MODERN DIGITAL SIGNAL PROCESSING TECHNIQUES APPLIED TO DOPPLER ULTRASOUND

by

Paul Ivan John Keeton

A thesis submitted to the University of Leicester
for the degree of Doctor of Philosophy

BTSP Research Group,
Department of Engineering,
University of Leicester

April 1997

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

MODERN DIGITAL SIGNAL PROCESSING TECHNIQUES APPLIED TO DOPPLER ULTRASOUND

by

Paul Ivan John Keeton

Declaration of Originality:

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in the Department of Engineering, University of Leicester, UK. All work presented in this thesis is original unless otherwise acknowledged in the text or by references. No part of it has been submitted for any other degree, either to the University of Leicester or to any other university.

Paul I. J. Keeton

April 1997

Acknowledgements

During the course of my research in Leicester I have received a lot of support from my family, friends and colleagues. I would like to thank them all for their help and encouragement over the past three years.

I would like to thank particularly Fernando S. Schlindwein for his guidance and help during the course of my research.

Thanks also to the University of Leicester for funding my research and also to the Leicester Royal Infirmary, particularly, Prof. D.H. Evans and Abigail Thrush for their help in providing me with data for my research.

Special thanks to the new friends that I have made whilst being at Leicester for making my time here a memorable one. I think out of all the mountains that I have climbed so far with them, my PhD may not have been the highest one but it was certainly the most 'technically difficult'.

I would also like to thank my family for giving me their moral and financial support throughout my research and for their encouragement every step of the way.

Last and by no means least, thanks to my future wife, Emma, for believing in me and encouraging me.

Contents

MODERN DIGITAL SIGNAL PROCESSING TECHNIQUES APPLIED TO DOPPLER ULTRASOUND

by

Paul Ivan John Keeton

ABSTRACT — Doppler ultrasound is used clinically to detect stenosis in the carotid artery. The presence of stenosis may be identified by disturbed flow patterns distal to the stenosis which cause spectral broadening in the spectrum of the Doppler signal around peak systole. This thesis investigates the ability of the short-time Fourier transform (STFT) and the autoregressive (AR) spectral estimators to perform time-frequency analysis of the non-stationary Doppler signal. Quantitative analysis of the degree of spectral broadening was measured using the spectral broadening index (SBI). A real-time system was developed using a modern DSP board combined with an IBM PC-compatible computer to analyse the Doppler signal in real-time using the STFT and AR algorithms.

The spectral estimators were compared using simulated Doppler spectra contaminated with noise over a range of signal-to-noise ratios (SNRs) and also real clinical Doppler signals recorded from both healthy subjects and patients with varying degrees of stenosis. The SBI was calculated using the mean and maximum frequency envelopes which were extracted from the STFT and AR sonograms using a threshold at -6 dB of the maximum magnitude component of each individual spectrum. The results of the analysis show a strong correlation between the indices calculated using the FFT and AR algorithms. A qualitative improvement in both the appearance of the AR sonograms and the shape of the individual AR spectra was noticeable, however, the estimation of SBI for short data frames is not significantly improved using AR.

The final section of this thesis describes the wavelet transform (WT) and illustrates its application to Doppler ultrasound with two examples. Firstly, it is shown how wavelets can be used as an alternative to the STFT for the extraction of the time-frequency distribution of Doppler ultrasound signals. Secondly, wavelet-based adaptive filtering is implemented for the extraction of maximum blood velocity envelopes in the post processing of Doppler signals.

Glossary

ACS	Autocorrelation sequence
ADC	Analogue-to-digital converter
AR	Autoregressive
ARMA	Autoregressive moving-average
CCA	Common carotid artery
CFT	Continuous Fourier transform
CVA	Cerebral vascular accident
CW	Continuous wave
CWT	Continuous wavelet transform
DAT	Digital audio tape
DFT	Discrete Fourier transform
DTFT	Discrete-time Fourier transform
DSP	Digital signal processing
DWT	Discrete wavelet transform
ECA	External carotid artery
ECG	Electrocardiogram
EEG	Electroencephalogram
FFT	Fast Fourier transform
FT	Fourier transform
GUI	Graphical user interface
Hm	Hamming
ICA	Internal carotid artery
ICFT	Inverse CFT
IDFT	Inverse DFT
IDTFT	Inverse DTFT
ISR	Interrupt service routine
JTFA	Joint time-frequency analysis
LSB	Lower sideband
LSI	Loughborough Sound Images
MA	Moving-average
MRA	Multiresolution analysis
PC	Personal computer
PI	Pulsatility index
PRF	Pulse repetition frequency
PSD	Power spectral density
PW	Pulsed wave
RAM	Random access memory
RBC	Red blood cell
RI	Resistance index
SBI	Spectral broadening index
SNR	Signal-to-noise ratio
SPTA	Spatial-peak, temporal average
STFT	Short-time Fourier transform
TIA	Transient ischaemic attack
USB	Upper sideband
WFT	Windowed Fourier transform
WT	Wavelet transform

1. Introduction

The aim of this study was to investigate the use of alternative signal processing tools for the analysis of Doppler ultrasound signals. This thesis is divided into 3 main sections: Firstly, the spectrum analysis of real-time Doppler ultrasound signals is introduced with a view to producing a real-time system; secondly, a comparison of spectrum analysis tools is made to establish whether modern spectrum analysis tools can improve upon the commercially accepted Fourier techniques; finally, the potential of using wavelets is discussed to assess their application to either the spectrum analysis of the Doppler signal or the post-processing of the joint time-frequency decomposition.

In this chapter a brief description of the physiology of the cardiovascular system is given and the background to Doppler ultrasound is outlined. The remainder of the chapter provides a summary of the limitations of Doppler ultrasound and safety considerations associated with the exposure of patients to ultrasound. Chapter 2 outlines the fundamentals of joint time-frequency analysis (JTFA) with particular emphasis on two spectrum analysis techniques, namely, the Fourier transform (FT) and autoregressive (AR) modelling, which have been used to analyse Doppler signals. Examples of JTFA using test signals are detailed and the results obtained using Fourier and AR modelling techniques are compared and contrasted. One of the important technical considerations for commercial systems is the implementation of real-time algorithms. This topic is addressed in chapter 3 and a real-time system using a commercial DSP board is developed. In chapter 4 the detection and diagnosis of stenosis using Doppler ultrasound is introduced. The relationships between the time-frequency profile of the received Doppler signal and the diagnosis of stenosis using qualitative and quantitative methods are discussed. The implementation of an off-line system for the JTFA of the Doppler signal and post-processing of the time-frequency profile for the quantitative analysis of stenosis using the ideas developed in the previous chapter is outlined in chapter 5. Chapter 6 examines the robustness of the signal processing tools to noise and the effect that noise has on the quantitative

analysis of simulated Doppler signals. In chapter 7 the signal processing and post-processing algorithms that have been developed in the previous chapters are used to analyse real clinical data for both healthy individuals and patients with carotid disease. Chapter 8 introduces a relatively new signal processing tool that is being increasingly used in the biomedical field, namely, wavelet analysis. The potential of applying wavelets to Doppler ultrasound is addressed. Chapter 9 brings this study to a close and summarises the key-points that this research project has achieved and future extensions to the work that could be pursued.

1.1 The cardiovascular system

The cardiovascular system is a complex network of vessels that act as a transportation system to maintain homeostasis within the body. The cardiovascular system deteriorates with age and cardiovascular disorders are particularly prevalent in older age groups (> 50 years). In the United States approximately 98 % of the 100 000 fatalities each year from cerebral vascular accident (CVA), otherwise known as stroke, are over the age of 50 (Carola *et al.*, 1992). The deterioration of the cardiovascular system can be accelerated due a poor diet, lifestyle and a lack of general fitness. Numerous disorders are associated with the atherosclerosis and the general deterioration of the cardiovascular system, these include: thrombosis, embolism, transient ischaemic attacks (TIAs), stroke, aneurysms and haemorrhage. Hence, it is very important to be able to observe the condition of vessels *in vivo* and to detect the onset of cardiovascular disease. The ability to do this non-invasively using techniques such as Doppler ultrasound, which is the subject of this thesis, has obvious advantages over invasive techniques such as angiography. Therefore, the advancement of quantitative methods for detecting cardiovascular diseases using diagnostic vascular ultrasound is essential.

1.1.1 Blood flow

Blood is not a homogeneous liquid and therefore cannot be treated as a Newtonian fluid. Instead blood is a colloidal solution which consists of approximately

55 % plasma and 45 % formed elements by volume, which include: erythrocytes (red blood cells); leukocytes (white blood cells) and thrombocytes (platelets). Plasma is normally around 90 % water, the remaining 10 % is made up from proteins and other trace chemicals. The red blood cells (RBCs) have a diameter of approximately 7 μm and constitute the major proportion by volume of the different formed elements within blood. Due to the non-uniform density of the RBCs within blood the haematocrit level varies from 36 % to 54 % (Evans *et al.*, 1989)[1].

1.1.2 Anatomy of the carotid branch

The main focus of this research is based on recordings made from the carotid branch, therefore, the following section gives a more in depth description of the anatomy of the carotid arteries and details the geometry of the vessels and typical velocity patterns that are expected in normal patients.

The carotid arteries are located in the neck, the left and the right common carotid arteries (CCAs) originate from the ascending aorta. The CCA has an average diameter of 6.2 mm[†]. The CCA bifurcates into the internal carotid artery (ICA) and the external carotid artery (ECA) at the upper part of the larynx. There are two sets of ICAs and ECAs originating from the left and right CCAs respectively. Immediately distal to the carotid bifurcation in the ICA is the carotid bulb. The carotid bulb at the bifurcation has a diameter of 6.2 mm[†] which tapers into the ICA with a diameter of 5.4 mm[†]. The ECA serves as the main blood supply to the face. The ICA is one of the main blood supplies to the brain along with the vertebral arteries, the ICA terminates at the 'circle of Willis' where it divides into the anterior and middle cerebral arteries.

Typically the CCA exhibits fully developed laminar flow with a flattened, symmetrical velocity profile, the normal maximum velocity at peak systole is of the order 0.4-1.2 ms^{-1}[§] . The ICA has disturbed flow in the carotid bulb exhibiting flow separation, the velocity profile in the ICA is initially predominately plug flow with

[†] Measurements taken from subjects between 20-29 yrs, figures obtained from Reneman *et al.* (1992).
[§] Figures obtained from Hedrick *et al.* (1995).

peak velocities of 0.6-1.0 ms^{-1}§. The ECA has typical velocities in the range of 0.6-1.2 ms^{-1}§. Reverse flow is not characteristic of the carotid arteries except in the carotid bulb. The presence of reverse flow is indicative of an abnormality provided the recordings are taken from a site at least 3 cm from the carotid bulb (Reneman *et al.*, 1992).

1.2 Diagnostic vascular ultrasound

In 1959 Satomura demonstrated that the velocities of the RBCs within blood could be tracked non-invasively by identifying the Doppler shifts from a backscattered ultrasound signal (Satomura, 1959). The non-invasive system used to generate the ultrasound beam and detect the backscattered echoes is called a Doppler velocimeter. The probe transmits and receives an ultrasound beam (figure 1.1) via piezoelectric transducers. The exact configuration of the velocimeter depends upon whether continuous-wave (CW) or pulsed-wave (PW) ultrasound is used. CW and PW ultrasound are discussed in more detail later in this chapter.

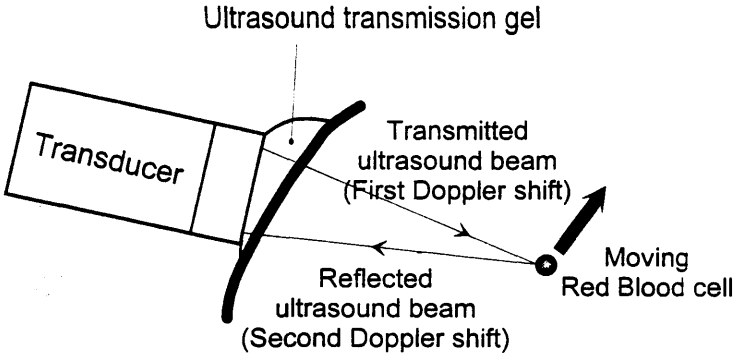

Figure 1.1 Simplified diagram illustrating the principles of Doppler ultrasound.

As already mentioned above the velocities of the RBCs can be tracked by extracting the Doppler shifts from the backscattered ultrasound signal. The beam is reflected or scattered as a result of acoustic impedance mismatches at boundaries between different surfaces within the body. The moving RBCs within a blood vessel backscatter the ultrasound beam and shift the frequency of the signal, this

phenomenon is known as the Doppler effect and is described below. The ultrasound beam has finite dimensions and therefore the backscattered signal contains echoes from a population of moving RBCs (as well as other echoes, for example, from tissue boundaries and the vessel walls). The received ultrasound signal therefore contains a multitude of Doppler shifts which ideally represent the histogram of velocities present in the vessel at a particular instant in time.

The ultrasound beam is focused at a specific depth in order to maximise the magnitude of echoes from the area of interest. Nowadays positioning of the sample volume is facilitated with Duplex scanning which gives the ultrasonographer simultaneous imaging of the examination area and tracking of the changing velocity profile as a function of time. Modern 'colour Doppler' systems superimpose the blood velocity components onto the grey-scale image to show the flow in the vessel.

1.2.1 Doppler effect

The Doppler effect is the perceived change in frequency of a signal that has been emitted by a moving point source as observed by a receiver positioned relative to the source. This phenomenon can be explained using a simple example: Two stationary receivers are positioned at A and B, with a point source between A and B emitting wavefronts at a constant frequency f as illustrated in figure 1.2. When the source is stationary each receiver is struck by wavefronts at a constant frequency f equal to the frequency of the source. If the source is moving towards A with a constant velocity v, then the wavefronts that reach A are closer together. As a result A perceives an increase in frequency as can be seen in figure 1.2. If the source is moving towards A then it is moving away from B, hence, B perceives a decrease in frequency since the wavefronts that reach B are further apart.

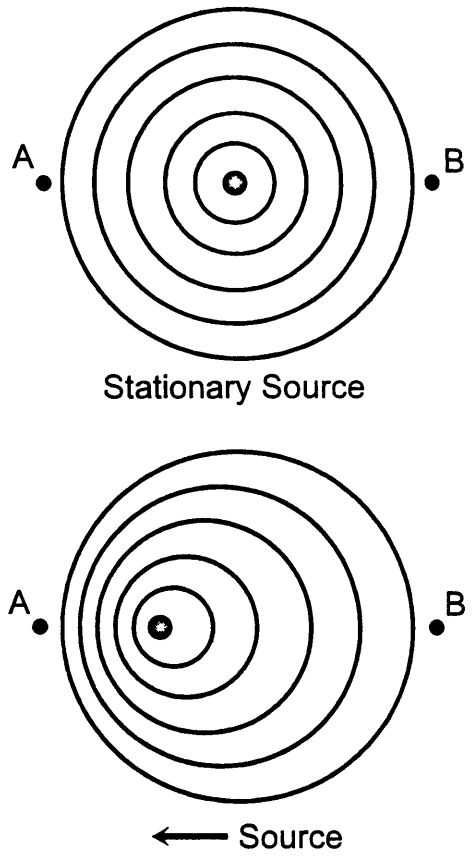

Figure 1.2 The Doppler effect: Stationary source ⇒ No change in frequency; Source moving towards (A) at a constant velocity ⇒ (A) sees a positive Doppler shift (B) sees a negative Doppler shift.

The apparent change in frequency is known as a Doppler shift Δf_D. The Doppler shift can be either positive or negative depending upon the relative motion of the source towards or away from the receiver respectively. The Doppler shift can be expressed using the well known Doppler equation [1.1]:

$$\Delta f_D \cong \frac{vf\cos\theta}{c} \qquad [1.1]$$

The same phenomenon is perceived if the source is stationary and the receiver is moving: if the receiver is moving towards the source then the receiver observes an increase in frequency; on the other hand, if the receiver is moving away from the source then a reduction in frequency will be observed. Therefore, in Doppler ultrasound their are two Doppler shifts: the first is caused by the Doppler velocimeter

acting as the source and the red blood cells acting as the targets; secondly, the ultrasound is scattered by the red blood cells in which case the red blood cells are acting as the source and the receiver is the Doppler velocimeter, this is illustrated in figure 1.1. The equation for Doppler ultrasound is a modified version of the original Doppler equation with a factor of two to account for the double Doppler shift [1.2]:

$$\Delta f_{DU} \cong \frac{2vf\cos\theta}{c} \qquad [\,1.2\,]$$

where: Δf_{DU} = change in frequency between transmitted ultrasound beam and received echo.

1.2.2 Continuous-wave and Pulsed-wave ultrasound

CW and PW ultrasound are the two types of systems used in diagnostic ultrasound. As its name suggests a CW ultrasound instrument emits a continuous ultrasound beam which is normally focused at a specific depth. The CW probe uses two transducers: one for transmitting the ultrasound and a separate transducer to receive the backscattered ultrasound. The PW ultrasound probe uses a single transducer which doubles up as the transmitter and receiver: the transducer emits a short ultrasound pulse and then waits for a period of time to receive the backscattered echoes. Using PW ultrasound it is possible to calculate the depth (l) of the target that reflected the ultrasound since the duration (t) between the emission of the ultrasound pulse and the reception of the echo can be controlled and the velocity of ultrasound (c) in tissue is known. Equation [1.3] shows the relationship between the depth of the target and the time taken to receive the echo. It is therefore possible to range gate the returning echoes so that only echoes from a specific depth are received, this helps to remove low frequency high intensity echoes from tissue boundaries and low frequency wall motion.

$$l = \frac{ct}{2} \qquad [\,1.3\,]$$

Depth measurement is not possible with CW ultrasound due to the continuous ultrasound beam. The main problem with PW ultrasound is that the maximum Doppler shift that can be measured using PW ultrasound is restricted to half the pulse repetition frequency (PRF). The PRF is dictated by the depth of the ultrasound scan, therefore, as the PRF is reduced the range of blood velocities that can be measured unambiguously is restricted. CW systems do not suffer from this problem.

1.2.3 Limitations of Doppler ultrasound

Ideally the backscattered signal from a blood vessel will contain a multitude of Doppler shifts whose frequencies are directly proportional to the velocities of the population of moving targets (mainly RBCs) within the vessel. The spectral broadening of the backscattered signal is therefore a measure of the spread of velocities within the vessel. Researchers have attempted to use spectral broadening as a diagnostic tool for characterising the type of flow within the vessel and to use it for the clinical assessment of arterial disease (Blackshear *et al.*, 1979). There are a number of considerations, however, that are implicit in practical Doppler ultrasound systems that affect this ideal relationship between spectral broadening and the spread of velocities. These factors distort the true frequency spectrum which reduces the sensitivity of spectral broadening parameters derived from the spread of Doppler shifts (Jones, 1993).

The Doppler ultrasound examination by its nature is non-invasive therefore it is only possible to make an estimate of the angle between the axis of the ultrasound beam and the velocity vector. The relationship between the magnitude of the Doppler shift and the velocity of the RBC can only be estimated. The Doppler angle is usually taken as the orientation of the vessel with respect to the ultrasound beam which can be deduced with the aid of Duplex scanning. This is an inaccurate assumption about the velocity vectors within the vessel, since, even in normal patients axial flow is generally not present, instead the flow is usually helical (Beach and Phillips, 1992). If the flow is turbulent then the velocity vectors within the vessel will be moving in

random directions and therefore it is impossible to get an exact relationship between velocity and frequency.

The width of the transducer is finite, therefore, the estimation of the velocity of a single target moving at a constant velocity through the ultrasound beam is subject to intrinsic spectral broadening due to the range of angles that are received by the Doppler velocimeter.

The velocity of ultrasound (c) in blood lies in the range 1540 - 1600 ms^{-1} (Evans *et al.*, 1989)[1]. The accuracy of the velocity estimated from the magnitude of the Doppler shift is therefore limited since the exact speed of ultrasound at the point of scattering by the RBC is not known exactly. This gives an error of approximately up to 4 % depending upon the choice of c.

The constitution of blood is made up of formed particles suspended in plasma. The motion of these particles is not independent of one another and therefore the direction of motion is constantly changing as they collide with one another. This also causes spectral broadening.

Doppler ultrasound examinations are carried out using a range of beam widths. The insonation of the vessel is dependent upon the beam width and the intensity of the ultrasound across the cross-section of the beam. If the beam is non-uniform then there will be a bias in the relationship between the magnitude of the Doppler shifts and the velocities of the RBCs. A narrow beam is usually focused on the centre of the vessel to identify the maximum velocity components. Using this type of beam the flow near the vessel walls will not be well represented, therefore, a narrow beam will have a frequency spectrum that is fairly narrow and biased towards the major velocity components which are typically found in midstream. A wide ultrasound beam will insonate the whole vessel more uniformly allowing for an unbiased representation of all velocities across the vessel lumen. The problem with a wide beam is that it picks up strong low frequency components as a result of wall motion, known as '*wall thump*'. A wide beam is also more likely to overlap another vessel.

The beam width is finite, therefore, the backscattered Doppler shifts are affected by the movement of the individual RBCs through the beam. If a RBC is accelerating then the frequency estimate will be smeared further.

The ultrasound beam is emitted from the Doppler velocimeter and is focused at a specific depth within the body. The frequency of the ultrasound beam that can be used in the examination is dependent upon the location of the vessel, since the attenuation of the ultrasound signal increases with frequency.

1.2.4 Safety considerations

Non-invasive diagnosis has obvious advantages over other diagnostic techniques, nevertheless, the use of ultrasound has received much attention with regards to the safety of prolonged exposure to ultrasound. The considered acceptable intensity (I_{SPTA}) of ultrasound beams during a normal examination is about 100 mW cm^{-2} (Hussey, 1975; Evans, 1989[1]). Despite these guidelines modern pulsed Doppler systems often far exceed these recommended power levels. Effects of prolonged ultrasound exposure include heating, streaming and cavitation (McDicken, 1991). In a controlled environment the heating effect can be used to benefit patients for example in physiotherapy. In the case of neonatal monitoring, however, these effects are obviously cause for concern.

1.3 Time-frequency profile of blood flow

The ideal time-frequency profile of the received Doppler signal is therefore proportional to the velocity of the population of targets moving within the vessel. The characteristic type of flow and the velocity of blood within a vessel depends upon the geometry of the vessel, the distance of the vessel from the heart and the physiology of the vessel, for example, the presence of valves. In some cases bi-directional flow is present (Gosling and King, 1974).

1.4 Summary

This chapter has outlined the principles of Doppler ultrasound and discussed some its limitations and safety aspects. In the next chapter the focus is directed towards analysis of the raw Doppler signal using JTFA.

2. Joint Time-Frequency Analysis

2.1 Introduction

The following chapter is intended to outline the theory behind the signal processing tools that are going to be implemented in the following chapters.

This text is not intended to give an exhaustive description of Fourier and autoregressive (AR) modelling techniques which have now been documented in numerous books (Marple, 1987; Oppenheim and Schafer, 1989). Instead a brief review of the key-points are summarised here, with a view to joint time-frequency implementation of these algorithms. The first section of this chapter explains why frequency analysis of the Doppler signal can provide an estimate of the velocity components within the vessel. This will be followed by a discussion of stationary and non-stationary signals. The next section looks at how the classical Fourier transform (FT) can be adapted to analyse both stationary and non-stationary signals. The last section of this chapter will discuss alternative modern signal processing tools such as AR modelling that have been developed in an attempt to improve upon the spectral matching ability and the resolution offered by the FT.

2.2 Spectrum analysis

In the time domain it is often difficult to interpret the data that is being collected. Identification of the frequency content of a signal is only possible for relatively simple signals and becomes impossible when several frequency components are present. The problem is compounded when the signal is masked by noise. It is often more convenient to transform the signal into the frequency domain to visualise its spectral content. In the case of Doppler ultrasound it is convenient to transform the Doppler signal into the frequency domain, since the velocities of the red blood cells (RBCs) are proportional to the Doppler shift frequencies as indicated by [2.1]. The

$$v \cong \frac{\Delta f_D c}{2 f \cos\theta} \quad \Rightarrow \quad v \propto \Delta f_D \qquad [2.1]$$

frequency spectrum of the Doppler signal ideally represents the histogram of the number of RBCs moving with a particular velocity as illustrated in figure 2.1.

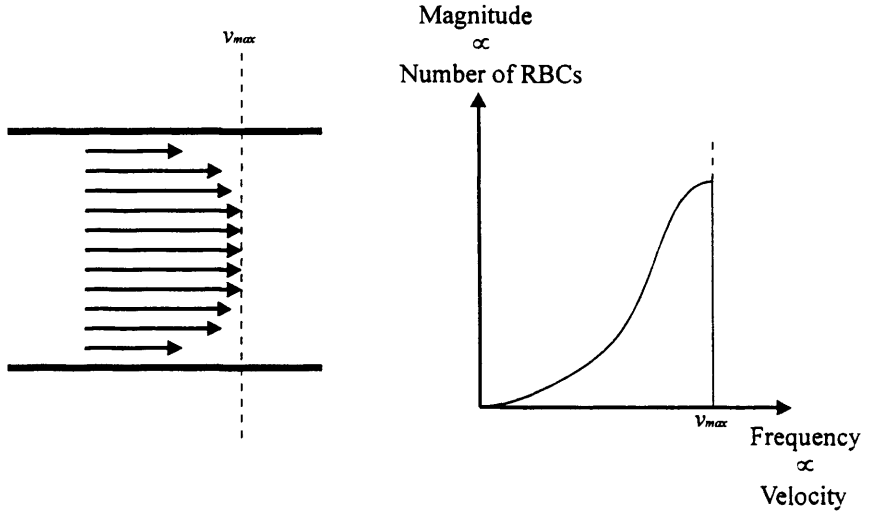

Figure 2.1 Relationship between velocity profile and frequency spectrum.

2.2.1 Stationary and non-stationary signals

Stochastic signals can be categorised into two groups: stationary or non-stationary. A wide-sense stationary stochastic signal is one whose first and second order statistical moments are time invariant (Marple, 1987). If the signal has time varying frequency characteristics then it is considered to be non-stationary.

The next section describes the classical Fourier transform and then explains why the conventional FT cannot be used for analysing non-stationary signals. A brief explanation of how the FT can be modified to analyse non-stationary signals is given and the limitations and trade-offs are examined.

2.2.2 The Fourier transform

The continuous Fourier transform (CFT) [2.2] evaluates the magnitudes of an infinite number of complex sinusoids which can be used to model a signal *x(t)*.

Sinusoids are optimally localised in frequency and therefore the magnitude of the coefficients $X(f)$ represent the contribution of each frequency to the overall signal. The function $X(f)$ is therefore the frequency domain representation of the time signal. The original time domain signal can be constructed using the inverse continuous Fourier transform (ICFT) [2.3].

$$X(f) = \int_{-\infty}^{\infty} x(t)\exp(-j2\pi ft)dt \qquad [2.2]$$

$$x(t) = \int_{-\infty}^{\infty} X(f)\exp(j2\pi ft)df \qquad [2.3]$$

where: $x(t)$ = continuous-time function, $X(f)$ = continuous-frequency function.

If the time domain signal is discretised at intervals of T (where: $f_{samp} = 1/T$) then the discrete version of the CFT is:

$$X(f) = \sum_{n=-\infty}^{\infty} x(nT)\exp(-j2\pi fnT) \qquad [2.4]$$

$$x(nT) = \int_{-1/2T}^{1/2T} X(f)\exp(2\pi fnT)df \qquad [2.5]$$

where: $x(nT)$ = discrete-time series, $X(f)$ = continuous-frequency function. Equation [2.4] is known as the discrete-time Fourier transform (DTFT). The discretisation of the time series forces periodicity of $(1/T)$ in the frequency domain. The discrete-time signal can be obtained from the function $X(f)$ using the inverse discrete-time Fourier transform (IDTFT) [2.5] by integrating over the period $1/T$.

The CFT and the DTFT assume that the time domain signal is defined from $-\infty$ to $+\infty$. In reality signals are not infinite, they have a definite start and finish point. If the signal is discretised then there are only a finite number of samples N. The discrete Fourier transform (DFT) [2.6] is a finite approximation to the DTFT, where the frequency domain is discretised as a result of the truncation of the time domain signal.

$$X\left(\frac{k}{NT}\right) = \sum_{n=0}^{N-1} x(nT)\exp\left(\frac{-j2\pi kn}{N}\right) \qquad [2.6]$$

$$x(nT) = \frac{1}{N}\sum_{k=0}^{N-1} X\left(\frac{k}{NT}\right)\exp\left(\frac{j2\pi kn}{N}\right) \qquad [2.7]$$

where: $x(nT)$ = discrete-time series, $X(k/NT)$ = discrete-frequency series. The notation for the discrete-time and discrete-frequency series is often simplified to just $x(n)$ and $X(k)$, this notation will be used from this point onwards. The DFT is computed for harmonic sinusoids of frequency k/NT. The discretisation of the frequency domain implicitly assumes that the time signal is periodic outside the known N discrete points. Since the frequency domain is periodic only a finite number of discrete orthogonal harmonics exist that can be used to describe the time domain signal. As a result of this implicit periodicity in the time domain pseudo discontinuities will exist at the boundaries of the period if the signal contains frequency components which are not exact multiples of $1/NT$. These discontinuities manifest themselves as spectral leakage in the spectral estimate. The inverse discrete Fourier transform (IDFT) [2.7] can be used to generate the time domain signal from the discrete frequency coefficients.

The analysis of stationary signals using the DFT is straightforward assuming the signal is adequately sampled according to the Nyquist criterion (Oppenheim and Schafer, 1975). If the signal is non-stationary, however, then it is futile to analyse the whole signal in one go since this will not provide any information about how the signal's frequency characteristics are changing with respect to time.

2.3 Joint time-frequency analysis

In order to obtain all the information about a non-stationary signal it is necessary to obtain not only the frequency information about the signal but also how the frequency content of the signal changes with respect to time. This is known as joint time-frequency analysis (JTFA).

2.3.1 Short-time Fourier transform

To observe what is happening in a non-stationary signal it is necessary to divide the signal up into shorter sections by multiplying it by a windowing function. If the signal is truncated into short enough sections then it can be considered to be stationary within the window. It is then possible to apply the DFT to extract the frequency spectrum of the windowed signal. This modified DFT is called the short-time Fourier transform (STFT) or windowed Fourier transform (WFT) [2.8] and is reasonably good at localising the spectral content in time.

$$X(k,n_0)=\sum_{n=0}^{N-1}x(n)w(n-n_0)\exp\left(\frac{-j2\pi nk}{N}\right) \qquad [2.8]$$

where: $w(n)$ = windowing function. By sliding the window over the signal it is possible to identify how the frequency characteristics of the signal are changing with respect to time, this is illustrated in figure 2.2.

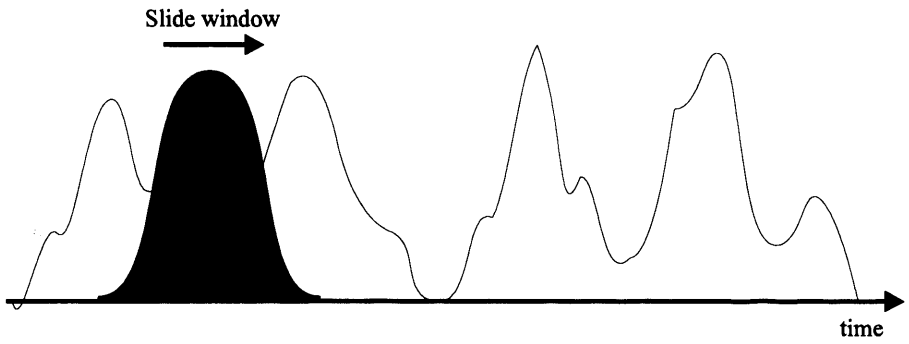

Figure 2.2 Illustration of the application of a sliding-window to a non-stationary signal in order to perform the STFT.

There are a number of limitations that are implicit in the definition of the STFT. The STFT suffers from the same problems as the DFT with respect to spectral leakage due to discontinuities at the boundaries of the windowing function. As for the DFT the resolution of the STFT in the frequency domain is limited to $1/NT$. When analysing highly non-stationary signals the number of data points that can be used to compute the STFT (N) is often small therefore the resulting frequency resolution is extremely poor. If only a few data points are available for the extraction of the

frequency spectrum then the variance of the spectral estimate compared to the true underlying spectrum is going to be high. In this situation several estimates of the frequency spectrum should be averaged together to improve the spectral stability of the estimate.

2.3.2 Fast Fourier transform

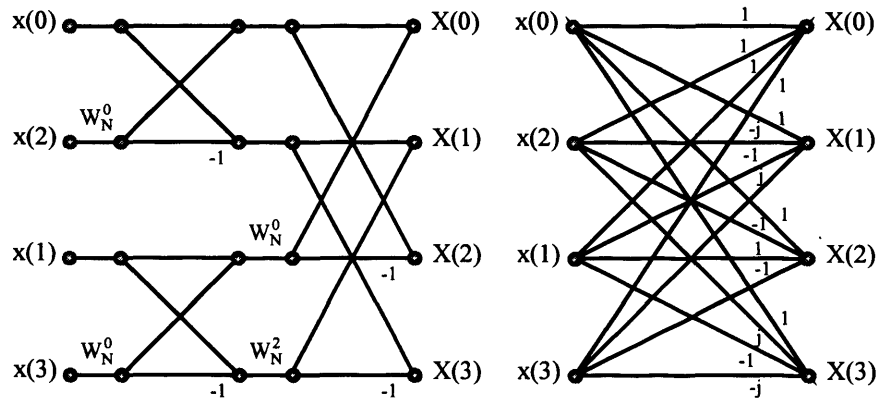

Figure 2.3 Comparison of the relative complexity of decimated-in-time DFT and the 4-point DFT.

The DFT can be solved using the computationally efficient fast Fourier transform (FFT) algorithm. The fast Fourier transform involves decimating the N-point DFT into a number of smaller DFT. In the limit the DFT can be simplified using radix-2 DFTs. The decimated DFT requires only $(N/2)log_2N$ operations (which includes: 6 additions and 4 multiplications) to solve the coefficients compared to the original DFT which requires N^2 operations (Oppenheim and Schafer, 1989). The comparison between the complexity of the N-point DFT and the efficient FFT routine is shown in figure 2.3. Hence the FFT has been widely used commercially in real-time systems.

2.4 Joint time-frequency analysis using the 'sliding-FFT'

The following section gives a review of the techniques that are available and illustrates the trade-offs that are required when performing time-frequency analysis using the STFT.

2.4.1 Stationary example

The first example consists of 4 sinusoids of different frequencies at 89, 202, 212 and 392 Hz respectively. The ideal spectrum of the signal is shown in figure 2.4. The signal is sampled at 1024 Hz. The following section illustrates the methods that can be used to estimate the frequency spectrum.

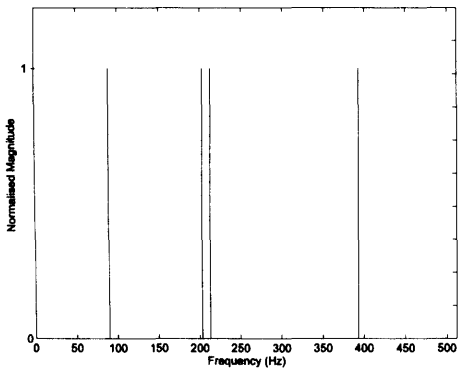

Figure 2.4 Ideal normalised frequency spectrum of example 1 comprised of 4 sinusoids at frequencies of 89, 202, 212 and 392 Hz.

