tones and one ground signal are sampled by the ELF-C31 DSP board, and digitally processed on the TMS320C31 DSP microprocessor to detect the fetal heart rate. The whole digital signal processing procedure takes less than .5 seconds to execute once in real time, and contains the four steps mentioned above.

Sampling the data requires the generation of the multiplex switch control signal, sampling the analog TDM signals at 8 kHz, demultiplexing the two TDM signals to eight separate discrete signals, determining the correspondence between the multiplexer signals and the seven acoustic fetal

heart signals, digitally filtering the seven fetal heart signals and computing their energy in terms of root-mean-square (RMS) values.

Since it was anticipated that best sensor for detecting the fetal heart rate could change from time to time in an unpredictable fashion, all seven signals were processed in parallel and subsequently evaluated with a figure of merit to rank the sensor performance at each time slice (every .5 sec). It was expected that of the seven acoustic signals, at least one would contain the true fetal heart tones since the fetus should not move out of the 12 cm diameter range covered by the sensor array. Therefore, all seven acoustic signals were processed by the fetal tone detection procedure, the autocorrelation algorithm, and finally ranked by the intelligent merit procedure. Since, at the time most of these algorithms were implemented, there was not sufficient training data to adequately train a neural net for signature matching, the Teager energy approach was adopted as the primary signature matching methodology for the detection of the fetal heart tones. The Teager energy signal for each channel was then processed by an autocorrelation, followed by peak picking, to determine a rate estimate for each channel. A figure of merit was also computed for each of these estimates. The fetal heart rate for the current half second was finally determined by comparing these merits.

3.3.2 Sampling Data

Sampling data is the first step of every digital signal processing system. In this application, this step was complicated by the need for demultiplexing,

and the necessity to maintain the best possible SNR. In this section, a detailed description of the sampling procedure for this project is presented. .

Multiplex Control Procedure

As mentioned in the last section, a TDM circuit was required in the system so that all seven acoustic signals could be used in the parallel redundancy technique for heart rate detection. It is necessary to synchronize the multiplexer switching with the sampling of the A/D. The frequency of the multiplexer switch control signal ω_{sc} has to be greater than or equal to the product of the maximum bandwidth ω_{max} of each input signals and the number of channels N_{chan} , i.e.,

$$\omega_{sc} \geq N_{chan} \omega_{max} .$$

Another consideration is that each acoustic signal has most of its energy within a frequency range from 20 Hz to 58 Hz since the pass band of the analog band pass filter is from 20 Hz to 58 Hz (and also as per the results of Pretlow) As mentioned in the last chapter, the Teager energy operator gives a better indication of signal energy if the sampling frequency is at least eight times as high as the maximum frequency of the signal. ω_{max} was thus chosen as 500 Hz. Since there are only four channels for each TDM channel, ω_{sc} was set to two kHz.

Sampling Procedure

For most analog interface circuits (AIC) on a DSP platform, there is an anti-aliasing filter for each A/D channel. The cutoff frequency of the filter is usually half of the sampling frequency. Such a filter can cause a major cross-channel interference effect on a TDM signal if the cutoff frequency of the filter is less than the switch control frequency. In addition, the higher the cutoff frequency is, the less cross-channel interference the filtered TDM signal has. For this study, the AIC was selected to sample the two TDM acoustic signals in parallel with an 8 kHz sampling frequency because it is the lowest available sampling frequency which was also convenient for the subsequent stages of processing.

Consequently, the AIC samples all four acoustic signals every two milliseconds with four successive points for each acoustic signal within each .5 ms time slot of the multiplexer. The third of these four points was selected as the sample for that acoustic signal within the two millisecond interval. The smoothing effect from the anti-aliasing filter was the main reason why the third point was chosen. The resulting demultiplexed signal has an effective sampling frequency at 500 Hz.

In order to minimize quantization noise effects, the programmable internal amplifier for each A/D channel was used to better adjust the dynamic range for the A/D interface. The resolution of the discrete signals was increased by frequently checking the average maximum amplitude of the

discrete signals and setting an appropriate dynamic range for the A/D interface. The average maximum amplitude of the discrete signals was assumed to be three -- four times higher than the energy of the signals in terms of RMS value. The RMS value of a discrete signal is computed as

$$S_{RMS} = \sqrt{\frac{1}{N} \sum_{n=0}^N S^2(n)} ,$$

where N is length of a frame of the signals. However, the dynamic range of the A/D interface should not be changed too often to keep the continuity of the discrete signal. Therefore, an average RMS value based on a ten-second long signal frame was adopted as the criteria of changing the dynamic range. This was computed as:

$$\bar{S}_{RMS} = \frac{1}{20} \sum_{i=0}^{19} S_{RMS}(i) ,$$

where $S_{RMS}(i)$ is an i th RMS value of a half-second long signal frame.

Digital Filtering Procedure

The demultiplexed acoustic signals are then filtered by a 124-order equal ripple digital bandpass FIR linear phase filter. The passband of the filter is

adjustable (via keyboard commands in the real-time system) with the options listed in table 3.1. The attenuation at 60 Hz for all filters was at least 45 dB, and the attenuation for stop bands was greater than 50 dB. In addition, the passband of all the filters was large enough to insure impulse durations were substantially less than the minimal time period between fetal heart tones. Figure 3.5 and figure 3.6 show the frequency response and the impulse response, respectively, for filter two (passband 20 Hz to 50 Hz). The other three filters have similar a response characteristic, except for the choice of cutoff frequencies. The digital filtering increases the SNR of the fetal heart tones by largely suppressing the non-fetal-heart-spectrum signals. After filtering, the acoustic signals were then decimated by a factor of two. The resulting signals then had an equivalent sampling frequency of 250 Hz.

Sensor Identification Procedure

As mentioned above, the multiplexer switching was synchronized to the A/D sampling. The AIC samples each acoustic signal in a fixed order for both A/D channels. However, the multiplexer control signal was initiated a small and uncertain amount before the beginning of the A/D sampling process. Within the SPOX environment, it was not possible to begin these two processes with a fixed timing relationship. Therefore, the acoustic signal sampled in the first multiplexer time slice could be any one of the four signals contained in that TDM signal. Therefore, without further logic, the seven demultiplexed

Table 3.1 124th Order Equal Ripple Digital Bandpass Filter with Sampling Rate at 500 Hz

Digital Filter Options	Pass Bandwidth(Hz)	60 Hz Attenuation(dB)
0	9 - 50	-65
1	16 - 50	-65
2	20 - 50	-65
3	25 - 50	-100
4	0 - 250	-100

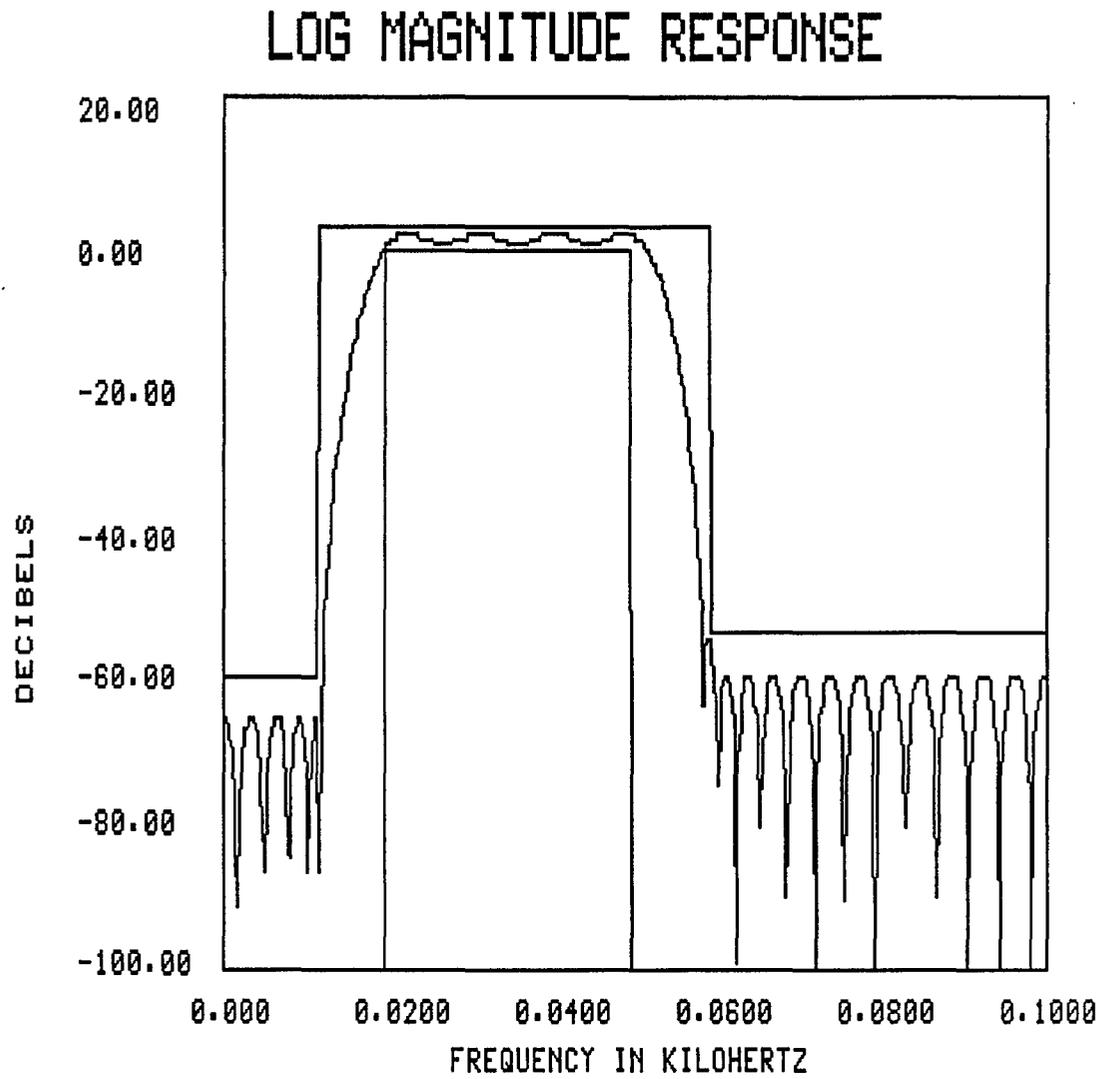


Figure 3.5 Frequency Response of Digital Bandpass Filter Two

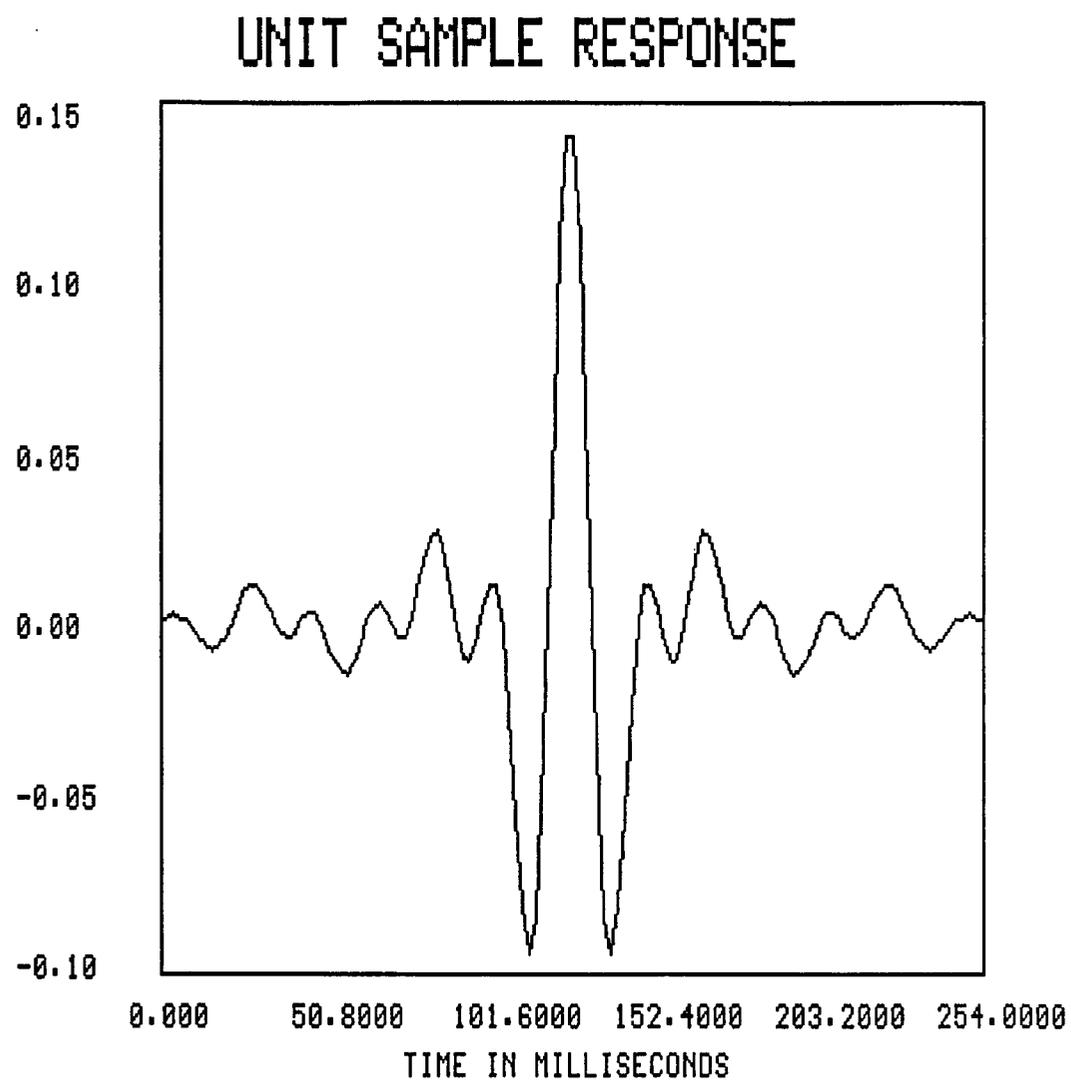


Figure 3.6 Impulse Response of Digital Bandpass Filter Two

signals cannot be identified as to the actual sensor number. To resolve this ambiguity, a ground signal was used in one multiplexer time slice in the second TDM signal (since only three sensor signals were needed for this TDM channel. The detailed procedure for associating multiplexer time positions with sensor signals is described in the next paragraph.

After bandpass filtering, the ground signal channel has nearly zero RMS value comparing to the seven acoustic signal channels. Thus, by comparing the RMS values of the four signals of the second TDM signal, the demultiplexed channel with the minimum RMS values is identified as the ground signal. Since the circular order among the four signals is certain before hand, the absolute identification of the signals is also certain once this reference position is determined. Additionally, since the ordering of the sensor signals is identical for both TDM signals, the sensor order in the first TDM acoustic signal is also determined by this same ground reference signal. This whole procedure was verified experimentally, by introducing sinusoidal signals on known individual sensor channels, and observing the PC display.

3.3.3 Teager Energy Operator

The Teager energy operator was chosen over the neural net as the primary methodology for identification of the fetal heart tones for this version of the monitor. At this stage of development, we did not have enough good acoustic samples of fetal heart tones to adequately train the neural net. In

addition, even with the powerful floating-point DSP processor (TMS320C31), there was not enough processing time for the neural net signature methodology to run the parallel redundance correction algorithm on all channels.

The Teager energy of all seven acoustic signals was first computed as described in chapter two. The output of each Teager energy operator at each sampling instant is the energy representation of the signal at that moment. Each of these energy signals was then slightly smoothed, using a four point boxcar average filter. Although the spectrum of each acoustic signal was restricted to a range of 20 Hz to 50 Hz, which is the optimum for fetal heart tones, non-periodic noise and maternal heart tones within this frequency range would still cause peaks in the Teager energy signals. However, the non-periodic noise can be eliminated by autocorrelation processing, as described in the next section. Maternal heart tones were suppressed from the fetal heart tones using the intelligent figure of merit, as described later.

3.3.4 Autocorrelation

Autocorrelation is a common algorithm to find the period of a periodic signal contaminated by noise. No matter which algorithm is used to identify the acoustic fetal heart tone, autocorrelation can also be used to determine the fetal heartbeat period and convert this to a heart rate in terms of beats per minute (BPM).

The autocorrelation of a stationary signal depends only on the time difference but not time itself. The fetal heart signals are considered as stationary signals within a short period due to their slowly changing tone. So are the Teager energy signals of the fetal heart tones E_n . Therefore, autocorrelation of the Teager energy of the fetal heart tones depends on time difference only, but not time itself. In other words, the autocorrelation of the Teager energy of the fetal heart tones $R_a(k)$ can be computed based on the Teager energy and its time shifted version using the formula :

$$R_a(k) = \sum_{n=0}^{N-k-1} E(n)E(n+k) ,$$

where N is frame length, and k is the time shift. The frame length has to be long enough to cover at least one period of the longest possible signal. For fetal heart tones, the longest expected period is one and half seconds (90 BPM). In addition, the frame length controls the degree to which the method can track variations in the heart rate. Autocorrelation computed with a long frame length will be more accurate for periods with little changes in heart rate, but will not be accurate if the heart rate changes rapidly. Conversely, an autocorrelation computed with a short frame length will be less accurate in periods with little changes in heart rate, but will respond better if the heart rate changes quickly. The frame length should not be too long to be unable to

track expected variations in the fetal heart rate, but also should not too short to lose accuracy in periods of unchanging heart rate. As a compromise, the frame length N was chosen as 750 samples (three seconds) for the tests performed with this version of the monitor.

The value of autocorrelation $R_a(k)$ represents the degree of correlation between the Teager energy signal and itself shifted by k sampling instants. For a periodic signal, $R_a(k)$ is a maximum value if k is equal to a multiple of the period, and is a minimum value if k is an odd multiple of half a period. Non-periodic noise sources do not have this effect on the autocorrelation, and therefore do not typically interfere with period estimation based on the autocorrelation. However, the signals for this study include not only periodic fetal heart tones but also other periodic noises with different intensities and periods than that of the fetal heart tones, such as maternal heart tones. This problem causes multiple peaks in the autocorrelation signal. The largest peak in the autocorrelation occurs at $k = 0$, and is equal to the energy of the signal in terms of its mean-square value. Subsequent peaks correspond to periodic energy bursts from either fetal heart tones or other acoustic noise. If the peak corresponding to fetal heart tones is found, the fetal heartbeat rate is then computed as

$$Rate = \frac{f_s \cdot 60}{k} \quad \text{beat/min} .$$

The next section describes the method used to discriminate peaks corresponding to fetal heart tones from those corresponding to periodic noise.

3.3.5 Intelligent Figure of Merit

Generally, fetal heart rates are in a range from 110 BPM to 180 BPM during the pre labor stage whereas maternal heart rates are typically within a range from 60 BPM to 110 BPM. In terms of *sample* values, the range for fetal heart rates is from 83 to 136, for a sampling rate of 250 Hz. As mentioned in the last section, autocorrelation peaks may represent periods of periodic signals. In this work peaks were defined as values larger than the immediately adjacent M points, with M typically 1. By searching for peaks in a range of 110-180 BPM instead of the full range (30-240 BPM), autocorrelation peaks representing the maternal heart signals are omitted. The peaks in this range are then converted to heart beat rates as mentioned above. However, there is only one rate which is the true fetal heart rate at each time instant. Therefore, an intelligent figure of merit was devised which could be computed for each potential rate to find the most reliable rate as the detected fetal heart rate for each half second. The merit reflects the reliability of each computed rate. The higher value the associated merit is, the more likely the corresponding rate is the true fetal heart rate. There are two primary factors that can effect the value of the merit: the normalized autocorrelation value of the peak, and the deviation of the rate corresponding to that peak from the short term average of previous fetal heart

beat rates. In addition, the average of the previous merits is used to weight the reliability of the rate deviation consideration.

The normalized autocorrelation value of a peak reflects the correlation degree between the processed fetal heart signal and its time shifted version at the shift value corresponding to the peak. The higher the autocorrelation value is, the more correlated the two signals are at that moment. Thus a larger normalized autocorrelation value leads to a higher figure of merit. Additionally, the heart rate generally cannot change too rapidly. Typically the rate does not change more than 28 beats per minute, from beat to beat [2]. Therefore, the rate deviation was set to 30 BPM, in the sense that the detected rate is considered to be unreliable if the rate is 30 BPM or more different from the short term average rate of previously detected rates. The average rate is a weighted combination of previously detected rates, using weight factors based on their reliability (Merit(i)) :

$$\overline{Rate} = \frac{1}{N_R} \sum_{i=1}^{N_R} Rate(i) Merit(i) .$$

The average of previous merits is an indication of the average reliability of of previously detected rates.

The intelligent figure of merit combines all these factors to yield one value. In particular, the merit is computed as:

$$Merit = \frac{R_a(k)}{R_a(0)} \left(1.0 - \overline{Merit} \frac{|Rate - \overline{Rate}|}{\sigma_R} \right),$$

where σ_R is the rate deviation, *Rate* is the detected fetal heart rate for this half second, $R_a(k)$ is the corresponding autocorrelation value, and overline stands for "average of the previous". *Merit* is equal to zero when the absolute value of the difference between *Rate* and average of the previous *Rates* is 30 BPM or more, which means the detected fetal heart rate is totally unreliable. The average of the previous *Merit* also shows the reliability of the average of previous *Rate*. The average of the previous *Merit* is equal to zero when all previous *Rates* are totally unreliable, and eliminates the effect of the rate deviation on the current *Merit* calculation. Typically the averages used in the merit calculations were computed over five frames (2.5 seconds).

A nominal *Merit* threshold was used to determine the reliability of the final detected *Rate*. The *Rate* with a figure of merit above the threshold was considered as a reliable *Rate*. Otherwise, the *Rate* was considered as a "drop-out," which is a unreliable detected *Rate*.

CHAPTER FOUR

REAL TIME IMPLEMENTATION

4.1 Introduction

In contrast with the other FHR monitoring systems in the world, the FHR monitor for this project uses acoustic pressure sensors to detect fetal heart tones. The monitor takes advantage of the floating point digital signal processor (TMS320C31) based platform (ELF-C31) to implement the complex algorithm introduced in the last chapter. The portable PC and the compact electronic support system make it feasible to monitor the fetus at home.

This chapter describes the real time software implementation of the fetal heart rate monitor. Section 4.2 gives an overview of the system software. Section 4.3 and section 4.4 describe in detail the two main parts of the system - the ELF-C31 side and PC side of the software. The PC-ELF handshaking and data communication techniques are explained in section 4.5. Timing and memory considerations are described in section 4.6.

4.2 Software Overview

The real time software implementation of the acoustic fetal heart rate

monitor consists of the PC system software and ELF-C31 DSP software, which run concurrently in real time.

The ELF-C31 DSP software processes all seven acoustic sensor signals for the detection of the fetal heart tones and derives the heart rates. The ELF-C31 software also accomplishes control of the TDM circuit, sensor identification and selection, and audio amplification of the fetal heart signal.

The PC system software is a multitask control/display system software. The system software has several options for displaying acoustic waveforms and rate information computed from the DSP algorithm. Through handshaking with the ELF-C31, the PC also controls the ELF-C31 DSP operation by keyboard. The details of keyboard controls are presented in the appendix. The system software also optionally communicates with the Corometrics unit through a serial port, and detects the push button signal through a parallel printer port.

