ESTIMATION AND PROCESSING OF FETAL HEART RATE FROM PHONOCARDIOGRAPHIC SIGNALS

Hany Elias Milad Bassil B.Sc.

Thesis submitted for the degree of

Doctor of Philosophy

to the Faculty of Science,

University of Edinburgh.

1991

ABSTRACT

With the advent of new fetal phonocardiographic transducers affording better signal qualities than hitherto possible, this non-invasive monitoring technique is explored with a view to long-term antenatal fetal heart rate (FHR) monitoring.

The analysis of fetal heart phonocardiographic (phono) signals in spectral composition, stationarity and type of interference noise sets the basis for the development of FHR estimation algorithms. For repeatability purposes, the fetal heart signals are simulated at varying signal to noise ratios (SNR). This enables the controlled generation of simulated fetal phono time series which are then used to assess the performance of the investigated algorithms. The computation cost of each algorithm is extracted and given in processor independent format, thus giving an idea of the applicability to real-time FHR processing .

A time domain beat-to-beat (BTB) FHR estimation algorithm, namely the "Pseudo-Inverse Filter", is given, and its performance is subjectively compared to the pulse correlation approach. The peusdo-inverse filter is shown to offer increased accuracy in the detection heart beat events. Contributions to block processing FHR estimation algorithms are also detailed in this work, notably the "Comb Filter", "Modified Block Autocorrelation" and the "Iterative Block Clean-up".

The performances of fourteen FHR estimation algorithms are compared using varying SNR phono signals. Those algorithms based on block processing are further subjected to accelerating FHR signals to assess their tracking performances. Using pre-determined FHR inclusion bands, two measurable performance quantities are obtained: The outlier error rate (OER) and the in-band standard deviation (IBSD).

Overnight recordings of phonocardiographic signals are processed in real-time using the "Comb Filter" on a stand-alone, VME based, signal processing facility. FHR estimates obtained from two second blocks are then transmitted to the Mainframe computer for further processing and display.

A novel post processing algorithm utilising an adaptive least mean squares filter in a forward/backward predictor configuration is employed off-line to detect and correct FHR outlier estimates. A number of software modules are also used to extract relevant heart rate parameters such FHR baseline, accelerations and variability.

ACKNOWLEDGEMENTS

I would like to thank Dr. James H. Dripps for his supervision and guidance, Prof. Peter Grant for his advice, and Dr. Ken Boddy for his medical assistance and expertise. I also thank numerous colleagues the department of electrical engineering for their productive support.

I dedicate this thesis to my parents as acknowledgement of their long lasting support.

ABBREVIATIONS

ADC:	Analogue to Digital Converter.
AGC:	Automatic Gain Control.
A.I.:	Artificial Intelligence
BPM:	Beats per Minute.
BTB:	Beat-to-Beat
BTBV:	Beat-to-Beat Variability
CF:	Comb Filter
CTG:	Cardiotocograph
CPU:	Central Processing Unit
DAC:	Digital to Analogue Converter.
DSP:	Digital Signal Processor
ECG:	Electrocardiogram
EPROM:	Electrically Programable Read Only Memory
FBPC:	Forward Backward Predictor Corrector
FECG:	Fetal Electrocardiogram
FFT:	Fast Fourier Transform
FHR:	Fetal Heart Rate
FHRV:	Fetal Heart Rate Variability
FIR:	Finite Impulse Response
FSE:	Fundamental Spectral Estimator
HR:	Heart Rate
HRV:	Heart Rate Variability
IBSD:	In-Band Standard Deviation
IFFT:	Inverse Fast Fourier Transform
IIR:	Infinite Impulse Response
LED:	Light Emitting Diode
LMS:	Least Mean Squares.
LVDT:	Linear Variable Differential Transformer

MAW:	Maternal Abdominal Waves
MBAC:	Modified Block Autocorrelation
MD:	Modulus Difference
NST:	Non Stress Test
OER:	Outlier Error Rate
PC:	Personal Computer
PHONO:	Phonocardiographic / Phonocardiogram.
PVDF:	Polyvinyl Fluoride.
RAM:	Random Access Memory.
SNR:	Signal to Noise Ratio.
TAS:	Time Averaged Series
US:	Ultrasound
VME:	Industry standard electronic bus.

CONTENT

CHAPTER ONE

1. INTRODUCTION

Recent technological advancements in microelectronics has enabled the use of powerful processing machines to a wide range of applications. Bio-medical engineering has been positively effected by these advancements with a widespread use of the new technology across the medical disciplines. The continuing processing power improvement afforded by modern machines has justifiably led to a matched demand increase in both quantity and reliability of bio-medical information. Generally speaking, each medical application requires both the optimisation of hardware configuration, and the development - or adaptation - of signal processing algorithms to suit its needs. In most cases, the development cycle of the application can be achieved on general purpose computers or flexible bus based systems (such as the industry standard VME bus).

The main concerns of this thesis are firstly the development of fetal heart rate (FHR) estimation algorithms, secondly the determination of reliability and accuracy of these algorithms and others investigated, and thirdly the subsequent processing of FHR results. This chapter discusses the medical relevance of antenatal FHR monitoring, current practice and some issues associated with the field. The use of the non-stress monitoring technique to assess fetal health is outlined, and the reliability of this approach as an indicator of fetal well-being or fetal compromise is discussed. Some

of the pitfalls associated with basing fetal health diagnosis on features extracted from FHR records are highlighted. An overview of the thesis concludes this chapter.

1.1. The Monitoring of Fetal Heart Rate

Historically, clinicians have listened to the fetal heart sounds in order to determine the heart rate. The term auscultation is used for this auditory FHR estimation procedure which was performed early on by applying the ear directly to the maternal abdomen but later made use of the stethoscope. A poor heart rate quantisation is obtained from auscultation and therefore only gross FHR changes can be correlated with fetal health.

The monitoring of FHR is believed to be essential in assessing the general health of the fetus. Due to cost and other considerations, only short term observations are currently possible, and therefore less confidence can be established between certain suspect FHR traces and the long term health of the fetus.

1.1.1. Current FHR Monitoring Practice

Due to its lower signal loss, the Doppler Ultrasound (US) currently dominates the field of antenatal fetal health assessment. Improved digital signal processing techniques such as autocorrelation has led to an increase in the reliability of FHR estimates. Consequently, a higher confidence can be attached to parameters derived from the continuous FHR trace[34, 43, 82]. Additional signals such as maternal uterine contractions, fetal body and breathing movements may be combined with the fetal heart rate trace for the dynamic assessment of fetal well-being. Other parameters including amneotic liquor volume, maternal hypertension, maternal blood pressure and fetal size make up the long-term biophysical profile of both patient and fetus.

The policy implemented for antenatal patient care varies from one establishment to another but mainly depends on the assessed state of both patient and fetus. 'High' risk patients are admitted into hospital for regular monitoring sessions which can be supplemented by ultrasound imaging. Although a reactive NST is currently accepted to indicate adequate fetal reserve for seven days, Phelan *et al*[84] suggest that more frequent testing would reduce fetal death rate.

1.1.2. Interpretation of FHR Patterns

The instantaneous value of fetal heart rate reflects the interaction of the various autonomic nervous control systems; parasympathetic and sympathetic. Blocking these in fetal lamb was shown to produce a drop in the short-term beat-to-beat variability from 1.6 to 0.6 ms with a similar effect on the medium term variability[16, 20]. The existence of variability in the FHR signal underlies the maturity and health of the fetal control system as its different nervous parts contrive to set the balance. External stimuli were found by various researchers to effect the character of the FHR: Walker *et al*[28] found that compressing the fetal head by applying pressure to the lower abdomen induced bradycardia while tachycardia can be brought about in response to sound or mechanical stimulation. A normal FHR trace is expected to display the following characteristics:

a - A baseline FHR within the range 120 to 160 BPM. This is expected to decline towards the end of gestation maturity[63].

b - Periodic episodes of FHR variability having amplitude in the range 5-20 BPM. There does not seem to be a general concurrence of opinion on the level of fetal heart rate variability (FHRV) which denotes fetal well-being, some of the suggested figures are: (>5 BPM)[63], (>10 BPM, lasting >60 s)[81] and (>20 BPM)[21].

c - A count of four or more significant baseline accelerations per hour. These are normally coincident with fetal activity and last a few seconds, (>10 BPM, >15 s)[81] , (>15 BPM, >20 s)[19, 73].

d - Certain acceptable patterns of minimal heart rate deceleration during uterine contractions[63].

Dawes *et al*[67] suggest that the mere presence of the high variability is not as good an indicator of fetal well-being as is the intermittent presence of variability. Findings by the same researchers suggest that approximately 84% of high FHRV fetal states last between 20 and 50 minutes[67]. This change of FHRV is associated by other researchers to human fetal states[1, 19, 68, 80], or fetal lambs[20].

Studies done by Chew *et al*[48] have shown that while normal FHR characteristics are accepted as good signs of fetal well-being, the absence of these does not immediately imply fetal compromise. Similar findings were also reported by other workers in the field[13, 19]. Epochs of flat tracings lasting about 20 minutes are commonly observed in normal healthy fetuses, especially during periods of low activity (e.g. deep sleep episodes). Spencer *et al*[88] found that quiet FHRV periods last 12-93 minutes (mean 25 minutes) and that 45 minutes of monitoring had to be performed to outlast 95% of low FHRV periods. In an experiment involving 634 patients, 63% of fetuses with non-reactive FHR tracing had normal outcome, Lawson *et al*[81] went on to implement a automated system which employs a variable observation period to monitor FHR: The monitoring epoch is extended until such time as an assuring reactive trace is obtained.

The slowing of the FHR is generally considered to be evidence of fetal asphyxia[63]. The persistence of this condition leads to acidosis which in turn causes varying degrees of fetal compromise (from brain damage to eventual death).

1.1.3. The Stress and Non Stress Tests

The Non Stress Test (NST) is so called because the patient is not subjected to any form of invasive stimulation in order to induce a fetal heart rate reaction. NST therefore, bases the fetal well-being assessment on the observation of the main FHR parameters detailed above together with the maternal uterine contraction activity. Sometimes NST is supplemented by ultrasound imaging and typical observation periods last about half an hour. Various score systems based on such derived parameters were devised to ascertain the fetal health state. NST is a relatively inexpensive approach for antenatal monitoring and has a good true negative catchment percentage. Positive NST outcomes however, have a poorer predictive value of fetal compromise, this can be partially attributed to the limited observation periods[13, 19, 21, 26]. The variable length observation period mentioned earlier[81] can shorten the average monitoring session and may lead to a percentage increase of the true positive results.

As alternative approaches to the passive NST, various stress based tests are available. These consist of monitoring the fetal heart rate patterns while inducing stress into the fetal environment. The oxytocin challenge test is such an example which evokes maternal uterine contractions thus stressing the fetus[21]. The reaction of FHR can then be used to detect abnormal patterns. Stress based tests are commonly used for high risk patients or consequent to a positive NST monitoring session which denotes possible fetal compromise[19].

1.2. Issues Related to Current Monitoring Practice

The main issues relating to current FHR practice are discussed in the following sections: These issues pertain to the reliability of FHR estimators, antenatal medical care, cost of providing antenatal monitoring and patient accessibility to care.

1.2.1. FHR Estimators Error Rate and Accuracy

As investigated later in detail (chapter four), all FHR estimators are prone to performance degradation when operating with poor quality signals. When one considers that the main objective of monitoring is to identify fetuses at risk, the importance of obtaining the best possible FHR traces becomes apparent. In addition, the cost of not implementing good discriminatory systems can be high, both in human and financial terms. Wrong interpretation of FHR traces can lead to either an intervention where none is needed, or to a vital delay in taking decisions. This FHR trace mis-interpretation was found most common when carried out by inexperienced staff who were unaware of the limitations of the monitors used[26, 65, 72, 76]. As Amato[83] remarked on the use of signal processing techniques:

"*Documentation of technique limitations must be understood in order to avoid potential inappropriate intervention*".

For example, in the case of the autocorrelation processing, the technique limitation may consist of rate halving, and in the case of time domain estimators a rate doubling can result when the true FHR is low[26].

The ability to correctly detect ominous patterns is strongly linked to the integrity of the estimating algorithm. By 'integrity' here, it is meant the ability to either present FHR estimates with a high degree of confidence, or indicate episodes of low confidence signals. As observed by Divon *et al*[32] and Fukushima *at al*[40] (1985)

when using autocorrelation based monitors: in the absence of fetal signals (or in the absence of the fetus itself!), fictitious 'fetal heart rate' traces can still be obtained. Furthermore, these were sometimes visually classified as good quality signals. Fukushima goes on to point out that:

"The possibility of missing data, while undesirable, is clinically less important than the addition of spurious data that may lead to inappropriate action".

Tromans *et al*[85] suggested that if more than 6 FHR outliers/minute are detected, then the record pertaining to the epoch in question should be reported. Dawes *et al*[70] noted that by rejecting records with low signal qualities, the number of observed decelerations was reduced by 42%. The latter observation will be confirmed by simulations in chapter four which analyses the performance degradation of the investigated FHR estimation algorithms when applied to poor signals.

1.2.2. Intra-observer and Inter-observer Differences

While the current literature suggests a discontent with the predictability value of the positive NST[13, 19, 81], it is difficult to see how any improvement can be obtained by using a manual system of parameter extraction from FHR. The problem lies with subjective influences when obstetricians interpret FHR patterns. In a study conducted by Van Geijn[79], fifteen distinguished obstetricians were asked to evaluate fetal heart rate traces, significant differences in opinion led to discrepancies in the selected course of action (table 1.1).

	No intervention	Intervention
case 1	4	11
case 2	10	5
case 3	7	8
case 4	7	8

Table 1.1 (Van Geijn[79]) Course taken by fifteen obstetricians when evaluating four different cases.

Flynn *et al*[43] suggested that the simple binary reporting of 'reactive' or 'non-reactive' FHR provided a better correlation with outcome than three established score based systems (Birmingham, Myer Menk/Fisher and the Cardiff scores). Other approaches made use of finer FHR trace report grading. All such attempts to increase the predictability of FHR derived parameters seem to be hindered by the inability of observers to be consistent in reporting these. Various workers point out the fact that as well as considerable disagreements existing between each other, observers had significant disagreements with themselves[34, 43, 52, 79, 82]. When inter-observer assessment of individual variables was scrutinised[82], the lowest level of agreement was found in the estimation of FHR oscillation frequency and the highest level of agreement was found in the estimation of the baseline FHR (the latter is readily available from the FHR trace). In addition, there was significant disagreement in the assessed FHRV level present in the signal. Borgatta[34] concludes on the subject:

"*The reliability of a test limits its predictive value for a given outcome; a test cannot have a correlation with an outcome that is higher than the reliability of the test itself*".

One obvious conclusion is that wherever possible, FHR trace parameters should be extracted by the computer. The increased reliability and consistency in deriving FHR parameters may enable the meaningful use of finer scores as fetal health indicators.

1.2.3. In-Home Remote Monitoring

The commercialisation of stand-alone computers has in recent years led to the widespread use of these in biomedical applications[25, 41, 85]. Although early biomedical implementations were limited to in-hospital use, recent reductions in the cost of computing technologies make possible the use of these at the patient's home. When coupled to the public telephone network via modems, stand-alone computers (or specially designed hardware) enable the antenatal monitoring of patients directly from their homes[53]. Remote monitoring is particularly applicable to cases where long-term or frequent monitoring sessions are required where the need for immediate intervention is not anticipated. Antenatal monitoring presents an ideal case for such a scenario. The major benefits expected from the application of home monitoring are:

a - The patient is relieved from the disruptions of having to travel for monitoring. These disruptions can effect family life or work routines.

b - The reduction of antenatal monitoring cost.

c - Increase in the accessibility of antenatal care, this is particularly beneficiary to patients in remote areas.

d - Longer monitoring sessions may be implemented thus providing a clearer indication of fetal status.

e - More frequent monitoring sessions.

Dalton *et al*[24] pointed out the cost advantages of implementing fetal home telemetry. When all costs were taken into account, Dalton found that the cost of a monitoring session was only 3% that of admitting the patient into hospital.

1.2.4. Phonocardiographic Long Term Monitoring

Although a clearer biophysical profile can be obtained from long-term monitoring, cost considerations and concern about possible over-exposure to ultrasound limits the use of this technique to short, in-hospital monitoring sessions. Phonocardiography is the modern electronic equivalent of the original passive listening technique. Being strictly non-invasive, phonocardiography is inherently more applicable to long term FHR monitoring. Talbert *et al*[66] (1985) developed a transducer based on a piezo-electric sensing element (beam), the dimensions of which are modified to achieve compliance matching to the maternal abdominal wall. Using this (TAPHO) transducer, Colley *et al*[1] (1986) claim a success rate of 100% in detecting fetal heart sounds from patients having a gestation age over 30 weeks. Phonocardiographic signals will be described in detail in the following chapter.

1.2.5. Other Derived Parameters

Additional parameters which are amenable to remote monitoring in the antenatal health care scenario:

a - Fetal heart rate: This has been discussed above.

b - Fetal breathing: Breathing movements detected at the maternal abdominal surface, or reported by the patient[52, 56, 57].

c - Fetal movement: Approach one - electronic registration involving phono transducers with good low frequency response, and approach two - binary reporting by the mother (i.e. whether fetal movement occurred or not). The detection of this signal may reveal the various fetal states (sleep/active)[52, 68].

d - Maternal uterine activity: The interaction of this signal with FHR is routinely used to assess the fetal well-being[63]. Other research suggests that preterm spontaneous labour may be predicted from uterine activity[33].

e - Other relevant maternal factors are: blood pressure and heart rate.

1.3. Thesis Overview

This thesis investigates the performance of various FHR estimation algorithms. Although the research concentrates on applying the estimators to phonocardiographic heart signals, many of the block based algorithms are equally applicable to ultrasound signals.

Chapter 2 is mainly concerned with phonocardiographic signals observed at the maternal abdominal surface. The source and character of each of these signals are detailed with emphasis on the fetal heart components. Both the relative amplitudes and spectral content of the fetal heart signals are detailed. Various transducing principles and technologies are then discussed including two phono transducer contributions. Analogue processing stages involved with the phono signal conditioning are also described.

Chapter 3 introduces the contributions to estimation algorithms, these are divided into two parts: Firstly estimators which base their FHR extraction on blocks of phono signals. These include the 'Comb Filter', the 'Iterative Block Clean-up' and the 'Modified Block Autocorrelation'. Secondly a beat-to-beat algorithm - the 'Pseudo-Inverse Filter' - which is compared subjectively to the industry standard 'Pulse correlation'. A normal and a modified inverse least squares (ILS) filter are also subjectively compared to pseudo-inverse filtering, the modified ILS version derives

its coefficients from an altered autocorrelation matrix to which artificial white noise is added. Most beat-to-beat algorithms make use of a semi-automated template forming method which is introduced.

The performance of fourteen estimators are compared in chapter four. Simulated phono fetal signals with added coloured noise of varying power are used to test each investigated estimator. Two resulting performance graphs under varying SNR are thus obtained: The first of these measures the outlier error rate and the second measures the variability introduced by the estimation algorithm. In addition to coloured noise, block processing algorithms are subjected to an accelerating FHR time series in order to determine their tracking performance. The susceptibility of algorithms to broad-band noise or signal non-stationarity is assessed, those algorithms which suffer significant degradation in performance are reported.

Results derived from long term phono fetal monitoring are presented in chapter 5, together with digital signal processing steps used to extract relevant parameters. A fetal heart variability index is proposed which attempts to emulate the visually assessed variability.

Chapter 6 summarises the work included in this thesis and draws some conclusions from the results obtained. Suggestions are made for relevant further work.

CHAPTER TWO

2. FETAL PHONOCARDIOGRAPHY

This chapter introduces phonocardiography (phono) as a means to non-invasive monitoring of fetal heart rate in preterm patients. The main phono transducer designs currently available are discussed, and two contributions: the optical and capacitive transducers are included. Other transducer technologies are also described briefly.

The definition of a transducer is "a device which converts a physical quantity to an electrical signal". In the context of this thesis 'physical quantity' represents pressure waves, displacement or electrical muscle activity. Velocity and acceleration are also relevant, but since they can be readily derived from displacement, the latter is used to include all. Principally three transducer technologies are identified with fetal monitoring: Ultrasound (US), Electrocardiography (ECG) and Phonocardiography (or phono). The latter is discussed first.

2.1. Historical Development

From ancient history it was accepted that a regular audible pulse emanated from the human chest. Methods of perceiving that pulse varied from applying the fingers to the chest wall to placing the ear directly on the chest. Until the 17^{th} century the

relationship between the heart and the emanating sounds were still not tied together, and claims to that effect made by W. Harvey (1628) were met with scepticism. Laennec (1781-1826) invented the stethoscope and described the occurrence of the first and second heart sounds. Early this century the audible sounds emanating from the heart were investigated in more detail. Luisada (1972) provided a detailed account of auscultation and phonocardiography, and described the relationship between heart activity and the signals observed[30]. Luisada also suggested a break from the traditional pursuance of the audible heart signals to the more general view of a pressure wave causing displacements. Although most of Luisada's work concentrated on adults, the same approach applies to fetal heart phonocardiography.

Auscultation - a method of detecting diseased lungs by listening to the sounds produced by the chest - has traditionally relied on audible sounds emanating from the two main heart phases constituting the heart beat cycle. The auditory system of humans, although used successfully in auscultation based diagnosis, is not ideally suited for detecting fetal heart pressure waves (HPW). The limited auditory range of man at low frequencies and the susceptibility of perception in the presence of additional noise can lead to discrepancies between clinical impressions and objective measurements. The following are found to be true:

a- In the presence of certain loud sounds, the ear is unable to detect fainter sounds that follow. This effect can be demonstrated by playing the recorded signal forward then backward[30].

b- Short term variation in FHR cannot be obtained from auscultation. The accuracy of the heart rate estimate is inversely proportional to the observation period; in order to achieve a one third beat per minute accuracy, fetal heart beats need to be averaged over three minutes.

Influenced by auscultation, diagnostic phonocardiography has until recently concentrated on the audible range of fetal heart wave frequencies. This, as will be demonstrated later, imposes an unnecessary limitation on the extraction of available signal power.

2.2. Phonocardiographic Transducers

Although still unproven, concern about the effects of long term monitoring of patients using invasive monitoring techniques, namely continuous wave ultrasound[50], has led to the investigation of alternative approaches. These include fetal Electrocardiography (FECG) and phonocardiography.

The term phono falsely implies that only audible sounds emanate from the maternal abdomen. In fact, as will be shown later, most of the signal power emanating from the fetal heart is contained in the spectral region below the human audible range. The phono term will hence be defined for our purposes to include displacement frequencies down to DC.

Two main advancements have spurred the viability of phonocardiography, these are:

a- The advancement in piezoelectric ceramic materials and more recently Ploy vinylidene flouride (PVDF) (Talbert *et al*[18, 66]).

b- Improvements in electronic components (low noise operational amplifiers) and design of the transducers have led to a new generation of high performance displacement detectors such as the inductive[51] and capacitive transducers, the latter is detailed in a following section.

Such new transducer developments have enabled the passive and continuous phonocardiographic monitoring of FHR[78]. Attempts at measuring the waves

emanating from the fetal heart rely on three measurement approaches which are detailed below:

a- Energy transfer[42, 66, 87] (matched compliance - matched loading transducers): this approach relies on finding a close match between the compliance of the maternal abdominal wall tissue and the compliance of the transducer. Experiments using a $5cm^2$ disc contact have estimated the average maternal abdominal compliance to be 3.5 mm/N with a five fold variation between patients. The compliance was also found to be inversely proportional to both gestational age and body mass index[42]. Therefore an energy transfer based transducer must ensure close matching between the compliances of the transducer used and the maternal abdominal wall. Studies have also shown that these natural compliance variations result in a proportional signal loss with this type of transducers[66]. The electrical analogy of this class of transducers is circuit impedance matching in order to extract (and hence observe) the maximum available power.

b- Observation (high compliance - minimal loading transducers): This class of transducers offers little or no resistance to the maternal abdominal surface, and produces a displacement related signal. Two sub-classes of observation transducers exist: The first includes those based on the derivatives of displacement (namely velocity and acceleration). They are characterised by a spring-mass combination - an inertial mass is used to create a tensile force which is measured by a strain gauge or a piezoelectric material. These velocity or acceleration transducers have the distinct disadvantage for this application of limited low frequency response. The second sub-class of observation transducers measure displacement. Examples of such transducers include the capacitive and inductive types. The voltage output of these is arranged to be linearly proportional to displacement. Velocity and acceleration are easily obtained from differentiating this output signal. In case of the inductive transducer - generally

known as the Linear Variable Differential Transformer (LVDT) - the position of the core (which is the sensing element) is converted to a differential sinusoidal output proportional in amplitude to the displacement. This output is then amplitude demodulated to obtain the required base-band signal. A variation on this design utilises the inductance value in an LC-oscillator circuit which transduces displacement by frequency modulating the oscillator output[51, 75].

The disadvantage of the 'observation' class of transducers is the need for a stationary reference platform (a mechanical zero) from which the maternal abdominal displacements can be detected. This is usually solved by using a combinations of two methods: (1) An annular guard ring is applied to the maternal abdomen using a double sided adhesive tape with movement being observed at the centre of the ring. Annular guard rings provide compact platform structures for phonocardiographic sensors, however the effects of such a circular cut-out on the maternal/transducer coupling transfer function is not known. (2) A belt is placed around the patient and transducer, this acts as the reference required by displacement transducers. The compliance of the belt is selected to be stiffer than the abdominal tissue.

c- Pressure (low compliance - maximal loading (stiff) transducers): The signal extracted by this class of transducers is proportional to the amplitude of pressure waves produced at the maternal abdominal surface. Transducers of this type need to be strapped tightly against the maternal abdominal wall using an adjustable belt. Such phono transducers are susceptible to corruption by maternal movements which can cause transducer saturation.

2.2.1. An Optical Transducer Design

The first contribution to transducer design is based on an optical modulation principle. Figure 2.1 illustrates this design, full circuit details are omitted for clarity. In terms of construction, the optical transducer is simple: an opaque box and membrane are used to enclose the sensor parts. These consist of a light emitting diode (LED) which is attached to the membrane and a photo transistor fixed to the top of the box. A high compliance membrane is selected to offer little resistance to oscillations at the patient's abdominal surface. The thin rubber membrane with the attached LED tracks the abdominal vibrations. The varying displacement between LED and photo transistor modulates the amount of light the latter receives and consequently its output current. The relevant equations for the principle operation are given:

ζ, the intensity of light at distance (r) from the emitting source is proportional to $\frac{1}{r^2}$.

The intensity is differentiated with respect to the distance r:

$$\delta\zeta = -2\,k\,\frac{1}{r^3}\,.\,\delta r \qquad 2.1$$

k is a constant dependant on the luminescence of the LED. The gain obtained from the optical transducer is much higher for small values of r, for linear operation of the transducer however, this distance must be kept high. Practically, the r is 15mm and δr has a maximum useful signal range of about 1mm.

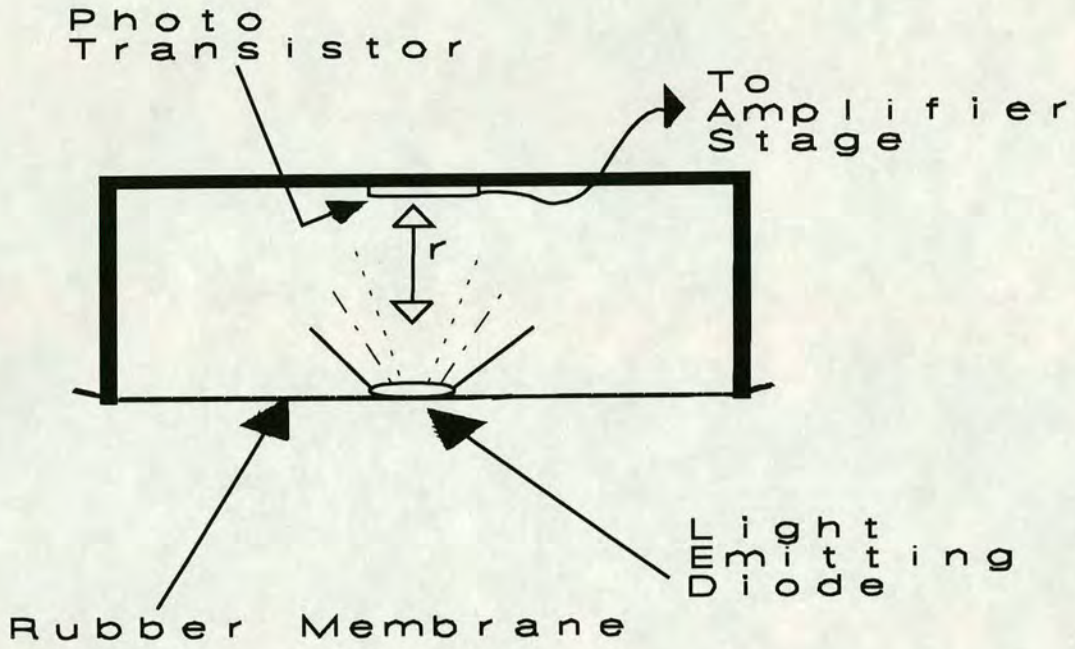

Figure 2.1 Schematic diagram of the optical transducer. The output of the photo transistor is taken through stages of DC rejection and high gain amplification.

Despite its simplicity and low cost, the optical transducer suffers from two main disadvantages; firstly (and more importantly) both the LED and the photo transistor introduce relatively large amounts of noise. This was found in some cases to have similar amplitude levels to signals from fetal HPW's. Secondly the overall signal gain is directly proportional to the brightness of the LED, this however also means higher power consumption.

2.2.2. A Capacitive Transducer Design

Strain gauges have been applied in the past to the detection of fetal breathing movements from the maternal abdomen[57]. Although successful, strain gauges have a low sensitivity and being resistive elements, suffer from noise generation. These problems effectively limit their use to fetal breathing detection - fetal HPW signals

are about 40dB lower in amplitude. In general, the main concerns with transducer instrumentation of fetal HPW are the low frequency performance[66], lack of sensitivity, additive noise, linearity and dynamic range. While linearity can be sacrificed in cases where the details of the HPW's are not being analysed (which is normally the case), the other parameters are crucial to the performance of the overall monitoring system. A variable inductance transducer was proposed which satisfies the above parameters[51, 75]. The capacitive transducer detailed below is a viable low-cost alternative with similar performance parameters.

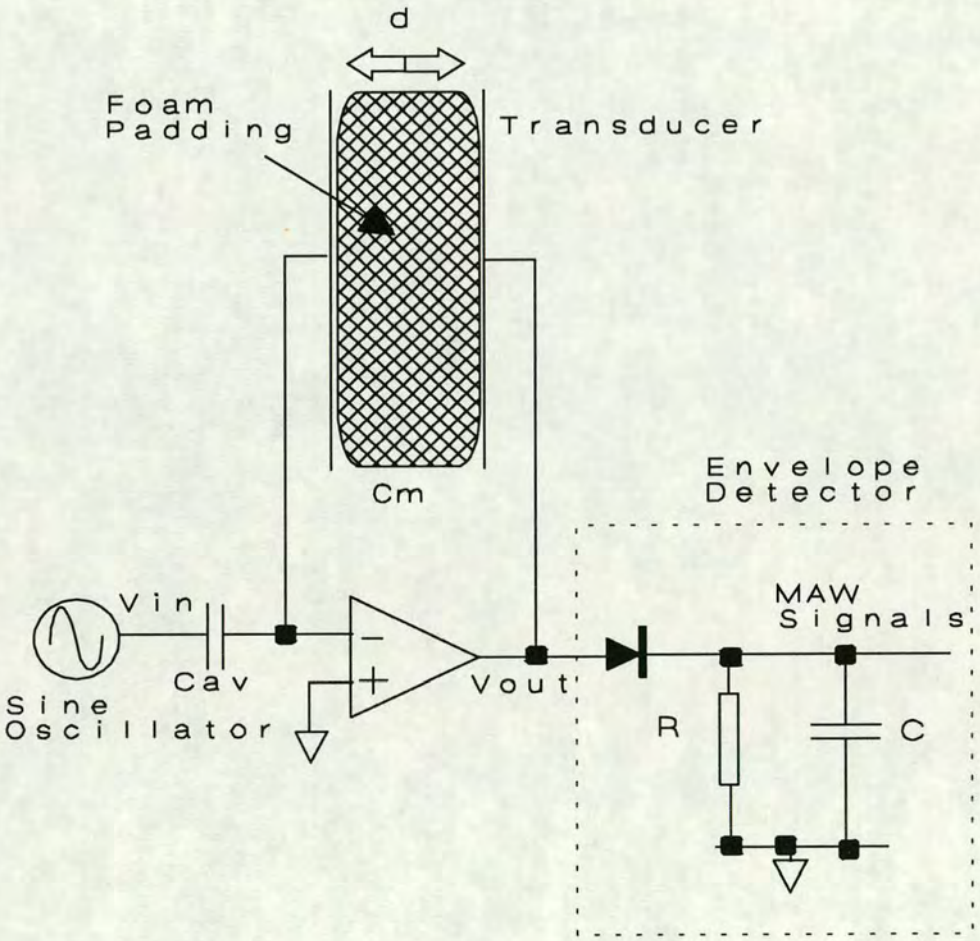

Figure 2.2 Schematic diagram illustrating the capacitive transducer which utilises the maternal abdominal displacements to amplitude modulate the sinusoidal oscillator output.

Presented in figure 2.2 is the sensor part of the capacitive transducer. The transducing element is formed from the relative motion of the two plates which form the capacitor Cm (top of figure). Using a sine wave oscillator at 100 KHz as input to the circuit, the transfer function of the amplifier is given:

$$\frac{Vout}{Vin} = \frac{Cav}{Cm} = \frac{Cav \,.\, d}{\varepsilon \,.\, A}$$

Where Cav is the input capacitor, ε is the permittivity of the foam material between the capacitor Cm plates, A is the area of the overlapping plates, and d the distance between the plates. The circuit output is:

$$Vout = \frac{Cav \,.\, d}{\varepsilon \,.\, A} \,.\, Vin \qquad 2.2$$

It can be seen that a linear relationship exists between the displacement d and the output voltage (eq. 2.2). An envelope detector stage is then employed at the output of the amplifier to produce the required signal. Filtering and gain stages have been omitted from the figure for clarity.

2.3. Continuous Wave and Gated Ultrasound

Currently the field of fetal heart rate monitoring is dominated by the use of ultrasound Doppler transducers. These radiate a high frequency ultrasound beam (2-6 Mhz) which is coupled to the maternal abdomen by a water based gel to ensure good coupling (mechanical impedance matching) to the maternal abdominal tissue. Discontinuities in refractive index in the path of the beam reflect a corresponding fraction of the ultrasound signal[77]. When such boundaries are in motion, the reflected US signal is Doppler shifted. Early US monitors and hand held ones use continuous wave US which is Doppler detected. Newer monitors employ a modified

(or gated) continuous wave approach to avoid the large backscatter component from the maternal abdominal surface.

By modifying the US crystal design, spatial selectivity can be achieved. For narrow US beams, high quality signals are obtained at the cost of increased difficulty in 'illuminating' the desired field. Alternatively, when the illumination beam is widened, interfering signals from undesirable sources are introduced. The quality of US signals obtained is related to the angle of heart illumination, and the rate of monitoring success is found to be patient dependant. In clinical practice, the wide beam US is found to reduce signal loss and is preferred to the narrow beam[24]. The current domination of ultrasound based fetal monitors is related to both the ease of use (wide beam) and lower signal loss obtained in routine monitoring[23].

2.4. Fetal Electrocardiography

This is a very successful monitoring technique when used during labour: a unipolar electrode clip is applied to the fetal scalp from which the electrical activity of the fetal heart can be detected. Used in this fashion, scalp FECG is an invasive monitoring technique which cannot be applied to the pre-term fetus. When the monitoring of fetal heart electrical activity is attempted from the abdominal surface, the FECG contains significant levels of maternal ECG and myoelectric muscle noise which can swamp that of the fetal heart[47]. The monitoring success rate is reduced to under 50% due signal loss or large interfering components. Newer processing approaches using singular value decomposition to decouple the signal from various interfering sources have reported encouraging results[86].

Since the Doppler US technique is, strictly speaking, an invasive process, and abdominal FECG has unacceptably low success rates, phonocardiography is therefore

the only method which currently offers a viable option for the long-term non-invasive continuous monitoring[1].

2.5. Maternal Abdominal Wave Signals

The various signals observed at the maternal abdomen are listed below in two categories, those appertaining to the patient (mother) and those appertaining to the fetus. Due to the low frequency components of these signals stretching into the sub-audible region of human hearing, the term 'maternal abdominal waves' (MAW) is introduced to include all observed displacement signals at the maternal abdominal surface.

Maternal Related MAW: The first and most significant maternal phono signals are those related to movement, such movements can easily cause the transducer to saturate and hence patients are asked to lie in a resting position during monitoring. Slow and continuous movements (as opposed to abrupt) were found not to be a problem. The second maternal component relates to breathing rate and character. Harmonics of the breathing fundamental (centred at 0.3 Hz) can constitute significant components up to 1.5 Hz, which can interfere with both fetal breathing and heart signals. Patients are therefore asked to breathe 'sinusoidally' when fetal breathing is being monitored. Strong maternal heart signals are present at the upper part of the abdomen and decay progressively as the transducer is moved to the normal fetal position (lower abdomen), in practice, these signals are rarely observed.

Fetal Related MAW: Similarly to above, fetal body movements are the largest observed phono components. During fetal activity these can cause the complete masking of other signals. Although fetal movements indicate a positive fetal health state, the computerised detection of these have not to date been extensively

investigated. A simple binary outcome is routinely reported (i.e. whether fetal movement has occurred or not). Two types of fetal movements are identified, these are fetal limb and gross body movements. Recently, there has been an increased interest in the detection of in utero fetal breathing activity, although such movements do not serve any immediate biological function, they are thought to prepare the fetal lung muscles for post-partum regular breathing[56, 57]. The normal range of fetal breathing movements is from 0.5 to 1.5 Hz (30 to 90 breaths/minute). These movements are non-sinusoidal and have a significant spread of harmonics. The real signals of interest here are the fetal HPW's which are also the smallest of all the above components, even with optimal transducer placement. The fundamental heart rate components are centred close to 2.3Hz (140 BPM), in this region however fetal heart signals are overwhelmed by constituents of the above described signals. Frequency components related to the heart beat have been observed from 10 to 80 Hz. The following section details the fetal heart signals.

The relative strength of the phono abdominal signals listed above varies significantly with the positioning of the transducer. These are listed in descending order of amplitude: maternal movement, fetal movement and maternal breathing, maternal heart, fetal breathing and lastly fetal heart. Fetal breathing and fetal heart signals are approximately inter-spaced by a power difference of 40dB.

In addition to the MAW signal sources emanating from mother or fetus, a further noise component is introduced by external vibrations which are due to the motion of persons in the vicinity of the monitoring area. The amplitude of these have been observed, in the extreme, to exceed the fetal heart component.

2.6. Fetal Heart Pressure Wave Signals

The traditional reliance on the audible sounds emanating from the fetal heart has, until recently, concentrated most phonocardiographic work on the upper part of the spectrum (45-100 Hz). As explained earlier, the normal fetal heart rate lies close to 2.3Hz and the observed higher frequency oscillations are not directly related to its harmonics. These are instead associated with vibrations caused by the contraction of the heart. Contrary to earlier belief which was based on the audible part of the spectrum, the heart waves are not solely caused by the heart valve leaflets clapping. Only part of the spectrum (specifically the higher frequency components) can be associated to valve activity[2]. Cardiac chambers and arterial vessels mechanically coupled by blood and tissue constitute a structure which may only be stimulated into vibration as a whole. This does not contradict the view held by Wriggers (1915) who claimed: "*Vibrations of different structures maintain their identity and can be identified in heart sounds*". Wriggers also stated that "*the sudden elevation of blood pressure produced by contractions sets many structures into vibration*". Once the cardiohemic structure has been set into vibration by the abrupt contraction, each structure may settle into its own mode of resonance. This is modified from the isolated mode of resonance by the mechanically coupled neighbouring organs. Experiments involving simultaneous imaging have shown the existence of sizable interval between actual mitral valve closure and the first heart wave[30].

Given the presence of substantial oscillations in the sub-audible region of the spectrum, the term 'pressure wave' is introduced to denote both the audible and the sub-audible vibrations under study. Such waves which are due to the fetal heart activity are called fetal Heart Pressure Waves (HPW). Therefore the first and second audible heart sounds constitute only part of the first and second fetal HPW's respectively.

The first fetal HPW is generated at the start of ventricular systole (contraction). Inward acceleration of the heart wall and the closure of the mitral and tricuspid valves cause the closely coupled cadiohemic system to vibrate. Although both the right and left ventricles contract simultaneously, the left ventricle is much larger in size and is the main contributor to the first HPW. The first pressure wave which encompasses the first 'heart sound' has a low frequency oscillatory shape and is commonly heard as a thump. The second HPW emanates at the start of the ventricular diastole (relaxation). It is caused by the semi-lunar valves stopping reverse flow of the arterial blood. This impulse change of kinetic energy induces the pressure waves in the arterial walls, the observed waves are characterised by their higher frequency composition and is identified by its characteristic clicking sound. The second HPW is usually weaker and shorter in duration than the first HPW.

Two other pressure waves (traditionally referred to as the third and fourth sounds) are well documented but are more difficult to isolate from the background noise. Currently there seems to be no prospect of basing routine clinical assessment on their detection. Additionally, murmurs are also present, these are thought to emanate from blood flowing from a restricted passage to a more dilated vessel. In mechanical terms, this situation is similar to strong air flow causing oscillations in large structures.

Figure 2.3 shows a waterfall short term spectrum of a typical HPW signal. Considerable signal energy is evident at frequencies down to 20 Hz while the resonance is centred at 30Hz. The spectral character of the first and second HPW's are evident, the first is more resonant in nature while the second has a broader spectral content. In certain cases of high body resonance or narrow-band filtering (fig. 2.4), the second HPW attains a pulse shape similar to the first, although smaller in amplitude. Also seen at 70Hz (fig. 2.3) is another component resonance which is classically identified with the fetal heart sounds. Its occurrence can be clearly seen

during first HPW pulses. Typically, the lower resonance mode is one order magnitude larger in amplitude than the resonance traditionally associated with the audible range.

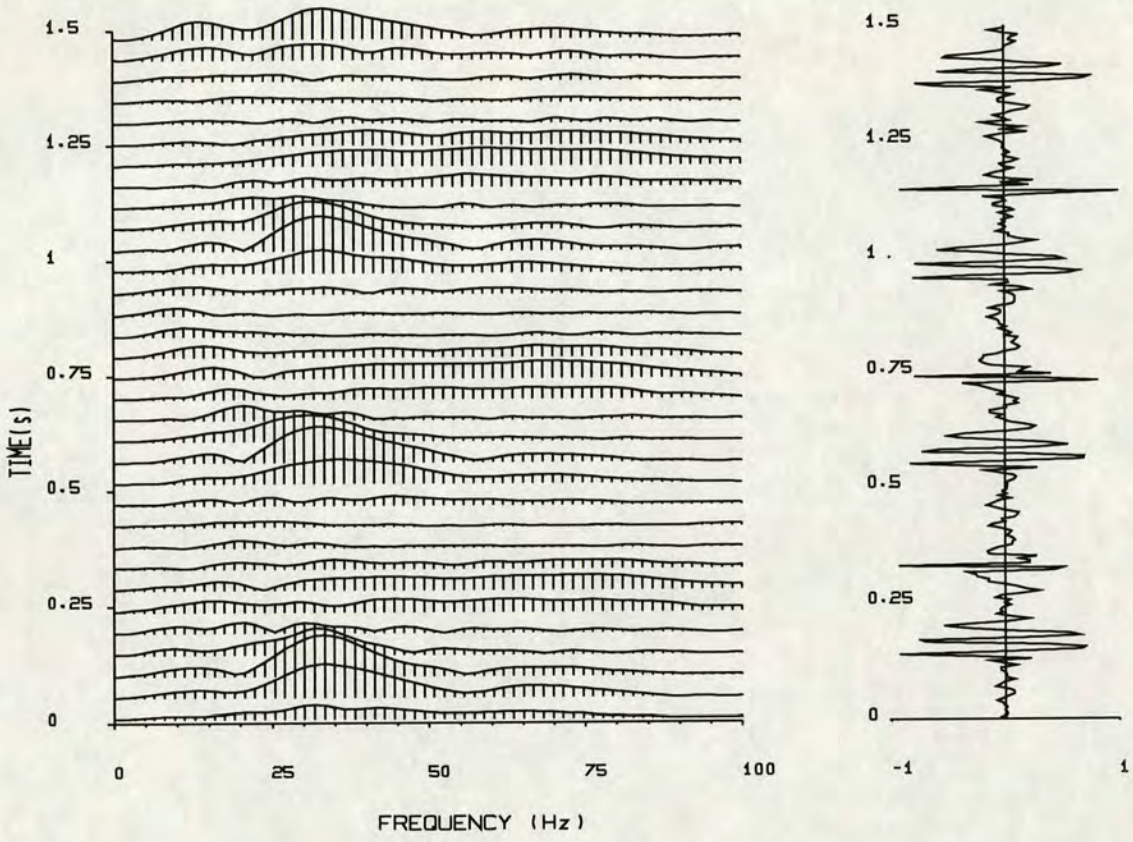

Figure 2.3 Linear spectral waterfall (left) and corresponding time series (right) of typical HPW signals. The spectral traces are obtained from 200ms Hanning windowed records overlapping by 160ms.

A 200ms window length is selected to contain all of the first HPW pulse, larger block lengths are not suitable for this quickly varying signal. The 80% overlap provides a degree of visual interpolation of the changing spectrum as well as ensuring that no spectral information is lost[45]. The industry standard Hanning window provides good sidelobe suppression while maintaining a relatively narrow peak[44, 46].

The example HPW time series given above (fig. 2.3) is an instance of a fetal phono recording. Other observations have shown that the first HPW resonance can lie in a wide range of the spectrum with 16 and 50 Hz being the lowest and highest centres of resonances recorded. The above spread in centre frequencies is probably attributed to the obesity of the patient - larger bodies have a lower frequency of resonance than smaller ones. The audible part of the HPW have been observed from 45 to 80 Hz. Sometimes the distinction between the audible and the sub-audible components is difficult unless observed from a good quality signal such as shown in figure 2.3. Earlier transducers based on piezoelectric crystals[66, 37] had a poor low frequency response and hence tended to suppress the main sub-audible HPW component.

The parameters of interest which are extracted from the HPW signals are: (a) The duration of the interval between the first and second HPW (T1), this interval is normally constant even during FHR change[62]. (b) The duration of the interval between the second and the first heart HPW of the following heart cycle (T2), and (c) The heart beat period (T), where T = T1 + T2.

2.7. The Conditioning of Fetal HPW Signals

The first stage of the fetal HPW signal conditioning consists of a band-pass filter which removes the unwanted MAW signals. The centre and bandwidth of this filter are selected using switchable, cascaded fourth order high-pass and low-pass Butterworth filters[36]. The bandwidth ranges between 15 and 40 Hz depending on the signal quality. When good quality signals are obtained, the filter bandwidth selected is not critical, for the example given earlier (fig. 2.3) the corner frequencies are 16 and 85 Hz.

The resulting band-pass filtered HPW's typically have an amplitude which is two orders of magnitude lower than the raw MAW signals (fig. 2.4 middle and top traces respectively). Therefore a set gain amplifier (not shown) followed by an automatic gain control (AGC) stage ensure a reasonable signal level. Beat-to-beat based estimators commonly use this version of the signal (labelled BPF in figure).

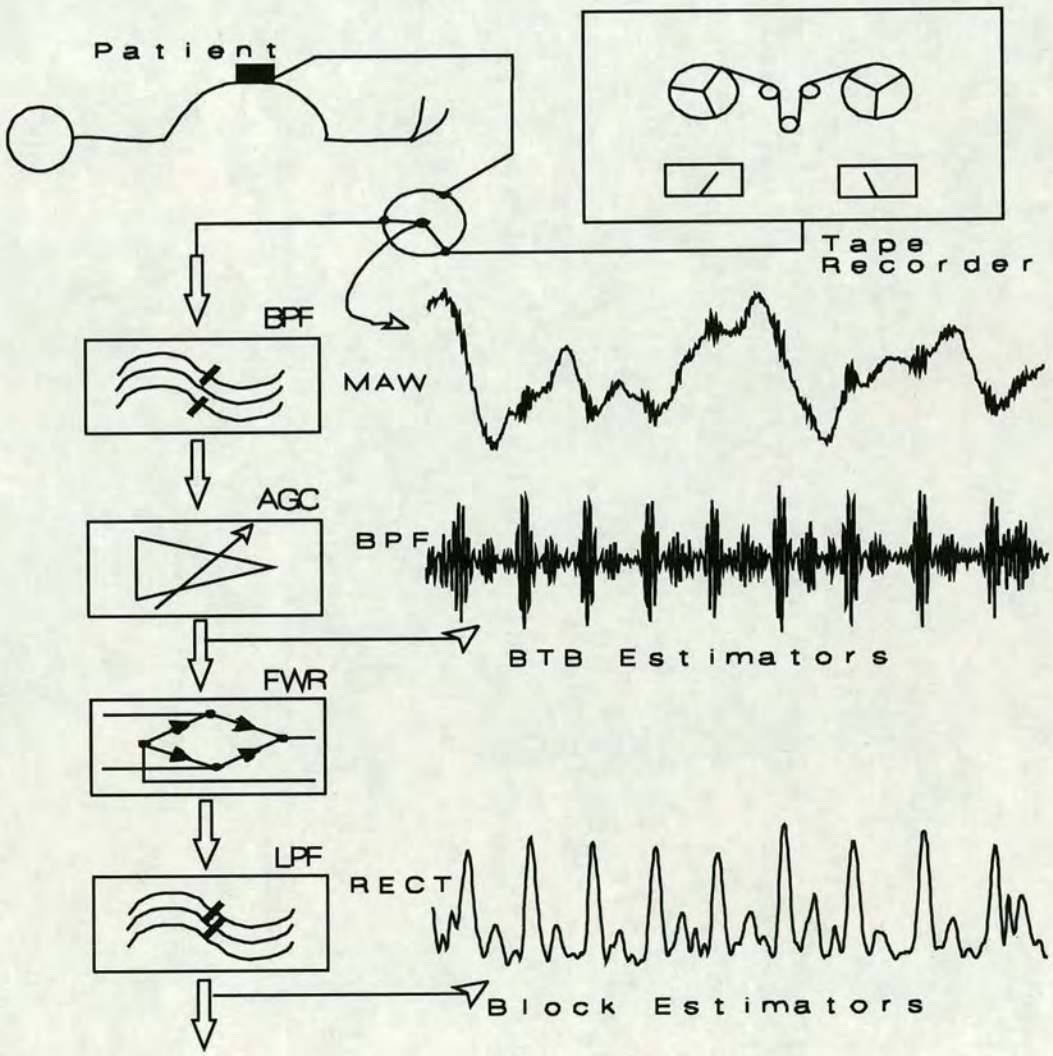

Figure 2.4 Illustration of the analogue processing stages employed in the conditioning of the maternal abdominal wave signals.

The signal is then passed through a full wave rectifier (FWR) followed by a 20Hz 4^{th} order low-pass filter which reconstructs the envelope of the heart beat pulses and limits the band-width to meet the Nyquist sampling rate. A final AGC stage (not shown) is employed prior to sampling, this signal (labelled RECT in fig. 2.4) is used by block processing algorithms for FHR extraction. The sampling rate commonly employed at this stage is 50Hz. Although this implementation utilises a 12-bit analogue to digital converter (ADC), 8 bits resolution (providing -48dB digitisation noise floor) is probably sufficient for block based estimators.

CHAPTER THREE

3. FETAL HEART RATE ESTIMATORS

Various signal processing techniques have been applied to FHR estimation, including analogue time domain, correlation and spectral based[38] estimators. The main objectives of these techniques are (a) to improve the accuracy of FHR estimates and, (b) to maintain reasonable performances during FHR non-stationarity.

It has also been noted[31, 34] that under certain operating conditions, additive noise can cause errors in the FHR estimation process. When noisy signals are processed, the resulting errors are observed as an FHR variability which can have a similar level to, or even exceed, that of the true physiological heart rate control. In addition, these traces look convincingly similar to real traces obtained from good signals[40]. This false variability may be mis-interpreted by the medical staff as a sign of fetal well-being[26]. It is therefore important for estimation algorithms to be able to indicate poor operating conditions.

All FHR estimators are classified as employing either 'Block' or 'Beat-to-Beat' (BTB) estimation. The contributions in this chapter span both block and BTB classes using fetal phono HPW signals to determine FHR. The first approach estimates the FHR from contiguous two seconds sample vectors using block estimators such as "Comb Filtering" (CF) and "Iterative Block Clean-up". The second approach based on BTB

estimators introduces a new "Pseudo-Inverse Filter" (PIF) algorithm and provides a subjective comparison with both correlation and a modified inverse least squares filter algorithm.

3.1. Block Processing Algorithms

Whether these algorithms base their estimates on: (a) Distinct signal blocks with no information feed-through from one block estimate to the next (such as FFT[38], Comb Filter[39], autocorrelation and modulus difference[11]). Or, (b) on a running FHR estimator (such as phase locked loop, continuous autoregressive spectral analysis or recursive demodulation[17]), the resulting estimates have an effective block length which can be specified. This block length has to remain as short as possible in order to retain some short term FHR variability and to localise in time the effects of signal non-stationarity. Typically, the block length is chosen to lie between one and four seconds.

Block based algorithms are generally robust to evolutionary HPW shape changes, this is mainly due to the envelope detection employed at the front end (detailed in ch. 2). In addition, and due to the coherent averaging inherent in block processing, noise immunity is gained. Therefore block algorithms can generally operate at relatively lower Signal-to-Noise Ratios (SNR) than can their BTB counterparts.

3.1.1. Comb Filtering

Dripps *et al*[69] have shown that when the spectrum of the demodulated ultrasound returns is observed over a period containing 2 or more heart beats, a large percentage of the signal power is distributed within the first few harmonics. The relative

distribution of power in these harmonics cannot be pre-determined and can change markedly within a short time period, or even from beat to beat due to the 'scintillation' effect from the ultrasonic target (fetal heart). The same effect (variation in harmonic distribution) is observed in phono HPW signals but to a lesser extent. Figure 3.1 demonstrates the power distribution variability within a record lasting five seconds. In this example the steady harmonic composition during the first two seconds (having a prominent second harmonic), changes towards the end of the observed epoch leaving most of power in the fundamental. This continuous change in the harmonic structure of the HPW signal is found to depend on the following factors:

a) The relative size of the second heart pulse wave when compared to the first HPW. It was noted that for some recordings of fetal phono signals, the second HPW can diminish significantly in amplitude. This is most commonly observed in lower quality recordings and can effect the dual pulse nature of heart beat and consequently its spectrum.

b) The current FHR, as well as changing the position of the constituent harmonics, has an indirect effect on their relative amplitude distribution. This factor is demonstrated in more detail later in this section.

c) Another more subtle effect is the shape of the HPW signal; for some patients, the feto-maternal body resonances are so accentuated that a spreading of both the first and second HPW's results. After Envelope detection, a highly resonant HPW signal produces relatively larger low frequency components. For other patients, both HPW's are short in duration thus resulting in a spread of the power into higher harmonics. The majority of patients observed however, have an HPW signal spectrum with approximately 90% of the power lying in the first three or four harmonics (fig. 3.1).