2.4.1.1 Data Frame length

The signal was analysed using a single FFT. The length of the FFT was varied to contrast the different resolutions inherent in the choice of data frame length. The signal was analysed using a 1024, 512, 256, 128 and 64-point FFT. The data frame was isolated from the rest of the signal by multiplying the signal by a rectangular (boxcar) window [2.9]. Figure 2.5 shows the frequency spectrum obtained using each FFT.

$$w(n) = \begin{cases} 1, & \frac{-N}{2} \le n \le \frac{N}{2} - 1, \\ 0, & otherwise. \end{cases} \qquad [2.9]$$

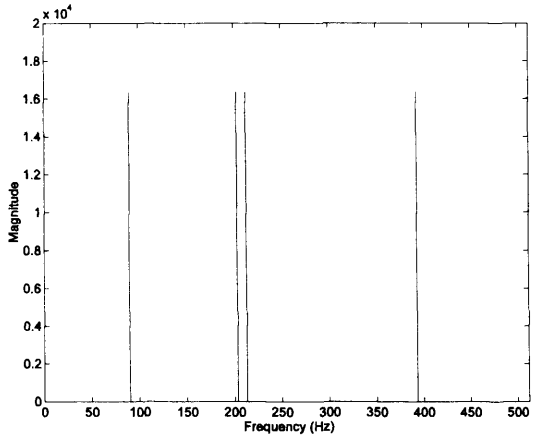

(a) 1024-point FFT

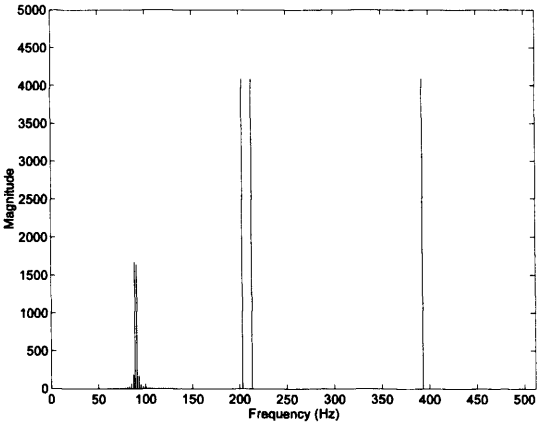

(b) 512-point FFT

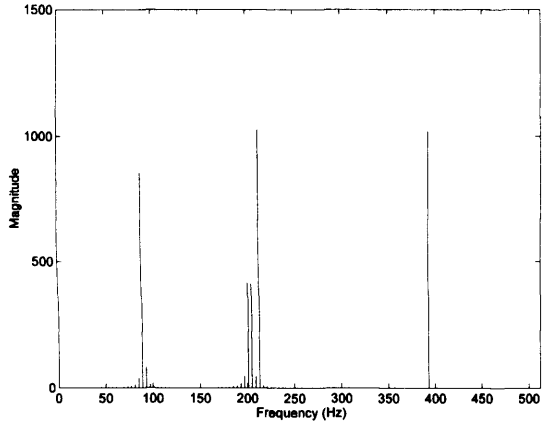

(c) 256-point FFT

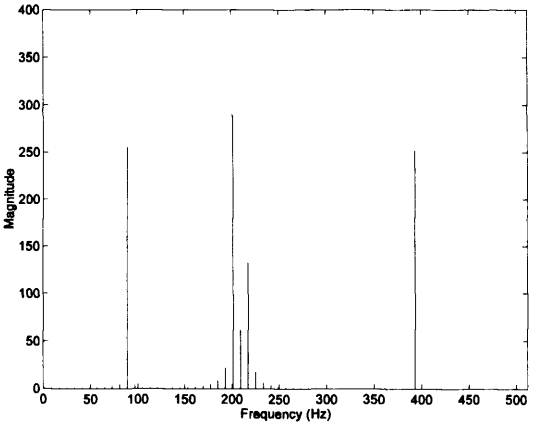

(d) 128-point FFT

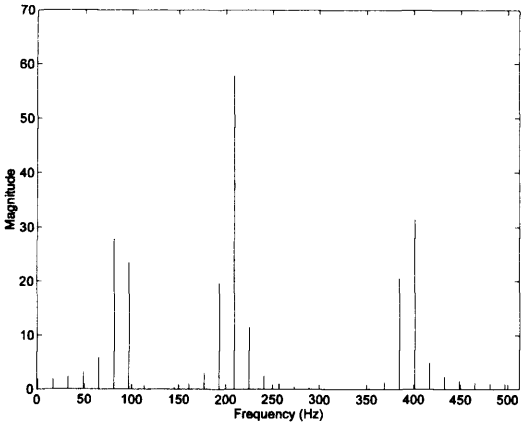

(e) 64-point FFT

Figure 2.5 FFT of example 1 signal containing 4 sinusoidal components at frequencies of 89, 202, 212 and 392 Hz.

It can be seen from figure 2.5 that for the 1024-point FFT all the sinusoidal components are resolved exactly since all the components of the signal are exact multiples of *1/NT*. The sampling frequency was 1024 Hz, therefore, the 1024-point FFT has a resolution of 1 Hz. The 512-point FFT has a resolution of 2 Hz in this case, hence, the FFT is unable to resolve the first sinusoidal (89 Hz) component exactly since it falls between two frequency bins. The FFT spreads the energy contained by this component to the neighbouring frequency bins this is known as spectral leakage. The spreading of the energy due to the 89 Hz sinusoidal component by spectral leakage reduces the magnitude of the peak so that it is smaller but more spread out than the other components in the signal. This makes it difficult to compare the relative magnitudes of the sinusoids. The resolution is still sufficient to successfully resolve the other components. The 256, 128 and 64-point FFTs have resolutions of 4, 8 and 16 Hz respectively. As the resolution deteriorates then the ability of the FFT to resolve the sinusoidal components is also reduced. In the case of the 64-point FFT it is not able to resolve any of the sinusoids exactly and the resolution is so poor that the two close sinusoids of 202 and 212 Hz can no longer be distinguished from one another. When the FFT is computed from a small data frame then the information obtained from the spectral estimate can be ambiguous particularly if the underlying spectrum is unknown. To interpret the frequency estimate accurately it is necessary to know what is happening between the frequency bins available, this can be achieved using zero padding.

2.4.1.2 Zero padding

The number of frequency bins in the frequency domain is limited by the number of data points used to compute the FFT. If a more complete spectrum is required then it is possible to zero pad the windowed signal by introducing zeros to the end of the data frame before applying the FFT. The effect of the zeros is to interpolate between the frequency bins that were obtained from the N points giving a more detailed observation of the frequency response. This results in smoothed spectral estimates due to the sin(f)/f frequency domain interpolation implicitly performed with zero-padding. Zero padding does not improve the resolution of the FFT it just fills in

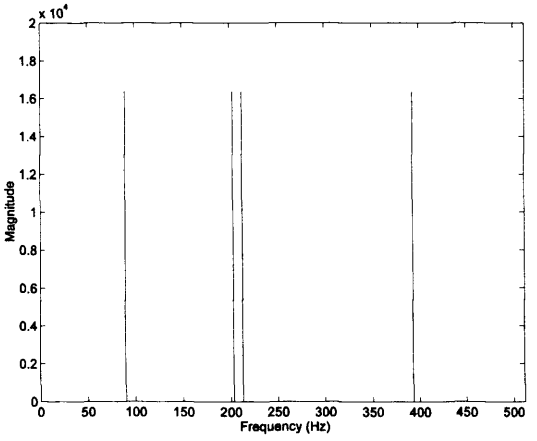

(a) 1024 data points

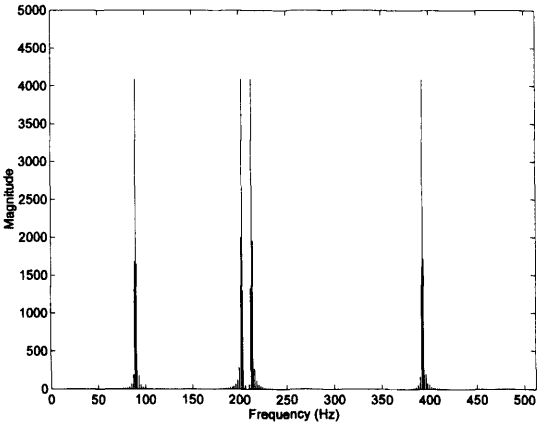

(b) 512 data points

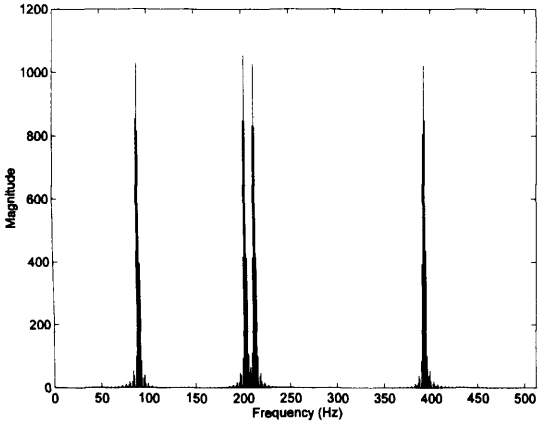

(c) 256 data points

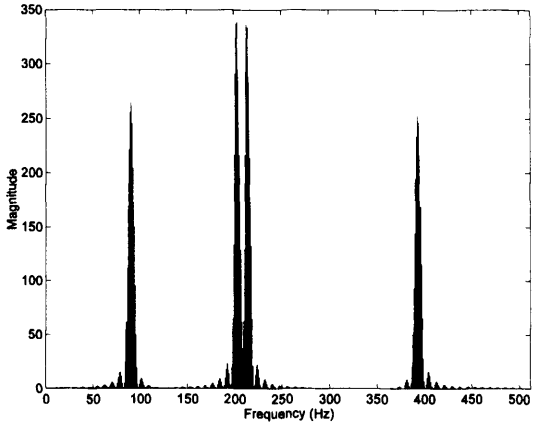

(d) 128 data points

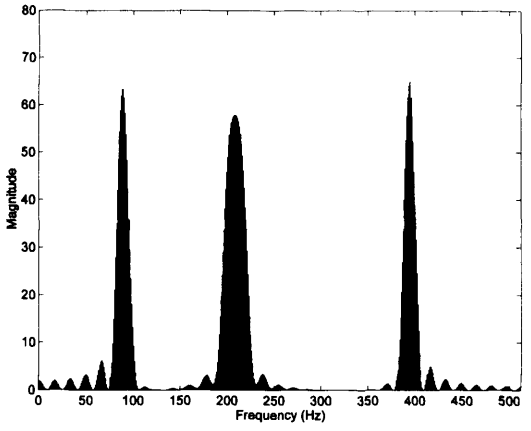

(e) 64 data points

Figure 2.6 1024-point FFT of example 1 with zero padded data frames.

the gaps to give a more complete picture of the frequency response. If the frequency resolution is not good enough then a longer window is required, however, the length of the window should not compromise the stationarity of the signal.

From the first example it can clearly be seen that there is a lack of spectral information available when the FFT is computed using a smaller data frame. This leads to ambiguity about the location of the spectral peaks in the frequency spectrum. Zero padding can be used to interpolate between the existing frequency bins and fill in the gaps. Each of the examples were zero padded up to 1024 points (with the exception of the data frame that already contained 1024 data points). The 1024-point FFT was then applied to the zero padded signals, the results are illustrated in figure 2.6.

It can be seen from figure 2.6 that zero padding provides a clearer picture as to the spectral content of the signal. It can also be seen that zero padding has not improved the resolution for the 64 point data frame, the 202 and 212 Hz sinusoidal components that could not be resolved using the 64-point FFT can still not be distinguished.

2.4.1.3 Anti-leakage windows

The windowing of the signal in time using the STFT creates discontinuities at the edge of the data frame if the signal components within the window are not multiples of *1/NT*. These discontinuities can be minimised by modifying the shape of the window (Harris, 1976). The simplest window is the rectangular window [2.9]. Alternative window types such as the Hamming or Hanning anti-leakage windows have tapered shapes to minimise edge effects. The window weights the components inside the data frame to reduce the discontinuities at the edge of the data frame this has the effect of reducing spectral leakage in the frequency domain. The suppression of side lobes is traded off against an increase in the width of the principal lobe which has the effect of reducing spectral resolution.

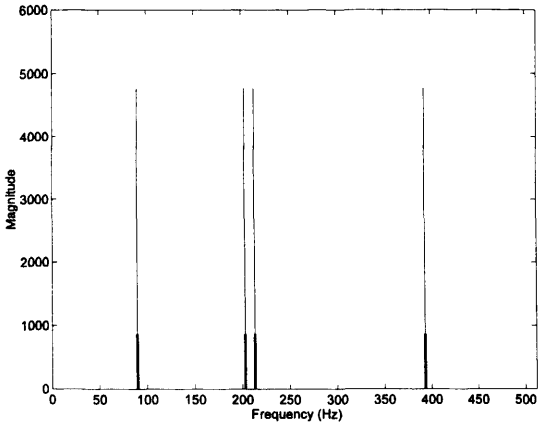

(a) 1024 data points

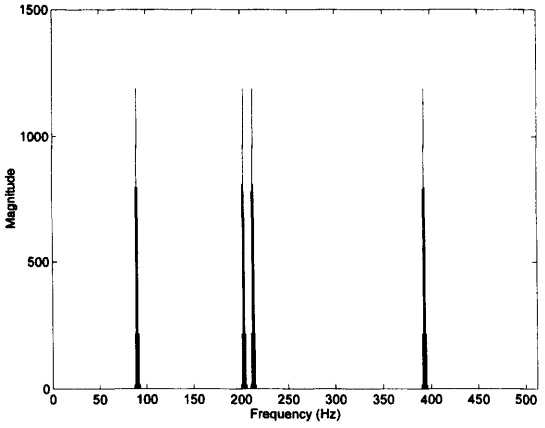

(b) 512 data points

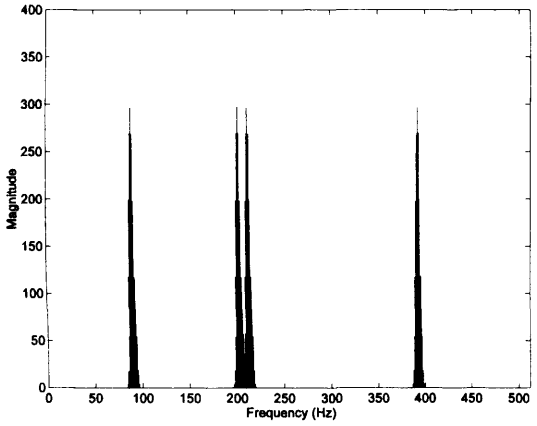

(c) 256 data points

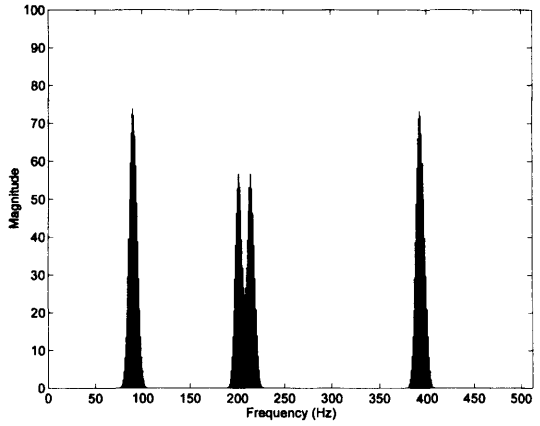

(d) 128 data points

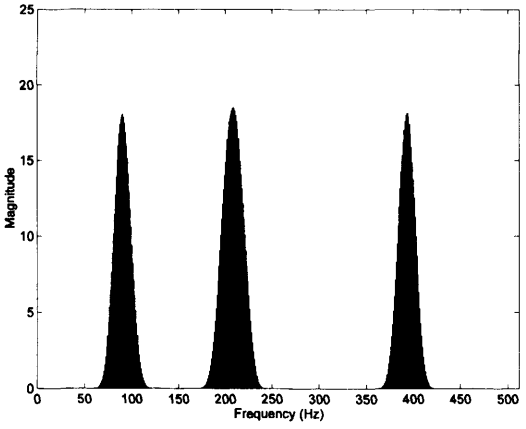

(e) 64 data points

Figure 2.7 1024-point FFT with a Hamming anti-leakage window and zero padding.

A Hamming anti-leakage window was applied to the signal to reduce the effects of spectral leakage for the 512, 256, 128 and 64-point data frames. The results are shown in figure 2.7.

It can be seen from figure 2.7 that the spectral leakage observed in figure 2.6 is suppressed by the application of the Hamming anti-leakage window. Figure 2.7 also shows the effect of the anti-leakage window on the resolution of the frequency spectrum: the two close peaks (202 and 212 Hz) that were totally separated using the 128-point FFT, figure 2.6(d), are now only partially resolved due to the application of the Hamming window, figure 2.7(d).

2.4.2 Spectral instability

The estimate of the frequency spectrum using the STFT produces statistically unstable results when the number of data points used to compute the FFT is small.

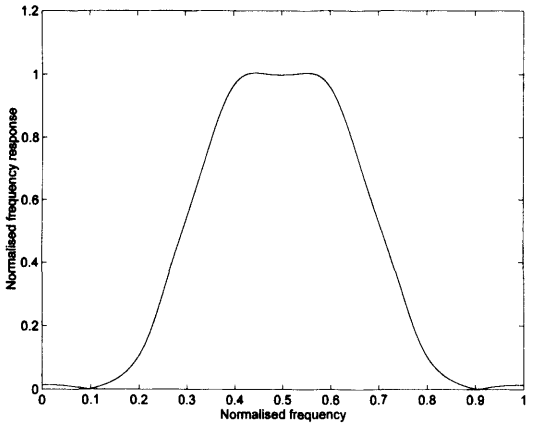

Figure 2.8 Normalised frequency response of filter used to generate example 2.

To illustrate spectral instability a second example is considered: The data is derived from a wide-band signal which was generated from white noise and shaped using a symmetrical filter as illustrated in figure 2.8. The wide-band signal was analysed using 256 point data frames and a Hamming window with zero padding up to 1024 points. The zero padded data frames were then analysed using a 1024-point FFT. The spectral instability of the frequency estimates can be appreciated from figure

2.9 in which several estimates are superimposed. It is interesting to join the discrete bins together to compare the shapes of the estimated frequency spectrum with the ideal filter response illustrated in figure 2.8.

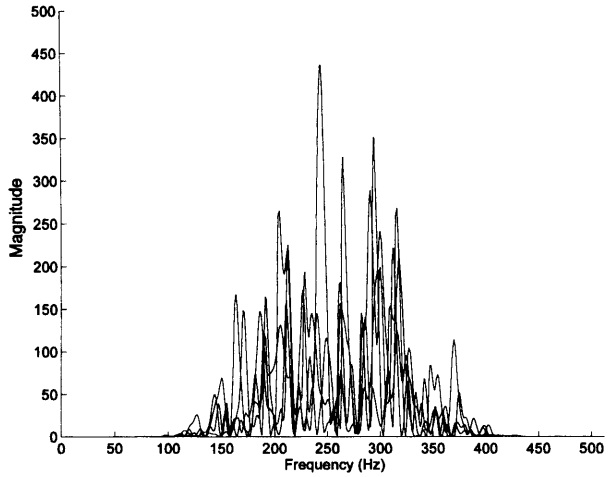

Figure 2.9 Spectral instability of the STFT: Analysis of example 2.

2.4.2.1 Overlapping the data frames

The whole signal contains 1024 data points and in this case it is known that the signal is stationary. When implementing the 512, 256, 128 and 64-point FFTs there are a number of positions that the data frame can be placed in to obtain the number of data points that are required for each FFT. There are two ways in which the data can be processed: non-overlapping and overlapping data frames. The temporal resolution is given by the length of the data frame, N. The number of independent frequency spectra that can be computed is equal to L/N, where L is the number of data points in the digitised signal. For a particular size of data frame analysing the whole signal using non-overlapping data frames is the quickest method of processing all the data. In real-time applications where the computational time to collect and process the data are critical considerations non-overlapping data frames are often the most viable solution. This topic will be discussed in more detail in the next chapter. By overlapping the data frames it is possible to increase the number of frames. The maximum number of frames is given by $L-(N-1)$ where the window is shifted by 1 sample each time. By overlapping the frames the same data is being used to generate successive frequency estimates, therefore the consecutive data frames are linearly

dependent. Sliding the frame over the signal means that all the time domain samples get to contribute equally to the overall spectral estimate, since if non-overlapping data frames are used and the samples at the boundaries of the data frames are weighted using anti-leakage windows then the frequency estimate will be biased. This is compensated for by overlapping the data frames which smoothes out the effects of the window producing a moving-average spectral estimate. For a Hanning anti-leakage window an overlap of 50 % is sufficient to produce unbiased consecutive frequency estimates (Welch, 1967). The first and last *N-1* samples at the boundaries of the signal cannot be equally weighted to estimate the frequency spectrum unless assumptions are made beyond the boundaries of the signal. These are known as edge-effects.

2.4.2.2 Spectral averaging

In an attempt to improve the spectral stability of the frequency estimate it is necessary to average together a number of frequency estimates. In this example the signal is stationary therefore it is possible to average together as many data frames as desired.

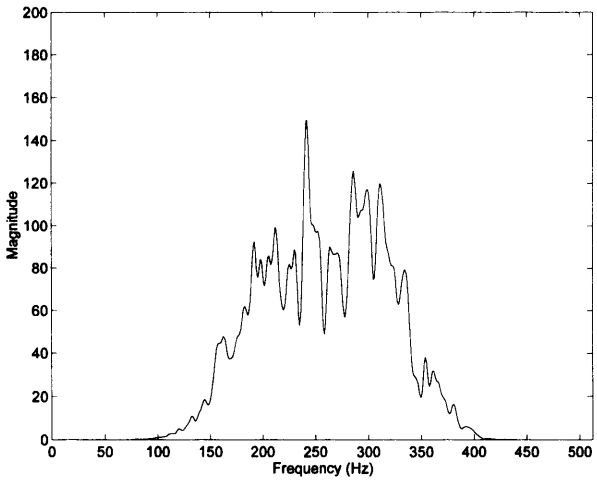

Figure 2.10 Averaged frequency estimate of example 2 (16 frames).

The results show a much closer correlation to the ideal filter response shown in figure 2.8.

2.4.3 Non-stationary example

When analysing stationary signals there is a degree of flexibility in the way that the data can be analysed. It is possible to choose the frame length; the degree of zero padding; the type of anti-leakage window and the amount of spectral averaging. If the signal is non-stationary then the problem becomes more difficult since the signal now needs to be resolved in both time and frequency. The solution is often a compromise between the resolution in the time and frequency domains. The next section illustrates these trade-offs using an example.

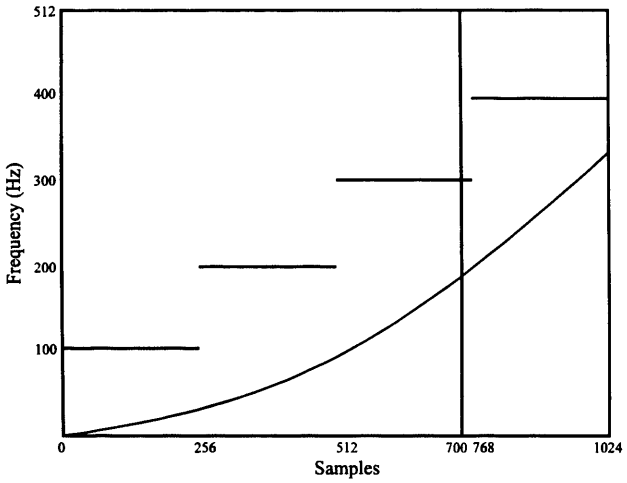

Figure 2.11 Ideal sonogram for example 3 containing 4 sinusoids with frequencies of 100, 200, 300 and 400 Hz separated in time, a non-linear chirp and an impulse at the 700th sample.

The third example is a non-stationary signal (1024 data points) that is comprised of a non-linear chirp ($sin(2\pi ft^3)$), 4 separate non-overlapping sinusoids of 100, 200, 300 and 400 Hz and an instantaneous spike at the 700th sample as shown in figure 2.11. The signal was analysed using a 1024-point FFT (figure 2.12) and then using a sliding 256-point FFT with data frames of 256, 128, 64 and 32 points. The data frames were multiplied by a Hamming anti-leakage window and zero padded up to 256 points. The window was moved by 1 sample each time, the results are illustrated in figure 2.13. In figure 2.13 the consecutive frames have been plotted in the form of a sonogram: horizontal axis is time; vertical axis is frequency; magnitude is plotted using an inverted grey-scale.

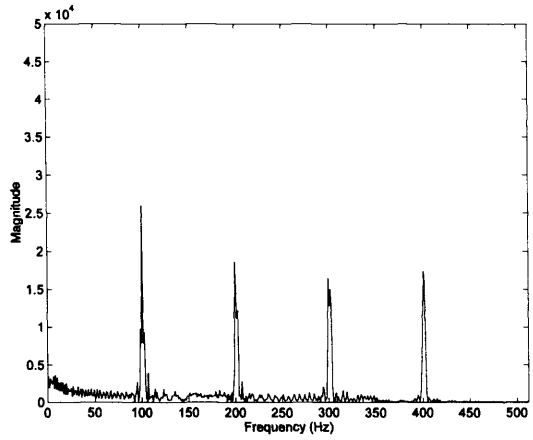

Figure 2.12 1024-point FFT of example 3.

It is now apparent from figure 2.12 that the 1024-point FFT is no longer the best means of analysing the signal. The frequency spectrum produces an overall estimate of the frequency content of the signal but provides no information as to which frequency components were present at a particular instant in time. It is possible to identify the four sinusoidal components, however, the chirp and the instantaneous spike cannot be clearly identified.

The sliding 256-point FFT in figure 2.13 reveals the non-stationarity of the signal: the four sinusoids are well resolved in frequency, however, their temporal resolution is extremely poor. The STFT is unable to localise the start and finish of the sinusoidal components and consequently there is a large overlap between the finish of one sinusoid and the start of the next when it can be seen quite clearly from the ideal sonogram in figure 2.11 that the sinusoidal components are separated in time. In addition the resolution of the chirp is poor and the instantaneous spike is not well localised in time. The sliding 128 and 64-point FFTs improve the time-frequency profile of the chirp and the localisation of the spike in time, however, the resolution in the frequency domain has deteriorated and hence the sharp resolution of the four sinusoids is compromised as a result of the improvement in temporal resolution. The results obtained using the 32-point data frame typify the limited frequency resolution of the FFT when the data frame contains only a few data points. The sinusoidal

components although practically separated in time are no longer concentrated in frequency along with the chirp. The benefits of a short data frame can be observed from the localisation of the instantaneous spike. This exemplifies the need for a trade-off in time and frequency resolution when using Fourier techniques.

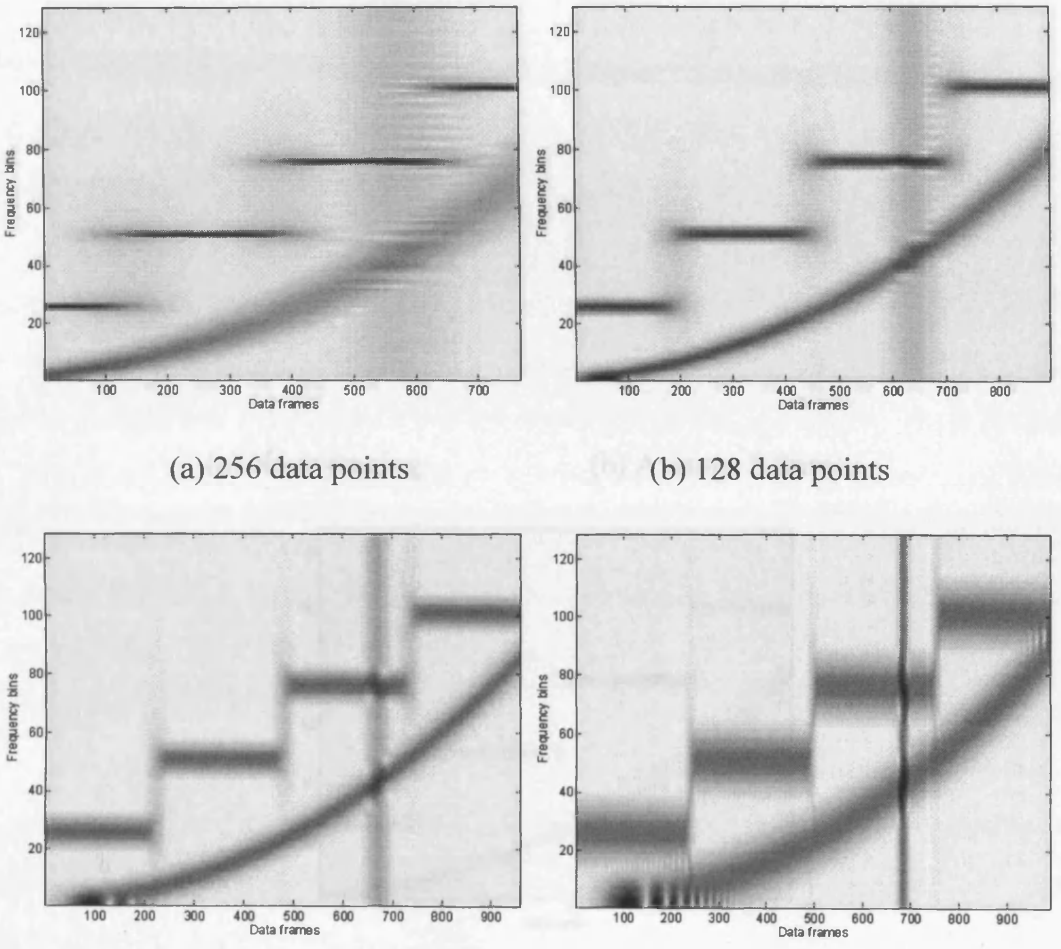

(a) 256 data points (b) 128 data points

(c) 64 data points (d) 32 data points

Figure 2.13 Sliding 256-FFT with zero padding. Window moved by 1 sample each time.

2.4.3.1 Spectral averaging of non-stationary signals

In the case of the stationary example it was possible to average together a large number of spectra to improve the stability of the frequency estimate. In the case of non-stationary signals care must be taken when averaging spectra, since, if too many spectra are used then the total duration of the frames used to compute the averaged

spectra may compromise the stationarity of the signal. In addition, averaging together several data frames helps improve the spectral stability of the underlying trends in the signal but transient events such as instantaneous spikes are flattened by the process of averaging.

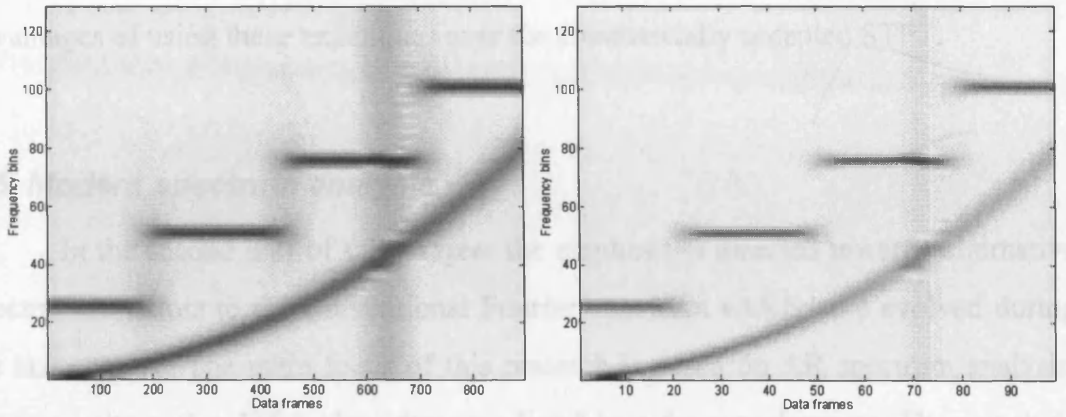

(a) No averaging (b) Average 8 frames

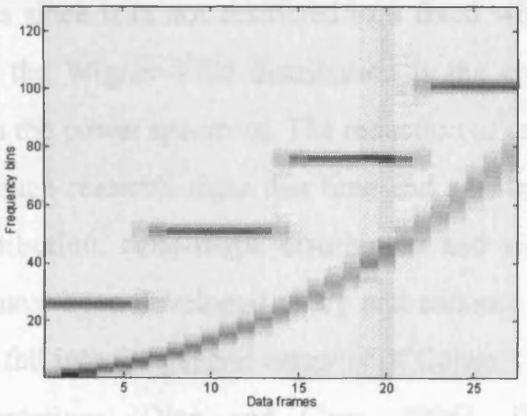

(c) Average 32 frames

Figure 2.14 Effects of spectral averaging on a non-stationary signal.

In figure 2.14 it can be seen that the effect of the averaging process reduces the magnitude of the frequency components which represent the instantaneous spike. Too much averaging leads to a reduction in the frame resolution and hence the smooth transition of the chirp as a function of time is lost. When 32 frames are averaged together the spike is hardly visible.

2.4.4 Summary of Fourier techniques

This section has given a brief description of the continuous and discrete Fourier transform and outlined how the DFT could be modified to perform time-frequency analysis of non-stationary signals using the STFT. The next section looks at alternative spectrum analysis tools that have been developed and looks at the potential advantages of using these techniques over the commercially accepted STFT.

2.5 Modern spectrum analysis

In the second half of this chapter the emphasis is directed towards alternative spectral estimators to the conventional Fourier transform which have evolved during the last century. The main focus of this research is based on AR spectrum analysis, however, alternative JTFA algorithms are listed here for completeness. These include the Wigner-Ville distribution introduced in 1948 which exhibits excellent joint time-frequency characteristics since it is not restricted to a fixed window like the STFT. The main drawback of the Wigner-Ville distribution is the cross-term interference which manifests itself in the power spectrum. The reduction of cross-term interference has been the focus of much research since this time and adapted algorithms such as the Choi-Williams distribution, cone-shape distribution and signal-dependent time-frequency distribution have been developed to try and enhance the power spectrum. These JTFA algorithms fall into the general category of Cohen's class or bilinear joint time-frequency representations (Qian and Chen, 1996). These time-frequency distributions have been applied to Doppler ultrasound and compared with the performance of the STFT by Fan (1994)[1,2], Smith *et al.* (1994), Guo *et al.*, (1994) and Cardoso *et al.*, (1996). Guo *et al.*, (1994) concluded that both the Choi-Williams distribution and AR modelling can generate good time-frequency distributions of Doppler signals. In the next section the AR process is studied in more detail.

2.5.1 Autoregressive spectral analysis

Fourier techniques assume that the signal is periodic outside the data frame which is a fairly large assumption and in the case of most practical signals is

unrealistic. AR spectral estimators are not so restrictive about the signal beyond the data frame. The only assumption is that the first and second order statistical properties of the signal remain constant outside the data frame (Marple, 1987). The AR system takes a generalised signal at the input *u(n)* and models the signal at the output *x(n)* as illustrated in figure 2.15.

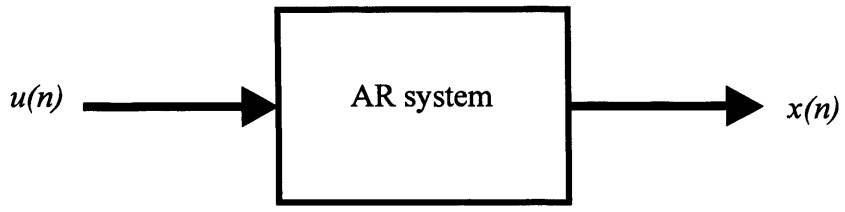

Figure 2.15 Model of the AR system.