4.3 A View From The ELF-C31

The ELF-C31 DSP software was coded using a DSP operating system -- SPOX. The advantage of the SPOX DSP operating system is that it is easy to efficiently manipulate data arrays and implement complex algorithms, such as digital filters and autocorrelations. An overview of the SPOX operating system is presented in section 4.3.1.

The ELF-C31 DSP software contains six main routines, including the main routine, the initialization routine, the sampling routine, the Teager energy routine, the neural network routine and the autocorrelation routine. All the routines are modular. Except for the main routine and the initialization routine, each routine has an initialization mode. During the initialization mode, each routine initializes arrays, streams, vectors, matrixes and variables which are used in the routine. Each routine is illustrated in detail in the following subsections, along with the data management. In the present implementation, the neural network routine is only called in the initialization mode, and is not called in the main loop since the Teager energy operator was chosen over the neural network as the primary fetal heart tone detection method for this research. In addition, the ELF-C31 DSP software has two operating modes, automatic mode and manual mode. In the automatic mode, the ELF-C31 software runs the parallel redundance algorithm to detect the fetal heart rate based on all seven acoustic signals. In the manual mode the ELF-C31 software detects the fetal heart rate based only on the selected signal. The default operating mode is the automatic mode, but a keyboard input can be used to switch to the manual mode. More details on mode switching are in the User's Guide in the appendix and in section 4.5.

4.3.1 SPOX DSP Operating System

SPOX is a highly modular and portable run-time environment that can be tailored for a variety of DSP hardware platforms and then integrated for these systems. As illustrated in figure 4.1, the SPOX operating system is organized into two distinct programming interfaces: the *application programming interface* (API) and the *system programming interface* (SPI). The DSP hardware is a TMS320C31 based DSP platform -- ELF-C31 DSP board produced by Atlantic Signal Processing Incorporation(ASPI). ASPI also offers a software system interfacing the SPOX SPI with the ELF-C31 DSP board. The ELF-C31 DSP software of the fetal heart rate monitor is implemented only through the SPOX API.

The SPOX API presents a standard set of high-level, processor-independent functions and objects. The objects and functions, used in the implementation of the ELF-C31 DSP software, are arrays, vectors, matrices, filters and optimized DSP math functions. With these modular objects, it becomes easy to implement complex algorithm. For example, the feedforward neural network procedure is simply a series of product and a summation of input vectors and weight matrices.

Arrays

A SPOX array object references a buffer (a block of memory) allocated by the system memory management routines, which is the base of all the other

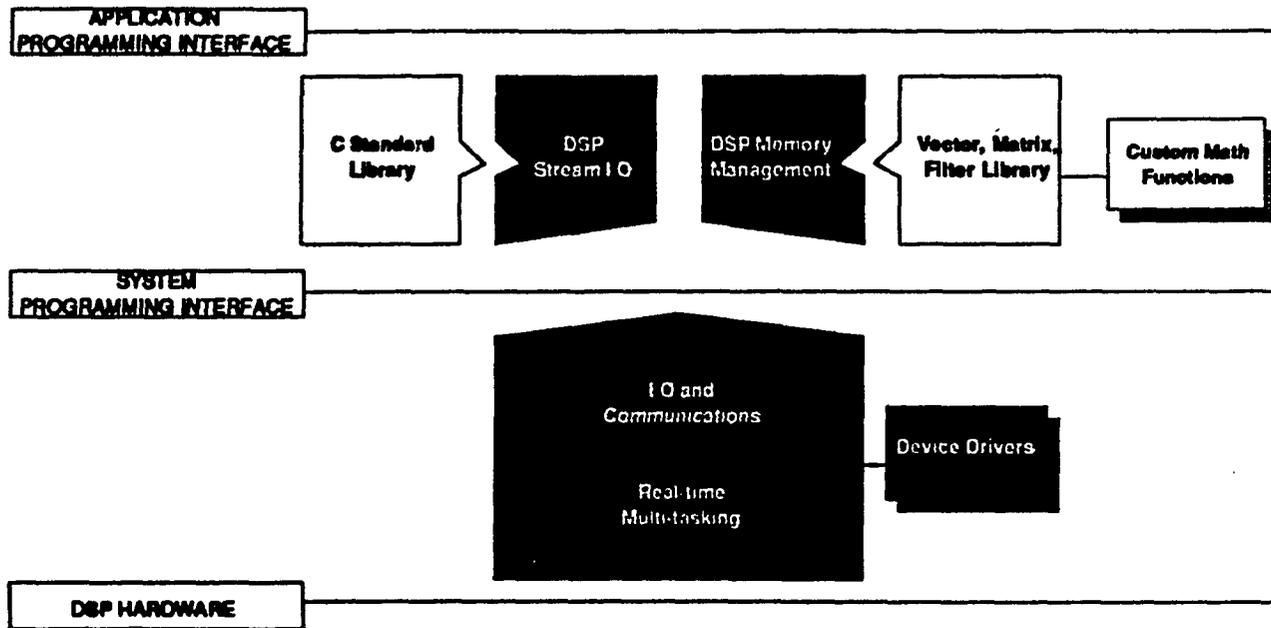


Figure 4.1 SPOX Software Architecture

objects. An array object consists of three major aspects, including size attributes, memory segments and optional attributes,

There are two size attributes associated with an array. The first is the *physical* size of the buffer allocated for the array. The second one is *logical* size indicating the number of bytes that are considered valid. The memory segment contains three values: *SG_CHIP*, *SG_SRAM* and *SG_DRAM*, which correspond to on-chip, off-chip SRAM, and off-chip DRAM memory segments. The optional attributes is a "C" structure consisting of two fields: *align* and *buf*. The *align* determines the boundary of the beginning address of the buffer.

Vector And Matrix

A SPOX vector is a "view" of the memory in an array object. A view imposes a structure on the array's memory buffer, including a type for the data and the number of elements in the memory. Since a vector object only views the memory in an array, more than one vector can be built on an array. Furthermore, more than one subvector can be built on a vector.

Unlike a mathematical vector, a SPOX vector has a base and stride in addition to a length. The base is the starting position of the vector in the array and stride is the offset between successive elements of the vector. The length is the number of elements in the vector. Refer to figure 4.2 as an illustration of these concepts. The length cannot exceed the capacity of the

underlying array. It can, however, be shorter. An element is the storage needed for the type of data stored in the vector.

The length, base and stride of a vector are stored in a view structure, defined as

```

struct  SV_View {
        Int    length;
        Int    base;
        Int    stride;
};

```

The view is dynamic and can be modified at any time using the *SV_setview()* function. Figure 4.2 shows how to access different sets of elements in an array by changing the view of a SPOX vector.

To sequentially access vector elements, use functions *SV_scan()* and *SV_next()*. These functions use cursors of type *SV_Cursor*. A cursor object encapsulates the current view of a vector for fast, sequential access. It is initialized by *SV_scan()*, and then used by *SV_next()* to step through the vector elements. In addition, more than one cursor can be initialized to a vector.

The two-dimensional analog of a vector is a matrix. Like the vector object, a matrix object is a way of viewing the memory in an array object. Instead of viewing the elements as a linear sequence, a matrix views the

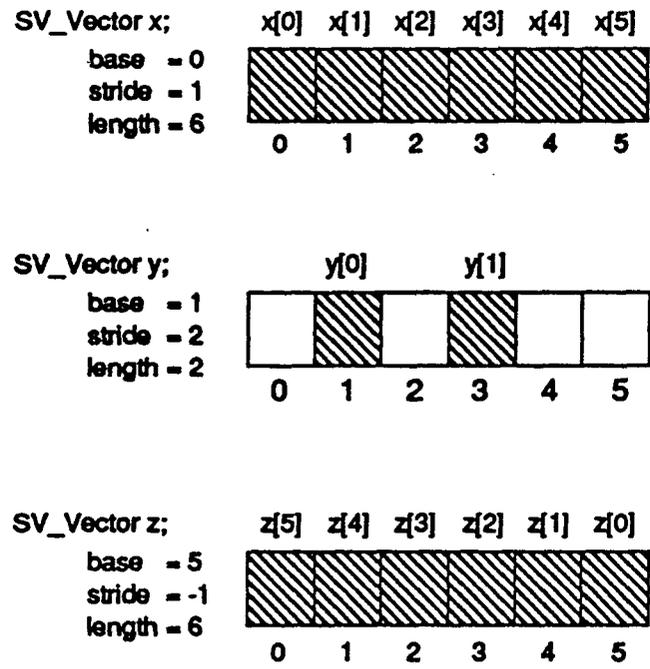


Figure 4.2 Multiple Vector Views

elements as a rectangular array. Figure 4.3 shows the layout of a 3×4 matrix. Note that a matrix object is laid out in row-major order.

Filters

A filter is a linear system with memory: for each input sample a filter produces an output sample. There are two general classes of digital filters: finite impulse response (FIR) filters and infinite impulse response (II) filters. The SPOX filter object encompasses the data and attributes needed to describe both types of filters.

The filter object contains state arrays, memory that is used to store the state of the filter. Vectors of coefficients used to implement the filter's transfer function are associated with a filter object using *SF_bind()*.

The attributes of a filter object describe the data type used by the filter and the filter's structure. Filters operate on *Float* values for this research. The structure of a filter determines how the filter is actually implemented. Currently, SPOX only provides direct-form filters. For direct-form filters, the state arrays are delay lines, simulated using circular buffers. The most efficient (in terms of execution speed) handling of the circular buffers occurs when the arrays are aligned. The *align* field of the state array must be on a power of two boundary such that the alignment is greater than the array length. For example, if the state array length is 12, the *align* must be 16.

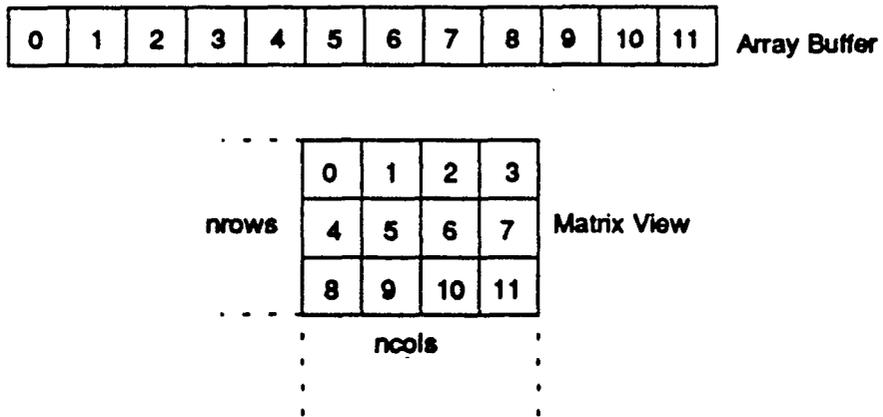


Figure 4.3 Matrix View of an Array

Ping-Pong Buffering

The sampling data management scheme is based on the double buffering technique known as "ping-pong buffering". In this technique two identical buffers are used. While one buffer is used to collect input data, the second buffer is used for processing. When the input buffer is completely full the buffers are swapped and the roles are inverted. In the ELF-C31 platform this operation is transparent to the user. However, certain attributes have to be determined in advance, such as the full scale range of the A/D convertor, data type and data range of the sampled data. The input/output streams and the corresponding I/O arrays must also be declared in advance.

4.3.2 Main Routine

The Main Routine shown in figure 4.4 constitutes the core of the ELF-C31 DSP software. After the Initialization Routine, the Main Routine enters a main loop and calls each of the other routines in sequence. Control parameters are first input and the loop repeats indefinitely.

The main loop starts with the Sampling Routine which samples all seven acoustic signals and transmits the acoustic signal from sensor four to the PC (for the automatic mode). Due to timing and memory constraints on the PC, only one acoustic channel can be saved. Sensor four was selected in the automatic mode because of its central location on the sensor belt. In the manual mode of operation, the manually selected sensor channel is transmitted

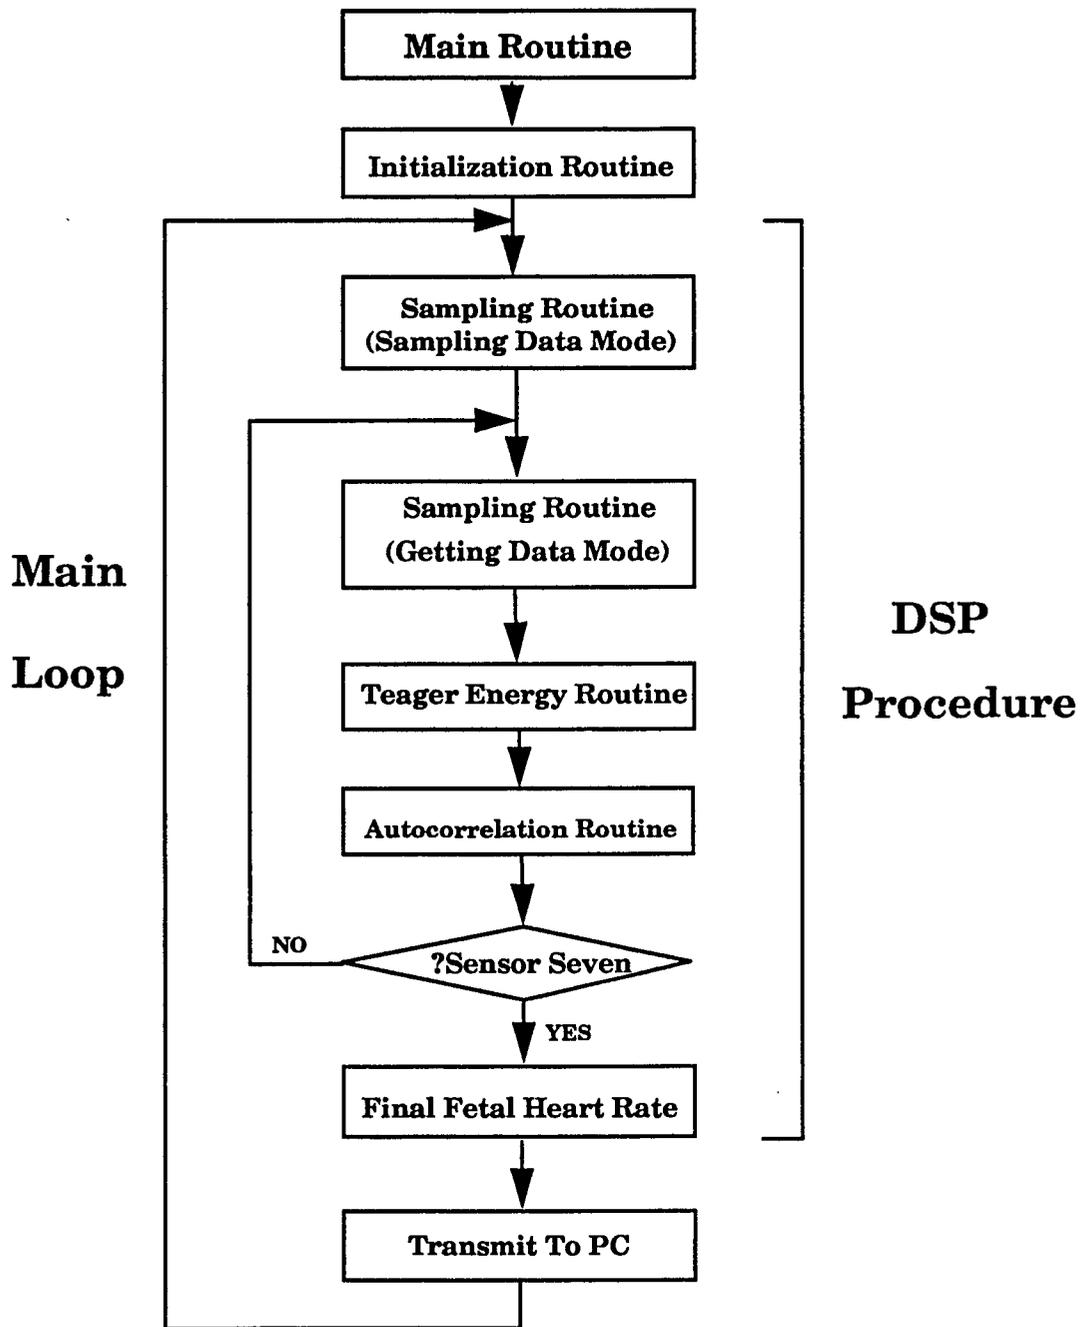


Figure 4.4 Main Routine

to the PC. In the automatic mode, all seven acoustic signals are then processed by the Teager energy operator. The Teager energy signals are passed to the Autocorrelation Routine. Based on the computed autocorrelation, one heart rate and its figure of merit are computed for each acoustic signal. Note that the "parallel" processing of the seven signals was actually done sequentially, due to the present hardware configuration. The final fetal heart rate is selected among these seven heart rates based on their merit values and transmitted to the PC.

4.3.3 Initialization Routine

On start up the Initialization Routine, shown in figure 4.5, calls each routine in an initialization mode and initializes several array objects, vector views and seven digital filter objects. In the initialization mode, the Neural Net Routine creates vectors and weight matrices for both the input layer and the hidden layer. The neural net weights are down-loaded from the PC and filled in the weight matrices. The Autocorrelation Routine creates and initializes the autocorrelation vector, the heart rate vector and the merit vector in the initialization mode.

To initialize the Sampling Routine, two vectors associated with both right and left A/D channels are created with a length of a half second of data for a sampling rate of 8 kHz (i.e., 4000 points each) and seven subvectors are created for seven acoustic signals respectively. By accessing the subvectors

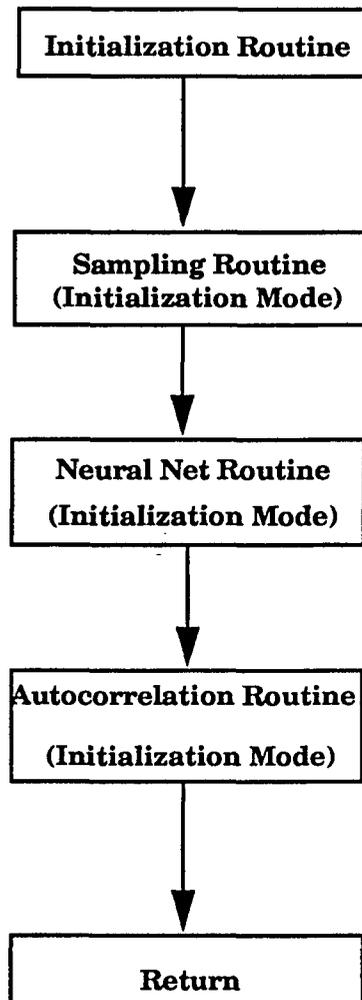


Figure 4.5 Initialization Routine

directly later in the main loop, the two TDM signals are demultiplexed to seven acoustic signals. Seven digital filter objects are also built for the acoustic signals respectively. One vector is created for the filtered signal with an effective sampling rate 500 Hz, and another subvector for the decimated signal with a factor of two is built on this vector. One global signal processing vector is built with a maximum length of six seconds sampling data at 250 Hz. The global vector is a data transferring station between the sampling routine and the other routines. Finally, one vector is built as the source memory buffer of the DMA queuing system. The default settings for AIC are also down-loaded along with those neural net weights. The sampling rate is set at 8 kHz. The dynamic range of the A/D convertor is initialized for a range of ± 2.8 volts. The data type and data range of the acoustic data are set as signed 16-bit integer and from -32,678 to 32,677 respectively. In addition, the ordering of the sensors is determined before entering the main loop.

After initializing all the other routines, the ELF-C31 software waits for a start command from the PC system software, and acknowledge the command reception by sending the start command back to PC. The ELF-C31 software then enters the main loop. The start command synchronizes the ELF-C31 software to the PC system software.

4.3.4 Sampling Routine

The Sampling Routine, as shown in figure 4.6, consists of two real-time operating modes besides the initialization mode, including the sampling data mode and the getting data mode. The Sampling Routine is first called in the sampling data mode by the main routine, and again in the getting data mode seven times before repeating the main loop.

In the sampling data mode, the Sampling Routine implements the sampling data procedure as described in chapter three, replays the fetal heart tones, and transmits the filtered acoustic signal from sensor four (or the selected channel if in manual mode) to the PC by the DMA queuing system (as described below in section 4.5). The fetal heart signal has to be interpolated to an 8 kHz rate before playback, since the DA rate is at 8 kHz (due to the necessity to generate the multiplexer control at this rate). Thus each sample of the acoustic signal is passed through a 31 point digital zero-order hold, and then lowpass filtered at 300 Hz before presentation to the DA channel. The resultant audio signal is valuable for positioning the sensor belt.

In the getting data mode, the Sampling Routine fills up the global signal processing vector for the following signal processing procedure. The global signal processing vector contains an acoustic signal with an effective sampling frequency of 250 Hz. It has a frame length of 3.0 seconds and a frame space of 0.5 second.