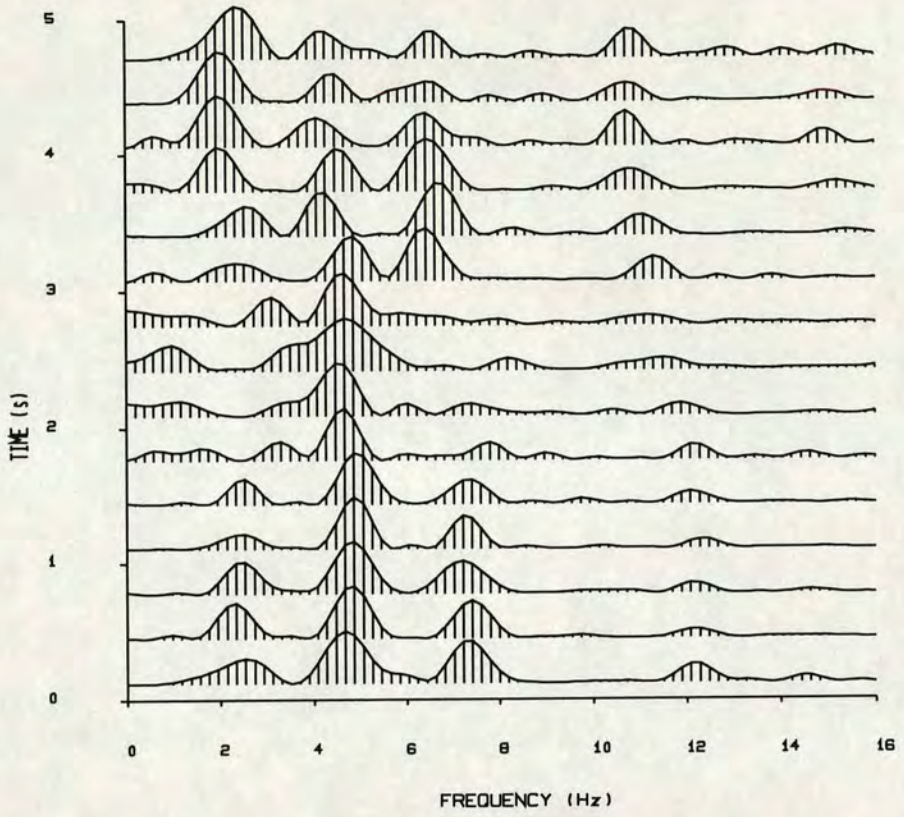

Figure 3.1 Spectral waterfall obtained from five seconds of HPW time series showing a typical harmonic amplitude evolution. The FFT block length is two seconds with a block overlap of 75%.

d) Accelerating (or decelerating) FHR; a fast change in FHR or simply an individual BTB period differing from its neighbours can, in the extreme, produce splitting of frequency domain harmonics. Figure 3.2 illustrates an example using an abrupt 180 degrees phase shift in an otherwise steady and continuous sinusoid. This results in peak splitting of the corresponding frequency domain component. If this sinusoid is the fundamental Fourier component of the HPW time series, then 180 degrees phase step corresponds to a 1/2 periodicity discontinuity. While such a large step is an unlikely event, a 1/4 period step change would produce the same effect on the second harmonic, a 1/6 period

change would produce the same effect on the third harmonic. The latter two cases and yet smaller step changes (similarly effecting higher harmonics) are commonly observed phenomena in time domain records and corresponding spectral degradation in the frequency domain. Such events rarely present themselves as 'clean' spectral splitting (such as in fig. 3.2), more frequently, smeared spectral peaks are noted.

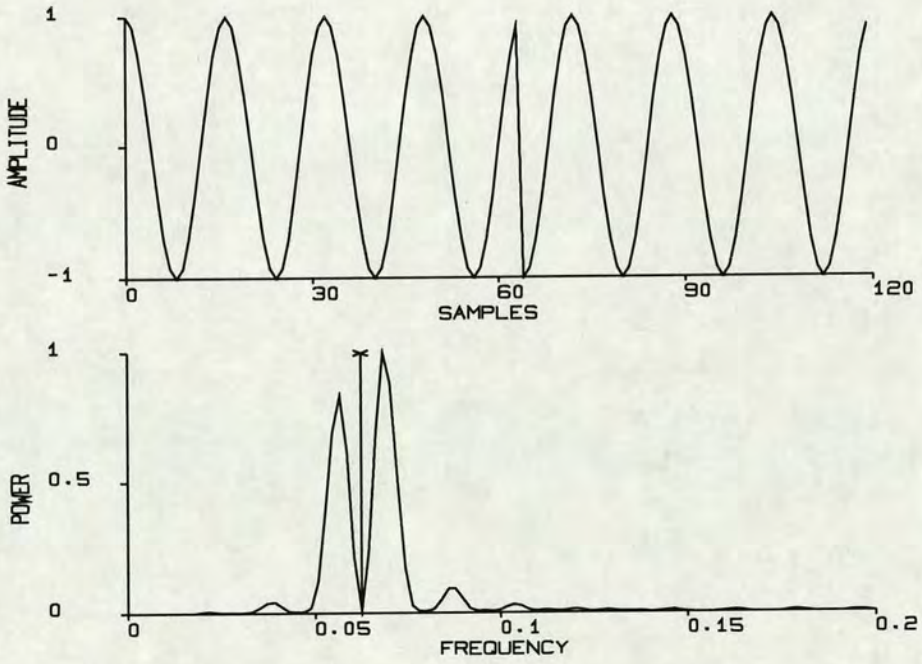

Figure 3.2 Effect on the spectrum of a simulated 180 degrees phase step change. The marker position (bottom) denotes the frequency of the sinusoid.

The interaction of the above mentioned effects prevents apriori assumptions about the harmonic signal structure of the HPW signals. This harmonic distribution must therefore be regarded as a dynamically evolving process. As a result, heart rate

estimators cannot rely on the relative amplitude of any given harmonic. Therein lies the difficulty of FHR estimation in the frequency domain.

None of the above mentioned factors may be taken independently when analysing an HPW harmonic distribution. In most cases, however, the amplitudes of the first and second HPW's and the feto-maternal resonances are, in the medium term at least, stationary. Therefore, given a particular instance of factors (a) and (c) described above, the dominant short term effect on the harmonic distribution is the FHR itself.

The effect of FHR on the harmonic distribution is demonstrated pictorially in figure 3.3 which depicts a simulated FHR deceleration from 200 to 90 BPM. About 10 seconds into the record (corresponding rate of 120 BPM) the second harmonic gives way, gradually, to the signal's fundamental and third harmonic. It must be pointed out here, that this harmonic distribution versus FHR (fig. 3.3) is particular to the instance of HPW relative amplitudes and pulse resonances, (factors (a) and (c) above).

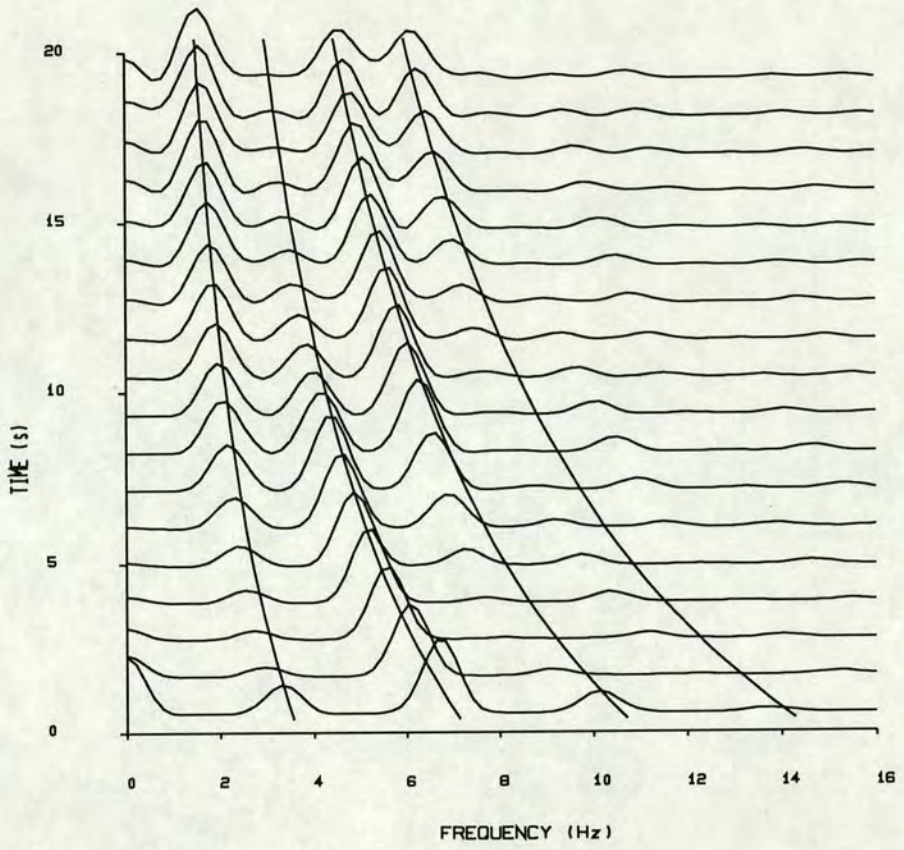

Figure 3.3 Spectral waterfall of a simulated decelerating FHR, starting at 200 BPM and ending at 90 BPM (3.3 to 1.5 Hz respectively). The FFT block length is three seconds, Hamming windowed with a block overlap of 67%.

In consequence to the above, a frequency domain FHR estimating algorithm which takes advantage on this harmonic spread without making apriori assumptions concerning the specific distribution may prove most effective. The comb filter[12] algorithm is proposed for this task[39, 69].

3.1.1.1. Comb Filter Algorithm

The Comb Filter (CF) introduced here assumes that the sampled HPW signal consists of two main components:

(1) A set of periodic impulses, the strongest of which is the fetal heart component and,

(2) A set of aperiodic signals which define the noise background when inspected in the frequency domain. This noise can be broadband or highly coloured.

The comb filter output, $T(W_p)$, is expressed in the continuous frequency domain as the sum of the harmonics of a signal with a fundamental frequency W_p[12]. The CF output is defined as:

$$T(W_p) = \int_{-Wc}^{+Wc} I(W,W_p) \cdot X(w) \; dW$$

Where Wc is the signal's cutoff frequency and X(w) is its corresponding power spectrum. The problem is now reduced to defining the kernel function I(W,Wp) which exhibits two principal properties:

(1) $I(W,W_p)$ must be defined such that $T(W_p)$'s output has a global peak when evaluated at W_{fhr}, the fundamental of the periodic HPW time series. This permits simple global peak detection to be carried out in the transform domain. To avoid making any assumptions about the distribution of harmonics in the signal, $I(W,W_p)$ itself will need to be periodic in W_p and its maxima lobes positioned at $k \cdot W_p$, k being an integer. In the case of normal comb filtering, this kernel function reduces to impulses at W_p intervals.

(2) When the function is 'looking' for a fundamental where no periodic signal exists, the output $I(W,W_p)$ should (ideally) sum up to a constant value depending on the noise floor. In the case where the kernel

function is a balanced window (i.e. with zero mean), the expectation of the output when I() is evaluated over white noise should be nil.

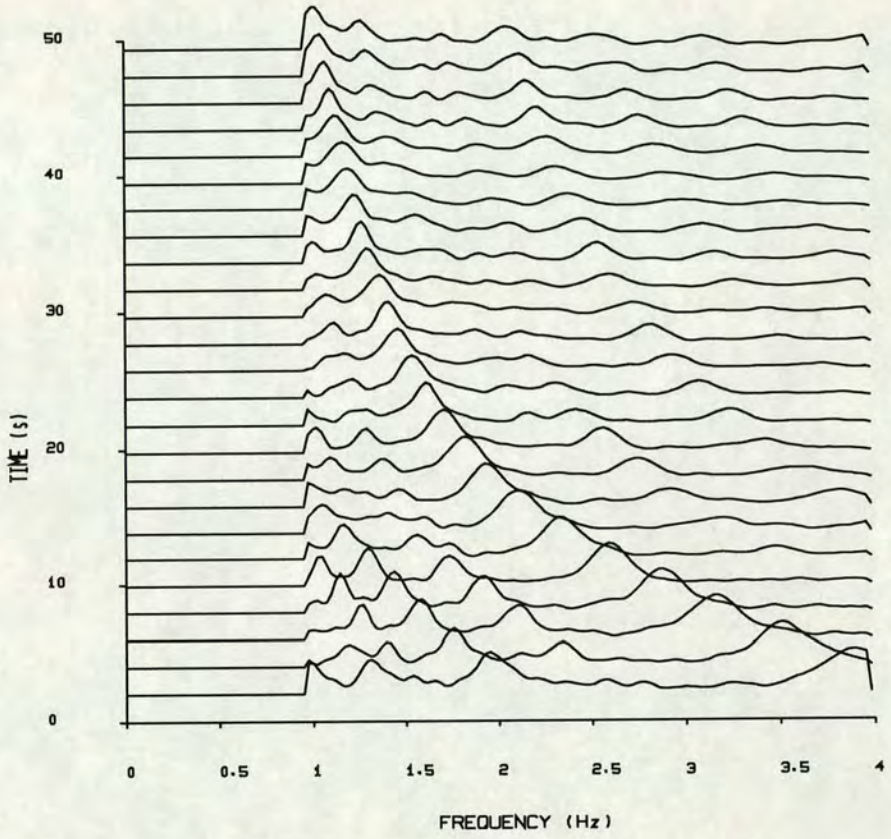

Figure 3.4 Comb filter output when applied to a 50 second simulated fetal heart trace deceleration from 4 to 1 Hz (240 to 60 BPM). A two second block duration is used with a 50% overlap.

Due to the cyclic nature of $I(W,W_p)$, the comb filter's output can contain significant power levels at both harmonic and sub-harmonic frequencies. The comb filter output of a noiseless decelerating FHR is depicted in figure 3.4. In the presence of noise the global peak could occur at a sub-harmonic (during high FHR) or at a harmonic (during low FHR). The use of harmonic tapering or restricting the number of these

(i.e. the number of teeth in the comb filter) can solve this problem. The latter option is utilised because it makes no apriori assumptions about the harmonic power distribution. Care must be exercised however, as a severe truncation of the number of harmonics can lead to harmonic error. In practice five harmonics are included in the kernel function $I(W,W_p)$, and the output of the comb filter is only computed over the valid FHR range (1-4Hz).

Given that $T(W_p)$ effectively estimates the HPW signal power by summing up the power at its fundamental and harmonics, the comb filter can also provide an estimate of the SNR with little computational overhead:

$$SNR(W_p) = \frac{T(W_p)}{P - T(W_p)}$$

Where P is the total power of the signal and noise. The obtained SNR estimate ($SNR(W_p)$) can be used as a measure of confidence in the selection of W_p as the heart rate. The discrete time version of the algorithm is given:

$$T(W_p) = \sum_{\psi=0}^{N-1} I(W_p, \psi) \; . \; X(\psi)$$

If the window applied at each of the W_p harmonics does not change, I() can be expressed as:

$$I(W_p, \psi) = \zeta(\psi) * \sum_{k=1}^{H} \delta(k \; . \; W_p - \psi)$$

Where:

H number of harmonics considered,

$\zeta(\psi)$ is the windowing function applied to each of the harmonics.

Further simplification of the comb filter computation leads to:

$$T(W_p) = 2 . \sum_{k=1}^{H} \sum_{\psi=-a}^{+a} \zeta(\psi) . X(k . W_p + \psi) \qquad (3.1)$$

where:

2.a +1 is the width of the kernel function applied to each harmonic.

The kernel function $I(\psi)$ can now be though of as a $\zeta(\psi)$ - impulse train which is superimposed on each of the harmonics.

Figure 3.4 illustrates the operation of comb filtering, with the resulting global peak consistently pin-pointing the underlying FHR. While this is found to be the case in most of the processed blocks, sudden bursts of interfering signals such as fetal or maternal movements can produce multiple local peaks in the CF output. In addition, recordings with poor signal to noise ratios can result in a large proportion of the estimates being obvious outliers (with large error margins) or marginally erroneous (errors which deviate from the true FHR by a small margin). Therefore the FHR extraction algorithm which post-processes the comb filter output must exercise some degree of artificial intelligence (A.I.) when selecting the most likely FHR peak.

When Comb filtering is tested using a synthetically generated signal with no additive noise (such as in fig. 3.4), 100% correct results are obtained, with any deviation from exact FHR being limited by the frequency domain sampling quantisation effect.

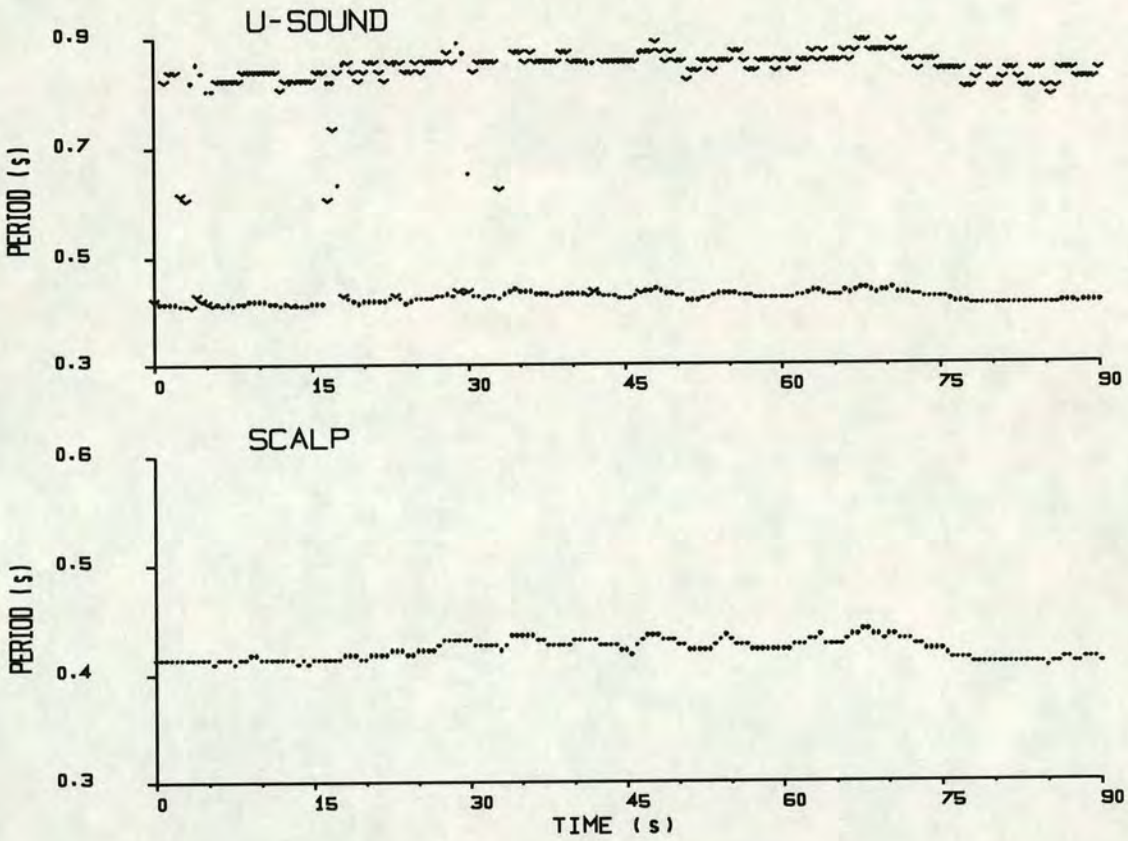

Figure 3.5 Comb filtering applied to simultaneous direct scalp FECG (bottom trace) and ultra sound records (top trace). Dots mark the location of the largest peaks and 'v' marks the location of second largest peaks (US case only). CF is applied to 2 second signal epochs using a transform size of 2048 with a 75% overlap.

To assess the CF performance against a gold standard, simultaneous recordings of ultrasound returns and high quality scalp FECG signals are used. A difference of more than one quantisation bin between the scalp FHR and the comb filter peak is regarded as an error in the US estimate. From the example of ultrasound record given in figure 3.5, the following results are obtained: (a) 15 sub-harmonic errors, (b) 3 definite errors, (c) 42 one bin quantisation differences (from the scalp simultaneous estimate) and (d) 105 exact matches. When the time records of the 3 definite erroneous US estimates (occurring between 15 and 20 seconds in figure 3.5) is closely inspected (fig. 3.6), a severe disintegration of the US signal is revealed.

Although the exact reason for this instance of corruption is not known, it is commonly related to fetal or US transducer movements. All definite errors observed in the comparison between US and scalp were due to similar time domain corruptions.

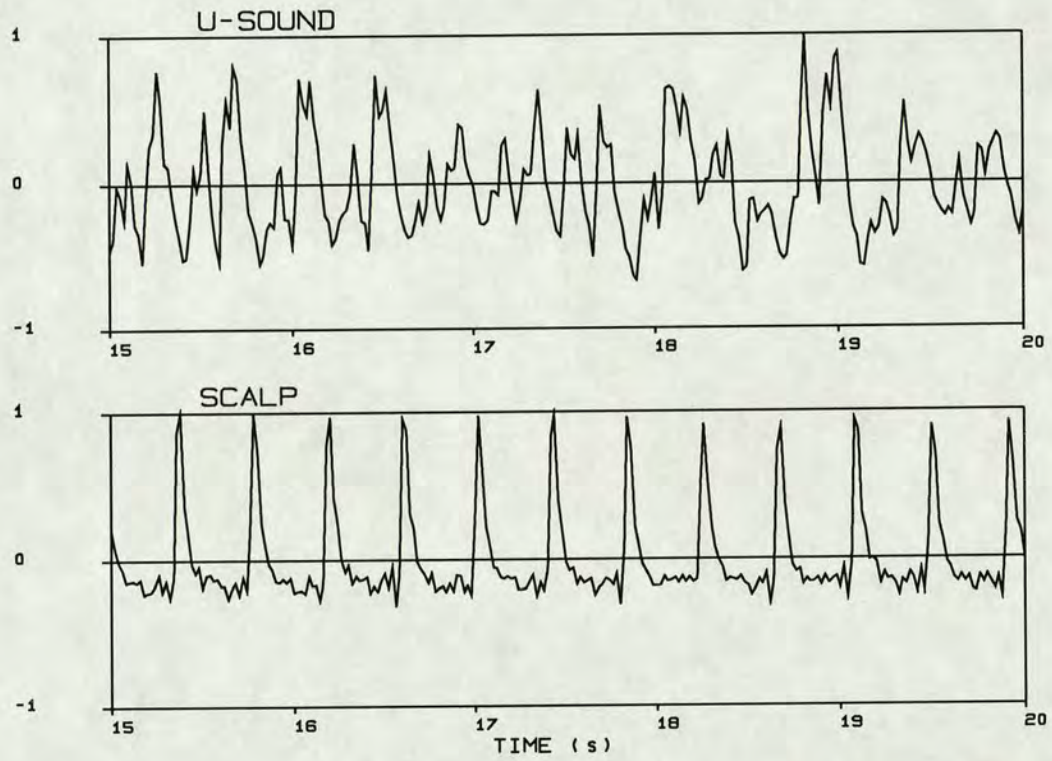

Figure 3.6 US and Scalp signals corresponding to the erroneous CF estimates output of figure 3.5 (15-20 seconds into the record). A reasonable signal quality from 15 to 17 seconds gets progressively corrupted by the interfering noise which is worst between 18 and 19 seconds.

In the case of the 15 sub-harmonic errors, the second largest peak in the CF output usually pin-points the true FHR. When the SNR of the estimated peaks is inspected (fig. 3.7), a clear correlation between the erroneous estimates and lower SNR values can be seen. Moreover, signal blocks resulting in sub-harmonic errors have amplitudes which are only slightly larger at the sub-harmonic position than at the true

FHR (fundamental peak). These sub-harmonic errors were caused by an US signal of marginal quality. The use of a carefully designed peak selection A.I. algorithm (discussed later) avoids the erroneous selection of most sub-harmonic errors. During epochs of good quality US signals and without recourse to the use of the A.I. peak selection technique, less than 3% of the obtained estimates were found to be in error.

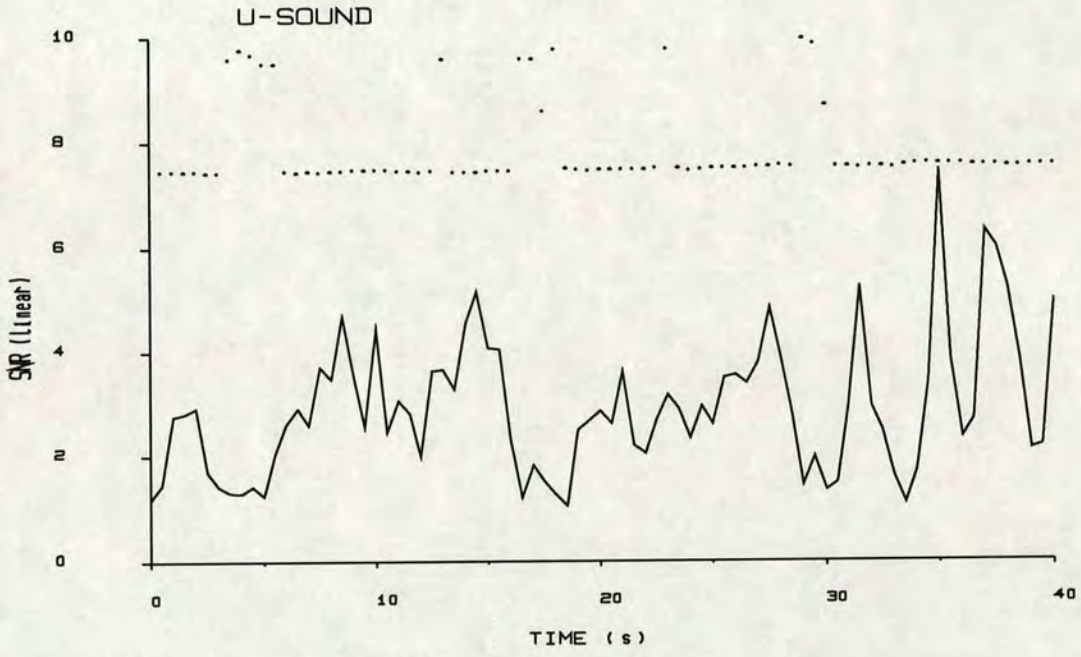

Figure 3.7 The solid line shows the first 40 seconds of signal to noise linear power ratio pertaining to the largest comb filter peak. The superimposed dots are the positions of the comb filter global peaks (as in fig. 3.5). Most errors occur in regions of low SNR.

Figure 3.8 demonstrates the effectiveness of the comb filtering algorithm in maintaining the global peak over the underlying FHR when applied to envelope detected real phonocardiographic signals. As shown earlier, no individual harmonic component can be relied upon to provide a consistent global peak for FHR estimation. The particular selected example (fig. 3.8) provides a fast deceleration which causes

spreading of the spectral components. This in turn causes the estimated SNR to drop at about 3.5 seconds into the trace. The deceleration depicted here is more severe than is observed on average, similar drops in estimated SNR however are frequently observed when spurious noise contaminates an otherwise good quality signal. In addition, the SNR values measured at 4 seconds into the trace may be classified as marginally unacceptable while those values measured at 9 seconds are classified as marginally acceptable. Figure 3.8 clearly shows the sub-harmonic peaks during the first 4 seconds of this example which give way to harmonic peaks in the period from 4 to 8 seconds.

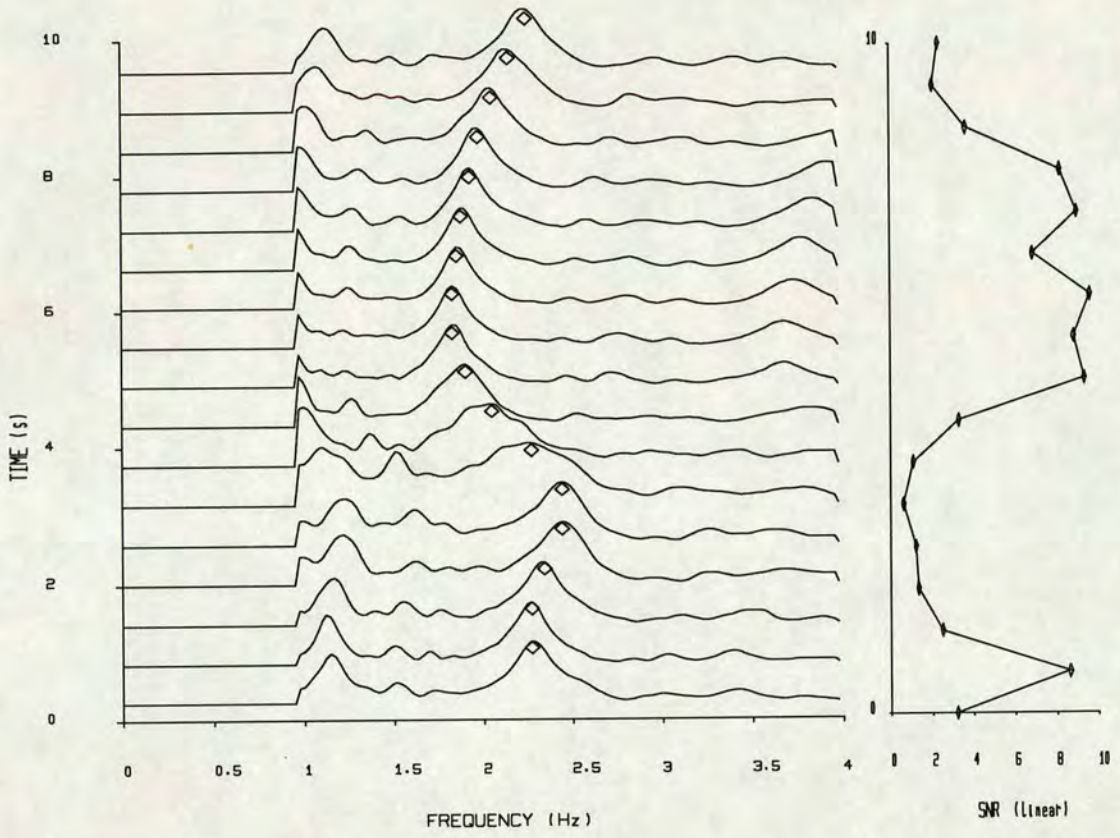

Figure 3.8 Comb filter output (waterfall on left) tracks the heart rate during a fast FHR deceleration. Block length is set to 2 seconds with a 75% processing overlap. The diamond markers superimposed on the waterfall traces denote the position of the global peak. The trace (right) shows SNR(W_p) which provides a measure of confidence in the corresponding global peak.

A more extensive investigation of the performance of the CF under varying noise levels is given in the following chapter, and results of long term monitoring using this technique are included in chapter five.

3.1.1.2. Comb Filter Kernel Functions

Alternative kernel functions ($I(\psi)$) based on different windows, $\zeta(\psi)$, were tried and tested in order to select the optimal kernel function. Varying operating conditions,

however, present diverse performance factors which must first be addressed. These factors are:

(a) Noise immunity: This is the most important factor in selecting shape and width of the window $\zeta(\psi)$. The wider the window, the higher the influence of neighbouring noise components on both the position and shape of the comb filter output peak.

(b) Comb filter peak width: (or peak flatness) effects the ability of the comb filter to resolve correctly the global peak position, the performance degradation (for a flat peak) is most acute in systems using fixed point arithmetic. Arithmetic noise arises from the fixed point computation of both spectrum and comb filter output. When this is superimposed on an otherwise flat-top peak, considerable estimate jitter results. It is therefore desirable to make use of windows which provide a fast roll-off from the harmonic peak position.

(c) Computational efficiency: This factor is relevant for stand alone FHR monitors which must operate in real time. Window functions based on multiplications or divisions would be computationally intensive. The triangular function is such an example (fig. 3.9). Simpler windows which are based on additions and subtractions require less processing power when using current general purpose microprocessors.

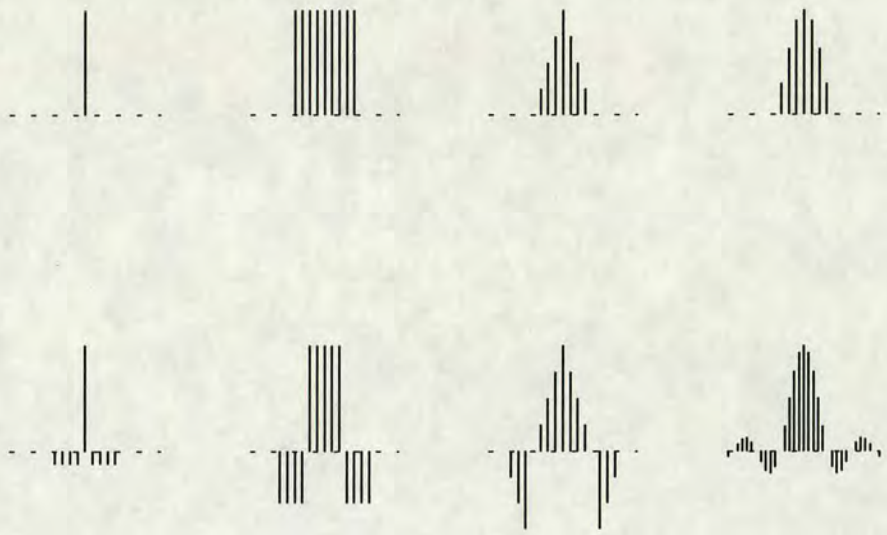

Figure 3.9 Some of the windows $\zeta(\psi)$ which were used as part of the kernel function $I(W_p,\psi)$. These include unipolar windows (top) and some balanced window functions (bottom).

A subset of the windows investigated is shown in figure 3.9. The distinguishing feature of these windows are: (a) whether they are balanced (resulting in a zero expected output when evaluated over white noise), (b) window shape (impulse, square, etc ..) and (c) the width of the central lobe.

Balanced windows (those with zero mean weights) can theoretically be employed to cancel the background white noise. In practice however, most of the observed noise is highly coloured. In consequence, balanced windows were found to have some major drawbacks: (1) Narrow balanced windows tend to suppress the underlying signal especially during FHR acceleration, and (2) wider ones are strongly influenced by coloured noise close to one of the harmonic (fig. 3.10) and (3) a split in one of the harmonics (as shown in figure 3.2) can lead to the total degradation of the comb filter

output. The resulting comb filter output is very similar to that produced in the presence of coloured noise contamination (fig. 3.10).

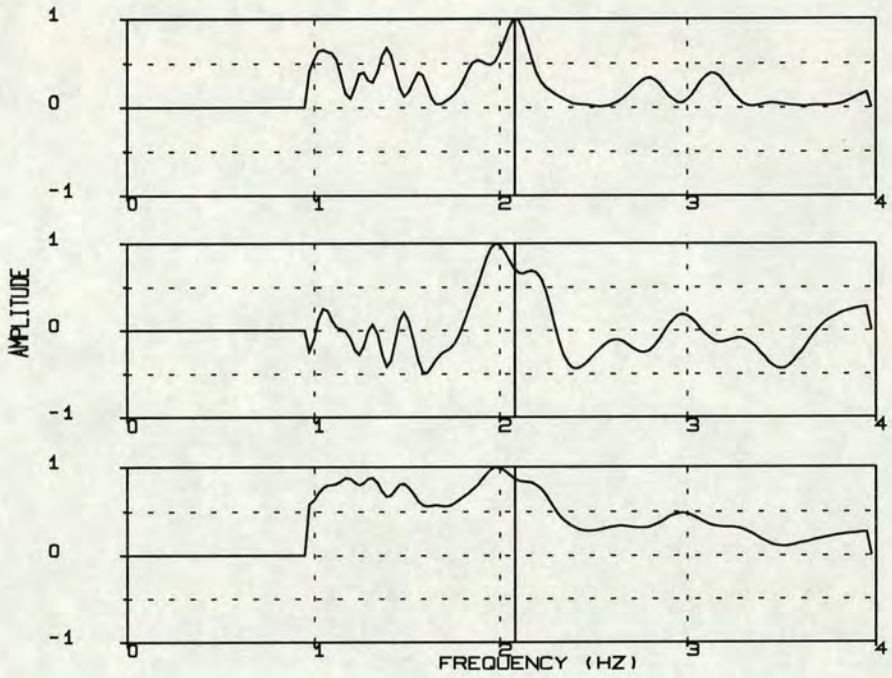

Figure 3.10 Illustration characteristic comb filter output peak degradation when subjected to highly coloured noise (-5dB w.r.t. signal power). The true peak position is denoted by the vertical line. The impulse window function produces the top trace (triangular window has a similar output), the middle and bottom traces are produced when a 'balanced square' and 'square' windows are used respectively.

The window choice can effect the global peak shape; the convolution of the window with the spectrum causes peak spreading with a consequent loss of resolution. Only the top trace in figure 3.10 escapes this degradation (impulse window is used). As mentioned earlier, a flat-top peak can, when combined with arithmetic noise, lead to FHR estimate jitter. Suitable windows which produce acceptable peak shapes are the impulse, triangular and the sinc functions.

When discussing window widths, the term 'narrow window' refers to a window which is narrower than the main spectral lobe in the Fourier transform of the padded time records. The converse applies to the term 'wide' window. There does not seem to exist an ideal width for all the types of windows investigated. While wide square windows cause global peak flattening, narrower ones have a better performance. Further narrowing of the square window eventually degenerates the latter into the impulse type which is treated separately here. Narrow triangular windows offer sharp peaks with the non-balanced ones providing a better performance during FHR changes. Narrow balanced triangular windows on the other hand suffer from significant performance degradation. The full Sinc function is equivalent to the impulse window when used to interpolate the spectrum. Both sinc and impulse functions offer good noise immunity and sharp peak roll-offs. The sinc function width must match that of the main spectral lobe when it is used to combine functions of windowing and interpolation (explained later).

Window Type	Advantages	Disadvantages
Impulse	Good noise immunity Sharp global peak roll-off Simple to implement	Requires large Fourier spectrum
Balanced Impulse	Reasonable noise immunity Sharp global peak roll-off Simple to implement	Assumes white background noise Requires large Fourier spectrum
Square	Simple to implement	Bias to neighbouring noise Flat global peak
Balanced Square	Simple to implement	Strong bias to neighbouring noise Flat global peak
Triangular	Reasonable noise immunity Sharp global peak roll-off Copes well with accelerations	Slight bias to neighbouring noise High computation cost
Balanced Triangular	Sharp global peak roll-off	Strong bias to neighbouring noise Multiple local peaks High computation cost
Clipped Sinc function	Good noise immunity	High computation cost
Full Sinc	Good noise immunity Amenable to interpolation	High computation cost

Table 3.1 Performance comparison of investigated window functions $\zeta(\psi)$.

As stated earlier, most of the observed noise when using fetal phonocardiography consists of bursts of highly coloured noise. It is therefore of little benefit to use balanced window functions which were initially introduced to deal with white background noise. Square windows were found to introduce peak smearing. Triangular windows, especially the non-balanced, while providing good global peak resolution by virtue of a sharp roll-off are computationally expensive. The full sinc window function, on the other hand, can be used to interpolate from low resolution spectra and therefore offsets the overall computational cost. Detailed later, is the full sinc interpolation method and the reduction in computing it affords. A summary of the characteristics of the windows discussed is shown in table 3.1.

3.1.1.3. Peak Selection Strategy

Ideally, the comb filter output would contain a distinct single global peak pin-pointing the FHR. In practice however, and during intervals of low SNR, multiple peaks are observed and an artificial intelligence algorithm is required to select the peak which most likely reflects the true FHR. Any peak selection algorithm must take into account the statistical FHR parameters such as current heart rate, trend and probability distribution. A heuristic artificial intelligence algorithm is designed to extract the most likely FHR by assigning 'favouring factors' to the comb filter peaks. The favouring factor is computed from an equation which is based on the following three criteria:

(a) Peak power: When operating above an acceptable SNR lower limit, the relative power contained within a peak is a good measure of the probability of this peak pertaining to the true FHR. Additive noise however, can introduce incorrect global peaks. While these can appear anywhere within the CF output, noise most commonly causes sub-harmonic errors and, slightly less often, harmonic errors.

(b) Closeness of the peak to previous FHR estimates: An exponentially weighted filter having a time constant between 2 and 3 estimates was found most successful. Longer time constants inhibit the correct tracking of fast or abrupt FHR changes, while shorter time constants can cause a chain of consecutive erroneous estimates.

(c) The statistical distribution of FHR is an important factor in effectively limiting the range of possible FHR values. The distribution can be preset (fixed) or can be allowed to evolve by continuous updating. It was found that the non-evolving type was adequate, but additional information pertaining to gestational age (apriori baseline value) can

increase the effectiveness of this factor. Severely curtailing a non-evolving distribution may lead to consistent harmonic or sub-harmonic errors.

Given the above criteria, if the algorithm is presented with a choice of multiple peaks of similar power, it should select the peak which is closest to the previous estimates. Possible outliers caused by coloured or spurious noise are thus rejected in favour of estimates which preserve the continuity of the FHR time series. In addition, if a non-evolving statistical distribution is employed, then the favouring weight of criterion (c) (above) would have to be less than that of criterion (b). Two cost functions Cp() and Cd() are required to implement the selection algorithm. The first relates recent FHR values and the second to the assumed statistical distribution. The equation for the favouring factor is given:

$$V(W_p,n) = T(W_p,n) \cdot Cp(\,|\,FHR_{avg}(n-1) - W_p|\,) \cdot Cd(\,|\,A - W_p|\,) \qquad 3.2$$

Where

n	is the block number (time index).
$V(W_p,n)$	the favouring factor at frequency W_p
$T(W_p,n)$	CF power estimate at frequency W_p (criterion (a)).
$FHR_{avg}(n-1)$	FHR average of previous estimates (criterion (b)).
A	centre of the FHR statistical distribution (criterion (c)).

The comb filter peak which attains the highest favouring factor from equation 3.2 is assumed to denote the true FHR. In order to reduce the computation load, only the three largest peaks are considered. A first order recursive filter is employed to obtain an exponentially running average of the FHR:

$$FHR_{avg}(n) = \frac{1}{2}\{\, FHR_{avg}(n-1) + FHR(n) \,\}$$

3.1.1.4. Limitations of Comb Filtering in Real Applications

Large amplitude low frequency noise generated by fetal or maternal movements can severely disrupt the normal spectral shape. This results in most of the power lying at low frequencies which tends to bias the CF output, and consequently, the peak selection algorithm. While the tracking algorithm is robust to short and isolated spurious corruptions, a chain of these was found to introduce temporary sub-harmonics locking. However, such bursts of erroneous estimates are readily identifiable and it may be possible for artificial intelligence programs to distinguish them.

While the range of FHR values accommodated by the processing spans more than one octave, the threat of both harmonic and sub-harmonic errors will exist. As described earlier, the majority of outliers seems to originate from low frequency noise. Consequently, it may be advantageous to raise the lower FHR search limit from 60 to 70 BPM. This is a reasonable proposition since 'fetal home remote monitoring' was never meant to cope with critical (and possibly fatal) heart rates which 60 BPM represents. A raised lower limit such as 70 BPM can still provide indications of ominous states of bradycardia as well as normal temporary FHR decelerations. The same reasoning can be applied to the upper FHR limit of 240 BPM although these higher frequencies were not seen to consistently bias estimates.

3.1.1.5. Interpolation Technique

When Fourier transforms are used for FHR estimation, the common approach to improving the quantisation is by zero-padding the time records before the transform is applied. This effectively provides an over-sampled frequency spectrum and is

equivalent to time domain oversampling in order to improve quantisation in correlation based processing. A quantisation of 1/3 BPM is suggested for research purposes which translates into a 8192 FFT size when the sampling rate is 50 Hz. This however constitutes a considerable computational load and is beyond the real-time capabilities of general purpose microprocessors (such as MC68000).

Since the interest in low quantisation is limited to the region of the estimated FHR, it seems wasteful to compute the whole spectrum. An efficient solution adopted for this work uses the binary search algorithm combined with an interpolation technique at the comb filter output stage. Having defined the region (interval) of interest which contains a peak with a smooth curve, the binary search algorithm iteratively bisects the interval and continues the search on the interval containing the larger of the two extremes. Figure 3.11 illustrates an example of a binary search for the global peak: The interval [x0 .. x1] is the initial region of interest. Two iterations lead to the consecutive range selections [x2 .. x1] and [x2 .. x3], the final choice is x3. The values at points x2 and x3 are computed using a full interpolation which is derived below:

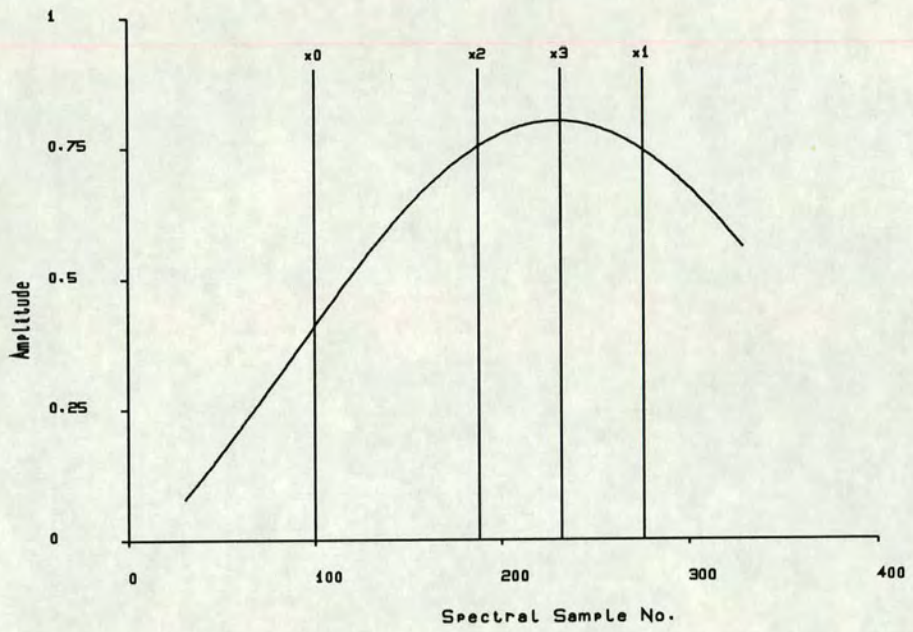

Figure 3.11 Example of the binary search locating the global peak.

Consider the discrete Fourier transform component at frequency $\omega_n = \frac{n \,.\, 2 \,.\, \pi \,.\, fs}{N}$, fs is the sampling rate:

$$X(\omega_n) = \sum_{k=0}^{N-1} x(k) \,.\, \theta^{-k.n} \qquad 3.3$$

$$\theta = e^{\frac{j.2.\pi}{N}}$$

Normally, all $X(\omega_n)$'s are computed at equally spaced frequency bins, where n is integer in the range [0 .. N-1]. For an FFT implementation, N is normally chosen to be a power of two (i.e. $N=2^L$, L is an integer). Equation eq. 3.3 also applies to spectral components at frequency ω_i, where i is not necessarily an integer:

$$X(\omega_i) = \sum_{k=0}^{N-1} x(k) \cdot \theta^{-k.i}$$

which is rewritten as:

$$X(\omega_i) = \sum_{k=0}^{N-1} \sum_{p=0}^{N-1} X(\omega_p) \cdot \theta^{p.k} \cdot \theta^{-k.i} \quad \text{where} \quad \sum_{p=0}^{N-1} X(\omega_p) \cdot \theta^{p.k}$$

is the inverse discrete Fourier transform of $\underline{X}(\omega)$ for the k^{th} time domain sample. Rearranging:

$$X(\omega_i) = \sum_{p=0}^{N-1} X(\omega_p) \cdot \sum_{k=0}^{N-1} \theta^{k.(p-i)}$$

$$= \sum_{p=0}^{N-1} X(\omega_p) \cdot Y(p-i) \qquad 3.4$$

The term Y(p-i) (eq. 3.4) is effectively the result of projecting vector $\underline{\theta}^i$ onto $\underline{\theta}^p$. Inspection of Y(p-i) shows that its value for a frequency difference of $\delta\omega$ (where $\delta\omega$=p-i) is dependent only on the frequency difference between p and i. Using equation 3.4, it is therefore possible to interpolate the spectrum to any required degree. Normally however, there is a need to improve frequency quantisation by a constant, pre-determined factor β. In that case, a vector $Y(\delta\omega)$ can be computed which serves to interpolate at all β intermediate bins.

Equation 3.4 can be viewed as the convolution of β zero-interleaved spectrum $X'(\omega)$ with the formed vector $Y(\delta\omega)$, the resulting fully interpolated spectrum is $X''(\omega)$:

$$X''(\omega) = X'(\omega) * Y(-\delta\omega)$$

While it is redundant to convolve $X'(\omega)$ (consisting mainly of zeros) with $Y(\delta\omega)$, in practice the implementation uses a β incremented index for the vector $\underline{Y}$.

The efficient method used for FHR extraction can now be given. An initial FHR estimate quantised to the nearest 11.7 BPM is obtained from a 256 point FFT (sampling rate is 50 Hz). Using the binary search and interpolation algorithms, the initial quantisation is iteratively narrowed down to 0.367 BPM to produce the final estimate (equivalent to 8192 point FFT). The computation costs of (a) applying a full 8192-point FFT and (b) the combination of the interpolation an binary search algorithms described above. The annotation $\alpha.O\{k\}$ denotes a computational cost of the order k, the constant multiplicand α is hardware dependant (normally omitted).

Cost of FFT $= O\{ N . lg_2(N) \} = 106496$

where N is the zero padded FFT size, N=8192 for a 0.367 BPM resolution,

Cost of interpolation $= O\{ n . lg_2(n) \} + O\{ n . lg_2(\beta) \} = 3328$

where n is the initial FFT size (256 point) and β is the interpolation factor, $\beta = \frac{N}{n} = 32$

Ratio of computation costs is: $\frac{106496}{3328} = 32$

The 32 fold theoretical improvement closely matches the measured 50 fold improvement.

3.1.2. Iterative Block Clean up

The motivation for this iterative approach stems from the duality of impulse trains in both time and frequency domains; only an impulse train in one domain produces a

corresponding impulse train in the other. In addition, observations of both time and frequency domains, during periods of acceptable SNR, reveal that for stationary interference noise (which excludes sudden fetal movements), noise components tend to lie at low amplitude levels relative to signal peaks.

This section introduces the Iterative Block Clean-up (IBC) which transforms the signal vector from the time domain into the frequency domain after setting to zero all samples that are unlikely to belong to the HPW signals. The cleaning up is achieved by effectively removing, in the time domain, all incoherent components as these are assumed to belong to background noise. Each IBC iteration cycle consists of two steps which are given:

1 - In the time domain, all samples falling below a pre-determined threshold are clipped to zero. Given the impulse train nature of HPW signals, low amplitude samples are assumed to belong to incoherent broadband background noise. The spectrum is then computed.

2 - In the frequency domain, the four largest peaks in the band 1 to 8 Hz are located. Peaks with power lying below 50% of the largest component are rejected. The frequency and phase of each remaining peaks are used to determine the corresponding periodicity and first pulse time delay relative to the start of the block. A pair of impulse trains having: (a) similar power levels, (b) same time delay and (c) periods which are harmonically related, are combined into a single impulse train which has a periodicity equal to the higher of the two. A pulse gating vector is then obtained for each of the remaining frequency peaks, this is weighted by the amplitude of the respective peak. The sum of all the resulting gating vectors is then convoluted with an appropriate rectangular window resulting in the desired gating

function. This gating function is then applied to the original HPW block.

The above sequence of steps is then repeated in the second and subsequent iterations.

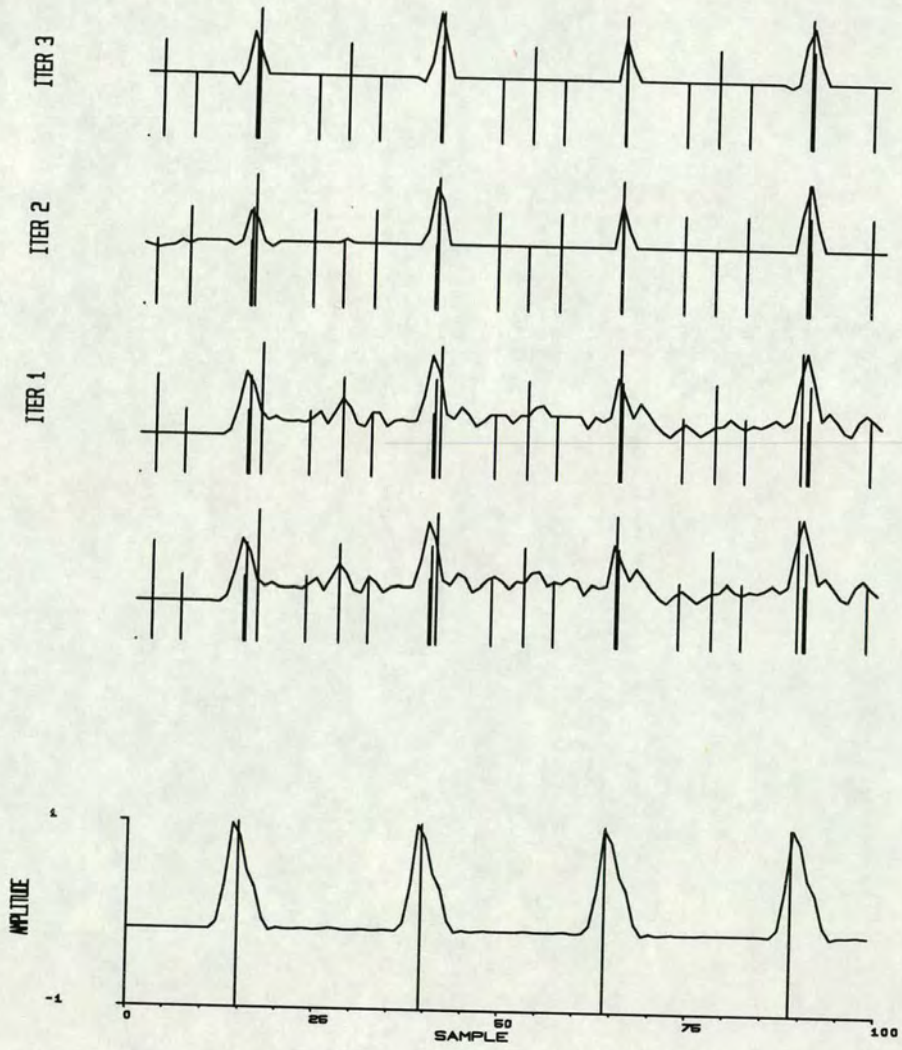

Figure 3.12 Bottom trace shows a noiseless rectified HPW time series (2 seconds). Noise is then added (SNR= 0dB, second trace up) and the three IBC iterations are applied with the resulting time traces ITER1, 2, 3 respectively. Superimposed on these traces are the impulse trains corresponding to the largest three peaks detected.

In step one of the iteration, the selection of an optimal threshold would require information about the signal amplitude as well as the noise level. In practice however, only an estimate of the HPW signal power obtained by using the total observed power. On average, using a normalised threshold value of 0.25 seems to give good results. Higher threshold settings may be required in cases where poor SNR signals are identified. The optimal number of iterations was found empirically to be 3 for good signal records and 4 for lower quality ones.

The IBC operation is illustrated pictorially in figure 3.12. The initial estimate of the fundamental (largest set of markers, trace labelled "ITER1") has both periodicity and delay errors. Such frequently observed bias would significantly increase the outlier rate in simple frequency domain estimators. Shown also (fig. 3.12) are the IBC second and third iterations (labelled "ITER 2" and "ITER 3"). These result in the gating out of most of the noisy time domain samples. The estimated FHR (or periodicity) is obtained from the impulse train with the most power after the third iteration. For comparison purposes this estimated impulse train is superimposed on the noiseless version of the signal (bottom trace fig. 3.12). The final estimated periodicity is 0.494s (simulated periodicity =0.5 s) and using "ITER 1" (equivalent to frequency estimation) would result in an estimated periodicity of 0.482 s. While the IBC can remove a substantial amount of non-coherent noise, nothing can be done to residual noise superimposed on the HPW peaks.

For a frequency domain peak found at bin k, the periodicity of the gating function T_k is expressed:

$$T_k = \frac{N}{k \, . \, f_S}$$

N is the FFT size and f_S is the sampling rate. The associated time delay ΔT_k:

$$\Delta T_k = \frac{\theta_k \cdot N}{2 \cdot \pi \cdot k}$$

θ_k is the phase of the peak at bin k, ΔT_k is rounded to the nearest integer. The performance of the IBC algorithm is investigated in more detail in chapter four.