The autoregressive model is a specific case of the more general autoregressive moving average model (ARMA). The general equation for ARMA models [2.10] is based on a causal digital filter that attempts to model a time sequence *x(n)* from a white noise input *u(n)*, using a unique finite series of filter coefficients a_k and b_k. The number of coefficients used in the filter is denoted by q and p for the non-recursive and recursive parts respectively.

$$x[n] = \sum_{k=0}^{q} b_k u[n-k] - \sum_{k=1}^{p} a_k x[n-k] \qquad [2.10]$$

The system can be simplified in two ways, either an '*all zero*' model is assumed leading to a moving average (MA) model or the '*all pole*' model is assumed which is called the AR model. The latter is of more interest to signal analysts since the solution of the AR coefficients can be found using a set of linear equations which are computationally easy to solve, whereas the MA and ARMA solutions are non-linear. This has led to much work in the field of AR modelling.

The transfer function *H(z)* for an AR process is obtained by assuming an '*all pole*' model, i.e. making all b_k=0 except for b_0=1. The AR equation can then be simplified, rearranged and z-transformed to produce *H(z)* for the AR filter [2.11].

$$H(z) = \frac{X(z)}{U(z)} = \frac{1}{1 + \sum_{k=1}^{p} a_k z^{-k}} \quad \text{where : } b_k = \begin{cases} 1 & k = 0 \\ 0 & 0 < k \le q \end{cases} \qquad [2.11]$$

The AR power spectral density (PSD) is given by [2.12], where σ^2 is the variance of the driving white noise input to the filter and T is the sampling period.

$$PSD_{AR}(f) = \frac{T\sigma^2}{\left|1 + \sum_{k=1}^{p} a_k \exp(-j2\pi fkT)\right|^2} \qquad [2.12]$$

The PSD can therefore be estimated if σ^2 and the coefficients a_k are known. Two methods for the solution of the filter coefficients are examined: The first is based on the Yule-Walker equations which uses a biased estimate of the autocorrelation sequence (ACS) to find a solution; The second method is Burg's algorithm (Kay and Marple, 1981) which calculates the coefficients directly from the data samples without estimating the ACS.

2.5.1.1 Yule-Walker AR coefficients

The autoregressive coefficients can be calculated using the Levinson-Durbin recursive algorithm [2.13 - 2.18] which provides an efficient solution to the Yule-Walker equations (Kay and Marple, 1981).

$$a_{11} = -R_{xx}(1) / R_{xx}(0) \qquad [2.13]$$

$$\sigma_1^2 = (1 - |a_{11}|^2) R_{xx}(0) \qquad [2.14]$$

where:

$$R_{xx}(k) = \frac{1}{N} \sum_{n=0}^{N-k-1} x(n+k)x(n) \qquad [2.15]$$

Recurse up to the desired model order, p:

$$a_{kk} = -\left[R_{xx}(k) + \sum_{l=1}^{k-1} a_{k-1,l} R_{xx}(k-l) \right] \Big/ \sigma_{k-1}^2 \qquad [2.16]$$

$$a_{ki} = a_{k-1,i} + a_{kk} a_{k-1,k-i}^* \qquad [2.17]$$

$$\sigma_k^2 = \left(1 - |a_{kk}|^2\right) \sigma_{k-1}^2 \qquad [2.18]$$

2.5.1.2 Burg's AR coefficients

To compare the results obtained using the Yule-Walker coefficients it was decided to calculate Burg's coefficients as well. The determination of Burg's coefficients is evolved directly from the data frame and does not involve estimating the ACS (Kay and Marple, 1981). This provides an interesting comparison with the results obtained using the Yule-Walker algorithm and it is a good error checking method to make sure that both algorithms are working correctly.

2.5.1.3 Model order

There is an element of choice when implementing the AR algorithms since the choice of model order (p) is not fixed. The selection of model order is important and determines how accurately the system can model the output sequence $x(n)$ (Akaike, 1974). If the choice of p is too low, then the system will not have enough poles to adequately model the signal. If the model order is too high then the system will be too biased towards the particular realisation of the process that it is modelling.

2.5.1.4 Estimating the power spectral density function

The coefficients can be used in two ways to extract the PSD of the signal: either the PSD can be calculated directly from the transfer function [2.12] or alternatively the PSD can be determined from the ACS [2.19]. It can be shown

(Marple, 1987) that in the limit the PSD of the ACS is exactly equal to the PSD determined from the transfer function of the AR system.

$$PSD_{AR}(f) = T\sum_{n=-\infty}^{\infty} R_{xx}(n)\exp(-j2\pi nkT) \qquad [2.19]$$

The problem of limited resolution as encountered with the FFT can be overcome by using the AR PSD [2.12] or [2.19]. In order to compute the AR PSD using [2.19] it is necessary to extrapolate the ACS beyond the p values obtained using the data sequence. This can be achieved using the principle of linear prediction [2.20]. It is not necessary to extrapolate the ACS to infinity since the ACS tends to zero and therefore the AR PSD can be approximated using a finite number of lags [2.21]. The approximated AR PSD [2.21] can be found by performing a fast Fourier transform on the extrapolated ACS which is computationally faster than evaluating [2.12].

$$\hat{R}_{xx}(n) = -\sum_{k=1}^{p} a_k\, R_{xx}(n-k) \quad \text{for: } |n| > p \qquad [2.20]$$

$$\hat{PSD}_{AR}(f) = T\sum_{n=-M}^{M-1} \hat{R}_{xx}(n)\exp(-j2\pi nkT) \qquad [2.21]$$

2.5.2 Autoregressive modelling of a non-stationary signal

The non-stationary signal analysed using the STFT is analysed again here and the results are compared with those obtained using the STFT. The signal was analysed using data frames comprising 256, 128, 64 and 32 data-points. The data frames were processed using the Yule-Walker AR algorithm: The recursive Levinson-Durbin equations (p = 20) were used to calculate the AR coefficients based on a biased estimate of the ACS. The AR coefficients were used to extrapolate the ACS to 128 points and the frequency spectrum was estimated using the FFT of the mirrored ACS. As for the STFT the signal was analysed by sliding the window over the signal by one sample each time, the results are illustrated in figure 2.16.

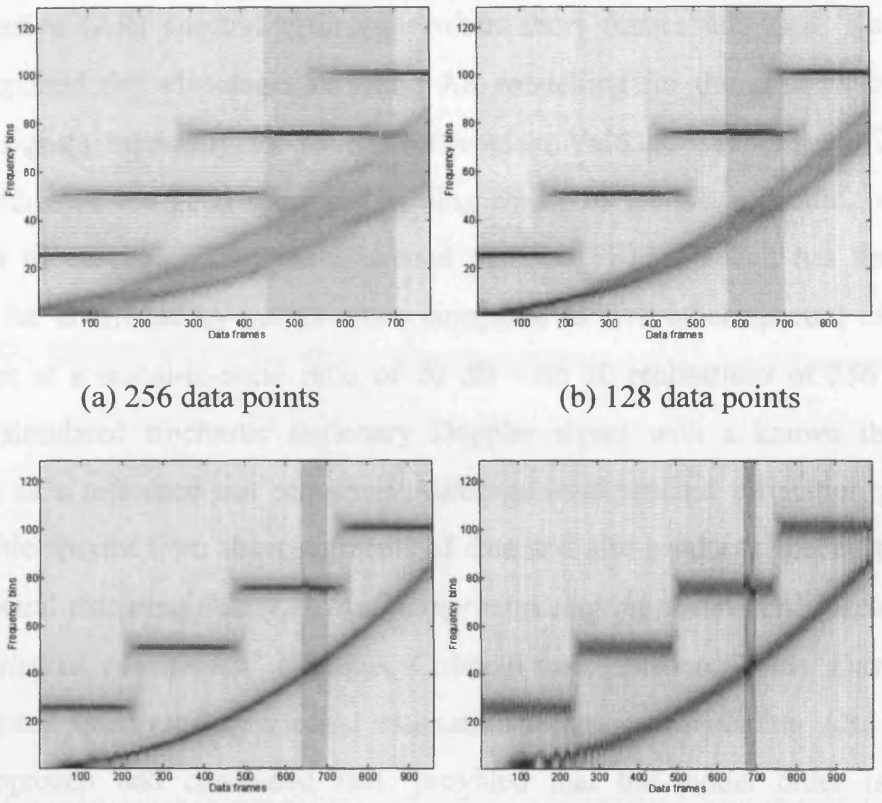

(a) 256 data points (b) 128 data points

(c) 64 data points (d) 32 data points

Figure 2.16 AR model order 20, PSD obtained using [2.21] with $M = 128$. Window moved by 1 sample each time.

In figure 2.16 it can be seen that in all cases the frequency resolution is better than the resolution provided by the STFT. The temporal resolution of the AR is not significantly better than the STFT, but it is clear that the AR algorithm is able to produce better results for the short data frame.

2.6 Comparison of FFT and AR properties

It has been suggested that in some situations a parametric approach such as autoregressive modelling (AR) may be better than the FFT approach for analysing Doppler signals. Marple (1977) demonstrated the better spectral resolution of autoregressive (AR) and autoregressive moving average (ARMA) models compared to FFT by separating two sinusoids of similar frequency. Kitney and Giddens (1986) stressed the better performance on spectral tracking and spectral resolution of

autoregressive (AR) spectral estimation when short frames are used. Kaluzynski (1987) reported the advantages of using AR modelling for the analysis of pulsed Doppler signals especially for short data lengths. Vaitkus, Cobbold and Johnston (1988) addressed the good spectral matching ability of the AR modelling approach compared to the FFT. They demonstrated that the FFT approach has the largest variance for all frequency values when compared to five other spectral estimation techniques at a signal-to-noise ratio of 20 dB with 20 realisations of 256 samples using a simulated stochastic stationary Doppler signal with a known theoretical spectrum as a reference test sequence. Autoregressive spectral estimation produces more stable spectra from short segments of data and also produces spectra that have good spectral matching ability, "*closely approximating the theoretical spectrum with good statistical consistency*" (Vaitkus, Cobbold and Johnston, 1988). David *et al.* (1991) tested three modern spectral estimation techniques, including AR, with the STFT approach and concluded that, provided that the model order is chosen appropriately, the modern techniques are superior to the traditional FFT-based approach.

The algorithms for performing AR are computationally more demanding than the FFT and real-time systems demand more computational power than that of a standard personal computer. Modern DSP boards combined with standard microcomputers allow a flexible software orientated approach to the implementation of real-time algorithms for analysing a wide range of signals. It has therefore been possible to implement more sophisticated 'modern' spectrum analysis tools over recent years for real-time applications such as Doppler ultrasound. In 1989 a real-time system based on AR for analysing Doppler ultrasound signals was described (Schlindwein and Evans, 1989).

2.7 Joint time-frequency analysis of the Doppler ultrasound signal

The techniques used so far have been applied to signals with known theoretical spectra. The conclusion of this chapter compares Fourier and autoregressive time-frequency analysis of the Doppler signal for a normal femoral artery.

A single heartbeat was analysed to illustrate the trade-off between time and frequency resolution. The sampling frequency of the signal is 10.24 kHz. The signal was analysed using a data frame containing 256, 128, 64 and 32 data samples. The signal was analysed using: a 256-point FFT with appropriate zero padding and a Hamming window; the AR process was implemented using the Levinson-Durbin equations with $p = 12$. The frequency spectrum was determined from the extrapolated (128-point) and mirrored ACS. The signal was analysed using a sliding-window with overlapping data frames. The window was moved along by 32 samples and then reapplied. The results are illustrated in figures 2.17 and 2.18 for the FFT and AR spectral estimators respectively.

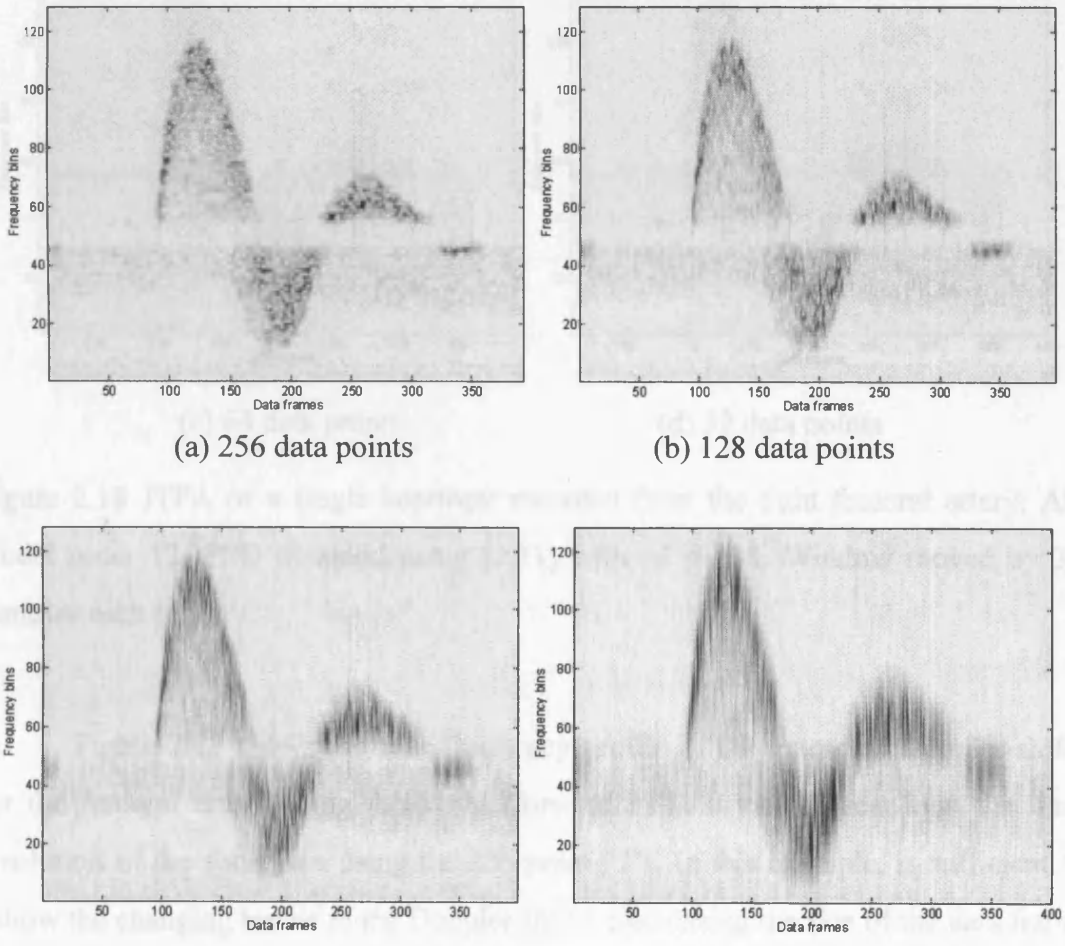

(a) 256 data points (b) 128 data points

(c) 64 data points (d) 32 data points

Figure 2.17 JTFA of a heartbeat recorded from the femoral artery; sliding 256-point FFT with Hamming window and zero padding. Window moved by 32 samples each time.

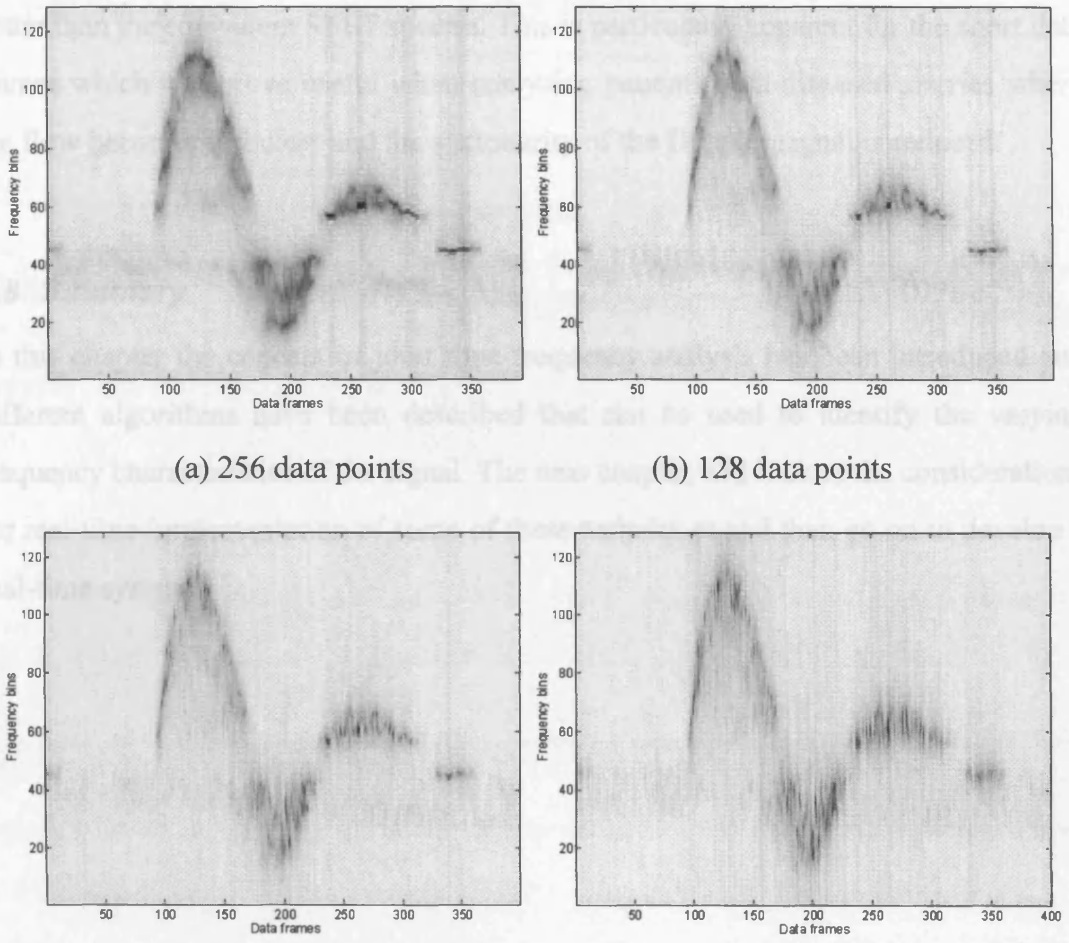

(a) 256 data points (b) 128 data points

(c) 64 data points (d) 32 data points

Figure 2.18 JTFA of a single heartbeat recorded from the right femoral artery; AR model order 12, PSD obtained using [2.21] with M = 128. Window moved by 32 samples each time.

Figure 2.17 shows the time-frequency profile of the extracted Doppler shifts for the femoral artery using the straightforward FFT. It can be seen that the time resolution of the sonogram using the 256-point FFT, in this example, is sufficient to follow the changing trends in the Doppler shifts. Decreasing the size of the data frame further improves the time resolution but the frequency resolution deteriorates. In the case of the 32-point FFT the reduction in frequency resolution almost causes the Doppler shifts to spread over the entire frequency spectrum. The spectra obtained using the AR algorithms, illustrated in figure 2.18, are similar to those obtained using

the STFT, with the exception that the frequency resolution of the AR spectra is visibly better than the equivalent STFT spectra. This is particularly apparent for the short data frames which will prove useful when analysing patients with diseased arteries where the flow becomes turbulent and the stationarity of the Doppler signal is reduced.

2.8 Summary

In this chapter the concept of joint time-frequency analysis has been introduced and different algorithms have been described that can be used to identify the varying frequency characteristics of the signal. The next chapter will look at the considerations for real-time implementation of some of these techniques and then go on to develop a real-time system.

3. Real-time Analysis

3.1 Introduction

In this chapter the merits of real-time processing are introduced. This is followed by a discussion about the requirements of a real-time Doppler ultrasound system. The remainder of the chapter details the real-time system that has been implemented in this research.

3.2 Real-time processing

Real-time systems provide a convenient cost-effective tool for performing clinical examinations. Off-line processing creates economical and logistical problems: Additional man-hours are required to collect the data and then process the information following the examination. Storage of raw data introduces excessive disk space requirements which are expensive. Processed data normally requires much less storage space than the original signal. In addition, if a problem occurred during the examination that was not detected whilst recording and was only apparent during the post-processing of the data it would be necessary to repeat the examination. It is important that the ultrasonographer is able to confidently perform the examination to minimise costs and to avoid unnecessary inconvenience and stress to the patient. The recordings can be affected by the patients movements or breathing. Depending upon the insonation of the vessel the recording may detect signals from neighbouring vessels. Examinations are facilitated nowadays by Duplex scanning which incorporates both a real-time B-mode imaging system and Doppler ultrasound system.

3.3 Real-time Doppler ultrasound using DSP32C

The initial part of the research was to set-up a system for real-time processing of Doppler ultrasound signals. The system was designed to compare the performance of modern signal processing techniques with the conventionally used fast Fourier transform (FFT). The limitations of the FFT have been underlined in the previous chapter and there is therefore a case for implementing alternative signal processing

tools which are computationally more demanding but may offer improved resolution and spectral matching capabilities. With advances in DSP technology it is now feasible to implement these algorithms in real-time relatively cheaply. The system was designed using a PC-486 DX2 66 MHz combined with an AT&T DSP32C digital signal processing chip mounted on a Loughborough Sound Images (LSI) board (LSI, 1991).

3.4 Requirements for a real-time Doppler ultrasound system

When designing a real-time system, the constraints that are put upon it are created by the physiological and practical conditions that the system will have to cope with and the output that is expected by the user.

3.4.1 Physiological conditions

The system should be able to track the velocities of the population of targets within the vessel as a function of time. Accurate extraction of Doppler shifts using current signal processing algorithms is extremely difficult since Doppler ultrasound signals are non-stationary, this is due to the pulsatile flow through the vessel. As was discussed in the previous chapter, to fully describe a non-stationary process it is necessary to analyse time frames that are shorter than the length of time over which the signal can be considered stationary. Typically blood flow can be considered stationary for short periods of time up to approximately 20 ms (Evans *et al.*, 1989[1]; Nichols and O'Rourke, 1990). The actual period of time depends upon the type of flow present in the vessel and may well be shorter than this if the flow is disturbed as a result of a stenosed vessel. If longer time frames are used then the sonogram will be smeared and the consecutive frames will not provide a detailed indication of how the velocities within the vessel are changing with respect to time. The minimum frame rate therefore needs to be 50 frames per second (fps). The restriction on the length of the data frame imposes a limit on the number of data points that can be used to estimate the frequency spectrum. Due to the statistical instability of the spectral estimator when analysing short data frames it is advised that several spectra should be

averaged together (Welch, 1967). The overall length of time used to produce an average spectrum should comply with the stationarity of the signal. Averaging data frames is dependent upon the processing power of the signal processor.

The conditions under which a Doppler ultrasound scan is performed dictate the maximum likely Doppler shift frequencies that are going to be present in the spectrum, this is governed by the Doppler equation. By substituting maximum expected values into this equation and taking the maximum cosine of the angle to be 1 (not possible in practice) the maximum Doppler shift is given by:

$$\Delta f_{D\max} \approx \frac{2 \cdot v_{\max} \cdot f_{\max} \cdot 1}{1580} \approx 21 \text{ kHz} \qquad [3.1]$$

where: v_{max} = 2 ms^{-1} (blood velocity), f_{max} = 8 MHz (ultrasound beam frequency). The value calculated in [3.1] varies dramatically depending on the conditions under which the scan is performed, the vessel being studied and the degree of stenosis. 21 kHz is seen to be an extreme case and the maximum Doppler frequency may well be substantially smaller than this value. The frequencies of the demodulated Doppler shifts typically lie within the audio range < 20 kHz and can be listened to using headphones. As a result of this broad frequency range a clinical system should incorporate a variable sampling frequency (f_{samp}). For frequencies up to 21 kHz the sampling rate should be at least 42 kHz (i.e. twice the maximum Doppler frequency) to avoid aliasing. In order to facilitate interpretation of the sonogram it is convenient for the sonogram to be displayed with a fixed frame rate regardless of f_{samp}.

3.4.2 Practical constraints

In order to produce a useful real-time system it is necessary to analyse all the data that is passed to the DSP board. If the physiological requirements cannot be met with the current technology then the system is not going to yield satisfactory results.

For real-time operation the most important criterion is that the system is able to process the data and display the results at the same rate as the data is being collected.

Based on the conditions laid out above, the next section describes the implementation of each stage of the real-time system. The system is divided into three sections pre-processing, processing and display.

3.4.3 Data

The data was stored on digital audio tapes (DAT) tapes with the forward and reverse signals on the left and right channels respectively.

3.5 Pre-processing - Heterodyne mixer

The heterodyne unit modulates the forward signal component, $f(t)$, and the reverse signal component, $r(t)$, onto a central adjustable heterodyne carrier frequency (f_c). The options for the heterodyne frequency are 500 Hz and 1, 2, 4 and 8 kHz which can be selected using a switch on the unit. There is also the facility to adjust the gain to optimise the dynamic range of the output. The following section gives a graphical representation of the combination of the separate forward and reverse flow components into a single signal $S(t)$. Figure 3.1 shows the frequency spectrum of the forward and reverse signals.

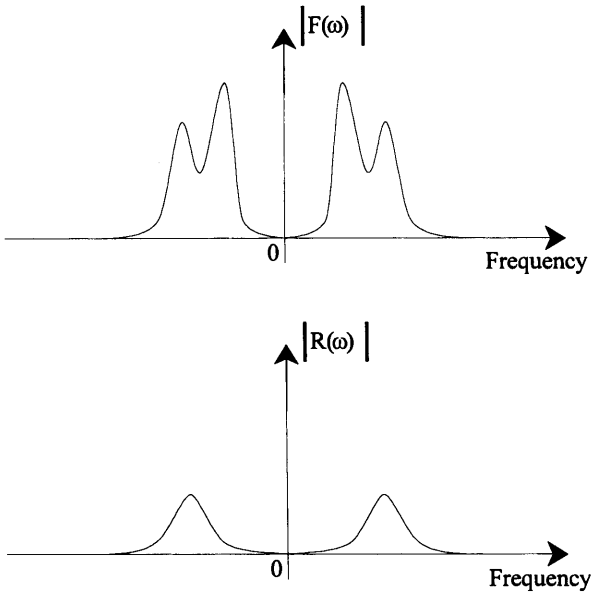

Figure 3.1 Representative frequency spectra of the forward and reverse signals.

The forward and reverse signals are modulated by $\omega_c = (2\pi f_c)$ as shown in figure 3.2. This generates signals with the forward and reverse flow modulated into the upper sideband (USB) and lower sideband (LSB) of each signal.

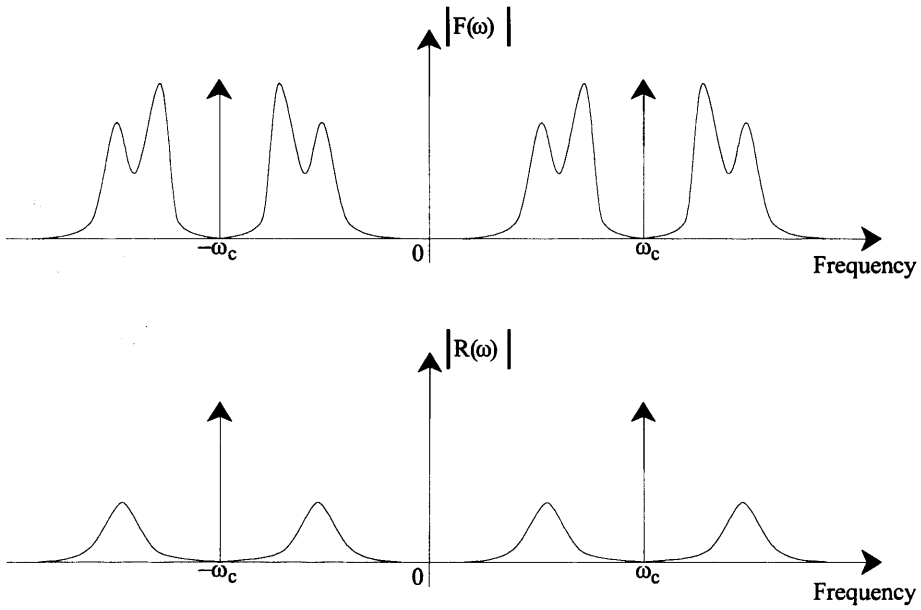

Figure 3.2 Forward and reverse components modulated by the carrier frequency ω_c.

The next stage is to remove the LSB of the forward signal and the USB of the reverse signal by high pass and low pass filtering respectively. The central carrier

frequency is suppressed by means of careful selection of the cut-off point of the filters during filtering. The filtered signals are shown in figure 3.3.

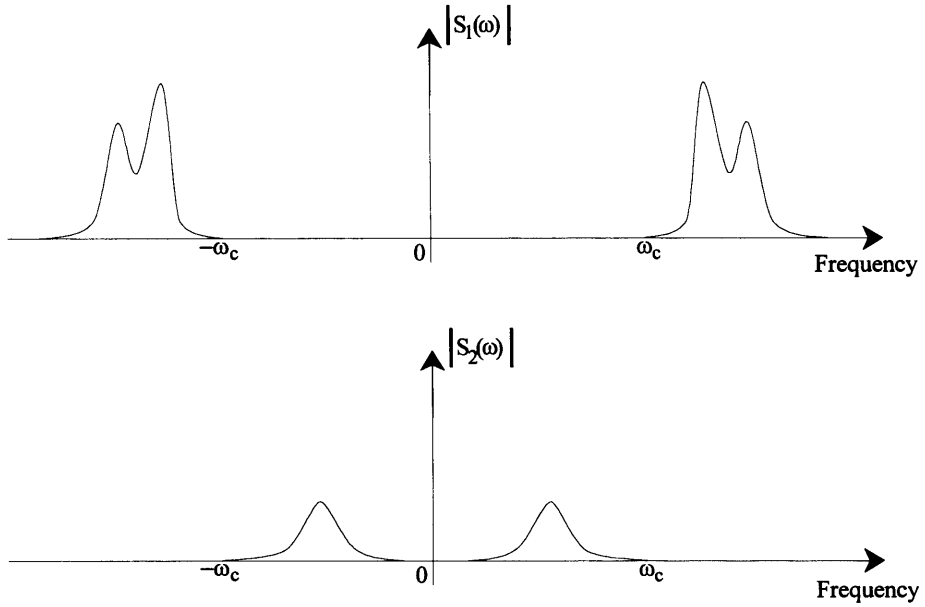

Figure 3.3 Filtered forward and reverse components with carrier frequency suppressed.

Following the filtering of the signals the two signals can be added together to generate a single signal as shown in figure 3.4. The signal $S(t)$ contains the forward wave (ω_f) in the USB of the carrier frequency and the reverse wave (ω_r) in the LSB of the carrier frequency [3.2].

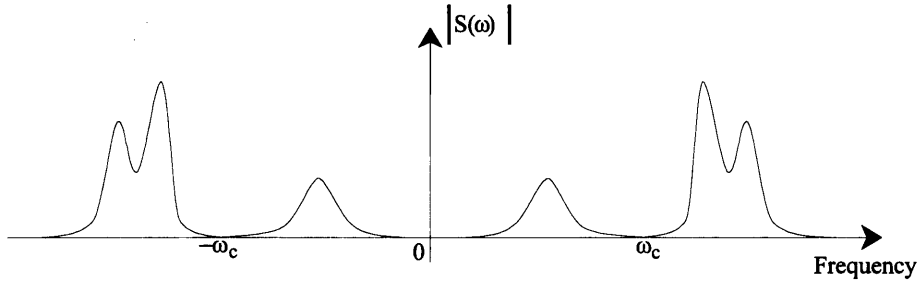

Figure 3.4 Graphical combination of the forward and reverse flow components about a central heterodyne carrier frequency.

$$S(t) = A_f \cos\left(\omega_c t + \omega_f t + \phi_f\right) + A_r \cos\left(\omega_c t + \omega_r t + \phi_r\right) \qquad [3.2]$$

The value of the heterodyne frequency is passed directly to the parallel printer port of the computer by means of a 3-bit code. The details of this heterodyne unit are further discussed in a separate technical note (Schlindwein *et al.*, 1996).

3.6 Processing - Data acquisition & signal processing using the DSP32C

3.6.1 Variable sampling frequency

The combined analogue signal is digitised and processed using the DSP board. The signal is sampled using a 16-bit analogue-to-digital converter (ADC) with a programmable sampling frequency defined by the user. This enables the operator to optimise the range of the spectrum analyser to be adequate to display the maximum Doppler shifts in the signal. The sampling frequency of the ADC can be altered for 5.12, 10.24, 20.48 or 40.96 kHz operation.

Since there is a variable sample rate then the data frame rate is going to vary also. One of the interesting qualities for the real-time system is to have a standardised output with respect to time. A frame rate of 80 fps was chosen which corresponds to a data frame length of 12.5 ms. There are two approaches for producing a standardised output with respect to time that is independent of sampling frequency and displays Doppler spectra at fixed time intervals: The first is the 'sliding' FFT, which uses the latest N samples, but updates the frames at every M (M<N) samples, while discarding the M oldest samples and including the newest M samples to the current frame, producing a sort of 'moving-average' of the spectra. The alternative is to have no overlap between data frames and to zero pad the short frames up to a standard size when a low sampling frequency is used. The latter approach has been used here because it produces spectra which use not more than 12.5 ms of data and because stationarity was a major concern.