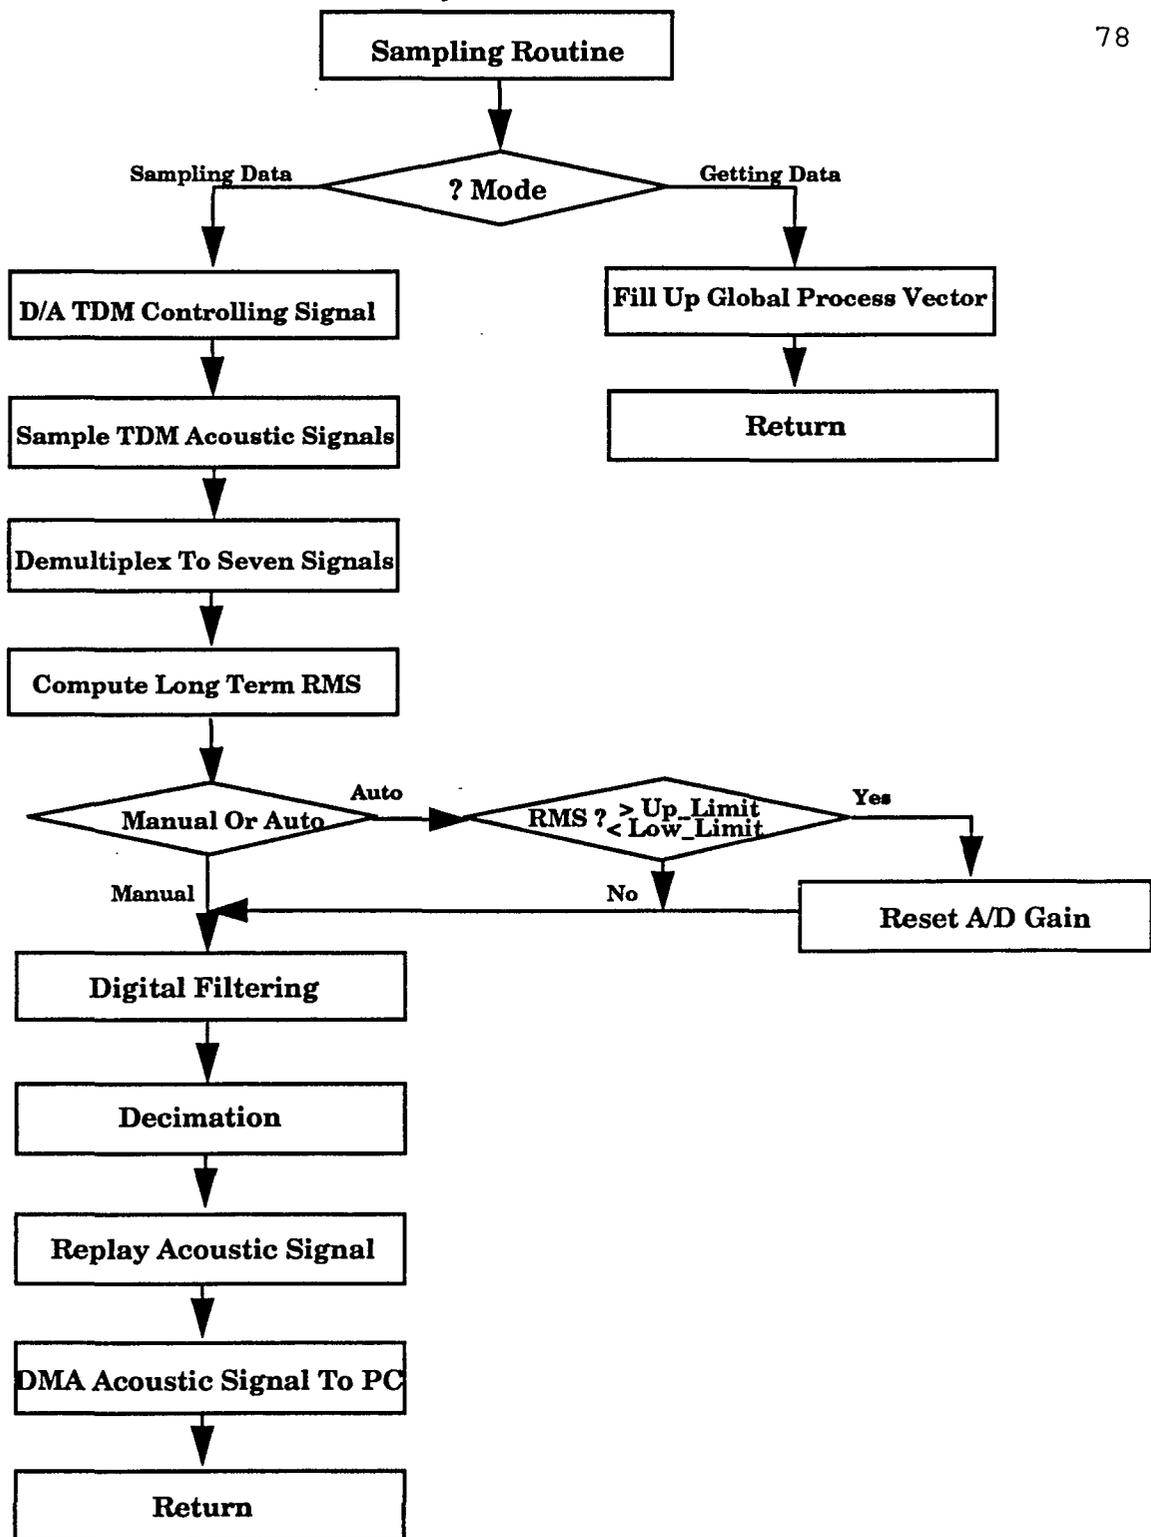


Figure 4.6 Sampling Routine

4.3.5 Teager Energy Routine

The Teager Energy Routine, as shown in figure 4.7, sequentially accesses elements of the global signal processing vector, and computes the Teager energy for each element as

$$E_n = S_n^2 - S_{n-1}S_{n+1}$$

The resultant signal is then filtered by a four-point FIR "boxcar average" filter. Two cursors are initialized to the global signal processing vector. One is for getting the acoustic signal data from the vector, and the other is for putting the Teager energy signal in the vector. Figure 4.8 shows an acoustic signal and its corresponding Teager energy signal.

4.3.6 Neural Network Routine

The Neural Network Routine implements the feedforward neural network computation in real time. It computes an output of the feedforward neural network for each data point in the global processing vector. In this way, the routine generates another signal at 250 samples per second, which emphasizes fetal heart tones, for the autocorrelation routine. In terms of a real time implement, the feedforward neural network is easily and efficiently implemented under the SPOX operating system whereas the previous version of the monitor did not have sufficient computational speed to implement this algorithm in real time.

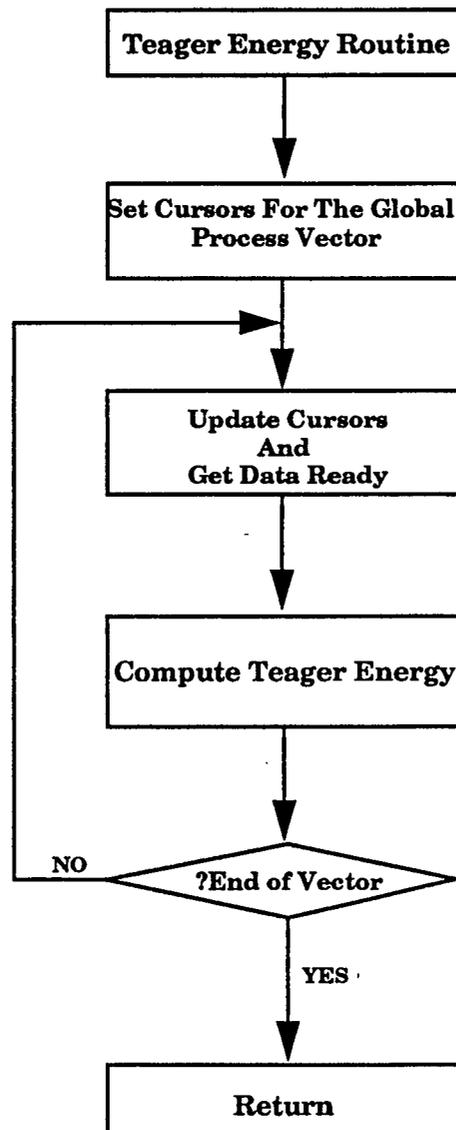


Figure 4.7 Teager Energy Routine

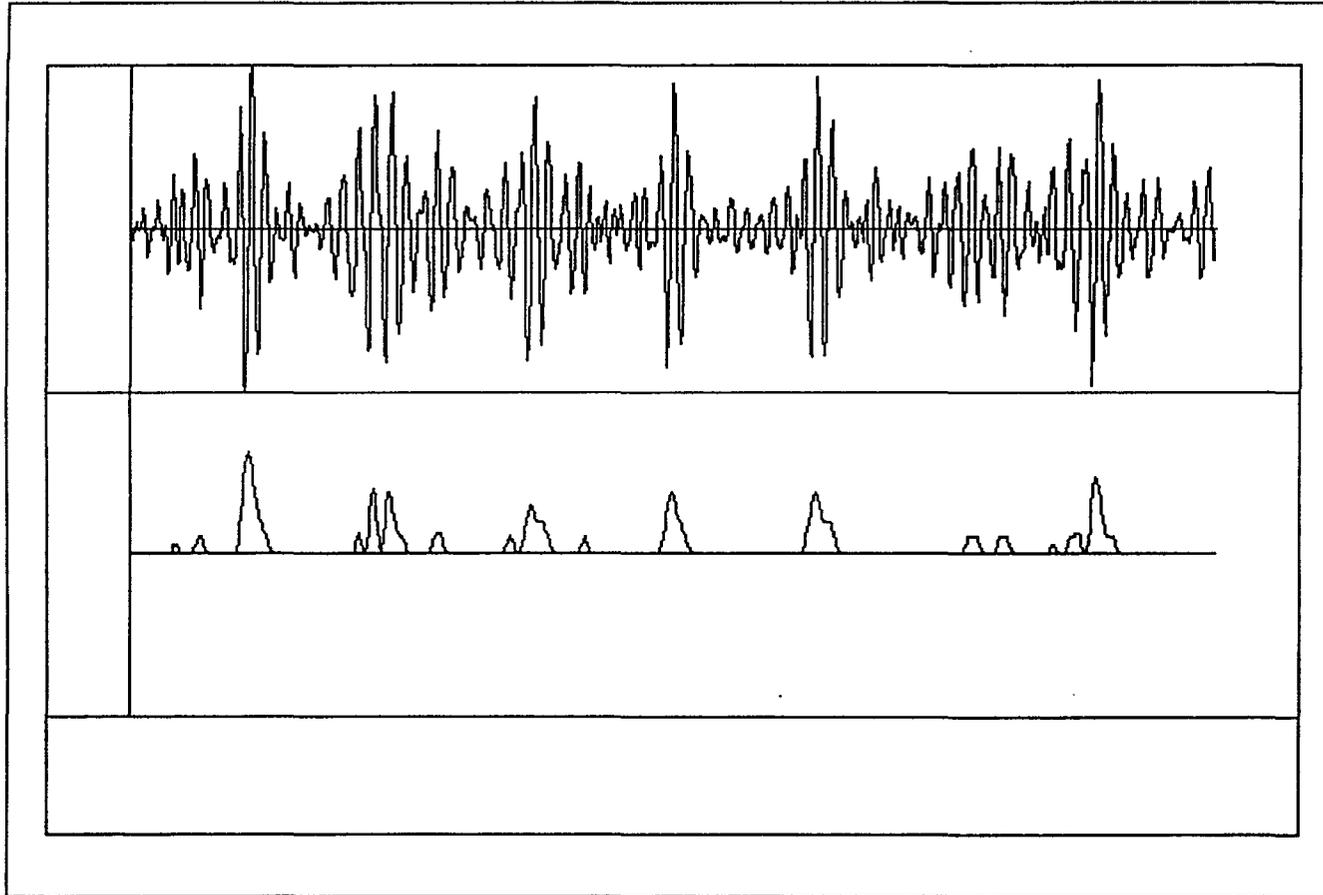


Figure 4.8 Three Seconds of Acoustic Fetal Heart Signal (Top) and Corresponding Teager Energy (Bottom)

With the SPOX DSP operating system, as described above in section 4.3.1, products between a floating-point vector and a floating-point matrix are efficiently done. This is the fundamental time-intensive computation for the feedforward neural network, as shown in chapter two. SPOX offers several optimized functions for operation on vectors and matrices. For example, one function called "SV_dotp" computes a dot product of two vectors. Another function called "SM_prodv" calculates a matrix-vector product. With these functions, the feedforward neural network algorithm is easily implemented in real time.

However, the Neural Network Routine was implemented primarily for future development of the fetal heart rate monitor, and was therefore not used for most real-time tests. As mentioned previously, we do not yet have enough training data to properly train the network. However, in limited testing of the Neural Network Routine, performance was good in detecting the fetal heart tones, sometimes better than for the Teager Energy approach. Figure 4.9 shows the same acoustic signal as shown in figure 4.8 and the corresponding neural network output signal.

4.3.7 Autocorrelation Routine

The Autocorrelation Routine shown in figure 4.10 implements the autocorrelation and intelligent figure of merit procedures as described in section 3.3.4 and section 3.3.5 respectively. The routine computes the

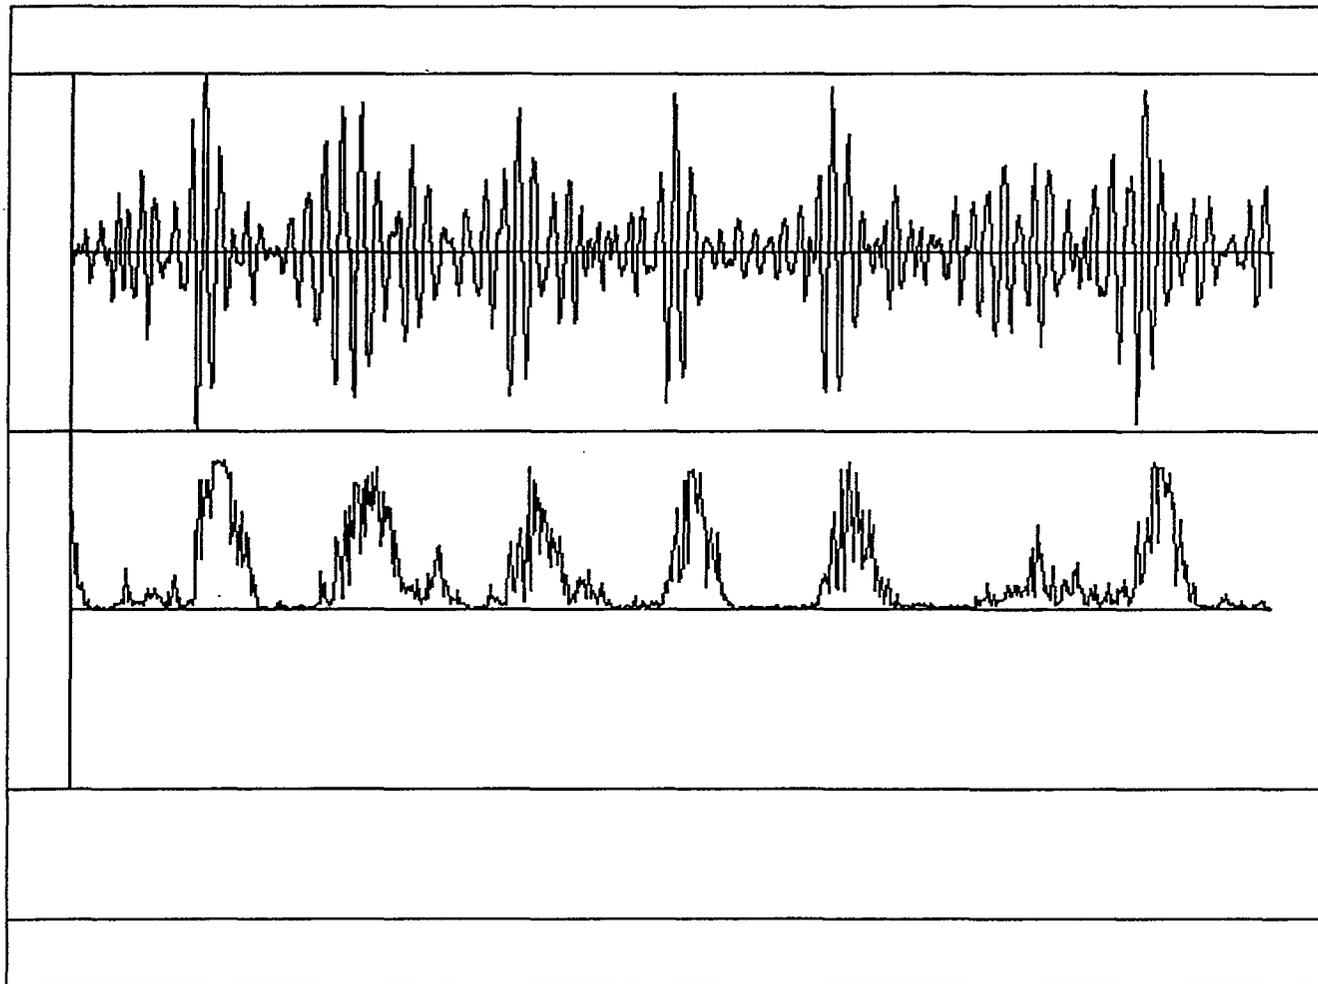


Figure 4.9 Three Seconds of Acoustic Fetal Heart Signal (Top) and Corresponding Neural Network Output (Bottom)

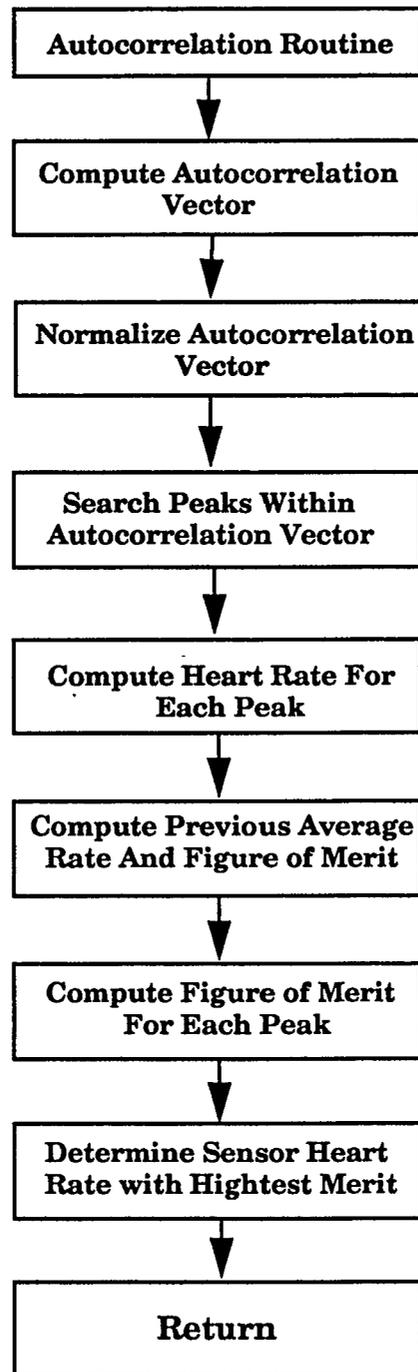


Figure 4.10 Autocorrelation Routine

autocorrelation of the Teager energy signal in the global signal processing vector, normalizes the autocorrelation vector by dividing by $R_a(0)$, searches for peaks from $R_a(83)$ to $R_a(136)$ in the autocorrelation vector, and computes the corresponding rate. The figures of merit are also computed for these autocorrelation peaks. The rate with the highest figure of merit is selected as the fetal heart rate for this frame of data. The fetal heart rate and its associated merit are stored for the final selection as described in section 4.3.2.

4.4 A View From The PC

The PC system software, shown in figure 4.11, has five options containing more suboptions (see more information in appendix -- User's Guide) and six main tasks during its real time operation. The five options are Corometrics Unit Only, Acoustic Unit Only, Both Units, One Channel Oscilloscope and Seven Channel Oscilloscope. The five main tasks are to receive and store both acoustic data and heart rate information, to display in real time both acoustic data and heart rate information, to communicate with the Corometrics unit through the RS232 serial port, to detect the closing of the push-button through LPT1 parallel printer port, and to control the monitor operation through keyboard entries. The PC system software consists of four main routines to implement these five options and six main tasks, including initialization routine, main loop routine, serial communication routine, and push-button routine.

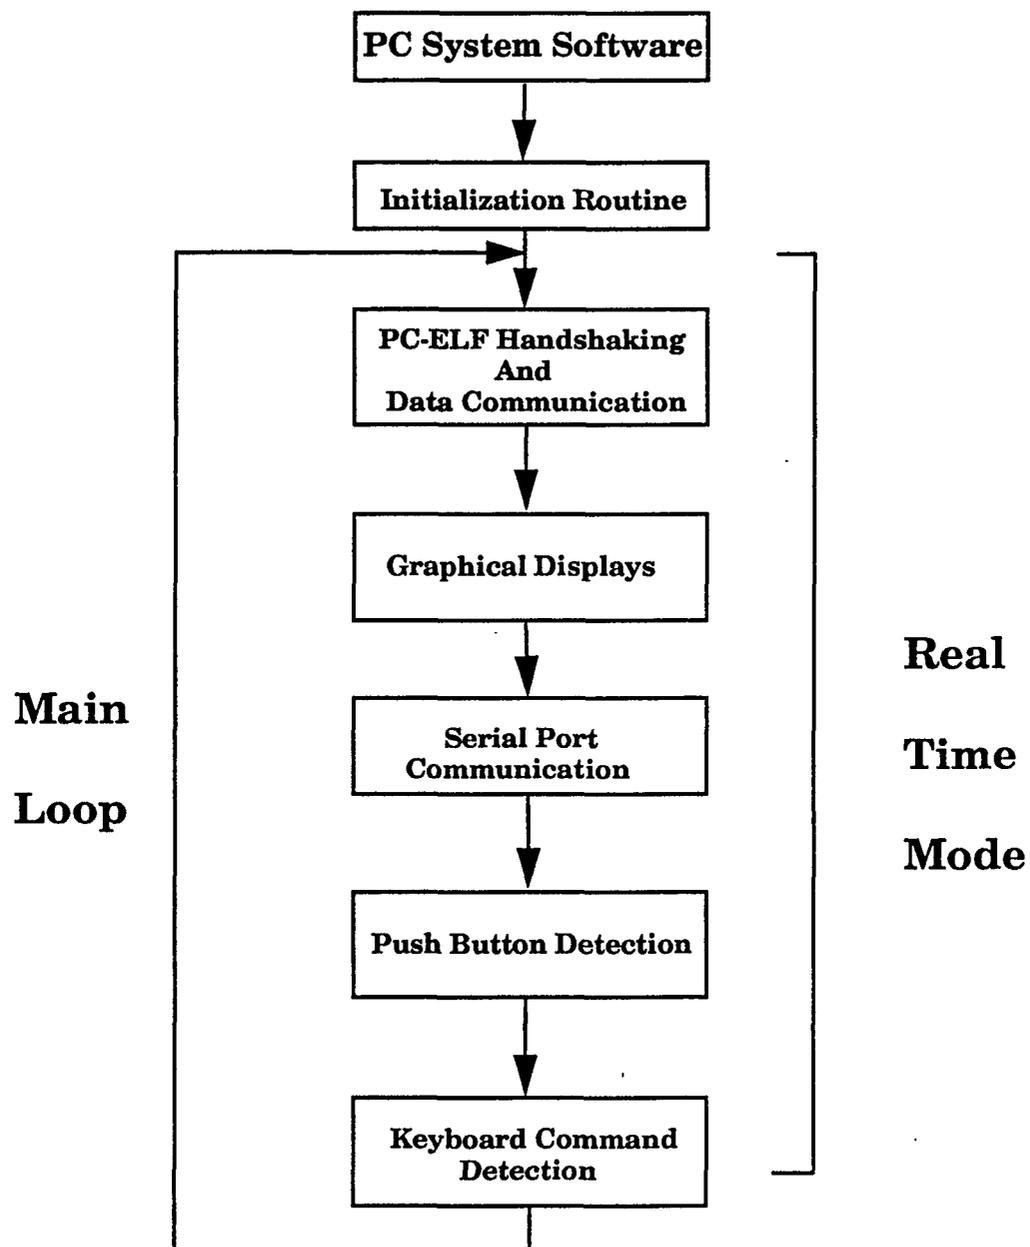


Figure 4.11 PC System Software Block Diagram

The PC system software also contains four main graphic displays, including one menu screen, one four-second oscilloscope screen, one seven channel four-second oscilloscope screen and one eighteen-minute heart rate chart. The menu screen graphically shows all five options and their corresponding keyboard commands. The heart rate chart shows an eighteen-minute-long fetal heart rate trace. Both oscilloscope screens are used to improve the front-end electronics support system and to tune up the final settings of the software system, and is simultaneously run with the heart rate chart in both the Acoustic Unit Only option and the Both Units option. However, only one of them is actually shown on the screen, and the other one is run in the background. It is easy to switch back and forth between these two graphic displays, using keyboard commands.

4.4.1 Initialization Routine

The Initialization Routine is the first step of the PC system software. A menu screen is first initialized on active page zero and displayed after starting the FHR monitoring system. After selecting an option, the ELF-C31 DSP software is downloaded to the ELF-C31 board, the corresponding graphic displays are initialized, and serial port communication or parallel printer port communication is initialized.