3.2. Beat-to-Beat Heart Rate Estimators

Very little work on the detection of individual fetal HPW's is reported in the literature although some success has been reported with an artificial intelligence approach employed in the time domain[35]. The problem with time domain approaches is that the phase linearity of the signal conditioning filters is crucial if pulse timing accuracy is to be preserved. This is particularly relevant if inter-pulse and intra-pulse timings are required. Observations of the fetal HPW signals indicated that the pulse shapes are not time invariant. Therefore, any algorithm designed to extract beat to beat timings must make use of a linear phase conditioning filter. In addition, it is necessary to select the correct signal feature to mark the occurrence of an HPW event[7]. Examples of valid features are: the first peak, largest peak, pulse 'centre of gravity' or maximum rate of change within the HPW.

3.2.1. Pseudo-Inverse Filter

An original contribution to time domain beat-to-beat estimators, namely the Pseudo-Inverse Filter (PIF), is detailed in this section. The advocated new algorithm derives the weights of a partially equalising filter which improves the ability to detect HPW occurrences. The pseudo-inverse filter algorithm is based on the following model:

The assumption is that oscillatory waveforms observed at the maternal abdomen, identified as the first and second HPW's, emanate from the fetal heart as a train of unipolar impulses. In travelling through the fetal and maternal abdominal structures (organs) to the abdominal surface and through the transducer, these impulses suffer attenuation and dispersion (Fig. 3.13). This implies that an observed HPW pulse is the impulse response of this described model.

Based on such a model, the transfer function between the impulse train source and the abdominal surface has a non-flat amplitude and a non-linear phase response.

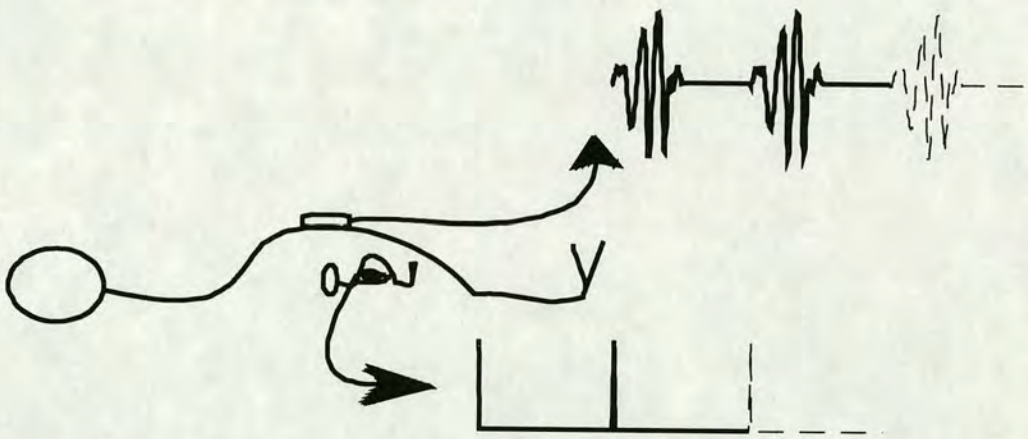

Figure 3.13 Illustration of both patient and fetus depicting the adopted model for HPW generation. Impulses initiated at the fetal heart (lower trace) are dispersed by the path to the transducer (upper trace).

Figure 3.14 shows a typical sampled first HPW reflecting the impulse response of the adopted model. While the second HPW is subjected to the same spatial filtering described above (through fetal and maternal organs), the underlying physical event has a different constitution. In consequence, both its time and spectral shapes are dissimilar to that of the first HPW (see chapter 2). Therefore a similar model but with

different dispersion characteristics must be assumed for the second HPW. The approach described here concentrates on extracting first HPW impulses from the maternal abdominal signals. A similar approach can be adopted for the second HPW extraction.

The advocated pseudo-inverse filter approach attempts to equalise both amplitude and phase responses of the first HPW. Therefore the reconstructed impulse train should, except for a processing delay imposed by the filter, pin-point the onset of systole events. While PIF processing is performed in the time domain, the derivation of the filter coefficients is achieved using a spectral estimate of the first HPW.

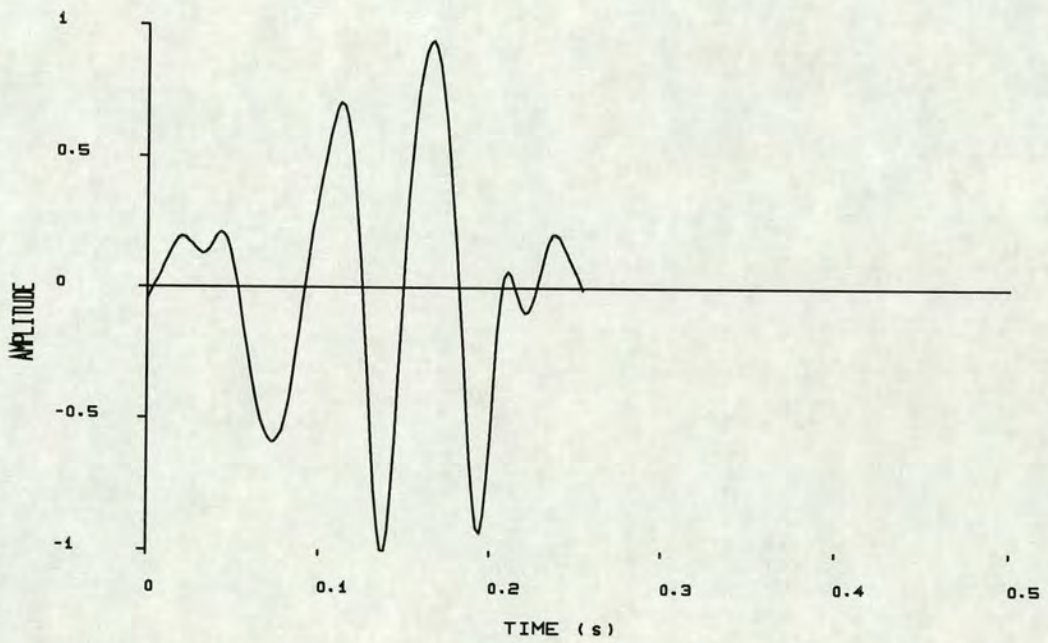

Figure 3.14 Example showing the assumed impulse response of fetal and maternal abdominal structures. This trace is obtained by averaging 10 consecutive first heart pulse waves.

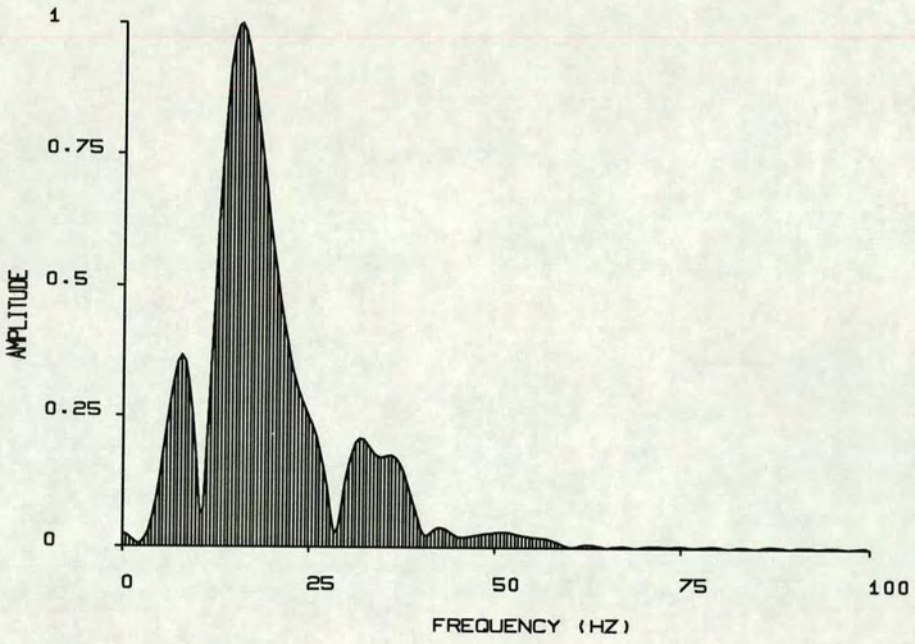

Figure 3.15 Linear spectral magnitude response of the assumed PIF model.

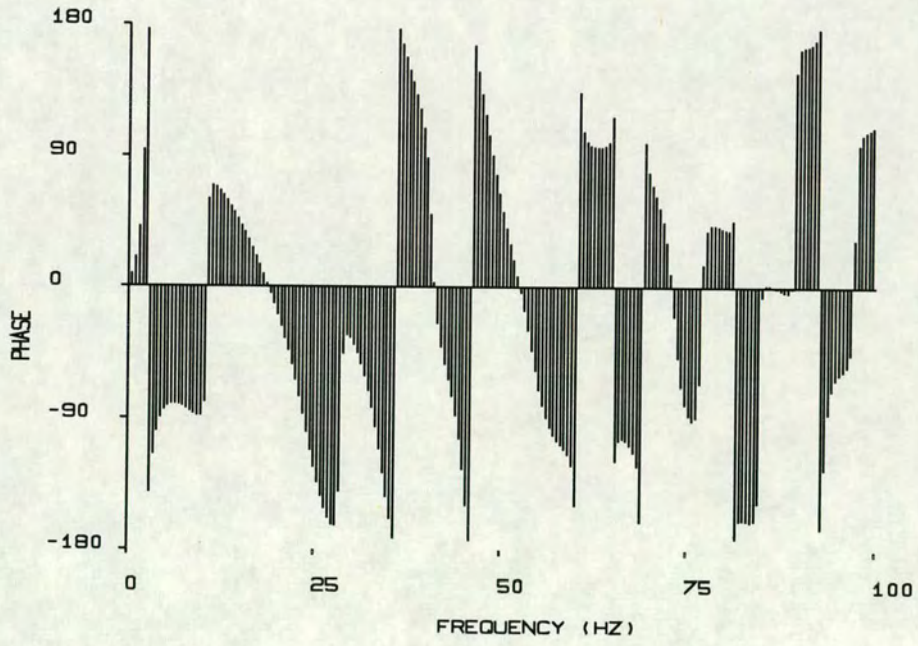

Figure 3.16 Phase response of the assumed PIF model.

In attempting to extract FHR on a beat-to-beat basis, signal processing algorithms can utilise two main approaches, correlation filtering or inverse filtering. All methods commonly construct a template which is used for subsequent processing. In the case of correlation, the sampled time series is correlated with the obtained template weights to provide a global peak per HPW event. In the inverse filter case, the coefficients are derived using the template. This requires a computationally intensive algorithm which is sensitive to both noise and spectral shape (Least squares detailed later).

The pseudo-inverse filter algorithm presented here provides a simpler and faster approach to obtaining an approximate solution to the inversion problem. The sensitivity of true inverse filters to signal non-stationarities can also be dealt with by programming robustness into the coefficient generation algorithm. The Pseudo-Inverse Filter algorithm is detailed:

3.2.1.1. Pseudo-Inverse Filter Algorithm

The first step in implementing this algorithm is to obtain a template of the first HPW - effectively a time averaged series (TAS) ("Template Forming" is dealt with later). The complex Fourier spectrum of this template is computed and both amplitude $|TAS(\omega)|$ and phase $Arg(TAS(\omega))$ vectors are extracted. The first vector is then normalised with respect to the largest component thus retaining only the relative amplitudes of the constituent components. The second vector, $Arg(TAS(\omega))$, provides the phase relationship that exists between the spectral components of the template. For the purposes of the model adopted, this time averaged series represents an averaged impulse response of the feto-maternal organ structures, and therefore its frequency spectrum shows the associated amplitude and phase distortions (Fig. 3.15

and 3.16). Averaging is necessary to avoid basing the model on a single HPW pulse which may contain either interfering signals (such as fetal movements) or excessive background noise.

In the case of a noiseless signal (consequently noiseless TAS template), the equalisation (deconvolution) is achieved by inverting the TAS spectrum (eq. 3.5).

$$H(\omega) = \frac{1}{TAS(\omega)^*} \qquad (3.5)$$

where $H(\omega)$ is the required inverse function and $TAS(\omega)^*$ is the complex conjugate of the Fourier transform of the template. A classical drawback with true inverse filters (eq. 3.5) is that they are only optimal in the complete absence of additive noise (where they are the Weiner optimum solution). A further limitation which is relevant to HPW signals, is the extreme sensitivity of inverse filters to non-stationarities involving low amplitude frequency components.

A compromise is proposed for the pseudo-inverse filter which takes into account inverse filtering, emphasis of background additive noise and mild signal non-stationarity. This is achieved by setting a threshold limit on the inverted spectrum of the template: Any amplitude falling below this predetermined threshold is not inverted, in fact it is suppressed even further. In order to avoid consequent large discontinuities in the pseudo inverse spectrum, a soft threshold transition between suppression and inversion is imposed on the transform. The resulting smooth pseudo-inverse spectrum is used to obtain the weights of an equalising FIR filter by inverse Fourier transformation. The smooth transition requirement between suppression and inversion is to limit the number of filter taps. A suitable function for a smooth transition over the decision boundary is the sigmoid function Sig() (eq. 3.6). The corresponding transfer function is shown in Figure 3.17. Before applying the sigmoid

function, the TAS(ω) elements are normalised with respect to the largest spectral component.

$$Sig(x) = \frac{1.0}{1.0 + e^{-x}} \qquad 3.6$$

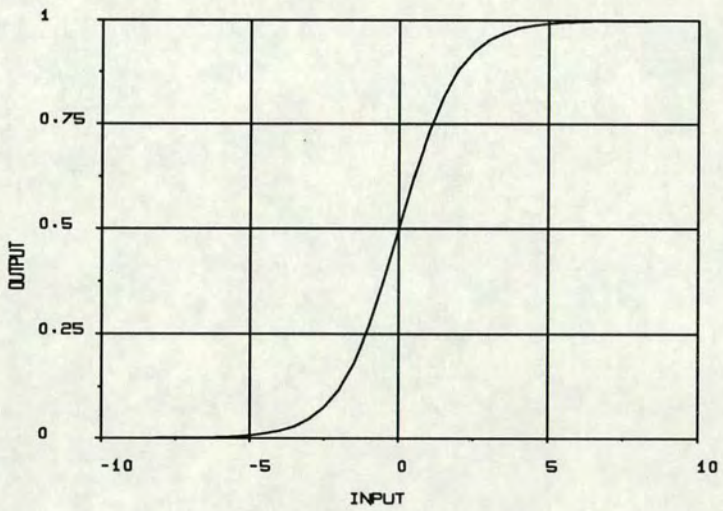

Figure 3.17 Sigmoid function mapping.

The equation of spectral pseudo-inverse mapping is given:

$$H(\omega) = \frac{Arg(TAS(\omega)^*) \,.\, Sig(G \,.\, (\,|TAS(\omega)| - T\,))}{|TAS(\omega)| + d} \qquad (3.7)$$

where,

Arg(): Argument function applied to the conjugate of TAS(ω);

Sig(): Sigmoid function employed to provide a smooth cut-off transition;

T: Threshold determining the position of the transition between inversion and suppression, it has a range of [0 .. 1];

G: Gain factor determining the slope (softness) of the sigmoid function, its useful range is [10 .. 90];

d: Small constant used to stabilise the denominator, typically d = 0.05

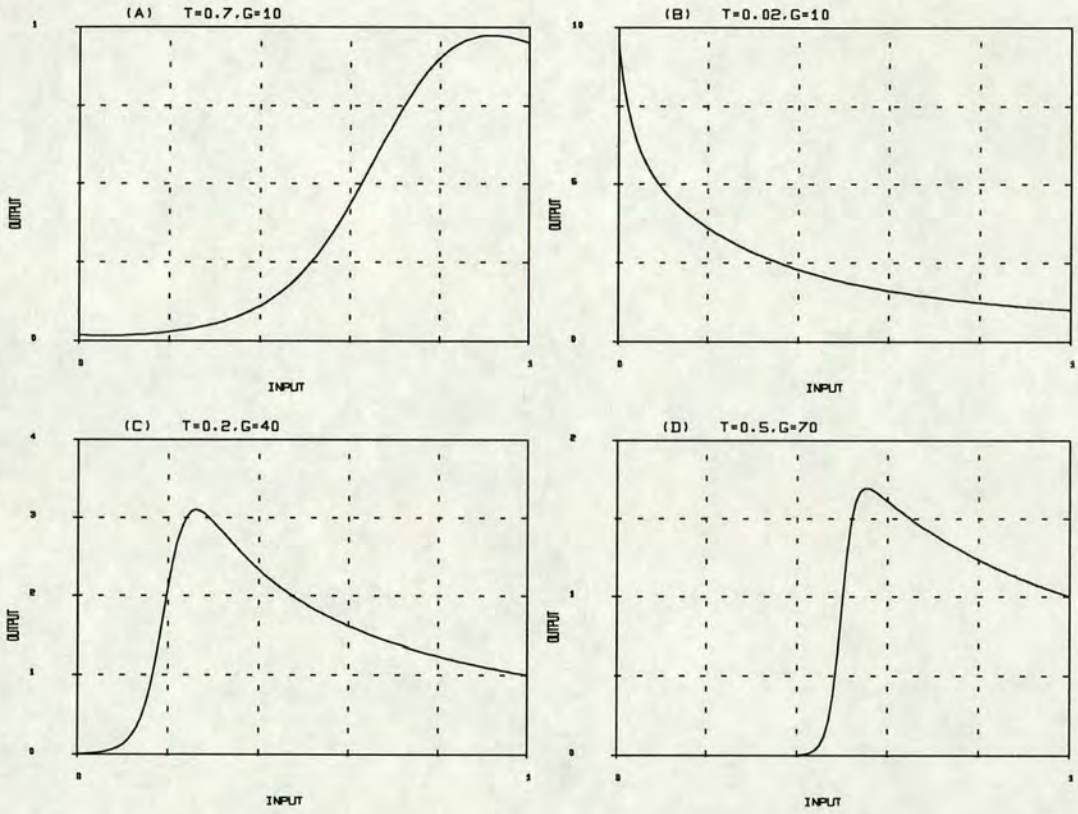

Figure 3.18 Typical magnitude mappings (eq. 3.7), d is set to 0.05. X and Y axis are the input and output magnitudes respectively for the different G and T parameter settings. The input amplitudes are normalised (range [0 .. 1]), and the Y axis have different scales.

Figure 3.18 illustrates four examples showing the effect of changing the T and G parameters on the magnitude mapping (eq. 3.7). For a medium value of T and a small value of G, for example 0.7 and 10 respectively (fig. 3.18.A), $|H(\omega)|$ approximates a correlator (which results from a 1:1 magnitude mapping slope). For small values of T and of G (0.02 and 10 respectively, fig. 3.18.B) the mapping leads to an

approximation to a true inverse filter. However, the above two examples are extreme cases, and normally a reasonably low threshold is selected (T=0.2) and a mid-range gain factor (G=40) (Fig. 3.18.C). This provides a mapping which effectively inverts all spectral amplitudes of the TAS template lying between 0.2 and 1, while spectral components with amplitudes lower than 0.2 are suppressed.

In cases of large interference noise components, the threshold T is raised accordingly (example: Fig. 3.18.D), although it will be shown in the following chapter that signals with high levels of noise are ultimately unusable for diagnostic purposes.

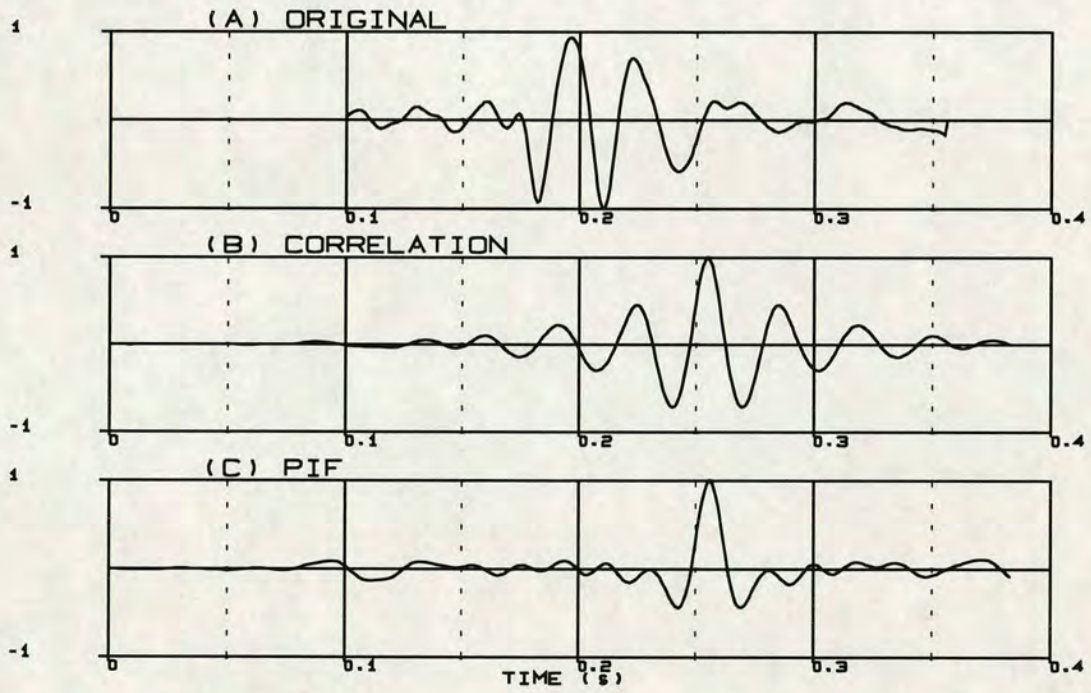

Figure 3.19 (A) Example of HPW pulse.
(B) Pulse correlation.
(C) Pseudo-inverse filter is used to provide partial equalisation.

Recall that H(ω) is the magnitude and phase of the desired (equalising) filter. Hence computing the inverse Fourier transform of H(ω) provides the impulse response of

the above filter (also the weights of its transversal filter realisation). These weights can now be convolved with the sampled HPW time series thus providing a partially equalised pulse train denoting the onset of the systolic intervals. Figure 3.19.C shows the PIF equalisation output in response to a single instance of noiseless first HPW signal. Some oscillations are still evident in the PIF equalised pulse (fig. 3.19.C). These are due to the partial equalisation of the signal spectrum. However, the PIF discrimination factor, in this case, is 8:1 compared to 2:1 for the corresponding correlator output (fig. 3.19.B). The discrimination factor is defined here as the peak impulse amplitude over the peak sidelobe amplitude.

The PIF equalised spectrum illustrated in figure 3.20.B, shows the smooth transition from equalised to suppressed regions of the spectrum. The spectrum of the equalised pulse (fig. 3.20.C) has linear phase characteristics.

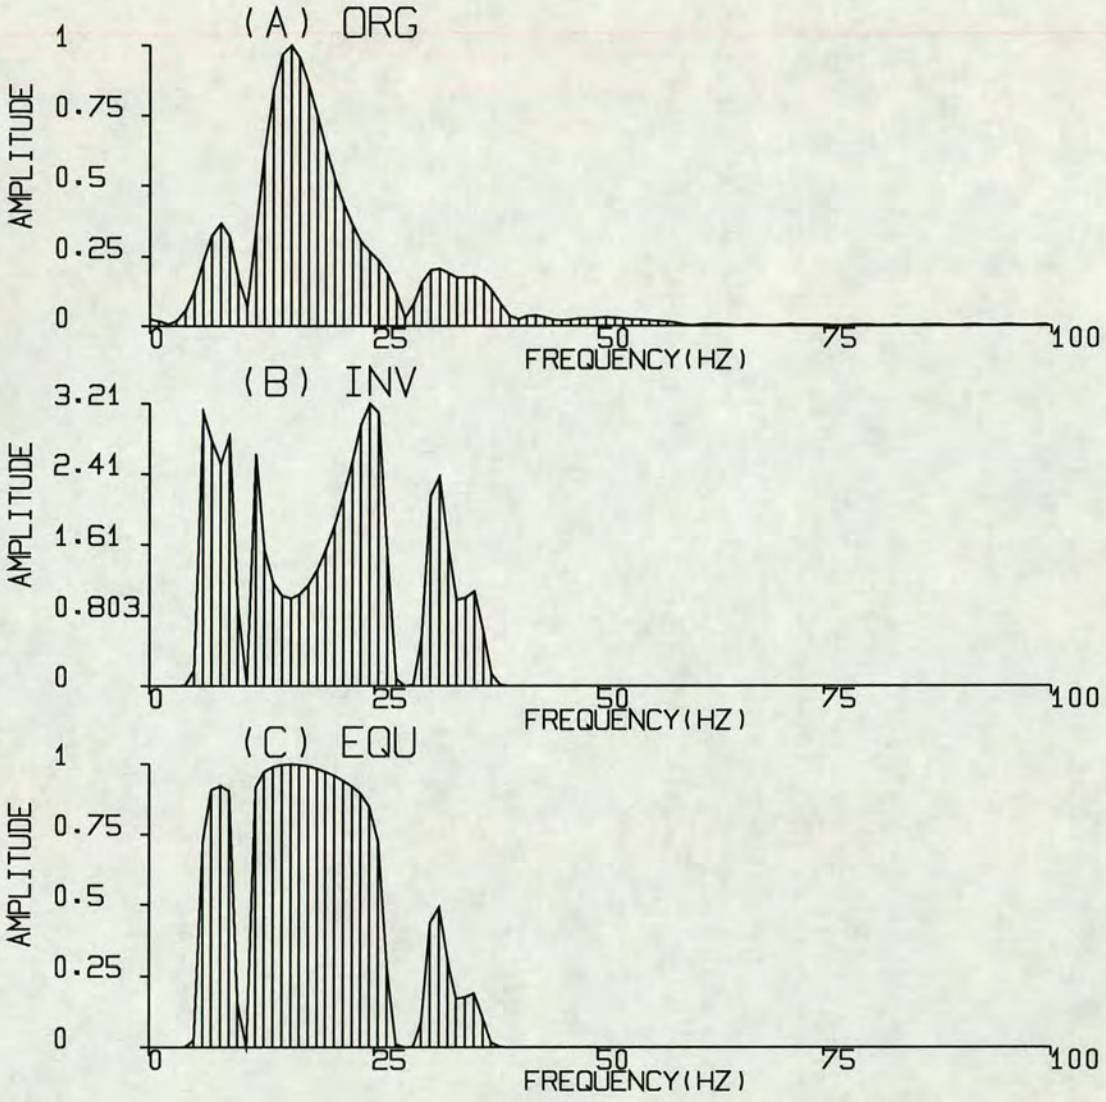

Figure 3.20 Top (A): Linear spectrum of the template; TAS(ω)
Middle (B): The pseudo-inverse spectrum; H(ω)
Bottom (C): Spectrum of partially equalised pulse.

Conflicting requirements of fully inverting the sampled HPW's, suppressing the interfering noise and limiting the number of PIF filter coefficients influence the setting of the T and G parameters. From experience of using the PIF algorithm, it was found that a gain factor G=40 provides a reasonably smooth transition from suppression to equalisation (hence limiting the spread of the PIF coefficients). It was

also found that a threshold setting of T=0.2 was large enough to take care of short term signal non-stationarities while leaving sufficient signal spectrum for equalisation therefore affording a better equalised impulse shape for a detection algorithm.

Applied to real sampled fetal HPW time series (fig. 3.21.A), the PIF equalisation technique demonstrates its capabilities in extracting the onset of systolic interval (fig. 3.21.B). A time delay due to the filtering is systematic and can be accounted for. The results obtained by correlation using the same TAS template are also shown (fig. 3.21.C).

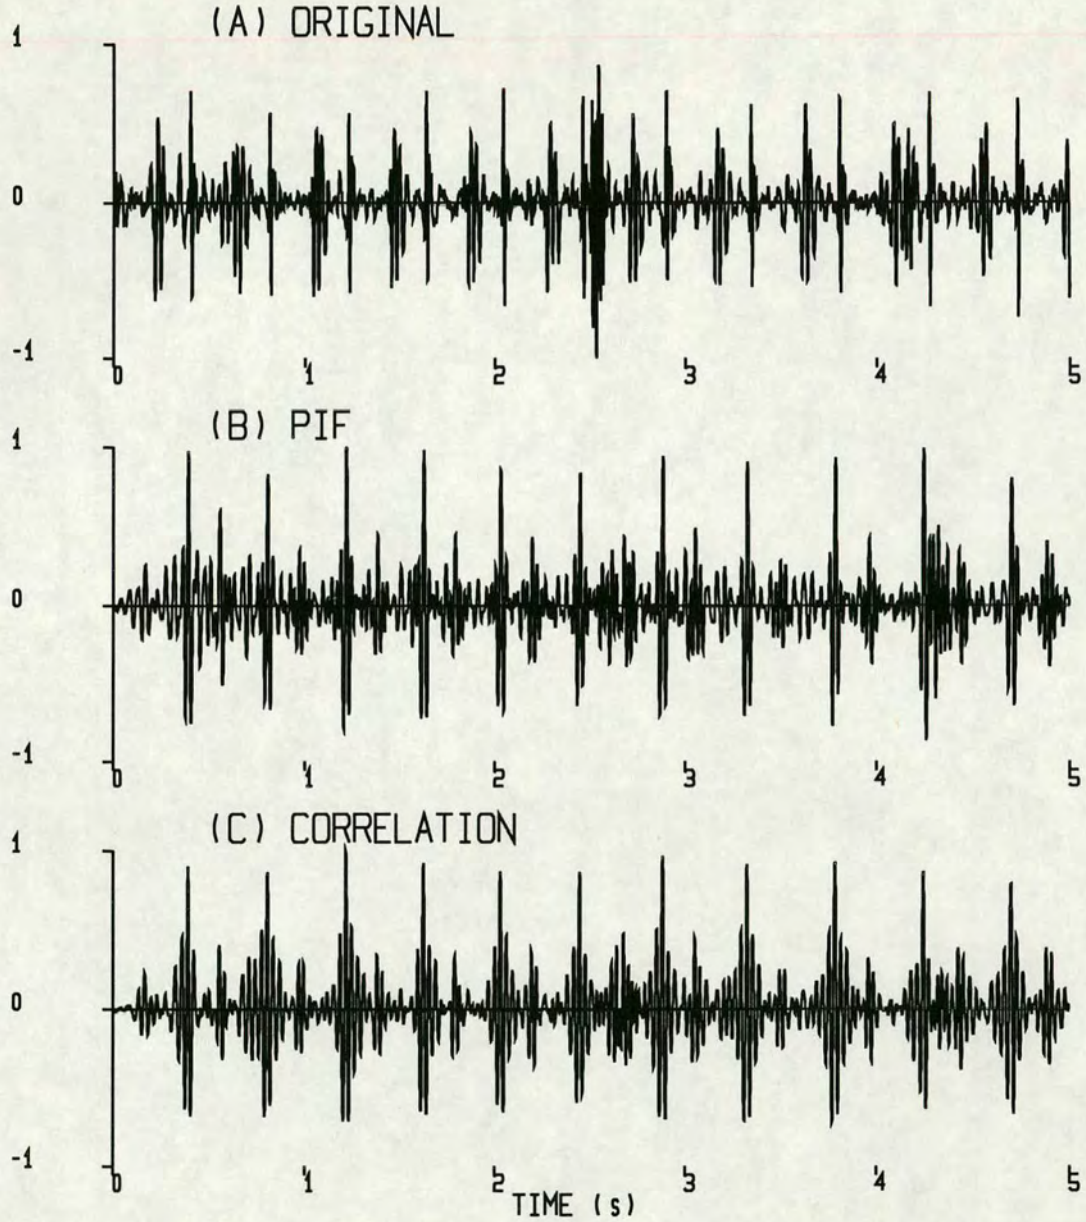

Figure 3.21 A real sampled fetal HPW signal (top trace), PIF equalisation of the first HPW (middle trace) and, correlation of the first HPW template (bottom trace).

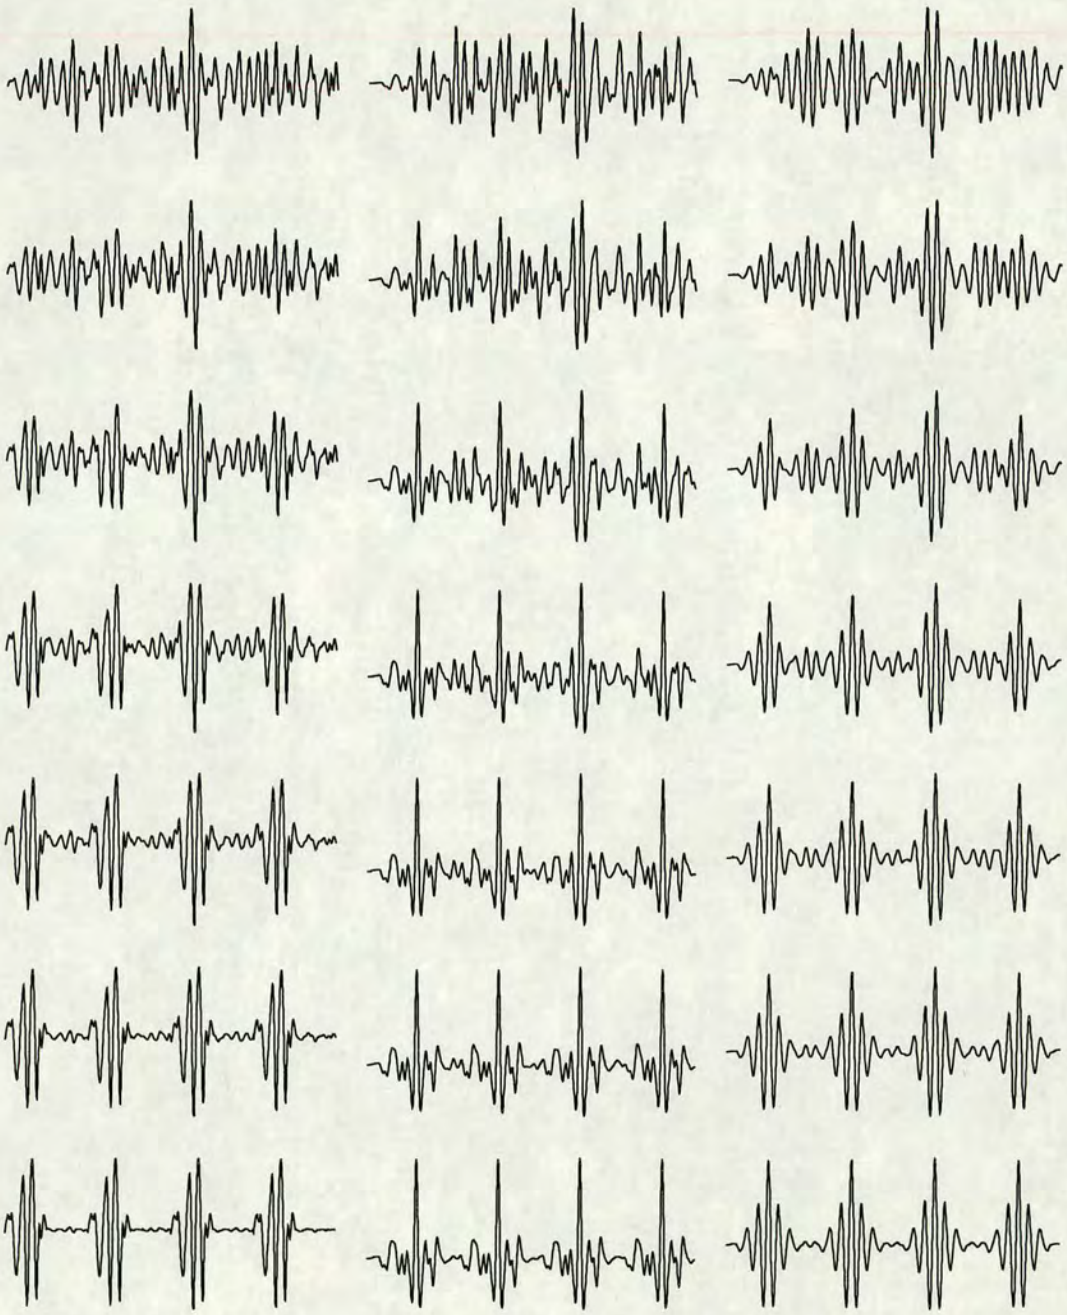

Figure 3.22 Simulated first HPW time series with increasing additive noise from bottom to top, the unprocessed trace is shown on the left side. The traces in the centre and on the right result from PIF and correlation processing respectively.

In order to justify the PIF equalisation approach, a pictorial subjective comparison with the correlation algorithm is included. This comparison uses a simulated HPW signal vector to which noise is added at progressively higher levels. Only a limited amount of pictorial examples can be shown, and therefore performance conclusions may not be drawn. A more formal simulation for the performance assessment of estimation algorithms is provided in chapter four. It is clear however from figure 3.22 that PIF provides single, prominent, local peaks for the detection of systole events. Correlation on the other hand, while improving the SNR, results in pulse spreading (ringing) when the TAS template is oscillatory in nature (which is normally the case).

3.2.2. Least Mean Squares Inverse Filter

Another possible equalising approach that utilises time domain filtering to extract the BTB timing from the sampled HPW time series is the inverse least squares (ILS) filter. The ILS algorithm uses the mean square error cost function[8] as criterion to formulate the coefficients of the inverse filter[9, 10]. A brief derivation of the algorithm is included, h(n) is the extracted filter coefficients vector:

Assuming an observed time series y(n) which is a dispersed version of the desired response d(n) (original signal), the ILS derived coefficients are used to filter y(n) in order to estimate d(n). Here $\hat{x}(n)$ and e(n) are the output estimate from the ILS filter and the error in the estimate respectively. In the fetal HPW case, y(n) represents the observed HPW signals at the maternal abdomen while d(n) is the assumed impulse train emanating from the fetal heart (according to the impulse generator model in figure 3.13). The objective here is to get a best estimate, $\hat{x}(n)$, of the heart beat impulse train of the assumed model. To this end the coefficients vector $\underline{h}$ of the inverse filter is derived:

$$\hat{x}(n) = \sum_{i=0}^{N-1} y(n-i)\, h(i)$$

Where N is the number of coefficients in the equalising filter.

$$e(n) = d(n) - \hat{x}(n)$$

$$= d(n) - \underline{h}^T \underline{y}(n)$$

The mean square error is:

$$\xi = E\{ (d(n) - \underline{h}^T \underline{y}(n))^2 \}$$

$$\xi = E\{ d(n)^2 \} + \underline{h}^T \underline{R}_{yy} \underline{h} - 2 \underline{h}^T \underline{R}_{dy} \qquad (3.8)$$

Where $\underline{R}_{yy}$ is the NxN autocorrelation matrix of the observed HPW signals, $\underline{R}_{dy}$ is the NxN cross-correlation matrix of y(n) (observed HPW signal) and d(n) (the desired signal). Minimising the quadratic equation 3.8 w.r.t. $\underline{h}$, leads to:

$$\underline{h}_{opt} = \underline{R}_{yy}^{-1} \underline{R}_{dy}$$

$\underline{h}_{opt}$ is the least squares solution to the filter weights coefficients. Substituting $\underline{h}_{opt}$ back into eq. 3.8 leads to the least mean squares error (LMSE) solution ξ_{opt}:

$$\xi_{opt} = E\{ d(n)^2 \} - \underline{h}_{opt}^T \underline{R}_{dy}$$

When $\underline{d}$, the desired response is a unit impulse, ξ_{opt} simplifies to:

$$\xi_{opt} = 1 - \underline{h}_{opt}^T \underline{R}_{dy} \qquad (3.9)$$

The resulting equation (3.9) is a performance indicator of the computed $\underline{h}_{opt}$. Three factors effect the performance of the designed ILS filter, in other words the attained level of LMSE (ξ_{opt}), these factors are:

a - The shape of desired impulse: Relaxing the constrains on the shape of the desired signal leads to an improved ξ_{opt} when the observed signal y(n) is composed from relatively narrow band components. The tradeoff here is between the duration of the desired pulse, $\underline{d}$, and the noise bandwidth of the equaliser.

b - The impulse position within the desired vector $\underline{d}$ is particularly important for non-minimum phase response systems for which the ILS filter is acausal[9]. The solution to this class of systems is not to restrict the position of the impulse (in vector $\underline{d}$) to the first sample. There does not exist a methodical approach for selecting the optimum impulse position in vector $\underline{d}$. The adopted solution, therefore, is to iteratively introduce different impulse delays and then to select the filter weights producing the smallest least square error, ξ_{opt}. This optimal delay was found to be 3N/4. In addition, ξ_{opt} proved (experimentally) to be insensitive to small differences in impulse position when the filter length was long.

c - The third of these factors is the filter length (N). This is related to classical FIR/IIR filter design where a truncated FIR filter is selected in order to approximate an IIR filter. This is achieved by setting the FIR coefficients to the IIR impulse response, the more coefficients are used, the better the approximation. Therefore if the system to be equalised contains zeros near unit circle in its transfer function, the corresponding inverse IIR filter would contain poles, this would lead to a long ringing impulse response and hence a large number of taps in the FIR realisation. In the case of equalising the first HPW, it was found that a filter length about twice the HPW pulse duration achieved a satisfactory level LSE (ξ). Increasing the filter length

further presents a case of diminishing returns as the other two factors (listed above) come into effect.

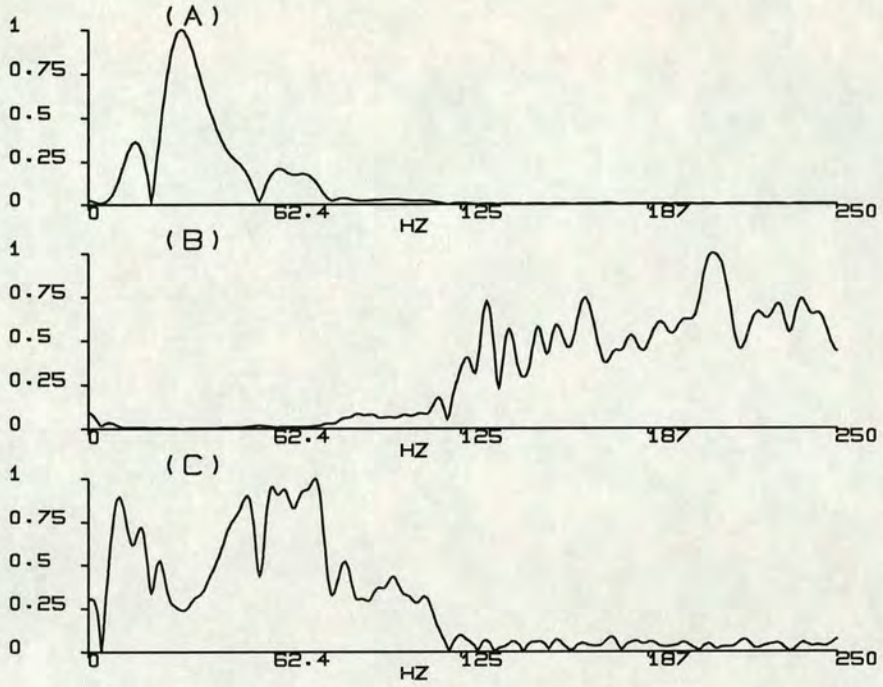

Figure 3.23 (A) The linear amplitude spectrum of system impulse response (same as that used for the PIF template)
(B) Spectrum when ILS attempts to equalise all spectral components.
(C) Equaliser spectrum when white noise of variance 0.2 is introduced into the auto-correlation matrix.

When ILS is applied to equalising the first HPW, the signal vector $\underline{y}$ is the same 64 samples, time averaged template used for the PIF realisation. The filter length therefore is set to 128 and the delay is set to 100 samples (200ms). For the purposes of this application, different pulse shapes were tried for the desired response, it was found a that an impulse response was adequate. Slightly better results were obtained when the desired response was set to the sinc function central lobe. The performance gain, however, was deemed not worthwhile. The correlation matrices $\underline{R}_{yy}$ and $\underline{R}_{dy}$

are formed and $\underline{h}_{opt}$ is then calculated. Recalling that vector $\underline{y}$ is an estimate of the impulse response of the feto-maternal system, the ILS algorithm assumes that the auto-correlation matrix $\underline{R}_{yy}$ includes the noise component, which implies apriori knowledge of the latter. Non-stationarities of both noise and signal render this assumption invalid. A compromise suggested here is to assume that the interference noise is white and has variance σ (normalised w.r.t. signal power). Adding the white noise term to the autocorrelation matrix results in:

$$\underline{R}_{yy} = \begin{bmatrix} r(0) + \sigma & \cdots & r(N-1) \\ \cdots & & \cdots \\ r(N-1) & \cdots & r(0) + \sigma \end{bmatrix}$$

where

$$r(k) = \sum_{i=0}^{N-1} y(k)\ y(k+i)$$

Figures 3.23.A shows the spectrum of the feto-maternal impulse response. The middle trace (fig. 3.23.B) shows the equaliser magnitude response when σ=0. At approximately 200 Hz, the ILS filter gain is over 3 orders of magnitude larger than the gain at 30 Hz where most of the signal power lies. When σ is set to 0.2, the equaliser magnitude response is shown in figure 3.23.C. It exhibits only a partial spectral inversion due to suppression of signal spectral components of low amplitude. Further level increases in assumed noise variance (σ) leads to the evolution of $H(\omega)$, the spectrum of $\underline{h}$, from the ideal inverse magnitude response (fig. 3.23.B) to the spectrum associated with the matched filter with:

$H(\omega) \approx Y(\omega)$; $Y(\omega)$ is the spectrum of $\underline{y}$.

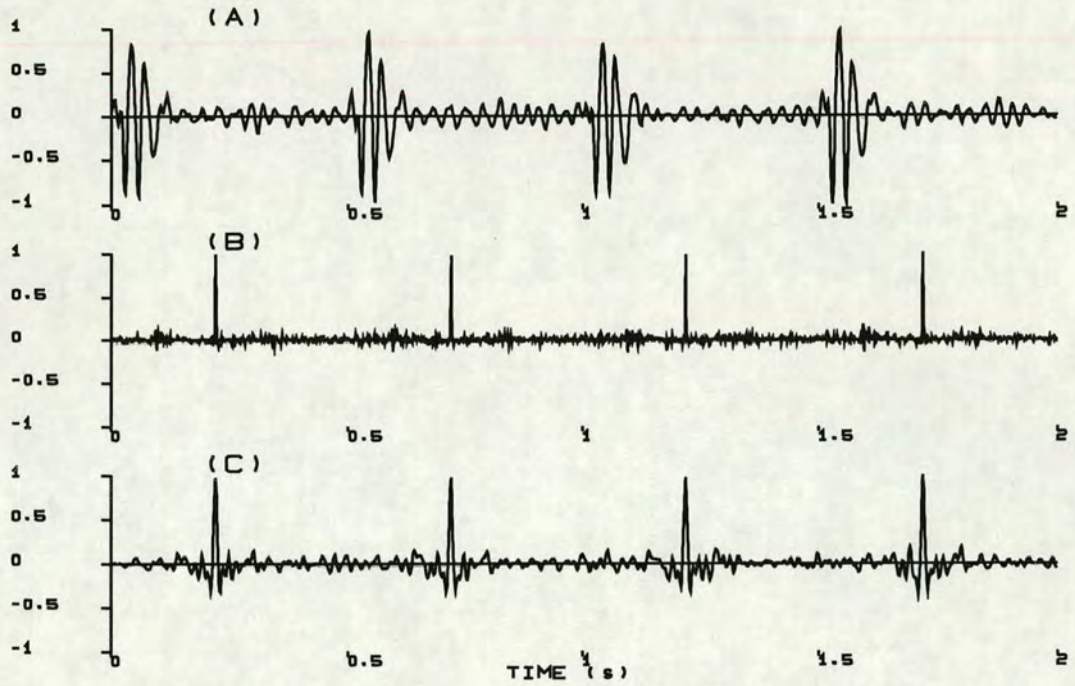

Figure 3.24 (A) Simulated signal having SNR= +20dB.
(B) Result of ILS filtering when the noise variance (σ) is set to zero.
(C) The bottom trace results from ILS filtering when σ=0.2.

When the two ILS filters derived above (with σ=0 and σ=0.2) are applied to a simulated HPW time series (fig. 3.24.A) to which coloured noise has been added at -20dB, traces (B) and (C) in figure 3.24 result respectively. It can be seen that when all of the spectrum is inverted, the obtained time series closely approximates a true impulse train function. As the level of noise is raised, only the portion of the spectrum containing the most power is equalised, consequently, some degree of pulse ringing is introduced (due to band-limited equalisation) (fig. 3.24.C). This latter case (σ=0.2) is comparable to PIF filtering.

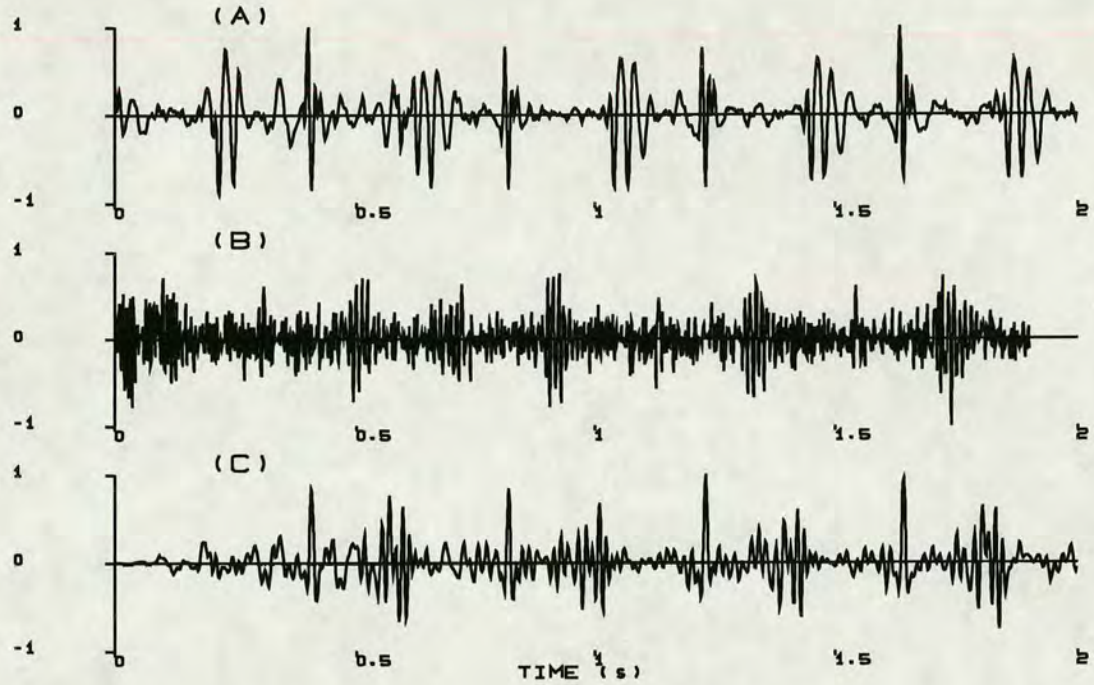

Figure 3.25 (A) Real HPW signals.
(B) Result of ILS filtering trace (A), assumes noise variance σ=0.0.
(C) Result of ILS filtering trace (A), assumes σ=0.2.

A real sampled fetal HPW signal (fig. 3.25.A) is used to test the ILS performance: Figure 3.25.B illustrates the result of ILS filtering when the assumed white noise variance (σ) is set to zero, figure 3.25.C is obtained when the assumption σ=0.2 is made. In using the averaged template to derive the filter coefficients, two processes have taken place: Firstly the broad-band HPW noise was cancelled by averaging, and secondly any residual noise in the template is indistinguishable from the genuine signal. The resulting performance illustrated in figure 3.25.B is unacceptable as most of the filter output pertains to wideband noise which has been over-amplified. As a result of assuming a white noise variance of 0.2, a substantial reduction in the wide-band noise is achieved (fig. 3.25.C). The inverse least squares filter solution in this example, however, remains unsatisfactory. Some additional improvement can be

obtained by raising the assumed white noise level σ. Unfortunately, the ILS approach requires too much expert intervention which cannot be easily automated.

3.2.3. Template Forming for Fetal Heart Pressure Waves

Many techniques can be applied for the selection of HPW pulses in order to form a template. The IBC method explained earlier, can for example be used to identify the HPW pulse positions within a 2 seconds vector. Other methods use correlation techniques or energy criteria[35] to identify the occurrence of HPW's. These methods however, require considerable artificial intelligence to be built into the peak selector. Described here is a robust technique which methodically selects HPW peaks while requiring little additional artificial intelligence.

In order to construct a useful template, it is necessary to select a portion of the signal (5-10 seconds long) which has a high SNR. Using an envelop detector, the signal amplitude time series is obtained. Fetal HPW's are assumed to lie under the peaks extracted from this time series. In addition, in the case of good quality signal records, the HPW related peaks are expected to have large amplitudes compared to other peaks present in the signal.

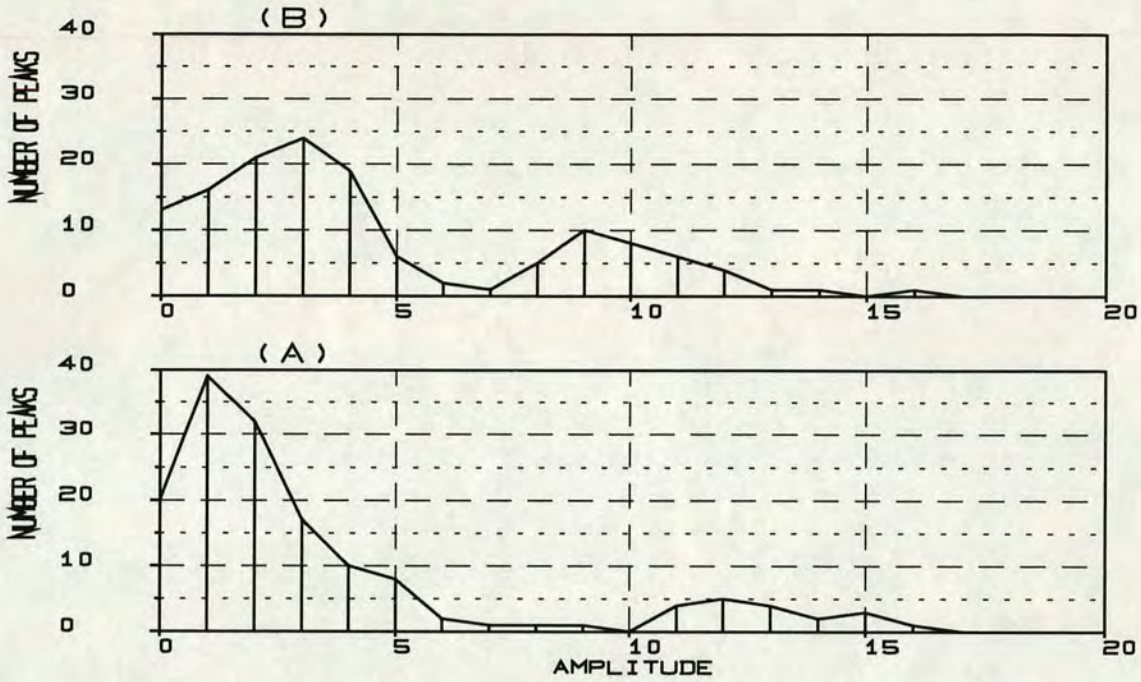

Figure 3.26 Two examples of peak distribution versus peak amplitude taken from 10 seconds records.