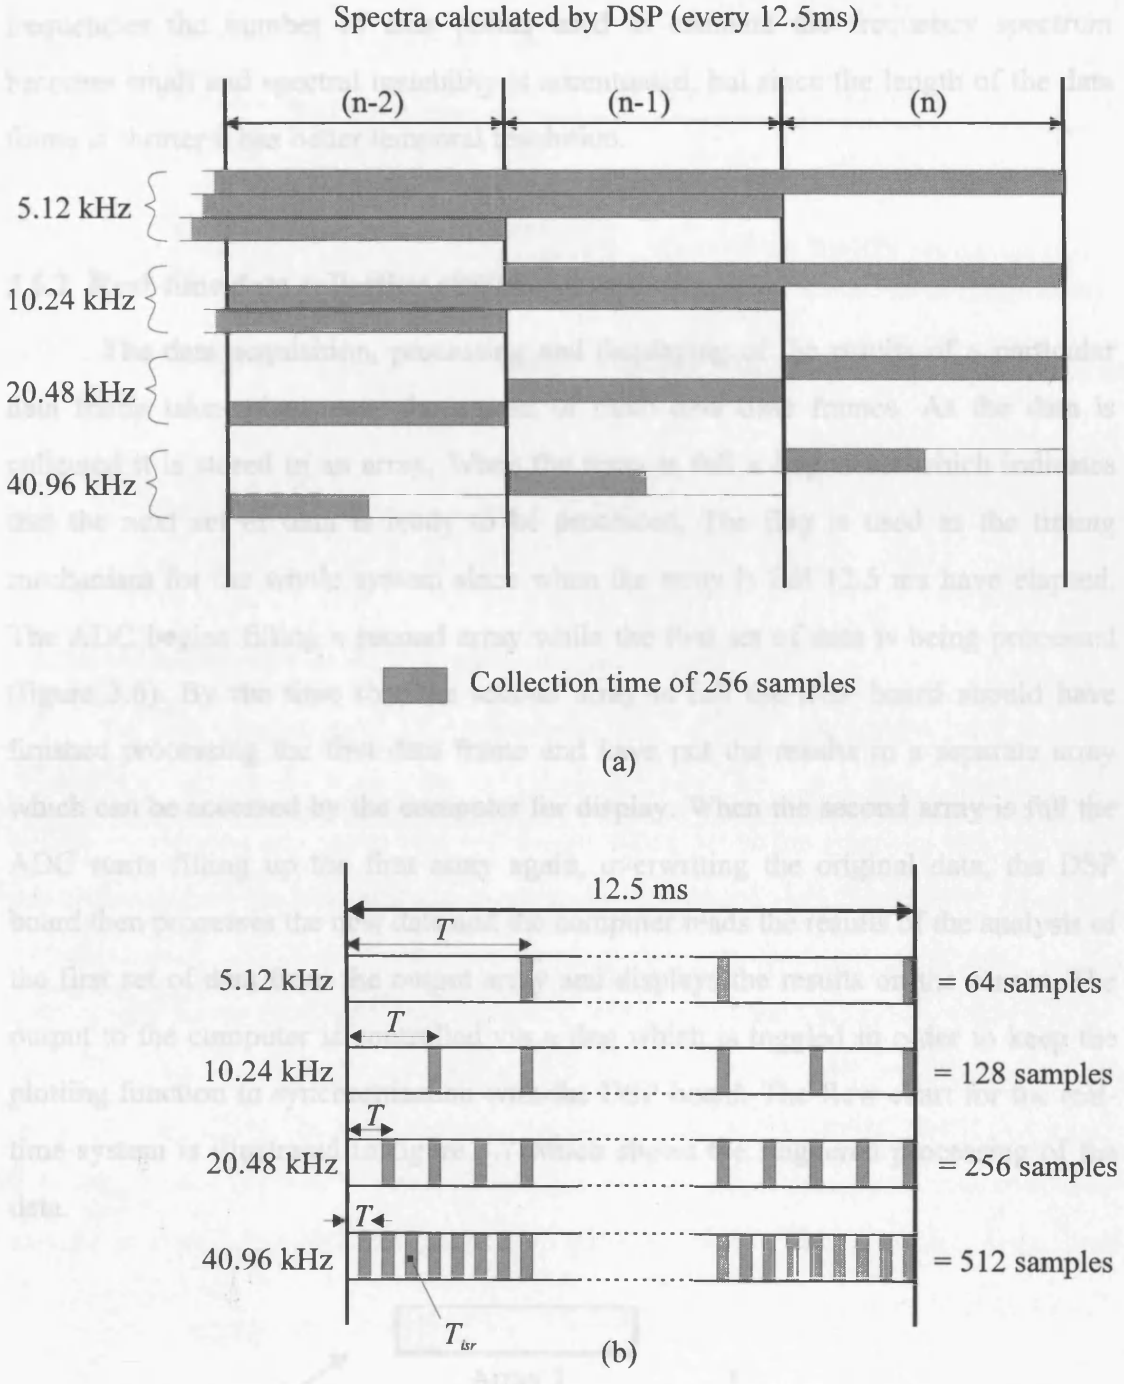

Figure 3.5 Comparison of the two approaches used to produce a standard frame rate at different sampling frequencies (a) Overlapping frames ⇒ number of samples = 256 (b) Non-overlapping frames ⇒ variable number of data samples.

Figure 3.5 shows the comparison between the two methods, the first method has the advantage that the spectral estimate is based on a reasonable size data frame for all sampling frequencies. In the case of the second method at low sampling

frequencies the number of data points used to estimate the frequency spectrum becomes small and spectral instability is accentuated, but since the length of the data frame is shorter it has better temporal resolution.

3.6.2 Real-time data collection and processing

The data acquisition, processing and displaying of the results of a particular data frame takes place over the course of three data time frames. As the data is collected it is stored in an array. When the array is full a flag is set which indicates that the next set of data is ready to be processed. The flag is used as the timing mechanism for the whole system since when the array is full 12.5 ms have elapsed. The ADC begins filling a second array while the first set of data is being processed (figure 3.6). By the time that the second array is full the DSP board should have finished processing the first data frame and have put the results in a separate array which can be accessed by the computer for display. When the second array is full the ADC starts filling up the first array again, overwriting the original data, the DSP board then processes the new data and the computer reads the results of the analysis of the first set of data from the output array and displays the results on the screen. The output to the computer is controlled via a flag which is toggled in order to keep the plotting function in synchronisation with the DSP board. The flow chart for the real-time system is illustrated in figure 3.7 which shows the staggered processing of the data.

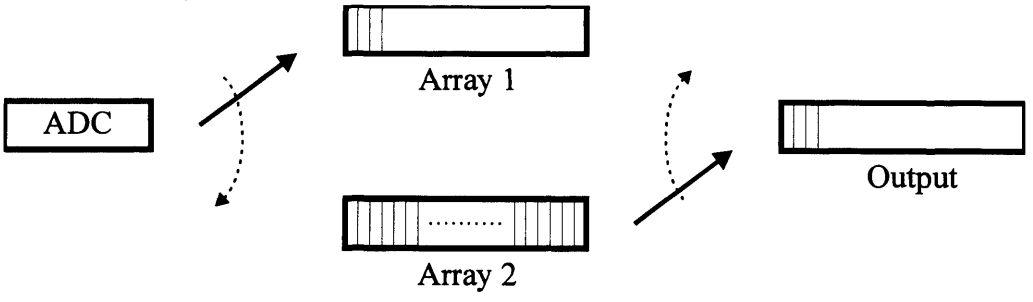

Figure 3.6 Simultaneous data acquisition and processing.

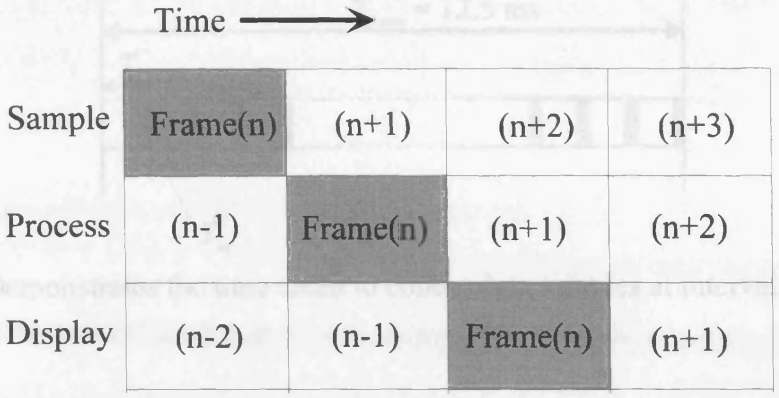

Figure 3.7 Flow chart of real-time system.

3.6.2.1 Simultaneous data acquisition and processing

As was mentioned in the previous section the DSP board must have finished processing a data frame before the next data frame becomes available otherwise data will be overwritten and the results will be corrupted. Since there is only one processor the processing of the data and the acquisition of the next data frame must share the processor time. The total time for collecting the new data and processing of the current data frame should be less than time taken to collect a single data frame (12.5 ms). The data is collected via an interrupt service routine (ISR) which stores new samples in a buffer. The ISR interrupts the processing of the current data frame every $T = 1/f_{samp}$ seconds (see figure 3.8) and jumps to a subroutine which reads the next sample in from the ADC and stores it in the input array. The total processing time T_{total} is given by:

$$T_{total} = T_{proc} + NT_{isr} \qquad [3.3]$$

where: T_{proc} = processing time for frequency analysis of a data frame, $N = f_{samp}T_{frame}$, T_{isr} = acquisition time of one sample. Therefore the total time available for processing is when $T_{total} = T_{frame} = 12.5$ ms.

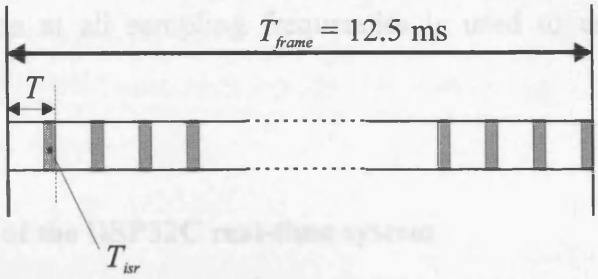

Figure 3.8 Demonstrates the time taken to collect data samples at intervals of $1/f_{samp}$.

3.6.2.2 Data processing

The system was designed with the choice of two signal processing algorithms, the first was the FFT which used zero padding to produce a data frame of 256 samples (see Table 3.1) and a Hanning window to reduce spectral leakage. The autoregressive modelling algorithm used the efficient Levinson-Durbin equations to evaluate the autoregressive coefficients. The autoregressive coefficients were used to extrapolate the autocorrelation series (ACS) to 128 points and then the function was mirrored around zero to produce an even sequence of 256 lags. An FFT was then used to estimate the frequency spectrum from the extrapolated ACS. Therefore, both algorithms have the same resolution for a particular sampling frequency.

Table 3.1 - Number of samples collected within a 12.5 ms data frame. Degree of zero padding required to obtain frames of 256 data points and produce spectra with 128 frequency bins.

Sampling frequency f_{samp} (Hz)	Number of data points	Zero padding
40960	512	-
20480	256	0
10240	128	128
5120	64	192

For the FFT algorithm when f_{samp} is 5.12, 10.24 or 20.48 kHz all the data is used to estimate the frequency spectrum. When f_{samp} is 40.96 kHz, it is possible to collect 512 new data samples within 12.5 ms. For the FFT algorithm a 256-point FFT is used to estimate the frequency spectra and therefore only the first half of the data frame is used, the remaining samples are discarded. In the case of the autoregressive

algorithm all the data at all sampling frequencies is used to estimate the biased estimate of the ACS.

3.6.3 Development of the DSP32C real-time system

The LSI software package includes a 'C' compiler for development of programs in a high-level language. The compiler can be used to generate executable programs that can be downloaded and run on the DSP board. This has the advantage of reduced development time. The alternative is to implement the algorithms directly in DSP32C assembly language which is more complex. The assembly language routines have the advantage of speed and efficiency over the 'C' compiler, which is an obvious advantage when considering real-time operation. The 'C' compiler has the advantage of being a high-level language and therefore it is easier to write and adapt routines without too much effort. The next section looks at the relative efficiency of the two software development techniques and compares the sacrifice paid for in code efficiency for improved speed in the development phase.

3.6.3.1 'DSP32C assembly' real-time FFT system vs. 'C' real-time FFT system

The performance of the program developed in 'C' was compared with a modified version of a spectrum analyser program written in DSP32C assembly language that was supplied with the DSP board by LSI. The program was modified to perform joint time-frequency analysis (JTFA) with the same time-frequency operation as the real-time system being developed in 'C'. Both systems used data frames of 12.5 ms with a Hanning anti-leakage window and zero padding up to a 256-point data frame. The frequency spectra were estimated using a 256-point FFT. The processing times for both systems were examined at each of the sampling frequencies. The execution times for each program are compared to contrast the speed of the two systems.

The DSP32C has a clock frequency of 50 MHz[‡] , with an instruction cycle time of 80 ns[‡]. Within each instruction the DSP32C can perform 4 operations: a single read operation; two operand fetches and a single write to memory. Due to the pipelined architecture of the DSP32C the fetch-multiply-accumulate-write procedure is not performed in a single instruction but is in fact sequentially processed over 4 instruction cycles. As a result in some cases wait states need to be inserted into the software in order for the appropriate memory address or register to have been updated, this effect is known as latency.

For a 12.5 ms data frame the DSP32C is able to perform 1.56E5 instructions (12.5E-3 ÷ 80E-9). The library of functions available with the DSP32C provides optimised algorithms for performing an FFT. A single 256-point FFT (_rffta) requires 5344[†] instructions and due to the latency effects 1612[†] wait states. The computational time required to perform the FFT can be calculated using:

$$T_{FFT} = IT_{in} + W/f_{clock} \qquad [3.4]$$

where: T_{FFT} = Total time to perform FFT, I = No. of instructions, T_{in} = Time to perform a single instruction (80 ns), W = No. of wait states, f_{clock} = DSP32C clock frequency (50 MHz). The total time required to perform the FFT is 0.460 ms which utilises approximately 3.7 % of the total time available. The total processing time is the accumulation of the time taken to acquire the next frame of the data via the ISR, process the current data frame and store the results in the output array. Hence, the overall processing time is between 6 % and 12 % of the total available time depending upon the sampling frequency as can be seen from figure 3.9. All the results in the next section are expressed as a percentage of this overall number of available operations.

[‡] Figures taken from 'Information Manual' for the AT&T WE® DSP32C Digital Signal Processor.
[†] Figures taken from 'Reference Manual' for the AT&T WE® DSP32 and DSP32C - Application Software Library.

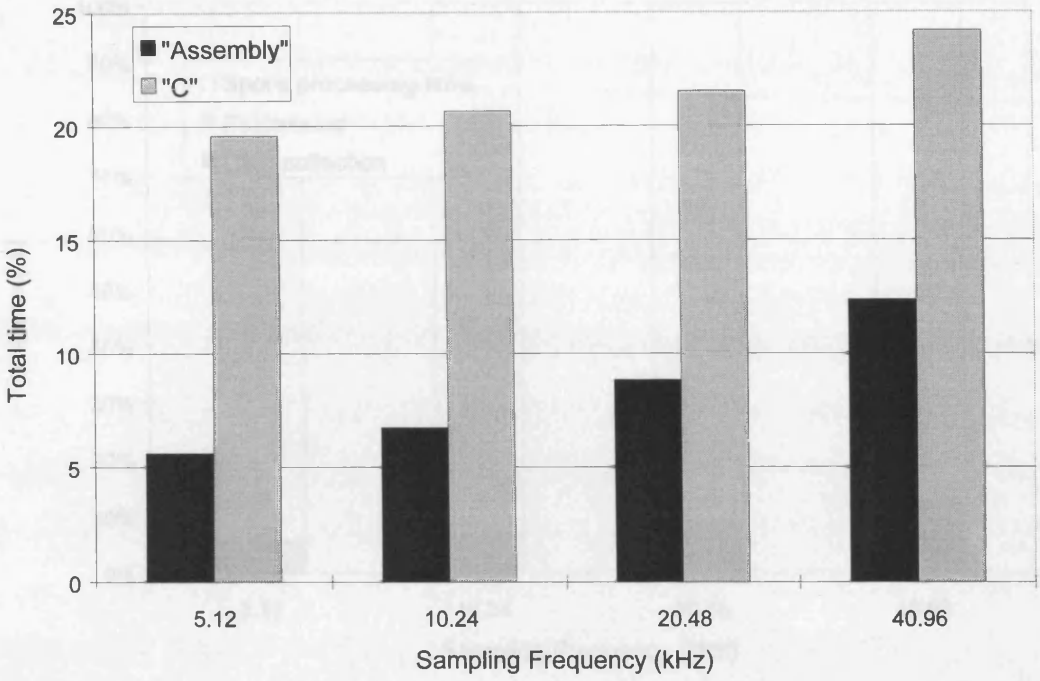

Figure 3.9 Comparison of the overall performance at different sampling frequencies of the FFT spectrum analysers developed in 'DSP32C assembly' and 'C'.

Figure 3.9 compares the overall performance of the two systems for the processing of the data at each sampling frequency. It can be seen that the performance of the system developed in 'assembly' is much better than the equivalent system developed using the 'C' compiler. For a sampling frequency of 5.12 kHz the 'assembly' system is approximately 3.5 times more efficient than the 'C' system at the same sampling frequency. Despite the improvement in performance obtained by using assembly over 'C' the programs developed in 'C' still only use less than 25 % of the available processing time at the highest desired sampling frequency (40.96 kHz). Therefore, there is scope to develop more sophisticated algorithms using 'C' and still maintain real-time operation.

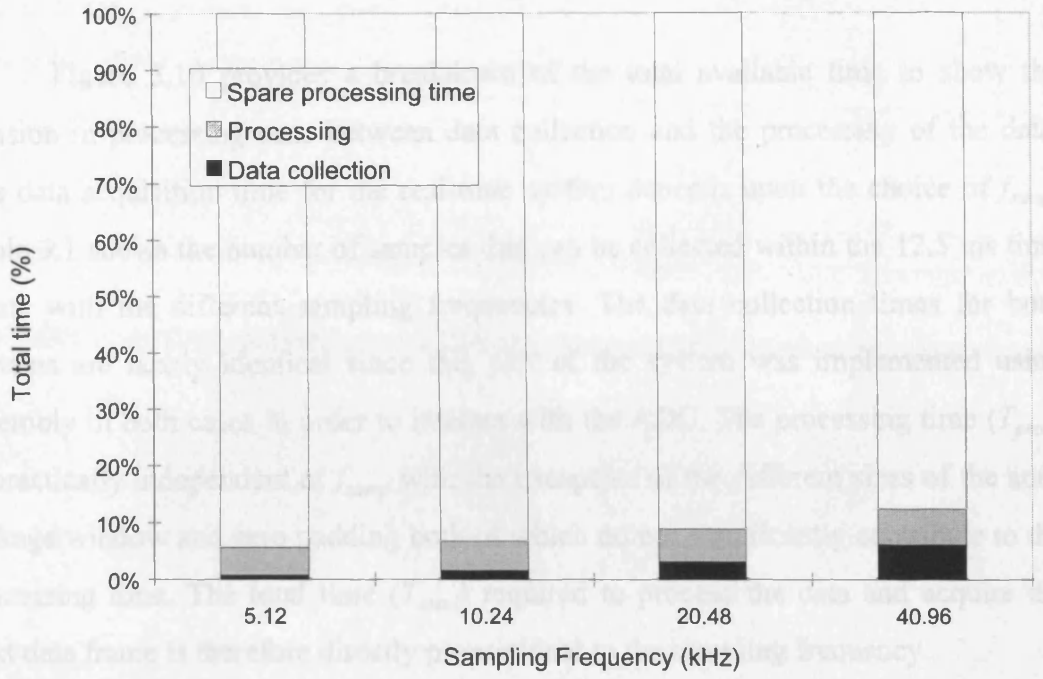

(a)

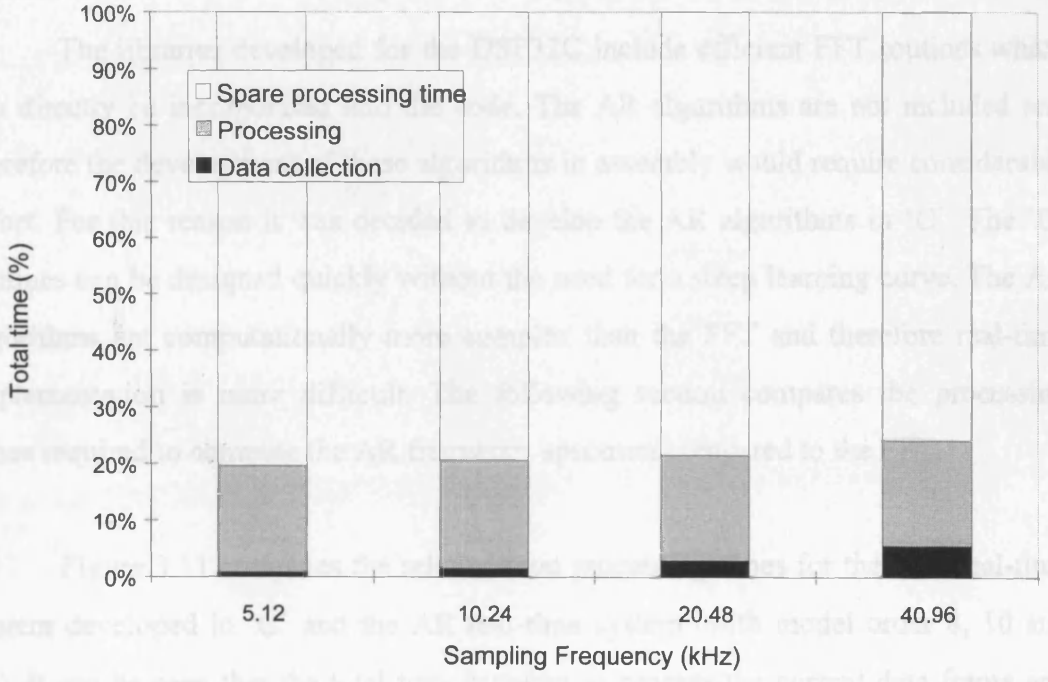

(b)

Figure 3.10 Breakdown of overall processing times into data acquisition times and processing times at different sampling frequencies for (a) 'DSP32C assembly' (b) 'C' FFT spectrum analysers.

Figure 3.10 provides a breakdown of the total available time to show the division in processing time between data collection and the processing of the data. The data acquisition time for the real-time system depends upon the choice of f_{samp}. Table 3.1 shows the number of samples that can be collected within the 12.5 ms time frame with the different sampling frequencies. The data collection times for both systems are nearly identical since this part of the system was implemented using assembly in both cases in order to interact with the ADC. The processing time (T_{proc}) is practically independent of f_{samp} with the exception of the different sizes of the anti-leakage window and zero padding both of which do not significantly contribute to the processing time. The total time (T_{total}) required to process the data and acquire the next data frame is therefore directly proportional to the sampling frequency.

3.6.3.2 Real-time FFT vs. AR spectrum analyser

The libraries developed for the DSP32C include efficient FFT routines which can directly be incorporated into the code. The AR algorithms are not included and therefore the development of these algorithms in assembly would require considerable effort. For this reason it was decided to develop the AR algorithms in 'C'. The 'C' routines can be designed quickly without the need for a steep learning curve. The AR algorithms are computationally more complex than the FFT and therefore real-time implementation is more difficult. The following section compares the processing times required to compute the AR frequency spectrum compared to the FFT.

Figure 3.11 compares the relative total processing times for the FFT real-time system developed in 'C' and the AR real-time system (with model order 8, 10 and 12). It can be seen that the total time required to process the current data frame and collect the next data frame is much greater for AR than the FFT. The processing time is also dependent upon the model order of the AR system. For 5.12, 10.24, 20.48 kHz it is possible to perform real-time analysis using the AR algorithms. At the maximum sampling frequency (40.96 kHz) real-time operation is compromised if a model order of 10 or 12 is used.

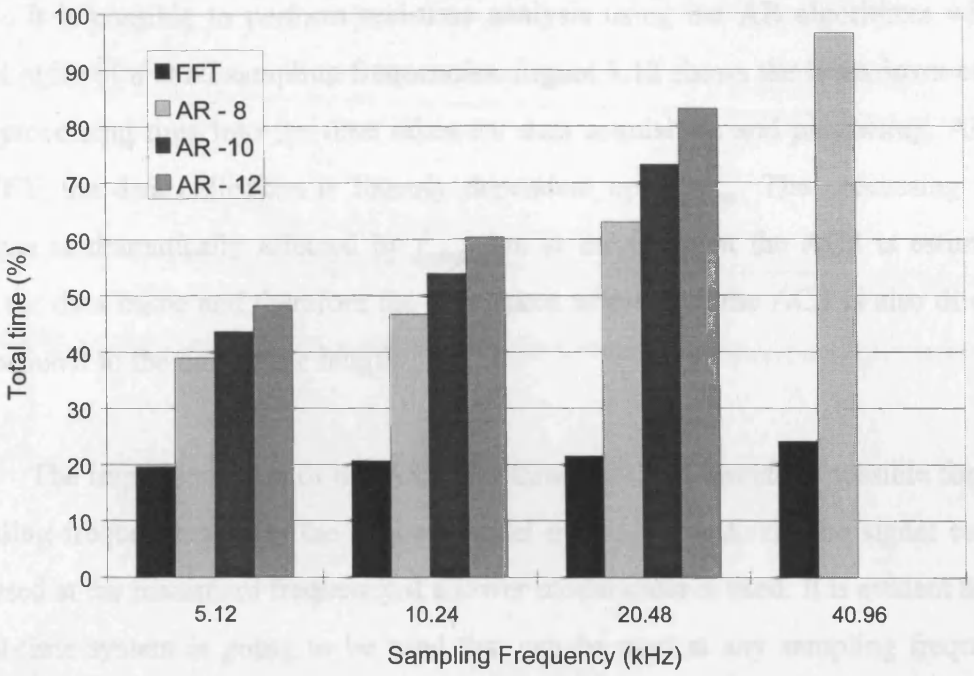

Figure 3.11 Comparison of overall processing times at different sampling frequencies for the FFT and AR spectrum analysers developed in 'C'.

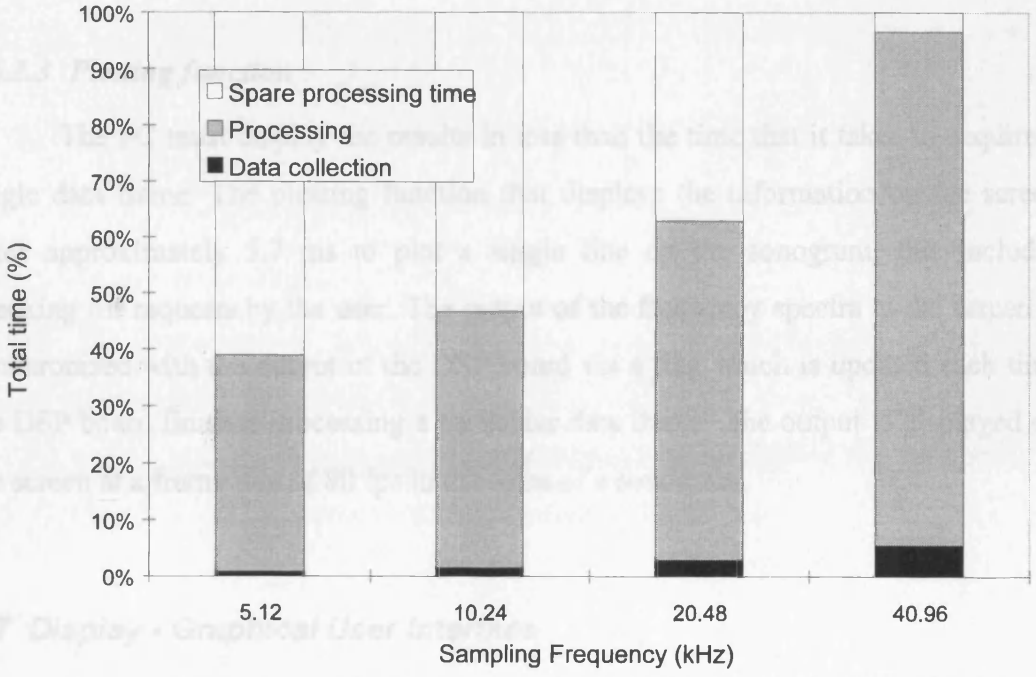

Figure 3.12 Breakdown of the overall processing times into data acquisition times and processing times at different sampling frequencies for the AR spectrum analyser (with model order = 8).

It is possible to perform real-time analysis using the AR algorithms with a model order of 8 at all sampling frequencies. Figure 3.12 shows the breakdown of the total processing time into the time taken for data acquisition and processing. As for the FFT, the data collection is linearly dependent upon f_{samp}. The processing time however is dramatically affected by f_{samp} due to the fact that the ACS is estimated from the data frame and therefore the time taken to estimate the ACS is also directly proportional to the data frame length.

The implementation of the AR algorithms in 'C' is therefore possible for low sampling frequencies up to the desired model order. Alternatively the signal can be analysed at the maximum frequency if a lower model order is used. It is evident that if a real-time system is going to be used that can be used at any sampling frequency without compromising the choice of model order, then it is necessary to develop some of the routines in assembly.

3.6.3.3 Plotting function

The PC must display the results in less than the time that it takes to acquire a single data frame. The plotting function that displays the information on the screen takes approximately 5.7 ms to plot a single line on the sonogram, this includes checking for requests by the user. The output of the frequency spectra to the screen is synchronised with the output of the DSP board via a flag which is updated each time the DSP board finishes processing a particular data frame. The output is displayed on the screen at a frame rate of 80 fps in the form of a sonogram.

3.7 Display - Graphical User Interface

The output was a graphical user interface (GUI) written in 'C' to display the sonogram of the data in real-time and a facility for saving the results to disk for further processing was provided. The most common format for displaying the spectra is using a sonogram, this is a two dimensional plot with time along the x-axis and frequency along the y-axis, a 128-tone colour scale is used to represent the magnitude

of the frequency components. The sonogram display on the PC shows 790 spectra on the screen at one time (corresponding to ~10 s of continuous data), the amplitude is divided into 128 different colours. Figure 3.13 shows a screen capture of the real-time spectrum analyser.

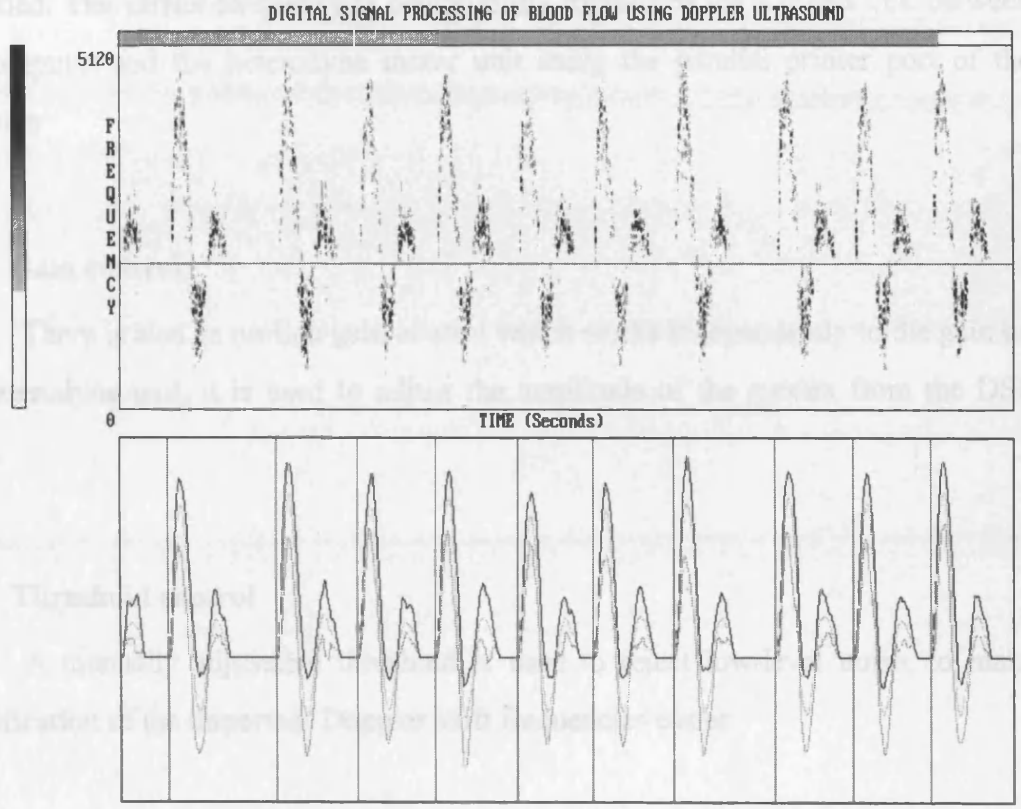

Figure 3.13 The output of the GUI displays the sonogram of a bi-directional Doppler signal processed using the real-time system described in this chapter together with the results of the heartbeat separation algorithm described in chapter 5. The grey-scale colour coding of the components in the sonogram is linear.

3.7.1 DSP algorithm

The user has the option of selecting the type of frequency analysis performed on the data, either FFT or AR and the sampling frequency prior to going on-line.

3.7.2 Heterodyne carrier frequency

The heterodyne carrier frequency can be modified on-line by selecting the desired frequency on the heterodyne unit. The position of the heterodyne frequency is plotted on the screen as a line so that forward and reverse flow can be easily identified. The carrier frequency is continuously monitored via a direct link between the computer and the heterodyne mixer unit using the parallel printer port of the computer.

3.7.3 Gain control

There is also an on-line gain control which works independently to the gain on the heterodyne unit, it is used to adjust the amplitude of the spectra from the DSP board.

3.7.4 Threshold control

A manually adjustable threshold is used to reject low-level noise, to make identification of the important Doppler shift frequencies easier.

3.7.5 Recording facility

There is also the facility to record the data which is controlled by the user. The only limit on the amount of data that can be recorded is governed by the amount of random access memory (RAM) the host computer has available. In order to save the data a '4 Mb RAM disk' was created which essentially works the same as a hard disk drive. The benefit of using this technique is that the access time for writing to the 'RAM disk' is significantly faster than writing directly to a hard disk. When recording is stopped the data is transferred to the hard disk for future use. All data is stored including low level spectral components regardless of the threshold value at the time of recording, this was done to enable extensive testing to be carried out on the same data using various threshold values and thresholding techniques. The file includes a 6-byte header (described in Appendix 1), which indicates the signal processing

technique used during recording, sampling rate and the heterodyne carrier frequency, for reproduction of the signal characteristics off-line.

3.8 Summary

So far in this thesis the text has described how digital signal processing tools such as the short-time Fourier transform and autoregressive modelling can be used to produce a time-frequency sonogram of the received Doppler echoes and a real-time system has been developed and its performance evaluated. In the next chapter cardiovascular disease is introduced and the way in which the joint time-frequency analysis of the Doppler signal can be used to detect and diagnose stenosis is described.

4. Cardiovascular disease - Stenosis

4.1 Introduction

Time-frequency analysis of the received Doppler ultrasound signal is the first stage in the overall diagnosis of cardiovascular disease. In this chapter the background behind the pathogenesis of atherosclerosis is presented. This chapter is focused towards ways in which Doppler ultrasound can be used to detect the presence of stenosis based on the effects that a stenosis has on blood velocity components within the vessel. The second part of the chapter identifies ways in which the time-frequency sonogram generated using the signal processing algorithms developed in the previous chapters can be analysed both qualitatively and quantitatively to diagnose and attempt to quantify a stenosis.

4.2 Stenosis

4.2.1 Mechanism for stenosis

The onset of atherosclerosis in the cardiovascular system is not uniform, some vessels are more prone to atherosclerosis than others for example the aorta, coronaries and carotids are primary sites for lesions to form. Lipid deposits accumulate at points in the vessel where stagnating flow is present for example bifurcations and bends in the vessel. There are a number of theories which attempt to explain the deposition of the lipids onto the endothelium which then act as a base for further lipid accumulation and calcification of the vessel: one theory is that the lesion is formed following injury to the endothelial surface of the vessel which then leaves the vessel vulnerable to atherosclerosis; another theory is that certain types of lipid are absorbed into the arterial wall for example cholesterol, thus, forming fatty deposits which can then act as a foundation for plaque formation (Born, 1992).

4.2.2 Effects of stenosis

The formation of a minor stenosis is not in itself a problem. The problems arise following secondary changes in the plaque or advanced calcification of the plaque so that it substantially or totally occludes the vessel lumen. Surgery on a stenosis is not generally considered until the diameter of the artery is reduced by more than 70 %. The results of the North American Symptomatic Carotid Endarterectomy Trial (NASCET) found that for stenosis greater than 70 % the best treatment was removal by carotid endarterectomy (Strandness, 1992).

Depending upon the shape of the stenosis, flow patterns are not dramatically affected by stenoses with less than 50 % reduction in diameter and are therefore difficult to detect (Reneman *et al.*, 1992; Beach and Phillips, 1992). As the degree of stenosis increases the volume flow rate is reduced, in the case of the carotid artery the blood supply to the brain can be affected. If some of the stenosis detaches itself and finds its way to the brain, this can cause serious problems, since the vessels in the brain are much smaller and the debris from the stenosis can totally block the vessel (thrombosis) leading to a possible transient ischaemic attack (TIA) or more seriously a stroke (Strandness, 1992; Born, 1992). Diagnosis of a significant stenosis (> 50 %) is fairly straight forward. It is important however to be able to differentiate between a 50 - 70 % stenosis and a stenosis with greater than 70 % reduction in diameter, since this could mean the difference between treatment with a course of aspirin or surgery (Strandness, 1992).