For both the Corometrics Unit Only and the Both Units options, serial port communication is initialized at COM1 serial port with 2400 baud rate and

no parity. For both the Acoustic Unit Only and the Both Units options, parallel printer port communication is initialized using LPT1 parallel port. The oscilloscope screens are initialized on active page one, and the heart rate chart is on active page zero. If two displays are initialized simultaneously, for example, the Acoustic Unit Only option and the Both Units option, the heart rate chart is the default visual page.

Right before entering the main loop, the PC sends a starting command to the ELF-C31 board and waits for an acknowledge from the ELF-C31 board.

4.4.2 Main Loop Routine

After the Initialization Routine, the Main Loop Routine steps through all five main tasks in sequence and repeats the loop every half second. For each task, a flag shows the status of the task. The Main Routine checks every flag, and executes the task when its flag is on.

The Main Loop Routine first checks if there are acoustic fetal heart signals or fetal heart rate information waiting at an I/O port. If the check is positive, it receives the data, stores the signals into a file with extension ".wav" and the rate information into a file with extension ".rat", and sets a data flag indicating that there are new data coming. The routine checks both the system time clock and the data flag. It plots the acoustic signal if the data flag is on. Rate information is updated every second. In the both Units mode, the routine receives a fetal heart rate from a Corometrics unit through the

serial communication as described below in section 4.4.3. The routine also checks the push-button signal (more detail given in section 4.4.4), and plots a marker on the heart rate chart if the button was pushed. The routine finally checks keyboard commands and executes them. More details of the keyboard commands are presented in the appendix -- User's Guide.

4.4.3 Serial Communication Routine

The serial communication between the Corometrics unit and the acoustic FHR monitor is one way communication. The Corometrics unit transmits the fetal heart rate information, and the Serial Communication Routine receives it and plots it on the heart rate chart. Therefore, it is important to describe how the Corometrics unit transmits the rate information.

The serial communication mode of the Corometrics Unit is set as the Update mode. In this mode, all information including the fetal heart rate is transmitted from the unit as soon as the information becomes available. The fetal heart rate is transmitted on a beat-to-beat basis. The transmitted data consists of a monitor type, response type, monitor ID, data field, and end of text. For this research, the heart rate information only was of interest. The transmitted heart rate has values listed in table 4.1. Note that the data field contains ASCII characters making up the heart rate.

The Serial Communication Routine starts by checking the response type. If the response type is 0x60, the routine starts receiving a heart rate until an

Table 4.1 RS232 Transmitting Format of Fetal Heart Rates In The Update Mode of Corometrics 116 Unit.

Field Description	Field Value
Monitor Type	2
Response Type	0x60
Monitor ID	0
Data Field	Fetal Heart Rate In ASCII Code
End of Text	0x03

end of text command is received. The received data field is then converted to a heart rate in decimal. The heart rate is then returned to the Main Loop Routine.

4.4.4 Push-Button Routine

The push-button detector in the electronics supporting system sets logic "low" when the button is being pushed, and "high" when the button is released. To send this two-state message to PC through the parallel printer port, the FHR monitor simulates the detector as a "printer" in one status and uses one of the status lines used by the "printer." Pin 10, which is the acknowledge status of the printer, is chosen to connect the detector. The remaining status line settings are listed in table 4.2.

The Push-Button Routine checks the status of the "printer", and plots the event trace on the heart rate chart. The status of the "printer" -- 0xC0 for a pushed button, is indicated by a red event trace at 60 BPM on the heart rate chart. When the button is released (0x80), the red trace is drawn at the 30 BPM position.

4.5 PC-ELF Handshaking And Data Communication

There are two basic ways to communicate between the Elf-C31 DSP board and the PC. One is through polling whereas the other is through direct memory accessing (DMA). In both ways a data exchange register is used as

Table 4.2 Parallel Printer Port Status Line Settings

Pin Number	Signal Name	Meaning	Logic Value
10	-ACK	Last character received	Active
11	-BUSY	Printer busy	Low
12	PE	Printer has no paper	Low
13	SLCT	Printer is online	Don't Care
15	-ERROR	Data transfer error	High
18 - 25	GND	Ground	Low

an output port to the PC. However, polling and DMA use different memory addresses as their data exchange register from the point of view of the ELF. The strategy used for controlling the transfer is different for the two cases. However, communication cannot be done in both ways simultaneously. The advantage of using the DMA controller is that the communication runs in parallel with other signal processing. In general, the DMA data transfer is much faster for transferring blocks of data. On the other hand, polling can be used to synchronize the Elf processing and the PC processing. Polling was used, for example, to synchronize the ELF main loop with the PC main loop as mentioned in section 4.3.3.

In this implementation, the sampled acoustic data are passed through the DMA to the PC and stored in a file named with extension ".wav". However, the detected heartbeat rates are transferred to the PC through polling and saved in a file with extension ".rat". In addition, keyboard commands are also passed to the ELF board through polling. Figure 4.12 shows the structure of the command register.

4.6 Memory And Timing

Memory and timing are two critical points in real-time system programming. For this research, memory was not a point of concern because there are a total of one Megabyte memory on the ELF-C31 board, more than enough for the fetal heart detection task. Table 4.3 lists the system buffers

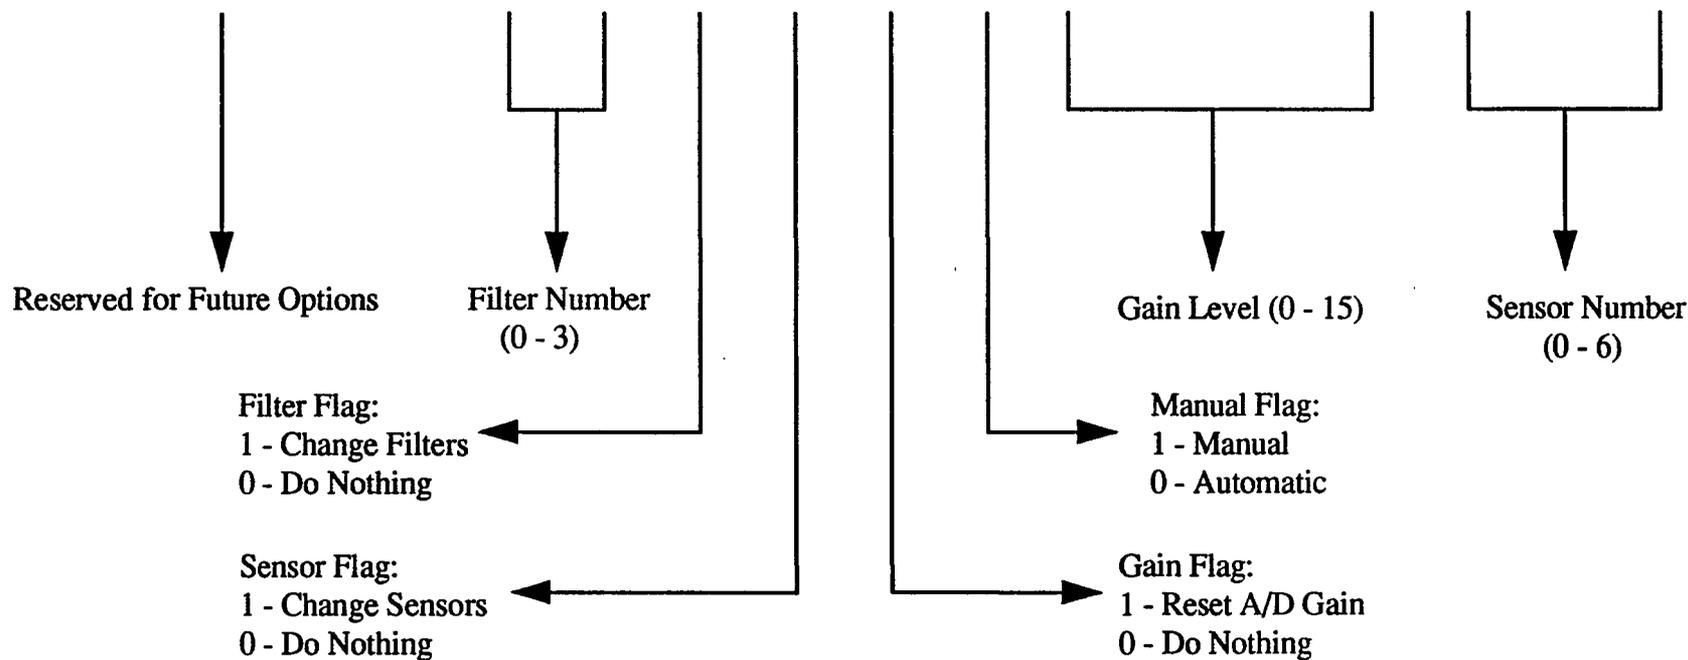
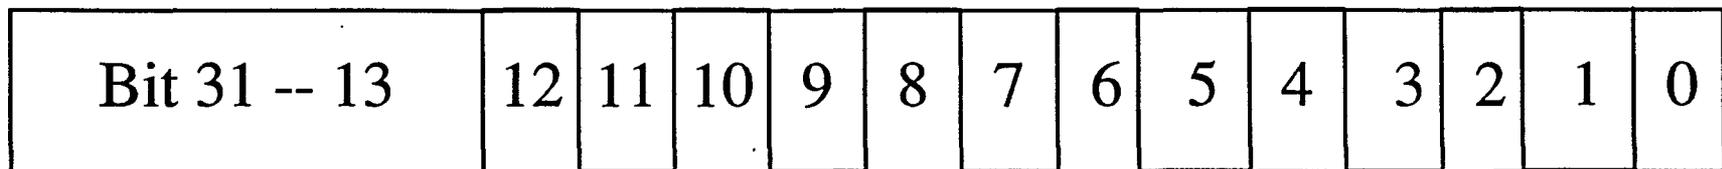


Figure 4.12 Structure of Keyboard Command Register

Table4.3 System Buffers and Memory

System Buffer	Buffer Description	Memory (element)
right_vec, left_vec	Ping-Pong buffers for both A/D channels	4 k for each one
out_Rvec	TDM control data buffer at right channel of D/A	4 k
out_Lvec	Play back data buffer at left channel of D/A	4 k
Proc_vec	Global signal processing buffer	750
Circul_?(1 - 7)	Data buffer for each sensor	750 for each buffer
autocor	Autocorrelation Buffer	750

and their corresponding memory. Timing was the major real-time concern for the signal processing. Compared to the previous FHR monitor, the current FHR monitor has much more computations and more additional features in real time operation, such as audio replaying and simultaneous monitoring of all sensors.

ASPI offers three functions to measure real-time code execution times. First, one function named "set_period" initializes a timer. Another function "start_timer" starts the timer operation. The other one "get_time" stops the timer and counts the elapsed clock cycles. Each clock cycle stands for 120ns. Therefore, the time needed for the code between "start_timer" and "get_time" is equal to the number of clock cycles times 120ns. Using this type of timing system, the operation of the system was observed. The timing results for various of the system operations shown in figure 4.5 are presented in table 4.4.

As shown in table 4.4, the fetal heart rate monitoring system can operate in real time. The time needed for processing one full main loop is 468.87ms, which is less than the available time -- 500 ms. These tests were made using a frame length of three second, and an autocorrelation search range of 110 to 180 BPM. Only 136.55 ms, which is one third of the available time, is consumed for implementing the detection algorithm for the fetal heart rates. Most of the time in the main loop is for the Sampling Routine in the sampling data mode, which contains complicated implementation of several features. The advantage of including accessory features is to reduce the physical size of the system. Also note that the Neural Network Routine takes about 100 times

Table 4.4 Timing Results For Each Routine Within The Main Loop

TASK DESCRIPTION	TIME(ms)
Sampling Routine In Sampling Data Mode	332.32
Sampling Routine In Getting Data Mode	0.462
Teager Energy Routine	3.329
Neural Network Routine	446.53
Autocorrelation Routine	15.766
Main Loop Operation	468.87
Total Available Time	500

longer than the Teager Energy Routine. It is impossible to implement a real-time system, using the neural network method operating simultaneously on all seven channels of data, with the current DSP power. In addition, the neural network method cannot operate on one channel of data either under the current data management. However, it is possible to change the data management to operate the neural network method on a single channel of data.

CHAPTER FIVE

EXPERIMENTS AND CONCLUSIONS

5.1 Introduction

The objective of the research is to develop and implement a real time signal processing algorithm for a portable acoustic FHR monitor. The real time implementation of the algorithm was presented in the previous chapter. The experimental study to validate the current version FHR monitor is presented in this chapter. Software verification was first conducted at a speech laboratory at O.D.U. Simulation tests were next performed at a NASA laboratory. Clinical tests were then performed at Norfolk General Hospital with the current version of the FHR monitor, and the output was compared with that of a Corometrics Model 115.

5.2 Technical Testing

To experimentally verify the correct operation of the FHR monitor prior to using expectant women, two tests were conducted to show the reliability of the monitor. The first test used a fetal heart beat signal reconstructed from previous sampled acoustic data. This testing could be used to validate the signal processing software. The second test, which used a signal generator,

amplifier, and transducer in a water-filled balloon to simulate an in utero fetus, was conducted at NASA Langley to show the reliability of the FHR monitor in terms of both hardware and software.

5.2.1 Software Verification

The software verification began with reconstructing a fetal heart beat signal. Previously sampled acoustic data from one of Pretlow's patients' records -- pt23_new.sig, was used as the signal.. One additional PC station with a TMS320C25 based DSP platform was used as a signal generator to reconstruct an analog fetal heart beat signal from "pt23_new.sig". The analog signal was then connected directly to the ELF-C31 board AD input (i.e., the analog electronics was not needed). The detected heart rate charts, computed with the Teager Energy method and the Neural network are shown in figures 5.1 and 5.2 respectively. The ultrasonic and LMS acoustic results [1] are shown in figure 2.4. By comparing these figures, it is clear that the new algorithms results in a heart rate trace comparable to the older results. In fact, at least for the Teager Energy method, the heart rate trace appears to be more similar to the ultrasonic trace than was the trace obtained with the LMS method. This test, plus other similar ones, was used to verify that the basic operation of the software.

FETAL HEART RATE

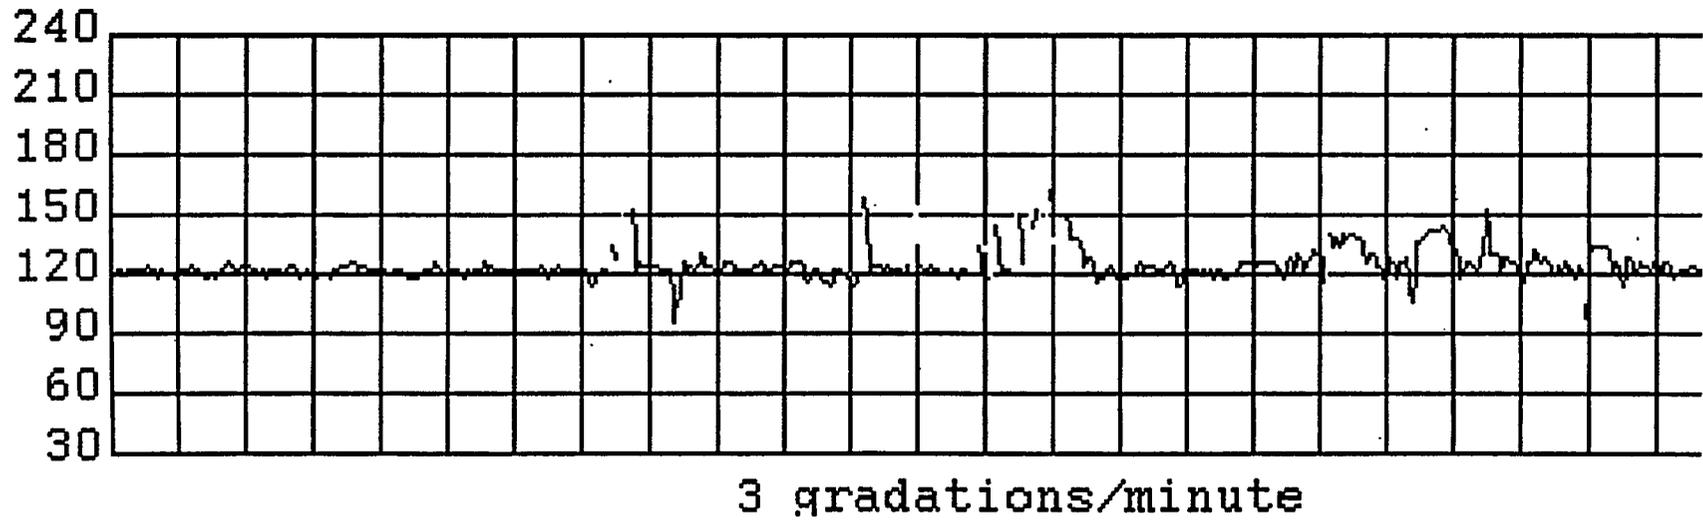


Figure 5.2 Fetal Heart Rate Chart For : Patient 23 (Neural Network)

FETAL HEART RATE

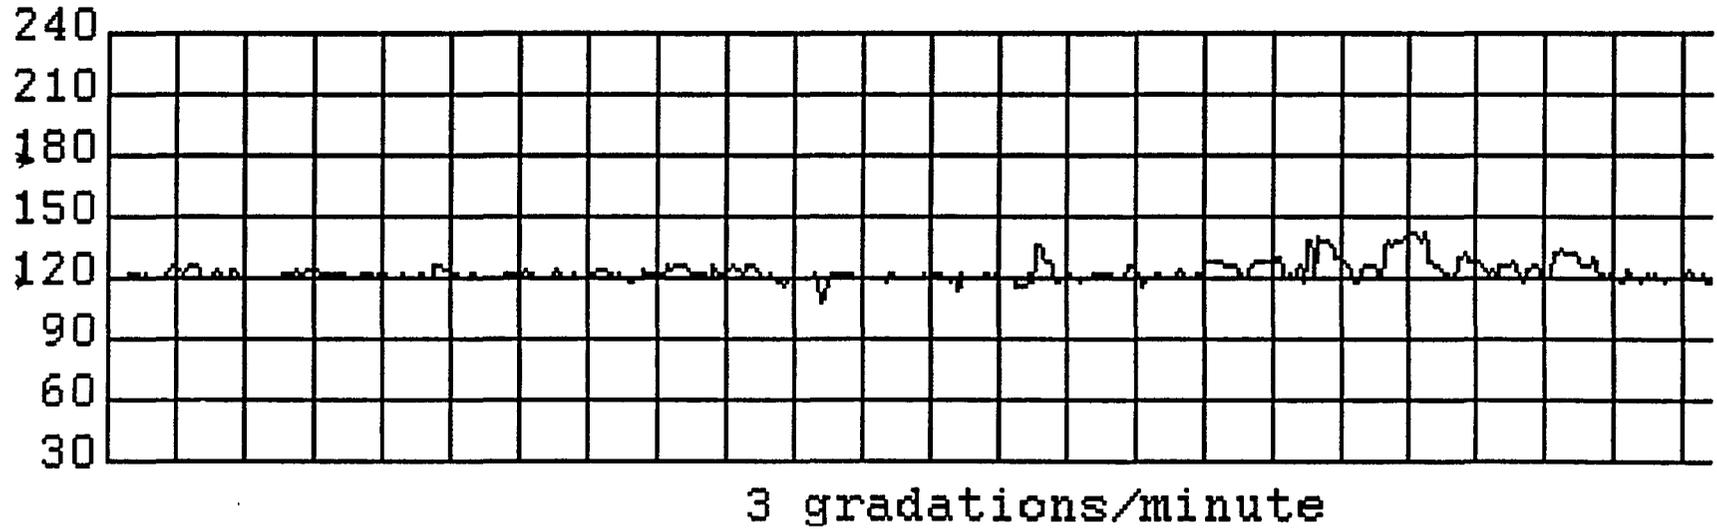


Figure 5.1 Fetal Heart Rate Chart For : Patient 23 (Teager Energy Operator)

5.2.2 Balloon Test

To verify the correct operation of the FHR monitor in both hardware and software, a simulation test, using a transducer in a water-filled balloon, was conducted at NASA Langley. The experimental arrangement for this test is shown in figure 5.3, comprising an acoustic coupling medium, the sensor belt, a sound source, a reference sensor, and supporting instrumentation. Detailed information is presented in [3].

The acoustic coupling medium comprises a water-filled balloon, made of 150- μm -thick vinyl and having a diameter of 0.356 *m*. The neck is opened to accommodate an aluminum tube, which seats the sound source. An o-ring presses the neck against the tube, thus providing a seal. The balloon is supported on a foam pad to reduce the effect of ambient vibration. The sensor belt is wrapped around the balloon with the sensors in contact with the vinyl surface. Such an arrangement simulates the maternal abdomen with fetus and amniotic fluid.