When the number of peaks versus their corresponding amplitudes are plotted, typical distribution shapes like those seen in figure 3.26 are obtained. The most relevant feature of these distributions is the commonly observed 'hump' at the higher end of the peak amplitudes. This feature mainly consists of HPW peaks. Two possible approaches used separately or combined can make use of the feature to determine the HPW peak events. These two approaches are:

a - All the peaks contributing to the hump feature of the peak amplitude distribution (fig. 3.26) are used to form the required template. The above feature can be seen from amplitude 10 to 17 in graph (A) and 7 to 15 in graph (B). The number of peaks thus obtained for trace (A) is 19, and the number manually counted HPW events is 22. In the case of trace (B), the corresponding numbers are 37 and 24 respectively. Therefore, only three errors have occurred in case (A) and 13 in case (B). Close inspection of the

peaks in trace (B) however, shows that 10 of these errors are related to the detection of two peaks for the same HPW event. This problem is easily rectified by identifying two closely spaced peaks.

b - The HPW peaks are assumed to occupy the distribution graph at the higher amplitude end. Taking the integral of the distribution function from right to left, i.e. from A_{max}=20 in both examples (fig. 3.26) and to an arbitrary lower amplitude A, results in an enclosed number of potential HPW's. The constraint on this integration is that the value of the integral should result in a number of HPW pulses which must correspond to an acceptable FHR range (i.e. 60-240 BPM). The integral is given:

$$FHR(A) = \frac{-1}{T} \cdot \int_{A_{max}}^{A} \psi(a)\ \delta a$$

and for the discrete case:

$$FHR(A) = \frac{1}{T} \cdot \sum_{A}^{k=A_{max}} \Psi(k) \qquad 3.10$$

Where FHR(A) is the estimated fetal heart rate when all peaks having amplitudes between A and A_{max} are considered as valid HPW's, T is the record length (seconds), ψ and Ψ are continuous and discrete peak distributions respectively. By varying A in equation 3.10, an acceptable FHR (1-4 Hz) with a corresponding A_{acc} is obtained. Most peaks encompassed in the amplitude range [A_{acc} .. A_{max}] can then be assumed to belong to valid HPW pulses. In practice, setting the acceptable FHR to 1.5 Hz (90 BPM) provides good results. At the worst case, this setting excludes 63% of the real HPW peaks when the true FHR is 4 Hz (240 BPM), and includes 33% erroneous peaks when the true FHR is 1 Hz (60 BPM).

A combination of the above methods can also be used as follows: an attempt at identifying the hump feature is made and the corresponding FHR is computed. If the obtained FHR is within a reasonable range (1.5 - 3 Hz) then all peaks enveloped by the hump feature are accepted as likely HPW events. If the obtained FHR is too high, the lower amplitude limit A is increased (resulting in estimated FHR decrease) until a 3 Hz rate is achieved otherwise A is lowered (resulting in estimated FHR increase) until the rate equals 1.5 Hz.

Some additional tidying up can be beneficially added to the methods described above. This involves deleting from the distribution graphs (fig. 3.26) all samples in amplitude bins which contain very few peaks. When this is applied to graph (A), it results in the deletion of all peaks having amplitudes 6,7,8,9 and 16, and when applied to graph (B), all peaks having amplitudes 6,7,13,14,16 are deleted.

Presumed HPW pulses pin-pointed by the previously selected peaks are then extracted, from the raw HPW signals (pre-rectification), into vectors of length L. L is set slightly larger than the required template length. The next process attempts to reject any pulses which are too 'different' from the averaged computed norm.

The cross-correlation $R_{i,j}(\tau)$ is computed for each vector pair of the selected pulses. $R_{i,j}(\tau)$ provides a global peak which denotes (a) the degree of correlation between the pair of vectors i and j, and (b) a shift adjustment $\tau(i,j)_{peak}$ which correctly aligns vector j with respect to vector i. A new figure which effectively measures the degree of correlation between an individual vector, the i^{th} for example, and all other vectors can be expressed as:

$$V(i) = \frac{1}{N-1} \sum_{j=1,\, j \neq i}^{N} R_{i,j}(\tau(i,j)_{peak})$$

$R_{i,j}(\tau(i,j)_{peak})$ is the maximum cross-correlation value between vectors i and j. V(i) is the cross-correlation of vector i with a template made-up from equally weighted combination of the other vectors. Those vectors with lower correlation values, V(i), are dismissed as 'false', 'noise contaminated' or 'non-conformant' instances of HPW pulses. Having removed these, a new value of V(i) is computed for each of the remaining vectors. A correcting shift factor $\delta\tau(i)$ is then computed, this correction provides a better reference point for each of the vectors.

$$\delta\tau(i) = \frac{-1}{2.(N-1)} \sum_{j=1}^{N} \tau(i,j)_{peak}$$

The template is then initialised to the vector which has a maximum V() value and built up by progressively from the remaining vectors in descending order of V(i). Each vector is added to the template using the shift which gives maximum correlation with the template.

CHAPTER FOUR

4. PERFORMANCE COMPARISON OF SEVERAL FHR ESTIMATORS

This chapter introduces some estimating algorithms and defines the performance criteria which are used to compare the investigated estimators. Two main performance criteria, the narrow-band noise performance and the tracking ability, are established then employed to extract the worthiness of each estimator. The susceptibility to broad-band noise is also checked and only the results of algorithms which were found to have significant differences in the performance are included. Similarly, real signals are utilised to detect some of the pitfalls associated with the estimators.

Two major factors play a role in determining the limitations of an estimator, the first factor relates to the performance deterioration when the fetal phono signals are subjected to increasing levels of noise. The second factor is the estimate degradation brought about by a changing fetal heart rate.

4.1. Simulations Design

In order to accurately measure the performance of the estimators detailed in this chapter, it was decided to simulate the fetal phonocardiographic signals. This enables

the generation of numerous HPW trace realisations with precise control over both the additive noise level and the heart rate.

Similar performance measurements are neither practical nor possible when using real fetal HPW signals. The difficulty is attributed to:

(a) Providing recorded phono signals with known SNR and heart rate (and accelerations).

(b) The requirement of a 'gold standard' signal which can be used to simultaneously check the accuracy of the estimates. This involves obtaining simultaneously recorded scalp FECG[32, 40, 70, 71], although this itself present additional difficulties[7, 40].

For the purposes of the simulations, both the signal and random additive noise are held stationary in order to facilitate the drawing of statistical significance. While similar stationarities cannot be confirmed from real observations of HPW signals, short term stationarity can be assumed when processing real data. In practice a specific performance will pertain to epochs of quasi-stationarity of the relevant factors.

In order to facilitate the reproducibility of results, all analogue signal conditioning steps are emulated digitally. Analogue processing modules such as full wave rectifiers and filters are cascaded to provide the equivalent analogue circuit functions.

4.1.1. Accuracy Considerations

Although the FHR is commonly expressed in beats per minute (i.e. in rate), in the case of time domain based estimators it was found easier to compute the periodicity of FHR. While this facilitates the simulations measurements below, real

implementations of time domain estimators convert periodicity estimates into heart rates (frequency domain). The conversion (taking into account the quantisation error) is given by:

f_{est} is in the range $\dfrac{1}{T_{hr} + \dfrac{T_s}{2}} \;..\; \dfrac{1}{T_{hr} - \dfrac{T_s}{2}}$

where f_{est} is the estimated FHR (Hz), T_{hr} is the measured FHR periodicity (s) and $\dfrac{T_s}{2}$ is the quantisation error. The above equation is rearranged to give:

$$= \frac{f_{hr} \cdot f_s}{f_s + \dfrac{f_{hr}}{2}} \;..\; \frac{f_{hr} \cdot f_s}{f_s - \dfrac{f_{hr}}{2}}$$

and now calculating the error difference in f_{est}:

$$\Delta f_{est} = \frac{4 \cdot f_{hr}^2 \cdot f_s}{4 \cdot f_s^2 - f_{hr}^2}$$

$$\approx \frac{f_{hr}^2}{f_s} \qquad \text{for} \qquad f_{hr}^2 << f_s^2$$

Therefore the time domain quantisation translates into a frequency domain quantisation error which is FHR dependant; this error is worse for high FHR's. The sampling rate for TD estimators is usually set to 500 Hz, and some resulting quantisation figures are given: 0.002Hz (0.12 BPM), 0.008Hz (0.48 BPM) and 0.032 (1.92 BPM) at FHR rates of 1, 2 and 4 Hz respectively.

Conversely, translating frequency estimates to time domain periodicities results in an estimate quantisation:

$$\Delta T_{est} = \frac{4 \cdot T_{hr}^2 \cdot \Delta f_s}{4 - T_{hr}^2 \cdot \Delta f_s^2}$$

Where T_{hr} is the heart rate periodicity, Δf_s is the frequency domain quantisation. For frequency domain block processing the sampling rate is 50Hz, the FFT size is set to 2048 for the simulations below, therefore Δf_s = 0.024 Hz (1.44 BPM). The above equation is then approximated by:

$$\Delta T_{est} \cong T_{hr}^2 \cdot \Delta f_s \qquad \text{when: } T_{hr}^2 \cdot \Delta f_s^2 << 4$$

The larger quantisation errors this time are obtained at low FHR's.

4.1.2. Phono Signal Composition

In order to realistically simulate fetal phono signals, a good phono recording which contains clear heart pulse waves is used. A template is formed from the first HPW pulses contained in the chosen record (see "Template Forming" chapter 3). The number of pulses forming the template is not critical since the character of the simulated HPW pulse is, generally speaking, a minor factor in determining the performance of the estimators. The resonance character of the pulses is chosen to represent the majority of observed fetal HPW's, a first HPW containing 2 or 3 oscillations is typical. In practice, as long as the selected phono record does not contain highly resonant HPW signals, the performance of the estimators tested below is not compromised.

4.1.3. Simulated Noise Synthesis

As detailed earlier (chapter 2), raw (unfiltered) phonocardiographic signals contain large low frequency components. Although some of these low frequency signals (down to heart rate itself) are expected to contain fetal heart components, this region of the spectrum is swamped by interference caused by body and breathing movements. The amplitude of these interfering signals diminishes progressively at higher frequencies, and are in practice filtered out. Other broad-band interfering signals, are also commonly observed in the raw signal and therefore the issue is reduced to determining both the centre frequency and bandwidth of the fetal heart pulse waves. In most cases it is necessary to model the phono interference noise only the region of interest, that is within the bandwidth of the employed band-pass filter.

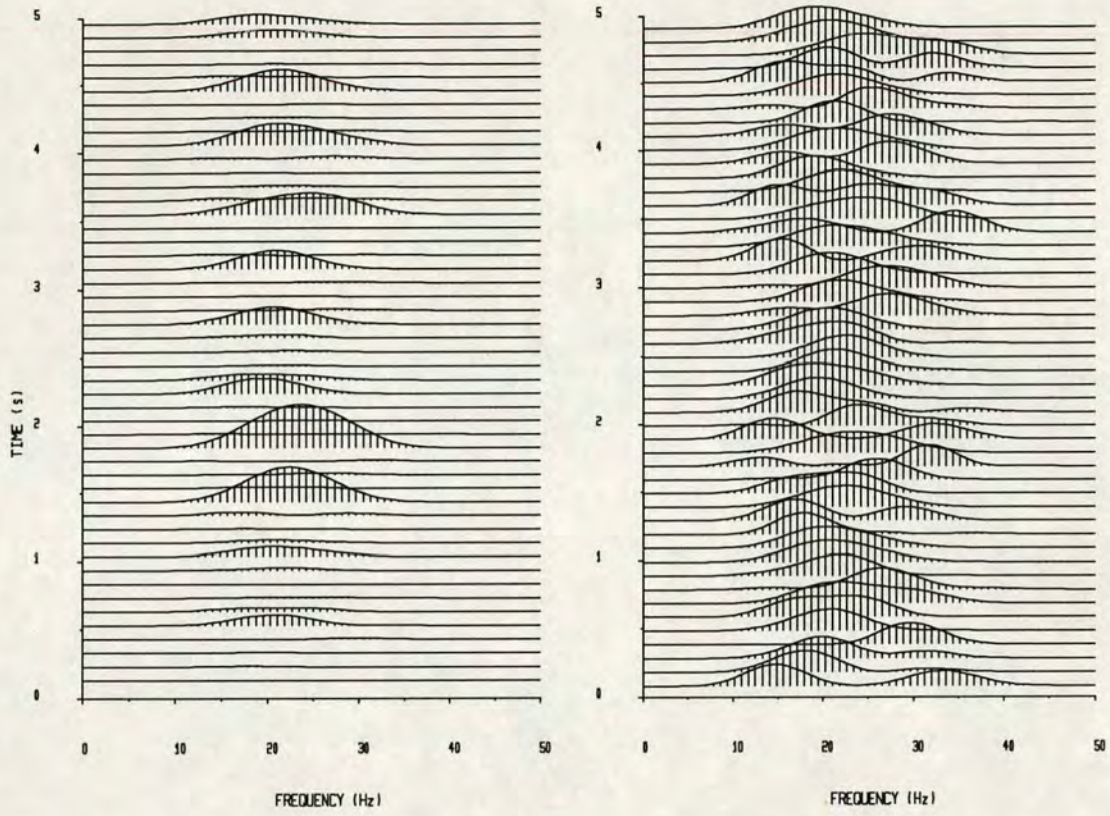

Figure 4.1 Power spectral waterfalls of the same HPW phono signal lasting 5 seconds (pertaining to trace A of figure 4.3). The spectral traces are obtained from 175ms consecutive time records to which the Hanning window is applied, the records overlap is set to 75ms. The waterfall traces on the right are auto-scaled (visually), this has the effect of exposing the spectral shape of the background noise.

Close inspection of many instances of fetal phono signals have revealed that the spectrum of the interference noise in the region of interest closely matches (in shape) that of the HPW signals. Shown in figure 4.1 (left) is a typical spectral waterfall of fetal HPW signals, this is selected from a high quality signal (corresponding to trace A fig. 4.3) to illustrate the following point: All of the traces of the waterfall diagram on the left have the same y-axis scaling, the spectral peak features (of which 11 can be seen) clearly identify the occurrence of the first HPW pulses. These peak features are inter-spaced by periods of low signal power. Hidden within these periods is the background interference noise which can be exposed by allowing each of the

waterfall traces to visually auto-scale (waterfall on the right of figure 4.1). Inspection of the latter reveals that within the region of interest (i.e. within the pass-band), the interference noise appears to be subjected to the same selective body resonances as the HPW pulses.

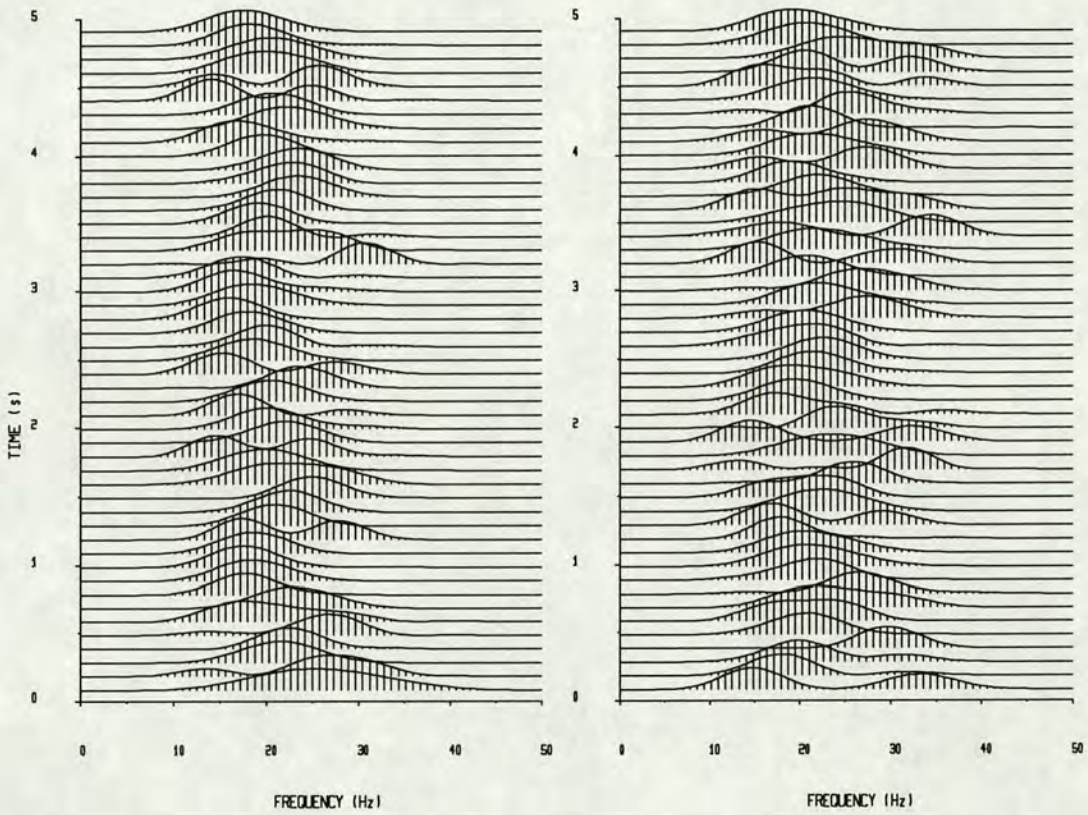

Figure 4.2 Left: Power waterfall spectrum of simulated interference noise.
Right: Power waterfall spectrum of real phono signals.
Both waterfall diagrams are obtained from 175ms consecutive time records to which the hanning window is applied, the records overlap is set to 75ms. The traces are visually auto-scaled.

It is therefore possible to model the phono interference noise by a white Gaussian noise source which is filtered using coefficients derived from the averaged HPW spectrum. The power level of the simulated noise can then be accurately controlled

before adding to the HPW time series. When compared to real signals, simulated noisy HPW signals look very similar (fig. 4.2). They do however differ in one major aspect: The simulated coloured noise lacks the dynamic variance observed in real phonocardiographic signals. In other words, simulated phono noise possesses long-term second order stationarity which cannot be assumed in the case of real phono signals. For simulations purposes however, the noise power stability is essential in obtaining valid performance measurements.

4.1.4. Signal to Noise Ratio Definition

Estimating the SNR for phonocardiographic signals is neither a trivial task nor a precise one. Record lengths of between 5 and 10 seconds were found most appropriate since quasi-stationarity assumptions may be invalid for longer epochs.

Observed HPW pulses can vary from clear, close to impulse shapes, to the more commonly observed 'ringing' pulse shapes which can last 200ms or more. Impulse shapes can aid the signal separation from noise while oscillatory shapes can complicate the separation. In addition, inter-pulse shape variations, some due to noise others related to genuine pulse shape evolutions are often noted. These consequently pose additional complications to the HPW signal pulse identification. Therefore, in order to facilitate HPW component separation, all non persistent inter-pulse changes are classified as noise components. As a general rule, while BTB estimators are prone to errors caused by HPW shape changes, block estimators using front-end envelope detection circuits are relatively immune to both small temporary and large evolutionary pulse shape changes.

The fetal heart rate contained in the record from which SNR is estimated, directly effects the available HPW power and in consequence the measured SNR. While the

heart rate has no effect on the performance of BTB estimators which depend on the HPW pulse power versus the background noise power, it has a proportional relationship with the SNR as seen by block estimators: I.e. doubling the FHR produces an equivalent 3dB increase in both signal power and SNR.

In order to avoid possible confusion caused by either of the two factors described above, it was decided to adopt the following simulation conventions for all the FHR estimators under test:

a - Given a constructed template, the pulse shape is fixed for the duration of the simulation. In these simulations, temporary pulse shape changes are modelled by additive coloured noise. The deviation between simulated and real HPW signals is only apparent at high SNR where the extent of inter-pulse changes for simulated HPW time series are much smaller than real HPW signals. However, this is not considered to be a severe restriction since high SNR phono signals are rarely observed.

b - When measuring the noise performance of various estimators, a fixed 120 BPM fetal heart rate is used.

In the case of real phonocardiographic signals, interfering noise is defined as the power of the background signal which excludes the fetal HPW components. This signal can be measured in vivo by moving the phono transducer away from the optimal pick-up region.

4.1.5. Noise Estimator

During FHR monitoring, a good awareness of both the error rate and FHR variance introduced by the estimating algorithm is crucial in guaranteeing the integrity of the

FHR trace. As shown later in this chapter, both error rate and variance depend on the prevalent background noise, therefore, the medium term performance can be deduced statistically from the estimated SNR. In the case of phonocardiographic signals, the difficulty is in attaining a reliable SNR estimate. This difficulty is exacerbated in cases of high feto-maternal body resonances which leads to continuous ringing of the HPW's. Trace (E) (of figure 4.3) is such an example of medium/high feto-maternal resonance where it is difficult to distinguish the HPW signal components from those related to background noise. Traces (A) and (F) are examples of very high and very low quality recordings respectively. Proposed below is an automated SNR estimation method which can be implemented on stand-alone monitors and does not require human intervention.

Figure 4.3 Six records, each of five seconds duration, of fetal phono signals showing a typical spread of heart pulse wave signal qualities. These are arranged from best to worst signal qualities (top trace to bottom). The respective φ values for the traces are: $\varphi(A)=1.40$, $\varphi(B)=1.29$, $\varphi(C)=1.19$, $\varphi(D)=0.99$, $\varphi(E)=0.89$ and $\varphi(F)=0.84$.

Close inspection of the signals shown in figure 4.3 reveals that the better quality traces (such as A and B) have distinct HPW pulses with clear (low amplitude) intermediate epochs. Progressively poorer signal qualities (traces D, E and F) have a persisting oscillations between the HPW pulses. Although in some cases, these continuous oscillations are associated with HPW pulses observed in patients with high body resonances, these signals are classified as being of poor quality. Moreover such HPW signals cannot be distinguished in the short term from additive noise. A primary prerequisite of an SNR estimator is the ability to distinguish between the impulse-like (good quality) signal and the noisy or highly oscillatory signal. This is achieved by

making use of the clear inter-pulse period observed in good quality signals. The proposed function $\varphi(\underline{x})$ is given:

$$\varphi(\underline{x}) = \frac{SD(\underline{x})}{E\{\ |x()|\ \}}$$

where $SD(\underline{x})$ is the standard deviation of the observed vector $\underline{x}$. For good front-end signal conditioning, an AGC is normally employed in order to maintain reasonable signal levels. In such a case, $E\{\ |x()|\ \}$ can be assumed be to constant, therefore $\varphi(\underline{x})$ depends on the standard deviation of the signal vector. $SD(\underline{x})$ has a global minimum value for equally distributed $|x()|$'s which increases progressively for impulse based vectors $\underline{x}$.

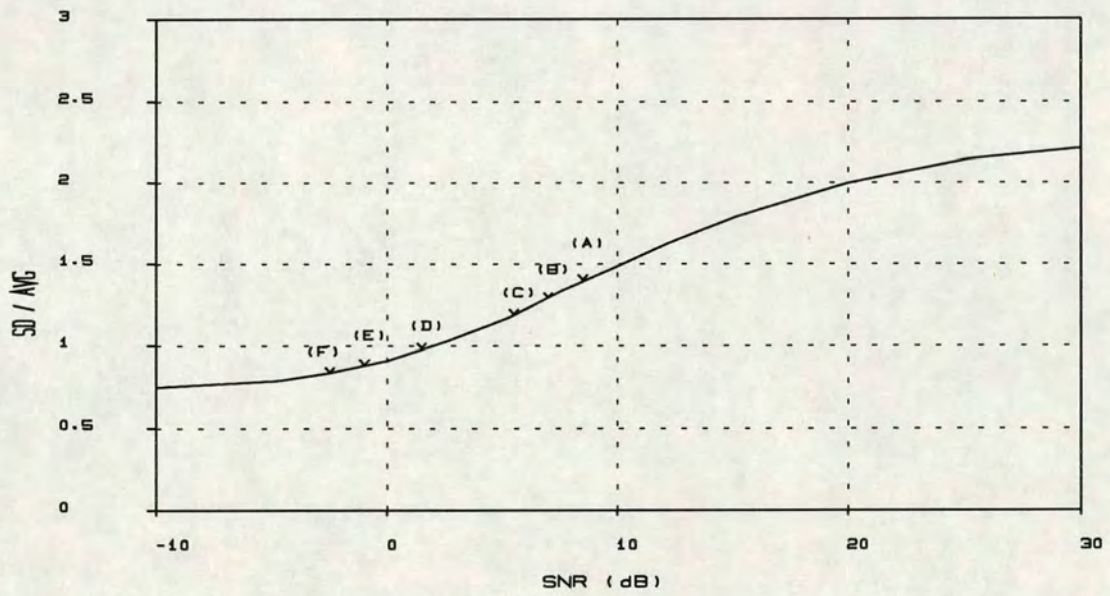

Figure 4.4 Plot of φ = (Standard Deviation / Average Magnitude) for different signal to noise ratios. The useful range for φ is [0.8 .. 2] corresponding to the SNR range of [-2 .. 20] dB. Superimposed are the φ results for the traces of figure 4.3.

The function φ() is applied to a range of simulated SNR values resulting in the mapping shown in figure 4.4. As shown, φ() provides a good discrimination factor in the critical SNR region of [0 .. 5] dB. Other likely functions such as the (maximum amplitude/average magnitude) were also tested. These were found to either fluctuate wildly, in cases of spurious noise, or not to provide enough discrimination at the critical SNR range where most observed signals are found to lie.

Superimposed on figure 4.4 are the φ() mappings applied to the signal vectors in figure 4.3. Traces A, B and C qualify as good HPW records, traces D and E are of poor quality while trace F is of unacceptable signal quality.

4.2. Estimation Algorithms and their Performance

An unbiased evaluation of the performance of different estimators is achieved by employing the same simulated HPW time series for all estimators. Some of the estimators, such as the pulse correlation algorithm, make use of the raw HPW time series while other algorithms, such as the Comb Filter, require an envelope detected version of the signal.

The tracking capability of an estimator and its noise performance are decoupled by setting a fixed FHR and changing the SNR or fixing SNR and varying the FHR. The performance criteria for the varying SNR case are discussed first:

4.2.1. Noise Performance Measurement

In engineering terms, it is desirable to evaluate at different SNR's, the underlying error rates and FHR variability falsely introduced by the estimation process[31]. The

simulation results presented here provide performance measures which can be used to assess particular estimating algorithms under given operating noise conditions. The different SNR's are obtained by mixing coloured noise of the correct amplitude with the ideal (noiseless) HPW time series.

Given an expected FHR and corresponding periodicity T_{hr}, an estimate inclusion band, $2.\delta T$ wide ($[T_{hr} - \delta T .. T_{hr} + \delta T]$), is defined which results in two performance criteria:

a - The Outlier Error Rate (OER) performance: Determines the fraction of erroneous FHR estimates when operating under a specific SNR (expressed as fraction or percentage). These errors are defined as the estimates falling outside the defined inclusion band limits, i.e. those T_{est} periodicity estimates outside the range $[T_{hr} - \delta T .. T_{hr} + \delta T]$.

b - The In-Band Standard Deviation (IBSD) performance: This is a measurement of the standard deviation of all estimates falling in the defined inclusion band, i.e. all estimates within the above range.

Given an operating SNR, the OER is a figure which must lie within an acceptable predetermined error rate limit, otherwise the signal quality is deemed too poor for monitoring purposes when using this particular estimator. While a standard upper limit for tolerated OER does not exist, it was decided to set an arbitrary limit of 10% for the purposes of this work. Similarly, IBSD varies with the operational SNR and has an upper limit which must not be exceeded by the estimator. To be of any use, the IBSD limit is set well below the true FHR variability which is medically reckoned to indicate fetal well-being. Typically, a variability of 10 BPM is accepted as a good fetal health indicator[48], and for the purposes of performance assessment, an IBSD threshold of 5ms (1.2 BPM) is selected. In cases where only long term baseline FHR

is of interest and the employed estimator is unbiased, it may be possible to ignore the IBSD performance.

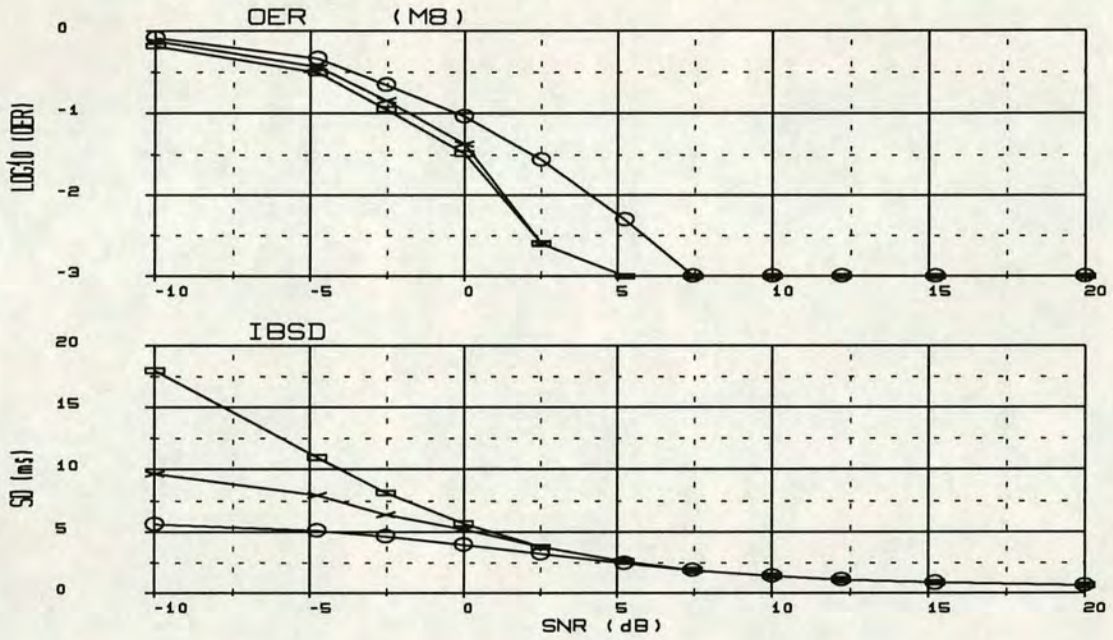

Figure 4.5 A typical estimator performance for a range of operating SNR's. The top graph gives the erroneous estimate rates outside the 10ms, 20ms and 40ms bands (circles, crosses and rectangles respectively). The corresponding standard deviation of estimates within the three bands (10, 20 and 40 ms) are shown in bottom graph.

If the inclusion band is widened when operating at a specific SNR, the outlier error rate can be expected to fall as more estimates are gathered in the inclusion band, this consequently ensues in a rise of the in-band standard deviation.

Typical measured performance graphs are shown in figure 4.5. The top graph shows the OER under increasing (left to right) SNR, likewise the corresponding IBSD (bottom graph). The same performance graphs are drawn for all the investigated estimators below. Three periodicity inclusion bands are depicted in the results of

figure 4.5, these are: (a) the ±10ms band which is marked in circles and has an equivalent FHR band width ±2.4 BPM, (b) ±20ms band is marked in crosses (FHR band ±4.5 BPM) (c) ±40ms band is marked in boxes (FHR band ±9.5 BPM). All bands are centred on 500ms periodicity (120 BPM). Frequency domain based FHR estimates are converted to corresponding fetal heart periodicities.

In the case of beat-to-beat algorithms, FHR estimates are obtained from the position of the processed HPW's within sectioned 0.5 second blocks. In order to avoid signal discontinuities, the expected HPW position is arranged to lie away from the edges of the 0.5 second blocks. For frequency domain based block algorithms, FHR estimates are obtained by searching through the valid FHR range (1 to 4 Hz). For time domain block algorithms, the periodicity is estimated from a search which covers heart periodicities from 0.25 to 1 seconds.

4.2.2. Tracking Performance Measurement

The biophysical fetal heart rate control system dynamically reacts to both external stimuli and internal conditions of the fetus[3]. As a result, the FHR time series is composed of: (a) baseline rate (long term) which is regularly subject to sustained and sometimes sudden accelerations (or decelerations), and (b) short term, small amplitude FHR variations. The latter components are present continuously in FHR traces of healthy fetuses and have minimal degradation effects on the performance of most estimators.

Block estimators effectively perform FHR averaging over the observation period, a block containing sudden and large FHR acceleration can corrupt the estimator output resulting in an increase of the expected OER. Hence, it is necessary to assess the tracking performance of the estimators when subjected to different acceleration rates.

For all the block estimators detailed below, the same results are obtained irrespective of whether the heart rate change is an acceleration or a corresponding deceleration. Consequently, in this chapter, the term acceleration does not distinguish between increasing or decreasing FHR.

When discussing tracking performance, beat-to-beat estimators are excluded because the detection of each HPW occurrence is independent of the recent FHR history. As a consequence the tracking performance of BTB estimators is the same as that obtained from constant FHR.

Four acceleration rates are used to measure the tracking performance of block estimators (2, 5 ,10 and 20 BPM/second). FHR is initialised to 240 BPM (upper FHR range) at the start of the simulation cycle and decelerated until 60 BPM is reached (lower FHR range). The cycle is repeated until enough estimates are obtained to draw statistical conclusions. Coloured noise at 0dB (as previously described) is then added to the time signal. As seen later, this chosen level of simulated SNR is close to the operational limit (SNR wise) of most block processing algorithms covered below. Finally, FHR estimates straddling the 60-240 BPM discontinuity in the deceleration cycle are discarded, and the difference between the estimated and the known FHR average is calculated for each block to reveal the errors in the estimator output.

4.2.3. A Simple Time Domain Estimator

This estimation approach emulates the old generation of fetal heart monitors which use time domain threshold triggering. A variety of analogue circuits such as trigger masking were commonly employed to restrict the number of false pulses assumed to pertain to fetal heart activity. For simulation purposes however, no attempt is made to

build any intelligence into this threshold detection algorithm: A single peak is extracted from blocks of 0.5 seconds duration.

Monitors based on time domain threshold detection are already known to have poor performances[26, 71]. Dawes[71, 73] provided extensive clinical comparisons of these with correlation based monitors and reported that on average threshold techniques have 4-8 times higher FHR signal loss. The signals used for these comparisons were however derived from the Doppler ultrasound returns which can be highly non-stationary. Slightly better results might be anticipated when this simple technique is applied to phono signals.

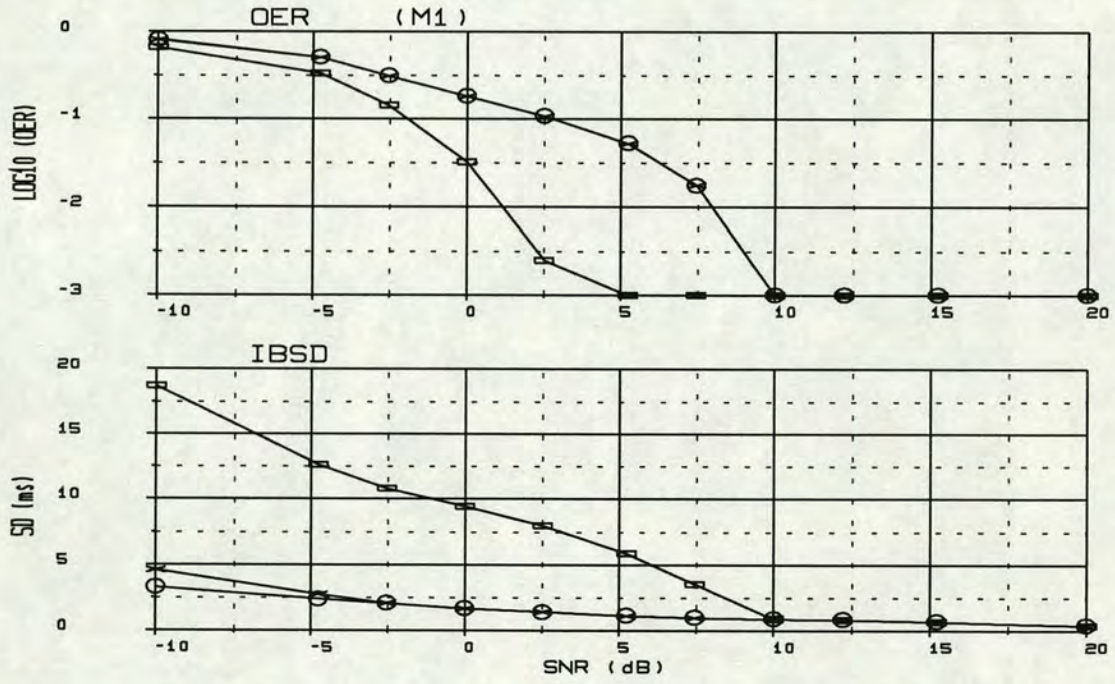

Figure 4.6 Performance of the 'Time Domain' estimator. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Compared to other estimators investigated later, the performance of this simple time domain estimator is found to deteriorate quickly with increasing noise (fig. 4.6). For example at 5 dB SNR (considered good quality phono signals), a sizeable pulse peak jitter is evident. This manifests itself as: (a) A high error rate when considering both ±10ms and ±20ms bands or, (b) a high IBSD which is too close to observed biophysical FHR variability when considering the ±40ms band.

4.2.4. Frequency Matched pre-Filter

In this approach, the HPW time series is processed using a frequency matched filter whose cut-offs were placed at the -3dB points of the template spectrum. For this particular HPW simulated signal, the 3dB points are located at 26 and 38Hz. The peak position is then determined from sectioned blocks of 0.5 seconds duration.

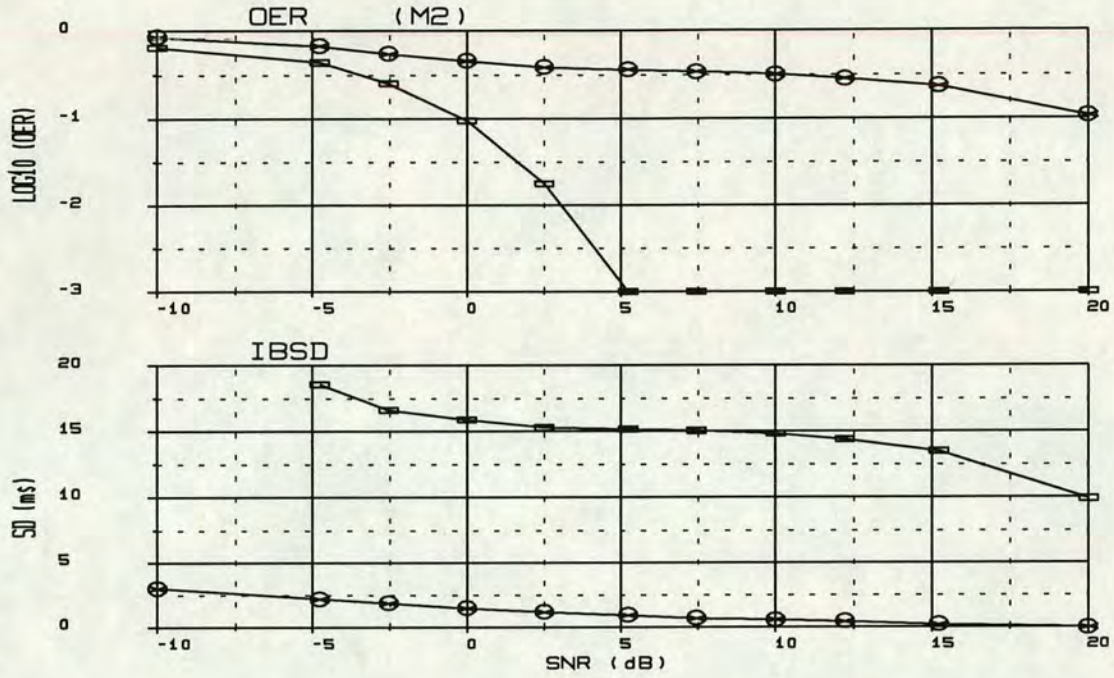

Figure 4.7 Performance of the 'Frequency matched pre-filter'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The performance obtained from band-pass filtering the signal is so poor (fig. 4.7) as to render it useless for short term FHR estimation. Figure 4.8 shows the band-pass filtered signal when the SNR is set to 10dB. The two prominent peaks of the displayed time series aided by the additive noise alternately possess the global peak thus resulting in step changes of the estimated fetal heart periodicity; the three most frequent periods observed are 468ms 500ms 532ms, 32ms being the time separation between the two prominent peaks. This global peak swapping explains the measured poor performance, even at high SNR. The Inter-pulse periodicity although frequently wrong, does not produce when averaged over a long period a biased error. Estimates thus obtained may therefore be used for FHR baseline long-term monitoring.

Some commonly observed fetal phono signals have similar or higher pulse resonances to that shown in figure 4.8 without the use of band-pass filters. In the context of

ultrasound, Manning[38] noted the problems associated with narrow band filtering: These include spurious noise 'ringing' which has a character indistinguishable from real fetal heart signals, and the lack of a consistent signal feature which can be used for pulse identification.

Figure 4.8 Band-pass filtered (26-38Hz, Q=2.7) simulated HPW time series (SNR=10dB).

Although a better performance can be obtained from this instance of simulation by selecting the most negative fetal HPW peak, the gain in performance is quickly eroded as the noise power is increased (fig. 4.9). The error rate of the ±40ms and the IBSD of the ±10ms and ±20ms bands are unchanged, these results are inherent to this type of estimator. Usually, the choice of a better peak 'polarity' is not available resulting in an average performance which is slightly worse than that shown in figure 4.9.

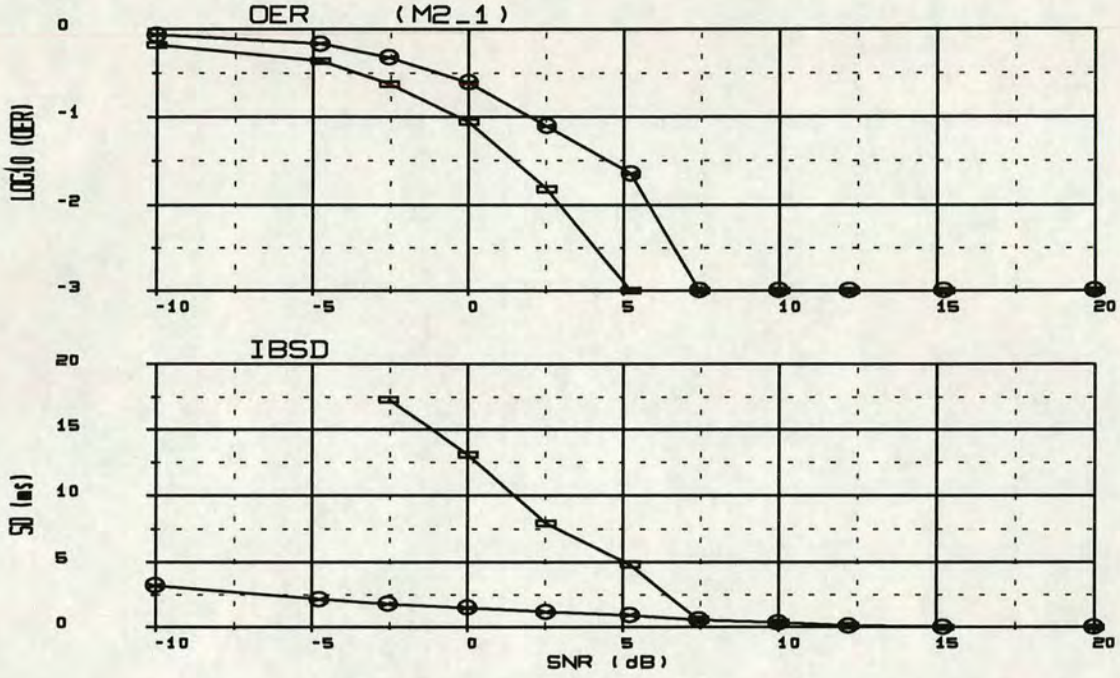

Figure 4.9 Performance of the 'Frequency Matched pre-Filter' when the most negative peak is selected. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

4.2.5. Template Pulse Correlation Estimator

The formed template vector $\underline{y}$ is correlated with the sampled HPW time series x(n).The resulting discrete correlation equation is given:

$$R_{xy}(n) = \sum_{k=0}^{L-1} x(n+k) \cdot y(k)$$

Where $R_{xy}(n)$ is the cross-correlation function and L is the template length. The correlator output is expected to be oscillatory, but with global peaks inter-spaced by the heart rate periodicity T_{hr}. The position of the global peak within 0.5 second blocks is assumed to pin-point the occurrence of the first HPW event.

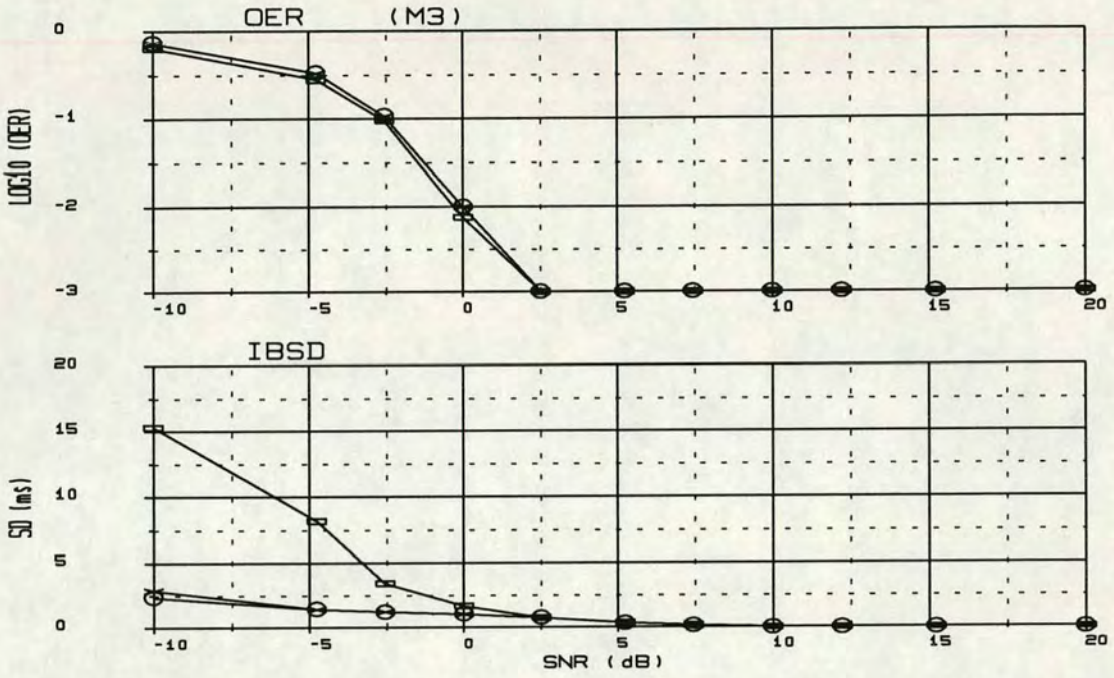

Figure 4.10 Performance of the 'Template Pulse Correlation Estimator'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

A good performance is obtained from pulse correlation (fig. 4.10). The algorithm can be applied to signals down to -2.5dB SNR, conditional on a 10% error rate being acceptable. Pulse correlation can easily be implemented on a stand-alone monitoring system. Its computing power requirement is proportional to the selected template length plus some overhead for control and template forming. Assuming a reasonable template quality can be obtained, the performance of the pulse correlation estimator is minimally effected by the presence of white noise of variance less than -10dB (w.r.t. signal).

4.2.6. Pseudo-Inverse Filter Based Estimator

Pseudo-inverse filtering was introduced in chapter 3 where it was shown subjectively to offer an improvement in the detection of HPW occurrences. The algorithm parameters gain (G) and threshold (T) are set here to 40 and 0.25 respectively. The setting of T=0.25 is a rather restrained choice and is employed to guard against HPW shape change and noise non-stationarities.

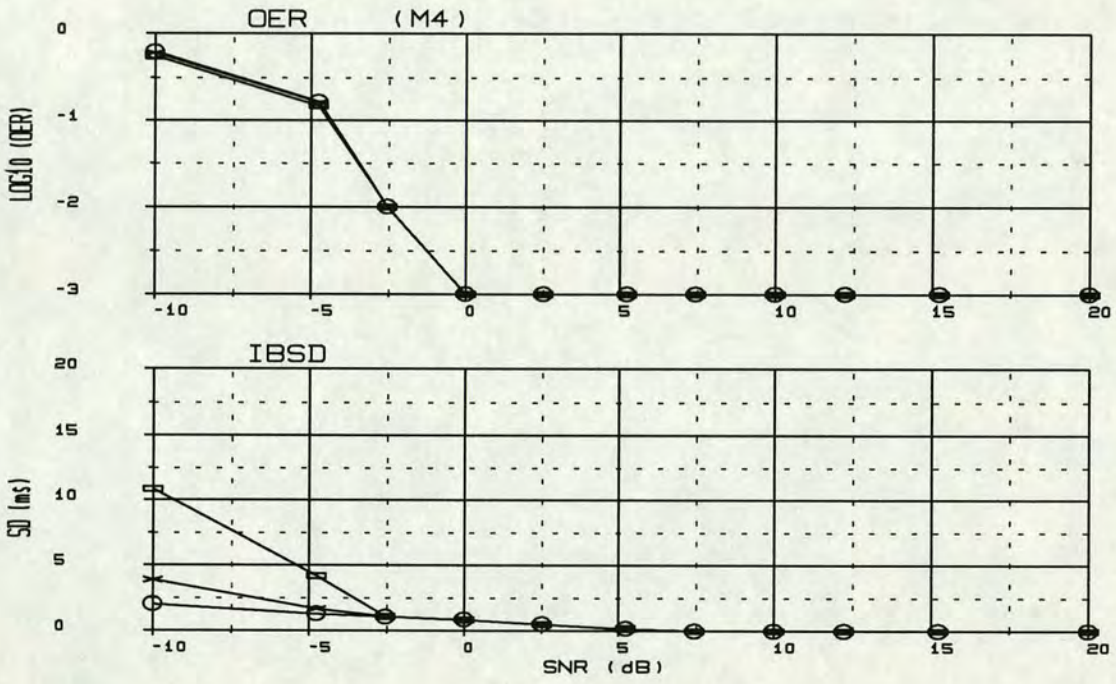

Figure 4.11 Performance of the 'Pseudo Inverse Filter Estimator' with parameters T=0.25 and G=40. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The exact reason for HPW shape evolution is not well understood. Small short-term variations are observed as well as long-term continuous changes. Some of the short-term pulse changes can be related to breathing movements, the main contributor to

HPW shape changes however is suspected by the author to originate from gross slow fetal body displacements. These modify the resonances of the feto-maternal body structures and consequently the HPW shape. The HPW pulse shape can usually be expected to evolve over long periods and therefore, steps must be taken in the derivation of fixed PIF coefficients to accommodate the possible pulse shape changes (for example setting T=0.25). The resulting performance graph (fig. 4.11) shows a 1.7dB improvement over the pulse correlation approach.

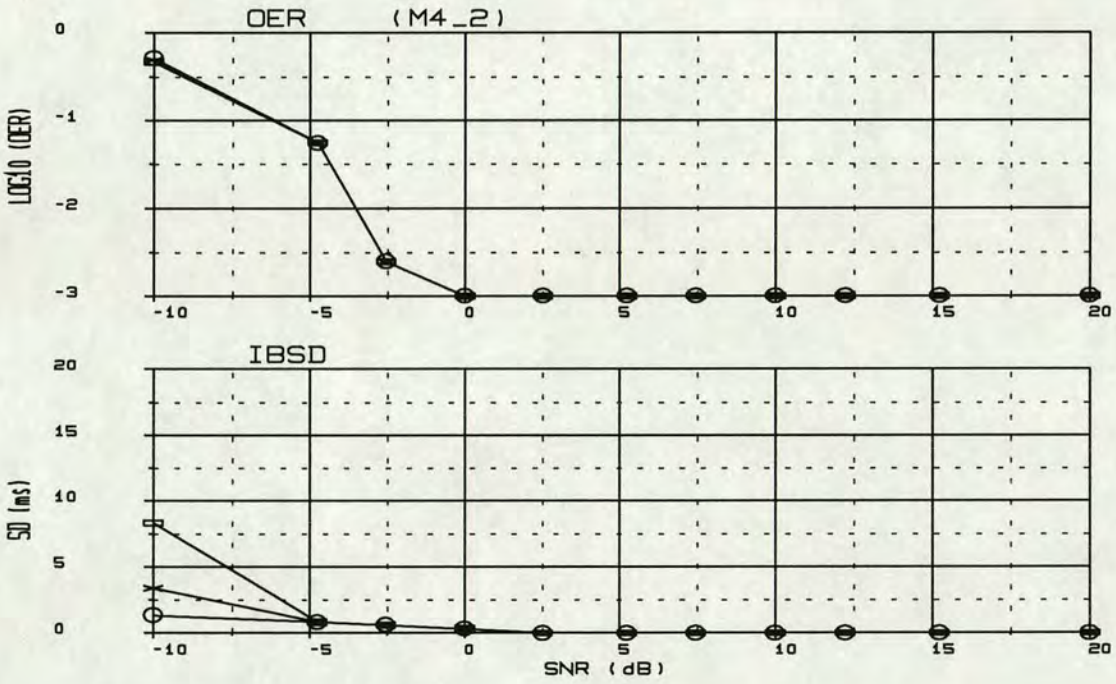

Figure 4.12 'Pseudo Inverse Filter' performance (T=0.15, G=60). The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

If continuously evolving PIF coefficients are employed (derived from a continuously updated template), a lower threshold T can be used to improve the estimates. For example, when the threshold parameter is lowered to T=0.15, another 2dB gain in performance is obtained (figure 4.12). Being at the operational limit of most BTB

estimators, this 2dB additional gain is very significant. It enables the PIF estimator to operate at SNR ranges where most other algorithms fail, practically, PIF processing reduces the signal 'down' time.

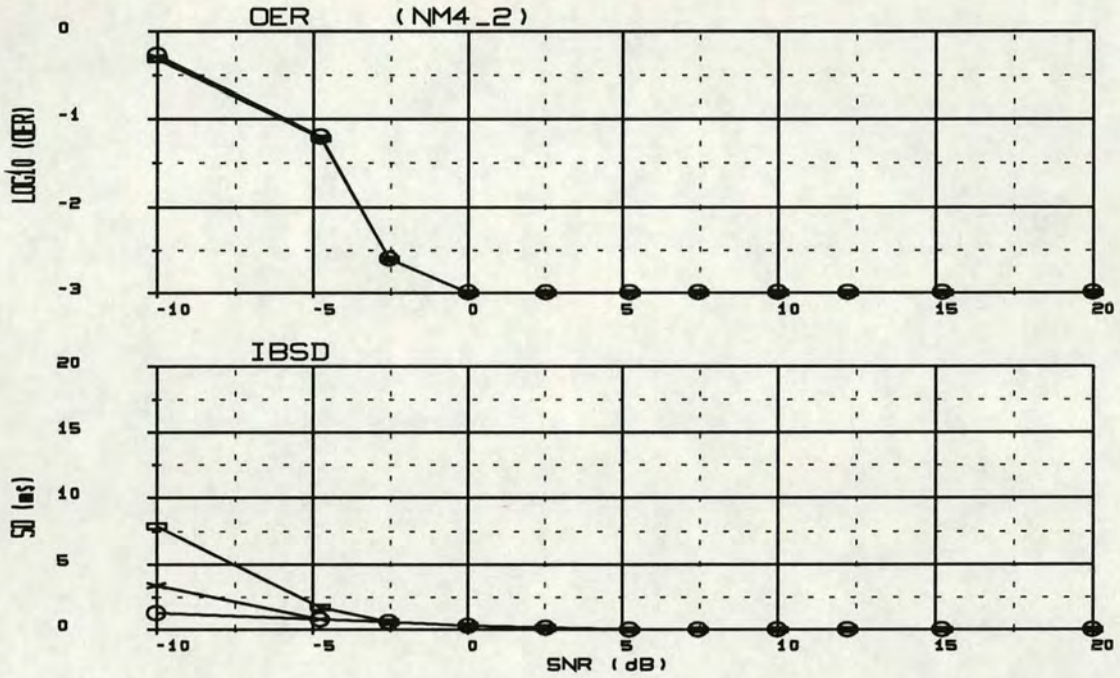

Figure 4.13 'Pseudo Inverse Filter' performance (T=0.15, G=60). Fixed white noise is added (-10dB w.r.t. signal power). The x axis measures the SNR w.r.t. the coloured additive noise The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

It is important when assessing the performance of algorithms involving some form of inverse operation to ensure that the integrity of their estimates (both in OER and IBSD terms) are not compromised in the presence of white noise in addition to coloured noise. The graphs in figure 4.13 demonstrate that the PIF performance is not disrupted by the additional presence of -10dB white noise (w.r.t. signal).

4.2.7. Block Autocorrelation Estimator

The correlation of two observed time series is widely used to assess the dependence of their respective generating processes. Computed at delay intervals (τ), correlation is traditionally used in communication systems to detect signals in buried noise. When the correlated signals are identical, the resulting autocorrelation reveals periodicities present within the observed signal. This block autocorrelation estimator utilises the envelope detected fetal HPW signals to extract the FHR estimates.