4.2.3 Detection of stenosis

Stenoses affect the pattern of blood flow in the vessel. Proximal to a stenosis the velocity profile will be normal, the velocity profile will depend upon the geometry of the vessel. If the artery is straight and smooth proximal to the stenosis then the flow is likely to be fully-developed laminar flow. The cardiovascular system, however, is highly complex and is made up of tortuous vessels that bifurcate and bend, therefore, 'normal' flow will depend upon the area of the cardiovascular system that is being studied. In the ICA for example the blood flow is predominately plug flow. At the

stenosis an increase in velocity is expected as the blood is forced through a narrower cross-section. Distal to the stenosis the flow may be turbulent and large velocity gradients exist as the blood flow decelerates and returns to normal, the presence of turbulence and velocity gradients are the best indication that a stenosis is present (Jones, 1993). Further downstream the velocity profile returns to normal. Figure 4.1 shows the changing velocity profiles through a stenosis, normal flow is considered to be parabolic, however, in many situations this will not be the case.

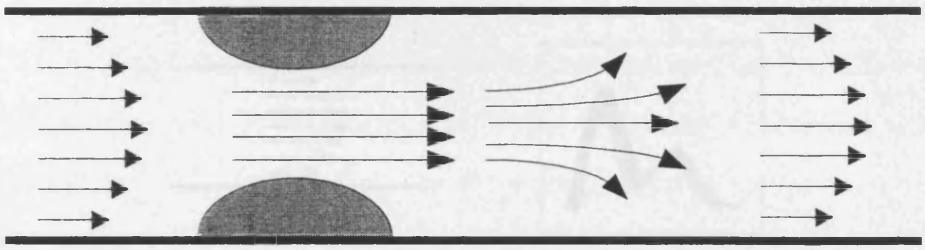

Figure 4.1 Changing velocity profile of blood flow across the vessel lumen as it is forced through a stenosis. 'Normal' flow is considered to be parabolic laminar flow in this example.

The type of flow detected using the Doppler velocimeter will depend upon the distance from the site of the stenosis to where the beam is insonating the vessel (Douville *et al.*, 1985). It is easier to detect the turbulence immediately distal to the stenosis rather than trying to penetrate the hard plaque at the site of the stenosis which will attenuate the ultrasound. Therefore the diagnosis and the quantification of the severity of the stenosis using Doppler ultrasound is usually indicated by the turbulence downstream rather than the increase in velocity at the site.

4.2.4 Effects of stenosis on velocity profile

In a healthy artery the velocity profile is usually fairly narrow since there are only a few low frequency components due to velocity vectors near the vessel wall. This creates a window beneath the time-frequency profile of the blood velocity components in the sonogram as illustrated in figure 4.2. The turbulence detected using the Doppler velocimeter is represented by a broadening of the frequency spectrum

around the systolic peak during the cardiac cycle. As the heart pumps the blood along the vessel there is maximum turbulence at this point. As the severity of the stenosis increases, the window in the sonogram below the velocity distribution is obliterated as low velocity components are also present (figure 4.2). If the stenosis causes eddy pools then there may even be a reverse flow component. The presence of a reverse flow component where it is not expected is a good indicator of stenosis (Evans *et al.*, 1989)[1], provided that it is not due to another vessel that has been encompassed in the examination.

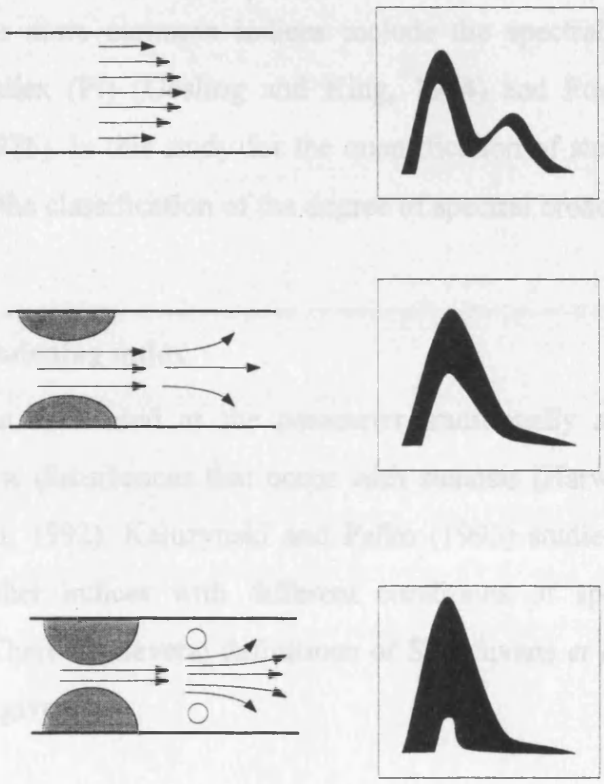

Figure 4.2 Effect of stenosis on the time-frequency profile (Hedrick *et al.*, 1996).

4.3 Qualitative analysis

The velocity profile during the course of the cardiac cycle is affected by the presence of stenosis and hence a qualitative assessment of the artery can be made based on the shape of the velocity profile. A number of features are used to assess the state of the artery, these include: spectral broadening, increased velocity at peak systole, diastolic flow and the presence of reverse flow. It is also possible for

subjective estimates to be made based on listening to the Doppler shifts, since, typically the range of Doppler shift frequencies lie in the audio range. Based on experience it is possible to differentiate between the sounds made by normal and diseased carotid arteries (Evans *et al.*, 1989)[1].

4.4 Quantitative analysis

There are a large number of indices that have been developed to measure various features of the sonogram for the prediction of arterial disease (Rittgers *et al.*, 1983). Some of the more common indices include the spectral broadening index (SBI), pulsatility index (PI) (Gosling and King, 1974) and Pourcelot's resistance index (Pourcelot, 1976). In this study for the quantification of stenosis the emphasis has been mainly on the classification of the degree of spectral broadening.

4.4.1 Spectral broadening index

The SBI has been used as the parameter traditionally associated with the measurement of flow disturbances that occur with stenosis (Harward *et al.*, 1986[1,2]; Labs and Fitzgerald, 1992). Kaluzynski and Palko (1993) studied the behaviour of SBI as well as other indices with different conditions of spectral analysis for simulated signals. There are several definitions of SBI (Evans *et al.*, 1989)[1], the one used in this work is given by:

$$SBI = \frac{f_{\max} - f_{mean}}{f_{\max}} \qquad [4.1]$$

where: f_{max} = the maximum frequency component in the spectrum above a threshold used to reject low level noise components and f_{mean} = the mean frequency of the spectral components that exceed this threshold (Brown *et al.*, 1982), figure 4.3 illustrates the position of f_{max} and f_{mean}.

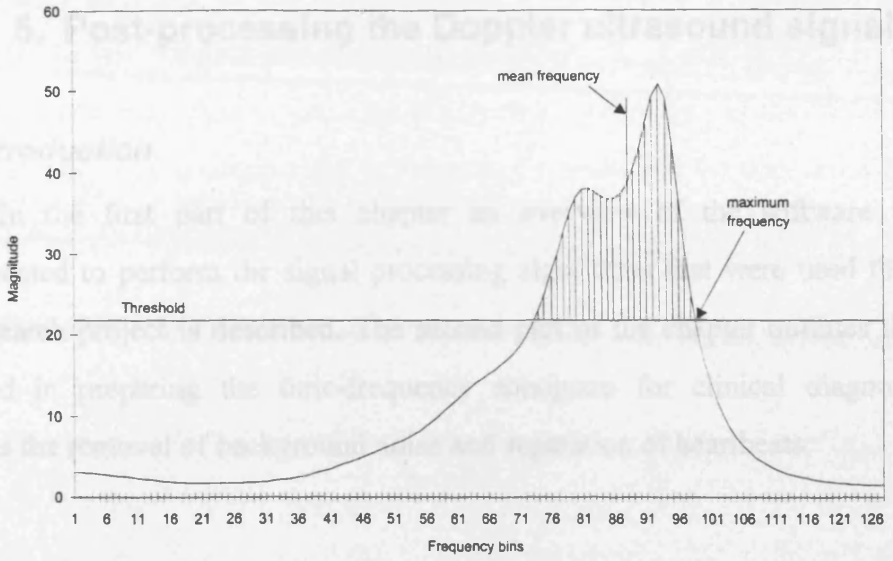

Figure 4.3 Determination of f_{max} and f_{mean} from frequency spectrum.

If a high grade stenosis is affecting the flow pattern, then it is possible to have simultaneous forward and reverse flow components as a result of turbulence. In this study forward and reverse flow are combined to produce a single signal. If forward and reverse flow components are present simultaneously then [4.1] is not adequate since it is only suitable for establishing the spectral broadening of a single wide band lobe. In order to overcome this problem the reverse flow was eliminated from the calculation of SBI and only the forward flow components were used to determine the f_{mean} envelope.

4.5 Summary

In this chapter the relationships between blood flow and the time-frequency sonogram and the effect that the presence of a stenosis has on the sonogram have been discussed. It has also been shown how the sonogram can be used to both qualitatively and quantitatively diagnose the condition of the vessel under observation. The next chapters develop these ideas and analyse the reliability of the spectrum analysis using both simulated and real data.

5. Post-processing the Doppler ultrasound signal

5.1 Introduction

In the first part of this chapter an overview of the software that was implemented to perform the signal processing algorithms that were used throughout this research project is described. The second part of the chapter outlines the stages involved in preparing the time-frequency sonogram for clinical diagnosis. This involves the removal of background noise and separation of heartbeats.

5.1.1 Off-line digital signal processing

The real-time system has already been discussed in chapter 3, this was used to assess the potential for real-time implementation of autoregressive algorithms using a modern DSP board compared to the computationally efficient FFT-based algorithms which are used in commercial systems. An off-line signal processing program was also developed to allow extensive testing of different algorithms without the constraints of real-time processing. The system was developed in 'C' and was designed to run on an IBM PC compatible computer. This allowed more flexibility in the design of the program than using the DSP32C as for the real-time system due to the restricted memory capacity on the DSP32C board. This system is used in the next two chapters for the analysis of simulated Doppler signals and real clinical data. A detailed breakdown of the operation of the computer-based spectrum analyser is given in Appendix 1.

5.2 Post-processing of the Doppler signal

5.2.1 Determination of frequency envelopes

The time-frequency sonogram obtained using any of the spectrum estimation techniques described above contains a large amount of information about the Doppler signal. In order to perform quantitative analysis of the results it is necessary to reduce the amount of information and pick out the features that can aid in the diagnosis of the

state of the artery. One of the ways that this can be done is to determine frequency envelopes. In this research the maximum frequency (f_{max}) and the mean frequency (f_{mean}) envelopes were used to calculate the spectral broadening of the Doppler signal. The envelopes were determined by calculating the mean and the maximum frequency components of an individual frequency spectrum above a certain threshold.

5.2.1.1 Implementation of the frequency envelopes

Envelope detection is achieved by determining f_{min}, f_{mean} and f_{max} for each frequency spectrum in the sonogram. The point at which the magnitude of the frequency spectrum first exceeds the specified threshold value is defined as f_{min} and the last time the frequency spectrum exceeds the threshold value is defined as f_{max}. The mean frequency, f_{mean}, is calculated based on the weighted average of the bins between f_{min} and f_{max} which are above the threshold [5.1]. These envelopes are calculated for the whole of the recorded signal and displayed on the screen.

$$f_{mean} = \frac{\sum_{n=1}^{N}(A_n \cdot f_n)}{N} \qquad [5.1]$$

where A_n = magnitude of frequency bin, f_n = frequency bin, N = number of frequency bins.

5.2.1.2 Noise removal

The implementation of a suitable threshold is required to separate the true Doppler signal from the background noise. Noise is inherently present in the raw data signal since the Doppler velocimeter will detect and amplify any signals that it receives during the examination regardless of their origin. Secondly, artefacts which may be due to the spectrum analyser, for example, spectral leakage in the case of the FFT contribute to noise since these leakage terms mask the true Doppler signal. It is therefore necessary to apply a threshold to the sonogram to remove low-level noise components which could bias the results. Two types of threshold were examined:

5.2.1.3 Global threshold

A global threshold was selected and applied to all the frequency spectra as described by Gibbons *et al.* (1981). The selection of the threshold level is based upon a subjective study of the Doppler sonogram by the operator and is fixed for all spectra. Figure 5.1 illustrates the application of the threshold to reject low-level noise components. Automatic threshold detectors have been used to remove low-level noise (Evans *et al.*, 1989)[2].

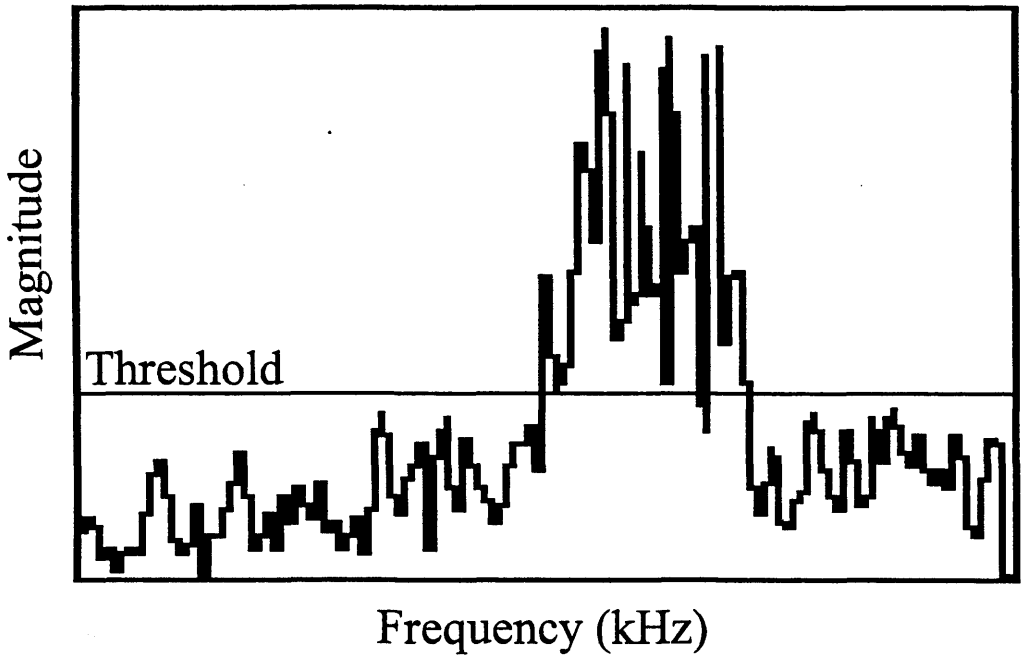

Figure 5.1 Simple threshold to reject low-level noise components.

5.2.1.4 Individual frequency spectrum threshold

This threshold is pre-defined and is measured with respect to the maximum component of each individual frequency spectrum as can be seen in figure 5.2. Different levels of threshold were examined: -3, -6, -9, -12, -15 and -18 dB in relation to the maximum magnitude component of the frequency spectrum (Harward *et al.*, 1986) [1,2]. The optimum threshold should be able to maximise the ability to reject noise but also retain as much of the Doppler signal as possible. Figure 5.3 shows the results

obtained using the different fixed threshold levels by superimposing the mean and maximum frequency envelopes on the FFT and AR sonograms.

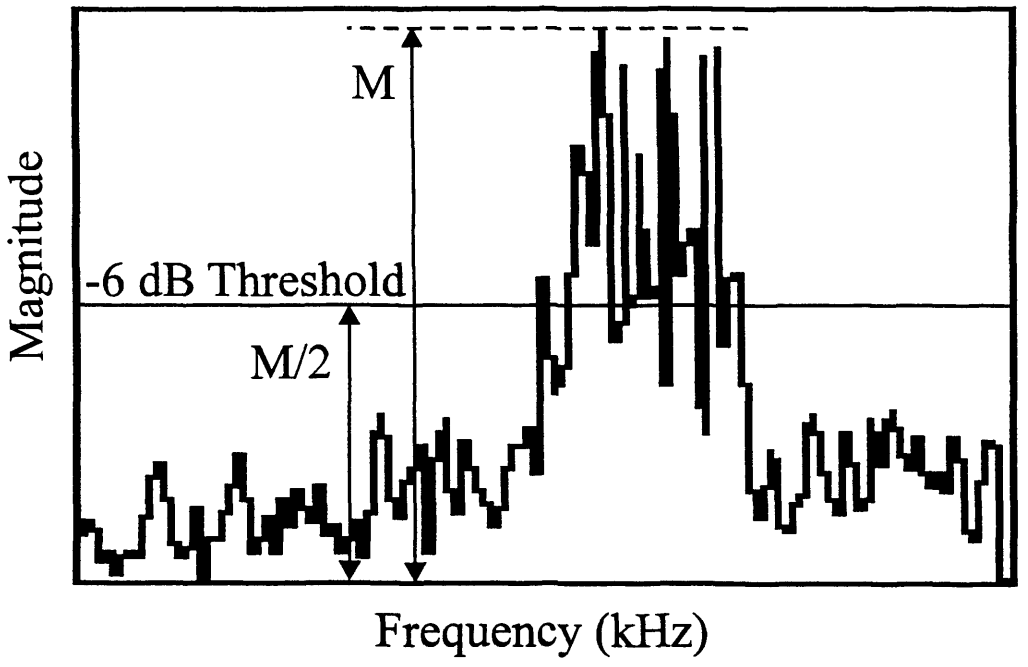

Figure 5.2 Implementation of the -6 dB threshold.

By comparing the respective thresholds for the FFT and AR sonograms in figure 5.3, it can be seen that the AR threshold is more robust than that of the FFT. This is due to the qualitative improvement in the AR spectral estimates over the FFT. As the threshold is reduced the distance between the maximum and mean frequency envelopes increases this is due to the tapered distribution typically observed in Doppler spectra. It was decided to use the -6 dB threshold as this gave the most sensible trade-off between background noise and maximum retention of the Doppler signal for both algorithms.

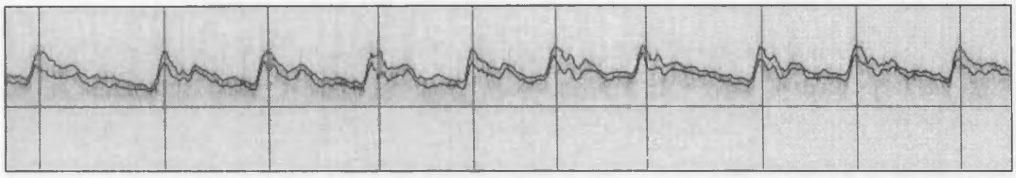

(a) -3 dB FFT

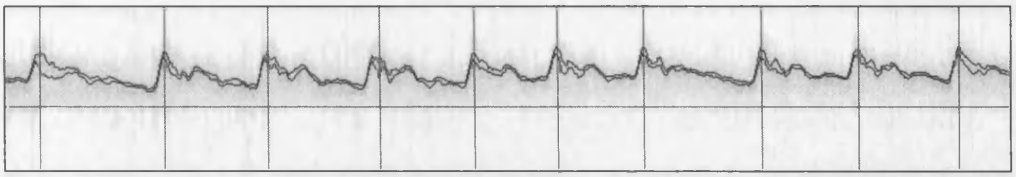

(b) -3 dB AR

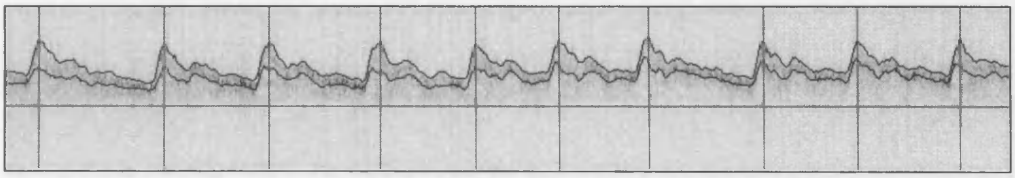

(c) -6 dB FFT

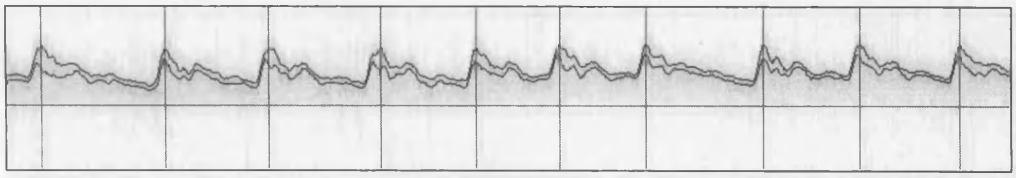

(d) -6 dB AR

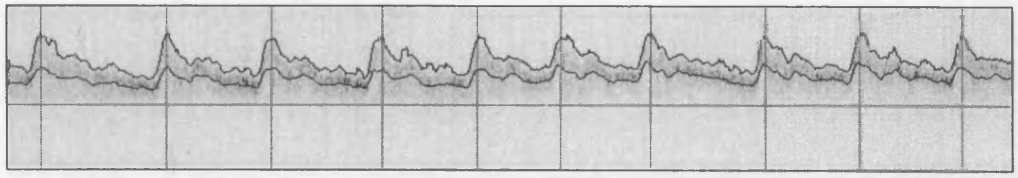

(e) -9 dB FFT

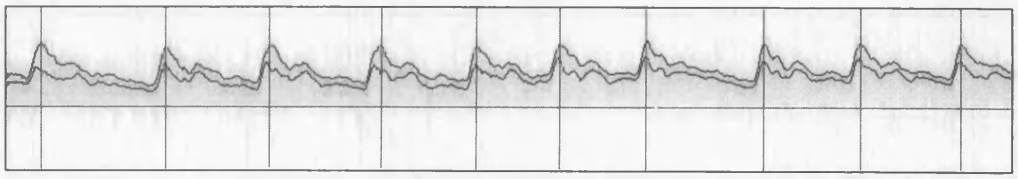

(f) -9 dB AR

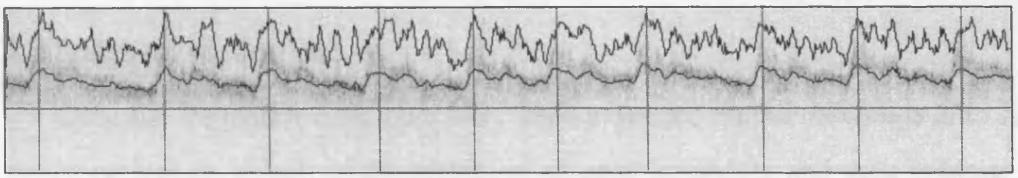

(g) -12 dB FFT

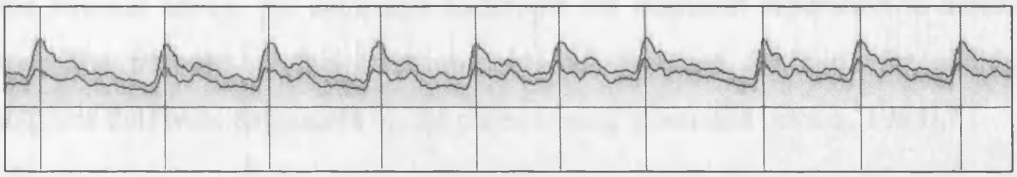

(h) -12 dB AR

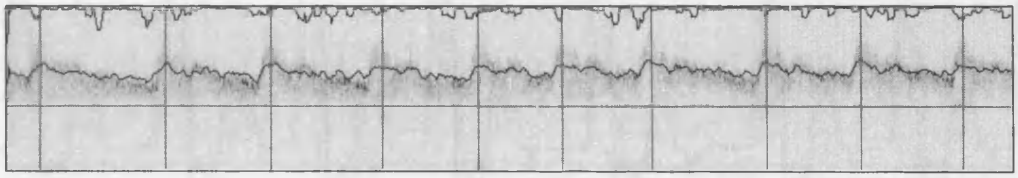

(i) -15 dB FFT

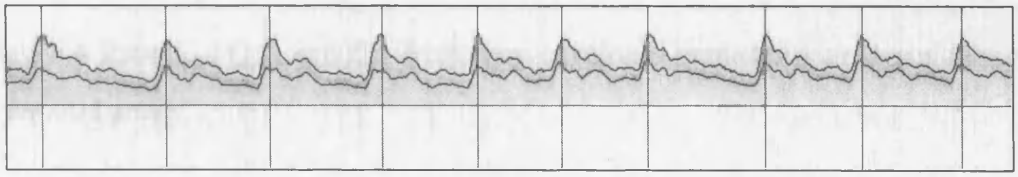

(j) -15 dB AR

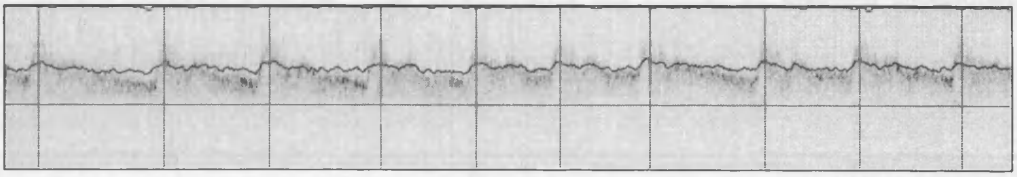

(k) -18 dB FFT

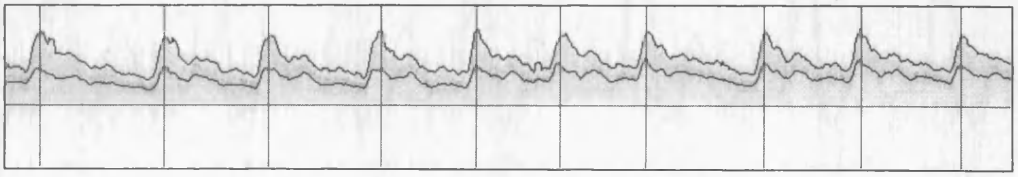

(l) -18 dB AR

Figure 5.3 Mean and maximum frequency envelopes superimposed on sonograms calculated using FFT and AR, threshold determined as 3, 6, 9, 12, 15, 18 dB below the maximum component of the individual spectra.

5.2.2 Heartbeat separation

Once the frequency envelopes have been extracted the second stage is to filter the envelope signals to reduce the effects of spectral instability and then separate the heartbeats. Figure 5.4 shows the raw frequency envelopes extracted from a sonogram of the femoral artery. An automatic technique for heartbeat separation is described below. The purpose of this technique was to eliminate the need for additional parameters that were dependent on the patient being examined (Evans, 1988).

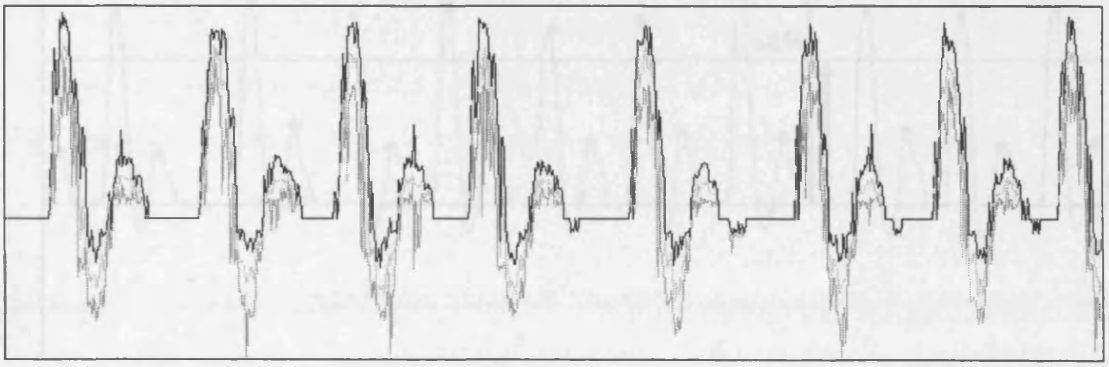

Figure 5.4 Raw f_{min}, f_{mean} and f_{max} frequency envelopes extracted from a sonogram of the femoral artery.

- The raw minimum, mean and maximum frequency envelopes were passed through a 5-point moving average filter to minimise the effects of spectral instability as illustrated in figure 5.5.

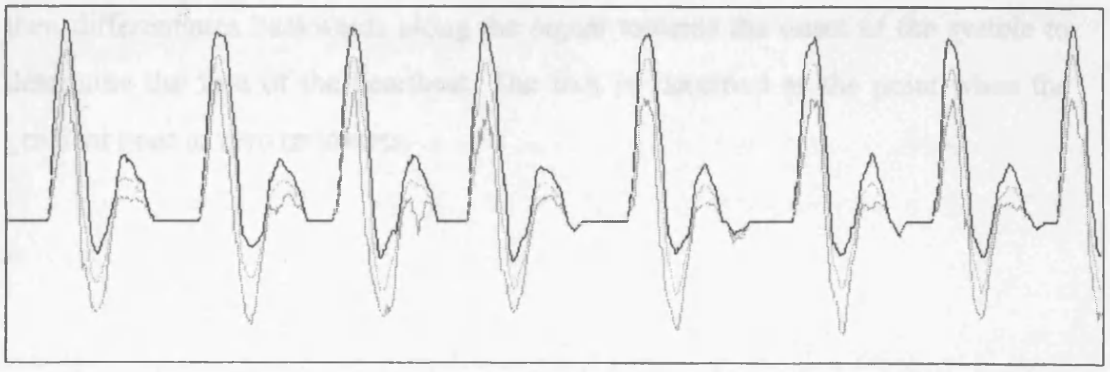

Figure 5.5 Filtered f_{min}, f_{mean} and f_{max} frequency envelopes.

- All local maxima in the f_{max} envelope were located by searching for a positive followed by a negative gradient as shown in figure 5.6.

- The largest peak is obtained from the maxima determined in the previous stage and the difference between this maxima and the heterodyne frequency is evaluated. A 66% threshold is used to reject peaks on the basis that they are not systolic maxima. These stages are illustrated in figure 5.6.

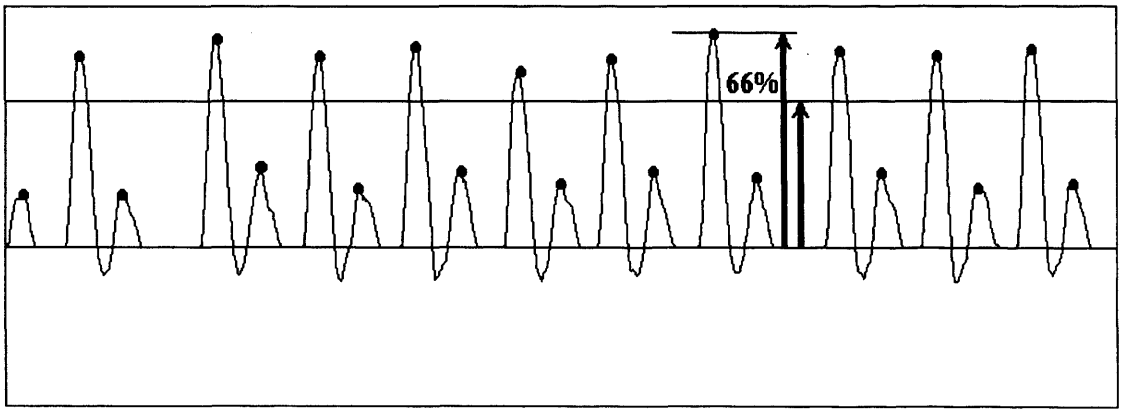

Figure 5.6 Identification of the systolic peaks by applying a threshold to reject non-systolic maxima.

- Rejection of peaks that follow within 0.5 seconds of a systolic peak being identified. This prevents multiple peaks being identified within a single heartbeat.

- Once the systolic peaks for all the heartbeats have been identified the algorithm then differentiates backwards along the signal towards the onset of the systole to determine the foot of the heartbeat. The foot is identified as the point when the gradient goes to zero or inverts.

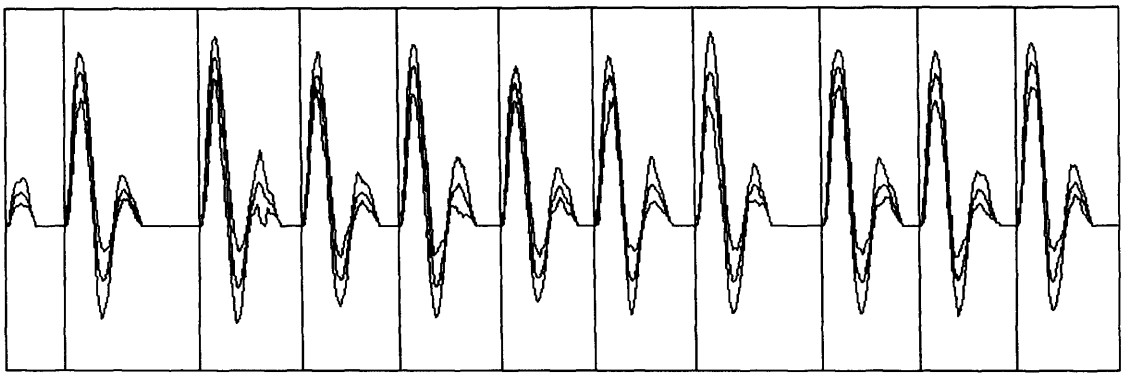

Figure 5.7 Separation of the frequency envelopes into the individual heartbeats.

Figure 5.7 shows the ability of the automatic heartbeat separation algorithm to successfully separate the individual heartbeats. The overall performance of the algorithm was, however, limited when dealing with extremely noisy envelopes.

5.2.3 Determination of spectral broadening index

Following the separation of the heartbeats it is then possible to estimate the spectral broadening index (SBI). There are two approaches for producing an estimation of the average value of a parameter derived from the sonograms:

5.2.3.1 Averaging heartbeats

The first method involves averaging spectra from different heartbeats, and then calculating the SBI from the averaged spectrum.

5.2.3.2 Averaging parameters

The second method involves calculating the SBI from the individual spectra from a single heartbeat and then averaging the values obtained over a number of heartbeats.

It was decided to use the second method here to estimate the SBI in order to avoid the inherent difficulties of time alignment of the first technique and the possibility of smearing the averaged spectrum if the alignment is not perfect. The SBI

was calculated at the systolic peak which is defined as being the maximum frequency component in the f_{max} envelope.

To conclude this chapter the performance of the two types of threshold described earlier are compared.

5.2.4 Comparison of the two types of threshold

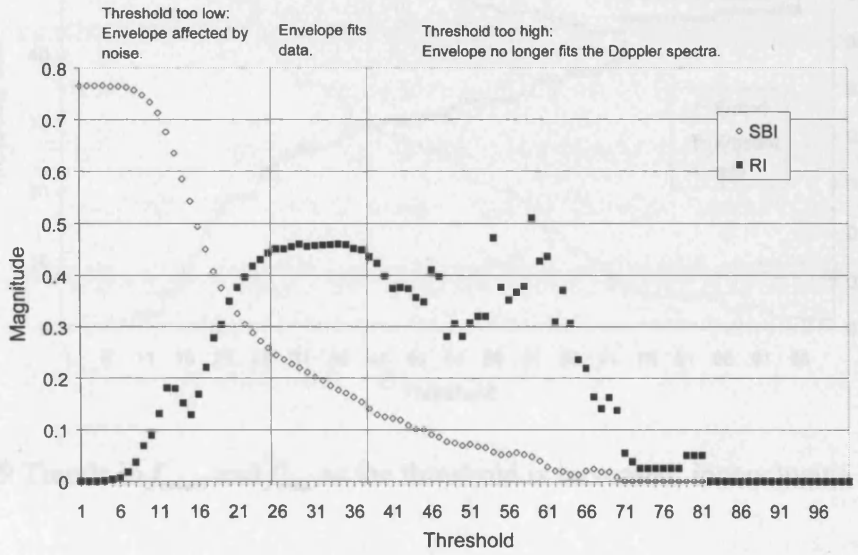

Figure 5.8 SBI & resistance index (RI) vs. threshold for an internal carotid artery.