In the present test, the sound source was a loud speaker suspended on the surface of the balloon. A preprogrammed function generator was used to simulate a fetal heart beat waveform. By repeating the heart beat, a continuous fetal heart beat signal was generated. The spectrum of this signal is shown in figure 5.4. The spectrum of this signal is quite similar to the spectrum shown in figure 2.3. By controlling the repetition frequency, a fetal heart signal with a known heart beat rate is simulated. Furthermore, a noise generator was used

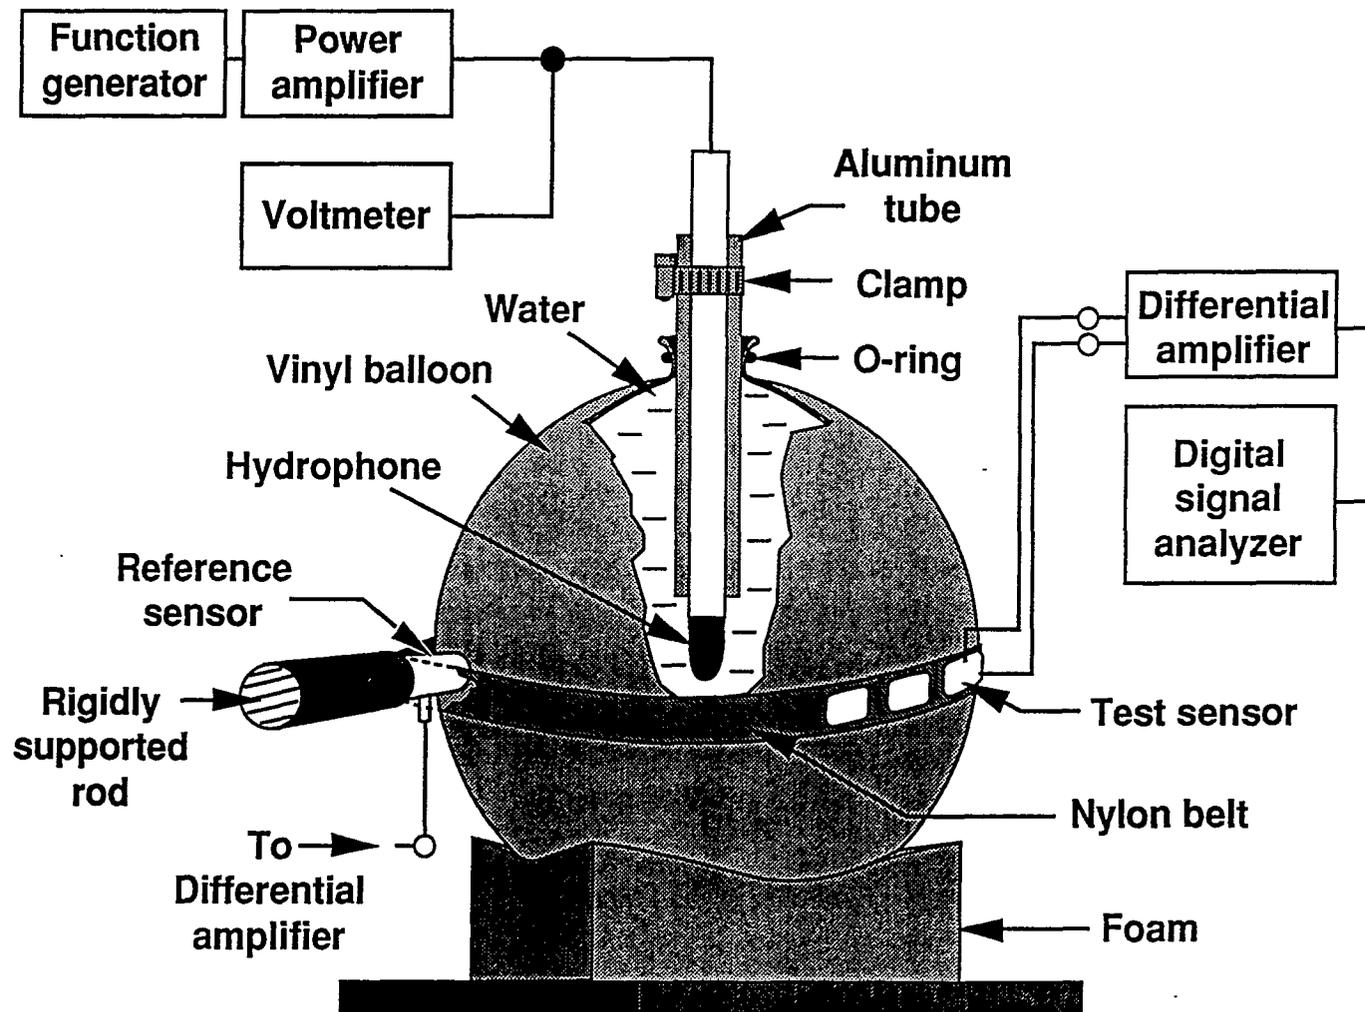


Figure 5.3 Simulation Experimental Set Up For The Balloon Test

W1 INST SPEC CH.B MAG INPUT MAIN Y: 133mV
Y: 150mV RMS LIN X: 16.00Hz
X: 0.00Hz + 200Hz LIN
SETUP W1* TOTAL : 309mV

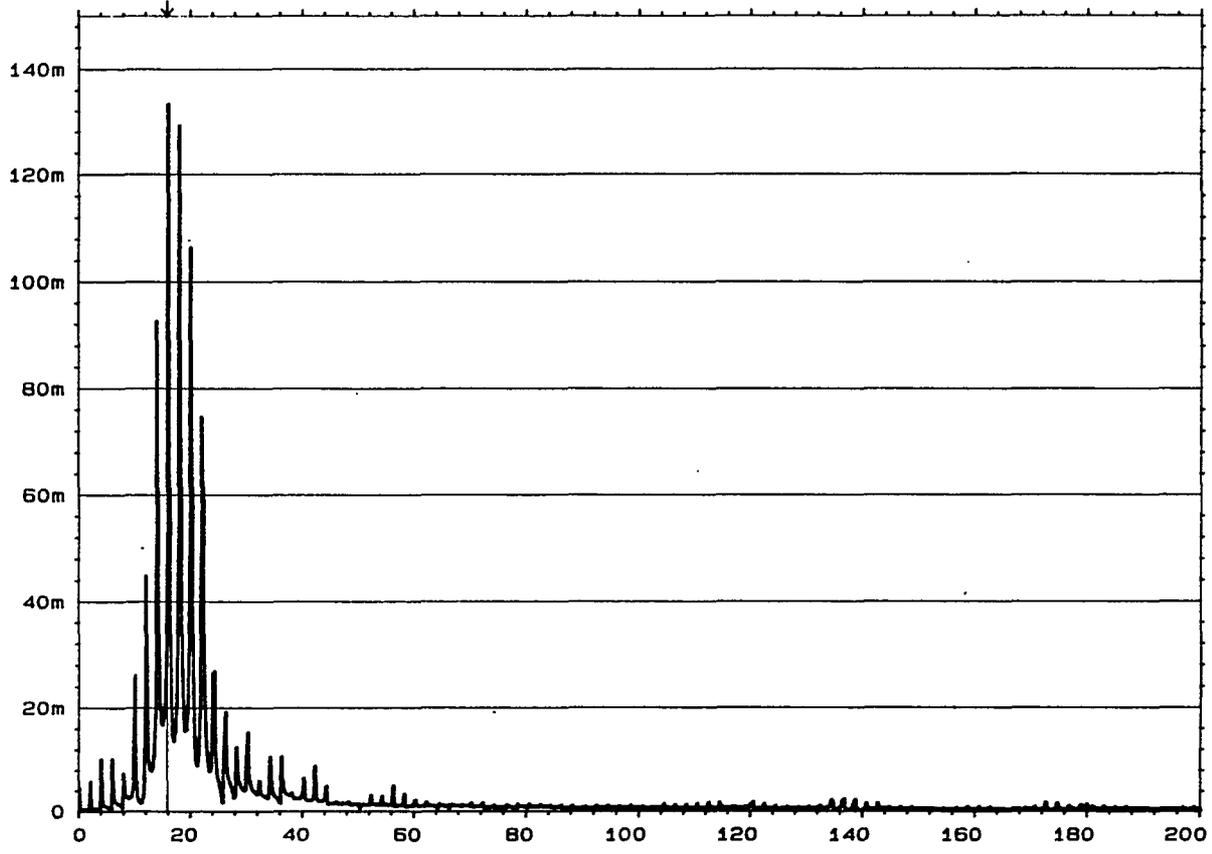


Figure 5.4 Spectrum of A Simulation Fetal Heart Signal in Balloon Test

to generate an additive random noise to simulate the acoustic noise encountered in actual situations.

The purpose of the balloon test was to verify the reliability of the current fetal heart rate monitoring system. Two tests have been conducted to show the reliability of the system in two different ways. The first test shows the consistency of the system with a fixed heart beat rate and the robustness of the system with a gradually increased noise level. The next test shows the consistency of the system with a varying heart beat rate under certain noise levels. The noise level was measured by the signal to noise ratio (SNR). The higher the SNR is, the lower the noise level is. Figure 5.5 and figure 5.6 show both test results.

From figure 5.5 and figure 5.6, it is clear that the monitor tracks the heart beat consistently at 120 BPM with SNR above 3 dB. The monitor still tracks the heart beat at a 120 BPM baseline with SNR below 3dB, but the artificial effects are increased due to large amplitude noise. In addition, the second test result also shows that the monitor tracks well from 90 BPM to 180 BPM with a noise level of both SNR 3dB and SNR 0 dB. However, the monitor does not track well if the noise level is increased for a SNR of -3 dB. In summary these tests showed that the new algorithm, hardware and software, performed well for these simulated conditions. These tests used the Teager energy operator for fetal heart tone detection.

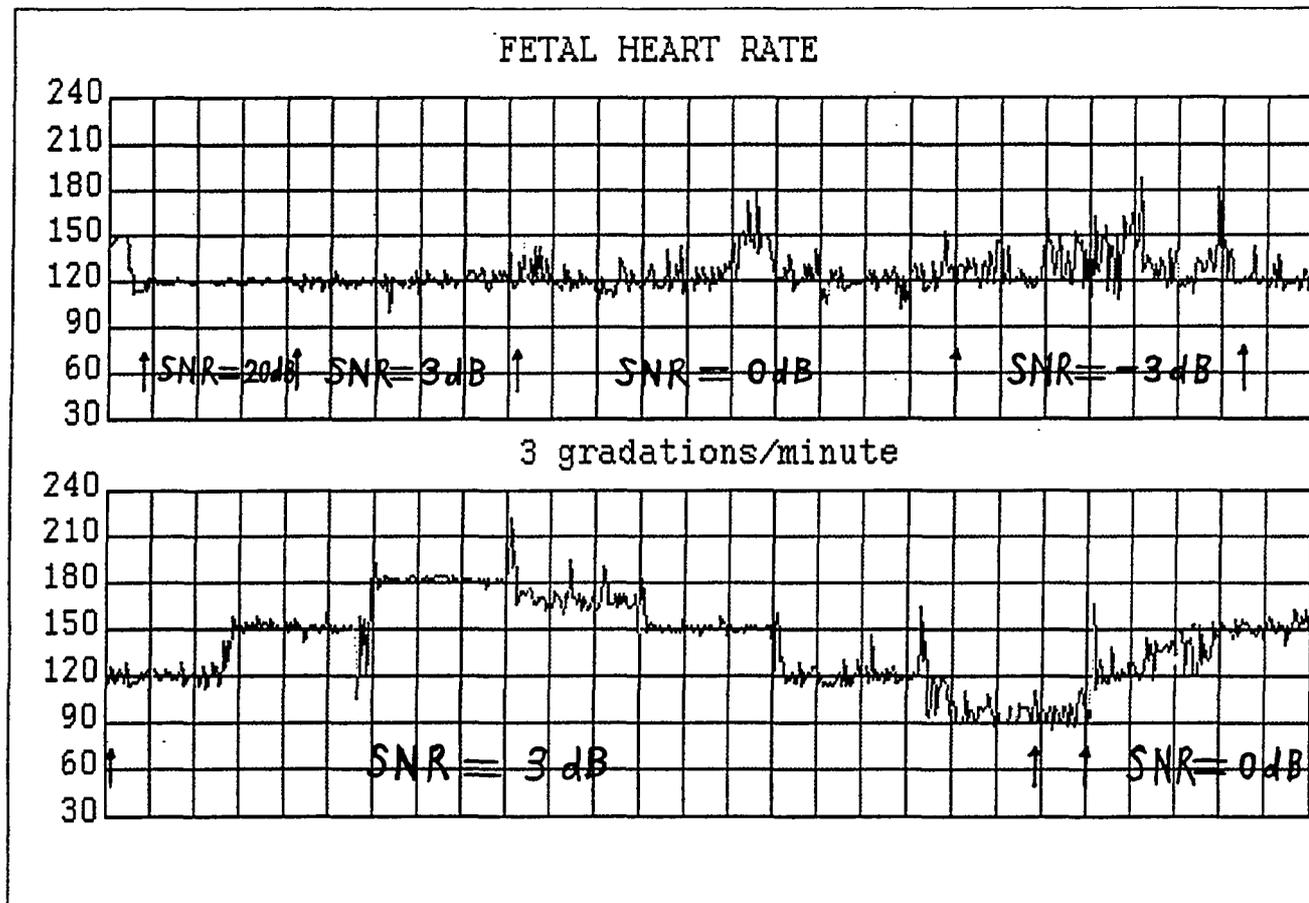


Figure 5.5 Balloon Test Results

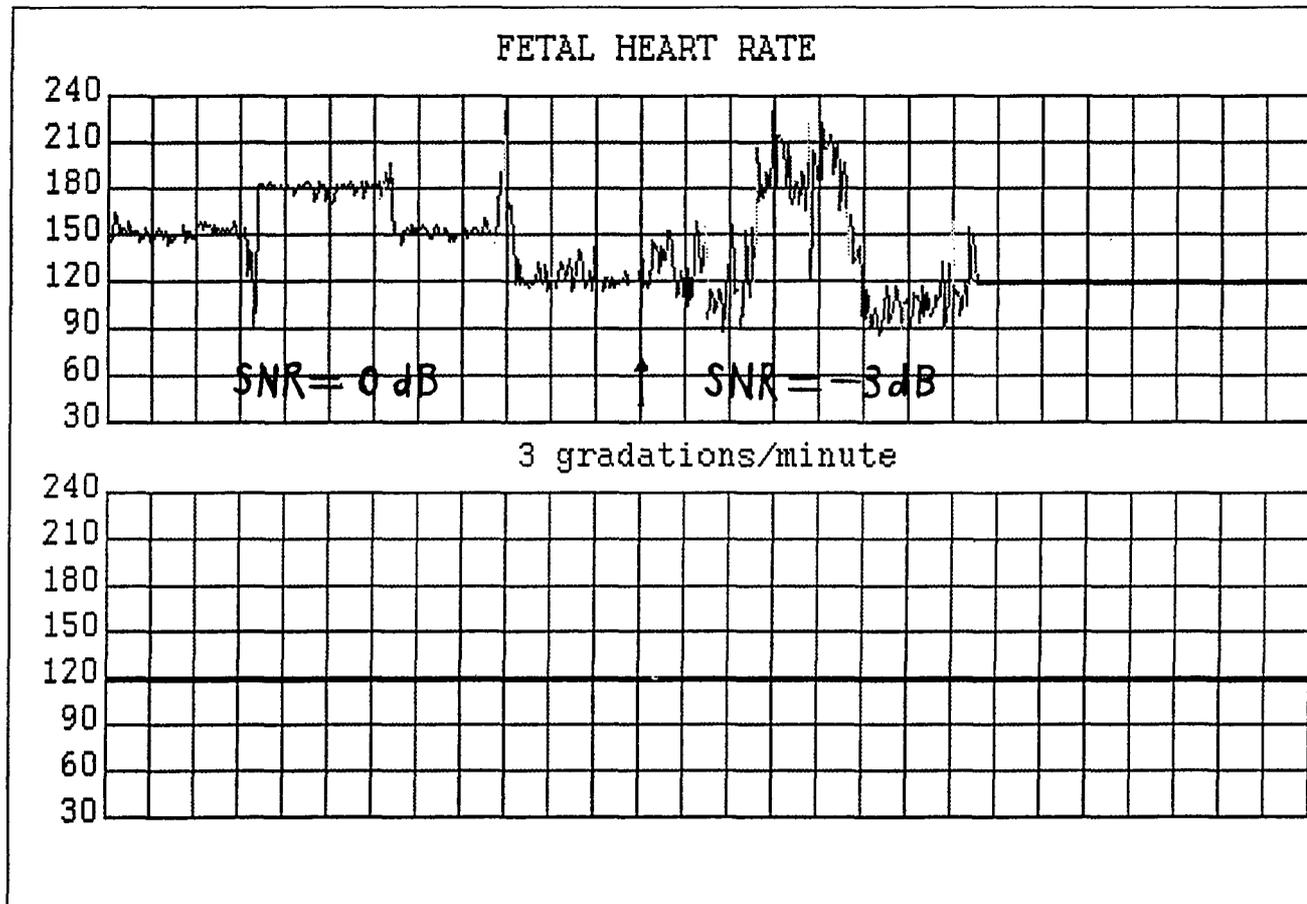


Figure 5.6 Balloon Test Results (Continued)

5.3 Clinical Testing

After experimentally showing that the monitor worked well in the laboratory, the monitor was then tested in a clinical setting. Three clinical tests were conducted at the EVMS Department of Fetal Maternal Medicine. Twelve patients participated in the clinical tests. They were high risk mothers who were being monitored for fetal well being by weekly ultrasound Non Stress Tests (NSTs).

5.3.1. Procedure

On arrival at the EVMS clinic patients were asked by the NST nurses whether they would be willing to participate in the acoustic monitor study. If they volunteered, they would take a regular ultrasonic NST before taking a twenty-minute acoustic NST. Both tests could not be carried out simultaneously since the sensor belt takes the same position as the ultrasonic sensors. The location of the fetal heart was then determined by means of an ultrasonic imaging device as per routine for the ultrasonic NST. A fetal stethoscope was next used to locate the point of the loudest heart tone in the area of the fetal heart indicated by the imaging device. The center sensor of the sensor array on the sensor belt, that is sensor four, was then positioned to the fetal heart tone point.

Once the sensor belt was positioned the acoustic FHR monitor system was powered up with concomitant initiation of the software. The electronics

support system was also powered up and the amplifier gain adjusted as necessary using the one channel oscilloscope screen of the monitor. The default operating mode was automatic mode, and the default digital filter had a pass band from 20 Hz to 50Hz. Once the FHR monitor was powered up, both acoustic signals from sensor four and fetal heart rate information were simultaneously recorded over the NST, a period of approximately 18-20 minutes.

Figures 5.7 through 5.11 show six-second acoustic fetal heart beat signals recorded from five patients under different operating conditions. Figures 5.12 through 5.16 show the average spectrums of these five acoustic fetal heart signals. These spectrums were averaged on ten consecutive frames of acoustic signals with a Hamming window having a length of 512 points. The spectrums, shown in figures 5.12 through 5.14, were based on filtered signals with a passband of 16 Hz to 50 Hz. Figure 5.15 shows the spectrum of a filtered signal with a passband of 20 Hz to 50 Hz. All the filtered signals mentioned so far are a combination of all seven sensor acoustic signals, and their corresponding heart rate charts are introduced in the next section. The figure 5.16 shows the average spectrum of a filtered acoustic signal only from sensor four with a passband 20 Hz to 50 Hz. However, its corresponding heart rate chart is not offered in the next section -- Ultrasound Comparative Study, since the ultrasonic record was not available at the time. Consequently, there is correlation among these average spectrums in a range of 20 Hz to 50 Hz.