Correlation based techniques were suggested by Farvet in the early 60's as a method for estimating FHR. At the time however, this proposition was ahead of the existing technological capabilities. With the advent of more powerful microprocessors, correlation based FHR monitors were was introduced in the early 1980's[32]. These offered large improvements (from 30% to 5% signal dropout) over their predecessor monitors which mainly relied on time domain threshold detection circuits[24, 32].

The discrete autocorrelation function $R_{XX}(k)$ is given:

$$R_{XX}(k) = \sum_{i=0}^{N-1} x(i) \, . \, x(i+k)$$

N is the block length and k is an integer. Other than at $R_{XX}(0)$, the global peak from the computed autocorrelation is expected to lie at a temporal shift equal to the heart beat periodicity.

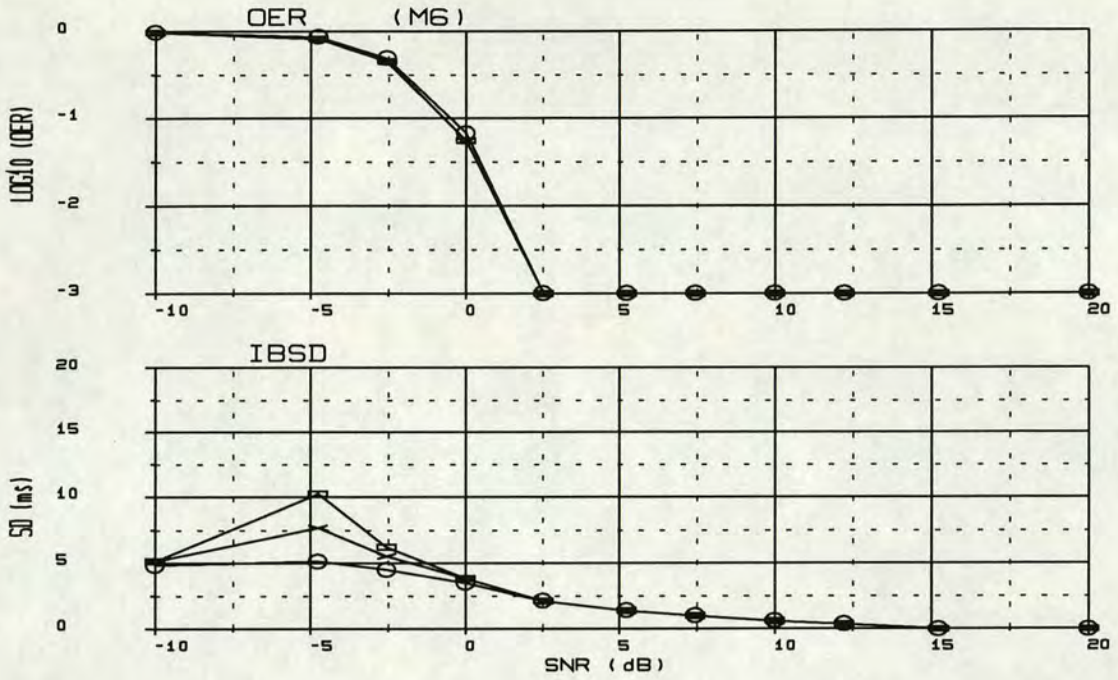

Figure 4.14 Performance of the 'Block Autocorrelation Estimator'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The rectified HPW time series is sampled continuously at 500Hz. Two second blocks (N=1000 samples) with a 75% overlap are correlated with the rest of the time series. This provides an unbiased correlation output from which a global peak is located at a temporal delay T_{est}, the peak search is limited to valid heart beat periodicities (i.e. from 0.25 to 1 seconds).

The noise performance of the autocorrelation estimator (fig. 4.14) has acceptable OER and IBSD down to -1.0dB. The OER plots of the three bands are almost merged suggesting the following two points: (a) when errors occur within the specified inclusion band, they are less than the smallest band (±10ms). And (b) when they are outside the specified band, they tend to occur outside are widest band (±40ms).

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 5.3ms, 45%	* Poor performance. * Strong bias to low frequency errors.
5 BPM/sec.	±40ms 5.9ms, 44%	* Tracking loss. * Low Frequency bias.
10 BPM/sec.	±40ms 9.4ms, 39%	
20 BPM/sec	±40ms 19.2ms, 48%	

Table 4.1 Acceleration performance of 'Block Autocorrelation'. The coloured noise SNR is set to 0dB.

As mentioned previously, unlike beat-to-beat estimators, block based estimators need to have their performance checked under the accelerating FHR. Table 4.1 illustrates a poor tracking performance obtained from the autocorrelation estimator. This is attributed to the collapse of the frequency harmonics when spectrum is computed from the accelerating FHR.

The computation cost of autocorrelation is of the order O(N.N) which increases dramatically for higher sample rates. This cost can be significantly reduced by computing the autocorrelation function from the inverse Fourier transform of the power spectrum, in which case the cost can be reduced to O($2.N.\log_2(N)$). A further reduction is possible by sampling the envelope detected signals (DC-20 Hz) closer to the Nyquist rate. Using a reduced size FFT, the power spectrum is computed, padded with zeros and inverse transformed. Thus an approximate two fold reduction in computational cost can be achieved.

Block autocorrelation was found to be immune to -10dB additive white noise.

4.2.8. Comb Filter

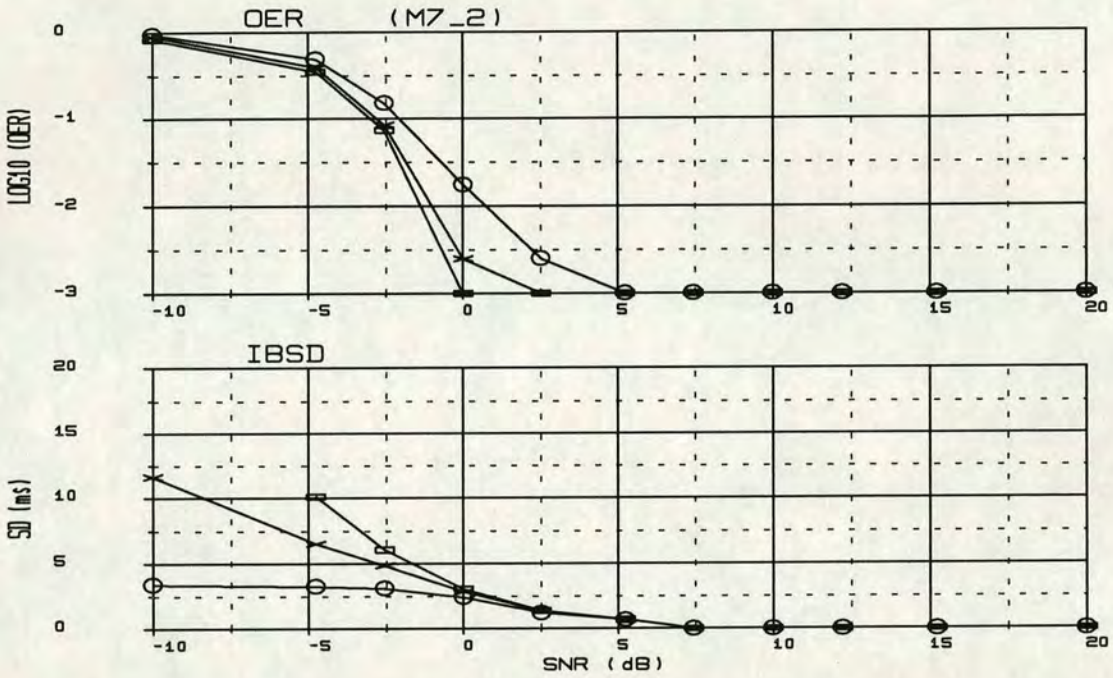

Figure 4.15 The performance of the 'Comb Filter' estimator. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Figure 4.15 illustrates the performance of the comb filter algorithm which was detailed in chapter 3. The noise performance for a constant FHR is about 1.5dB better than that obtained from the traditional autocorrelation technique. The advantage of comb filtering becomes even more apparent when the FHR tracking capability is considered. Table 4.2 shows that comb filtering can maintain reasonable error rates at relatively high FHR accelerations. In addition, quite a few of the errors observed here are due to harmonic and sub-harmonic errors (as compared to random errors). Such harmonically related errors can be easily trapped and corrected by post-processing algorithms. Although an acceleration rate of 10BPM produces an OER above the suggested 10% limit, the comb filter manages to track the heart rate albeit with

marginal deviation from the true block FHR average. This was termed in table 4.2 as 'degenerate tracking'.

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 9.4ms, 7.8%	* Some harmonic/sub-harmonic errors.
5 BPM/sec.	±40ms 9.7ms, 7.3%	* Some harmonic/sub-harmonic errors.
10 BPM/sec.	±40ms 11ms, 12.5%	* Degenerate tracking: harmonic errors at high FHR, sub-harmonic errors at low FHR.
20 BPM/sec	±40ms 17ms, 28%	* Tracking loss. * Bias to low frequency errors. * Bias to sub-harmonic errors.

Table 4.2 Acceleration performance of the 'Comb Filter'. The coloured noise SNR is set to 0dB.

4.2.9. Envelope Detector-Estimator

This algorithm makes use of the envelope detected HPW signals which are commonly derived for most block estimators. Estimates are extracted from 0.5 second blocks by selecting the global peak position. The performance of the Envelope detector-estimator is included here (fig. 4.16) to provide an indication of the measure of improvement achieved by block based algorithms.

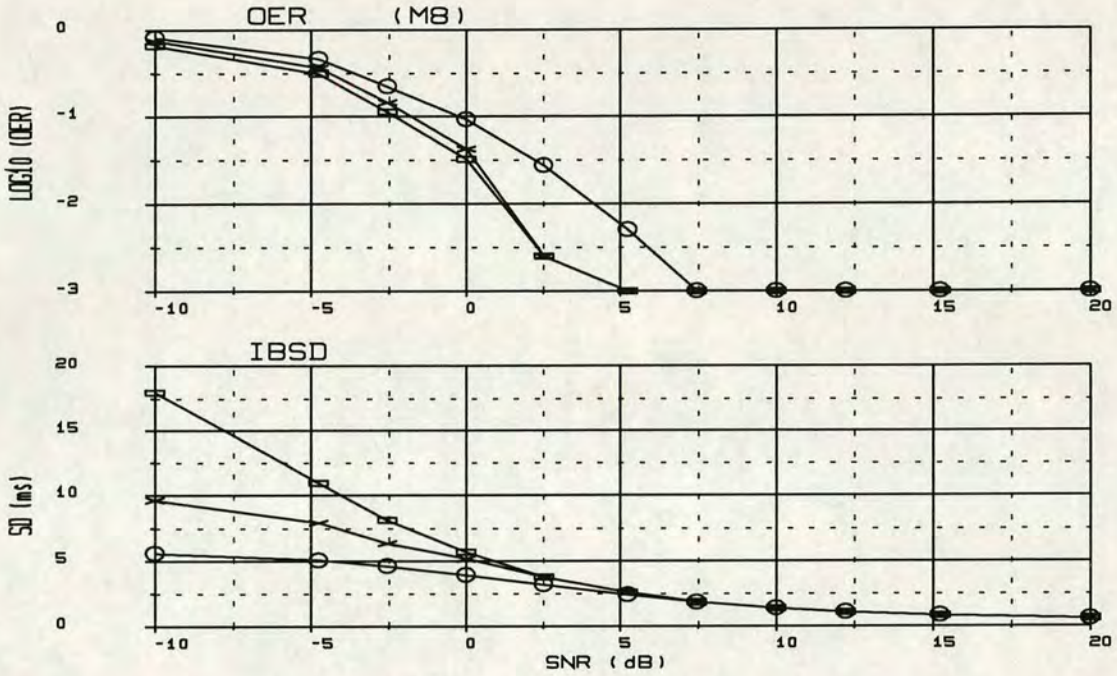

Figure 4.16 Performance of the 'Envelope Detector-Estimator'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Compared to other time domain algorithms, the limiting performance factor of the 'Envelope Detector-Estimator' seems to be the variability introduced by the HPW peak jitter. At 0dB SNR, the in-band standard deviation reaches the maximum accepted level. If long term baseline FHR was required however, then using ±40ms band the operational SNR range stretches down to -2.5 dB.

4.2.10. Cepstrum Estimator

Cepstral analysis was classically used in deconvolution processing and separation of signal products, and was initially proposed by Bogert *et al.* in 1963 for echo detection purposes. A new vocabulary was also defined to describe the relationship between the logarithmic spectrum and its Fourier transform (cepstrum). Homomorphic

applications of this technique described by Oppenheim *et al*[4] include liftering (filtering of spectrum) and spectral editing.

For the purposes of FHR estimation, signals processed using cepstral analysis are assumed to be formed from the convolution of an impulse train time series with the impulse response of the system. In the case of fetal heart phonocardiographic signals, the impulse train models fetal heart beat events and the impulse response appertains to the feto-maternal body resonances. The observed time signal x(t) is expressed as the convolution:

$$x(t) = i(t) * h(t)$$

where i(t) is the impulse train of the heart beats and h(t) is the impulse response. Cepstral analysis involves two steps these are: (a) computing the power spectrum of two seconds of observed time series x(t):

$$X(\omega) = I(\omega) \,.\, H(\omega)$$

which substitutes the convolution of the signals i() and h() in the time domain by the multiplication of their respective spectra in the frequency domain. Taking the logarithm of both sides:

$$\log \{X(\omega)\} = \log\{I(\omega)\} + \log\{H(\omega)\} \qquad 4.1$$

Step (b) involves computing the power inverse Fourier transform of equation 4.1 above:

$$FFT^{-1} [\log \{X(\omega)\}] = FFT^{-1} [\log\{I(\omega)\}] + FFT^{-1} [\log\{H(\omega)\}]$$

$$\chi(t) = \iota(t) + \eta(t) \qquad 4.2$$

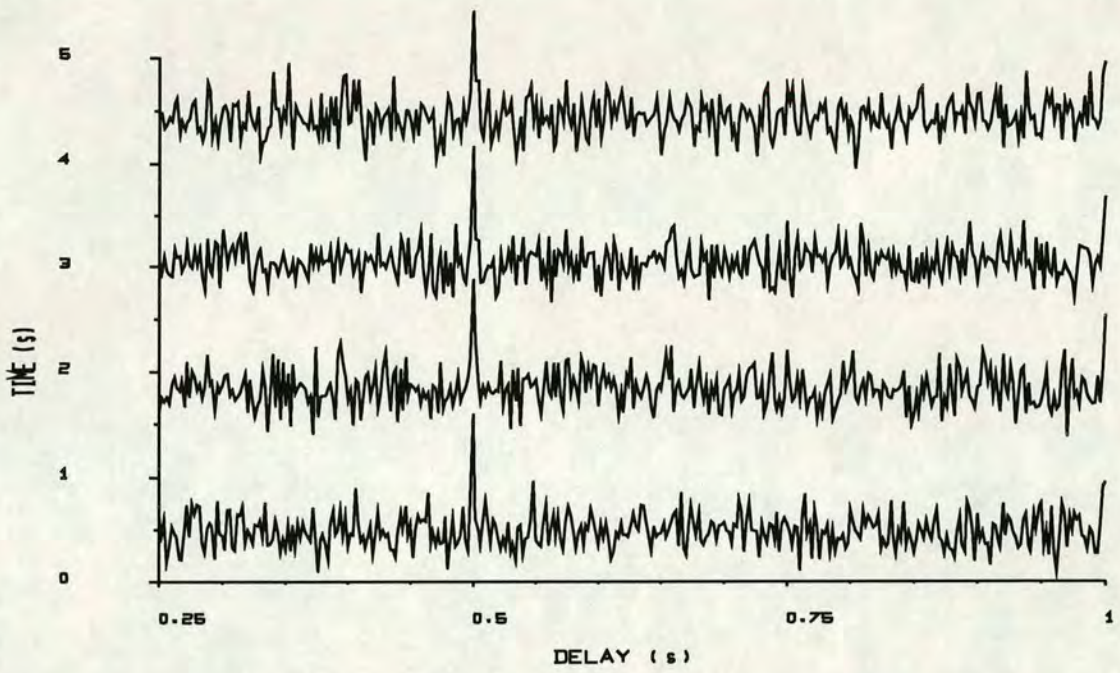

Figure 4.17 Output of the cepstrum, $\chi(t)$, obtained from 2 seconds blocks with an overlap factor of 50%. The simulated periodicity is 0.5 seconds and SNR = 0dB.

$\chi(t)$ in equation 4.2 is the addition of two components, $\iota(t)$ and $\eta(t)$, which are the power inverse Fourier transforms of the logarithmic spectrums of i(t) and h(t) respectively. As long as the two components in $\chi(t)$ can remain distinct in the cepstral domain, the detection of the impulse i(t) can be obtained (fig. 4.17). To aid spectral separation, it is desirable that the impulse response h(t) have a smooth spectrum. Such a case is seen centred around 25 Hz in figure 4.18. When inverse Fourier transformed into cepstral domain, an impulse response with a smooth spectrum results in a component lying close to zero shift.

Figure 4.18 shows a typical log spectrum taken from a block of two seconds phono signal when FHR is 2Hz. The harmonic structure of the signal extends well into the low signal amplitudes. The inset of the same figure details some of the harmonic components; the peaks A, B and C are equally spaced by 2Hz. When inverse

transformed, these peaks contribute to a 0.5 seconds pulse train. Also seen in figure 4.18 are the contributions to that same 0.5 seconds impulse train from components below -40dB (75-200 Hz). These have harmonic structures which are equally large to those from the region of the main signal power (15-50 Hz).

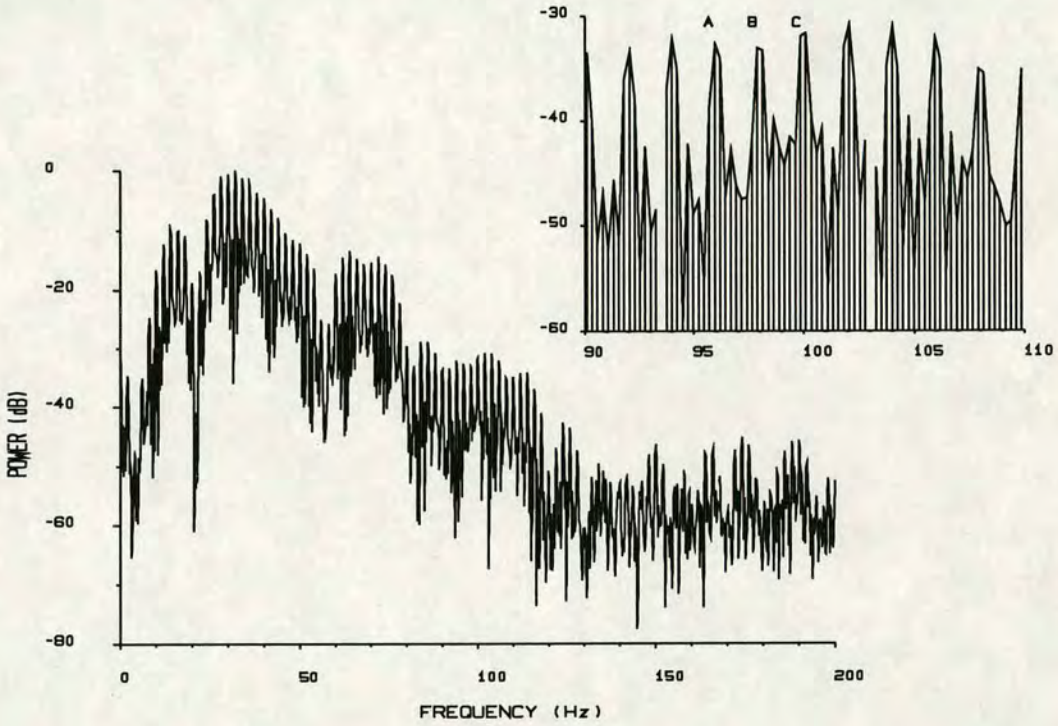

Figure 4.18 Logarithmic spectrum obtained from a 2 seconds block of fetal phono signals, the heart rate is 2Hz and SNR=-40B (coloured noise).
Main: reveals the underlying harmonic structure of the impulse train.
Inset: expanded spectral region from 90 to 110 Hz showing spectral detail which extends into the lower signal amplitudes (@120-200 Hz).

In comparison to the logarithmic spectrum, the power spectrum severely curtails the influence of low amplitude harmonics. Therefore the pulse train obtained from the inverse Fourier transform (corresponding to the signal autocorrelation) is mainly composed from harmonics within the main signal power (20-40 Hz). The structure of these harmonics can be severely disrupted by the presence of narrow band noise

which is frequently present in fetal heart phono recordings. This results in the autocorrelation having a poor performance (as shown later) when applied to blocks of the raw sampled HPW signals, instead, it is usually applied to the demodulated HPW signals.

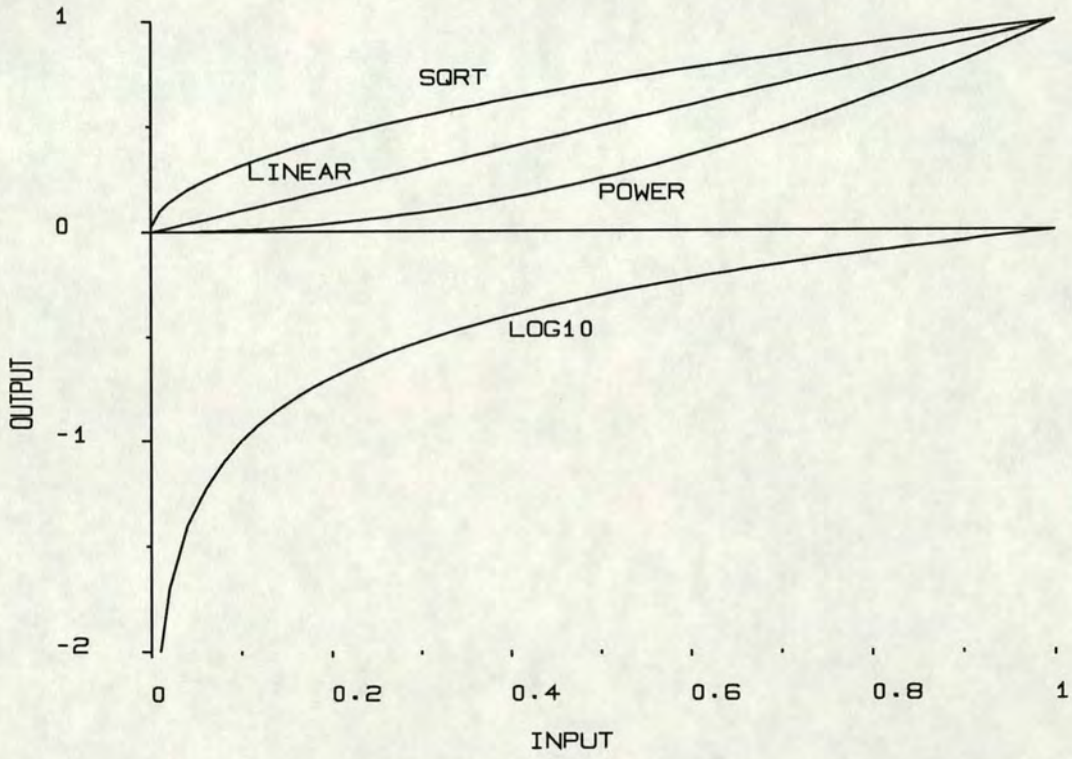

Figure 4.19 Mapping functions on the normalised spectrum magnitude. The output of LOG10(amplitude) is truncated at -2.

In order to further understand the mechanism of cepstral analysis, the effect magnitude mapping performed by the logarithm operator is compared to linear and power magnitude (fig. 4.19). It can be seen that the logarithm mapping significantly emphasises low signal amplitudes which the power spectrum suppresses. Clearly the gain obtained in cepstral analysis strongly relies on the ability of the logarithm operator to emphasise the harmonic and null structures at very low signal levels.

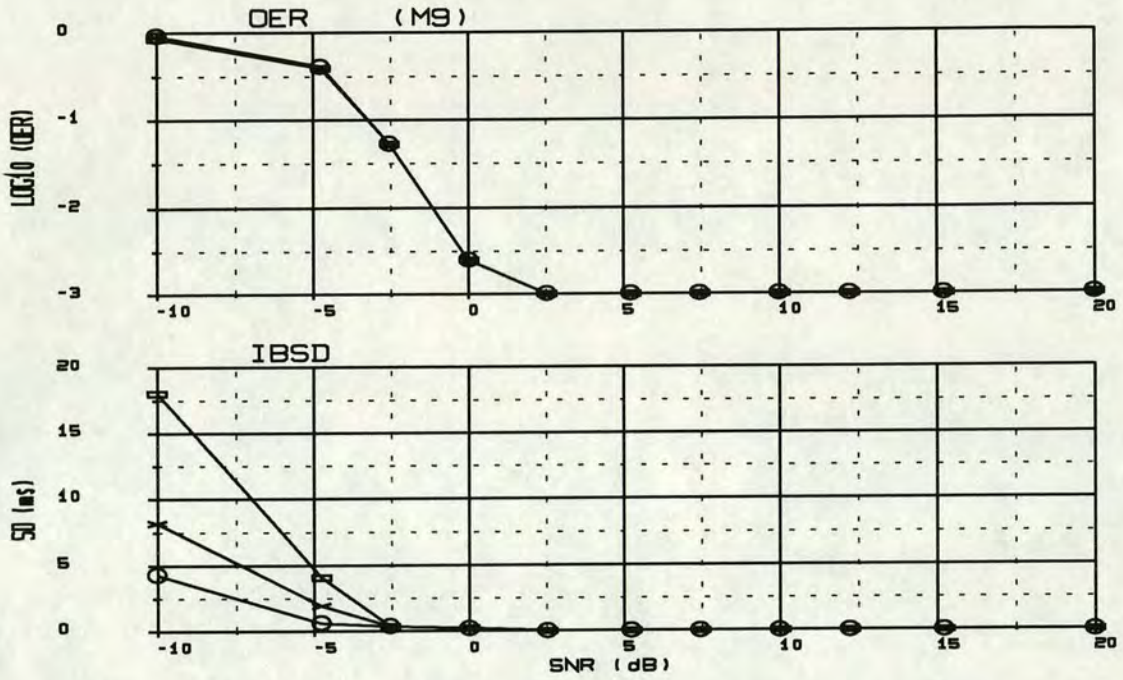

Figure 4.20 The performance of the 'Cepstral Estimator'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

A good OER performance is obtained from cepstral analysis when applied to simulated HPW signals with coloured noise (fig. 4.20). In addition, the IBSD introduced by estimate jitter within the inclusion band (±20ms for example) is kept reasonably low and therefore has the ability to detect genuine low FHR variance.

The nulls observed in the frequency spectrum (fig. 4.18) correspond, for that block, to zeros placed close to unit circle in the Z domain. These zeros are expected to move away from the unit circle in the presence of broadband noise. This noise as well as 'swamping' the harmonic and null structures at low amplitudes, contributes to a random signal component across the spectrum. This results in the disturbance of the cepstral impulse train obtained from the inverse Fourier transform.

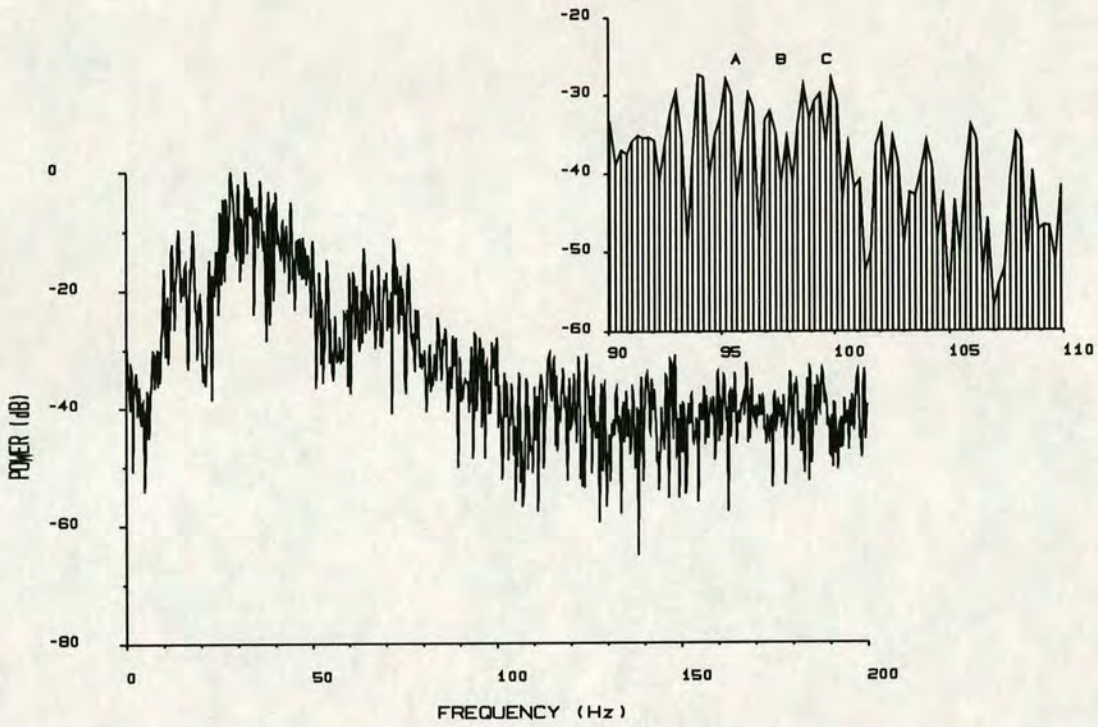

Figure 4.21 White noise which has a 0.01 variance (-20dB) w.r.t. signal power is added to the same signal block which produced figure 4.18.
Inset: expanded spectral region from 90 to 110 Hz, compare with inset of figure 4.18.

The effect of broad band noise on the logarithmic spectrum can be clearly seen when figures 4.18 and 4.21 are compared: The low signal power (@120-200 Hz) is swamped and the harmonic/null structures are totally decimated. When corrupted by broad band noise, the spectrum below the three markers A, B and C portrays no harmonic structure (fig. 4.21 inset).

While pure white noise is not commonly observed in fetal phono signals, low amplitude broad-band noise is usually present and is the performance limiting factor for cepstral analysis. Another relevant factor is the digital quantisation noise which is -48dB for an eight bit optimally conditioned signal. In practice however, levels close to -24 dB are often observed as the maximum gain range of the AGC circuitry is reached. It can be concluded that cepstral analysis performance is dependant on the

type of interference noise present in the signal, this being worst for broadband interference noise (fig. 4.22).

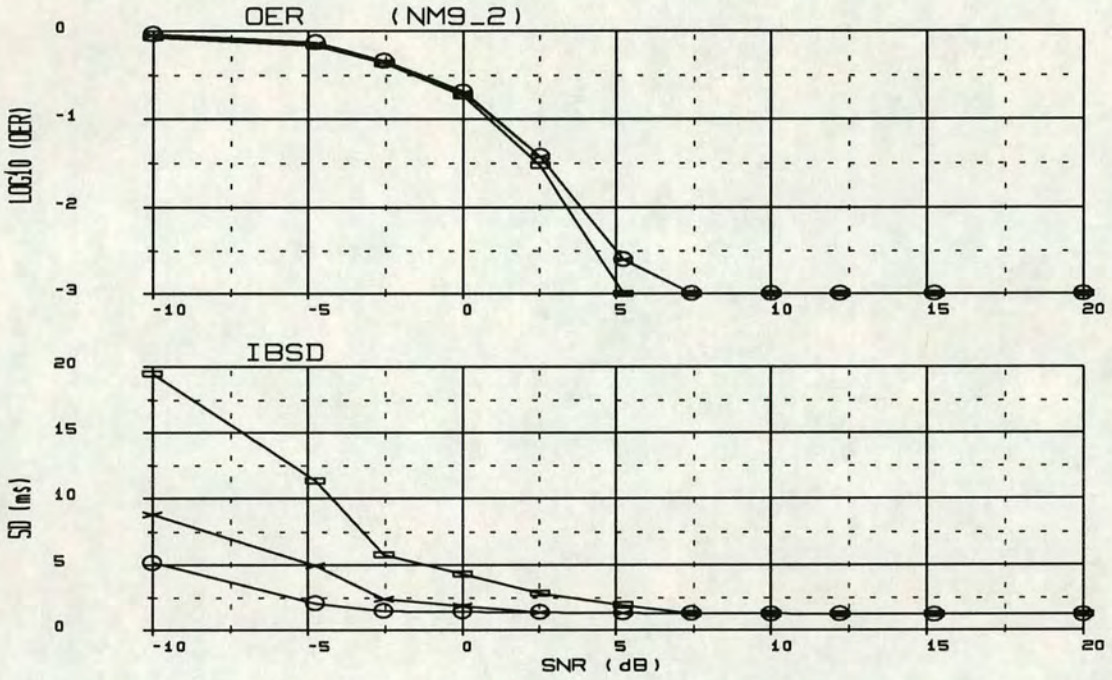

Figure 4.22 Cepstral analysis with a fixed power additive white noise -10dB w.r.t. signal power. The x axis measures the SNR w.r.t. the coloured additive noise. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 6.9ms, >50%	* Poor Tracking performance. * Tracking loss at low acceleration rates. * All errors are random, no bias.
5 BPM/sec.	±40ms 12ms, 48%	
10 BPM/sec.	±40ms 16ms, 47%	
20 BPM/sec	±40ms 22ms, >50%	

Table 4.3 Acceleration performance of Cepstrum processing. The coloured noise SNR is set to 0dB.

The cepstrum estimator is evaluated using HPW signals containing FHR accelerations (table 4.3). A sizeable degradation in performance is obtained even for the lowest simulated accelerations (2 BPM/second), moreover, the erroneous estimates were mainly random, few of them were harmonically related. While high acceleration rates are rare events from which the cepstrum estimator can recover, lower acceleration rates (less than 5BPM/second) are very commonly encountered in cyclic FHR changes. These form an important part of short term FHR variability measure.

No computational reduction is possible when using cepstral analysis. Firstly reducing the HPW time series sample rate results in the loss of most of the harmonic structure which forms a sizeable portion of the signal cepstrum (fig. 4.18). Secondly the time domain padding required to adequately expose the harmonic structures of the spectrum cannot be reduced otherwise aliasing introduced into the cepstrum domain. The aliasing is a product of the strong non-linearity created by the logarithmic operator which increases the frequency content of the inverse Fourier transform.

4.2.11. Raw Signal Block Autocorrelation

The performance of cepstral processing was shown (above) to degrade significantly in the presence of white noise. This degradation stems from the application of the logarithmic operator on the spectrum. Presented first in this section is the most commonly used of the spectral mappings (shown in figure 4.19), namely the power spectrum. The inverse transform of which results in the time domain block autocorrelation. Compared to the envelope detected block autocorrelation described earlier, this approach bases its estimates on the autocorrelation of unprocessed HPW signal blocks.

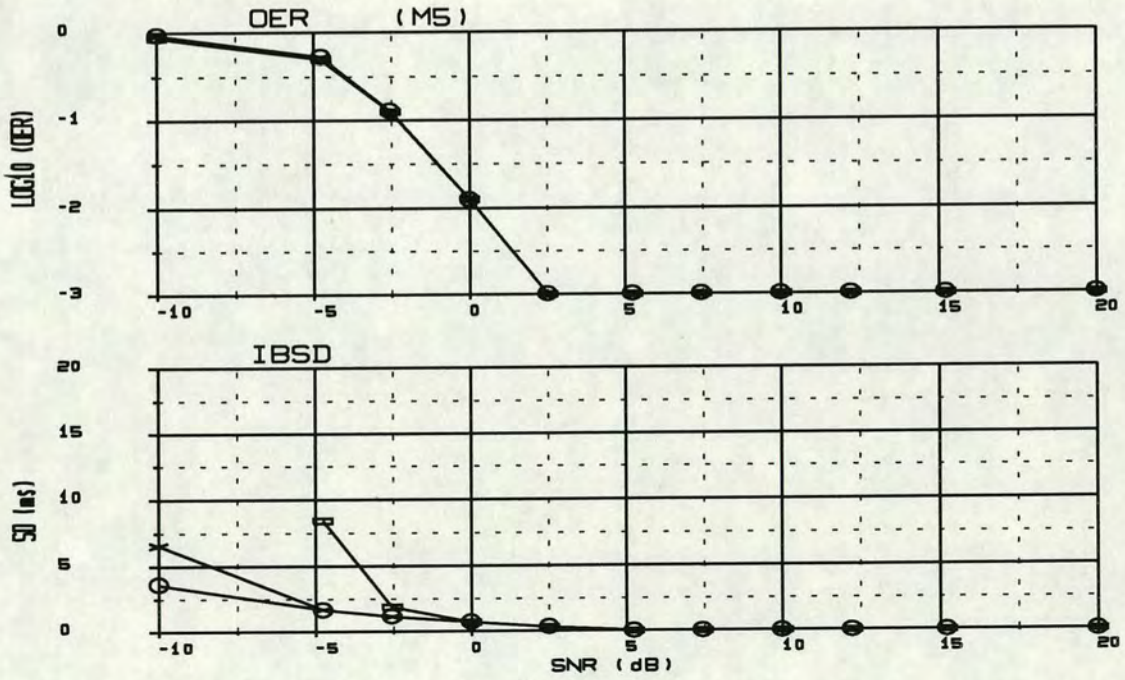

Figure 4.23 'Raw Signal Block Autocorrelation' performance. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The good performance obtained here is not representative of the real phono HPW monitoring (fig. 4.23). The main reason is that the simulated HPW signal does not contain any beat-to-beat variability. The effect of FHR change on the frequency spectrum was explained in chapter 3 "Comb Filtering". The raw signal autocorrelation is considerably more sensitive to FHR change, a hypothetical example is given: Assuming a high resonance HPW centred at 25Hz (40ms oscillation periodicity), when the autocorrelation of a block straddling a step periodicity increase of 20ms is computed, a null is obtained at the expected block average periodicity. This is due to the 180 degrees misalignment caused by the FHR change.

As expected table 4.4 shows the poor performance from the raw signal autocorrelation, even at the lowest acceleration rates, the OER is too high for any practical implementation of the algorithm. Even when simulated at a higher SNR, the performance of this algorithm is very poor.

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 7.2ms, 42%	* All errors are random. * High variance.
5 BPM/sec.	±40ms OER > 50%	
10 BPM/sec.	±40ms OER > 50%	
20 BPM/sec	±40ms OER > 50%	

Table 4.4 Acceleration performance of 'Raw Signal Block Autocorrelation'. The coloured noise SNR is set to 0dB.

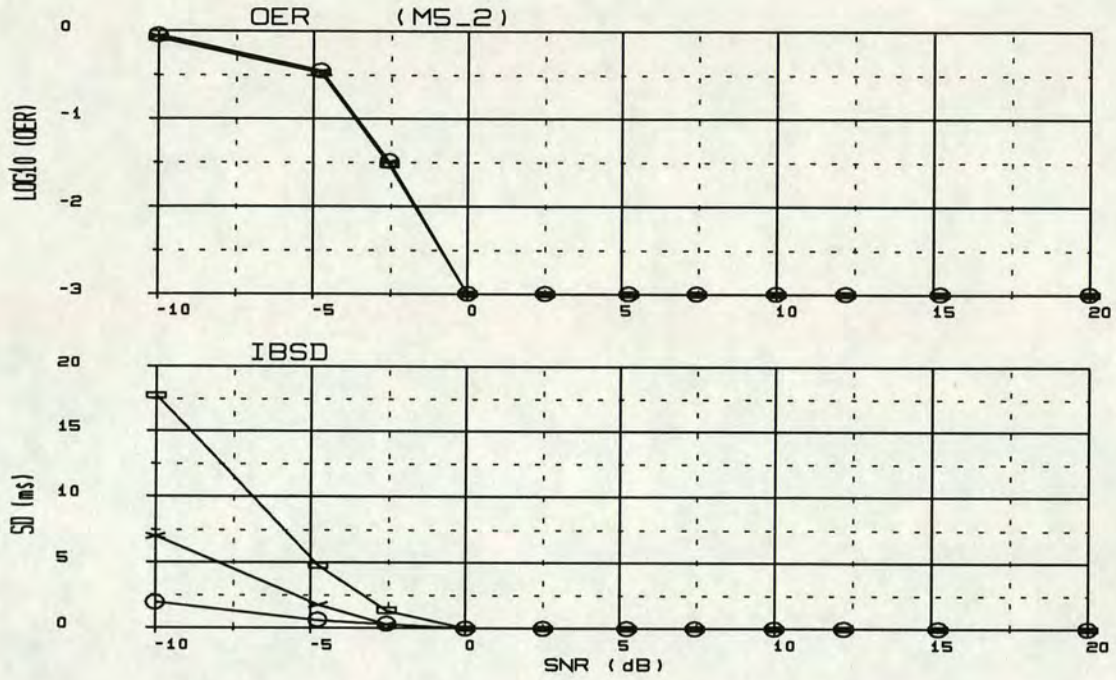

Figure 4.24 Performance of inverse Fourier transform of SQR spectral mapping. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Another possibility for spectral mapping is the square root (SQRT) operator (fig. 4.19) which has an amplitude transfer function akin to the logarithmic mapping. Shown in figure 4.24, the SQRT based estimator out-performs the normal autocorrelation processing by about 1dB. While mildly emphasising the low spectral components, SQRT mapping limits the effect for very low signal levels. It can therefore be expected that SQRT would be less susceptible than cepstral analysis to broad-band interference noise. This was found to be the case with less than 1.0 dB degradation in performance. There was not a sufficiently large improvement in the acceleration performance of the SQRT mapping to warrant further investigation. The OER for 2 BPM/second acceleration was measured at 31% (full results are not included).

The performance of the linear spectral mapping (fig. 4.19) falls somewhere between the performance of normal autocorrelation, and of the SQRT mapping shown above. Additional mappings which take the n^{th} root of the spectral magnitude were also tested. The results which are not included here show that although the steady FHR estimates can be improved on, the acceleration OER of this family of spectral mappings remained at unacceptably high levels.

4.2.12. Modulus Difference

Initially suggested for voice pitch extraction[6], the Modulus Difference (MD) algorithm was adapted by Manning[11] for FHR estimation. Expressed for stationary FHR signals as:

$$MD(\Delta t) = E\{ |x(t) - x(t+\Delta t)| \} = \int_{t=-\infty}^{+\infty} |x(t) - x(t+\Delta t)| \, \delta t$$

The modulus difference output should reach a minimum when the temporal shift Δt, precisely matches the inter-pulse interval. For signals which do not possess long-term stationarity, shorter signal vectors (of length N) are employed in the discrete modulus difference. The MD output is computed over the valid fetal heart periodicity; i.e. Δt from 0.25 to 1.0 seconds, and the shift producing the minimum MD is assumed to pin-point the fetal heart periodicity. The discrete modulus difference equation is given:

$$MD(t,\Delta t) = \frac{1}{N} \sum_{i=0}^{N-1} |x(t+i) - x(t+i-\Delta t)|$$

MD is a non-linear processing technique which greatly complicates analytical approaches. Compared to the square error introduced by correlation however, the MD

output has an error which is weighted modulus- linearly. It can therefore be expected that MD will have a worse performance than correlation at high SNR and a better one at low SNR, the cross-over point being at 0dB.

The main motivation behind employing this algorithm as opposed to correlation was to avoid costly multiplications by replacing them with simple additions and subtractions[11]. Newer microprocessors however, such as TMS320, MC56000 or DSP32 have reduced the multiplication operation cost and rendered it equivalent to that of performing additions. As such, the updated modern computation cost of the modulus difference algorithm is very similar to that of non-optimised correlation processing.

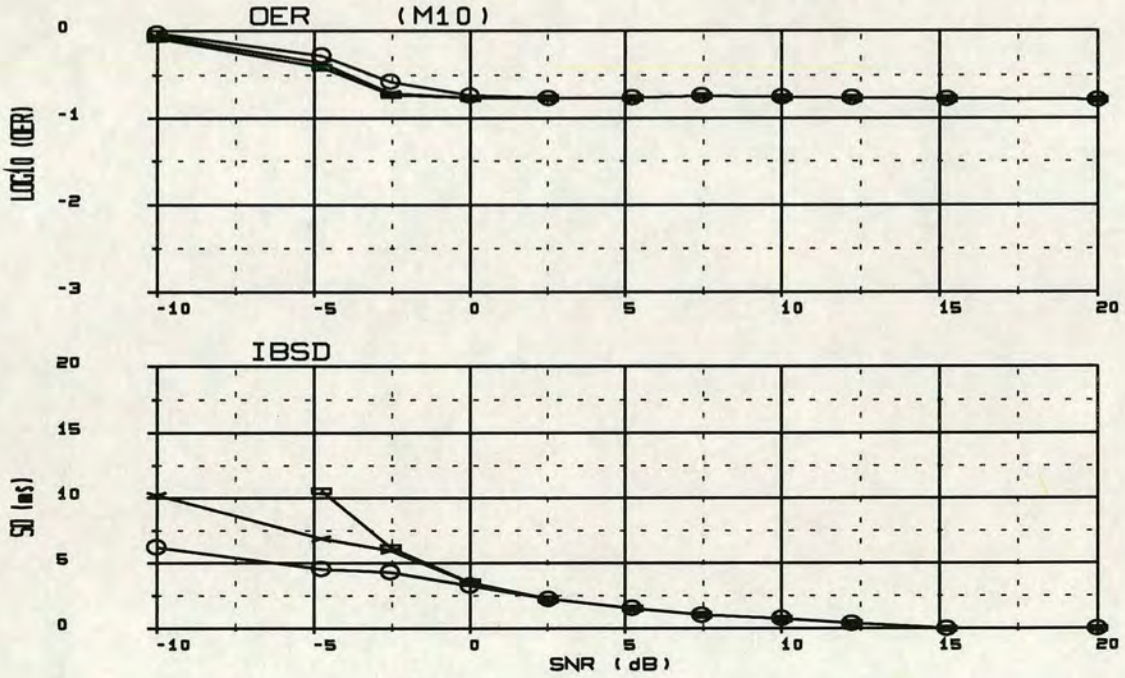

Figure 4.25 Performance of the 'Modulus Difference'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

An error rate of 1 in 6 can be seen in figure 4.25 suggesting a poor performance from the modulus difference algorithm even at high levels of signal to noise ratios. A close inspection of the MD output (fig. 4.26) reveals that even at high SNR's most of these errors are due to the cyclic nature of MD which sometimes results in estimates having twice the inter-pulse period (1 second in this case).

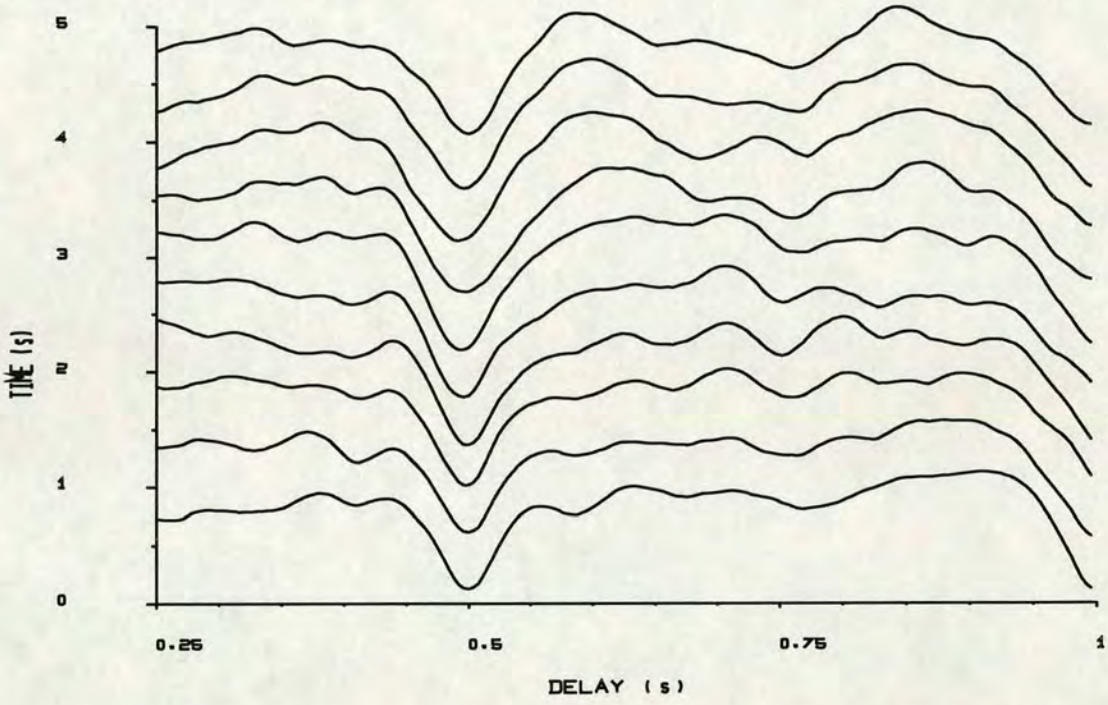

Figure 4.26 Modulus difference output from a simulated, steady 120 BPM (0.5 seconds periodicity) phono signal. Block length is 2 seconds with a 75% overlap, the SNR is 0dB.

This kind of sub-harmonic error is very similar to that observed in correlation processing, the situation can be rectified by simply applying a bias to the MD output. Figure 4.27 shows the effect on the performance when linear amplitude bias (0.8-1.0) is applied to the MD output. The error rate at high SNR drops from 1 in 6 to 1 in 30.

In addition, all errors below 0dB SNR (fig. 4.27) remain due to estimating twice the heart pulse periodicity.

In practice, two factors aid the production of the correct global MD minimum: Firstly, evolutionary changes in HPW pulse shape result in the comparison of increasingly dissimilar pulses for larger Δt shifts. And secondly, gradual changes in heart rate favours the production of a global MD minimum at smaller shifts. These two factors are particularly relevant to Δt shifts which are multiples of the true inter-pulse periodicity.

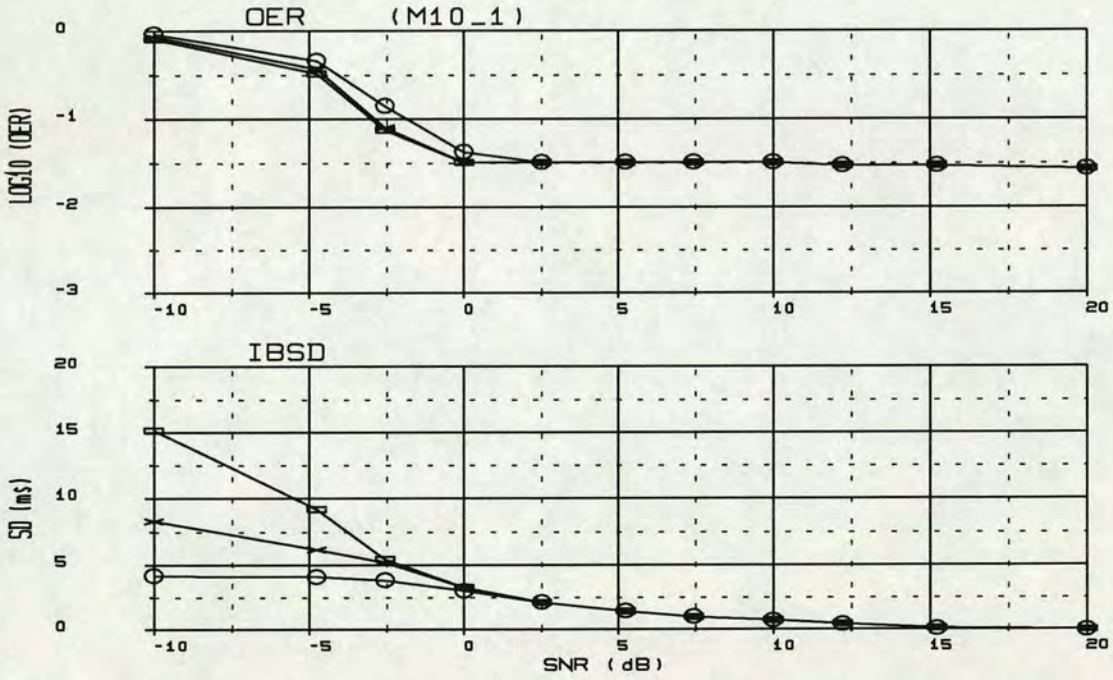

Figure 4.27 The performance of the 'Modulus Difference' estimator when a simple bias is introduced across the MD output vector. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Despite artificially introducing a linear bias (fig. 4.27) and the natural favouring of lower Δt shifts, in practice an artificial intelligence algorithm is required to trap the

dominant double periodicity errors. The same effect is commonly observed when using monitors which are based on threshold or correlation techniques.

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 6.1ms,6.9%	* Good performance. * Some sub-harmonic errors.
5 BPM/sec.	±40ms 10.3ms, 5.5%	* Good performance. * Some sub-harmonic errors.
10 BPM/sec.	±40ms 16.3ms, 8.3%	* Reasonable tracking. * Deviates from correct estimates.
20 BPM/sec	±40ms 22ms, 34%	* Reasonable tracking. * Large deviation from correct estimates.

Table 4.5 Acceleration performance of 'Modulus Difference'. The coloured noise SNR is set to 0dB.

The MD is found to have a good FHR tracking ability (table 4.5), the OER remains reasonable at FHR acceleration of 10 BPM/second and while an acceleration of 20 BPM/second produces an error rate 34%. These errors are mainly due to estimates which manage to track the FHR acceleration trend but have sightly large deviations.

4.2.13. Fundamental Spectral Estimator

The envelope detected time series is sampled at 50 Hz and gathered in blocks of 2 seconds duration. The Fundamental Spectral Estimator pads the data record with zeros and then applies the Fourier transform. The size of zero padding used depends on the quantisation required. The spectrum is then searched from 1 to 4 Hz to locate the fundamental heart beat component.

The resolution, defined as the ability to resolve two closely spaced frequency components, achieved from Fourier transforms is inversely proportional to the observation time. In the absence of noise, only FHR related components are present in the spectrum and therefore spectral resolution (between signal and noise components) does not effect the estimator performance. In cases of noisy signals however, random spectral components reduce resolution, in the extreme to the order of 0.5 Hz (30 BPM). In most cases of good quality signals, the limitation is the frequency quantisation when attempting to locate the peak position. This is inversely proportional to the transform size; Δf=12mHz for a 4096 transform size (0.73 BPM).

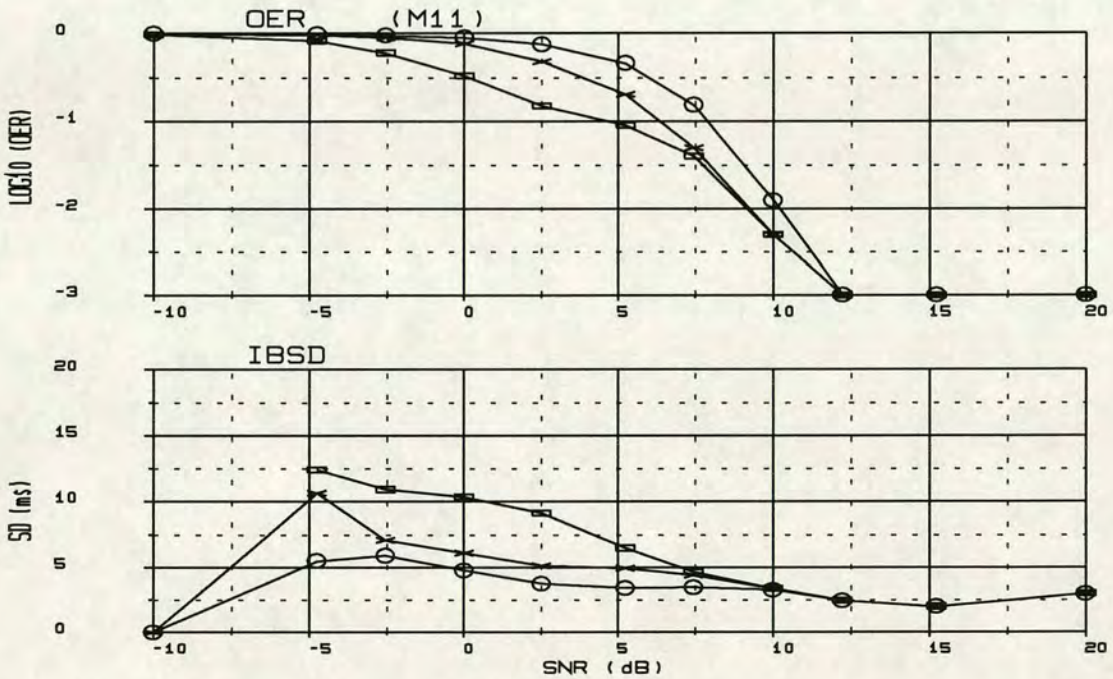

Figure 4.28 Performance of the 'Fundamental Spectral Estimator'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles). Estimates obtained from two seconds blocks.