It was found that the positioning of the global threshold was fairly ambiguous and a range of values exists over which the threshold could be chosen. It has been observed by Schlindwein *et al.* (1994) that if a simple threshold is implemented then there is no obvious plateau on the SBI *vs.* threshold plot where the SBI can be considered constant. Figure 5.8 shows the SBI and the resistance index (RI) (Pourcelot, 1976) for a real Doppler signal, the signal was recorded from an internal carotid artery. It can be seen that there is a range of thresholds over which the value of RI is constant but the SBI is changing (RI has been included here to illustrate the range of thresholds over which the envelopes used to determine the diagnostic parameters fit the signal, i.e. the range of thresholds where noise is rejected with

minimal loss of the true Doppler signal). The reason why there is no plateau for SBI is due to the trends in the mean and maximum frequency envelopes: As the threshold is increased, f_{max} decreases and f_{mean} increases as can be seen from figure 5.9. The uncertainty caused by the positioning of the threshold results in an inconclusive SBI value.

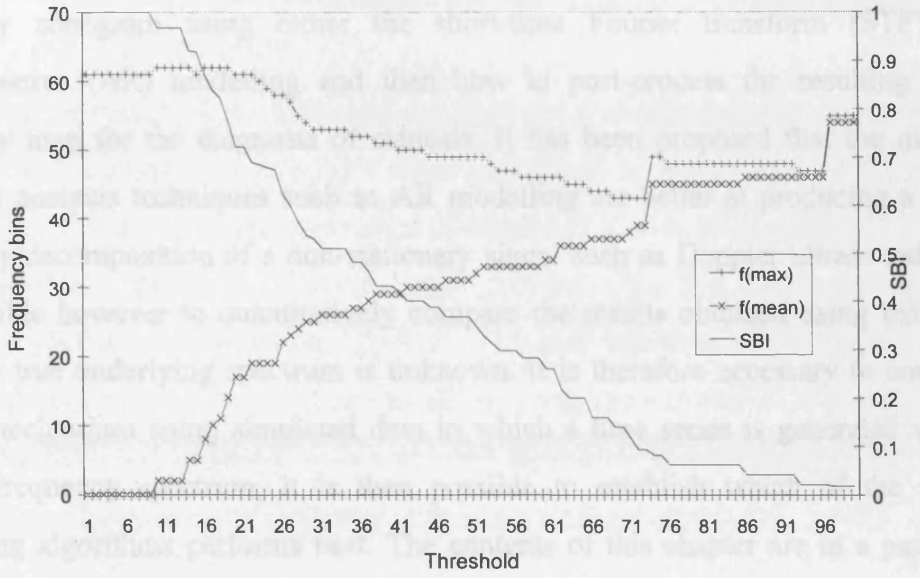

Figure 5.9 Trends in f_{mean} and f_{max} as the threshold is increased; inconclusive SBI.

In the case of the threshold based on the individual spectra there is no ambiguity caused by the choice of threshold since the threshold is defined relative to the magnitude of the individual frequency spectrum and is not fixed at a certain level. This significantly reduces the sensitivity of the SBI to gain and also produces a single estimate of SBI for a particular recording.

5.3 Summary

This chapter has outlined the techniques required to post-process the time-frequency sonogram. In the next chapter simulated data is used to compare the accuracy of the FFT and AR algorithms in the presence of noise.

6. Spectral broadening of simulated Doppler signals

6.1 Introduction

The previous chapters have described the stages necessary to take a raw ultrasound signal from its original time-domain form and to convert it into a time-frequency sonogram using either the short-time Fourier transform (STFT) or autoregressive (AR) modelling and then how to post-process the resulting time-frequency map for the diagnosis of stenosis. It has been proposed that the modern spectrum analysis techniques such as AR modelling are better at producing a time-frequency decomposition of a non-stationary signal such as Doppler ultrasound. It is not possible however to quantitatively compare the results obtained using real data since the true underlying spectrum is unknown. It is therefore necessary to compare the two techniques using simulated data in which a time series is generated with a known frequency spectrum, it is then possible to establish which of the signal processing algorithms performs best. The contents of this chapter are in a paper by Keeton *et al.* (1997)[2].

This chapter is divided into two parts, firstly, the relationship between the spectra obtained using FFT and AR was studied using the spectral broadening index (SBI), described in chapter 4, for simulated symmetrical wide band signals using data frames containing 64 and 256 samples. Secondly, three filters were used to simulate Doppler spectra typically found around the systolic peak which represent signals taken from a healthy carotid artery and signals that represented diseased carotid arteries. Gaussian noise was added to the simulated Doppler signals over a range of signal-to-noise ratios (SNRs). The accuracy and robustness of the SBI for both the FFT and AR spectra were compared for a range of SNRs, since estimation of the frequency spectrum is complicated by the presence of noise (Kay, 1979; Scott *et al.*, 1987).

It has been reported that typical clinical SNR levels are in the range of 0-20 dB (Forsberg, 1991). If the signal processing algorithm is not stable in the presence of noise then the variance of the spectral estimate is going to be large. Parameters

derived from these spectra will have a large variance making it clinically infeasible to produce a statistically significant estimate since the number of heartbeats that need to be sequentially recorded for averaging the SBI will be too large. Therefore, one of the aims of this research is to establish the magnitude and variance of SBI in the presence of noise at SNRs present in a clinical environment.

6.2 Methods & Discussion

The behaviour of the SBI obtained from spectral estimates based on the FFT approach was compared to that from spectra produced by autoregressive modelling. The first algorithm was a straightforward 256-point FFT with no overlap and no anti-leakage window. The second algorithm was identical to the first except that the data was multiplied by a Hamming window. The AR algorithm was based on an implementation of the Yule-Walker equations using the recursive Levinson-Durbin approach with model order 12 as described by Schlindwein and Evans (1989). Both techniques produced spectral estimates with 128 frequency bins.

The signals used in this study were generated using filtered Gaussian white noise which has a uniform theoretical spectrum from zero to half the sampling frequency. The shape of the frequency spectrum of the signals was determined by filtering the Gaussian white noise using a set of digital filters. Using this method a series of wide band signals with varying bandwidths and frequency response characteristics were generated. The ideal shape of the filter was approximated using an 8-order infinite impulse response (IIR) Yule-Walker digital filter whose coefficients were determined using a least squares fit to the desired filter shape. The design was performed using MATLAB. A Gaussian white noise sequence of N samples was filtered and the first 256 filtered samples were discarded so that transients caused by the filter would not be processed. The impulse response of the filter was used to determine the number of samples that should be rejected to ensure that the initial transients had decayed sufficiently within the first 256 samples. Figure 6.1 shows the frequency response of the symmetrical band-pass filters that were used to create the wide band signals.

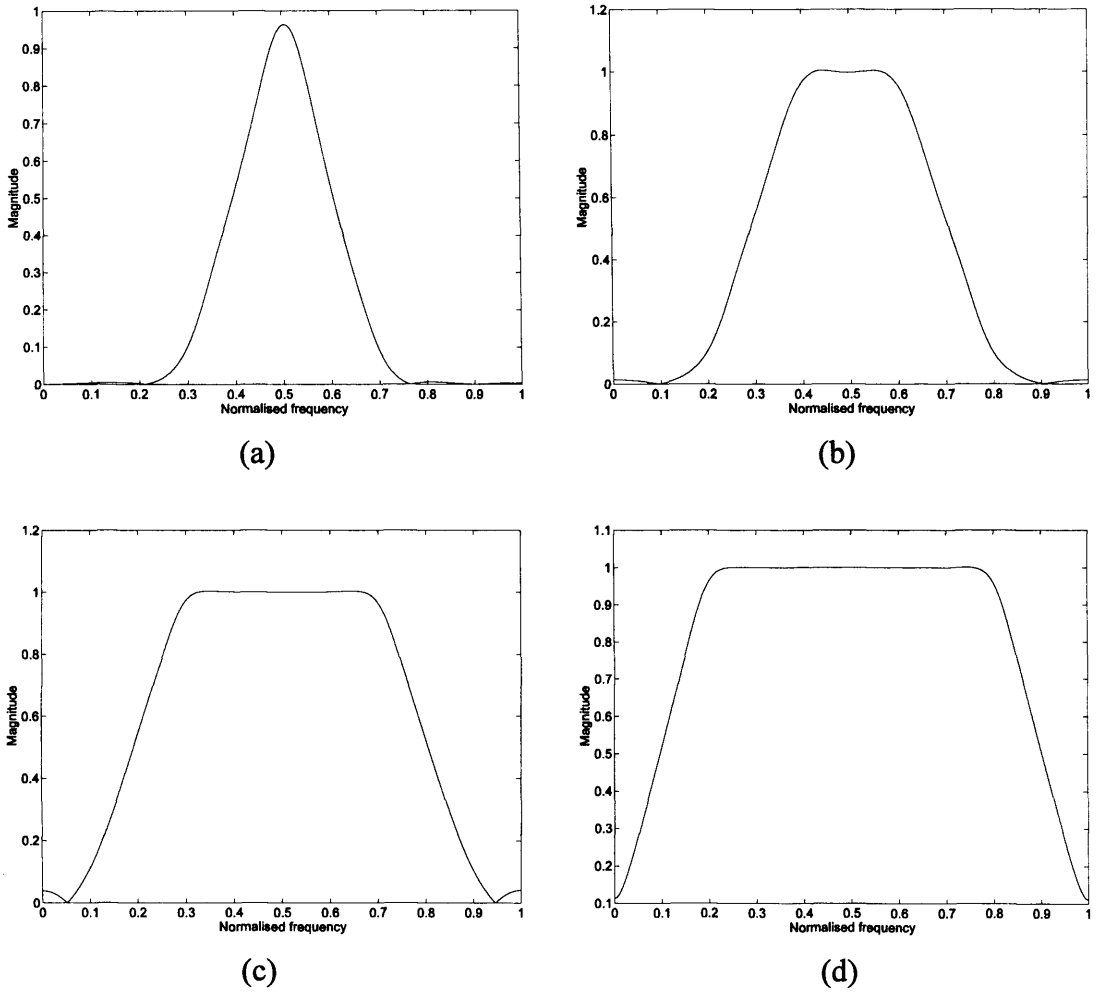

Figure 6.1 Frequency response of filters used to generate symmetrical wide band signals. Normalised 6 dB bandwidths for each filter are (a) 0.2 (b) 0.4 (c) 0.6 (d) 0.8.

The calculation of SBI was performed using [4.1]. The f_{mean} and f_{max} frequency envelopes were determined using the -6 dB threshold described in chapter 5 in which the level of the threshold was measured with respect to the maximum magnitude component in each individual frequency spectrum.

6.3 Comparison of FFT-SBI & AR-SBI

Preliminary tests were carried out to establish the relationship between FFT-based SBI (with and without an anti-leakage window) and AR-based SBI by comparing the magnitude and variance of the estimates. The SBI was calculated from 64 separate spectral estimates obtained using a 64 point data frame and a 256 point

data frame using each of the algorithms described above. The results obtained were plotted against one another and are shown in figure 6.2.

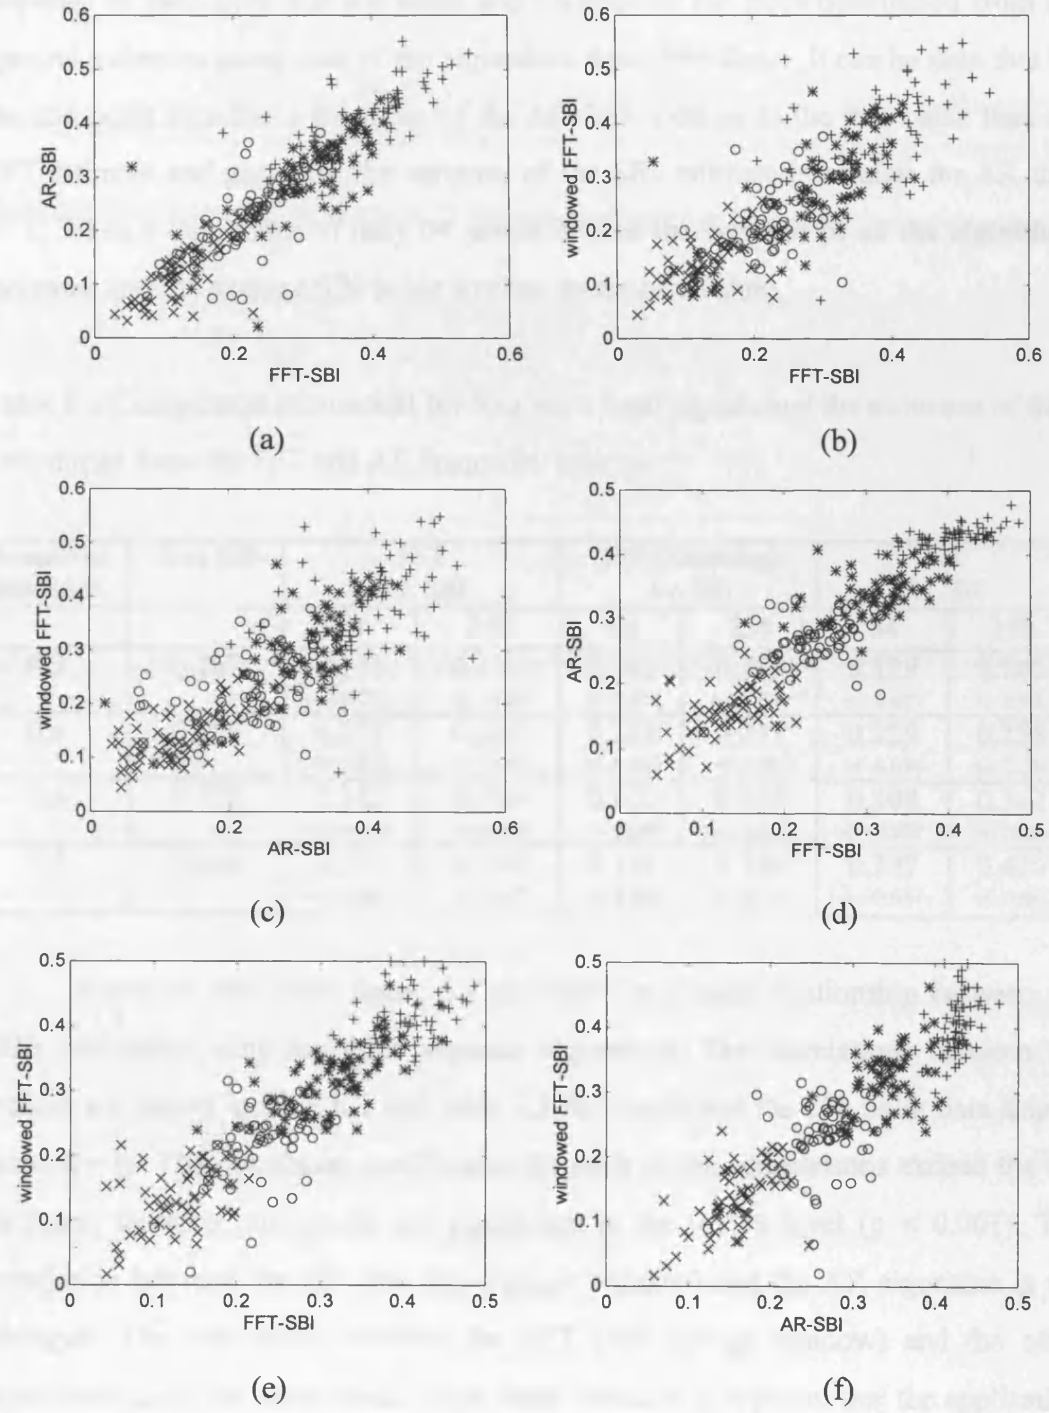

Figure 6.2 Comparison of FFT-SBI, AR-SBI and windowed FFT-SBI for four symmetrical wide band signals (a),(b),(c) 64-point data frame (d),(e),(f) 256-point data frame.

Normalised Bandwidth	
×	0.2
o	0.4
*	0.6
+	0.8

Table 6.1 shows the exact value of the SBI determined from the frequency response of each filter and the mean and variance of the SBIs determined from the spectral estimates using each of the algorithms described above. It can be seen that for the 256 point data frame the mean of the AR-SBI is closer to the true value than the FFT estimate and generally the variance of the SBI estimate is smaller for AR than FFT. When a data frame of only 64 points is used the variance of all the algorithms increases and the average SBI is not so close to the true value.

Table 6.1 Comparison of true SBI for four wide band signals and the estimates of SBI determined from the FFT and AR frequency spectra.

Normalised Bandwidth	True SBI	FFT Av. SBI		FFT (Hamming) Av. SBI		AR Av. SBI	
		64	256	64	256	64	256
0.2	0.167	0.133 +/- 0.053	0.131 +/- 0.043	0.141 +/- 0.048	0.123 +/- 0.060	0.129 +/- 0.059	0.146 +/- 0.023
0.4	0.286	0.227 +/- 0.062	0.237 +/- 0.053	0.223 +/- 0.066	0.211 +/- 0.081	0.229 +/- 0.076	0.258 +/- 0.054
0.6	0.375	0.312 +/- 0.076	0.333 +/- 0.038	0.300 +/- 0.089	0.322 +/- 0.068	0.308 +/- 0.080	0.342 +/- 0.026
0.8	0.444	0.378 +/- 0.081	0.396 +/- 0.053	0.379 +/- 0.090	0.394 +/- 0.064	0.397 +/- 0.074	0.426 +/- 0.031

It can be seen from figure 6.2 that there is a linear relationship between the SBIs calculated using the three separate algorithms. The correlations between the indices are shown in table 6.2 and table 6.3 for the 64 and the 256 point data frames respectively. The correlation coefficients for each of the comparisons exceed the 0.1 % value, therefore the results are significant at the 0.1 % level ($p < 0.001$). The correlation between the FFT (no anti-leakage window) and the AR algorithm is the strongest. The correlation between the FFT (anti-leakage window) and the other algorithms gives the worst result. From these results it is apparent that the application of an anti-leakage window does not improve the estimation of SBI, in fact the variance of the SBI when using FFT and an anti-leakage window is larger than that of the SBI obtained for FFT spectra with no anti-leakage window.

Table 6.2 Matrix of Pearson Product-Moment Correlation Coefficients for 64-point data frame.

	AR	FFT (Hamming)	FFT
FFT	0.90	0.84	1.0
FFT (Hamming)	0.83	1.0	
AR	1.0		

Table 6.3 Matrix of Pearson Product-Moment Correlation Coefficients for 256-point data frame.

	AR	FFT (Hamming)	FFT
FFT	0.93	0.91	1.0
FFT (Hamming)	0.90	1.0	
AR	1.0		

6.4 Effect of AR model order on SBI

When implementing the AR algorithm there is a certain degree of choice in the selection of model order for the AR model. There are various techniques for establishing the 'best' order for a particular process. For Doppler signals Kaluzynski (1989) and then Schlindwein and Evans (1990) suggested that the AR model order (p) should be around 12. The effect of using different model orders was tested using the same symmetrical wide band signals that were used to compare FFT-SBI and AR-SBI. Figure 6.3 shows the result of model order on the estimation of AR-SBI. If the model order is less than about 8 then there are not enough poles to adequately model the signal and therefore the spectral estimates are not representative of the true spectrum. If the model order is too high then the filter attempts to model the local fluctuations in noise as well as the underlying signal producing a biased estimate (Akaike, 1974). In addition the computational time to calculate the autoregressive coefficients using the Levinson-Durbin recursive algorithm is proportional to p^2, therefore, in real-time applications of the AR algorithm it is sensible to avoid unnecessary computation. For AR model orders between 10 and 16 the variance of AR-SBI is smaller than that of FFT-SBI for all cases.

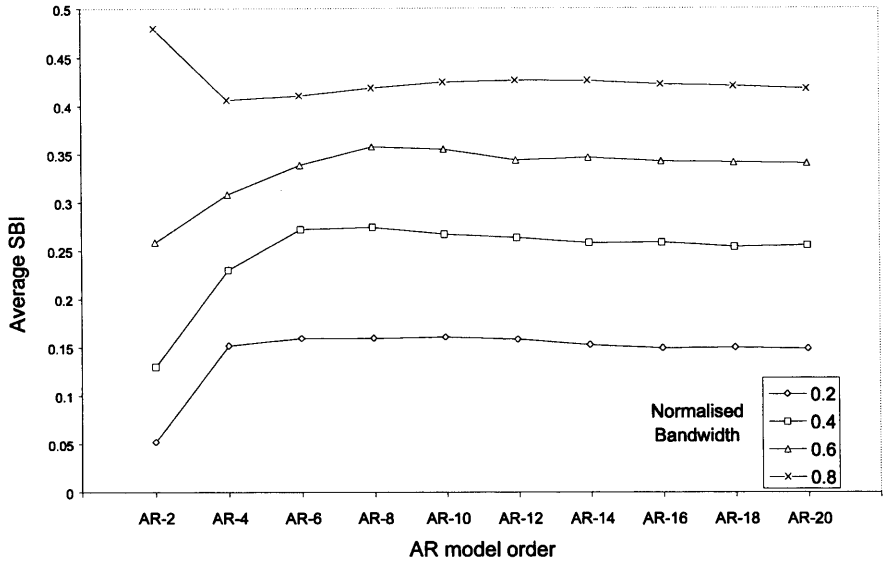

Figure 6.3 Effect of AR model order on the estimation of SBI for four symmetrical wide band signals.

6.5 Estimation of SBI from simulated Doppler signals

Having established that the FFT and AR algorithms produce comparable indices for symmetrical signals the second stage of the study was to look at the ability of the algorithms to produce estimates of SBI in the presence of noise and with asymmetric spectra, which are more comparable to Doppler signals. Simulated Doppler signals were created using digital filters modelled on the typical frequency spectra obtained around the systolic peak of the cardiac cycle (Hedrick *et al.*, 1995). Three types of filter were used to simulate different degrees of spectral broadening. Figure 6.4 shows the frequency responses of the filters that were used to generate the simulated Doppler signals. The first filter has only slight spectral broadening and is characteristic of a healthy ICA carotid artery, where plug flow is present. The second two spectra simulate the effect of a stenosis which would lead to spectral broadening as the flow would be disturbed. A second Gaussian white noise signal was used to add noise to the simulated Doppler signals to generate noisy signals with SNR from +10 dB down to -10 dB.

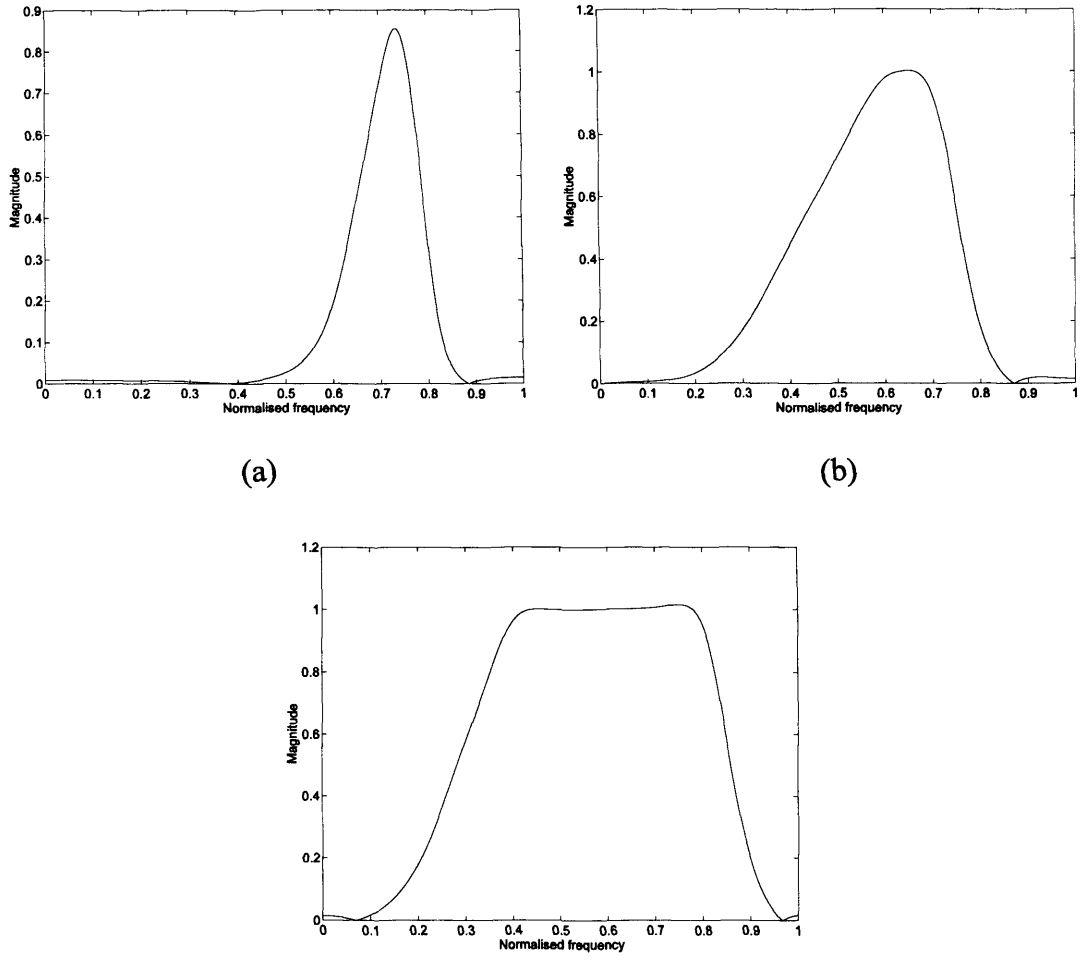

Figure 6.4 Frequency response of the filters used to simulate the Doppler spectra

The signals were processed using the FFT and AR algorithms using data frames of 64 and 256 samples. Figure 6.5 shows examples of typical spectra obtained using each of the algorithms with a data frame of 256 samples.

It can clearly be seen that the AR spectral estimate in figure 6.5(c) is closer to the shape of the filter, refer to figure 6.4(b), than either of the FFT spectra. All three spectral estimators produce spectra with a varying degree of ripple in the pass band, figure 6.6 shows the variance in the spectral estimates for the simulated Doppler spectra, figure 6.4(c), using AR with a data frame of 256 points.

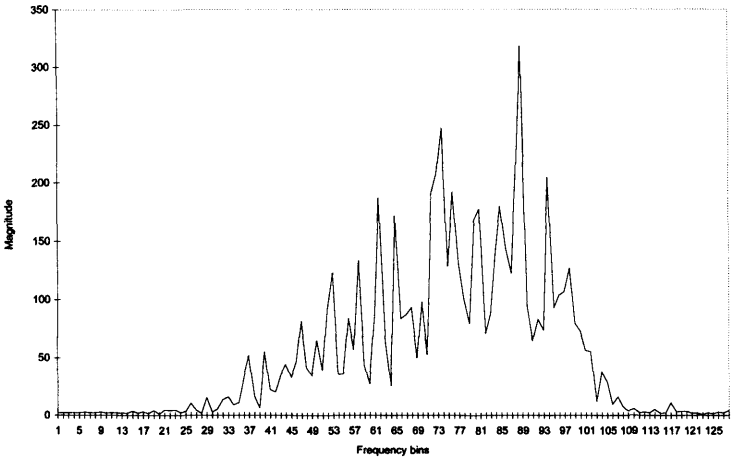

(a)

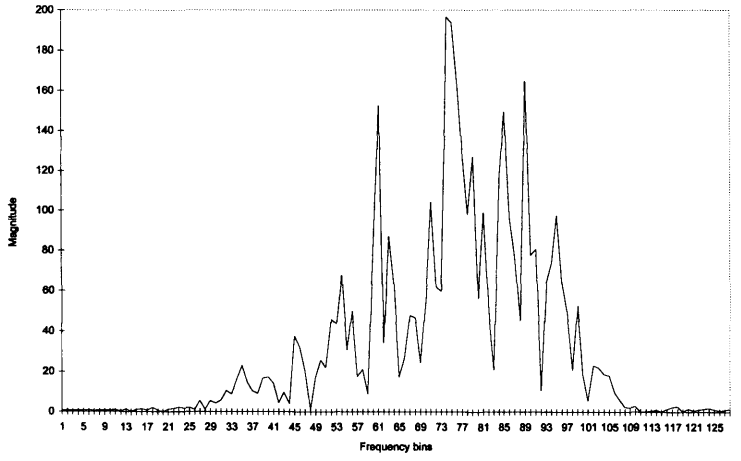

(b)

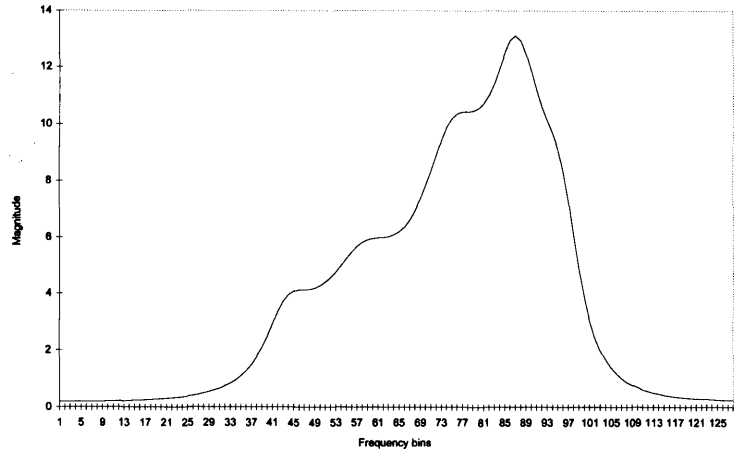

(c)

Figure 6.5 Frequency response for simulated Doppler spectrum. (a) FFT (b) FFT-Hamming window (c) AR model order 12. Signal generated using filter with frequency response shown in figure 6.4(c).

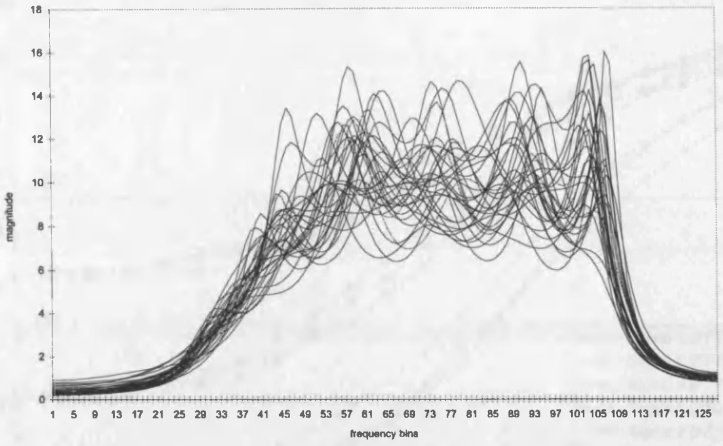

Figure 6.6 Spectral estimates for simulated Doppler spectra using AR (model order 12) on consecutive 256 point data frames.

The results show that although there is a high variance in the pass band, the bandwidths of the estimated spectra are stable. The average SBI and the variance was estimated from 128 and 32 frames for each signal using the 64 and 256 sample data frames respectively. The -6 dB threshold was used to determine SBI from the simulated Doppler spectra. The results are shown in figure 6.7 and figure 6.8.

It can be seen that for high SNRs the average SBI is constant for both FFT and AR. Figure 6.7 also shows the true SBI for each of the filters revealing that all of the algorithms underestimate the actual value, this is particularly true for the 64 point data frame. When the 256 point data frame is used the estimate of SBI is better, with AR-SBI producing the most accurate result. The variance of the SBIs for the short data frames is approximately the same for both the FFT and AR, however the AR algorithm performs marginally better than the FFT for a 256 point data frame. The ability of the FFT to perform as well as AR particularly for the short data frames can be attributed to the fact that the SBI is solely dependent upon the rise and fall of the wide band frequency spectral estimate, the variance across the pass band does not affect SBI as long as the ripple is uniform. Therefore even though the individual spectral estimates for FFT are qualitatively worse than AR, when the results are averaged over a number of frames the FFT is able to closely match the shape of the filter of figure 6.4(c). Figure 6.9 shows the averaged spectra for all the algorithms using 64 and 256 samples.

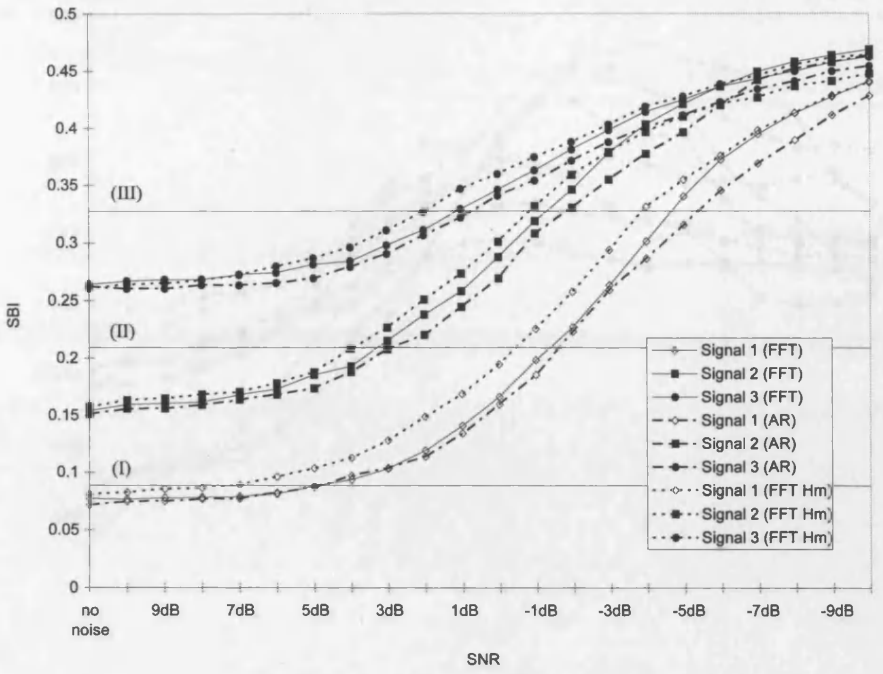

(a)

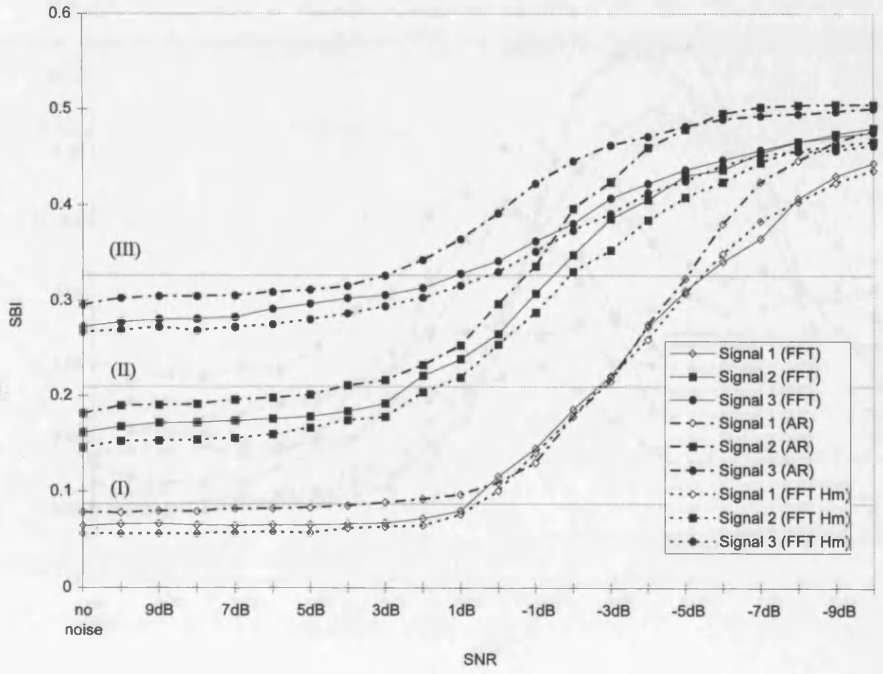

(b)

Figure 6.7 Average SBI for three simulated Doppler signals in noise (a) 64 point data frame (b) 256 point data frame. True SBI for simulated Doppler spectra indicated using the lines I, II & III for the three asymmetric filters in figure 6.4(a,b,c) respectively.