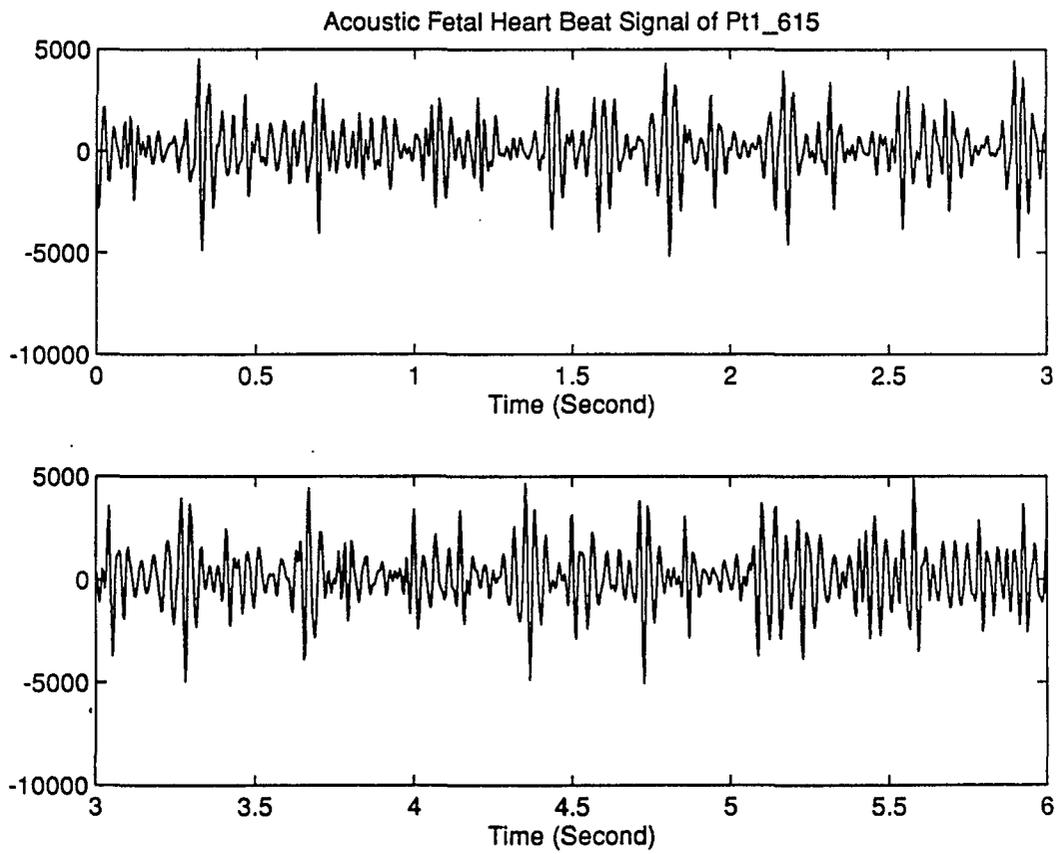


Figure 5.7 Six Second Acoustic Fetal Heart Beat Signal of Pt1_615

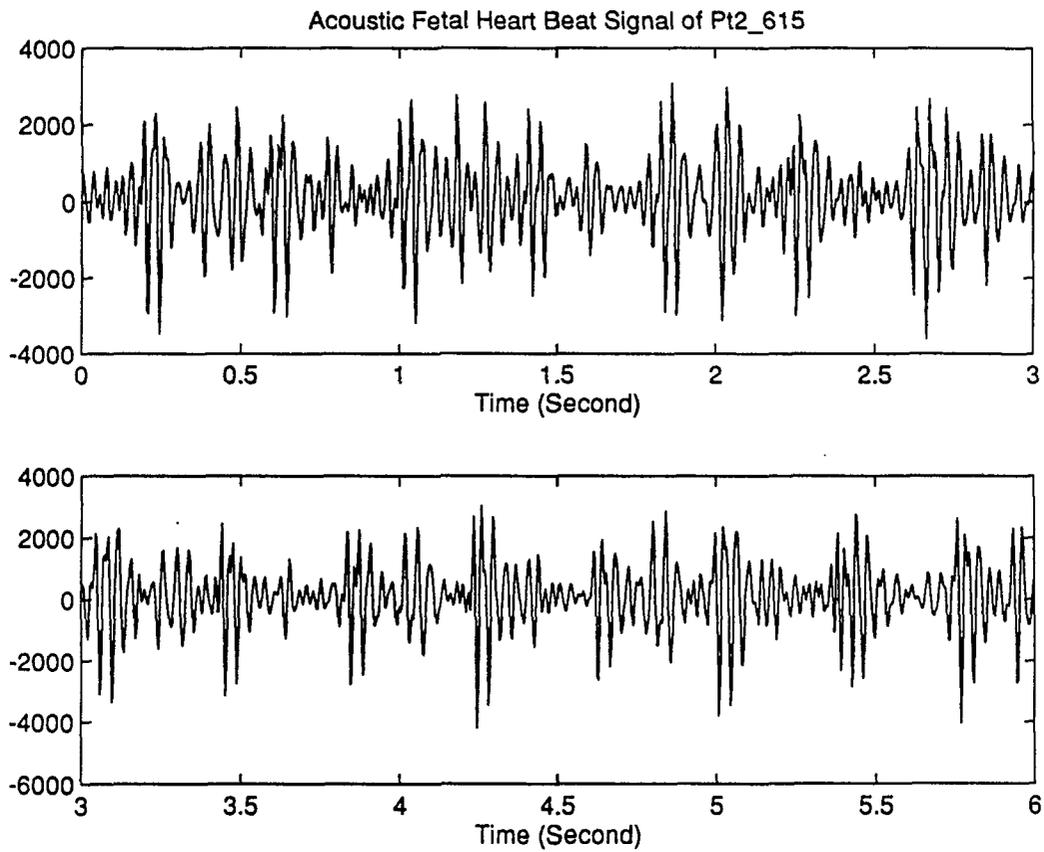


Figure 5.8 Six Second Acoustic Fetal Heart Beat Signal of Pt2_615

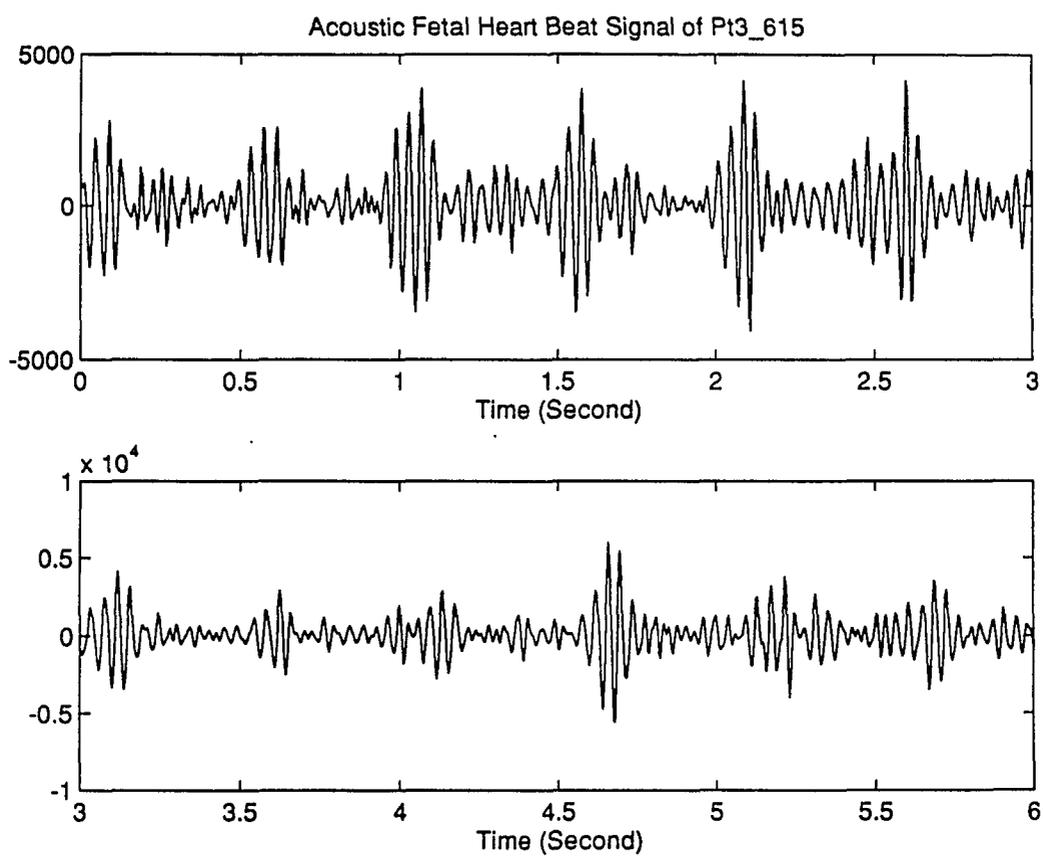


Figure 5.9 Six Second Acoustic Fetal Heart Beat Signal of Pt3_615

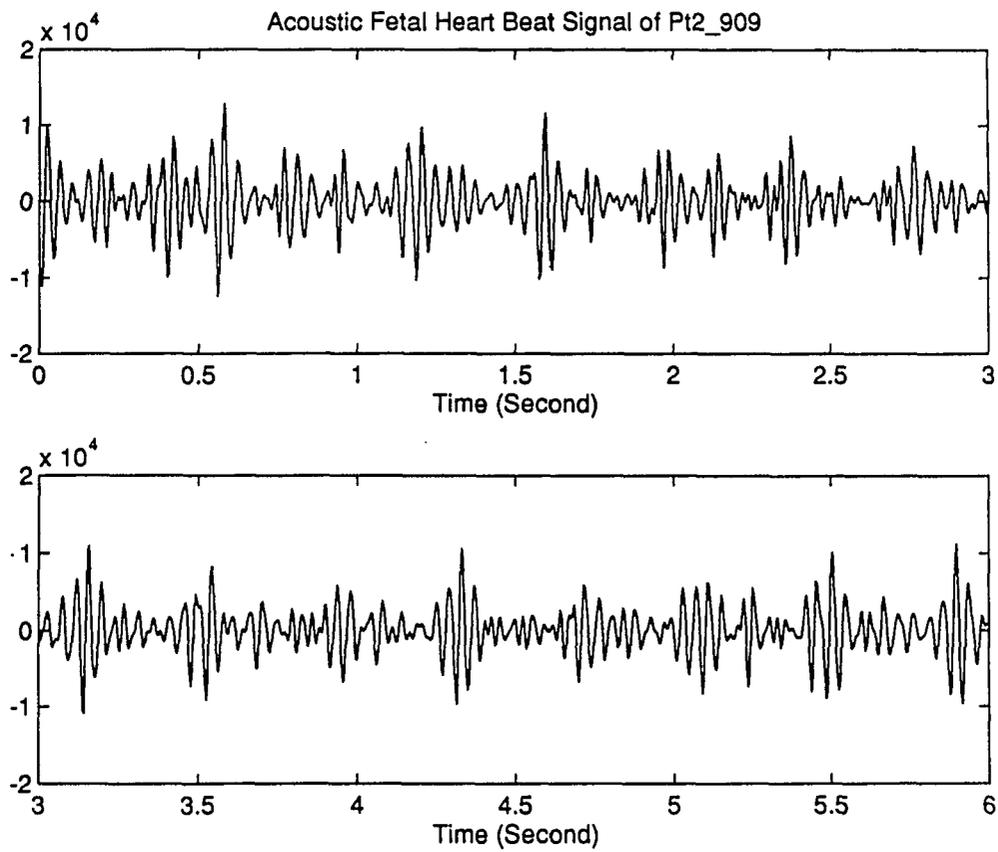


Figure 5.10 Six Second Acoustic Fetal Heart Beat Signal of Pt2_909

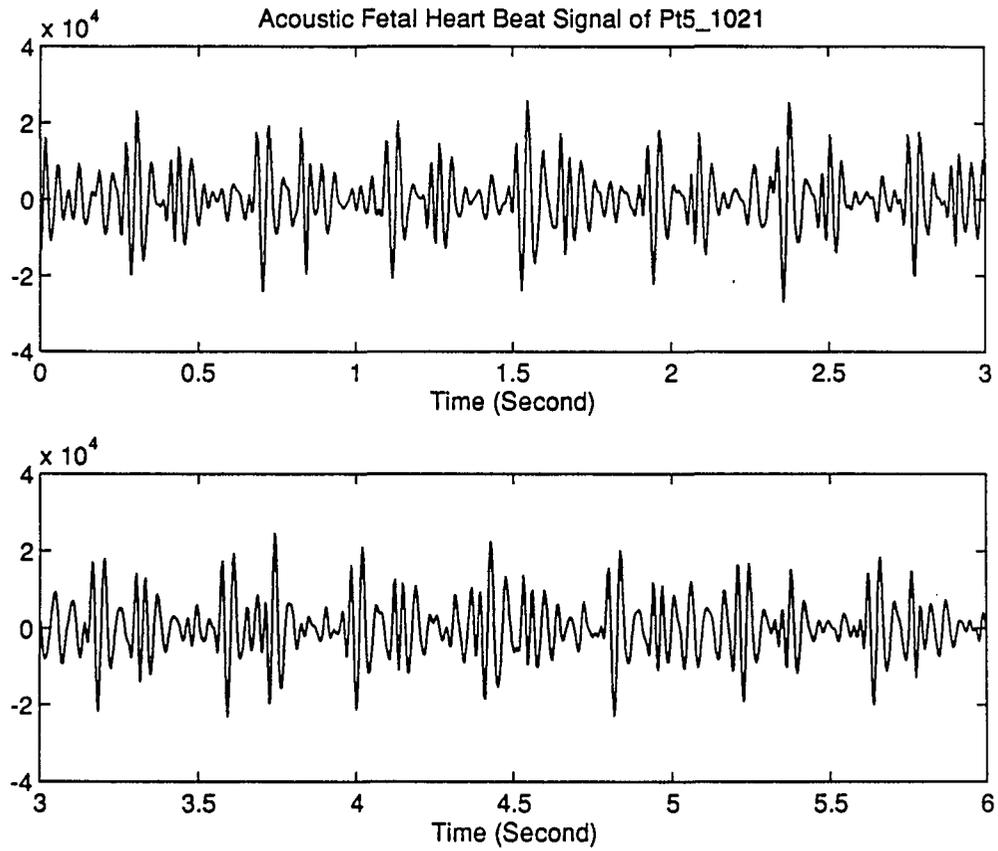


Figure 5.11 Six Second Acoustic Fetal Heart Beat Signal of Pt5_1021

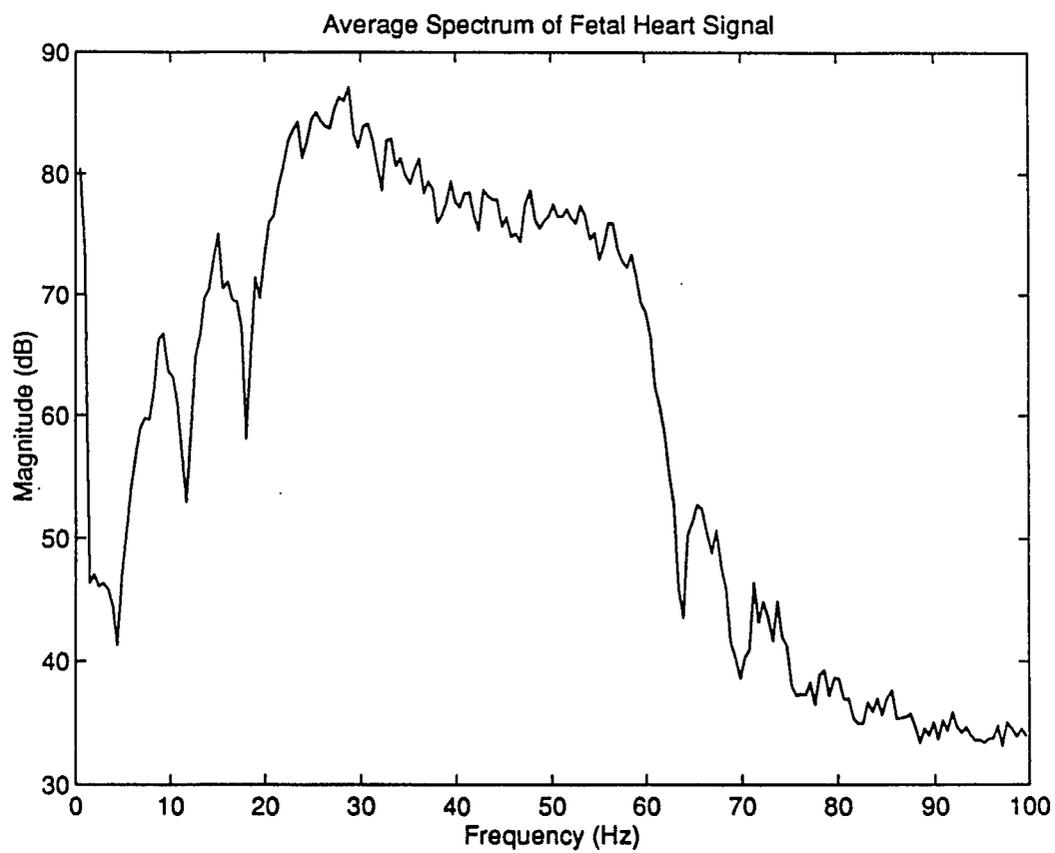


Figure 5.12 Average Spectrum of Fetal Heart Signal From Pt1_615

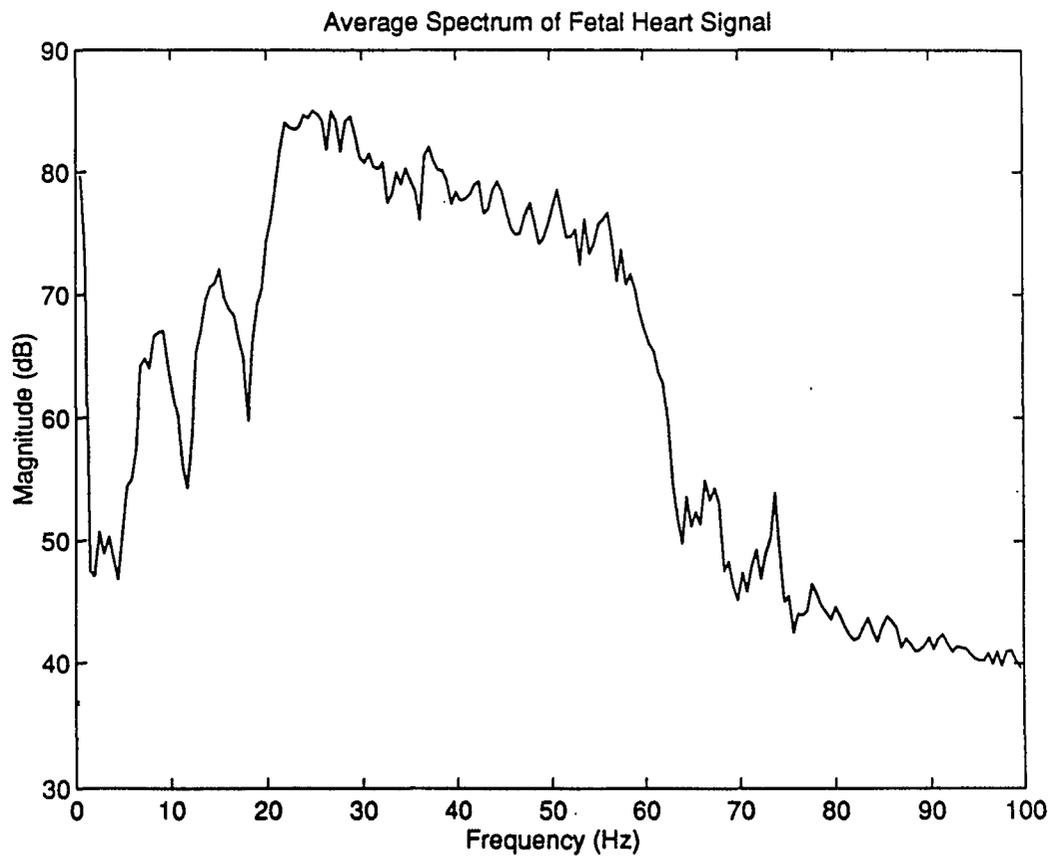


Figure 5.13 Average Spectrum of Fetal Heart Signal From Pt2_615

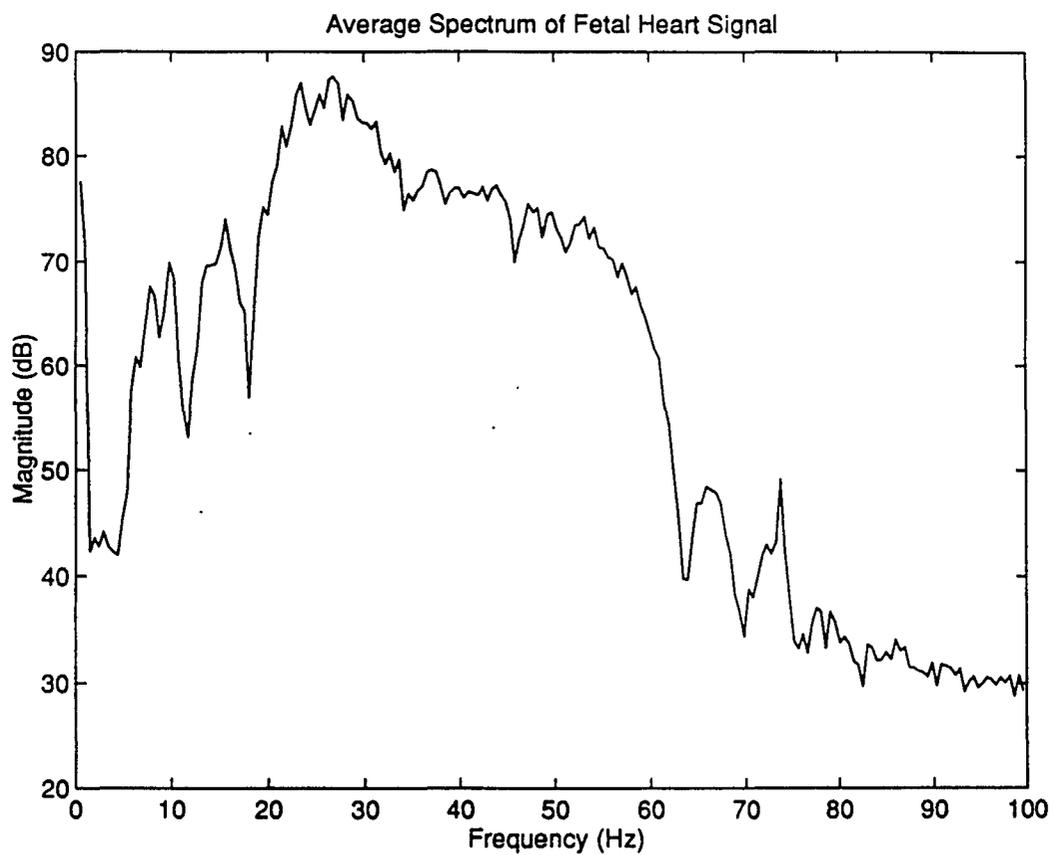


Figure 5.14 Average Spectrum of Fetal Heart Signal From Pt3_615

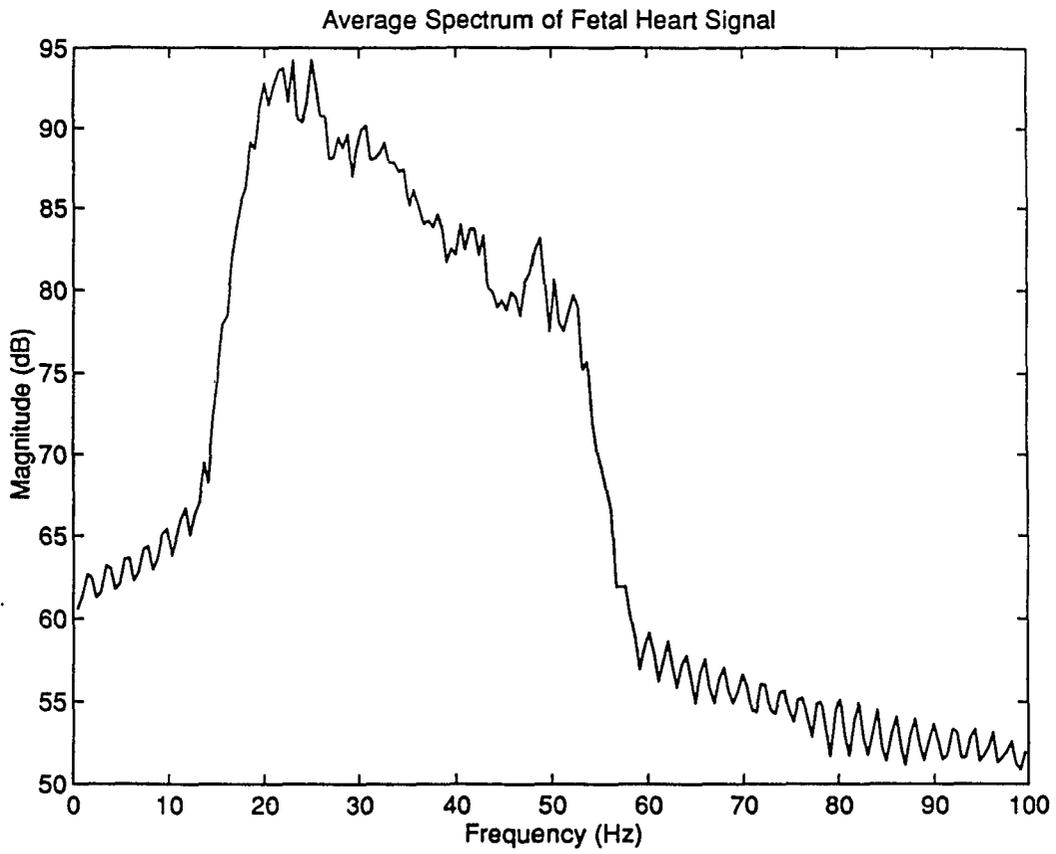


Figure 5.15 Average Spectrum of Fetal Heart Signal From Pt2_909

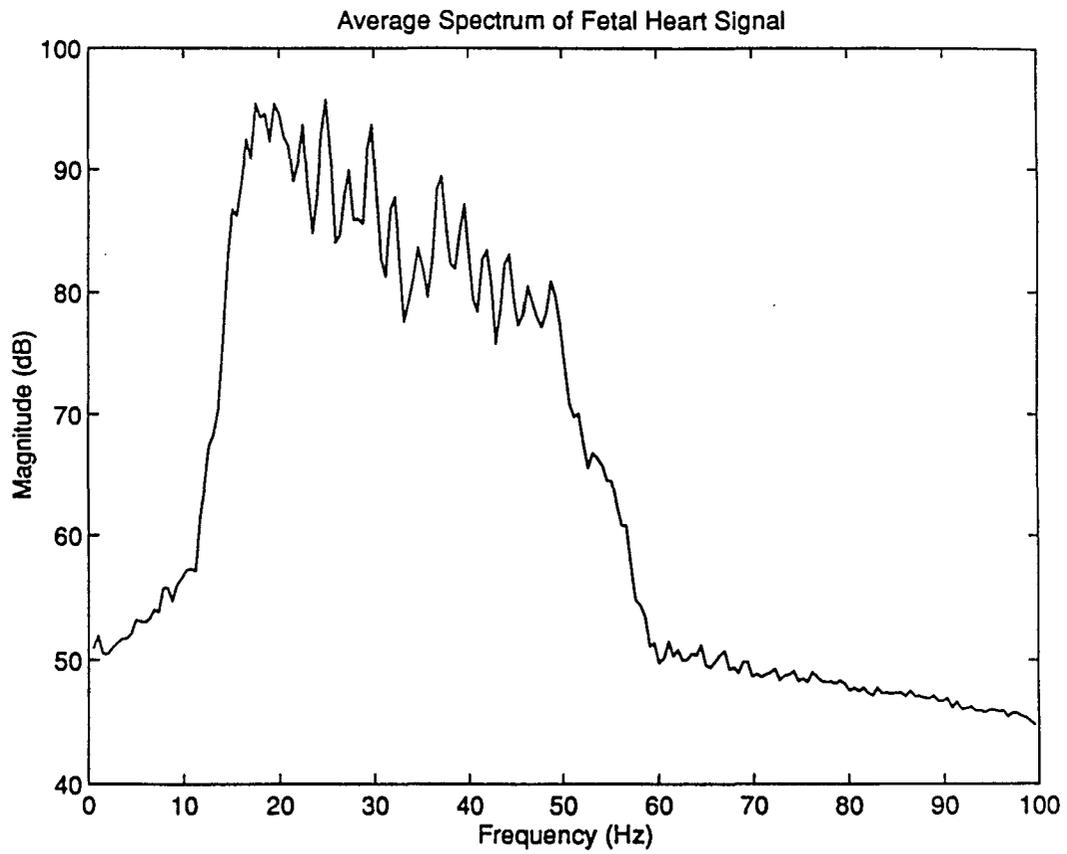


Figure 5.16 Average Spectrum of Fetal Heart Signal From Pt5_1021

5.3.2. Ultrasound Comparative Study

The comparative acoustic and ultrasonic heart rate charts for four patients are shown in figures 5.17 through 5.25. Due to the great width of the acoustic sensor belt, it was impossible to locate both the ultrasonic sensor and the acoustic sensor belt simultaneously. According to [12], fetuses always have their own distinctive heart rate patterns. During two consecutive NSTs, a fetus should have a similar heart rate patterns in terms of baseline and acceleration. A comparison was made for each patient between ultrasonic and acoustic heart rate records.