Commonly observed coloured noise which occurs close to the fundamental component induces 'peak pulling' which manifests itself in the performance graphs as: (a) FHR variance (fig. 4.28 bottom) for small interference signals or (b) as errors (fig. 4.28 top) for larger ones. A rigourous inspection of the distribution of the estimates, reveals an increasingly strong bias towards lower frequencies as the level of interference noise is raised. This phenomenon will now be investigated.

When conditioning the HPW signals, selective filtering rejects most of the noise leaving components which lie in the spectral region of the HPW. The application of analogue non-linear circuits (such as the full wave rectifier) causes the re-introduction of wide-band noise. This effect is analysed using white Gaussian noise $\psi(t)$ of variance σ^2, and simulated using 100Hz band-limited white Gaussian noise which is sampled at 500 Hz. Taking the square of the signal $\psi(t)$ has the Fourier equivalent:

$$\psi(t)^2 \quad <==> \quad \Psi(\omega) * \Psi(\omega) \qquad 4.3$$

Which is the auto-convolution of a random, flat spectral distribution. Taking the expectation of the convolution (eq. 4.3):

$$E\{ \Psi(\omega) * \Psi(\omega) \} = \sigma^2 \quad \text{for } \omega = 0$$

$$= 0 \quad \text{otherwise}$$

the vector $\Psi(\omega)$ is not infinite, and the convolution integral (eq. 4.3) in this case has a finite observation length which is equivalent to random walk processes[5], these have the following variance:

$$\vartheta(\omega) = \frac{2.\pi - \omega}{2.\pi} \cdot \sigma_\theta^2 \qquad 4.4$$

Where σ_θ^2 is the variance of $\psi(t)^2$ with the mean removed.

Equation 4.4 results in a triangular shape power spectral density centred on 0 Hz and 200 Hz wide, a simulation of the above analysis is shown in Figure 4.29. The DC term, σ^2, is considerably larger than the contribution of equation 4.4 and is therefore removed to aid visual presentation, in practice the signal is AC coupled to get rid of the large DC component.

The application of the square root operator (to obtain full wave rectified samples from $\psi(t)^2$) is a further non-linearity. This can be expected to cause further cross-spectral spreading due to the introduction of intermodulation products. In practice, the distributions of the spectral components of both the squared and full wave rectified signals seem very similar (except for a scaling factor) (fig. 4.29).

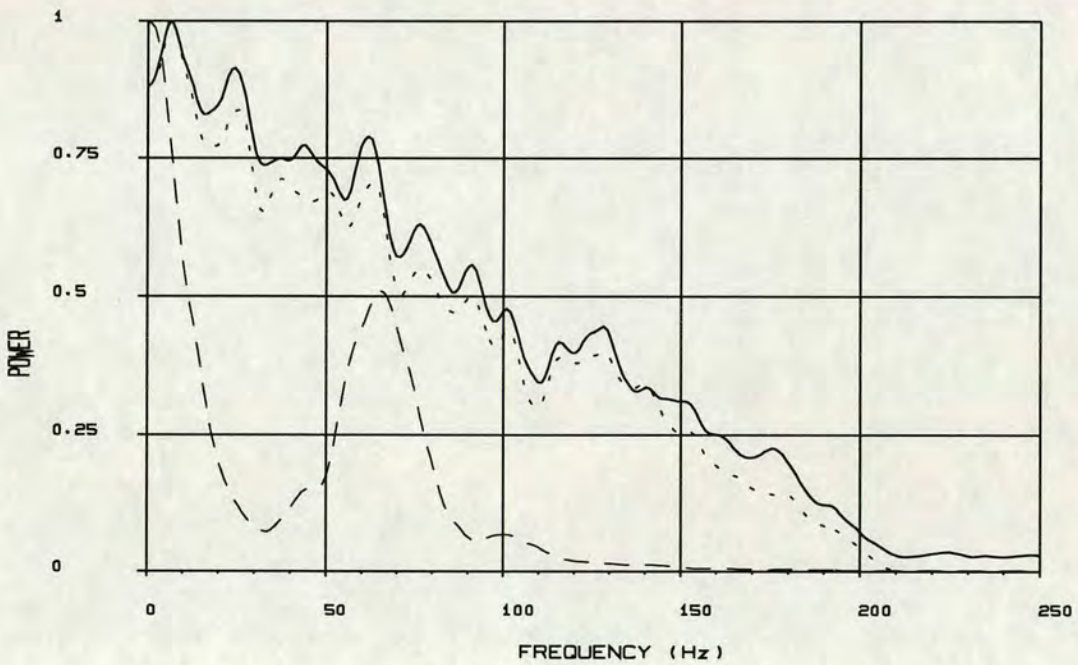

Figure 4.29 Power spectral density (DC component removed) obtained by spectral averaging over 40 seconds of: (a) the square of white Gaussian noise band limited to 100 Hz (dotted trace), (b) Full wave rectified white Gaussian noise (continuous trace) and (c) full wave rectified coloured noise (dashed trace).

When spectral averaging is performed on full wave rectified simulated phono noise (fig. 4.29), two main features appear: (a) a component centred at 65 Hz which corresponds approximately to twice the centre frequency of both the observed noise and HPW's. This component is normally filtered out. (b) The other feature is the low frequency decaying spectral component which if extrapolated, would expose approximately a 25 Hz noise span with a triangular shape. The latter feature is assumed to originate from the cross products of terms existing within the band-pass filter which spans the HPW resonances (26-38 Hz).

In order to obtain the expected power spectral density (smooth trace), spectral averaging of the full wave rectified noise was employed in figure 4.29. In practice, the noise power varies from block to block resulting in spectral component power variance. These low frequency components do not form a steady feature which can be accounted for when selecting the fundamental heart beat component from a single block. On the contrary, observations of the valid FHR region show peaks which appear randomly across the spectrum. These peaks have a higher probability of occurrence at the lower end of the spectrum.

While the same low frequency noise is present in all spectrum based algorithms, the fundamental spectral estimator bases its estimates on the power of the FHR fundamental which lies at the lower end of the spectrum. This region of the spectrum was shown above to be most contaminated with noise, hence the bias in the observed results towards lower frequencies.

In practice the presence of none of the harmonics (including the fundamental) components can be relied upon. This is due to two major effects: (a) the dual pulse nature of HPW signal (2 beats per cycle) and (b) changes in FHR causes smearing or splitting of the harmonic peaks (the effect is more pronounced at higher harmonics). The former effect can be alleviated by suppressing the second HPW (by appropriate

filtering) thus obtaining an pulse time series composed solely of the first HPW (simulated in this case).

As shown in table 4.6 the 'Fundamental Spectral Estimator' has a poor tracking ability even at the lowest FHR accelerations rates (2 BPM/second).

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 14.9ms, 50%	* Poor tracking. * Strong low frequency bias.
5 BPM/sec.	±40ms 16.4ms, 44%	
10 BPM/sec.	±40ms 16.7ms, 37%	
20 BPM/sec	±40ms 18ms, 38%	

Table 4.6 Acceleration performance of the 'Fundamental Spectral Estimator'. The coloured noise SNR is set to 0dB.

4.2.14. Iterative Block Clean-up

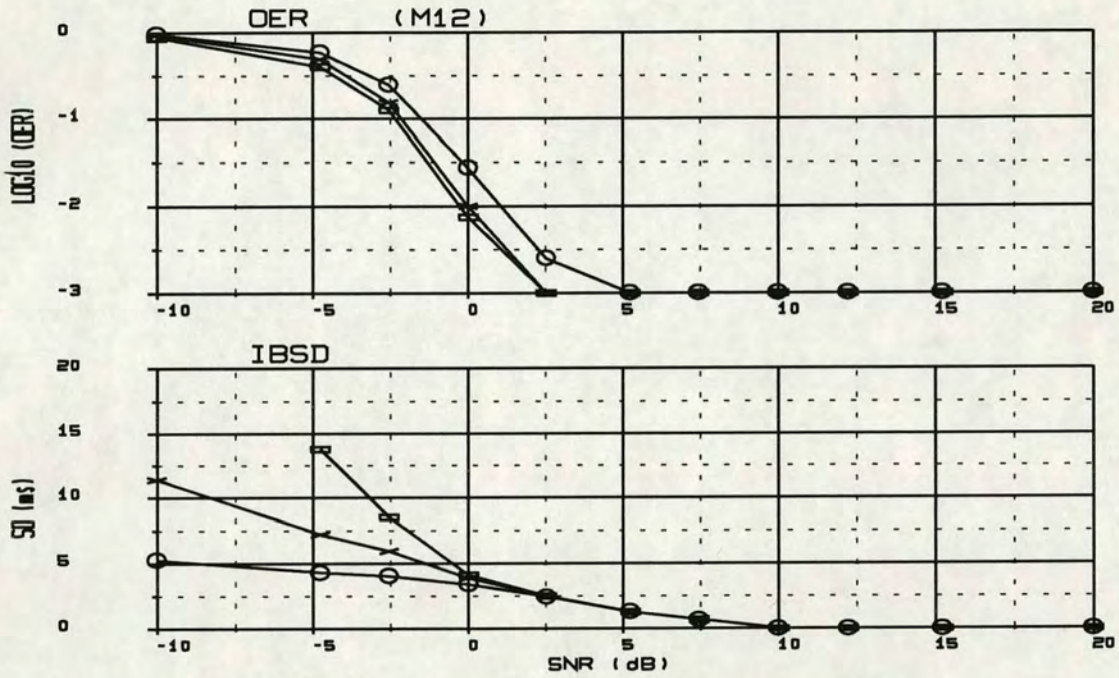

Figure 4.30 Performance of the 'Iterative Block Clean-up'. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The 'Iterative Block Clean-up' (IBC) which was introduced in chapter 3 has a reasonable OER performance (fig. 4.30). However, a rather high IBSD is introduced by the estimator. On inspection of the erroneous estimate distribution, with the exception of a few harmonically related errors, the majority of errors were random. This lack of bias makes the IBC algorithm suitable FHR baseline monitoring.

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 8.6ms, 16%	* Reasonable performance. * Some harmonic error.
5 BPM/sec.	±40ms 7.6ms, 13%	* Reasonable performance. * Some harmonic error.
10 BPM/sec.	±40ms 10.9ms, 7%	* Reasonable performance.
20 BPM/sec	±40ms 15.7ms, 14%	* Tracks well. * Large deviation from correct estimate.

Table 4.7 Acceleration performance of the 'Iterative Block Clean-up'. The coloured noise SNR is set to 0dB.

The IBC is amenable to computation cost reduction; the first three iterative steps can be performed at low FFT size to locate approximately the HPW positions. A final stage comprising zero-padded FFT is utilised to improve the quantisation. Because of the relatively high estimate variance in this case, improving the quantisation by FFT padding is a case of diminishing returns. Table 4.7 shows that IBC is robust in the presence of FHR accelerations up to 20BPM/second.

4.2.15. Inverse Least Squares

The inverse least squares filter was discussed in chapter 3. White noise at -7dB was included in the signal autocorrelation matrix in order to build broad-band noise immunity into the filter coefficients.

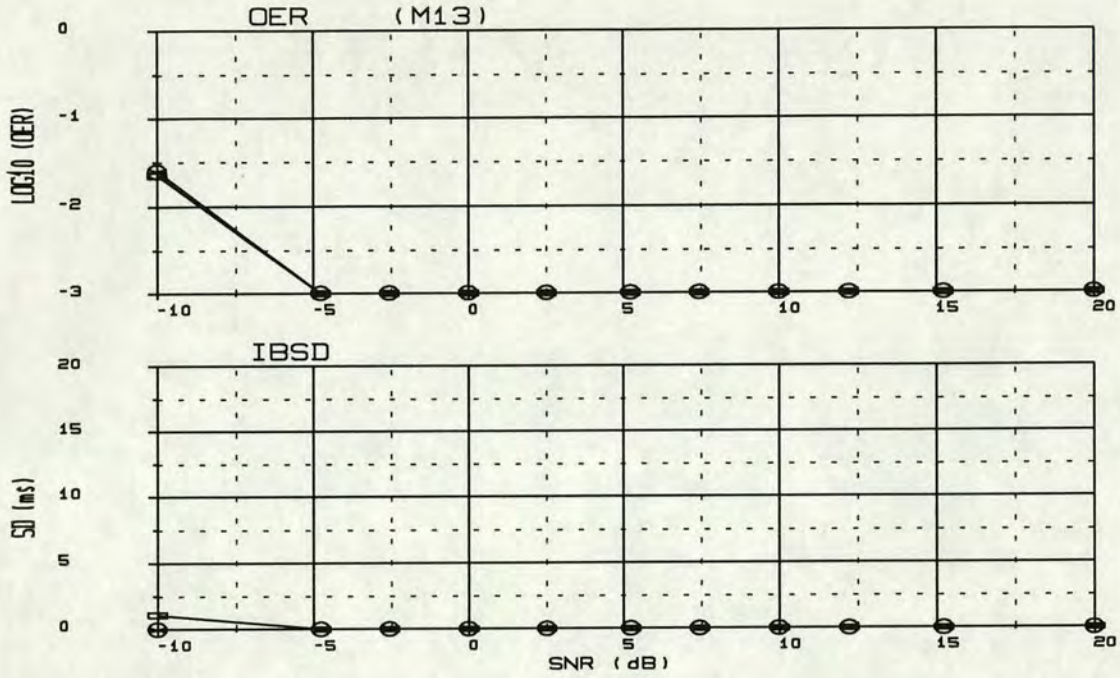

Figure 4.31 Performance of the 'Inverse Least Squares' estimator. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

Excellent performances are obtained from the ILS filter approach when applied to HPW signals corrupted with simulated coloured phono noise. The graphs in figure 4.31 suggest that the SNR operation range of the ILS stretches to below -10dB. Due to the inverse nature of ILS however, it is essential to check both its white noise and real signal performances. Figure 4.32 shows that despite a slight degradation due to the introduction of -10dB white noise, ILS can still maintain a 10% OER down to -7.5dB SNR. The ILS performance when applied to real signals however reveals a more severe deterioration in the performance (example given in chapter 3). This discrepancy is attributed to the low variance of the simulated noise in the band of interest when compared to real observations.

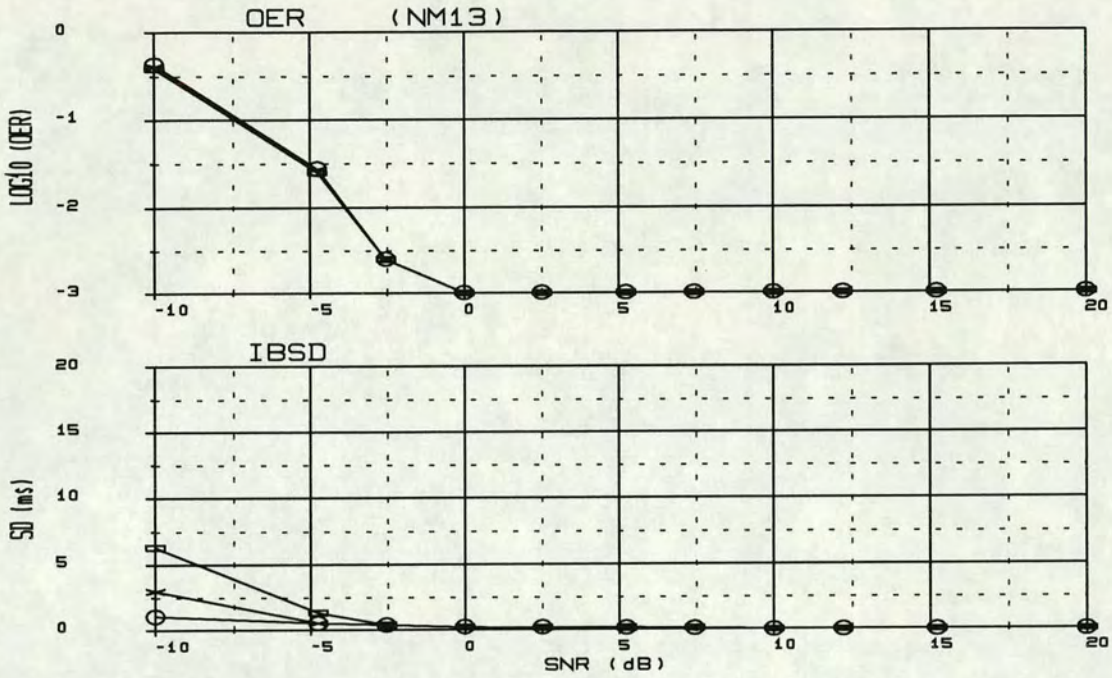

Figure 4.32 'Inverse Least Squares' performance in the presence of fixed power additive white noise -10dB w.r.t. signal power. The x axis measures the SNR w.r.t. the coloured additive noise. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

4.2.16. Modified Block Autocorrelation

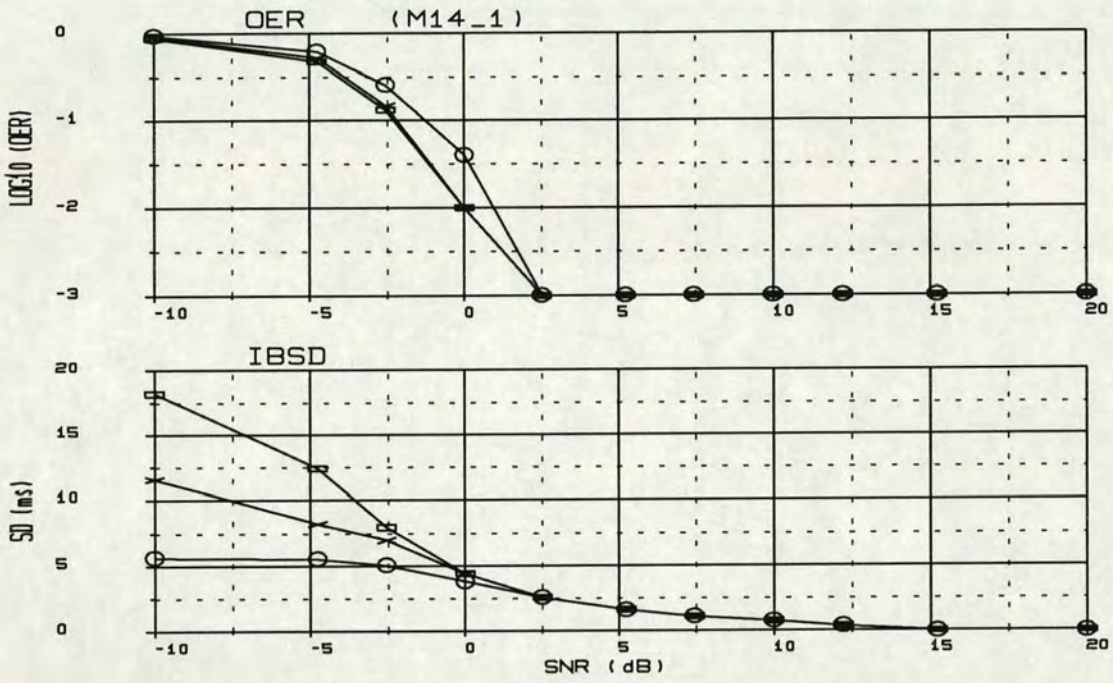

Figure 4.33 Performance of the 'Modified Block Autocorrelation' estimator. The three measured bands are: ±10ms (circles), ±20ms (crosses) and ±40ms (rectangles).

The modified block autocorrelator (MBAC) is derived from the block autocorrelation algorithm. The deviation from the latter is that the time domain output is obtained from the inverse Fourier transform of the linear spectrum rather than the power spectrum. About 1dB improvement in OER over block autocorrelation is obtained. The estimate variance however, is marginally worse (fig. 4.33).

A variety of spectral magnitude mappings were tested (graph results not included), these were the same as the different mappings used with the raw signal block processing (fig. 4.19). Cepstral processing which utilises logarithm spectral mapping has a very poor performance when applied to rectified HPW signals. Other mappings

which extract the n^{th} root from the spectral magnitudes have similar performances to MBAC.

Acceleration Rate (BPM/sec).	SNR=0dB Band, SD and OER	Comment
2 BPM/sec.	±40ms 6.6ms, 8.3%	* Sub-harmonic errors.
5 BPM/sec.	±40ms 8.6ms, 5.5%	* Sub-harmonic errors. * Deviates from correct estimates.
10 BPM/sec.	±40ms 14ms, 5.7%	* Reasonable tracking. * Large deviation from correct estimate.
20 BPM/sec	±40ms 20.5ms, 21%	* Reasonable tracking. * Large deviation from correct estimate.

Table 4.8 Acceleration performance of 'Modified block autocorrelation'. The coloured noise SNR is set to 0dB.

The real advantage of MBAC is apparent when the FHR accelerations performances are taken into account (Table 4.8), these show the MBAC ability to track signals even at high acceleration rates. Although the FHR acceleration of 20 BPM/second (shown above) shows a 21% OER, these erroneous estimates track the accelerations but have an error deviation which fall outside the measured ±40ms band.

The cost of calculating MBAC is equivalent to cost of computing block autocorrelation from the spectral domain. There is however the added cost of computing the square root of the power spectrum.

4.3. Estimators Comparison

This section compares, in a tabulated format, all the algorithms investigated in this chapter. The computation cost of each implementation is based on extracting FHR

estimates at the rate of two per second. The exact time taken by a particular algorithm is not disclosed as it heavily depends on the processor used: Modern DSP microprocessors, such as the MC56000, have significantly reduced the cost of multiplications as well as enabling parallel arithmetic operations to take place. Whereas in the case of general purpose microprocessors, such as the MC68000, the cost of these operations is approximately one order of magnitude higher. Therefore the computational cost is given using the dominant operation as the indicator of cost and is expressed as "of the order of" and annotated as O{number of operations}. A FIR filter for example has a computational cost equal to O{L.K} where L is the filter length and K is the number of samples to be filtered. The fact that each filter operation involves both a multiplication and an addition is not included.

The tabulated comparisons below include the acceleration performance of the block algorithms. In the case of BTB estimators, the performance under accelerating FHR is assumed to be the same as that under constant FHR. Compared to real implementations of the algorithms, no attempt was made to build intelligence into the peak extraction algorithms other than limiting the choice of FHR to the valid range of 60 to 240 BPM. The reason is that the success rates achieved by artificial intelligence algorithms are dependant on the error types introduced by the estimating process. In the case of each estimator, the occurrence of easily detectable errors such as harmonically related errors was highlighted. The parameters effecting the computation cost are listed:

L Template length (seconds) employed by some BTB estimators.

f_{sh} Sampling rate normally used for BTB estimators, usually 1KHz.

f_{sl} Sampling rate normally used for block estimators, usually 50 Hz.

C.Cost The order of computation cost (per second) for implementing a particular algorithm.

N FFT size used in block estimators, usually 4096.

Estimator	SNR, Band, SD and OER C.Cost	Advantages (+), Disadvantages (-)
Time Domain	3dB, ±20ms 1.5ms, 10% O{0}	+ Low processing cost. - Poor performance.
Frequency Matched pre-Filter	O{0}	+ Low processing cost. - Poor OER which is pulse shape dependant. - High estimate variance.
Template Pulse Correlation Estimator	-2.5dB, ±40ms 3.5ms, 10% $O\{L \cdot f_{sh}^2\}$	+ Good OER. + Easy to implement. - High C.Cost.
Pseudo-Inverse Filter. G=60, T=0.15	-6.2dB, ±40ms 3.0ms, 10% $O\{L \cdot f_{sh}^2\}$	+ Very good OER performance. + Low estimate variance. + Not prone to broad-band noise. + Sensitivity control is built into the algorithm. - Normally operated sub- optimally. - High C.Cost.
Envelope Detector-Estimator	+0.5dB, ±20ms 5.0ms, 3.2% O{0}	+ Low processing cost. - High estimate variance.
Inverse Least Squares	-8.5dB, ±40ms 1.0ms, 10% $O\{L \cdot f_{sh}^2\}$	+ Excellent OER. + Very low estimate variance. - Poorer wide-band noise performance. - Poor real signal performance. - Very sensitive to utilised template. - High C.Cost.

Table 4.9 Beat-to-Beat based estimators. The middle column shows the best performance obtained while the OER is less than 10% and the IBSD deviation of the 'in band' estimates is less than 5ms. All BTB based estimators are immune to FHR accelerations.

When the BTB estimators are compared (Table 4.9), it can be clearly seen that both the 'Time Domain' and 'Frequency Matched pre-Filter' provide very poor OER, this is attributed to the lack of prominent signal feature which can be used to pin-point HPW events. This observation is corroborated with observed real phono signals. The 'Inverse Least Squares' estimator although implying a very good performance from

the simulated signals, in practice was found too sensitive to the template from which the filter coefficients were derived.

The 'Envelope Detector-Estimator' is found to perform reasonably well for the small processing effort required to obtain the estimates. Both 'Template Pulse Correlation' and 'Pseudo-Inverse Filter' have similar performances unless the template is continuously updated. In that case, PIF offers a 3.7dB SNR improvement over pulse correlation. Both of the latter estimators have high processing cost and thus require high performance DSP for the filtering stage.

Table 4.10 presents the performance of block based estimators. Both 'Raw Signal Block Autocorrelation' and 'Cepstrum' have a poor FHR tracking performance, in addition, cepstrum processing seems to suffer from broad-band noise reducing its steady FHR performance threshold to -1.2 and 1.0 dB in the presence of -20 and -10 dB white noise respectively. 'Block Autocorrelation' has a good OER but shares with the 'Fundamental Spectral Estimator' a poor FHR tracking performance. The remaining four estimators have good OER. Of these, the 'Iterative Block Clean-up' and the 'Modified Block Autocorrelation' have good FHR tracking capabilities. And both 'Comb Filter' and 'Modulus Difference' have similar performances. The 'Comb Filter' is harder to implement but computationally less expensive than the 'Modulus Difference', especially for newer processors.

Estimator	SNR, Band, SD and OER C.Cost	Advantages (+), Disadvantages (-)
Raw Signal Block Autocorrelation (M5)	-2.3dB, ±40ms 2ms, 10% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Immune from broad-band noise. - Very poor FHR tracking.
Block Rectified Autocorrelation (M6)	-1.0dB, ±40ms 4.5ms, 10% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Good OER performance. + Immune from broad-band noise. + Low C.Cost (FFT/interpolation). - Poor FHR tracking.
Comb Filter (M7)	-2.5dB, ±20ms 4.6ms, 7.9% $O\{f_{sl} \cdot \log_2(f_{sl})\}$	+ Good OER performance. + Immune from broad-band noise. + Good FHR tracking. + Low C.Cost (via interpolation).
Cepstrum (M9)	-3.5dB, ±40ms 2.0ms, 10% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Very good OER performance. + Excellent low estimate variance. - OER is prone to broad-band noise. - High C.Cost. - Very poor tracking of FHR accelerations.
Modulus Difference (M10)	-2.5dB, ±40ms 5ms, 6.3% $O\{f_{sh}^2\}$	+ Good OER performance. + Immune from broad-band noise. + Good FHR tracking. + May be amenable to C.Cost reduction. - High C.Cost for modern DSP's.
Fundamental Spectral Estimator (M11)	+6.3dB, ±20ms 4.5ms, 10% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Easy to implement. + Immune from broad-band noise. - Poor OER performance - Estimates are prone to low frequency bias. - Poor FHR tracking.
Iterative Block Clean-up (M12)	-1.5dB, ±10ms 3.7ms, 10% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Reasonable OER. + Immune from broad-band noise. + Reasonable FHR tracking. + C.Cost can be very low.
Modified Auto-Correlation (M14)	-1.0dB, ±40ms 5.0ms, 2.5% $O\{f_{sh} \cdot \log_2(f_{sh})\}$	+ Reasonable OER. + Immune from broad-band noise. + Good FHR tracking. - High FHR variance.

Table 4.10 Block based estimators. The middle column shows the best performance obtained while the OER is less than 10% and the standard deviation of the 'in band' estimates is less than 5ms.

4.3.1. Comparison of Block and Beat-to-Beat Estimators

It is clear from the two tables above that of the estimators investigated, BTB estimators provide a lower estimate variance than block estimators. This quality is a prerequisite for FHR variability monitoring. The 'Pseudo-Inverse Filter' estimator is selected for its all round performance. In addition, the low OER afforded by PIF processing enables the use of phono signals with SNR down to -6dB. This performance level is not attained by any of the block based estimators. The use of BTB estimators, specifically 'PIF', requires: (a) The continuous updating of the template to offset HPW pulse changes and (b) the derivation of short term SNR estimate which is used in deriving the PIF coefficients. Two factors favour the use of block estimators, firstly the computation cost can be significantly lower and secondly BTB estimators require considerably more powerful artificial intelligence algorithms to enable their proper operation.

CHAPTER FIVE

5. POST PROCESSING AND RESULTS

The signal processing development facility and hardware configuration are detailed in this chapter. Results obtained using this facility from overnight phonocardiographic recordings[78] using the Talbert and Southall[18, 66] transducer are presented. The algorithm implemented for FHR estimation is the comb filter estimator which was discussed in chapter 3. A variety of parameters are then extracted from the FHR time series. These parameters may be averaged over the whole observation period or displayed in conjunction with the FHR baseline in continuous trace format to reveal the dynamic biophysical profile variables.

5.1. Hardware Configuration

A four channel, frequency modulated recorder enables the simultaneous registration of multiple transducers with good low frequency reproduction[78]. The recorded fetal MAW signals facilitate the testing of both analogue circuits and software modules.

The need for flexibility in the hardware configuration led to selection of a VME based bus system. This consists of a rack with a backplane to which module cards conforming to the VME bus protocol can be attached. VME cards can operate in

either `bus master` or `bus slave` modes depending on their function. Each card is assigned a base physical address which other cards (placed on the VME backplane) may access. Bus communication is orchestrated by a prioritised VME bus controller which also handles interrupt oriented data transfers. As long as the various cards adhere to the bus timing specifications, the VME bus can be thought of as an extension of the memory address space.

Figure 5.1 shows a block diagram of the three VME boards used with this project. They consist of a main processor card, a data acquisition card and an additional RAM card. The main card has a MC68000 microprocessor running at 8Mhz, 128K byte of RAM, two serial ports (RS232), a monitor in EPROM and the VME bus controller/arbitrator. This main card is responsible for configuring the VME system, organising data transfers and signal processing. The memory card (fig. 5.1) contains 512K byte of battery backed static RAM which is used to: (a) Store FHR estimates, (b) keep hardware resident library routines/modules, (c) hold pre-computed data vectors and (d) act as a silicon tape recorder which stores fetal MAW signals for fast re-processing. The data acquisition card has 16 channel of differential analogue inputs with a sampling bandwidth 10,000/N (N is the number of channels used), the ADC has 12 bit resolution and the capability of interrupting the main processor at the end of the conversion cycle. Two DAC channels which are mainly used to display continuously updated traces on the oscilloscope are also available on this board. The analogue conditioning unit (fig. 5.1) contains various modules which include high pass filters, low pass filters, full wave rectifiers and gain amplification stages.

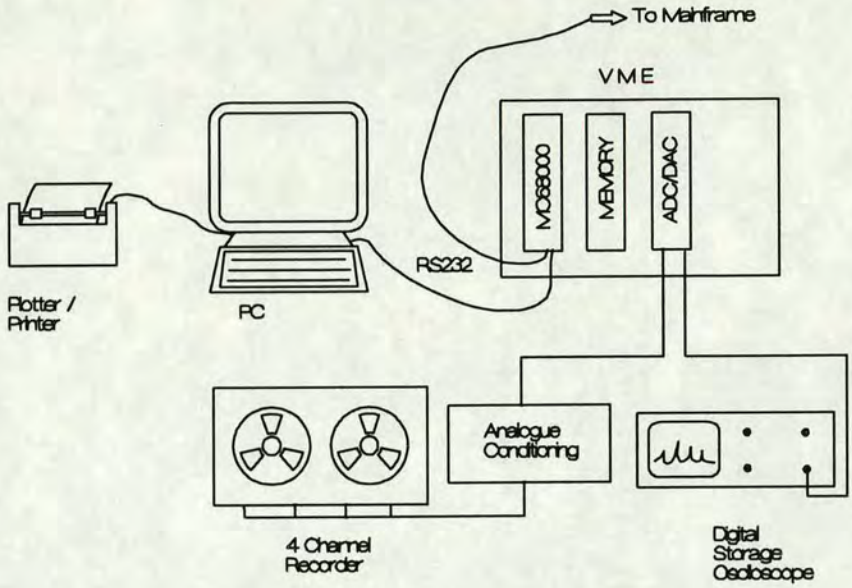

Figure 5.1 Block diagram of the system hardware configuration.

5.2. Software Development

Personal computers (PC) based on either UNIX or MS-DOS operating systems provide suitable platforms for software development. Programmes written in the 'C' language are initially developed and tested on the host system where good display, plotting and debugging facilities are available. Specific signals can be used to test signal processing modules, and algorithm parameters are passed to modules via their argument lists. The performance of the processing can then be scrutinised for the results thus obtained. For example:

COMB -f signal.in -o comb.out -n 100 -N 4096 -s 50 -d

uses the file `**signal.in**` as input and produces the comb filter output in the file `**comb.out**`. The block size is set to 100 samples (`**-n 100**`) and FFT size is set to 4096 (`**-N 4096**`). The sampling rate is assumed to be 50 Hz (`**-s 50**`) and the programme is run in debugging mode (`**-d**`).

Once proven working, the module programmes are incorporated into software libraries. A controlling shell programme is then cross-compiled in 68000 machine code and linked into a single object code file. This is then downloaded to the main processor board on the VME system for execution. The FHR time series estimates are sent back to the PC for further processing.

5.3. Least Mean Squares Predictor and Outlier Detector.

Residual estimate errors (outliers) should be removed from the FHR time series before any analysis can be carried out[49]. If left unchecked, these errors induce (a) a large bias in the short term FHR average and (b) artificial FHR variance whose level exceeds that of natural biological rhythms. The algorithm described here assumes that the FHR outliers have no cyclic patterns (i.e. are randomly distributed).

The least mean squares (LMS) adaptive FIR filter algorithm operates in an error feedback mode to track the Wiener optimal solution[58, 59]. The governing equations are:

$$\hat{x}(n) = \underline{h}^T(n) \cdot \underline{y}(n)$$

$$\underline{e}(n) = d(n) - \hat{x}(n)$$

$$= d(n) - \underline{h}^T(n) \cdot \underline{y}(n)$$

Where $\hat{x}(n)$ is the output of the filter at time n, d(n) is the desired signal, the vector $\underline{h}(n)$ is the adapted filter weights, and e(n) is the output error. Taking the expectation of the error squared and differentiating w.r.t. $\underline{h}(n)$:

$$\frac{\delta e(n)}{\delta \underline{h}(n)} = -2\, E\{ \underline{y}(n)\, e(n) \} \qquad 5.1$$

Equation 5.1 shows that the gradient vector of the error surface can be estimated from the expected value of $\{\underline{y}(n)\ e(n)\}$. The least mean squares algorithm utilises this derived gradient estimate to recursively seek out the global minimum error. The recursion on $\underline{h}(n)$ is given:

$$\underline{h}(n+1) = \underline{h}(n) + 2\,\mu\,\underline{y}(n)\,e(n)$$

The convergence constant μ controls the adaptation speed; a small μ value slows the adaptation time while a larger μ creates significant time discontinuities in the FIR weights $\underline{h}()$. The characteristic convergence of the LMS is determined by (a) the non stationarity of the signal $\underline{y}()$ and (b) the eigenvalue spread of its autocorrelation matrix. More detail on the derivation of the LMS algorithm and performance comparisons with faster converging adapting filter structures such as the recursive least squares is given in the literature[10, 58, 60]. Commonly, μ is set to:

$$\mu = \frac{\alpha}{N \,.\, P}$$

where $0 < \alpha < 1.0$, N is the number of filter weights and P is the power of the signal.

Used in a forward predictor configuration, the LMS filter weights converge towards the optimal filter for the observed signal. Any unpredictable noise in the input signal is thus observed as the error e(). In terms of FHR processing, the estimates are fed into the LMS predictor which tries to track the underlying FHR time series trends. In such a configuration, random erroneous FHR estimates cannot be predicted and result in large errors.

The novelty in the LMS outlier detector compared with traditional adaptive filter implementations lies in the way estimated outliers are handled. Input FHR samples producing errors exceeding a predetermined error threshold level are: (a) flagged as erroneous and (b) replaced by the predictor's output $\hat{x}()$ plus a small error margin. For

example, if the desired signal d(n) has an error of +50 BPM, which is a very unlikely FHR step increase, the input sample y(n) is replaced by the predicted output $\hat{x}(n)$ plus a limited error (2 BPM in the direction of the outlier). Therefore having decided that an estimate is erroneous, the best policy to adopt is to replace it with the LMS forecasted value thus avoiding the propagation of an outlier through the filter structure. A limited error margin must always be added to the LMS output when replacing an assumed outlier otherwise the filter gets locked into its current state and genuine FHR step changes are then missed.

The forward predictor configuration (discussed above) is quite successful in detecting some FHR outliers. However, it has a poor performance in tracking FHR accelerations and decelerations. This is related to its inability to 'observe' signs indicating the onset of fast FHR changes. A similar performance is obtained when a backward predictor configuration is implemented.

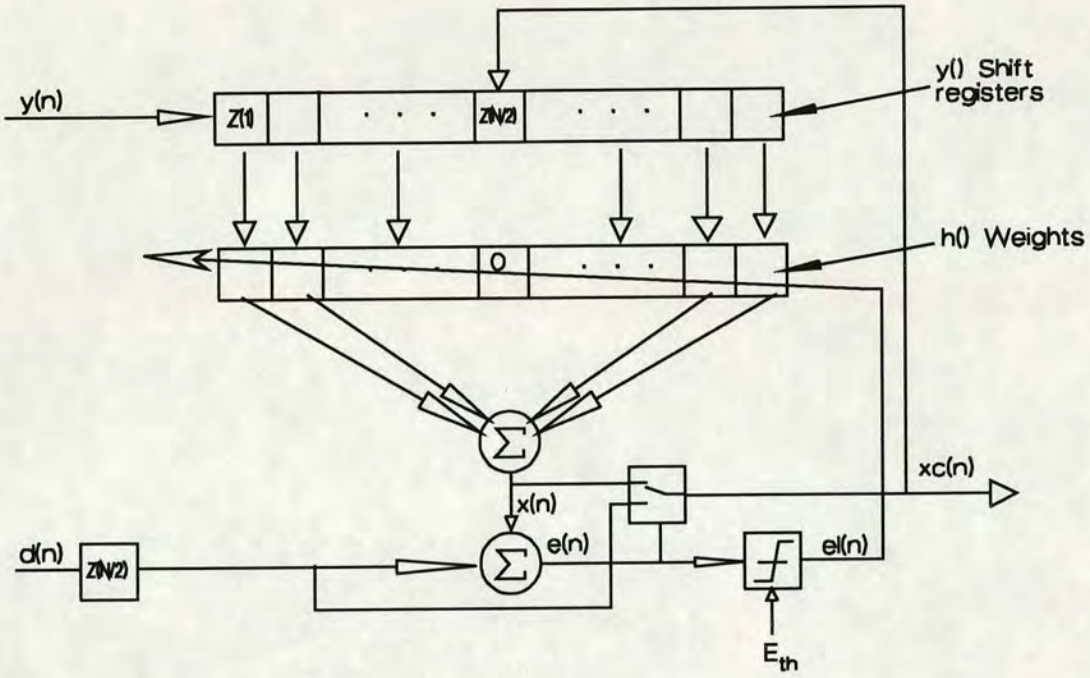

Figure 5.2 Block diagram of the combined forward and backward LMS configuration implemented for outlier detection. The tap weight h(N/2) is set and held at zero.

A suggested LMS structure which combines both forward and backward predictors is depicted in figure 5.2 where: el(n) is the clipped error, Z(L) denotes a time delay of L sample periods, E_{th} is the error threshold chosen for the configuration stage and xc(n) is the corrected FHR estimate. If the error e(n) is lower than the predetermined threshold E_{th}, xc(n) is obtained from the delayed $\underline{d}(n - N/2)$, otherwise it is taken as the addition of $\hat{x}(n)$ plus a limited error value (2 BPM). In either case xc(n) is fed back into the FIR structure. The centre tap $\underline{h}(N/2)$ is set to zero and is not updated. The forward backward predictor corrector (FBPC) (fig. 5.2) offers considerable improvement over the normal LMS forward predictor, especially during fast FHR changes, as a result of its ability to 'look' ahead of the current time sample.

While most outliers are detected and corrected using the FBPC implementation above, the occurrence of a large erroneous FHR estimate, 220 BPM for example,

when the neighbouring baseline is 120 BPM can lead adjacent correct estimates to be flagged as erroneous. The presence of the 220 BPM outlier in the lower half of the y() shift registers (1 .. N/2) causes the filter to predict high FHR values which results in the incorrect labelling of the preceding sample(s) as outlier(s). Once the 220 BPM outlier reaches the centre tap however, it is both detected as outlier and corrected.

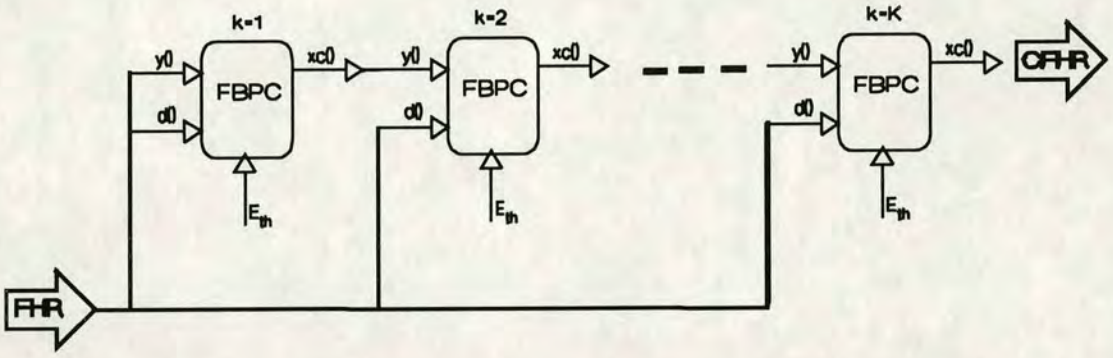

Figure 5.3 Cascade implementation of the 'Forward Backward Predictor Corrector'.

The solution to such large outliers disrupting adjacent estimates is to cascade a number of FBPC stages (fig. 5.3). The error threshold E_{th} is set high for the first stage and progressively lowered to the desired level at the K^{th} (last) stage. The benefit of using the cascade of predictor corrector stages is illustrated using an example: At the first FBPC stage a large outlier at time (n) causes its preceding estimate d(n-1,1) to be falsely detected as an outlier and corrected to xc(n-1,1). The erroneous sample d(n) is then detected and corrected to a value xc(n,1) closer to the true underlying trend. A later j^{th} FBPC stage having the hind sight of the corrected outlier xc(n,1) predicts an output $\hat{x}(n-1,j)$ which is close to d(n-1,j) and hence the FBPC(j) selects d(n-1,j) rather than $\hat{x}(n-1,j)$ as its output xc(n-1,j). A further FBPC

stage using the true estimates adjacent to the error at (n), produces an output $xc(n,j+1)=\hat{x}(n,j+1)$ which is more consistent with the background FHR trends.

In practice, it was found that 4 to 7 FBPC stages were sufficient to detect the all observed large outliers and most of the smaller ones. The implementation in this case makes use of 7 FBPC stages, and the error threshold is progressively set to 40, 30, 20, 10, 10 and 10 BPM.

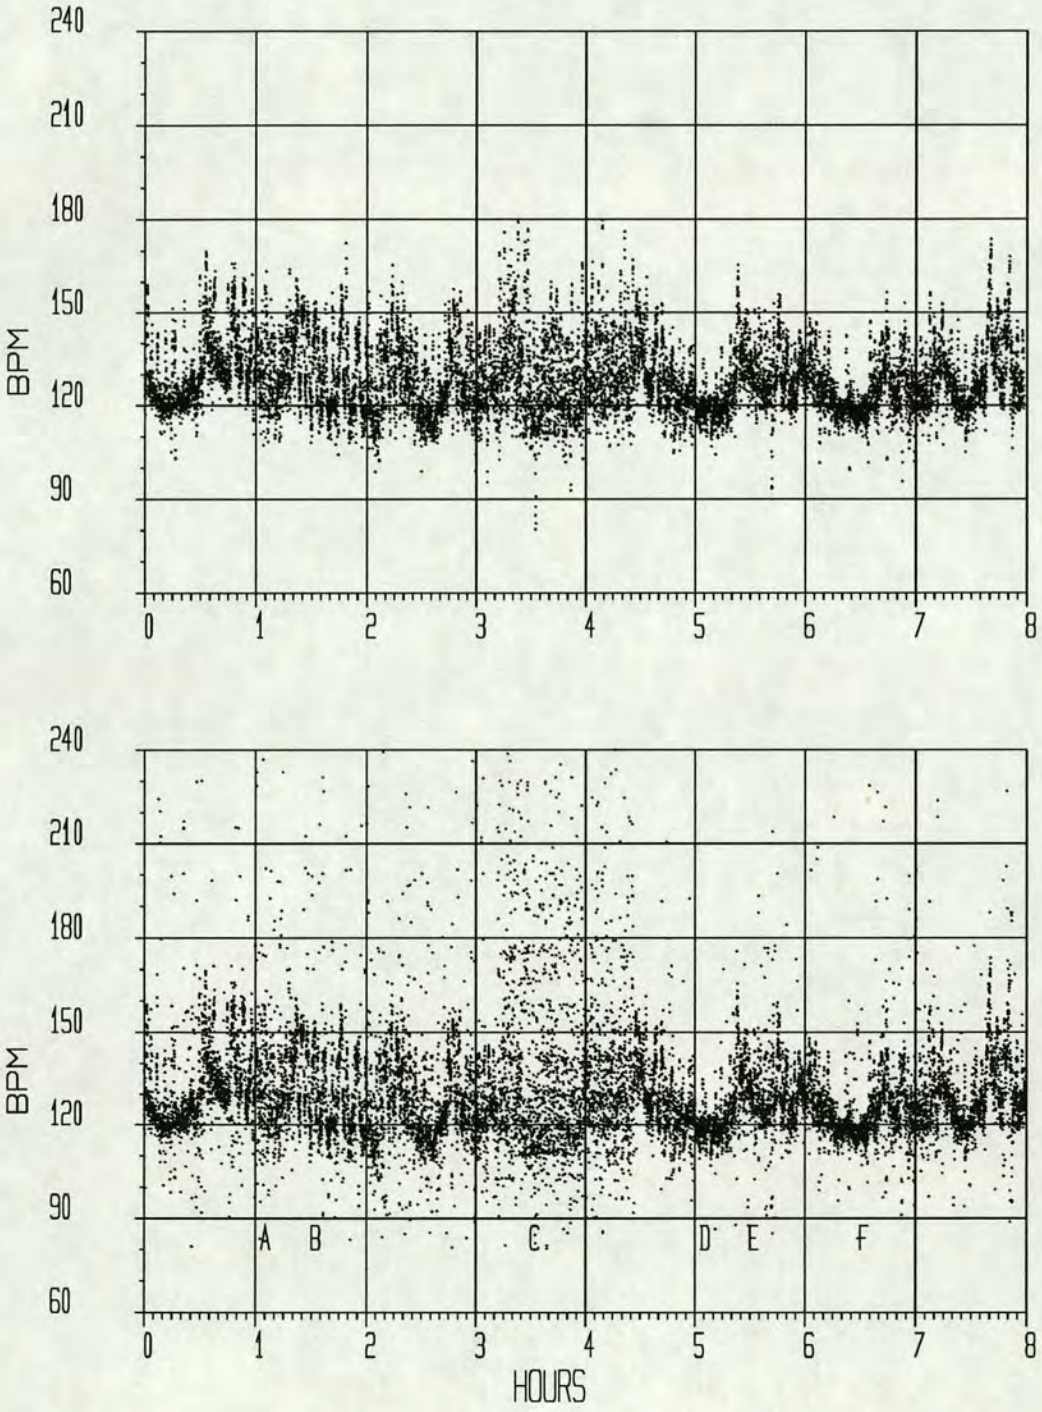

Figure 5.4 Bottom: Results of eight hours of continuous fetal phono processing, each dot represents the average FHR of a two second epochs.
Top: The result of using seven FBPC stages to correct the FHR time series (shown at the bottom).

Figure 5.4 displays 8 hours of FHR time series obtained from applying comb filter processing, each estimate is obtained from 2 second block. A seven stage FBPC is then applied to the time series and a resulting OER average of 14.6% is obtained. This high OER is not evenly distributed over the recording, markers A,B,C,D,E and F have been superimposed on figure 5.4, these have respective OER rates of 24.5%, 10.5%, 33%, 0.5%, 3% and 2%.

The regions pin-pointed by the markers B,C and E (in fig. 5.4) are expanded to show records of 6 minutes duration (fig. 5.5), these reveal episodes of average, high and low error rates respectively. Large errors for example at time=3.55 hours (fig. 5.5, middle trace) are corrected while true large FHR changes (at time 3.54 hours) are tracked without modification. In addition, corrections as instituted by the FBPC visually follow the underlying FHR trend.

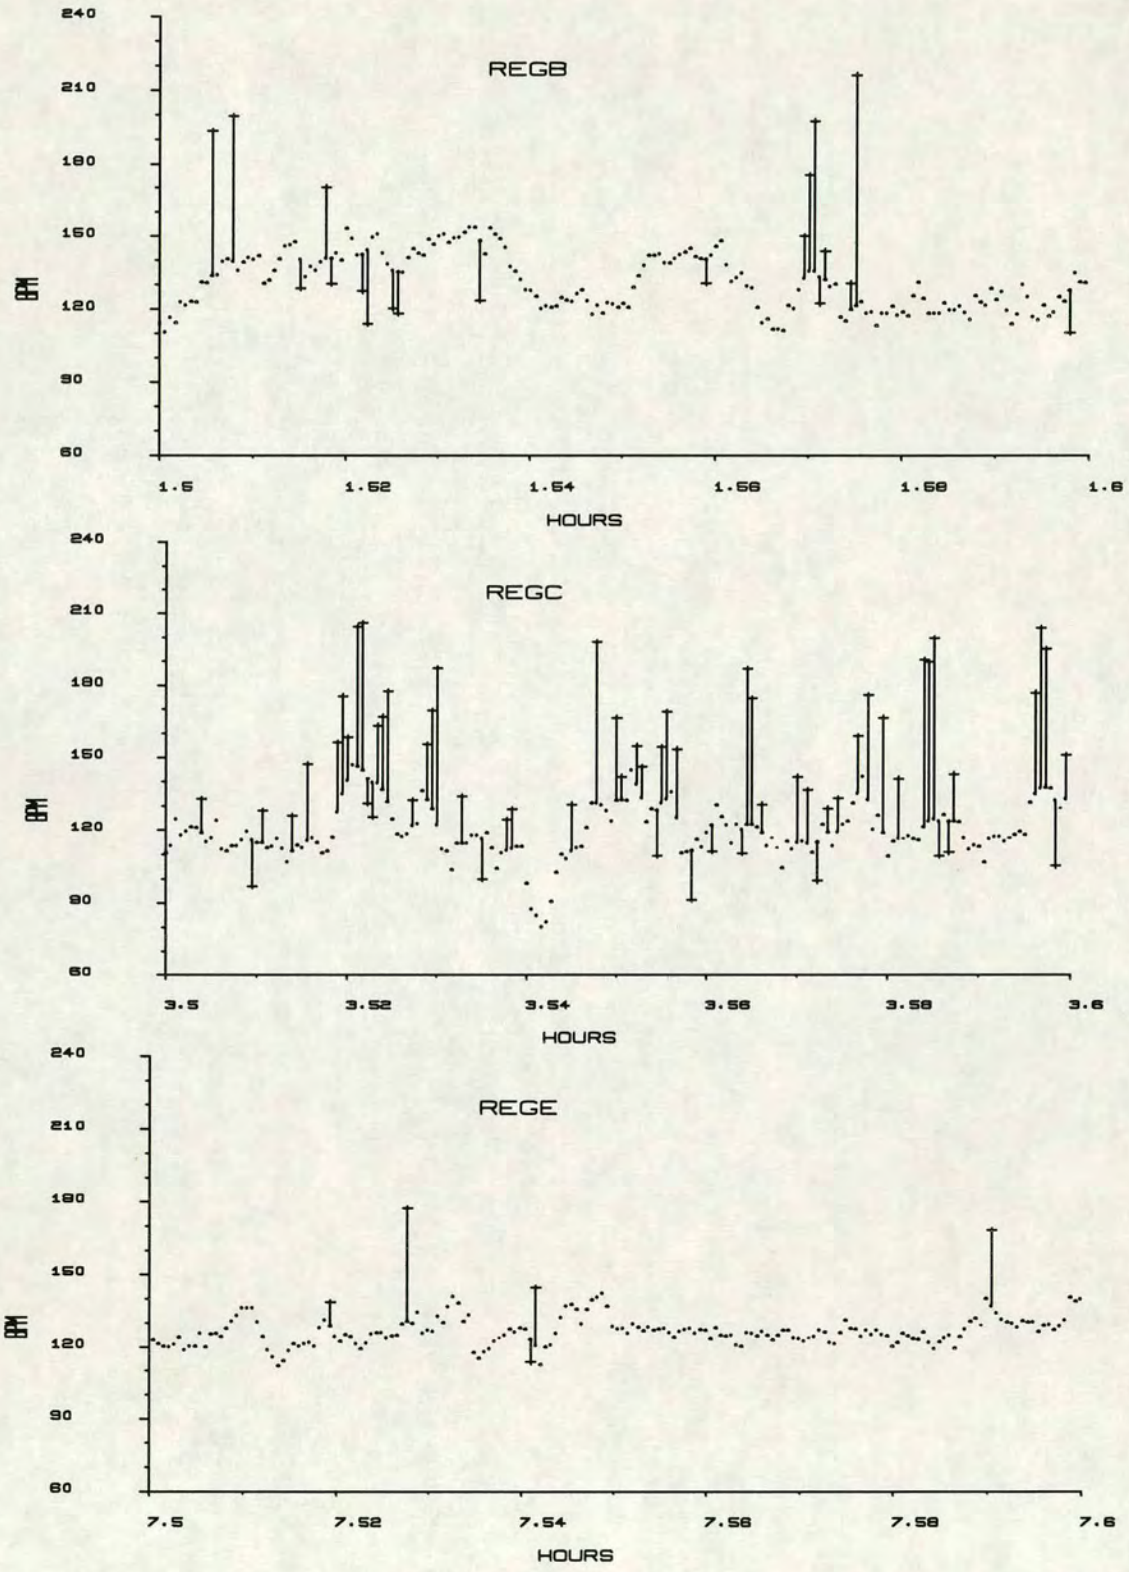

Figure 5.5 Corrections instituted by the LMS predictor, (expanded records from figure 5.4). Top, middle and bottom traces correspond to regions B, C and E and have respective estimated OER's of 10.5%, 33% and 3%.

5.4. The Extraction of Fetal Heart Rate

Currently the main FHR parameters associated with fetal monitoring are (a) long term FHR (or baseline), (b) medium and short term variabilities and (c) accelerations and decelerations. The 'non-stress test' (NST) refers to such parameters extracted from the passive monitoring of fetal heart rate. This section applies digital signal processing techniques to eight hours of continuous FHR data in order to extract these and other parameters. Although the derivation of the parameters described below is processed off-line, the algorithms can be integrated into stand alone monitors to provide continuous on-line FHR parameter update.