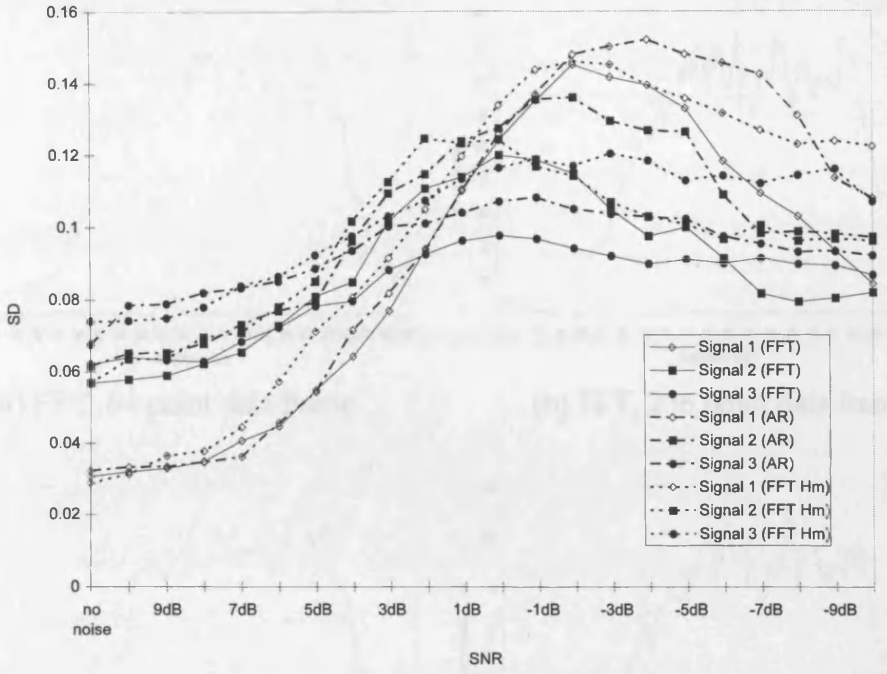

(a)

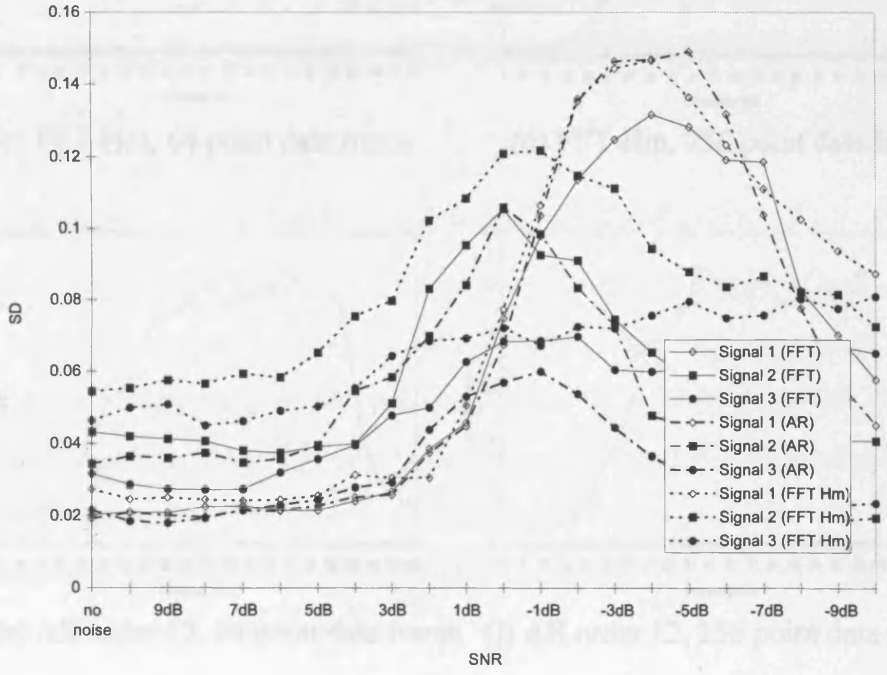

(b)

Figure 6.8 Standard deviation of SBI for three simulated Doppler signals in noise (a) 64 point data frame (b) 256 point data frame.

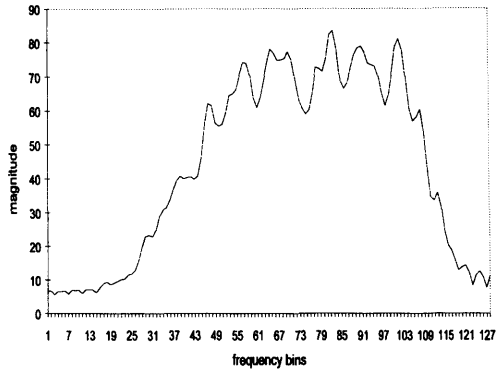

(a) FFT, 64 point data frame

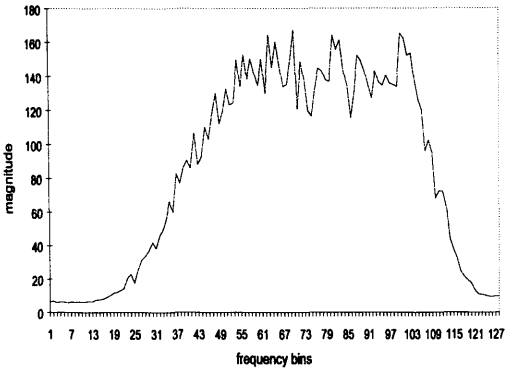

(b) FFT, 256 point data frame

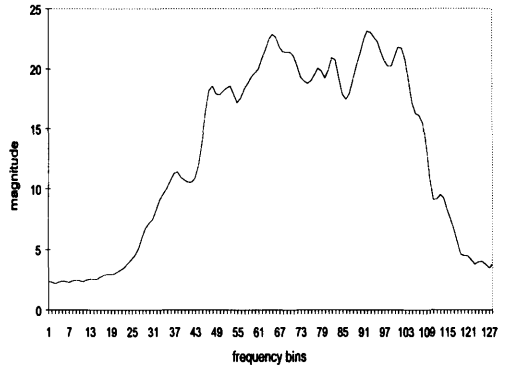

(c) FFT-Hm, 64 point data frame

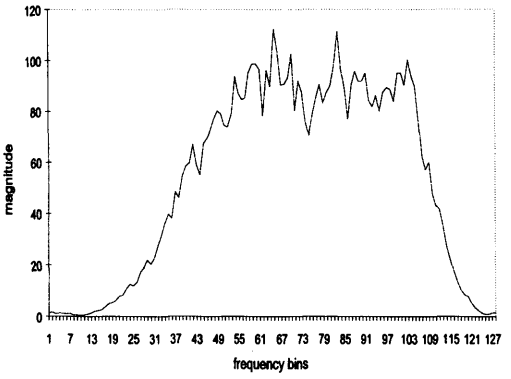

(d) FFT-Hm, 256 point data frame

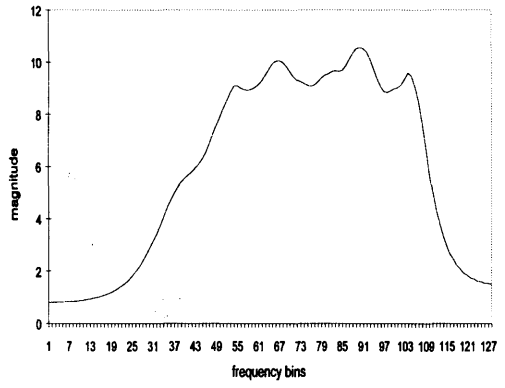

(e) AR order 12, 64 point data frame

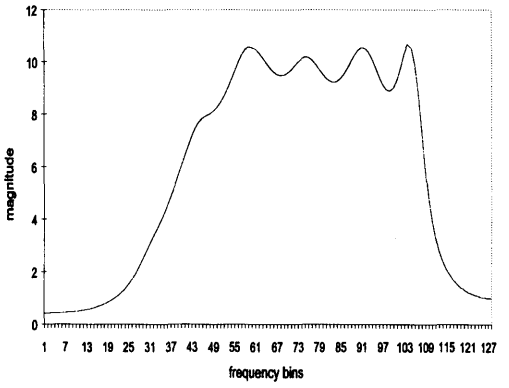

(f) AR order 12, 256 point data frame

Figure 6.9 Averaged frequency spectra for a stationary simulated Doppler signal.

When the level of noise approaches the threshold the SBI starts to increase. The SNR at which the SBI starts to drift is dependent upon the bandwidth of the simulated Doppler signal. For signals with a high degree of spectral broadening the noise will start to cause the SBI to drift at higher signal-to-noise ratios than signals with a lesser degree of spectral broadening. This is due to the larger relative amplitude

of the noise for a broad band signal compared to a narrow band signal. It can be seen from figures 6.7 and 6.8 that the magnitude and variance of the SBI are stable when the SNR is higher than 3 dB for the 256 point data frame (7 dB for the 64 point data frame). Therefore, in a normal clinical environment where the SNR is around 10-20 dB (Forsberg, 1991) the results of SBI can be considered stable. In a very noisy environment where the SNR is approaching 0 dB the results will be affected by the noise.

6.6 Conclusions

The results of this study have shown that there is a strong correlation between the estimates of SBI using both the FFT and AR algorithms. The actual value of SBI is more accurate with AR for a 256 point data frame although for both data frame sizes (64 and 256 samples) all the algorithms underestimate the true value. The use of an anti-leakage window on the FFT algorithm does not reduce the variance of the spectral broadening.

The average SBI for all algorithms is constant in low-level noise down to SNRs of 3 dB for the 256-point data frame and 7 dB for the 64 point data frame (figures 6.7 and 6.8), it is therefore possible to estimate the SBI down to these SNRs. The variance of the estimates for both FFT and AR at high signal-to-noise ratios remains fairly constant until the level of noise reaches the threshold where the SBI is measured. When the level of noise is comparable to the threshold the variance of the FFT and AR SBIs increase and the average SBI starts to increase as a result of noise contamination, as expected. The SNR at which the estimate of SBI tends to drift is dependent upon the degree of spectral broadening present in the signal. It should be noted that the magnitude of SBI tends to approximately 0.5 at low SNR.

For AR model orders greater than 10 the average value of AR-SBI is stable. For AR model orders between 10 and 16 the variance of AR-SBI is smaller than that of FFT-SBI.

Despite the fact that the results of this study have shown that there is a distinct qualitative improvement in the spectral estimates for AR compared to those of the FFT, the ability of the AR model to produce stable spectra does not result in a significant improvement in the determination of SBI over the FFT approach. For short data frames the FFT is able to produce similar SBI results to AR.

6.7 Summary

In this chapter the accuracy of the STFT and AR have been compared using a known wide band signal that was modelled on the typical spectra obtained around peak systole. It has been shown that the spectra obtained using the AR techniques are smoother than the STFT spectra and that they are closer to the true spectrum. The stability of the algorithms to produce spectra in the presence of noise was also studied to compare the stability of the SBI at various signal-to-noise ratios. In the next chapter the algorithms are tested using real clinical data.

7. Spectral broadening of clinical Doppler signals

7.1 Introduction

In this chapter ten real cases are studied to correlate the behaviour between the results obtained using the simulated data described in the previous chapter and real clinical Doppler ultrasound scans.

Real clinical data from both healthy and diseased patients was used to compare the spectra obtained using the FFT and the spectra obtained using AR modelling to give a range of Doppler signals with different degrees of spectral broadening. As for the study of simulated data described in the previous chapter the spectral broadening index (SBI), described in chapter 4, was used to quantify the broadening of the frequency spectrum. The SBI was calculated from both the FFT and the AR-based spectra at peak systole and also by averaging the SBI value of up to five spectra starting at peak systole. The correlation between both FFT-SBI and AR-SBI and the degree of stenosis is presented for 10 documented cases. In all cases more than 20 heartbeats were processed, with up to 5 spectra per heartbeat, totalling more than 1000 estimations of the systolic spectrum and SBI for each approach.

7.2 Off-line analysis

For the purposes of this comparison the real-time system described in chapter 3 was not used to analyse the Doppler signals. Instead, the system was taken off-line so that the same digital data could be analysed using both the FFT and AR algorithms. The off-line system is described in chapter 5 and again in more detail in Appendix 1. Ten Doppler ultrasound scans were digitised and stored on disk, each scan was comprised of at least 20 sequentially recorded heartbeats. Table 7.1 contains details of the recordings that were used in this study. All the recordings were taken from internal carotid arteries (ICA) of (i) patients with known diseased arteries ranging in severity of stenosis and (ii) healthy subjects. A subjective estimate of the degree of stenosis

was made at the time of the examination using colour Doppler images of the artery (both axial and cross-sectional views).

Table 7.1 Internal carotid artery (ICA) Doppler ultrasound data used in evaluation of AR against FFT.

Signal	Degree of Stenosis (Approx. % of lumen occluded)	Sample Frequency (kHz)
1	0 (Slight wall thickening)	10.24
2	0	10.24
3	20	10.24
4	25	10.24
5	40-50	10.24
6	50	10.24
7	65	5.12
8	>70	10.24
9	80	10.24
10	>80	10.24

The signals were acquired by skilled personnel at the Leicester Royal Infirmary from patients that had been admitted for colour Doppler ultrasound scans. Doppler ultrasound recordings were recorded and stored on digital audio tapes (DAT) tapes. A small sample volume positioned in the centre of the vessel was used to assess the flow through the vessel.

For patients with stenosis the site of the scan was around 2 diameters distal to the stenosis. For normal subjects the recordings were taken from a position distal from the carotid branch, which corresponded to the point where the data collection was made for the patients with stenosis. A continuous wave Doppler velocimeter, with a 4 MHz focused 1.5 mm beam width, was used for data collection. The beam transfixed the vessel axis at an angle of around 50°.

The conditions for the digital signal processing of the Doppler signals were identical to the specifications of the real-time system described in chapter 3. The spectral frame rate was set at 80 frames per second, corresponding to a data frame

length of 12.5 ms irrespective of the sampling rate. The signals were analysed using non-overlapping data frames. The choice of sampling frequency was dependent upon the observed bandwidth for each signal and is listed in Table 7.1. The signals were analysed using an straightforward 256-point FFT, the data was multiplied by a Hamming window. Three AR algorithms were implemented: the first algorithm was based on the Yule-Walker equations using the recursive Levinson-Durbin approach with model order 12 to calculate the autoregressive coefficients (a_k), the PSD was computed using [7.1]; the second algorithm again used the Levinson-Durbin recursive algorithm to calculate a_k, the PSD was calculated using [7.2]; the third AR algorithm used Burg's algorithm for calculating a_k, the PSD was computed using [7.1]. All four algorithms produced spectral estimates with 128 frequency bins.

$$PSD_{AR}(k) = \frac{T\sigma^2}{\left|1 + \sum_{n=1}^{p} a_k \exp(-j2\pi knT)\right|^2} \qquad [7.1]$$

where: T = sampling period, σ^2 = variance of the driving white noise input to the AR filter.

$$\hat{PSD}_{AR}(k) = T \sum_{n=-M}^{M-1} \hat{R}_{xx}(n) \exp(-j2\pi nkT) \qquad [7.2]$$

where: $\hat{R}_{xx}$ = estimated (extrapolated) autocorrelation sequence, for more details refer to chapter 2.

7.3 Results & Discussion

Figure 7.1(a-d) shows the output for several heartbeats using each of the four methods described (FFT, Yule-Walker AR [7.1], Yule-Walker AR using [7.2] and Burg AR) and figure 7.1(e-h) shows a typical spectrum obtained from around the systolic peak using each of the algorithms. It can be seen that there is no significant difference between the frequency spectra obtained using the different AR algorithms.

There is a distinct qualitative improvement however in the frequency spectra obtained using the AR algorithms over the FFT.

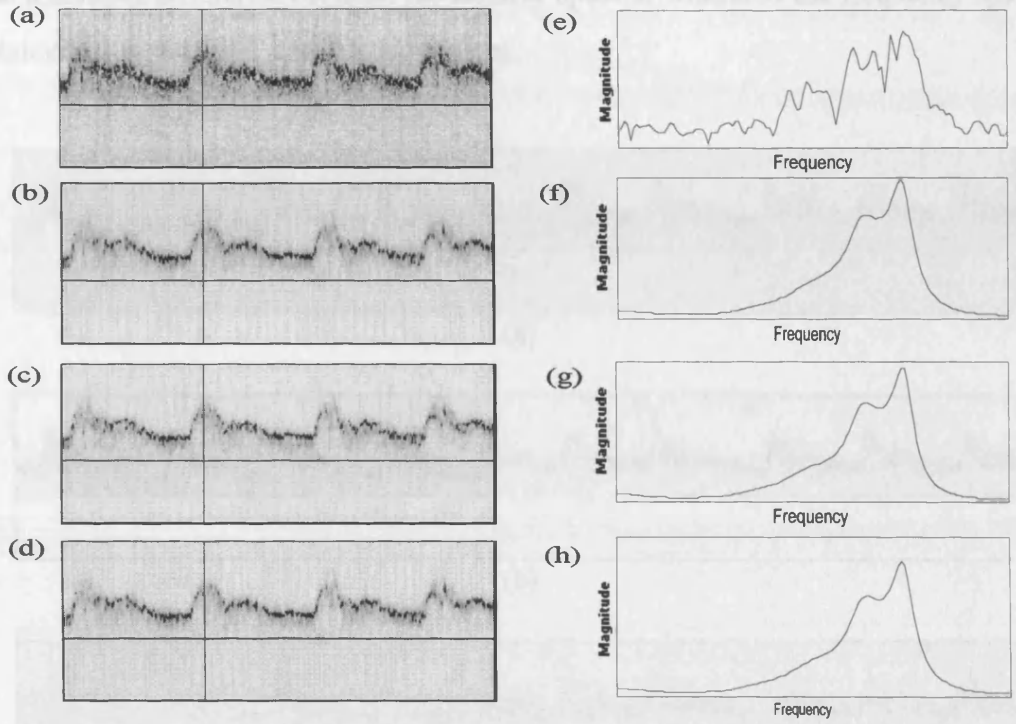

Figure 7.1 Sonograms of internal carotid arteries (ICA) using (a) FFT, (b) AR with Yule-Walker approach and eqn [7.1] for PSD, (c) AR with Yule-Walker approach and eqn [7.2] for PSD, (d) Burg and eqn [7.1] for PSD, (e)-(h) Frequency spectra taken from respective sonograms around the peak systole of the cardiac cycle (as marked with the vertical line on (a)-(d).

The -6 dB threshold described in chapter 5 was used to extract the f_{mean} and f_{max} frequency envelopes. Figure 7.2(a,b) shows the raw maximum and mean frequency envelopes extracted from the sonograms obtained using the FFT and AR algorithms respectively. The envelopes have been superimposed on the original sonograms to show the ability of the -6 dB threshold to successfully reject low-level noise and retain the true Doppler signal. The two envelopes are then low-pass filtered using a 5-point

moving-average filter. It can be seen from figure 7.2(c,d) that there are significant differences between the envelopes obtained using the FFT and AR algorithms. These differences can affect the positioning of the systolic peak which is defined as the maximum frequency component of the f_{max} envelope in each cardiac cycle. There is also a distinct difference between the relative spectral widths of the frequency spectra obtained using the FFT and AR algorithms.

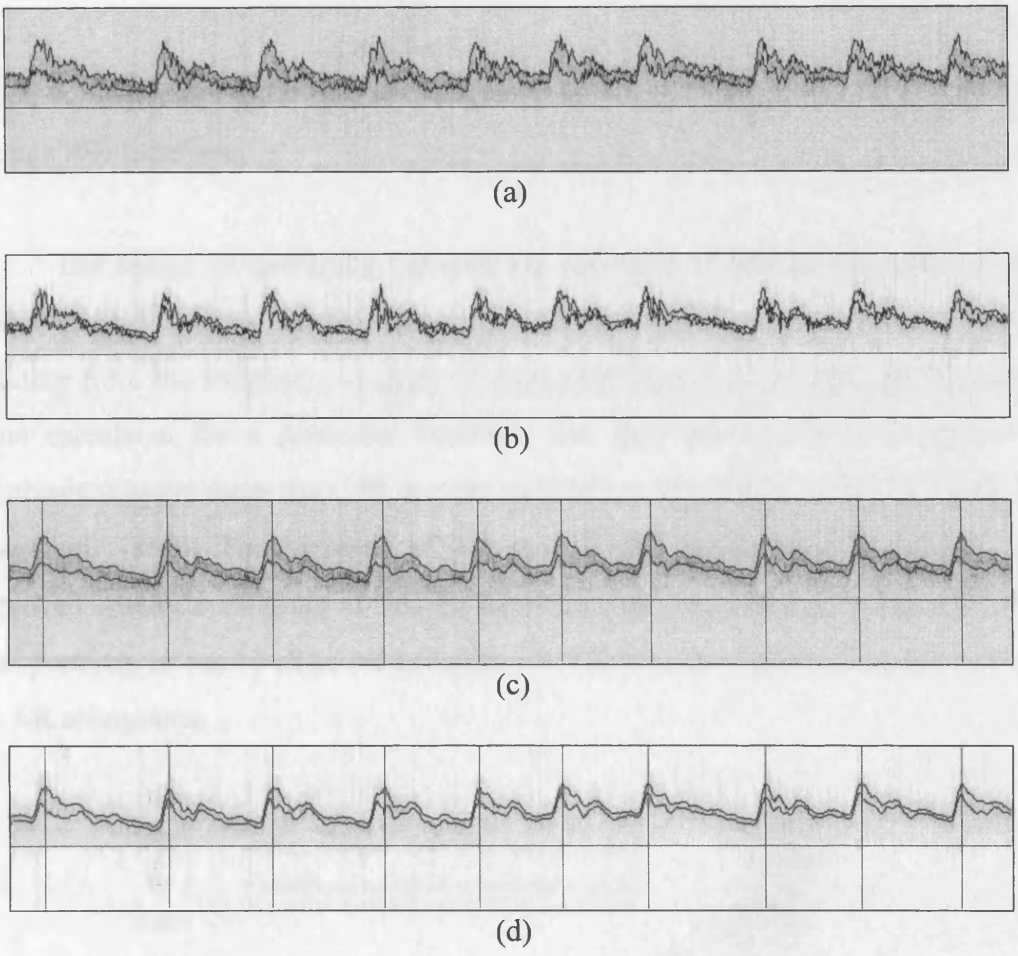

Figure 7.2 Raw maximum and mean frequency envelopes extracted from the (a) FFT (b) AR sonograms using the -6 dB threshold. Figures (c) & (d) show the low-pass filtered envelopes of (a) & (b) respectively and the comparison between the location of the systolic peaks for each heartbeat based on the FFT and AR sonograms.

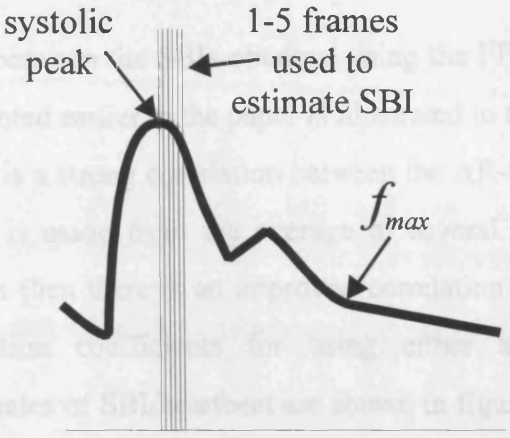

Figure 7.3 Illustrates the consecutive frames after the systolic peak used to produce an average SBI/heartbeat.

The results of averaging between 1-5 estimates of SBI around each systolic peak (figure 7.3) were studied to assess whether the effects of spectral instability resulting from the frequency analysis of short data frames can be reduced. The mean value calculated for a particular heartbeat was then averaged over a number of heartbeats (always more than 20) in order to obtain a statistically valid SBI (Labs and Fitzgerald, 1992). The variance of the overall SBI for a particular patient was improved when the estimate of SBI for each heartbeat was obtained using more than one spectrum, as can be observed in figure 7.4. This was the case for both the FFT and the AR sonograms.

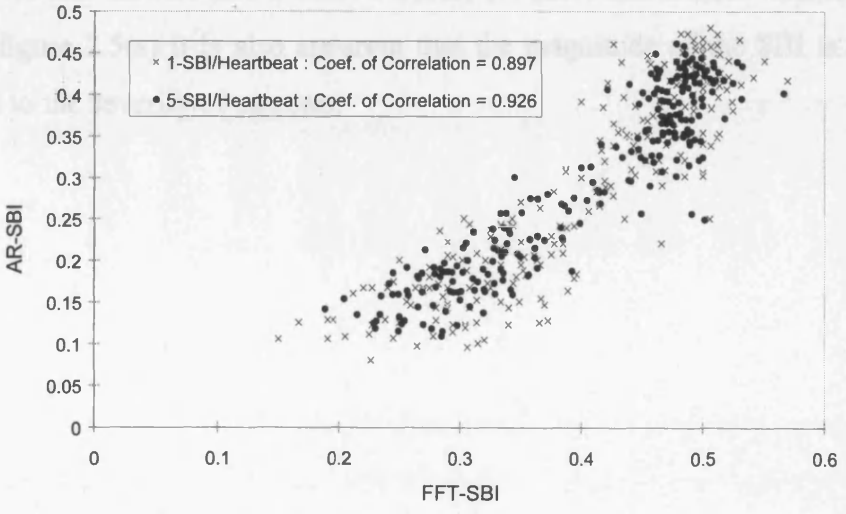

Figure 7.4 Scattergram of FFT-SBI versus AR-SBI for 1-SBI/Heartbeat and 5-SBI/Heartbeat.

The correlation between the SBIs obtained using the FFT and AR algorithms for the 10 cases documented earlier in the paper is illustrated in figure 7.4, from where it can be seen that there is a strong correlation between the AR-SBI and the FFT-SBI. If the estimate of SBI is made from the average of several estimates taken from around the systolic peak then there is an improved correlation between AR-SBI and FFT-SBI. The correlation coefficients for using either a single estimate of SBI/heartbeat or 5 estimates of SBI/heartbeat are shown in figure 7.4. The correlation coefficients for each of the comparisons exceed the 0.1% value, therefore the results are significant at the 0.1% level ($p < 0.001$).

Having established that there is a strong correlation between the FFT-SBI and the AR-SBI, the magnitude and the standard deviation of the overall SBIs calculated using 5 estimates of SBI/heartbeat are compared in figure 7.5. The magnitude of the AR-SBI is significantly smaller than that of the FFT-SBI which is concurrent with the relative spectral width trends seen in figure 7.2(c,d). The standard deviations of the AR-SBI estimates are generally larger than the FFT-SBI. The smaller magnitude of the AR-SBI results in a larger coefficient of variation as can be seen in figure 7.5(c). The SBIs calculated using the FFT algorithms are therefore generally more stable. They are not, however, necessarily more accurate, as reported by Keeton *et al.* (1997)[2] who concluded that the AR-SBI is more accurate for short data frames with simulated data. From figure 7.5(a) it is also apparent that the magnitude of the SBI is directly proportional to the severity of stenosis.

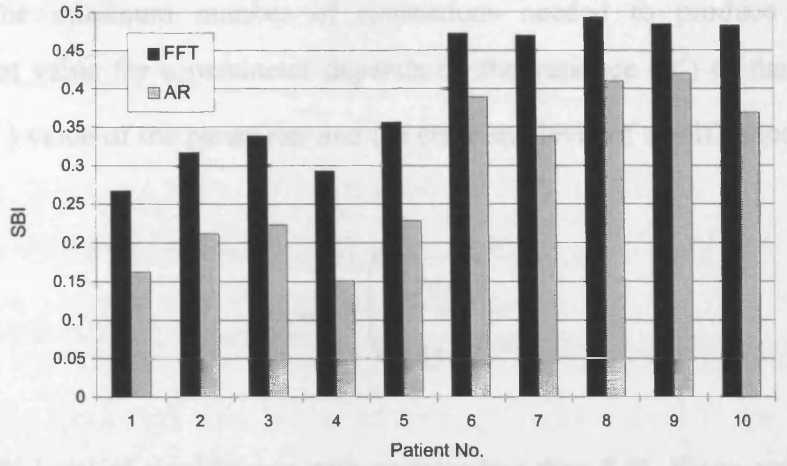

(a)

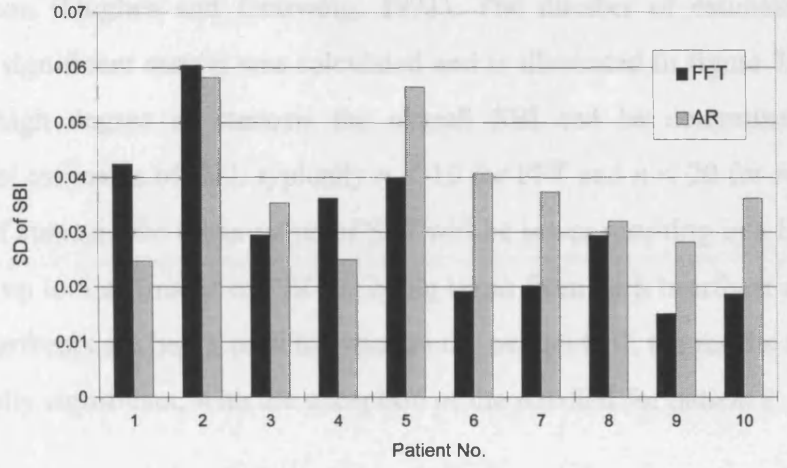

(b)

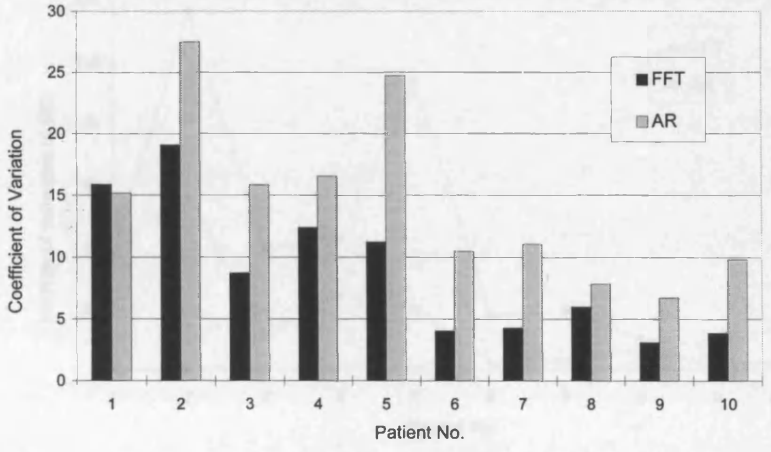

(c)

Figure 7.5 (a) Magnitude of SBI (b) Standard deviation (c) Coefficient of variation for 10 documented cases using FFT and AR sonograms based on 5 estimates of SBI/heartbeat. Increasing severity of stenosis from left to right based on a subjective study of colour Doppler ultrasound scans.

The minimum number of estimations needed to produce a statistically significant value for a parameter depends on the variance (σ^2) of the estimates, the mean ($\bar{x}$) value of the parameter and the error and level of significance desired and is given by:

$$n_{5\%,\,95\%} = \left(\frac{1.96\sigma}{0.05\,\bar{x}} \right)^2 \qquad [7.3]$$

for a 95 % level of significance with an error less than 5 %, if one assumes a normal distribution (Hughes and Graiwoig, 1971). The number of estimations needed to produce significant results was calculated and is illustrated in figure 7.6. For patients with a high degree of stenosis the overall SBI can be determined using fewer individual estimates of SBI: typically $n < 10$ for FFT and $n < 20$ for AR. For a lesser degree of stenosis the mean value of SBI will be lower resulting in a higher value for n. Since up to 5 estimates of SBI are being taken from each heartbeat and a minimum of 20 heartbeats are being used to estimate the overall SBI, the results for all cases are statistically significant, with the exception of the AR-SBI for patient 2.

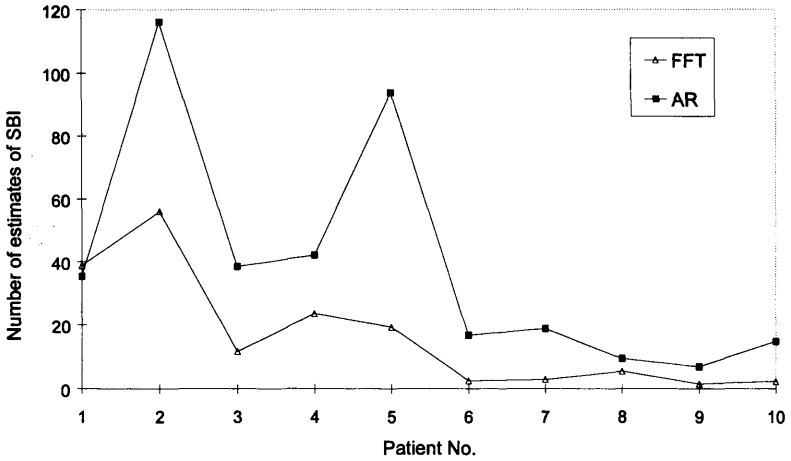

Figure 7.6 Sample size required to produce a statistically valid estimate of SBI for 10 documented cases using FFT and AR sonograms.. Results are based on 5 estimates of SBI/heartbeat, therefore, No. of heartbeats = No. of estimations/5. Increasing severity of stenosis from left to right based on a subjective study of colour Doppler ultrasound scans.

7.4 Conclusions

Based upon the results obtained for the two techniques, it is clear that the AR approach can help improve the quality of the spectra (figs. 7.1(e-h)). This is particularly true when there are too few samples for the FFT to be able to produce reasonable spectral estimates due to its inherent limitations. The AR sonograms do not suffer from some of the intrinsic problems that affect the FFT based spectral estimation and hence there is a distinct qualitative improvement in the visualisation of the dynamic flow over the cardiac cycle (figs. 7.1(a-d)). With the sample cases examined it was found that the AR algorithms all produced similar sonograms, and all the AR algorithms produce smoother sonograms than the FFT approach.

A strong correlation was observed between the value of the FFT-SBI and the AR-SBI. The mean value of the FFT-based SBI is larger than that of the AR-based SBI and the variance of the FFT-SBI is smaller than that of the AR-SBI.

If only one spectrum per heartbeat is used for the estimation then the number of heart beats required in order to produce a statistically significant value for SBI (with a 95 % level of significance) is perhaps too large for clinical practice particularly when the degree of stenosis is mild. If five spectra around the peak systole are used for the estimation then, for all cases studied here (with the exception of one AR-case), 20 heartbeats was more than adequate to produce a statistically significant value for SBI. As the severity of the stenosis increases fewer heartbeats are required.