The purpose of this experimental study was to clinically test the reliability of the new portable acoustic FHR monitor for real time fetal heart tone detection and heart rate derivation. There is a good apparent correlation between the ultrasonic and acoustic heart rate records in three of four patients whose data is presented. (In the third record shown, figure 5.22, the general correlation is also good, but there are many "drop-outs," as indicated by the dotted lines in the tracings.) In addition, the ultrasonic NST of pt5_1021 is unavailable at this moment. A comparison study between the acoustic monitor and the "gold" standard -- ECG monitor is being conducted at Tarzania Regional Medical Center in Los Angeles, under the direction of Dr. Barry Schifrin, a leading authority in electronic fetal monitoring.

AA0

COROMETRICS MEDICAL SYSTEMS
© 1985

09911

4483AA0

COROMETRICS MEDICAL SYSTEMS
© 1985

09912

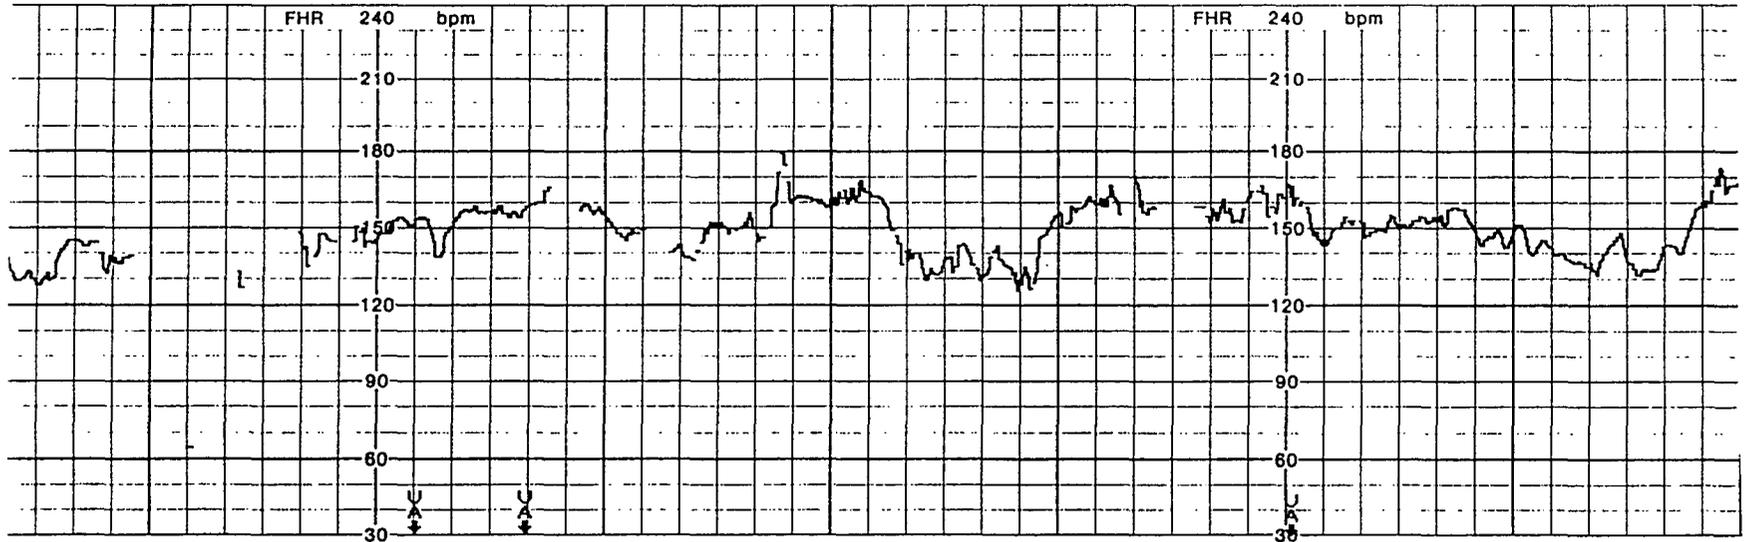
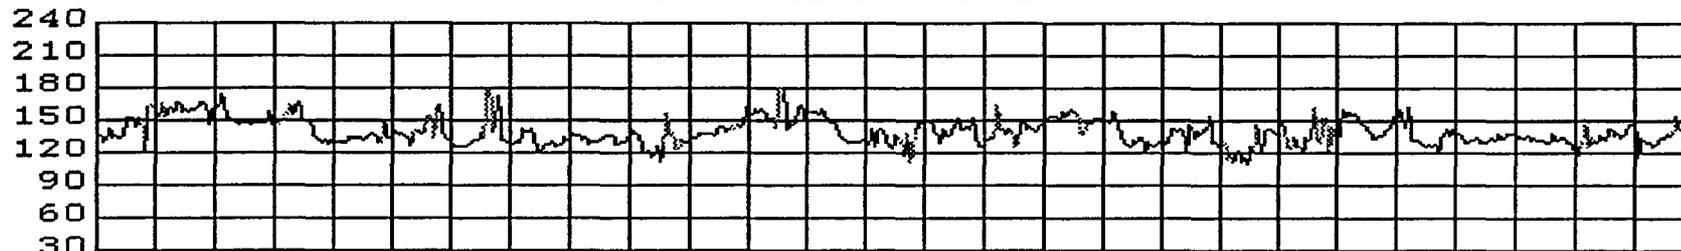


Figure 5.17 Ultrasonic NST of Pt1_615 (10 gradations / minute)

FETAL HEART RATE



.3 gradations/minute

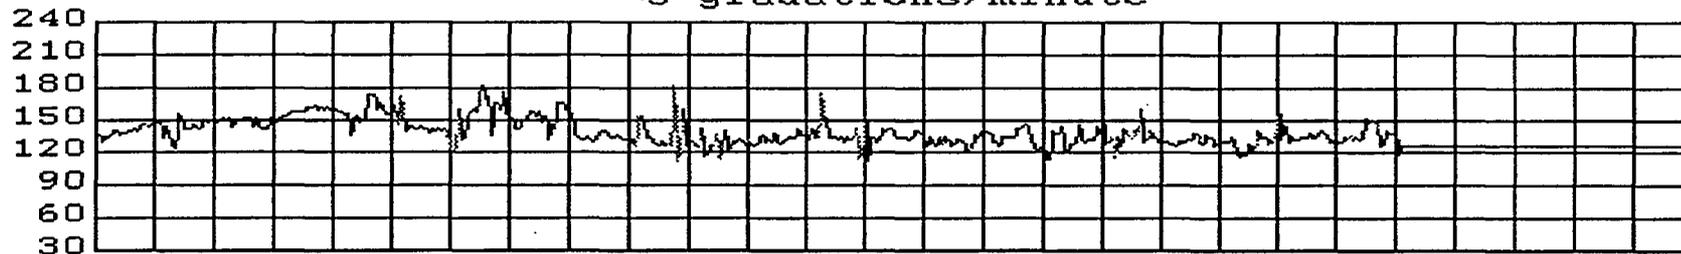


Figure 5.18 Acoustic NST of Pt1_615

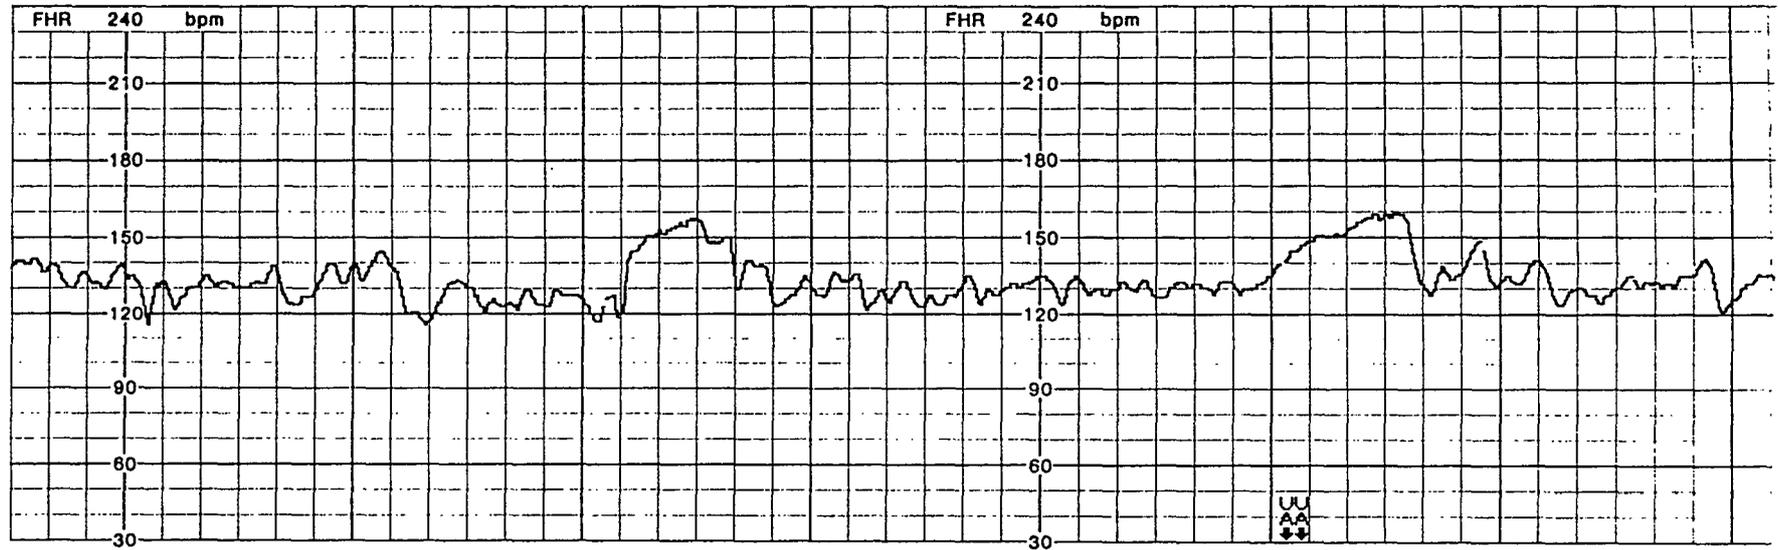


Figure 5.19 Ultrasonic NST of Pt2_615 (10 gradations / minute)

FETAL HEART RATE

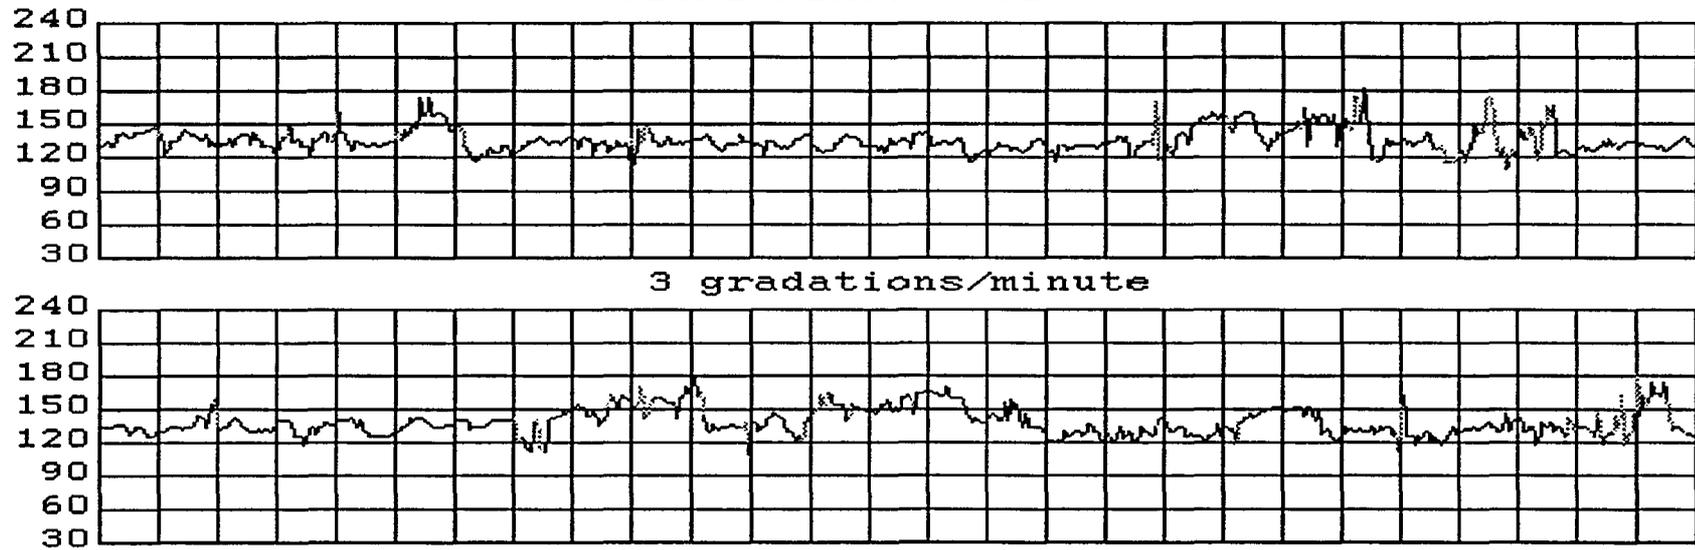


Figure 5.20 Acoustic NST of Pt2_615

4483AA0

 COROMETRICS MEDICAL SYSTEMS
© 1985

09930

4483AA0

 COROMETRICS MEDICAL SYSTEMS
© 1985

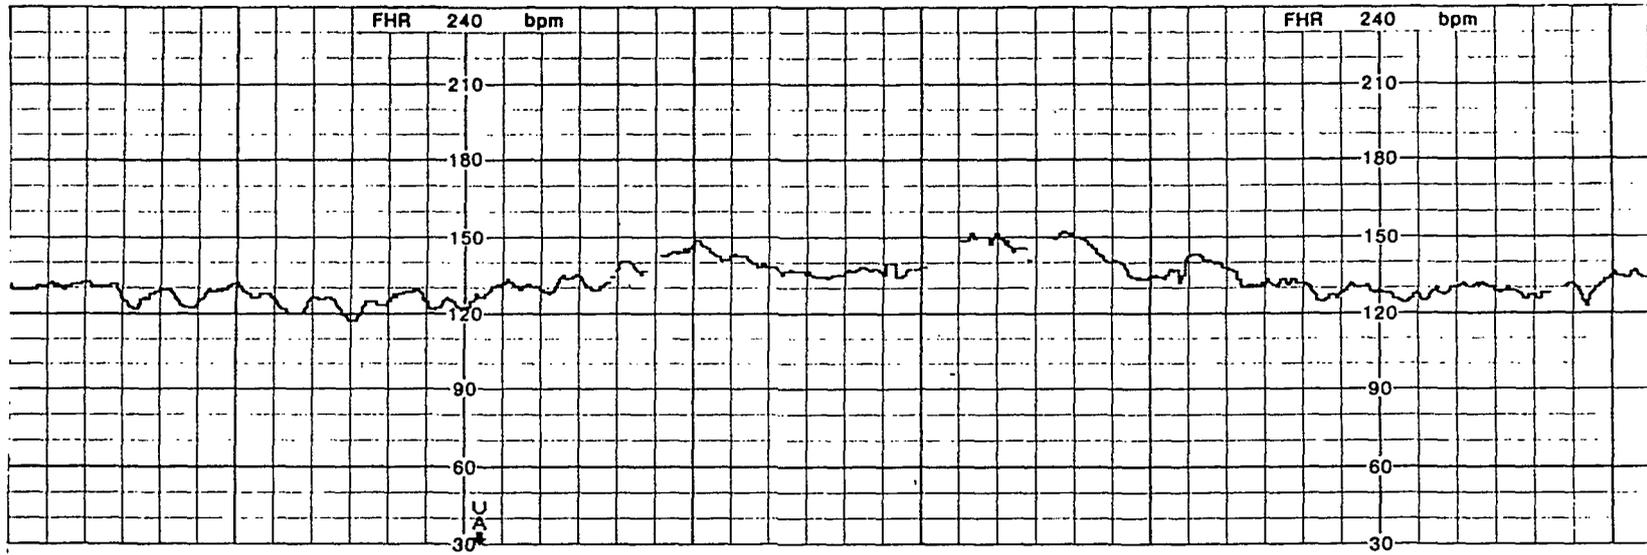


Figure 5.21 Ultrasonic NST of Pt3_615 (10 gradations / minute)

FETAL HEART RATE

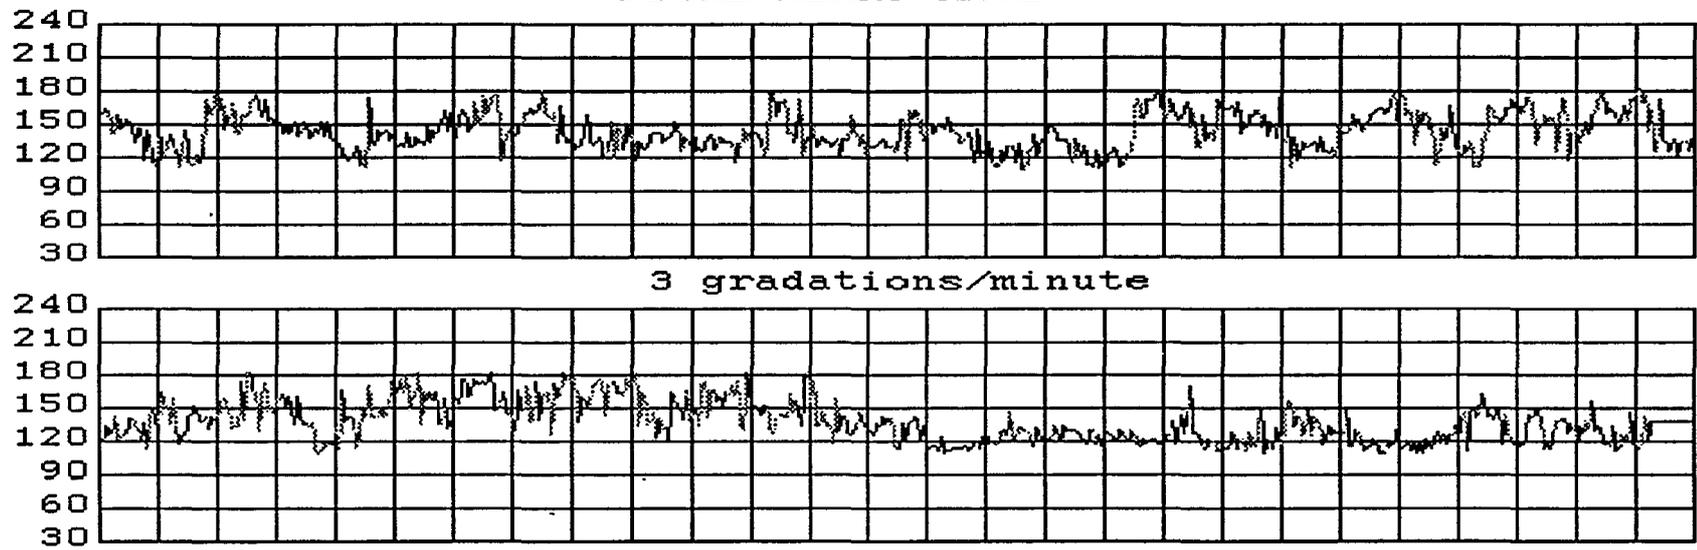


Figure 5.22 Acoustic NST of Pt3_615

COROMETRICS MEDICAL SYSTEMS, INC.
56553

4305AAO  COROMETRICS MEDICAL SYSTEMS, INC.
56554

4305AAO  COROMETRICS MEDICAL SYSTEMS, INC.
56555

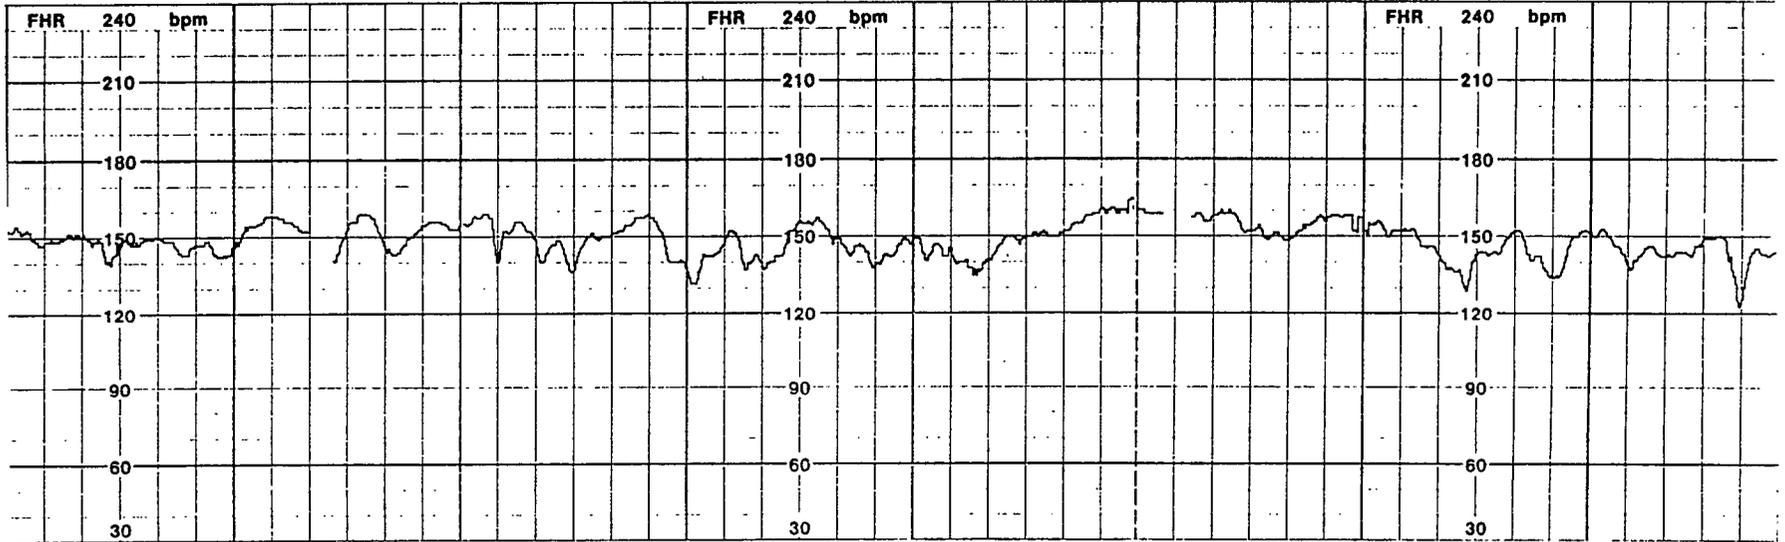
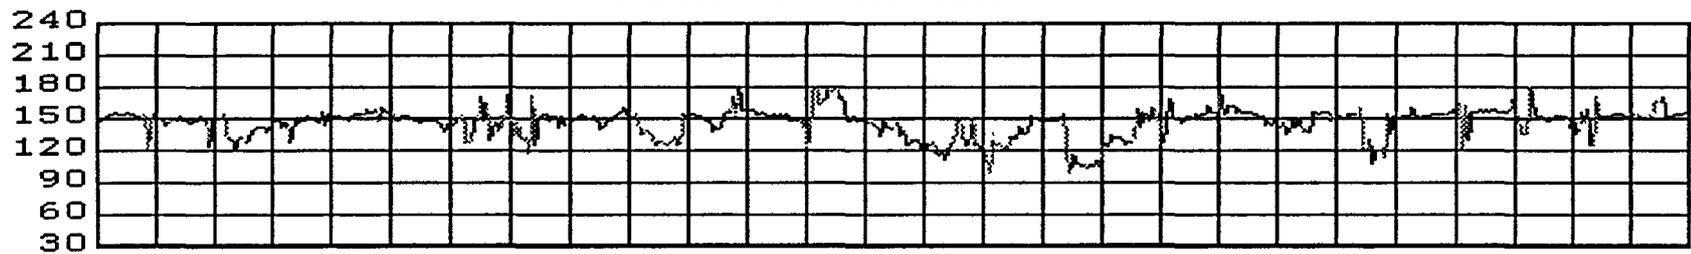