5.4.1. FHR Baseline

Baseline is the most commonly used FHR parameter in the assessment of antenatal fetal health[13]. Baseline is generally accepted as the low-pass filtered FHR trace. Currently however, no standards exist which define the frequency components to be included in the baseline time series. Dalton *et al*[54] suggested a first order recursive filter having an approximate cut-off frequency at 1mHz. This filter is applied successively in the forward then backward direction leading to the cancellation of non-linear phase shifts present in the filter. Such phase shifts would otherwise cause some of the FHR components to suffer from group delay effects.

A normal fetal heart trace has a rate between 120 and 160 BPM, and is combined with other trace characteristics to form an overall picture of the fetal health status[63]. When taken into account[13], the baseline trace is described as either (a) a 'reactive trace' when variations in both the medium and long term are observed or (b) a 'non-reactive (or flat) trace' when the FHR trace is featureless.

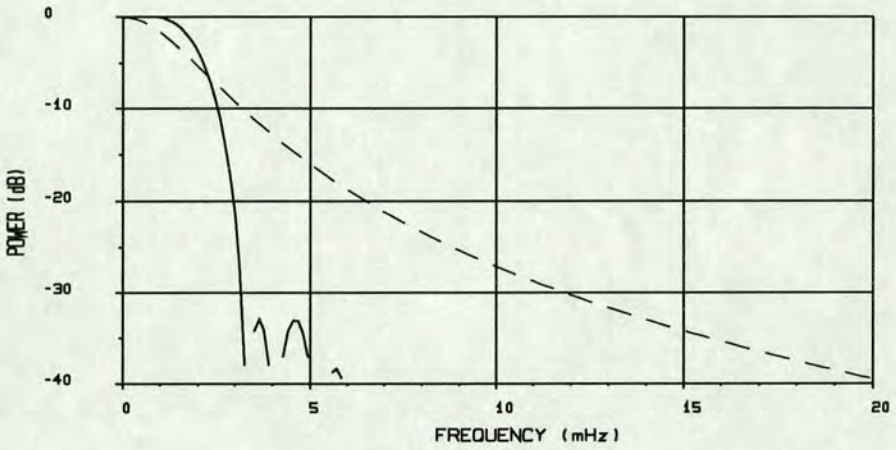

Figure 5.6 Frequency response of the first order filter proposed by Dalton *et al*[54] (dashed trace) and FIR filter proposed here (solid line).

While the filter suggested by Dalton[54] appears to be adequate when inspected visually, it has a poor transition between the pass and rejection bands (fig. 5.6) resulting in a section of the spectrum where FHR components are only partially included in the baseline signal. The filter advocated here is a simple FIR derived from a seventh order Butterworth filter. The filter design uses the frequency transfer function synthesis approach[9] and deviates from the standard Butterworth prototype by having a zero (linear) phase shift across the spectrum. The tap weights thus obtained from the inverse Fourier transform are acausal, therefore a time shift equivalent to half of the filter length is introduced (in this case 127 taps). Figure 5.6 shows that from an engineering point of view the FIR has a sharper transition

between the FHR frequency components included in the baseline signal (pass-band) and those excluded (stop-band). Using the bilinear transform method, standard frequency responses can be obtained from higher order IIR filter configurations[9]. These as well as producing better transfer functions are computationally more efficient than FIR designs.

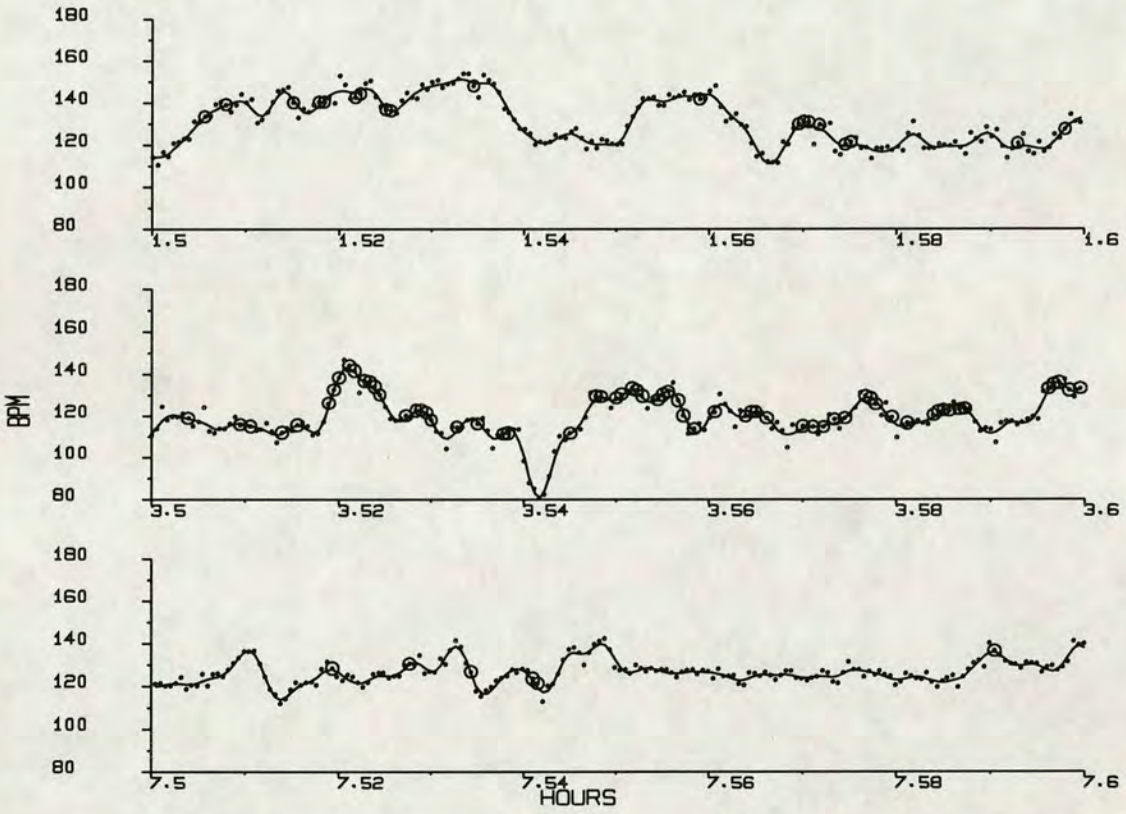

Figure 5.7 Corresponding to bands B, C and E of figure 5.4: Short term baseline obtained from low pass filtering at 67mHz (T=15 seconds). The dots show the estimates prior to filtering, those enclosed in circles are corrected estimates and the solid line is the filter output. The corrected estimates are seen to be very close to the baseline.

In the absence of standard cut-off frequencies for the baseline generation, two results are included here. Firstly the short term baseline (fig. 5.7) which has a cut-off at 67mHz (or a time constant T=15 seconds). Secondly the long term baseline which is

filtered using a cut-off at 1.7mHz (T=10 minutes); this is obtained by initially low pass filtering the FHR estimates to 17mHz (T=1 minute), sub-sampling, and then low pass filtering the resulting time series (fig. 5.8).

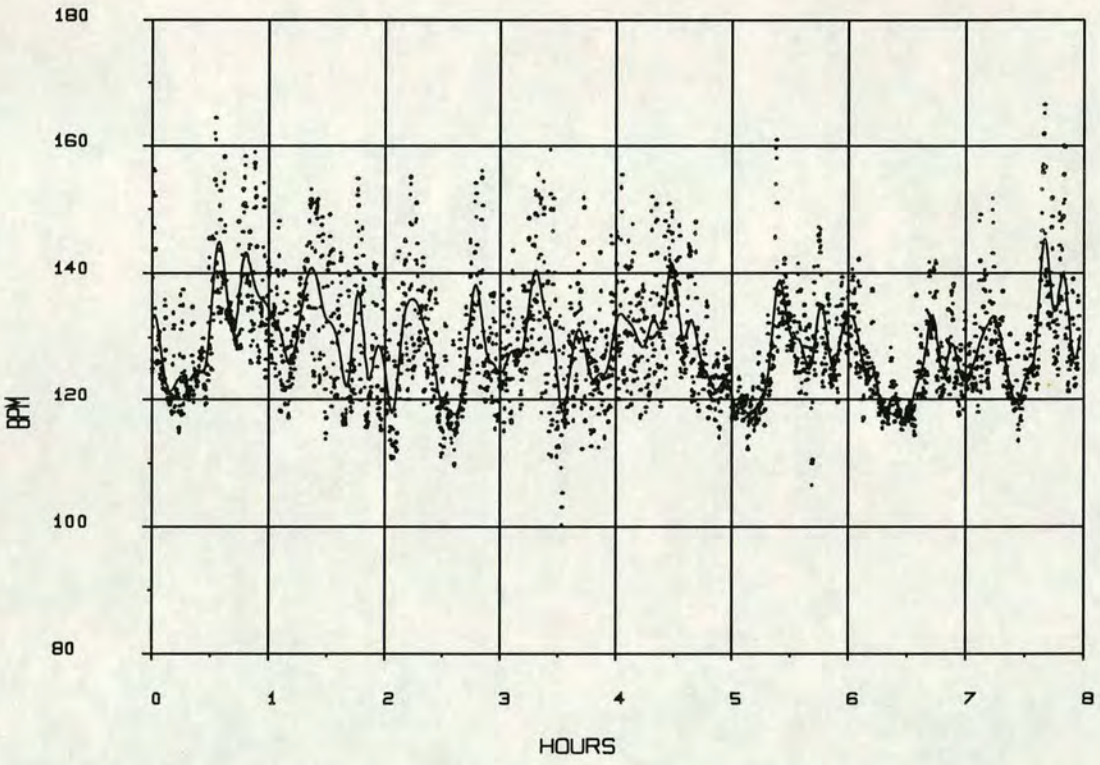

Figure 5.8 Medium term baseline (dots, T=1 minute) and long term FHR baseline (solid trace, T=10 minutes).

5.4.2. FHR Accelerations and Decelerations

FHR accelerations have been related to fetal activity[19] such as breathing and movement and are indicative of fetal well-being. Accelerations are understood to be short term FHR deviations from the baseline time series. There does not seem to exist a formal definition (in the engineering sense) for either the duration or the depth of

deviation which qualify these as significant. Duration figures of a few seconds and amplitudes in the region of 15 BPM have been used[19]. The approach taken here is to band-pass filter the FHR starting from the cut-off frequency of the baseline extraction filter to the maximum frequency components desired.

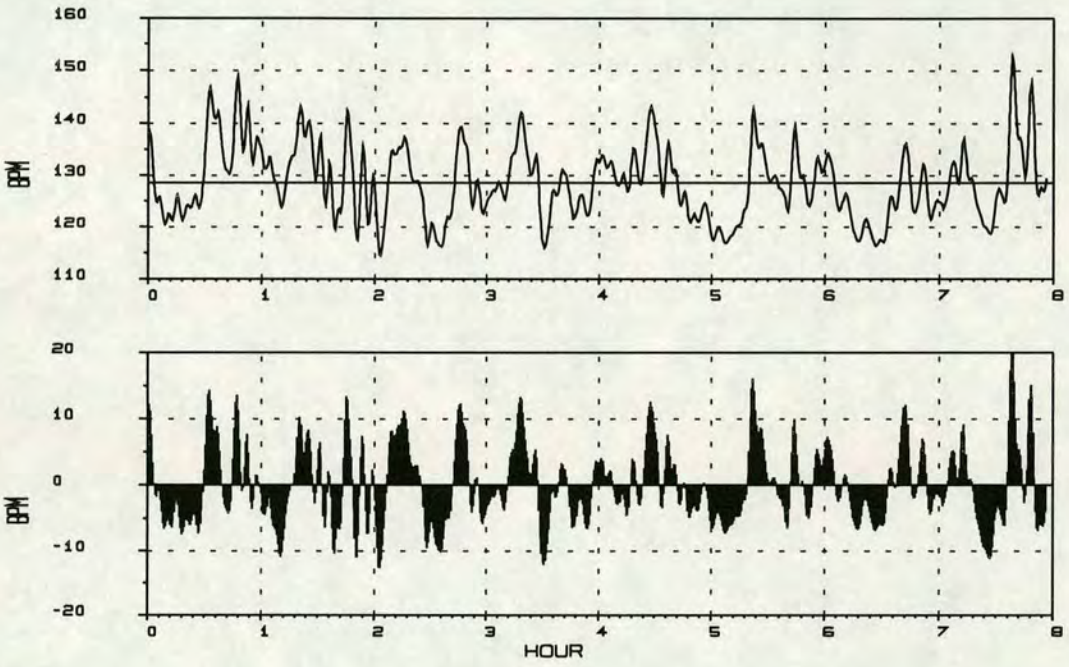

Figure 5.9 Top: 3 mHz baseline (T=5 minutes), the horizontal solid line (at 128.6 BPM) shows the FHR average for the whole 8 hours.
Bottom: FHR Band pass filtered from 0.185 mHz to 3mHz (T=90 .. 5 minutes) to reveal accelerations and decelerations from the 90 minute baseline.

5.4.3. FHR Variability and the Variability Amplitude Index

Studies performed on both fetuses and newborns have demonstrated a strong correlation between heart rate variability and fetal behavioural state[25, 60, 68]. Such switching between high and low FHRV is linked to periods of high fetal activity or

fetal resting state respectively and was found by Dawes *et al*[67] to be a better indicator of fetal well-being than other indicators. While current clinical practice relies on fetal monitors based on block estimators to assess the heart rate variability (medium term variability), most research is concentrating on the extraction of FHRV indices from beat-to-beat measurements[20, 22, 27, 29, 64] (short term variability). The relationship between the medium and short term variabilities has not to-date been resolved although the variation of each has been independently linked to fetal behavioural states. Medium term variabilities with periodicities from a few seconds to a few minutes are commonly utilised in the assessment the fetal health[73]. A normal variability reading is universally accepted as a better indicator of fetal well being state than is a low variability reading of fetal compromise[19, 21, 81]. This is attributed to the often prolonged 'quiet' state of the fetus which can last for over an hour. It has been reported[88] that after 45 minutes of fetal rest 95% of fetuses returned to normal FHRV levels.

The short term non-stationarities of the US Doppler signal (Manning[38]) render the detection of individual heart beat events useless for beat-to-beat measurements[40, 73]. Most of BTB variability measurements are therefore limited to scalp ECG or phonocardiography as suggested by Trierweiler *et al*[21]. In fact, the artificial variability introduced by some phono BTB based estimators such the Pseudo-inverse filter (detailed in ch. 3) was shown to be less than that introduced by most block estimators (ch. 4). It can therefore be concluded that unless an important improvement in the registration of external abdominal FECG can be obtained, fetal phono may be the only practical method for BTB FHRV measurement. Otherwise, only medium term FHRV is available from block estimators.

Four performance criteria for FHRV indices have been suggested by Parer *et al*[15], these are:

A - Short term FRHV should have an amplitude linearly dependant on the BTB interval difference.

B - Long term FHRV should be proportional to the amplitude of FHR variations.

C - When applied to sinusoidal FHR, long term FHRV should have an amplitude proportional to the frequency of FHR variations.

D - Long term FHRV should be independent of short term FRHV.

Parer also provided a comparison of indices commonly used in BTB variability[20, 27, 29]. One index proposed by Huey[55] was found to satisfy all of the above criteria. Detwiler *et al*[74] showed that the random error in the variability measurement introduced by the limited observation period can be reduced to ±10% with 95% confidence if the observation period is extended to five minutes. This result however relied on simulated stationary variability which is not commonly observed in real fetal heart rate time series.

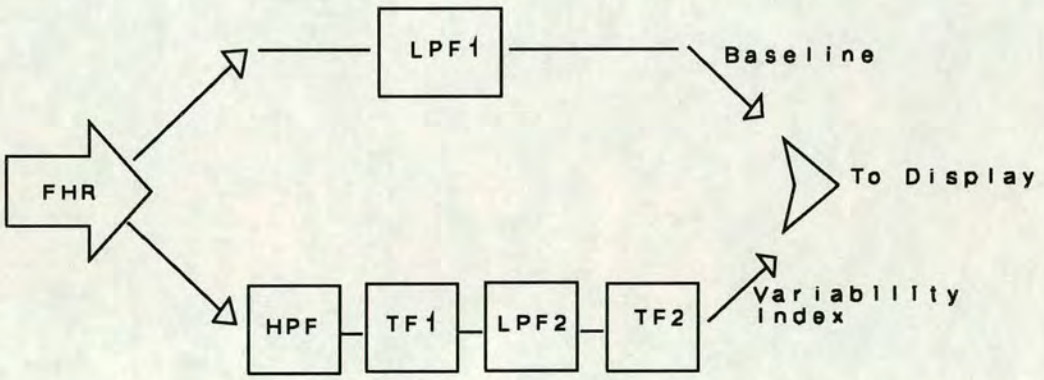

Figure 5.10 Processing stages utilised in obtaining FHR baseline and variability indices. TF1and TF2 are dependant on the type of index.

Figure 5.10 shows a typical block diagram for extracting FHR variability indices, in the case of variance measurement for example; HPF is a high pass filter with a corner frequency at 17mHz, TF1 is the square operator and LPF2 a five minute block average, TF2 would have a unity transfer. In the case of standard deviation measurement, TF2 is set to the square root operator. A proposed variability measurement which provides a linear index approximating the visually assessed FHR trace is the 'Variability Amplitude Index' (VAI).

The VAI parameter is a continuously sampled index which can be displayed in conjunction with baseline traces to provide a visually interpolated FHRV block average. The VAI index extraction is based on using the magnitude operator combined with a gain factor for the TF1 stage. The gain compensates for the peak to peak amplitude loss when the output of TF1 is low pass filtered. LPF2 is a zero phase-shift low pass filter with five minutes time constant, TF2 is not used in the derivation of VAI.

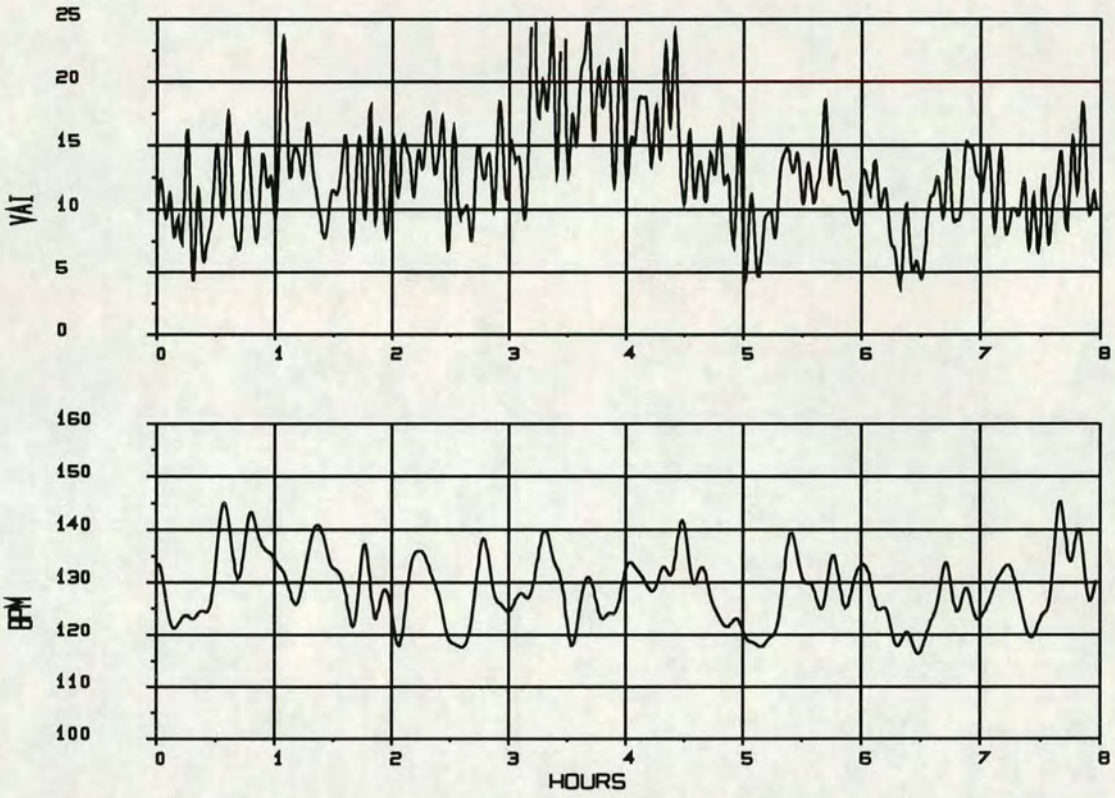

Figure 5.11 Bottom: FHR long term baseline 1.7 mHz (T=10 minutes). Top: 1minute variability which is low pass filtered with a time constant of five minutes.

An added benefit of the VAI as an index is that a single reading from the VAI trace (fig. 5.11) provides the variability measurement as assessed by clinicians with an effective averaging over the neighbouring five minutes of FHR. Also seen in figure 5.11 is the correlation between medium term FHRV and long term baseline, this correspondence is better displayed in scatter diagram format in figure 5.12. The results illustrated in figure 5.12 shows a positive correlation between the medium term VAI and long term heart rate, this agrees with the long term variability index proposed by Van Geijn *et al*[64, 68].

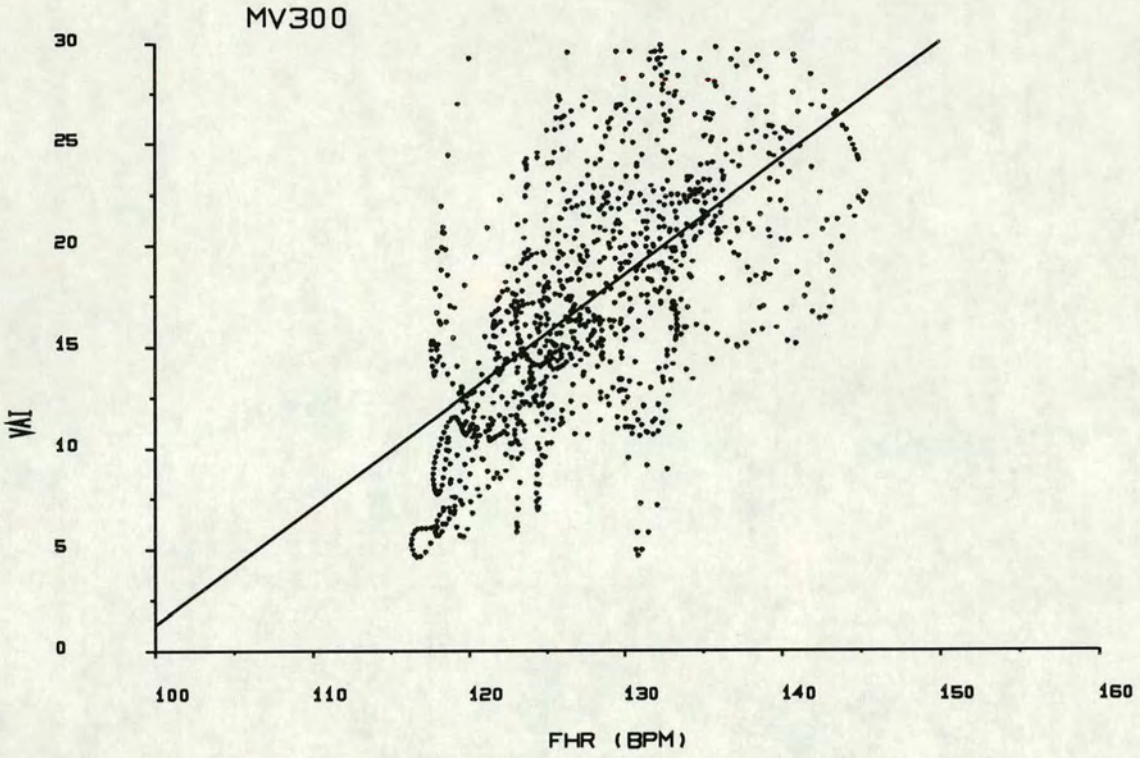

Figure 5.12 Plot showing the relationship between the medium-term VAI and the FHR baseline. A least squares straight line is fitted to the data.

5.4.4. Rate and Spectral Distribution

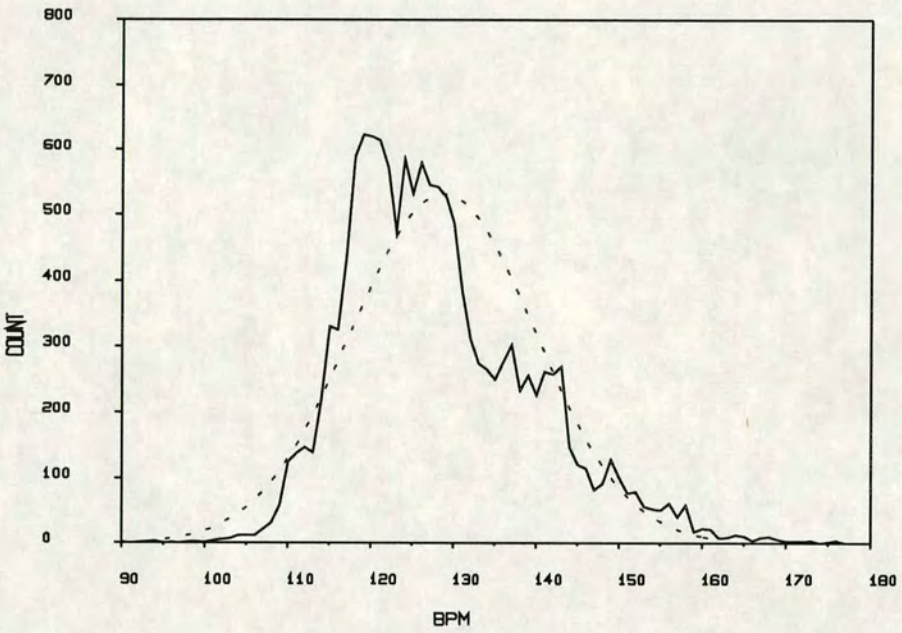

Figure 5.13 Distribution of FHR estimates (solid trace) with a superimposed normal Gaussian distribution (dotted line).

Close inspection of the heart rate distribution (fig. 5.13) reveals significant skewness. This suggests an inhibiting mechanism which prevents the heart rate from falling below a certain minimum. The same control mechanism is not seen at high FHR where the FHR distribution has a longer tail.

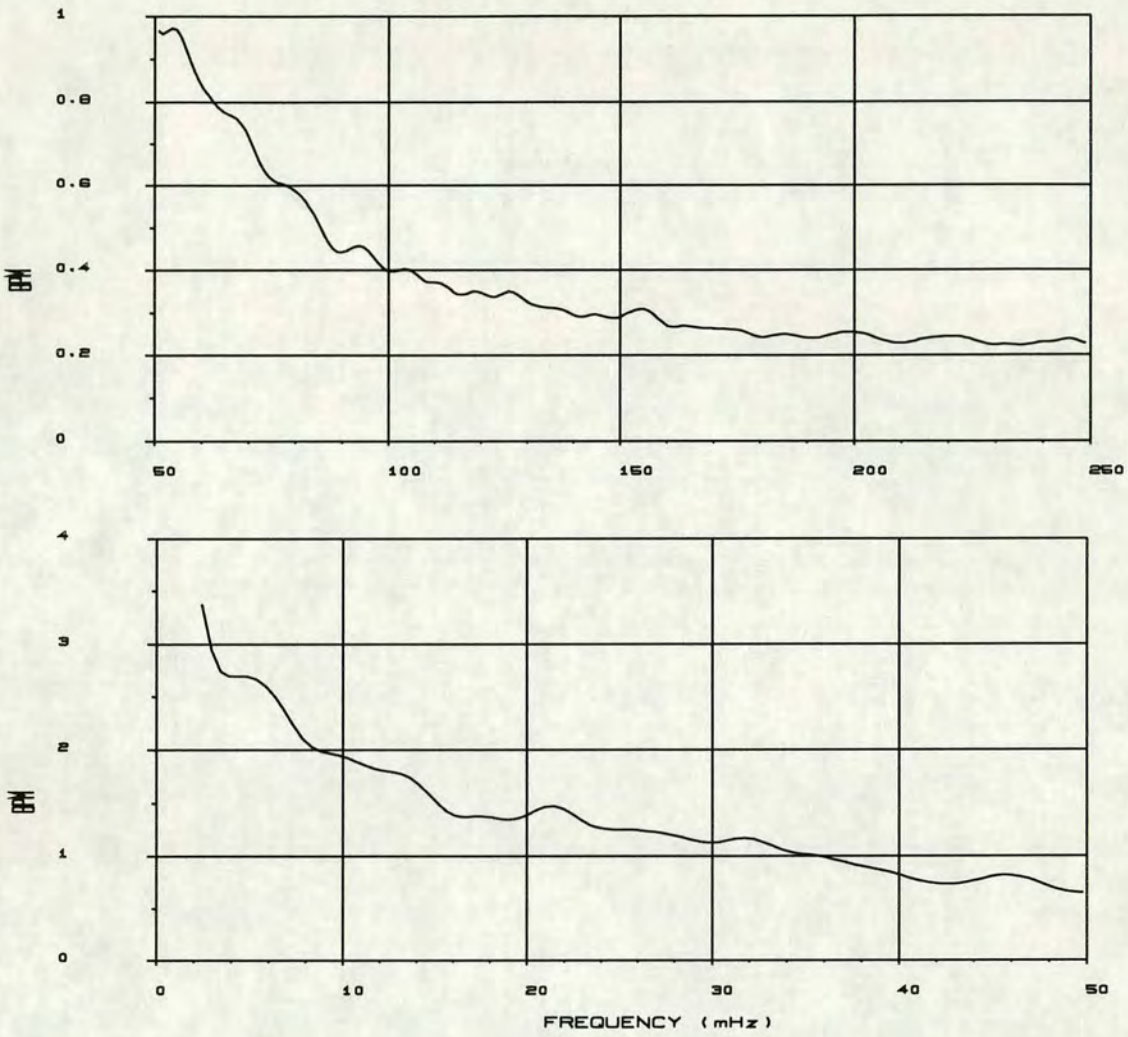

Figure 5.14 Spectral averaging over the full 8 hours, the Hanning window is applied with a 50% block overlap, the graphs are displayed linearly. Top: High frequency spectral components (50 to 250 mHz). Bottom: Lower frequency components (2 to 50 mHz). Note the difference is vertical scale between the two graphs.

The averaged spectrum of figure 5.14 exhibits no persisting frequency components when the FHR time series is observed over a long period. Close inspection of time record sections however, reveals the presence of short duration oscillations in the FHR. Figure 5.15 displays an expanded medium-term FHR trace which shows large amplitude oscillations. In this case, the oscillations start with a periodicity of 1.5

minutes, which then extends to 3 minutes and eventually to 5 minutes (at 1 hour, 1.5 hours and 2 hours into the recordings respectively).

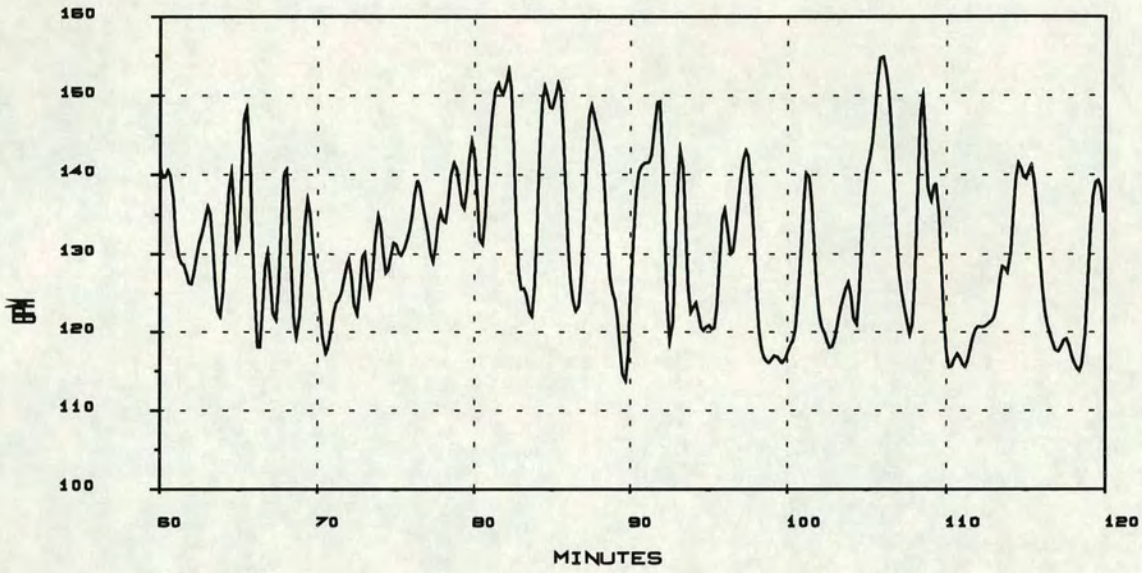

Figure 5.15 Example of observed medium term oscillations. This record corresponds to the epoch 1-2 hours of figure 5.4.

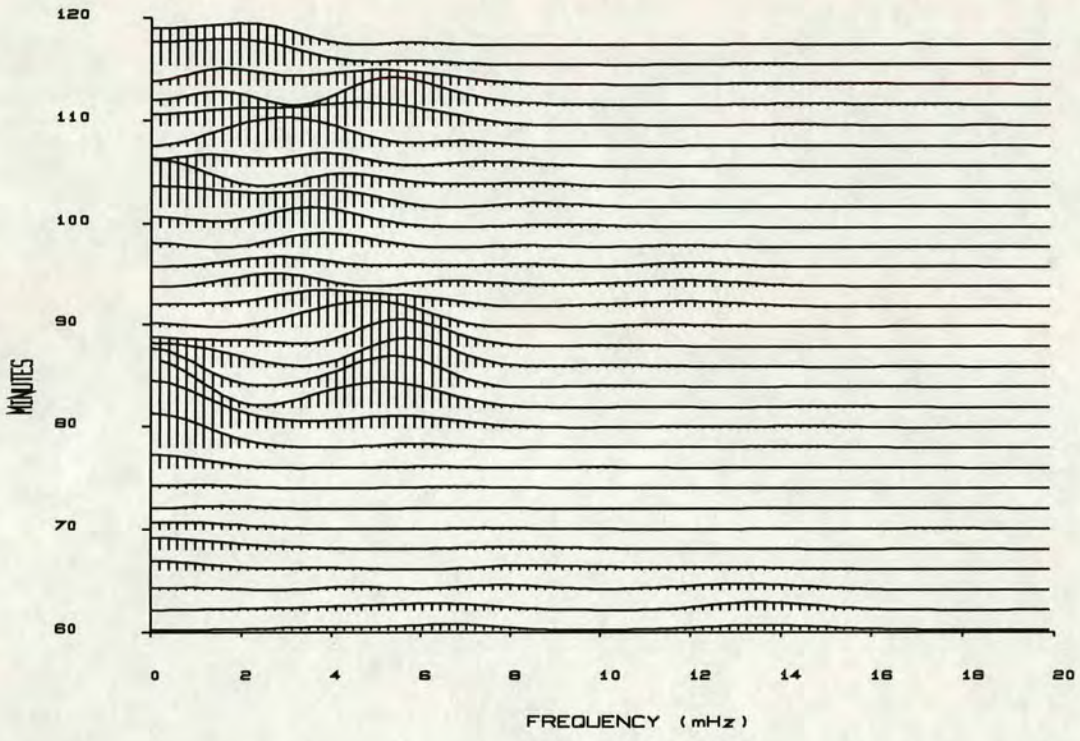

Figure 5.16 Power waterfall spectrum of the FHR time series (average FHR removed) corresponding to the epoch 1-2 hours of figure 5.4. This is obtained from 10 minute blocks with a 75% overlap, a Hanning window is used.

Figure 5.16 reveals the short term spectral presence of FHR rhythms. In this particular epoch, these were more pronounced than observed elsewhere. The above three figures show that while rhythms can be identified, these are relatively short lived and tend to meander through the spectrum. Therefore, when averaged over a long enough period, a smooth spectral content is obtained.

5.5. Results from Other Patients

Results from processed three overnight recordings[78] are included here. The three main FHR parameters are given in the graphs: (a) Long term baseline (T=10 minutes), (b) the variability amplitude index include components from 17 to 250mHz

(1 minute to 4 seconds), and (c) accelerations/decelerations using a band-pass filter from 3.3 to 66 mHz (5 minutes to 15 seconds). For presentation purposes, only one hour of these parameters is included in the graphs. No diagnosis of the fetal health state is attempted here other than noting that the results from the first patient (whose eight hour recording was used above) have reasonable acceleration count, VAI level and baseline range (fig. 5.17). The results for the second patient depicted in figure 5.18 show a low VAI, low accelerations count with corresponding small amplitudes, the baseline lies within an acceptable range. The results depicted in figure 5.19 (third patient) show normal baseline range and VAI index and accelerations count.

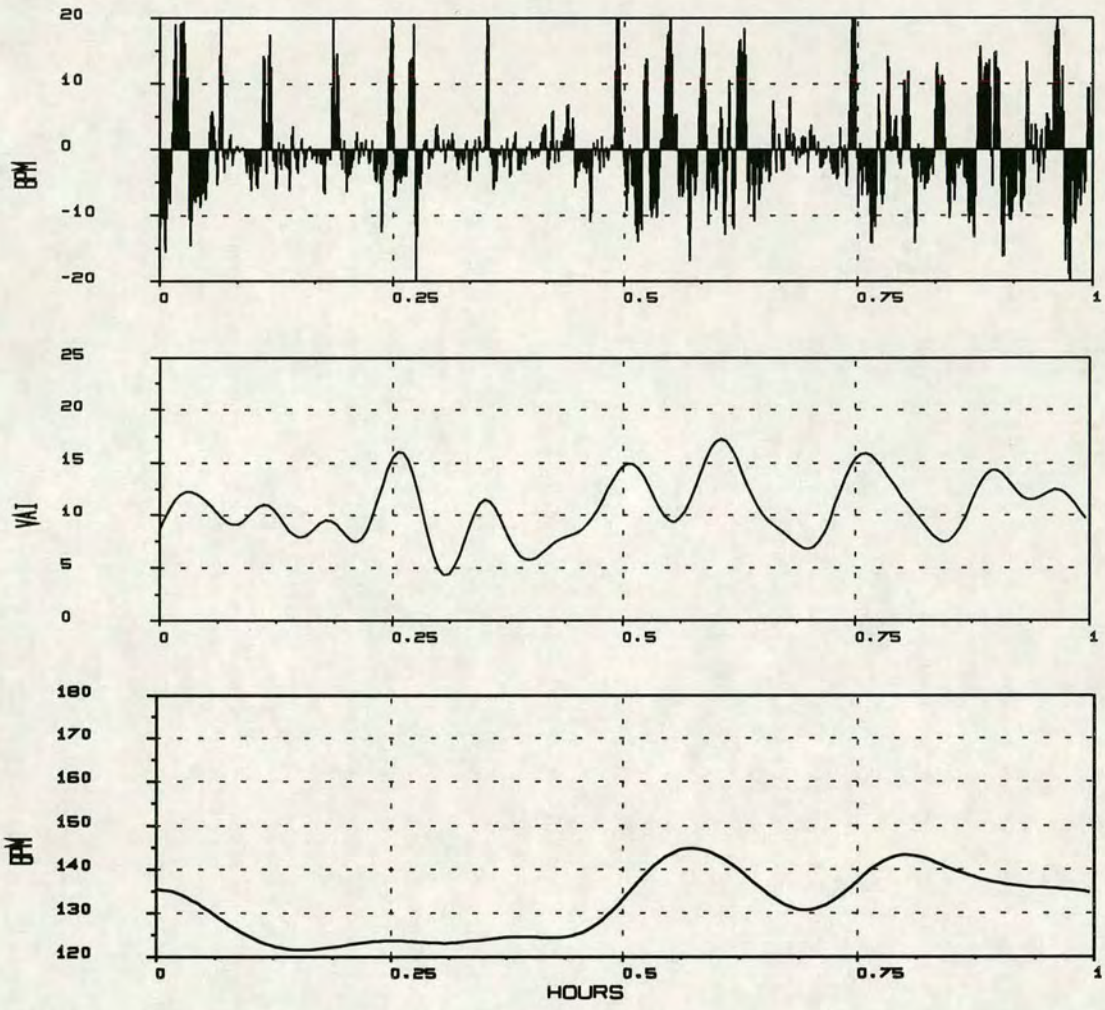

Figure 5.17 Bottom: Long term baseline filtered with a 10 minutes time constant. Middle: 1minute variability which is low pass filtered with a time constant of five minutes.
Top: Accelerations including FHR components in the band 3.3 to 67mHz (5 minutes to 15 seconds).

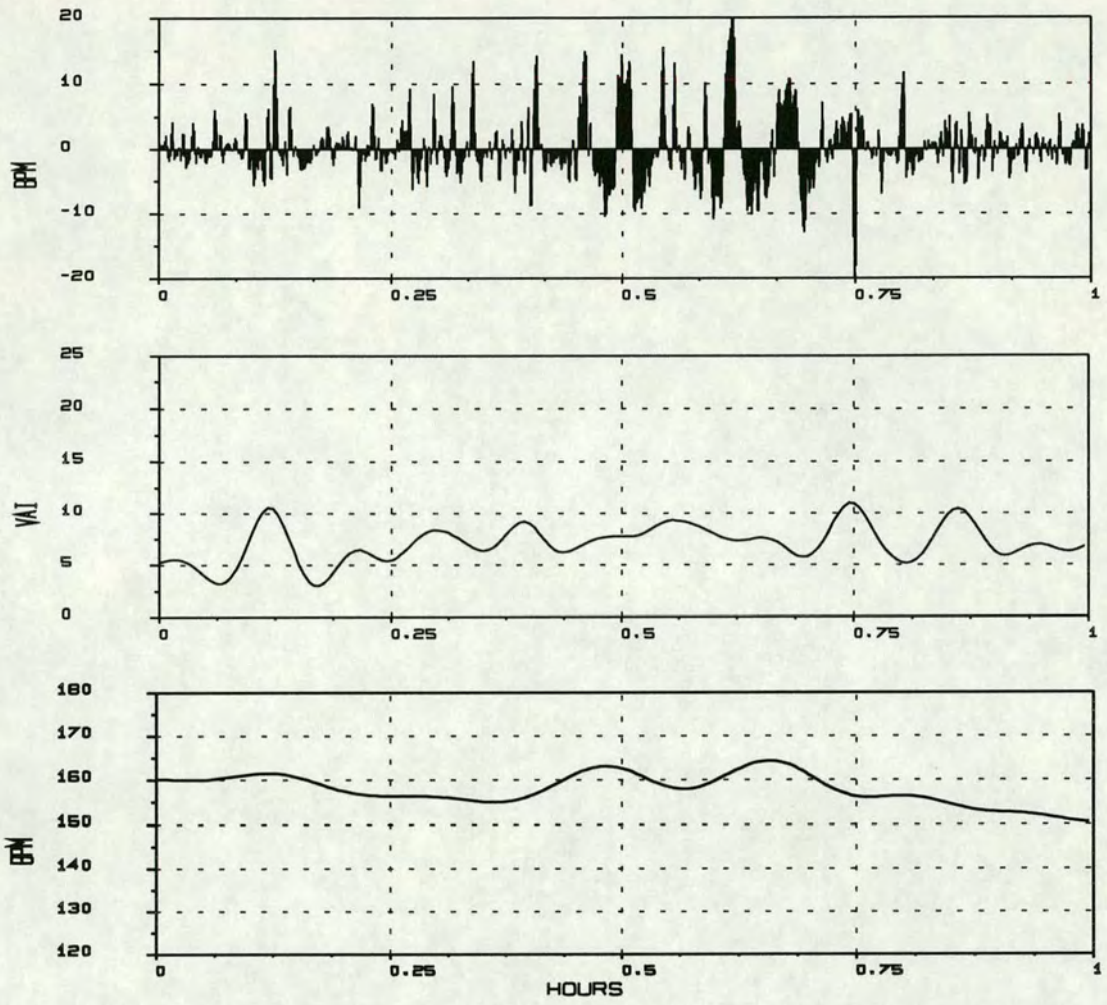

Figure 5.18 Bottom: Long term baseline filtered with a 10 minutes time constant. Middle: 1minute variability which is low pass filtered with a time constant of five minutes.
Top: Accelerations including FHR components in the band 3.3 to 67mHz (5 minutes to 15 seconds).

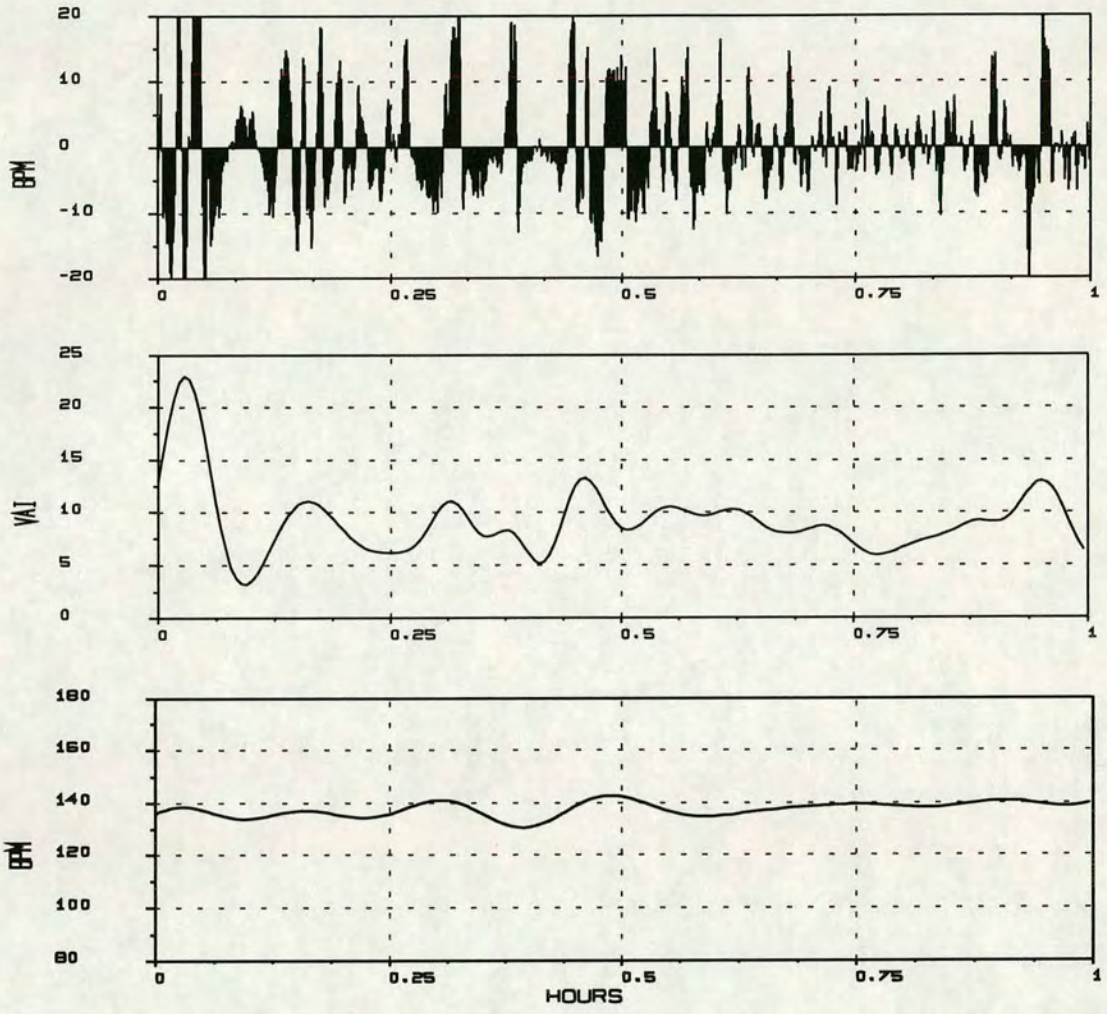

Figure 5.19 Bottom: Long term baseline filtered with a 10 minutes time constant. Middle: 1minute variability which is low pass filtered with a time constant of five minutes.
Top: Accelerations including FHR components in the band 3.3 to 67mHz (5 minutes to 15 seconds).

CHAPTER SIX

6. CONCLUSION

This thesis introduced the relevance of fetal heart rate monitoring and its related parameters which are visually assessed in current routine clinical practice. Pit-falls associated with this field were also highlighted: (a) Some are brought about by technique limitations, which suggests that there is still room for improvement in the signal processing algorithms employed in the derivation of FHR estimates. (b) Others were found in the subjective observer influences where significant inter- and intra-observer differences were identified. This latter case leads to the deduction that, wherever possible, parameter extraction should be automated.

Although most of the heart sound recordings[66, 78] exhibited the main signal component at about 55 Hz, much larger components down to 15 Hz were observed with newer transducers developed in this work. The signals found in this sub-audible region look more promising and merit further attention. The conditioning of fetal heart phonocardiographic signals (fetal heart pressure waves) was detailed: The band-passed HPW signals are used for beat-to-beat estimators while envelope detected versions are applied to block estimators. Two original contributions to fetal phonocardiographic transducers - optical and capacitive - were given and other transducing technologies discussed.

Compared to frequency domain estimators, the 'Comb Filter' algorithm makes no apriori assumption about the distribution of spectral components as such an assumption was shown to be invalid. The comb filter was then applied to ultrasound signals and the results were compared with the gold standard signal (scalp FECG). The main estimate errors were found to be harmonically related. Such errors can be detected with relative ease and the remaining overall error rate was found to be less than 3%.

The 'Iterative Block Clean-up' technique was introduced. This relies on transforming the signal vector from the time domain to the frequency domain while removing all assumed non-coherent signal components. The 'Modified Block Autocorrelation' algorithm was also given, this is similar to the normal block autocorrelation technique except that its 'correlation domain' is derived from the inverse Fourier transform of the linear spectrum. This latter algorithm was found to improve on the outlier error rate while slightly increasing the in-band standard deviation.

A novel time domain beat-to-beat estimator was also introduced: The 'Pseudo Inverse Filter' (PIF) was shown to enhance the ability to detect HPW occurrences while maintaining wide-band noise immunity. Although normally operated sub-optimally, this algorithm was shown to provide almost 1.7dB SNR improvement over the industry-standard 'Pulse Correlation'. When implemented optimally, PIF stretched its performance lead to 3.5dB. The inverse least squares was also investigated but required a modification in its autocorrelation matrix in order to build broad-band robustness into its derived filter weights. However, the performance of the latter was found to be poor when used to filter real fetal heart phono signals.

A simulation was then constructed to assess the performance of fourteen estimators. Firstly, those based on time domain beat-to-beat estimators such as: 'Time Domain Estimator', 'Frequency Matched pre-Filter', 'Envelope Detector-Estimator', 'Pulse

Correlation', 'Inverse Least Squares' and 'Pseudo-Inverse Filter'. Of these, both the 'Pulse Correlation' and the 'Pseudo-Inverse Filter' were found to have similar performances the latter having 3.5dB SNR advantage when used optimally. Secondly, those estimators based on signal blocks of two seconds duration: 'Block Autocorrelation', 'Comb Filter', 'Cepstrum', 'Raw Signal Block Autocorrelation', 'Modulus Difference', 'Fundamental Spectral Estimator', 'Iterative Block Clean-up' and the 'Modified Block Autocorrelation'. Both 'Cepstrum' and 'Block Autocorrelation' estimators were found to have disappointing results, especially during FHR accelerations. The 'Fundamental Spectral Estimator' was found to have a high outlier error rate which was related to the distribution of noise introduced by the envelope detector (employed in the conditioning of fetal HPW signals). Both 'Modulus Difference' and 'Comb Filter' out-performed all other block estimators and were found to have similar performances. The latter was chosen for the reduction in computation cost it affords.

Accurate real time FHR monitoring has been achieved from passive phono signals. The implementation used the 'Comb Filter' algorithm based a VME stand-alone system. Results were then processed off-line to reduce the outliers present in the estimates. A new design based on the stocastic gradient LMS adaptive filter was implemented to detect FHR outliers. The 'Forward Backward Predictor Corrector' (FBPC) structure made use of a cascaded combination of forward and backward LMS predictor configurations with the output of each stage being used to replace estimated errors. Corrected FHR estimates were shown to visually track the underlying FHR trends even at poor signal to noise ratios. In addition, corrected estimates were found not to introduce artificial short term FHR variability.

Further processing was employed to extract relevant parameters such as FHR baseline, accelerations and the 'Variability Amplitude Index'. The latter parameter

was proposed to approximate the amplitude variability as assessed visually by clinicians. Other information such as the FHR amplitude distribution or the correlation between long-term baseline and medium-term variability provide a great wealth of information. Whether these can be used to enhance antenatal fetal monitoring remains to be proven.

The ability to obtain reliable FHR estimates from phonocardiographic signals is ultimately a useful step towards the routine, long-term monitoring of patients.

6.2. Future Work

Cepstral analysis, although unsuccessful in the form investigated here, may produce better results if 'complex cepstrum' processing is implemented. This would involve the 'editing' of the logarithmic complex spectra[4, 61] in order to deconvolve the HPW impulse response.

In this work, the comb filter made use of a peak selection algorithm in order to obtain one FHR estimate from each signal block. Outliers in these results were later 'corrected' using the FBPC. A possible improved implementation would couple the FBPC with the comb filter output stage; in this case, assumed outliers could be replaced by the nearest peak to the predictor output. This approach should result in a closer estimate of the underlying short-term FHR variability. Another possible modification involves using the estimated comb filter SNR to determine the convergence of each FBPC stage.

The comparison of the estimators could benefit from additional testing using real simultaneous recordings of fetal phono and scalp FECG. Increased repeatability could be obtained if the two signals were sampled and stored digitally. In addition, blocks

producing erroneous estimates could be analysed in detail to reveal the failure mechanisms of individual estimators.

References

1. Colley N., Talbert D.G. and Southall D.P.,
"Biophysical Profile in the Fetus from a Phonographic Sensor",
Eur. J. Obstet. Gynecol. Reprod. Biol., 23: pp 261-266, 1986

2. Leech G.J. and Leatham A.,
"Correlation of Heart Sounds and Valve Motions",
edited W. Bleifeld,
In: "Evaluation of Cardiac Function by Echocardiography",
Springer-Verlag, pp 89-96, Berlin, 1980

3. Cibils L.A.,
"Clinical Significance of Fetal Heart Rate Patterns During Labour",
Am. J. Obstet. Gynecol., Vol. 125, pp 290-305, 1976

4. Oppenheim A.V., Schafer R.W. and Stockham T.G.,
"Nonlinear Filtering of Multiplied and Convolved Signals",
IEEE Trans. Audio & Electroacoustics, VOL. AU-16, No. 3, 1968.

5. Chatfield C.,
"The Analysis of Time Series, An Introduction",
Thier edition, Chapman and Hall, 1984.

6. Ross M.J., Shaffer H.L., Cohen A., Freudberg R. and Manley H.J.,
"Average Magnitude Difference Function Pitch Extractor",
IEEE Trans. ASSP, VOL. ASSP-22, No 5, pp 353-362, 1974.

7. Kitney R.I. and Pompelman O.,
"The Beat-by-Beat investigation of Cardiovascular Function. Measument, Analysis and Applications",
Oxford Science Publications, 1987.

8. Strang G.,
"Linear Algebra and its Applications",
Academic Press Inc., 1976.

9. Roberts R. A. and Mullis C.T.,
"Digital Signal Processing",
Addison-Wesley Publishing Company Inc., 1987.

10. Grant P.M., Cowan C.F.N., Mulgrew B. and Dripps J.H.,
"Analogue and Digital Signal Processing and Coding",
Chartwell-Bratt Ltd, 1989.

11. Manning G.K.and Dripps J.H.,
"Comparison of Correlation and Modulus Difference Processing Algorithms for the Determination of Foetal Heart Rate from Ultrasonic Doppler Signals",
Med. & Biol. Eng. & Comp., 24, pp 121-129, 1986.