Figure 5.23 Ultrasonic NST of Pt2_909 (10 gradations / minute)

FETAL HEART RATE



3 gradations/minute

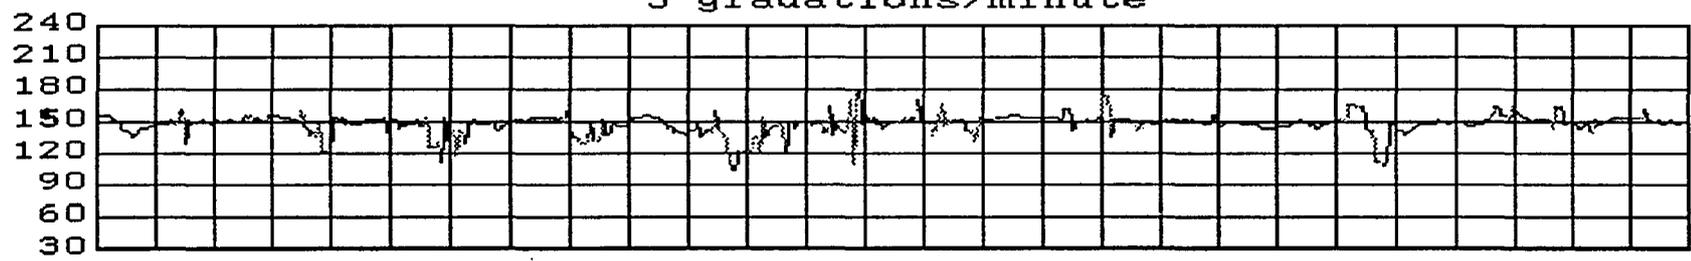
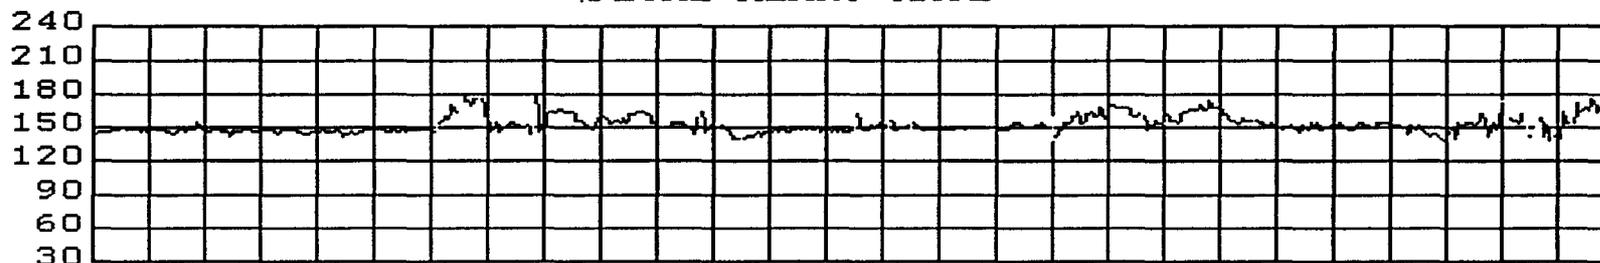


Figure 5.24 Acoustic NST of Pt2_909

FETAL HEART RATE



3 gradations/minute

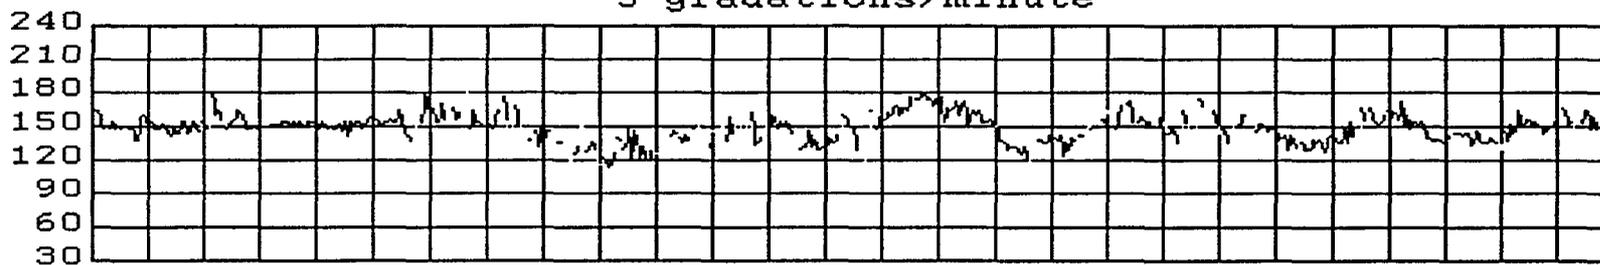


Figure 5.25 Acoustic NST of Pt5_1021

5.4. Conclusion

This research has implemented a portable working system for the real time detection of fetal heart tones from a noise contaminated acoustic signal with subsequent derivation of heart rate. Compared to a commercial ultrasonic monitoring unit, the current acoustic fetal heart rate monitoring algorithm has better performance than the previous linear predictor algorithm on the same previously recorded acoustic fetal heart signal. The acoustic FHR system also compares favorably with the ultrasonic monitor for heart rate generation in the context of the fetal nonstress test. Certain problems remain to be solved.

Algorithm Performance

1). Good apparent correlation in three out of four valid comparative studies, especially more correlation shown in the latest test than in the early tests, indicates a high degree of accuracy for detection of fetal heart tones by the Teager energy and autocorrelation combination algorithm.

2). The algorithm is highly time efficient requiring only one third of available processing time to implement a more complex algorithm. However, more time was consumed in data management, and the remaining time was also used to implement accompanying functions.

3). The FHR system has been implemented with a portable unit. The remaining processing time was utilized to implement accompanying functions

in software, such as acoustic fetal heart beat playback, to reduce the physical size of the system.

4). An automatic sensor scanning mode was successfully developed in the FHR monitor. The fetal heart rate detection performance under the automatic mode is better than that under the manual mode (single sensor detection mode). However, the reason for this result is unknown at this time.

5). Problems encountered:

The most significant problem encountered with the FHR system was the inability to detect the fetal heart rate accurately if a mother had a large abdomen with a lot of amniotic fluid. Since the fetal heart tones are much weaker in these two cases, the pre-processing gain is not large enough to keep the SNR at the minimum required level (at least zero dB as described above). Other problems encountered include sensitivity of the FHR monitor to maternal heart tones and sensor belt movement.

Future Development Potential of The Fetal Heart Rate Monitor

The portable acoustic fetal heart rate monitor has several potential development areas:

1). A neural network could be well trained with a large training space of good fetal heart tones obtained by the current monitor. The neural network could also be trained to ignore any type of known noise, such as maternal heart tones or sensor belt movement. The fetal heart tone detection could be done by the neural network. In addition, the parallel redundance algorithm could also be

used with the neural network detection algorithm by using more than one TMS320C31 DSP processors in the real time system.

2). The hardware system of the FHR monitor could be optimized by designing a more compact instrument instead of a PC-based system. The number of A/D channels could be increased to seven in order to reduce the complexity of the front-end electronic support system and the software implementation.

3). A fetal heart rate pattern recognition system could be further developed to help mothers prevent fetal distress at home by themselves. This would enable greater surveillance of high risk pregnancies and provide convenience for the mothers.

A portable non-invasive acoustic FHR monitoring system has been developed and shown to work. A commercial medical instrument is becoming more feasible than ever. Home use of the acoustic FHR monitoring system could be a reality in the near future. Continued effort should bring this idea to a useful reality sooner.

BIBLIOGRAPHY

1. Pretlow, R. A. (1991), "Signal processing methodologies for an acoustic fetal heart rate monitor," Master Thesis, Old Dominion University.
2. Ifeachor, E. C. and Jervis, B. W., (1993) Digital Signal Processing. A practical approach, (Addison-Wesley Publishing Company, Wokingham, England).
3. Zuckerwar, A. J., Pretlow, R. A., Stoughton, J.W. and Baker, D.A., (1993), "Development of a Piezopolymer Pressure Sensor for a Portable Fetal Heart Rate Monitor," IEEE Transactions On Biomedical Engineering, vol. 40, No. 9, pp. 963-969
4. Kaiser, J. F. (1990), "On a simple algorithm to calculate the 'energy of a signal'," ICASSP-90. pp. 381-384, Albuquerque, New Mexico.
5. Talbert, D. G., Davies, W. L., Johnson, F., Abraham, N., Colley, N., Southall, D. P., (1986), "Wide bandwidth fetal phonography using a sensor matched to the compliance of the mother's abdominal wall," IEEE Transactions On Biomedical Engineering, vol. BME-33, no. 2.
6. Hewlett-Packard GmbH, (1970), Operating and Service Manual for Model 8020A Cardiotacograph, Hewlett-Packard GmbH, Boblingen, Germany.
7. Jenssen, H., (1980), "Fetal systolic time intervals after paracervical block during labor," Acta Obst. Gynecol. Scand., vol. 59, pp. 115-121.
8. Kobayashi, K. and Yasuda, T., (1981), "An application of PVDF-film to medical transducers," Ferroelectrics, vol. 32, pp. 181-184.
9. Nagel, J., (1986), "New diagnostic and technical aspects of fetal phonocardiography," Eur. J. Obst. Gynecol. Reprod. Biol., vol. 23, pp. 295-303.
10. Talbert, D. G., Dewhurst, J. and Southall, D. P., (1984), "New transducer for detecting fetal heart sounds: Use of compliance matching for maximum sound transfer," Lancet, pp. 426-427.

11. Manning, T., Flores, C., Hon, E. and Davidson, E., (1985), "Limitation of autocorrelation in fetal heart rate monitoring," American Journal of Obstetrics and Gynecology, vol. 153, pp. 685-692.
12. Schifrin, B. S., (1990), Exercises In Fetal Monitoring, (Mosby-Year Book, Inc., St. Louis).

APPENDIX

USER'S GUIDE

The acoustic fetal heart rate monitor is a research project sponsored by NASA which will lead to a valuable medical instrument. Hardware of the monitor includes an IBM compatible PC, a TMS320C31 based ELF DSP platform, a front-end electronics box and an acoustic sensor belt.

After connecting all the hardware correctly, as shown in figure I, turn on the power, and start the monitor by typing the following commands under a DOS prompt :

```
☞ c:>cd fhm\pc←
```

```
☞ c:\fhm\pc>fhm←
```

A command menu is then displayed on the PC screen. There are six primary options listed on the menu. They are:

C -- Corometrics Unit Only

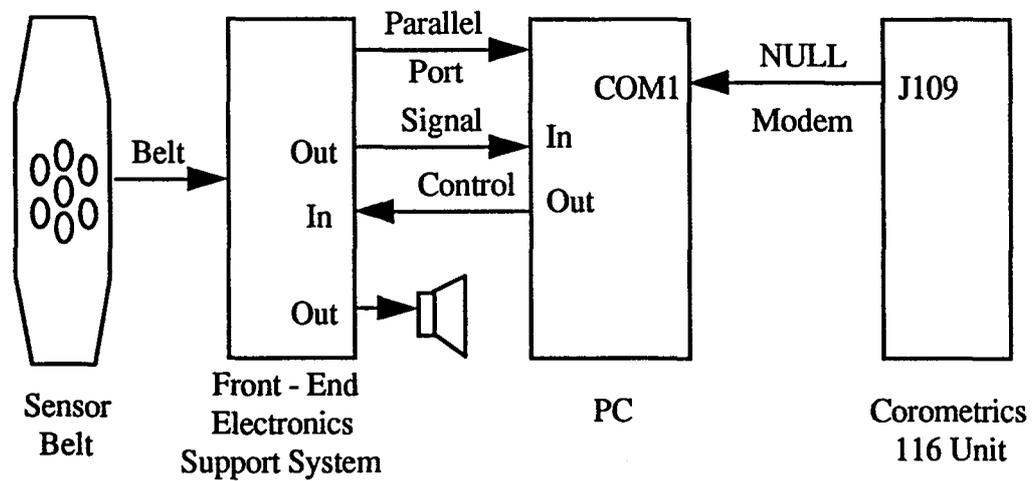
A -- Acoustic Unit Only

B -- Both Units

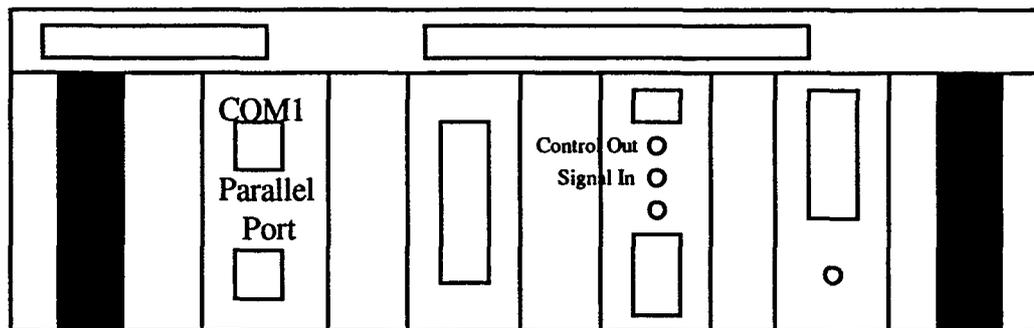
O -- One Channel Oscilloscope

S -- Seven Channel Oscilloscope

Q-- Quit



(a) System Connection



(b) PC I/O Configuration

Figure I System Hardware Connection

Each option is selected by pressing the keyboard letter listed in front of the option. For example, to select "Acoustic Unit Only" option, press "A" with the keyboard. To quit the monitoring, press "Q". Note that this will work with both upper and lower case. We next briefly describe each of these options.

Corometrics Unit Only

This option is used when fetal heart monitoring is only based on the Corometrics 116 unit. The PC serial port (COM1) has to be connected to either J109 or J110 of the Corometrics unit. The serial communication speed of the Corometrics unit has to be set at 2400 baud.

The only information acquired by the PC is the fetal heart rate information. After pressing "C" at the command menu, a strip chart type recording is drawn on the PC screen. The rate chart has a range from 30 beats per minute (BPM) to 240 BPM, and has two strips, each depicting up to 9 minutes of rate information. The rate information is also displayed as a numerical value in the middle of the bottom screen every second. If a recording is longer than 18 minutes, the tracing is redrawn on the oldest part of the chart.

In addition, there are three of four common suboptions for the rate chart.

Save

Pause

Quit

To select one of them, simply press the capital letter, which is the first letter of the desired option. For example, "S" is for the "Save" option. The "Save" option is used to begin saving the rate information to a file. The "Pause" option here is to stop saving the rate information to the file. The "Quit" option quits the current option and returns the user to the command menu mentioned above.

Acoustic Unit Only

The "Acoustic Unit Only" option is the primary mode used with the monitor. This option has two operating modes and two graphic modes. The two graphic modes are rate chart mode and one-channel oscilloscope mode. The two operating modes are automatic mode and manual mode. In the automatic mode, which is the default operating mode, the monitor scans all seven acoustic sensors and selects the best detected result. In addition, the monitor automatically sets the optimal A/D gain level. In the manual mode, the user can set the monitor at the desired AD gain level, and monitor the desired sensor. In this option, both a rate chart is presented (as described for the "Corometrics Unit Only" option), and also a one-channel oscilloscope is available to view acoustic fetal heart tones of the selected sensor. Note that in the automatic mode sensor four is always displayed with the oscilloscope option, even though the indicated sensor numbers may vary.

On entering the "Acoustic Unit Only" option, patient's initial has to be given to generate the proper file name for the NST record. The automatic

operating mode is the default operating mode of the system. However, the operating mode can be changed later (as described in the "One Channel Oscilloscope" section). If the operating mode is manual, the scanned sensor has to be selected from one to seven, and the A/D gain level has to be chosen from sixteen levels. The first zero level is ± 2.80 volts, and there is a 1.5 dB attenuation between two consecutive level. The final fifteenth level is ± 0.21 volts. Typically choose sensor 4 (center sensor) and gain level 8. If the system is in the automatic operating mode, the default starting scanned sensor is sensor four, and the default starting gain level is the zero level (± 2.80 volts).

After all settings are made, the rate chart is the same as the one mentioned in the "Corometrics Unit Only" option first displayed. Note that there are four options available, as indicated by four words typed at the bottom of the display:

Save

Pause

Oscil

Quit

Each option is invoked by typing the first letter of the corresponding word. Save, Pause, and Quit have the same functions as mentioned for the Corometrics unit, with the exception that the Save option causes two files to be saved. Note that the default condition is for the Save mode to be entered automatically. This will be described more below. In addition, there is an

additional Oscilloscope option (Oscil) which can be used to view acoustic data, and also to modify the recording conditions. To switch to the one-channel oscilloscope, press "O" on the keyboard. As for the rate chart mode, there are several suboptions available in the one-channel oscilloscope mode, including changing the operating mode (Manual versus Automatic). These options are presented in detail in the following "One Channel Oscilloscope" section.

One Channel Oscilloscope

The "One Channel Oscilloscope" is used to view acoustic fetal heart tones of the selected sensor under the selected A/D gain level. Since this option is used to view the acoustic fetal heart tones, this option is available when fetal heart monitoring is based on the acoustic unit. That is, when either the "Acoustic Unit Only" option or the "Both Units" option is selected. However, this option is also available at the command menu. This option has two operating modes and two viewing modes. The two operating modes are the same as mentioned in the "Acoustic Unit Only" section. The two viewing modes are the static viewing mode and the dynamic viewing mode. The static viewing mode "freezes" two seconds of the acoustic fetal heart signal (by hitting the "P" key). The dynamic viewing mode continuously displays two seconds of acoustic fetal heart signal, with a "flow-mode" update every half second. Note that new information enters from the right and flows to the left.

The "One Channel Oscilloscope" option has three suboptions, including changing the operating mode, changing the pass band of a digital filter, changing the viewing mode and quitting. If this option is chosen from the main command menu, the operating mode ("M" or "A") has to be first chosen (as described in "Acoustic Unit Only" section). The default viewing mode is the dynamic mode.

There are several options available from the oscilloscope display, as listed below. (However, due to screen size limitations only prompts are listed for Pause, Continue, and Quit on the viewing screen.). The options are

- "P" -- Pause , that is freeze the display;
- "C" -- Continue, switch back to the dynamic mode;
- "Q" -- quit this option and switch back to either the command menu or the rate chart mode;
- "M" -- switch to the manual operating mode;
- "A" -- switch to the automatic operating mode;
- "S" -- change the pass band of the digital filter.

Note there are five filter options available:

- 1 9-50 Hz
- 2 16-50Hz
- 3 20-50Hz
- 4 25-50 Hz
- 5 filter off

Typically (and also the default) is 20-50 Hz. Each filter is an FIR linear phase 124th order digital filter.

If "M" is pressed to invoke the manual mode, there is another active option keyboard list:

"F" -- increase the channel number;

"B" -- decrease the channel number;

"T" -- increase the A/D gain level;

"D" -- decrease the A/D gain level.

These options can be used to view in more detail the operation of the various sensors.

Both Units

In this option, the monitor combines all the suboptions from the "Acoustic Unit Only" and the "Corometrics Unit Only". The rate chart here simultaneously displays both fetal heart rates from the acoustic unit and the Corometrics unit. The fetal heart rates from the acoustic unit are displayed in the top strip, and those from the Corometrics unit are in the bottom strip. Each strip is 9 minutes long. If the recording is longer than 9 minutes, the tracing is redrawn on the oldest part of the chart. The numerical heart rate is the one from the acoustic unit.

Seven Channel Oscilloscope Mode

This option creates a "flow-mode" oscilloscope display of all seven sensor signals. The gain level is automatically adjusted. The bottom trace is an average of all seven sensor signals. Thus channel seven is not explicitly shown. The only options available are Pause, (to freeze the display), Continue (to return to the dynamic flow-mode), and Quit, to return to the main menu.

Notes on file saving

In the Acoustic Unit only, or the Both units option, as mentioned above, file saving options are available. These are very important for record keeping to evaluate and improve the monitor. Whenever file saving is invoked with the acoustic unit in use, two files will be saved. The user will enter only the patients' initial with three or four letters. A file name is automatically generated by combining patients' initials and the NST date (i.e., The record of C.A. Norton's NST on 11/12 will have a name "CAN1112" with a proper extension). The extension is also automatically generated, as mentioned below to differentiate the two files.

**.rat*

This is an ASCII file which contains a table of computed heart rates, time information, status (0 is for reliable decisions, 1 for unreliable), and selected sensor number. If both units are used (Corometrics plus acoustic

unit), then two heart rates are listed (acoustic unit is Rate1 and Corometrics is Rate2.)

**.wav*

This is a binary file containing samples of the acoustic heart beat signal from the selected sensor. The sampling rate is 250 Hz. This file also has a header used for later processing. These files will be saved in the *C:\FHM\PC* directory. The files must each be uniquely named to prevent overwriting of old files. The files can be copied to floppies and removed from the hard disk to prevent overfilling the capacity of the disk.

Notes on Print Rate Charts

For the current version of the software, there is no direct approach to print the fetal heart chart. A PCX file, containing the fetal heart rate chart image, has to be first obtained, and a graphic editor, such as Microsoft Power Point or a word processor, is needed to print the fetal heart chart on the Laser printer.

To get the PCX file, take the following steps:

```
c:>cd \fhm\show←
```

```
c:\fhm\show>pcxgrab /I←
```

```
c:\fhm\show> display←
```

The testing date and the patient's initial are requested by the program. A heart rate chart is then displayed on the screen. However, the color scheme

in this display is not very good for hardcopy on a black and white laser printer. Therefore, a second heart rate chart, identical except for the trace color, is displayed after the "Esc" key is pressed. To grab the current display into a PCX file, press the keys "Alt" and "G" simultaneously. Two consecutive beeps will be heard to indicate that the current display is grabbed to a file with a name "grab_xx.pcx." Once the PCX file is created, we can print the display using Microsoft Power Point. (This could also be done using Word Perfect 6.0 or Hijack software.) Here we only give basic directions if using Power Point.

On entering MS Power Point, use a mouse to click the "Insert" and choose "Picture." Once the PCX file is inserted into the Presentation, highlight the image box by pointing the mouse to the image and clicking the mouse. Eight small black buttons are shown around the image box. Drag these buttons to resize the image. To print the image, choose "Print" from the "File" menu.