12. Trevor J. Terrell, 1980.
"Introduction to digital filters".
Macmillan Press Ltd, 1980.

13. Fournie A., Grandjean M., Reme J.M., Pontonnier G.,
"Significance of flat tracings, accelerations and Fisher's score.",
Commission of the Euro. Communities, Perinatal Monitoring, 3rd Progress Report.

14. Bassil H.E, Dripps J.H, Boddy K,
"Aspects of Phonocardiographic Fetal Heart Monitoring"
2nd Portuguese Congress on Biomed. Eng., pp 10.4.1-10.4.5, 1990

15. Parer W.J., Parer J.T., Holbrook R.H., and Block B.S.B.,
"Validity of Mathemetical Methods of Quantitating Fetal Heart Rate Variability",
Am. J. Obstet. Gynecol., Vol. 153, pp 402-409, 1985.

16. Zugaib M., Forsythe A.B., Nuwayhid B., Lieb S.M., Tabsh K., Erkkola R., Ushioda E., Brinkman C.R., Assali N.S.,
"Mechanisms of Beat-to-beat Variability in the Heart Rate of the Neonatal Lamb",
Am. J. Obstet. Gynecol., Vol. 138, pp 444-452, 1980.

17. Visser G.H.A., Goodman J.D.S., Levine D.H. and Dawes G.S.,
"Diurnal and Other Cyclical Variations in Fetal Heart Rate Near Term",
Am. J. Obstet. Gynecol., Vol. 142, pp 535-544, 1982

18. Talbert D.G., Dewhurst J. and Southall D.P.,
"New Transducer for Detecting Fetal Heart Sounds: Use of Compliance Matching for Maximum Sound Transfer",
Lancet, Vol. 1, pp 426-427, 1984

19. Nochimson D.J., Turbeville J.S., Terry J.E., et al.,
"The Nonstress Test",
Obstet. Gynecol., Vol. 51, pp 419-421, 1978

20. Dalton K.J., Dawes G.S. and Patrik J.E.,
"Diurnal, Respiratory, and Other Rhythms of Fetal Heart Rate in Lambs",
Am. J. Obstet. Gynecol., Vol. 127, pp 414-424, 1977

21. Trierweiler M., Freeman R. and James J.,
"Baseline Fetal Heart Rate Characteristics as an Indicator of Fetal Status During the Antepartum Period",
Am. J. Obstet. Gynecol., Vol. 125, pp 618-623, 1976.

22. Organ L.W., Hawrylyshyn P.A., Goodwin J.W., Quilligan J.E.,
"Quantitative Indices of Short- and Long-Trem Heart Rate Variability",
Am. J. Obstet. Gynecol., Vol. 130, pp 20-27, 1978.

23. Solum T.,
"A comparison of the Three Methods for External Fetal Cardiography",
Acta Obstet. Gynecol. Scand., Vol. 59, pp 123-126, 1980.

24. Dalton K.J. and Currie J.R.,
"Fetal Home Telemetry Made Simple",
J. Obstet. Gynaecol., Vol. 6, pp 151-154, 1986.

25. Cahill S.J. and McClure G.,
"A Microcomputer-Based Heart Rate Variability Monitor",
IEEE Trans. Biomed. Eng., Vol. BME-30, No. 2, pp 87-93, 1983

26. Bobitt J.R.,
"Abnormal Antepartum Fetal Heart Rate Tracings, Failure to Intervene, and Fetal Death: Review of Five Cases Reveals Potential Pitfalls of Antepartum Monitoring Programs",
Am. J. Obstet. Gynecol., Vol. 133, pp 415-421, 1979

27. Angel E.S., Fox H.E. and Titlebaum E.L.,
"Digital FIltering and Fetal Heart Rate Variability",
Comp. Biomed. Research, Vol. 12, pp 167-180, 1979

28. Walker D., Grimwade J. and Wood C.,
"The Effects of Pressure on Fetal Heart Rate",
Obstet. Gynecol., Vol. 41, pp 351-354, 1973

29. Yeh S.Y., Forsythe A. and Hon E.H.,
"Quantification of Fetal Heart Beat-to-Beat Differences.",
Obstet. Gynecol., Vol. 41, pp 355, 1973

30. Luisada A.A,
"The Sounds of the Normal Heart",
Warren H. Green Inc, 1972.

31. Wheeler T., Cooke E., Murrills A.,
"Computer Analysis of Fetal Heart Rate Variation During Normal Pregnancy",
Br. J. Obstet. Gynaecol., Vol 86 pp 186-197, 1979.

32. Divon M.Y., Torres F.P., Yeh S.Y. and Paul R.H.,
"Autocorrelation Techniques in Fetal Monitoring",
Am. J. Obstet. Gynecol., Vol. 151, pp 2-6, 1985

33. Katz M. and Gill P.J.,
"Initial Evaluation on an Ambulatory System for Home Monitoring and Transmission of Uterine Activity Data",
Obstet. Gynecol., Vol. 66, pp 273-277, 1985

34. Borgatta L., Shrout P.E. and Divon M.Y.,
"Reliability and Reproducibility of Nonstress Test Readings",
Am. J. Obstet. Gynecol., Vol. 159, pp 554-558, 1988

35. McDonnell J.T.E.,
"Knowledge-Based Interpretation of Foetal Phonocardiographic Signals",
Colloquium on the Application of Artificial Intelligence,
IEE Proc., Vol 137, No 5, pp 311-318, 1990

36. Clayton G.B.,
"Operational Amplifiers",
Second edition, *Butterworth & Co (pulishers) Ltd*, 1983,

37. Nagel J.,
"New Diagnostic and Technical Aspects of Fetal Phonocardiography",
Eur. J. Obstet. Gynecol. Reprod. Biol., 23: 295-303. 1986

38. Manning G.K.,
"Signal Processing for Ultrasonic Foetal Monitoring",
PhD Thesis, Univ. Edinburgh, 1987.

39. Bassil H.E. and Dripps J.H,
"Real Time Processing and Analysis of Fetal Phonocardiographic Signals",
Clin. Phys. Physiol. Meas., Vol. 10, Suppl. B, pp 67-74, 1989.

40. Fukushima T, Flores C.A, Hon E.H, Davidson E.C,
"Limitations of Autocorrelation in Fetal Heart Rate Monitoring",
Am. J. Obstet. Gynecol., Vol. 153, pp 685-692, 1985.

41. Gough N.A., Dawson A.J. and Tomkins T.,
"Antepartum Fetal Heart Rate Recording and Subsequent Fast Transmission by a Distributed Microprocessor-Based Dedicated System."
International J. of Biomed. Comp., 1986

42. Cohen D., Timbs A.E. and Dalton K.J.,
"Measurement of Compliance of the Maternal Abdominal Wall in Pregnancy",
Eur. J. Obstet. Gynecol. Reprod. Biol., 23: 267-272. 1986

43. Flynn A.M., Kelly J., Matthews K.,
"Predictive Value of, and Observer Variability in, Several Ways of Reporting Antepartum Cardiotocograms",
Br. J. Obstet. Gynaecol., Vol. 89, pp 434-440, 1982

44. Welch P.D.,
"The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms",
IEEE Trans. on Audio and Electroacoustics, AU-15, no. 2, 1967.

45. Allen J.B.,
"Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform",
IEEE Trans. ASSP, Vol. 25, NO. 3, 1977.

46. Harris F.J.,
"On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform",
Proc. IEEE, Vol. 66, NO. 1, 1978.

47. Carter M.C., Gunn P. and Beard R.W.,
"Fetal Heart Rate Monitoring Using Abdominal Fetal Electrocardiogram",
Br. J. Obstet. Gynaecol., Vol 87, pp396-401, 1980

48. Chew F.T., Drew J.H., Oats J.N., Riley S.F. and Beischer N.A.,
"Nonstressed Antepartum Cardiography in Patients Undergoing Elective Cesarean Section",
Am. J. Obstet. Gynecol., Vol. 151, pp 318-321, 1985

49. Dalton K.J, Denman D.W, Dawson A.J. and Hoffman H.J,
"Ultradian Rhythms in Human Fetal Heart Rate: a Computarised Time Series Analysis".
Int. J. Bio-Medical Comput., 18, pp 45-60, 1986.

50. Campbell S., Trickey N. and Whittle M.J.,
"Report of the Royal College of Obstetricians and Gynaecologists Working Party on Routine Ultrasound Examination in Pregnancy",
R. Coll. Obstet. Gynecol., App. II, pp 50-67, 1984

51. Goovaert H.G.,
"Instrumentation for the Quatitative Assessment of Fetal Respiratory Sinus Arrhythmia",
PhD Thesis, Free University of Amsterdam, 1989.

52. Manning F.A., Baskerr T.F., Marrison I. and Lang I.,
"Fetal Biophysical Profile Scoring: a Prospective Study in 1184 High Risk Patients",
Am. J. Obstet. Gynecol., Vol. 140, pp 289-294, 1981

53. Dalton K.J., Dripps J.H., Manning K. and Currie J.R.,
"CEUSPEC-a Computerised System for Fetal Home Telemetry",
Int. J. Biomed. Comp., Vol 18, pp 145-153, 1986

54. Dalton K.J. and Dawson A.J.,
"Baseline: A Computer Method of Calculating Baseline in Fetal Heart Rate Recordings",
Int. J. Biomed. Comp., Vol 15, pp 311-317, 1984

55. Huey J.R., Paul R.H., Hadjiev A.A., Jilek J. and Hon H.E.,
"Fetal Heart Rate Variability: an Approach to Automated Assessment",
Am. J. Obstet. Gynecol., Vol 134, pp 691-694, 1979

56. Boddy K. and Dawes G.S.,
"Fetal Breathing",
Br. Med. Bull., Vol. 31, No 1, pp 3-7, 1975.

57. Boddy K. and Mantell C.,
"Observation of Fetal Breathing Movements Transmitted Through Maternal Abdominal Walls",
Lancet II, pp 1219-1220, 1972.

58. Honig M.L. and Messerschmitt D.G.,
"Adaptive Filter Structures, Algorithms and Applications",
Kluwer Academic Publishers, 1984.

59. Cowan C.F.N. and Grant P.M.,
"Adaptive Filters",
Prentice Hall Signal Processing Series, 1985.

60. Widrow B., McCool J.M., Larimore M.G. and Johnson R.,
"Stationary and non Stationary Learning Characteristics of the LMS Adaptive Algorithm",
Proc. IEEE Vol 64, pp 1151-1162, 1976.

61. Tribolet J.M.,
"A New Phase Unwarpping Algorithm",
IEEE, Vol ASSP-25, No 2, pp 170-177, 1977

62. Goodlin R.C., Girard J., and Hollmen A.,
"Systolic Time Intervals in the Fetus and Neonate",
Obstet. Gynecol., Vol 39, No 2, pp 295-302, 1972

63. "Fetal Heart Patterns and their Cilinical Interpretations",
Beard R.W. and Finnegan T.S.,
Sonic Aid Publication, 1974

64. Van Geijn H.P., Jongsma H.W., de Han J. and Eskes T.K.A.B.,
"Analysis of Heart Rate and Beat-to-beat Variability: The interval Difference Index",
Edited by Van Geijn H.P.,
"Studies on Fetal and Neonatal Baseline Heart Rate Variability",
Krips Repro Meppel, ch. 3 pp 99-118, 1980

65. Lawson G.W., Belcher R., Dawes G.S. and Redman C.W.G.,
"A Comparison of Ultrasound (With Autocorrelation) and Direct ECG Fetal Heart Rate Detector Systems",
Am. J. Obstet. Gynecol., Vol. 147, pp 721-722, 1983

66. Talbert D.G., Davies W.L., Johnson F., Abraham N., Colley N. and Southall D.P.,
"Wide Bandwidth Fetal Phonography Using a Sensor Matched to the Compliance of the Mother's Abdominal Wall".
IEEE Trans. Biomed. Eng., Vol. BME-33, No 2, pp 175-180, 1986.

67. Dawes G.S., Houghton C.R.S., Redman C.W.G. and Visser G.H.A.,
"Pattern of the Normal Human Fetal Heart Rate",
Br. J. Obstet. Gynaecol., Vol. 89, pp 276-284, 1982

68. Van Geijn H.P., Jongsma H.W., de Haan J., Eskes T.K.A. and Prechtl H.F.R., "Heart Rate as an Indicator of the Behavioral State" *Am. J. Obstet. Gynecol.*, Vol. 136, No. 8, pp 1061-1066, 1980.

69. Dripps J.H., Manning G.K. and Zhu X. "A Signal-Processing Research Facility and its Application to the processing of Phonocardiographic Signals for the Heart Rate Estimation", *Eur. J. Obstet. Gynecol. Reprod. Biol.*, 23, pp 281-288, 1986.

70. Dawes G.S, Moulden M. and Redman C.W.G., "Limitations of Antenatal Fetal Heart Rate Monitors", *Am. J. Obstet. Gynecol.*, Vol. 162, No. 1, pp 170-173, 1990.

71. Dawes G.S., Visser G.H.A., Goodman J.D.S. and Redman C.W.G., "Numerical Analysis of Human Fetal Heart Rate: the Quality of Ultrasound Records", *Am. J. Obstet. Gynecol.*, Vol. 141, pp 43-52, 1981

72. Lawson G., Dawes G.S. and Redman C.W.G., "A Comparison of Two Fetal Heart Rate Ultrsound Detector Systems", *Am. J. Obstet. Gynecol.*, Vol. 143, pp 840-842, 1982

73. Dawes G.S., Redman C.W.G. and Smith J.H., "Improvements in the Registration and analysis of the Fetal Heart Rate Records at the Bedside", *Br. J. Obstet. Gynaecol.*, Vol. 92, pp 317-325, 1985

74. Detwiler J.S., Jarish W. and Caritis S.N., "Statistical Fluctuations in Heart Rate Variability Indices", *Am. J. Obstet. Gynecol.*, Vol. 136, pp 243-248, 1980

75. Goovaert H.G., Rompelman O. and van Geijn H.P., "A Transducer for the Detection of Fetal Breathing Movements", *IEEE Trans. Biomed. Eng.*, Vol. 36, No 4, pp 471-478, 1989.

76. Spencer J.A.D., Belcher R. and Dawes G.S.,
"The Influence of Signal Loss on the Comparison Between Computer Analysis of the Fetal Heart Rate in Labour Using Doppler Ultrasound (with Autocorrelation) and Simultaneous Scalp Electrocardiogram",
Eur. J. Obstet. Gynaecol. Reprod. Biol., Vol. 25, pp 29-34, 1987

77. Solum T.,
"Antenatal cardiotocography",
Acta Obstet. Gynecol. Scand., Suppl. 96, 1980

78. Southall D.P., Colley N. and Talbert D.G.,
"Overnight Recordings of Fetal Heart Sounds and Fetal Movements",
Eur. J. Obstet. Gynaeco. Reprod. Biol., Vol. 20, pp 85-86, 1985

79. Van Geijn H.P.,
"Fetal Monitoring Present and Future: the Evaluation of Fetal Heart Rate Patterns",
Eur. J. Obstet. Gynecol. Reprod. Biol., Vol. 24, pp 117-119, 1987

80. Mondalou H.D., Freeman R.K. and Braly P.,
"A Simple Method of Fetal and Neonatal Heart Rate Beat-to-Beat Variability Quantification: Perliminary Report",
Am. J. Obstet. Gynecol., Vol. 127, pp 861-868, 1977

81. Lawson G.W., Dawes G.S. and Redman C.W.G.,
"Analysis of Fetal Heart Rate On-line at 32 Weeks Gestation",
Br. J. Obstet. Gynaecol., Vol. 91, pp 542-550, 1984

82. Lotgering F.K., Wallenburg H.C.S. and Schouten H.J.A.,
"Interobserver and Intraobserver Variation in the Assessment of Antepartum Cardiotocograms",
Am. J. Obstet. Gynecol., Vol. 144, pp 701-705, 1982

83. Amato J.C.,
"Fetal Heart Rate Monitoring",
Am. J. Obstet. Gynecol., Vol. 147, pp 967-969, 1983

84. Phelan J.P., Cromartie A.D. and Smith C.V.,
"The Nonstress Test: the False Negative Test",
Am. J. Obstet. Gynecol., Vol. 142, pp 293-296, 1982

85. Tromans P.M., Sheen M.A. and Beazley M.,
"Feto-maternal Surveillance in Labour: a New Approach with an on-line Microcomputer".
Br. J. Obstet. Gynaecol., Vol. 89, pp 1021-1030, 1982

86. Vandewalle J. and Callaerts D.,
"Singular Value Decomposition: A Powerful Concept and Tool in Signal Processing",
Edited by McWhirter J.G.,
in "Mathematics in Signal Processing II",
Clarendon Press, pp 539-559, 1990

87. Vermarien H. and Van Vollenhoven E.,
"The Recording of Heart Vibrations: A Problem of Vibration Measurement on Soft Tissue",
Med. Biol. Eng. Comput., Vol. 22, pp168-178, 1984

88. Spencer J.A.D. and Johnson P.,
"Fetal Heart Rate Variability Changes and Fetal Behavioural Cycles During Labour",
Br. J. Obstet. Gynaecol., Vol. 93, pp 314-321, 1986

Publications

Bassil H.E. and Dripps J.H,

"Real Time Processing and Analysis of Fetal Phonocardiographic Signals",

Clin. Phys. Physiol. Meas. Vol. 10, Suppl. B, pp 67-74, 1989.

Bassil H.E, Dripps J.H, Boddy K,

"Aspects of Phonocardiographic Fetal Heart Monitoring"

2nd Portuguese Congress on Biomed. Eng. pp 10.4.1-10.4.5, 1990

Both papers reprinted here.

Clin. Phys. Physiol. Meas., 1989, Vol. 10, Suppl. B, 67–74. Printed in UK

Real time processing and analysis of fetal phonocardiographic signals

Hany E Bassil and James H Dripps

Department of Electrical Engineering, King's Buildings, Edinburgh University, Mayfield Road, Edinburgh EH9 3JL, UK

Abstract. The monitoring of fetal heart rate (FHR) is commonly used in assessing the general health of the fetus. Although certain periodic cycles may be indicative of fetal problems, only short term observations are routinely employed in clinical practice. This is due to cost considerations, inconvenience to the patient and concern about long term ultrasonic monitoring. Therefore only a low confidence assessment can be established between detected rhythms and the health of the fetus. The technique advocated in this paper makes use of an inexpensive, non-invasive phonocardiographic (phono) transducer which facilitates safe long-term patient monitoring. A variable comb filter applied to the frequency domain is used in order to take full advantage of the harmonic content of fetal heart signals. Real time estimation of FHR has been achieved on pre-recorded phono signals lasting eight hours. Recordings with a reasonable signal quality were analysed and some of the results are given. Advanced signal processing techniques followed by Artificial Intelligence (AI) algorithms reduce the number of erroneous estimates during periods of low signal to noise ratio (SNR). The resulting FHR time series is stored on the host computer for further processing, display and parameter extraction. This paper outlines the processing steps involved.

1. Introduction

Recent developments such as the piezoelectric beam (Talbert *et al* 1984, 1986) and newer PVDF film transducers have enabled the passive phonocardiographic monitoring of FHR. Heart sound signals analysed in this paper were obtained from 8 h overnight recordings using the Talbert and Southall TAPHO transducer. Typically, heart beat signals had an amplitude of 10 mV while other interfering noises were about two orders of magnitude larger.

2. Method

Phono transducer signals are susceptible to corruption by pressure waves from maternal and fetal movements. These movements generate output signals much larger than the fetal heart sound signals and can sometimes cause the transducer to saturate. Patients are therefore monitored while resting. Any residual noise, including maternal and fetal breathing, is reduced by frequency matched filtering of heart sounds. The signals were initially recorded on a standard Racal frequency modulation tape recorder without analogue conditioning. Eight hours of signal can be captured at the slowest tape speed (16/17 IPS). It is very useful during algorithm development to be able to replay the same data. Figure 1 shows the various signal processing stages which are explained in detail.

2.1. Analogue pre-processing

The first stage of the FHR processor consists of analogue conditioning circuitry. Signals are passed through a 45 to 65 Hz bandpass filter which greatly enhances the SNR of the fetal

0951-7979/89/0B0067 + 08 $02.50

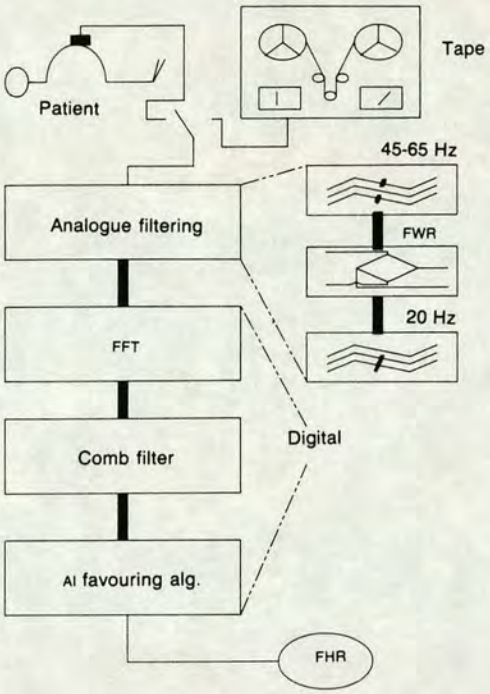

Figure 1. Processing stages of continuous fetal heart rate monitoring. FWR = full wave rectifier; FFT = fast Fourier transform; AI Favouring Alg. = artificial intelligence favouring algorithm; FHR = fetal heart rate.

heart sounds. Even after bandpass filtering, noise interference (such as movements) can reduce the SNR to below 0 dB. In such conditions, simple threshold detection algorithms will fail to identify heart beat sounds from the background clutter. The output of the bandpass filter is passed through a full wave rectifier (FWR) followed by a 20 Hz low-pass filter thus reconstructing the heart beat pulses (figure 2(*a*)). The resulting signal is sampled at 50 samples per second and divided into 2 s blocks (100 samples). A 256 point zero-padded FFT is then applied to the 100 sample record. Figure 2(*b*) shows a typical frequency spectrum; note the dominance of the second harmonic (this was the case in about 90% of the blocks analysed). However, observations of the spectrum have shown that the presence of neither the fundamental nor the second harmonic can be relied upon. Moreover, determining the FHR from a single peak can lead to a biased estimate in the presence of coloured noise. Hence there is a strong incentive to consider a span of harmonics when estimating the position of the fundamental harmonic (and hence the FHR). To this end a comb filter was employed.

2.2. Comb filter

Dripps *et al* (1986) have shown that a large percentage of the power is distributed within the first few harmonics. Hence an algorithm which makes use of this spread in power may prove most effective. The comb filter was chosen for this purpose. The captured signal is considered to consist of two main components:

(1) a set of periodic signals, the strongest of which is the fetal heart component; and

(2) a set of non-periodic signals which define the white noise background to the signal when viewed in the frequency domain.

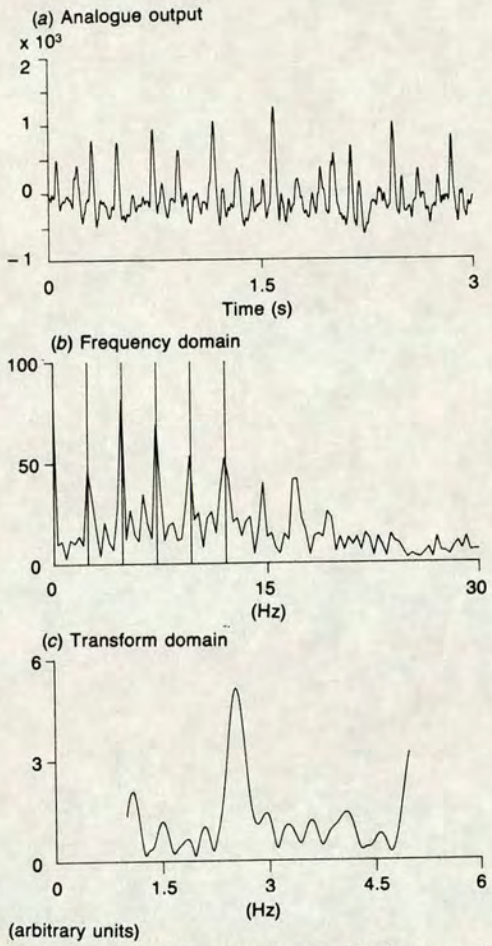

Figure 2. (*a*) Full wave rectifier reconstructed heart beat pulses. (*b*) Frequency domain reveals the harmonic structure of the signal. (*c*) Comb filter output has a global peak at 2.5 Hz (= 150 BPM).

The comb filter can be expressed in the frequency domain as the sum of all the harmonics of a signal with a fundamental frequency W_p.

$$T(W_p) = \int_{-W_c}^{+W_c} I(W, W_p) \star X(W)\, dW \tag{1}$$

where W_c is the cut-off frequency of the signal.

The problem is now reduced to defining the kernel function $I(W, W_p)$ which exhibits two principal properties:

(1) $I(W, W_p)$ has to be defined such that the global output of the function $T(W_p)$ occurs on the fundamental of the periodic signal. In order to avoid making assumptions about the distribution of harmonics in the signal, $I(W, W_p)$ itself will have to be periodic in W_p with maxima lobes at $n \star W_p$ (n being an integer).

(2) When the function is looking for a fundamental where no periodic signal exists, the output should sum to zero. In other words, when evaluated over white noise the output of the function should have an expectation value of zero. This restriction is necessary because of observed low frequency noise in phono signals.

Due to the cyclic nature of $I(W, W_p)$, the output can result in sub-harmonic errors, even in cases of high SNR. The use of a tapering algorithm or restricting the number of harmonics (i.e. the number of teeth in the comb filter) can solve this problem. The latter option is used because it requires no *a priori* assumptions about the harmonic distribution of the signal. On the other hand, severe truncation of the number of harmonics can induce harmonic error. In practice five harmonics are included in the kernel function $I(W, W_p)$.

This algorithm also provides an estimate of the SNR with little computational overhead.

$$\text{SNR}\ (W_p) = \frac{T\ (W_p)}{[P\ -\ T\ (W_p)]} \tag{2}$$

where P is the total power of the signal. This quantity is used as a degree of confidence that can be attached to the output of this transform. The discrete time version of the algorithm is given as:

$$T\ (W_p) = \sum_{n=1}^{H} \sum_{k=-a}^{+a} X\ (W_p + k{\star}\mathrm{d}W) \star I\ (k{\star}\mathrm{d}W) \tag{3}$$

where H is the maximum number of harmonics considered. $2{\star}a$ is the width of the kernel function, $\mathrm{d}W$ is the resolution of the frequency spectrum. Here the kernel function $I\ (k{\star}\mathrm{d}W)$ is assumed to be the same for all the harmonics.

The global peak of the comb filter domain is expected to pinpoint the FHR (figure 2(*c*)). In most cases this was true but sudden bursts of noise can produce multiple local peaks, and care must be taken when selecting the most likely FHR peak.

The above phonocardiographic processing was tested against a high quality scalp FECG signal where the global peak in the comb filter domain was assumed to indicate the true FHR. A resulting error rate of less than 3% was observed. Moreover, in the cases of erroneous estimates, the second highest peak in the comb filter domain usually indicated the true FHR.

2.3. Peak selection

Ideally, the output from the comb filter will be a trace with a single global peak denoting the position of the FHR. In practice, multiple peaks are observed and the program has to select the peak which most likely reflects the FHR during that 2 s block. Any peak selection algorithm should take account of statistical parameters such as current FHR trend and distribution. An artificial intelligence algorithm was developed for this purpose. This algorithm will favour the peaks according to the three criteria: (1) the signal power of the peak, (2) the closeness of the peak to the previous FHR estimate, and (3) the position of the peak relative to the recent statistical distribution of estimated FHR (this can be updated during processing). For example, if the algorithm is given the choice between multiple peaks of similar power, it will select the one closest to the previous estimates. Outliers caused by highly coloured noise are thus rejected in favour of estimates which preserve the continuity of the FHR time series. Two cost functions are needed for this AI selection algorithm, the two governing equations are given:

$$V\ (F_n,t) = P_n \star C_d(\text{FHR}(t-1)-F_n) \star C_p(A-F_n) \tag{4}$$

$$\text{FHR}(t) = F_x;\ \text{such that}\ V(F_x,t) = \text{MAX}\ (V(F_n,t)) \tag{5}$$

where $V()$ is the favouring the peak F_n will get, P_n is the power at rate F_n, $C_d()$ is the cost function relating to previous FHR estimate namely ($\text{FHR}(t-1)$), C_p is the cost function relating to the expected average FHR A. Using these techniques, the error rate is reduced to less than 1% when the signal quality is good.

2.4. Interpolation

Reducing the quantisation error in FHR estimates to acceptable limits (i.e. 1/3 BPM) suggests oversampling in the case of time domain correlation processing or performing large zero-padded FFTs in the case of frequency domain processing. This constitutes a considerable increase in the computational load and therefore may only be applicable in off-line processing. An efficient solution adopted in this work uses interpolation at the comb filter output stage. This interpolation technique gives an exact measure of the comb filter magnitude (not an approximation) and therefore does not compromise the location of the peak position.

Initially, an approximate location of the peak indicating the FHR estimate is determined to the nearest ±10 beats per minute (BPM). Then a binary search algorithm using the interpolation technique is used to reduce the quantisation error to the required size. It was felt that ±0.2 BPM was a small enough quantisation for purposes of this research. A 50-fold reduction in computational load resulted from implementing this binary search/ interpolation technique.

Programs were written in 'C' language and tested on the UNIX operating system where good display and debugging facilities were available. The resulting program modules are then linked and the machine code was downloaded to a VME based signal processing facility. Real time FHR extraction was performed by a MC68000 processor running at 8 MHz. A menu driven style of programming was adopted for development, algorithm parameters could easily be modified and optimised iteratively.

The resulting FHR time series is then sent to a host mainframe for further processing. Residual estimate errors (outliers) should be removed before any analysis can be carried out. Assuming that outliers in the FHR time series have no cyclic patterns (i.e. are randomly distributed), then a least mean square (LMS) prediction-error filter may be used to track the underlying signal. When the LMS prediction and a new FHR sample differ by more than a predetermined limit, the sample is assumed erroneous. In such a case the LMS output is used to replace the suspected outlier. Useful information such as dominant frequencies, auto-correlation and cross-correlations etc can then be extracted.

3. Results

The research is currently concentrating on long term and short term FHR time series analysis. The FHR spectrum (with outliers removed) is inspected and dominant frequencies are extracted using digital filtering techniques (a typical long term spectrum is displayed in figure 3). Three main bands of frequencies were identified in the regions of 3, 27 and

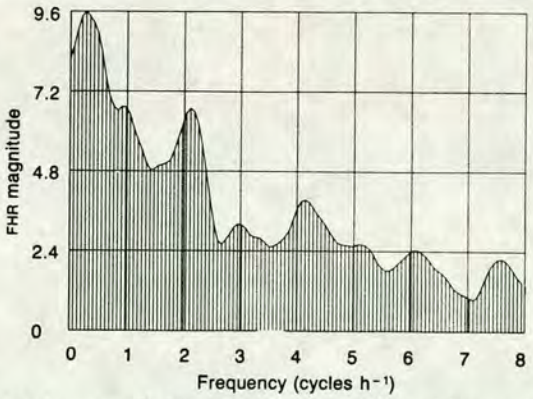

Figure 3. Typical long-term fetal heart rate spectrum obtained from an 8 h recording.

70 min, these are separated in figure 4 (*a, b* and *c*). It was also observed that in most cases a strong correlation existed between the average rate and the variance of the short term FHR

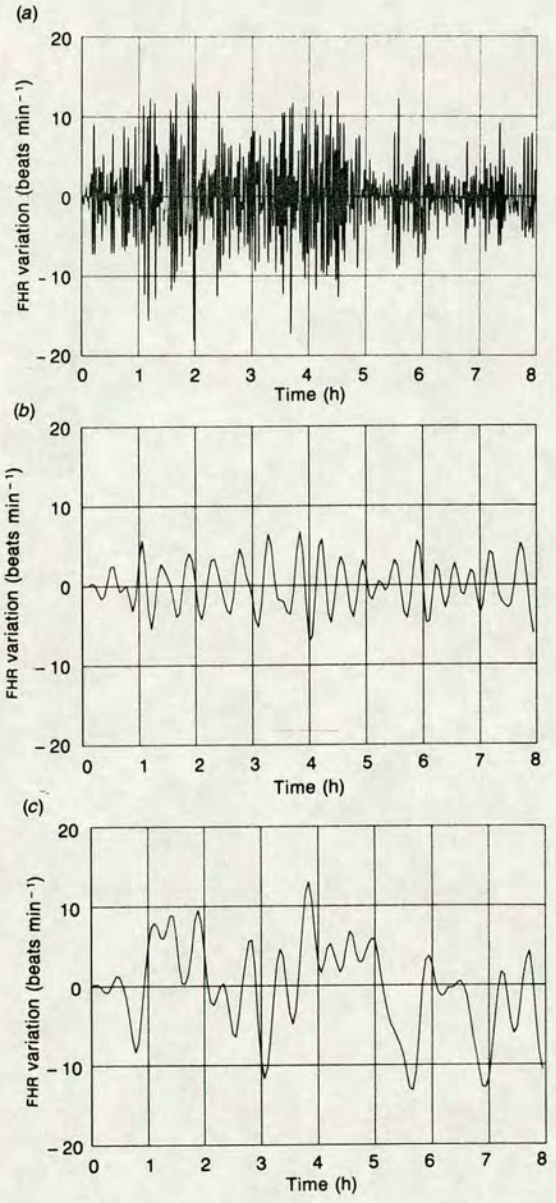

Figure 4. (*a*) 3 minutes variations extracted from the fetal heart rate (FHR) time series. Bandpass filtered from 2 to 4 min (Q factor = 1.5). (*b*) 27 minutes variations extracted from the FHR time series. Bandpass filtered from 20 to 35 min (Q factor = 1.77). (*c*) Long variations extracted from the FHR time series. Bandpass filtered from DC to 35 minutes.

(described in alternative format by van Geijn *et al* (1980)). Figure 5 is a typical scatter plot of FHR variance against FHR average, various block lengths were tried including 32 and 256 s (figure 5 (*a*) and 5 (*b*) respectively). There is an interesting difference in outlines of the allowed regions of the two plots. The existence of points with high average and low variance in the short block calculations (32 s) led us to conclude that besides the correlation of

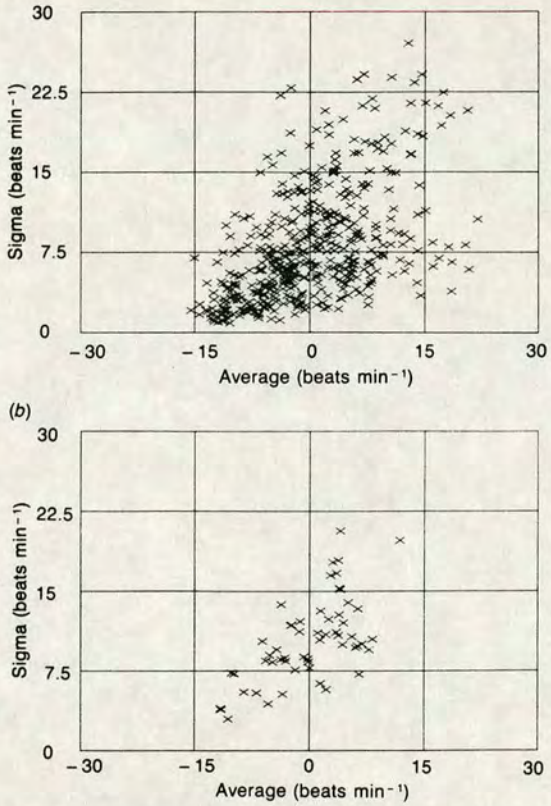

Figure 5. (*a*) Scatter plot of sigma against average, 32 s block length was used. Note the presence of samples in the high average−low sigma region. (*b*) Scatter plot of sigma against average, 256 s block length was used. Note the absence of samples in the high average−low sigma region.

average and variance, there also exists a lag in the coupling between mean rate and variance. This lag only appears to exist during positive trends (accelerations) in FHR data and not during negative trends (decelerations). In other words, as the FHR increases the variance of the signal does so too but only after a delay. In addition as the FHR decreases the variance also decreases without any delay, this is deduced from the absence of high variance and low average in all block sizes. The delay associated with the positive FHR trend is roughly estimated at one minute. A model (figure 6) is proposed for the physiological control

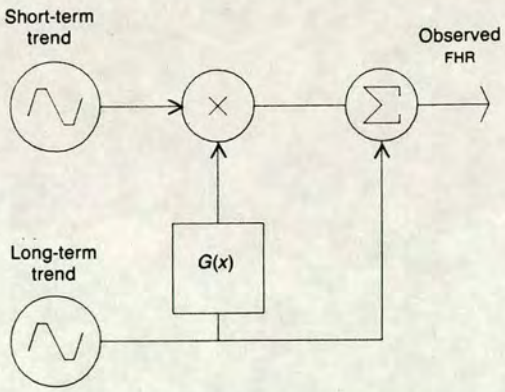

Figure 6. Proposed physiological model for the control of fetal heart rate variance. $G(x)$ is a non-linear delay filter.

underlying FHR variations where the non-linear element $G(x)$ incorporates the variance delay for positive FHR trends.

4. Discussion

Accurate real time FHR tracking has been achieved from passive phono signals. Although most of heart sound recordings exhibited a signal at about 55 Hz, centre frequencies as low as 35 Hz were observed. This shift in the resonant frequency is probably related to the patient's obesity: mechanically, larger bodies have lower frequency of resonance. Initially selection of the bandpass filter for each patient may prove necessary to solve this problem, alternatively one could make use of adaptive tracking filters. The ability to derive accurate FHR estimates in real time from phonocardiographic signals is ultimately a useful step towards routine, long-term patient monitoring. In the short term it is providing a good basis for the study of FHR rhythms which in time could lead to a better understanding of fetal cardiovascular dynamics and the discovery of some early indicators of fetal well-being.

Acknowledgements

We wish to thank the Scottish Home and Health Department for the financial support of Medical Signal Processing Facility (Jan 86 - Dec 87) contract: R/NMD/2/2/C349 and Biolife Ltd for support under contract: AR200/904295. The 8 hour tape recordings of fetal heart sounds were made by N Colley (Cardiothoracic Institute, Royal Brompton Hospital, London).

References

Dripps J H, Manning G K and Zhu X 1986 A signal-processing research facility and its application to the processing of fetal phonocardiographic signals for heart rate estimation *Eur. J. Obstet. Gynecol. Reprod. Biol.* **23** 281-8

van Geijn H P, Jongsma H W, de Haan J and Eskes T K A B 1980 Analysis of heart rate and beat-to-beat variability: Interval difference index *Am. J. Obstet. Gynecol.* **138** 246-51

Talbert D G, Davies W L, Johnson F, Abraham N, Colley N and Southall D P 1986 Wide bandwidth fetal phonocardiography using a sensor matched to the compliance of the mother's abdominal wall *IEEE Trans. Biomed. Eng.* **BME-33** 175-80

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* **1** 426-7

ASPECTS of PHONOCARDIOGRAPHIC FETAL HEART MONITORING

Hany E. Bassil, James H. Dripps and Kenneth Boddy*

Department of Electrical Engineering
University of Edinburgh.

*Department on Obstetrics and Gynaecology
University of Edinburgh

Abstract

This paper considers the conditioning of signals derived from maternal abdominal surface movements (MAM) due to internal pressure waves. The source of interest here is the fetal heart. Correlation is performed as a reference technique for subjective comparison. A new processing method utilising frequency domain equalisation of the MAM signal is advocated in this paper. This method provides filter weights which may be used to equalise, in the time domain, the transient oscillatory waveforms corresponding to the first and second heart sounds. The resulting waveform may be used to estimate accurately heart rate, variability and systolic intervals on a beat-to-beat basis. MAM sample blocks of five seconds duration are processed on a Personal Computer (PC). The resulting timings are compared with the corresponding manually extracted results.

Introduction

Historically, clinicians have listened to the fetal heart (auscultation) to determine fetal health. Ultrasound based electronic FHR monitoring is currently widely used for fetal health assessment. The ultrasound machines provide a continuous FHR trace using correlation processing techniques.

Although a clearer biophysical profile can be obtained from long-term monitoring, patient acceptability, cost considerations and concern about possible overexposure to ultrasound limit monitoring sessions to about 30 minutes. Phonocardiography is the modern electronic equivalent of the original passive listening technique. Being strictly non-invasive, phono-cardiography is inherently more applicable to long term FHR monitoring, although major difficulties remain with obtaining usable signals on a routine clinical basis. However, the estimated lower cost of phonocardiographic monitoring equipment will facilitate self monitoring in the patient's home.

Significance of FHR traces

When interpreting FHR plots, the main features of interest to experts are short term variations (beat to beat variability), medium term variations (of the order of minutes) and the underlying (baseline) FHR trends.
In the case of short term variations, it is generally believed that rate variability in excess of 10 BPM is a positive indicator of fetal wellbeing. With medium term variations, the clinician expects to see an average of at least 4 accelerations per hour from a healthy Fetus. Regarding long term variations, the baseline FHR should remain within upper and lower limits determined by the gestational age of the fetus. These indicators, when combined with other patient information, help physicians to assess Fetal health.

The MAM Fetal Heart Signals

The auditory system of humans, although used successfully in auscultation based diagnosis, is not ideally suited for detecting pressure waves due to heart activity [1].

The inability of humans to detect frequencies below 30 Hz has resulted in the processing of these signals being concentrated on frequencies above this threshold of human auditory perception. The use of the resulting sounds for diagnostic purposes has traditionally constituted the major part of Phonocardiography.

Figure 1 shows the short term spectrum of a typical fetal heart MAM trace in the form of a waterfall diagram. Considerable amounts of energy is evident at frequencies down to 10 Hz. Moreover, the total power in this lower

region of the frequency spectrum is about three orders of magnitude larger than the power associated with the classical heart sounds in the band 45 - 200 Hz, the traditional audible band identified with auscultation.

Generation of Heart Sounds.

Classically the heart beat is identified from two audible sounds in the heart cycle. The timing between successive heart sounds is used to extract:

a. the duration of the interval between the first and second heart sounds (T1),

b. the duration of the interval between the second sound of a given beat and the first heart sound of the next beat (T2),

c. The heart beat period (T) where T = T1 + T2

This paper is mainly concerned with the conditioning of the first heart sound with a view to extracting the overall timing T. As shown earlier the part of the spectrum used for auscultation contains only a small portion of the total available power in the wider region from 5 Hz to 200 Hz. For this reason, we will adopt the term fetal Heart Pressure Wave (HPW) to denote both the audible and sub-audible parts of the signal.

The two main fetal HPW's detected at the maternal abdominal wall are referred to as the first and second HPW's.

The first HPW is generated at the start of ventricular systole (contraction). Inward acceleration of the heart wall and the closure of the mitral and tricuspid valves causes the structure (heart and its surroundings) to vibrate. Although both the right and left ventricles contract simultaneously, the left ventricle is larger than the right and is therefore the main contributor to this first HPW.

The second HPW emanates at the start of ventricular diastole (relaxation). It is mainly caused by the semi-lunar valves stopping reverse flow of the arterial blood. This second wave is usually weaker and generally contains higher frequencies than the first HPW.

Two additional pressure waves (the third and fourth sounds) are well documented but are more difficult to isolate from background noise, particularly in the fetus and currently there seems to be no prospect of basing routine clinical assessment on their detection.

FHR estimation

There are two main methods of FHR estimation: block processing and beat to beat processing. The two methods are explained below:

(a) Block estimation

This approach depends on extracting the average FHR from a short block of fetal HPW samples. Block sizes typically used vary from 1 to 4 seconds. These methods are generally robust to evolutionary changes in the HPW shape and can operate at relatively low Signal-to-Noise Ratios (SNR's) in comparison with the beat-to-beat methods.

Described elsewhere [2], block processing of fetal phonocardiographic signals is typically preceded with a simple analogue band-pass filter with -3dB points at 45 and 65 Hz in order to improve SNR. Even after this enhancement, SNR's close to 0 dB are often encountered due to interference (e.g. maternal and fetal movements). Under such conditions, simple time domain methods such as thresholding will fail to reliably identify the first and second heart sounds in the high background of noise. The filtered HPW's are then passed through a full wave rectifier which enables the reconstruction of the envelope of the HPW signal by the succeeding low-pass filter. The low-pass filter also acts as an anti-aliasing filter for the signal which is then sampled, digitised and passed to the digital signal processor.

Some techniques which have been applied to estimate FHR are: Comb filtering [2] Autoregressive Spectral Analysis [4] and Correlation [3]. Due to the coherent averaging achieved by block processing, noise immunity is gained but beat-to-beat variations in FHR are lost.

(b) Beat to beat estimation

Very little work on the detection of individual fetal HPW's is reported in the literature although some success has been reported [7] with an Artificial Intelligence approach employed in the time domain. The problem with time domain approaches is that the phase linearity of the signal conditioning filters is crucial if pulse timing accuracy is to be preserved. Observations of the fetal HPW signals indicates that the pulse shapes are not time invariant. Therefore any algorithm designed to extract beat to beat timings must ensure a linear phase relation in the conditioning filter.

Correlation processing of the HPW's induces spreading of the pulse as shown in figure 3 and does not guarantee a detectable global peak if the HPW shape is to any degree non-stationary. It has also been noted [5] and [6] that additive noise can cause errors in the FHR estimation process which exceed the true physiological FHR variability.

Determination of the timings (T, T1 and T2) from the raw signal presents another set of problems: In addition to the usual problem of false triggering on noise spikes, there is also the design problem of selection of the correct signal feature on which to trigger such as first peak, centre of gravity or maximum acceleration within the HPW [7].

Proposed frequency domain inverse filtering.

The method advocated in this paper attempts the equalisation, in the frequency domain, of both the phase and amplitude of the first HPW. The assumption is being made here that the observed oscillatory waveform identified as the first HPW originated as a single unipolar pulse at an instant in time which is being taken as the start of the systolic interval (T1). In travelling to the maternal abdominal surface and through the transducer, this pulse suffers dispersion. In other words, the transfer function of the abdominal medium will have a non-flat amplitude and a non-linear phase response. This is reflected in the amplitude and phase components of the Fourier Transform of the HPW.

The objective of the equalisation process is not only to reconstruct the original pulse, but also to have the peak of the reconstructed pulse pinpoint the onset of the systolic interval.

The procedure is described below. A time averaged first HPW is obtained from a portion of the recording where the first HPW is clearly visible and discernible from the second HPW. This averaged signal is then transformed into the complex frequency domain. In an ideal case where noise is absent from both the above Time Averaged Series (TAS) and the individual HPW's to be equalised, the spectrum would be inverted using the following equation:

$$H(w) = \frac{1}{X(w)^*}$$

where H(w) is the required inverse function and X(w) is the spectrum of the TAS.

The problem with the true inverse filter above is that equalisation based on multiplication of S(w) + N(w) by H(w) would lead to the unacceptable emphasis of background noise, where S(w) and N(w) are the signal and noise spectra respectively. This is the classic problem with inverse filtering of low signal-to-noise ratio parts of the signal spectrum. A compromise is required between true inverse filtering and emphasis of background additive noise. This compromise is achieved by thresholding the spectrum of the TAS: the amplitude of any frequency that falls below a certain predetermined threshold is not inverted, in fact it is reduced even further. In order to avoid sudden discontinuities in this pseudo inverse, a soft threshold is imposed on the transform. A suitable function for the purpose is the sigmoid function. The resulting "inverse" becomes:

$$H(w) = \frac{A(X(w)^*) * Sig(G*(|X(w)|-T))}{(|X(w)| + d)}$$

A(): is the argument extraction function applied to the conjugate of X(w);

Sig(): is the sigmoid function,

$$Sig(x) = \frac{1.0}{1.0 + e^{-x}}$$

T: Threshold set for determining the level of noise rejected.

G: A gain factor determining the slope (softness) of the sigmoid function;

d: A small constant used to stabilise the denominator typically d = 0.05;

Figure 2 shows the modulus of H(w) (|H(w)|). The transfer function H(w) is inverse transformed and the weights of the equalising time domain filter are obtained (equivalent to the impulse response time series). When these weights are convolved with the captured signal, the first HPW's should be transformed into a pulse train where each pulse denotes the onset of another systolic interval.

Data Acquisition

The signal is obtained from a capacitive displacement transducer with a conversion ratio of 10 Volts/mm. Patient isolation is achieved using an optical isolator following amplification of the signal (to avoid isolator noise). The isolator output is passed through a fourth order band-pass filter having a passband from 20Hz to 150Hz. The signal is then sampled at 500Hz using a 12 bit ADC. The ADC is controlled from an 68000 microprocessor in an industry standard STE bus system. The sampled signal is then transferred to an IBM PC compatible Toshiba T1600 (286 laptop) with a math co-processor for signal processing, display and plotting.

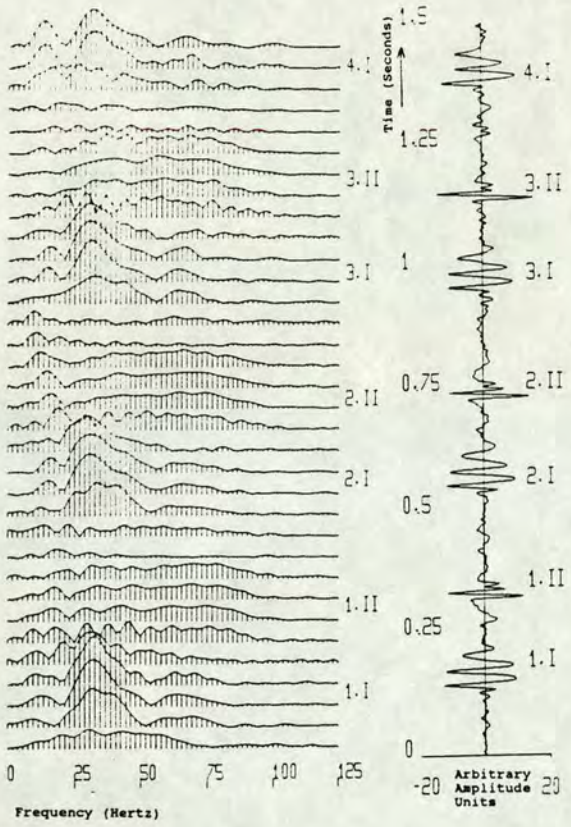

Figure 1: Spectral waterfall diagram of the fetal heart pressure wave detected at the abdominal wall. Three complete heart cycles are apparent from the trace. The linear spectrum is obtained from a sliding 160ms block with an overlap of 120ms. For clear identification of the first and second fetal pressure waves, they have been marked for each of the heart cycles (1.I,1.II for the first and second HPW's of the first heart cycle etc ...)

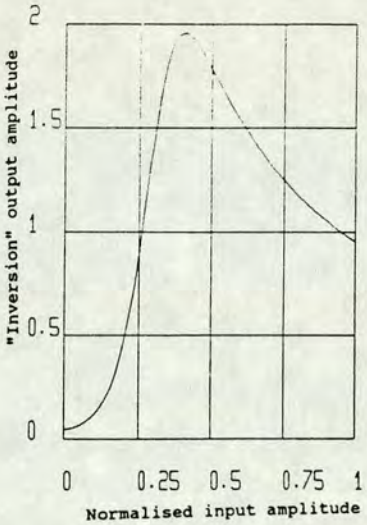

Figure 2: The transfer function of the advocated "inversion function". The gain factor (G) is 20 and the threshold (T) is 0.25 . Note how a low input amplitude (x axis) is suppressed even further while input amplitudes greater than the threshold are "inverted".

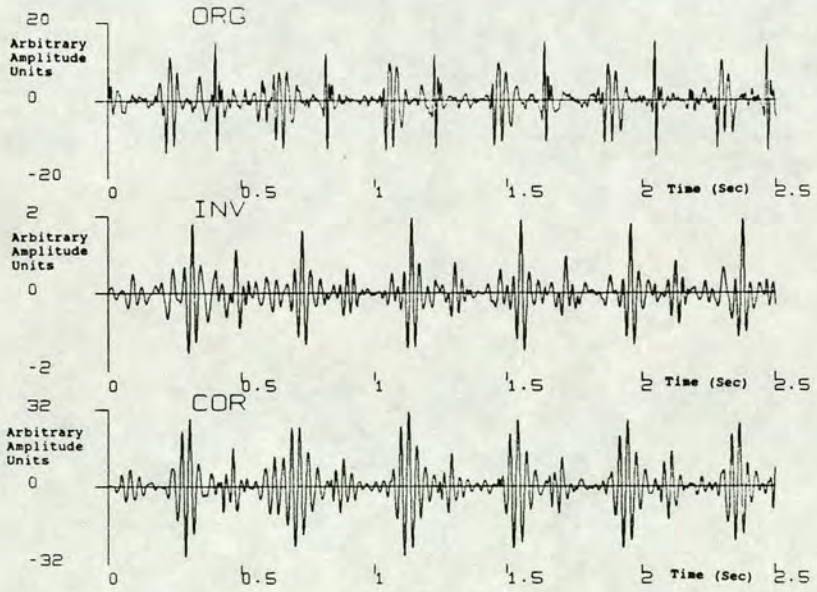

Figure 3: The original fetal HPW (top trace). The result of convolving the original signal with the time domain taps of the "inverse filter" (middle trace). The result of correlating the time-averaged first HPW with the original signal (bottom trace).

Results

The original signal depicted in figure 3 (top trace) shows 6 complete of fetal heart beats. Points worthy of mention are:
(a) The first HPW's are longer in duration than the second HPW's, consisting typically of three oscillations.
(b) In this recording the second HPW's are unusually clear and well defined but in general they are smaller than the first HPW's and in some cases they may be indistinguishable from background noise (hence the concentration on the first HPW's.
(c) The difficulty of extracting a unique timing feature in this trace is obvious.
The averaged time series used for the inverse filter and the correlation in figure 3 (middle and bottom traces respectively) was obtained from the summation of ten successive first fetal HPW's from this record using visual alignment.
Figure 3 (middle trace) shows the result of inverse filtering the original HPW using the same inverse parameters for gain and threshold as in figure 2. The processed first HPW's now provide a series of uniquely identifiable global peaks.

The same time averaged first HPW was correlated with the original signal to provide a comparative processed signal. The usual problem of waveform spreading in time is apparent and this technique fails to provide an identifiable stable global peak.

The technique advocated in this paper appears to provide a unique indicator of the fetal first HPW for timing purposes. Work remains to be done in ascertaining the variance of beat-to-beat estimates based on this method in comparison with existing methods.

Acknowledgements

This work was sponsored by the Advanced Informatics in Medicine (AIM) under project A1032 (telemedicine). We would like to thank the AIM Telemedicine (Edinburgh) team for consultations.

References

[1] Luisada A A,
"The Sounds of the Normal Heart",
Warren H. Green INC, 1972.

[2] Bassil H E, Dripps J H,
"Real Time Processing and Analysis of Fetal Phonocardiographic Signals".
Clin. Phys. Physiol. Meas., Vol 10, Suppl. B, pp 67-74, 1989.

[3] Manning G K, Dripps J H,
"Comparison of correlation and modulus difference Processing Algorithms for the determination of Foetal heart rate from ultrasonic Doppler Signals",
Med. & Biol. Eng. & Comp., 24, pp 121-129, 1986.

[4] Manning G K, Dripps J H,
"Autoregressive Spectral Analysis Techniques for the Determination of Fetal Heart Rate",
Biocomp '83, BES/MRC Meeting, Cambridge, 1983.

[5] Wheeler T, Cooke E, Murrills A,
"Computer Analysis of Fetal Heart Rate Variation During Normal Pregnancy",
Br. J. Obstet. Gynaecol., Vol 86 pp 186-197, 1979.

[6] Murrills A, Wilmshurst T, Wheeler T,
"Ante-natal measurement of beat-by-beat fetal heart rate variation",
2nd International Conference: Fetal and Neonatal Physiological Measurements.
Session I. 1984.

[7] McDonnell E, Dripps J, Grant P,
"The Knowledge-based detection, regeneration and classification of Fetal Heart Sounds",

Colloquium on the Application of Artificial Intelligence Techniques to Signal Processing, IEE Digest Number 1989/42, 1989.