7.5 Summary

This chapter has completed the discussion about the merits of using alternative AR techniques compared to the STFT. The remainder of this thesis is devoted to looking at alternative signal processing tools that have, as far as we know, so far not been applied to Doppler ultrasound and to discuss their potential applications.

8. Wavelets

8.1 Introduction

The previous chapters have discussed techniques that have been used in Doppler ultrasound for over a decade and has provided a quantitative assessment of the effects of using different signal processing algorithms to analyse the Doppler signal. This chapter introduces a new tool which has emerged in the signal processors toolbox and which is finding an ever increasing range of applications in biomedical engineering. Wavelets were introduced in the early 1980's and are still in the experimental stage of development. This chapter looks at ways in which the wavelet transform could be applied to Doppler ultrasound. The contents of this chapter have been published in a paper by Keeton and Schlindwein (1997)[1].

8.2 Wavelets

The Fourier transform (FT) is excellent at extracting the frequency information from stationary signals. The short-time Fourier transform (STFT) is a modified version of the FT which has been developed to perform time-frequency analysis of non-stationary signals. There are a number of limitations inherent in the STFT which have been discussed in detail in chapter 3. One of the fundamental limitations of the STFT is the selection of a fixed window for the analysis of the signal. This limits the ability of the STFT to be able to detect the different features of the signal. "*If we look at a signal (or function) through a large 'window', we would notice gross features. Similarly, if we look at a signal through a small 'window', we would notice the small features.*" (Graps, 1995). Therefore, by fixing the size of the window we are limiting the resolution with which we can study the signal. The wavelet transform addresses the problem of fixed resolution by using base functions that can be scaled. These 'wavelets' act in a similar way to the windowed complex exponentials that are used in the STFT except that with the wavelet transform the length of signal being analysed is not fixed. "*The result in wavelet analysis is to see both the forest and the trees*" (Graps, 1995). Figure 8.1 illustrates some of the more common wavelets. At first

glance it is apparent that wavelets are better suited to analysing transient signals, since they are well localised in time, whereas sinusoids extend over all time. Wavelets are not however just well localised in time, they are also well localised in frequency, although not as well as sinusoids. The property of time and frequency localisation is known as compact support and is one of the most attractive features of wavelet analysis.

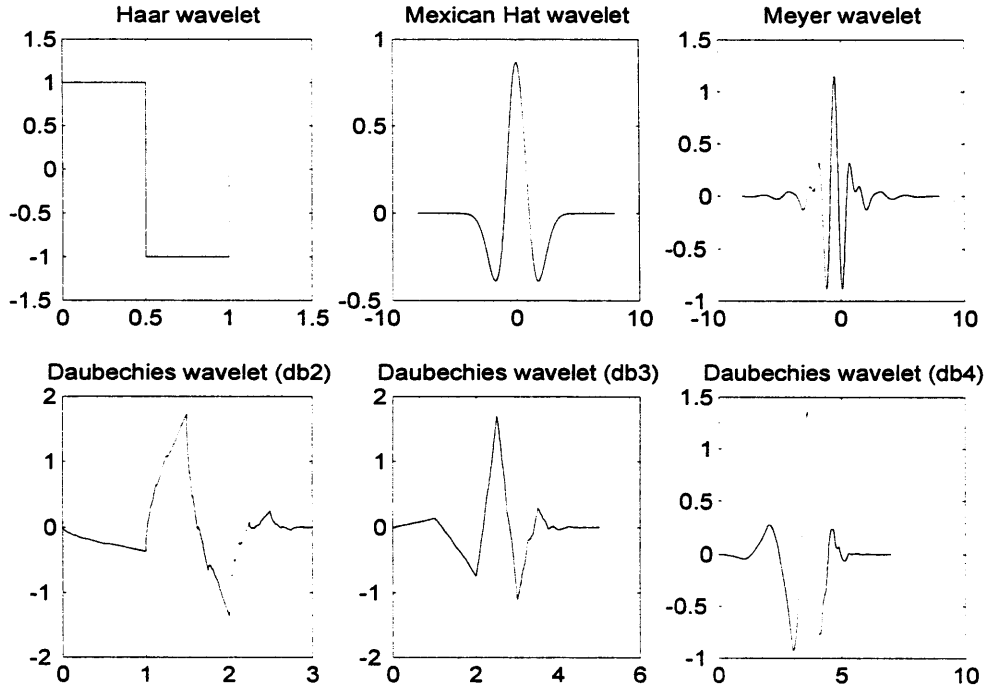

Figure 8.1 Examples of common wavelets.

Wavelets have already been used in biomedical applications for the analysis of the ECG (Gramatikov and Georgiev, 1995; Thakor and Sherman, 1995) and the EEG (Thakor and Sherman, 1995). In this research the potential of using wavelets for the analysis of blood flow is explored. The variable time-frequency resolution of wavelet analysis makes wavelets well suited to the analysis of Doppler blood flow signals: a compressed wavelet to analyse the rapid time varying frequency components around the systolic peak and a dilated wavelet to analyse the slower frequency changes during diastole. In addition, the wavelet has similarities with the shape of the Doppler waveform and offers compact support which is ideal for modelling the nonstationary characteristics of the received Doppler signal. Before looking at specific applications of wavelets to Doppler ultrasound the next section describes the wavelet transform.

8.2.1 Continuous wavelet transform

The continuous wavelet transform (CWT) [8.1] compresses or dilates and translates a 'mother' wavelet and correlates it with the signal at all times and scales.

$$CWT(a,b)=\int_{-\infty}^{\infty} x(t)\frac{1}{\sqrt{a}}\Psi\left(\frac{t-b}{a}\right)dt \qquad [8.1]$$

where: $\Psi(t)$ is the wavelet, b is the translation factor and a is the dilation factor. The output of the transform shows the correlation between the signal and wavelet as a function of time. If the signal and wavelet are a good match then the correlation between the signal and the wavelet is high resulting in a large coefficient. The choice of wavelet depends upon the application. The concept of scale in the CWT is analogous to the inverse of frequency in the FT. When the wavelet is highly compressed it extracts the localised high frequency details of the signal. When the wavelet is fully dilated the length of the wavelet is more comparable to the length of the signal and therefore it extracts the low frequency trends of the signal.

Figure 8.2(a) shows a chirp signal with a sinusoid and an instantaneous spike superimposed. Figure 8.2(b-c) shows the results of analysing the signal using the STFT and the CWT respectively. The results of the Fourier analysis illustrate the trade-off between time and frequency resolution. If the sinusoid is going to be resolved with any degree of accuracy then the time resolution of the chirp and the localisation of the instantaneous spike are compromised. If the time resolution is improved then the frequency resolution becomes inadequate to resolve the sinusoid and the chirp. The CWT pinpoints the instantaneous spike and shows the sinusoid and chirp. The CWT also shows the oscillatory behaviour in the time-scale diagram, this corresponds to the relative phase difference between the wavelet and the signal. The CWT does not resolve sinusoidal components as well as Fourier analysis, its ability to localise transient events however, makes it an attractive alternative to the STFT.

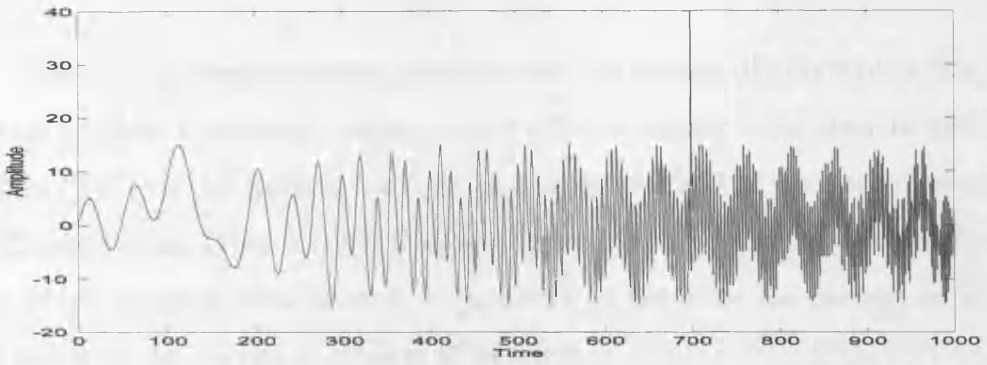

(a)

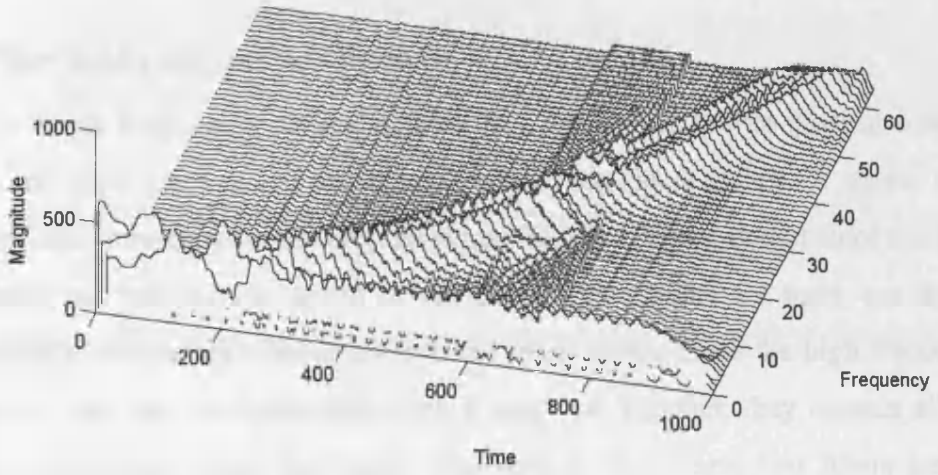

(b)

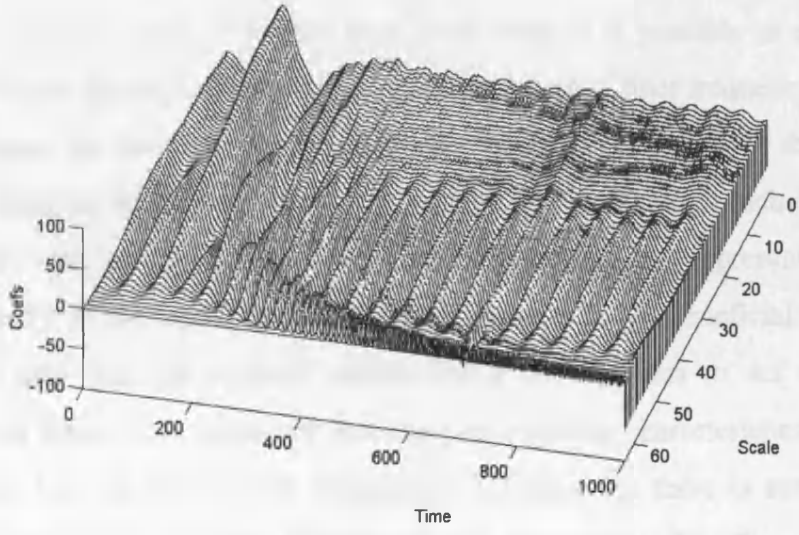

(c)

Figure 8.2 (a) Signal composed of a chirp with a sinusoid and an impulse superimposed. (b) STFT analysis using a sliding 128-point FFT. (c) CWT decomposition. Scale is the inverse of frequency.

The CWT is computationally intensive and the amount of information that the transform provides is extremely large. A more efficient routine is the discrete wavelet transform (DWT) which looks at the signal at specific dilations of the mother wavelet. The efficiency of the DWT is better than that of the FFT. To understand the DWT and to see where wavelets come from it is necessary to introduce the concept of filter banks which are the discrete equivalent of wavelets.

8.2.2 Filter banks and subband coding

A single filter has a particular frequency response and after passing a signal through the filter some of the information within the signal is lost. A single filter cannot be used to reconstruct the original signal from the filtered output since once the information has been lost it cannot be retrieved. If two filters are used, one which retains the low frequency information and the other which keeps the high frequency information they can be designed in such a way that, together, they contain all the original information within the signal. The outputs from these two filters can be combined to reconstruct the original signal. The set of filters used to split the spectral information in such a way is known as a filter bank. It is possible to extend this concept to further decomposition of the signal into finer and finer frequency bands. In order to manage the data that is output from the filters it is necessary to downsample the outputs from the filters otherwise as the signal is passed through each level of the filter bank the amount of data is doubled. Generally the idea is to represent the signal more efficiently. If the signal is decomposed n times it is not beneficial to have 2^n times more data than the original signal. For a set of filters to act as 'perfect reconstruction filters' it is necessary that they have special characteristics since real filters do not have perfect cut-off frequencies and therefore there is some overlap between adjacent filters in order to retain all the information. The filters that were developed in the early 1980's to perform this task are known as quadrature mirror filters (Strang and Nguyen, 1996).

8.2.3 Discrete wavelet transform

The discrete wavelet transform is derived from these filters and is based on a multiresolution decomposition of the signal to give a coarser and coarser approximation to the original signal by removing high frequency detail at each level of the decomposition as illustrated in figure 8.3(a). The original signal can be reconstructed by adding together the approximation and detail at the lowest level to reconstruct the approximation at the higher level, this can be repeated until the original signal is retrieved. In order to reconstruct the original signal it is therefore only necessary to keep the lowest level approximation and the detail coefficients at each level. The time-frequency representation of the decomposition is given in figure 8.3(b). It can be seen that the low frequency trends are well localised in frequency and the high frequency components are well localised in time but not in frequency.

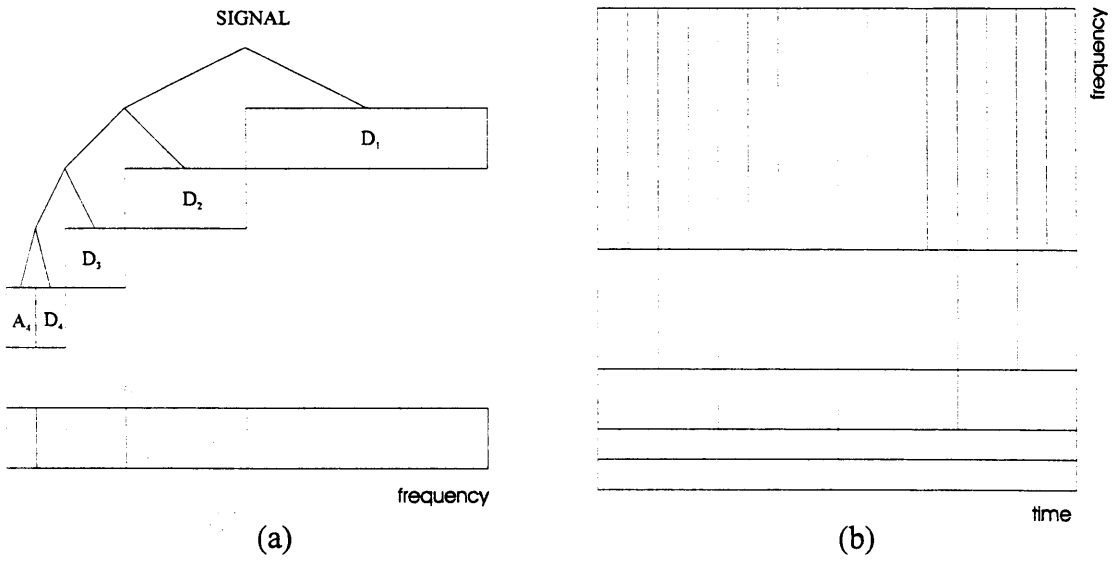

Figure 8.3 (a) DWT decomposition of a time varying signal. (b) Time-frequency representation using the DWT.

8.2.4 Wavelet packet analysis

Wavelet packet analysis is an extension of the DWT and it turns out that the DWT is only one of the many possible decomposition's that could be performed on the signal. Instead of just decomposing the low frequency component of the filter bank each time it is possible to decompose the high frequency component as well. It is

therefore possible to subdivide the whole time-frequency plane into different time-frequency tilings as can be seen from figure 8.4.

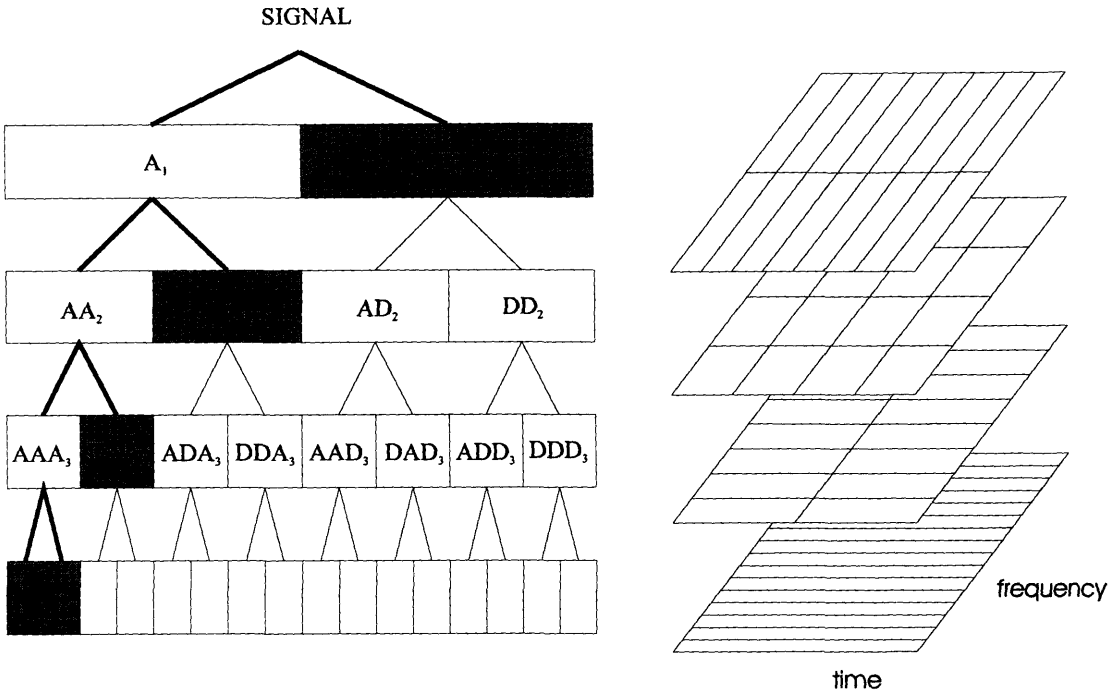

Figure 8.4 (a) Total decomposition of a time varying signal using wavelet packet analysis. (b) Time-frequency representation for each level of the decomposition.

The advantage of wavelet packet analysis is that it is possible to combine the different levels of decomposition in order to achieve the optimum time-frequency representation of the original signal.

8.3 Application of the wavelet transform to real signals

The concept of being able to totally decompose a signal and then perfectly reconstruct the signal again is nice, but it is not particularly useful by itself. In order to make use of this tool it is necessary to manipulate the wavelet coefficients to identify characteristics of the signal that were not apparent from the original time domain signal. The next section outlines how the wavelet transform could be used to analyse Doppler ultrasound signals: Firstly, the extraction of a time-frequency sonogram of

the Doppler signal using wavelet analysis compared to the conventional FFT is considered. Secondly, wavelets are used in the post-processing of the sonogram to extract information that can be used for clinical diagnosis.

8.3.1 Time-frequency analysis of the Doppler signal

Conventional time-frequency analysis of the Doppler signal involves processing the signal using data frames of fixed duration. The maximum size of the data frame is dependent upon the stationarity of the signal. The Doppler signal is non-stationary and therefore a time frame of < 20 ms is used for analysis. If a longer time frame is used then the transient behaviour of the blood flow becomes blurred. Figure 8.5(a,b) shows typical FFT sonograms for two Doppler signals, the first signal is taken from the internal carotid artery and the second from the femoral artery. The sonograms show the periodic heartbeats and within each beat it is possible to visualise the systolic and diastolic flow as the heart contracts and then relaxes. The same signals were analysed using wavelet packet analysis (db3 wavelet, see figure 8.1), the results are shown in figure 8.5(c,d).

From figure 8.5 it can be seen that the time-frequency decomposition using wavelet packet analysis produces a representation of the signal that resembles the STFT sonogram. The difference between the sonogram obtained using the STFT and the wavelet packet time-frequency representation is that wavelet packet analysis does not analyse the signal with a fixed resolution as is the case for the STFT. Instead a flexible time window is used which enables efficient modelling of the signal based on a minimum energy criterion which allows optimisation of the time-frequency tiling. The wavelet decomposition coefficients contain all the energy of the original signal and therefore the time-frequency decomposition can be used to perfectly reconstruct the original signal, something which cannot always be achieved using the STFT. The low energy coefficients can be removed from the signal using thresholding techniques to improve the quality of the time-frequency decomposition since they do not contribute a significant amount to the original signal.

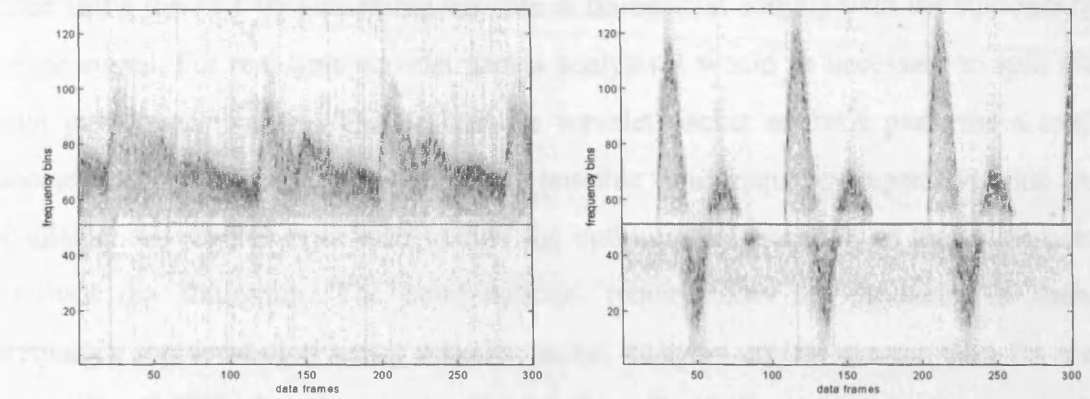

(a) (b)

(c)

(d)

Figure 8.5 Time-Frequency decomposition of Doppler ultrasound signals. (a) STFT - internal carotid artery. (b) STFT - femoral artery. (c) Wavelet packet analysis - internal carotid artery. (d) Wavelet packet analysis - femoral artery.

Doppler ultrasound systems process the acquired data in real-time. This is done using the FFT by processing the data in frames that comply with the stationarity of the signal. For real-time wavelet packet analysis it would be necessary to split the data into longer frames. This is because wavelet packet analysis performs a total decomposition on a signal so that all the possible time-frequency representations are available. An algorithm is used to find the optimum combination of these levels to produce the sonogram. The computational requirements for producing a time-frequency representation using wavelet packet analysis are far greater than for the conventional FFT, therefore the implementation of real-time wavelet packet analysis would not be straightforward.

The effect of analysing such a large time segment is that the memory requirements for storage of the decomposed data becomes extremely large. The results shown in figure 8.5 were generated using MATLAB on an IBM PC-computer with 20 Mb RAM. It was only possible to obtain 5 levels of decomposition which is equivalent to 32 (2^5) frequency bins. This resolution is 4 times worse than that available using the STFT and therefore it is not possible to make a direct comparison between the two time-frequency representations.

Despite the practical limitations imposed by wavelet analysis the results show interesting features which warrant further investigation. Wavelet packet analysis provides a novel way to model the Doppler signal compared with conventional spectral estimators such as the STFT. As far as we know this is the first time that wavelets have been used for the joint time-frequency analysis of Doppler signals.

8.3.2 Post-processing of the time-frequency decomposition

Following time-frequency analysis of the Doppler signal a clinical diagnosis is made to determine whether or not the vessel is diseased. This decision is based on a subjective study of the sonogram and further quantitative analysis. One kind of quantitative analysis is performed using frequency envelopes extracted from the sonogram. The envelopes are normally obtained by selecting a threshold to reject low

level noise components. Figure 8.6 shows the extracted mean and maximum frequency envelopes for the two examples over a number of cardiac cycles.

8.3.3 Denoising the frequency envelopes

The raw frequency envelopes extracted from a sonogram are usually corrupted by noise (figure 8.6) due to statistical instability in the frequency spectrum. To use the envelopes for quantitative analysis it is necessary to remove the noise to identify the underlying trends. Conventionally a 5-point moving average filter is used to denoise the envelopes as illustrated in figure 8.6. Discrete wavelet analysis can be used to clean the noisy envelopes. A signal can be represented efficiently using a relatively small number of wavelet coefficients assuming an appropriate selection of wavelet is made. It is therefore possible to discard small wavelet coefficients that do not contain a significant amount of information about the overall signal. The coefficients that are kept are used to reconstruct the original signal. Figure 8.6 shows the effect of denoising the raw envelopes using this technique. This wavelet-based adaptive approach has advantages over a fixed filter because the wavelet denoising process does not indiscriminately attenuate specific frequencies. Wavelet denoising assumes that the signal-to-noise ratio (SNR) is significantly large, that is, the noise coefficients are small compared to the coefficients used to model the signal. This seems to be the case for the removal of noise from the frequency envelopes as figure 8.6 shows.

8.3.4 Compression

Computer storage of patient data is invaluable since data can be transferred electronically in a few seconds. Efficient compression of data ensures minimal storage space and quicker data transfer. At present storage of Doppler ultrasound recordings is made by producing a grey scale hard copy of several cardiac cycles. By using 2-D wavelet image compression it may be possible to digitally store recorded sonograms in a patient database for future reference using much less storage space. High compression ratios have been achieved using wavelets for image compression and the FBI are using wavelets to reduce storage space required to store their database of fingerprints.

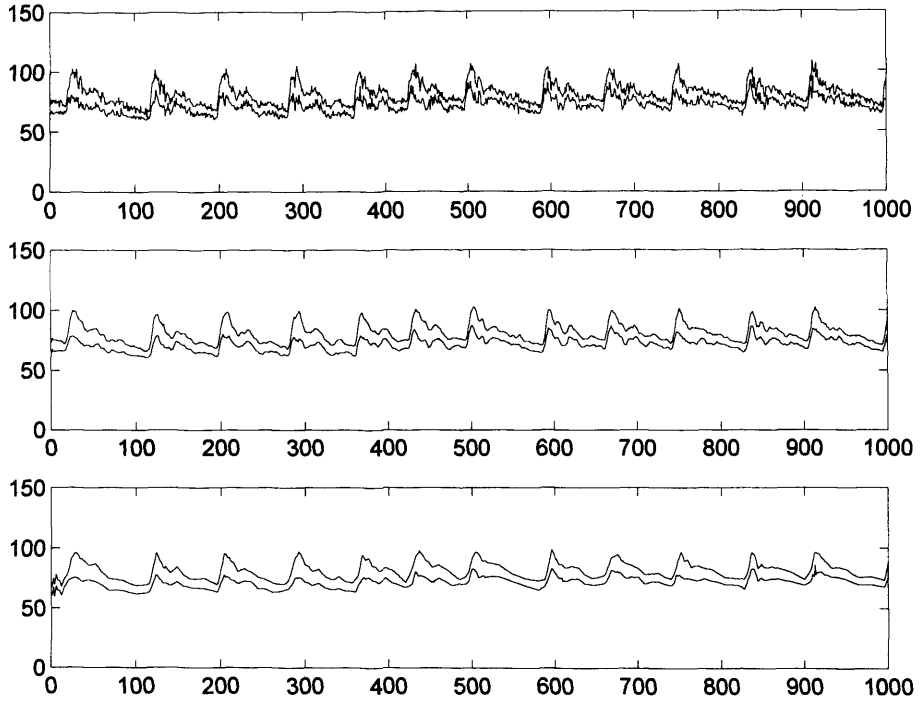

(a) **Top**: Raw envelopes. **Middle**: 5-point MA filtered envelopes. **Bottom**: Wavelet denoised envelopes.

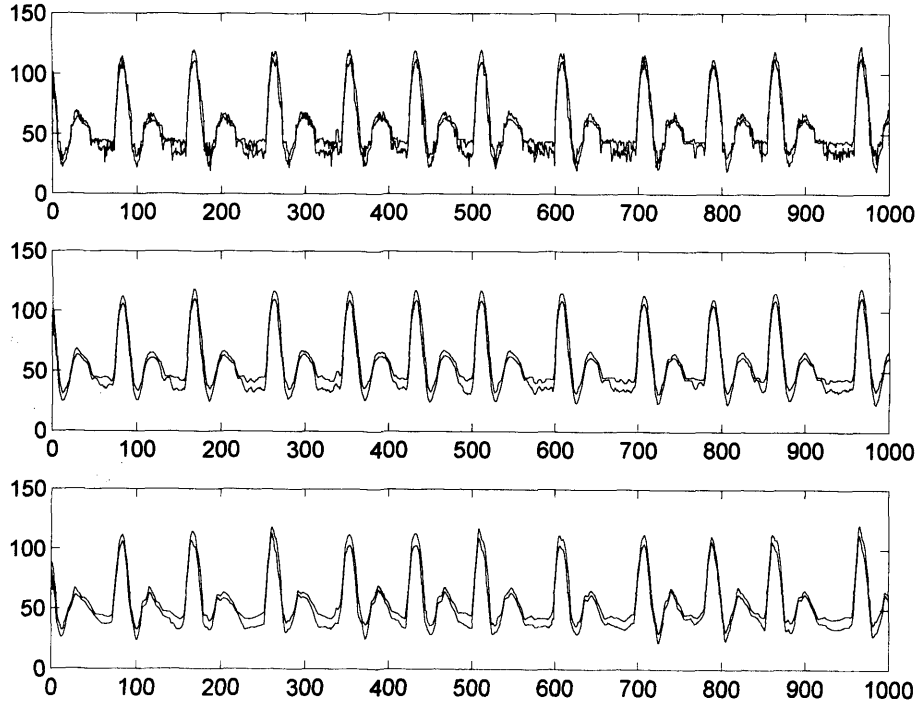

(b) **Top**: Raw envelopes. **Middle**: 5-point MA filtered envelopes. **Bottom**: Wavelet denoised envelopes.

Figure 8.6 Envelopes extracted from the STFT-sonograms. (a) Internal carotid artery. (b) Femoral artery.

8.4 Conclusions

Wavelet analysis provides an interesting alternative to conventional Fourier methods. The compact support inherent in wavelets makes them an attractive choice for time-frequency analysis. Fourier analysis has a fixed time-frequency resolution. This means that there is a trade-off between the resolution of transient events and underlying trends in the signal. Wavelets are better suited to resolving transient events since wavelet analysis incorporates the concept of scale into the equation, which gives the analyst the flexibility to look at the time domain signal at different resolutions: A compressed wavelet for analysing high frequency detail and a dilated wavelet for detecting underlying trends.

Wavelet packet analysis of the Doppler signal can provide a time-frequency decomposition similar to that of the FFT. The advantage of wavelet packet analysis over the FFT is the ability to optimise the time-frequency decomposition of the sonogram. The computational overheads for wavelet packet analysis are far greater than for FFT analysis of the Doppler signal making real-time implementation difficult.

The DWT can be used to denoise the frequency (blood velocity) envelopes extracted from the sonograms by manipulating the wavelet coefficients. Large wavelet coefficients contain significant information about the signal while small coefficients are expendable. If a signal is contaminated with noise and the SNR is high, then it is possible to remove the noise by implementing a threshold to remove small coefficients.

8.5 Summary

This chapter has presented the fundamental ideas behind continuous and discrete wavelet analysis. Examples of the application of wavelets to Doppler ultrasound were given and potential areas for further research have been outlined. This concludes the various aspects of the research that has been carried out. The next chapter provides a summary of all the key-points that have been raised during the course of this project.

9. Conclusions

This thesis can be separated into three distinct sections: In the first section real-time frequency analysis was addressed and the implementation of a real-time system was discussed; The second section extensively examined the relationship between classical and modern spectral estimators for analysis of simulated and clinical Doppler signals; The third section looked at the wavelet transform which is a relatively new signal processing tool in the signal analyst's toolbox. The potential of using the wavelet transform in the field of diagnostic vascular ultrasound was investigated and suggestions as to its potential areas of application were put forward. This chapter serves to bring together the different aspects of the research that have been covered during the course of this research project and to highlight the key points that have been observed in the preceding chapters.

In chapter 3 a real-time system was developed using a DSP32C DSP board combined with a PC 486 DX2 66 MHz to compare the implementation of the FFT and the AR algorithms. The system mixed forward and reverse flow signals onto a single signal centred around a suppressed programmable carrier frequency via an independent analogue heterodyne mixer unit. The output was displayed on the computer screen in a sonogram format. Initially the FFT system was used to contrast the speed of implementing the system in 'DSP32C assembly' or 'C'. It was found that the assembly routines were several times quicker than the same routines written using 'C'. The real-time FFT analyser developed in 'C' was however sufficiently quick to enable more sophisticated algorithms to be implemented. The development time for programs written in 'C' is much quicker than for programs written in 'DSP32C assembly' and therefore the AR routines were written in 'C'. The system had a variable sampling frequency that could be selected by the operator and included 5.12, 10.24, 20.48 and 40.96 kHz sample rates. It was found that if a model order of 8 was used then the system could operate in real-time at any of the sampling frequencies. For a higher model order (10 or 12) real-time operation was compromised at the maximum sampling frequency. If real-time operation is desired for a model order

greater than 8 it would be necessary to develop some of the functions in DSP32C assembly.

Chapter 4 looks at how the joint time-frequency analysis of the Doppler signal can be used to aid in the diagnosis and the quantification of stenosis. The methods used in this research have been based around the spectral broadening of the frequency spectrum around the systolic peak. The measure of spectral broadening was estimated using the spectral broadening index (SBI). In chapter 5 a post-processing system was developed to enable calculation of the SBI for a series of sequentially recorded heartbeats. A threshold was used to reject low level noise and to enable the extraction of frequency envelopes from which the SBI could be computed. Two types of threshold were examined, the simple global threshold was found to create ambiguity in the choice of threshold since there was no plateau in the SBI versus threshold plot where the SBI could be considered constant. This was found to be due to the trends in the mean and maximum frequency envelopes and was consistent with results found in other published research. A second threshold was used which was set at a level relative to the maximum component in the individual frequency estimates. The threshold was tested at -3, -6, -9, -12, -15, -18 dB levels. The optimum threshold is defined as being the level at which the background noise is rejected with maximum retention of the true Doppler signal. The threshold was tested using sonograms generated using the FFT and AR algorithms, it was found that the AR spectra were more robust than the same spectra obtained using the FFT. For comparison of the relative magnitudes of the SBI estimated from the FFT and AR sonograms the threshold was set at the -6 dB level.

In chapter 6 the behaviour of the spectral broadening index (SBI) derived from wide band spectra obtained using autoregressive modelling (AR) compared to the SBI based on the fast Fourier transform (FFT) spectra was investigated. Simulated Doppler signals were created using white noise and shaped filters to analyse spectra typically found around the systolic peak and to assess the magnitude and variance of AR and FFT-SBI for a range of signal-to-noise ratios. The results of the analysis show a strong correlation between the indices calculated using the FFT and AR algorithms.

Despite the qualitative improvement of the AR spectra over the FFT the estimation of SBI for short data frames is not significantly improved using AR.

Chapter 7 investigates the behaviour of the spectral broadening index (SBI) derived from spectra obtained using autoregressive (AR) modelling compared to that of SBI based on fast Fourier transform (FFT) analysis of clinical Doppler ultrasound scans. Doppler signals from internal carotid arteries of patients with normal and diseased vessels with up to 80% stenosis were analysed. A threshold at -6 dB of the maximum magnitude component of each individual spectrum was implemented to reject low-level noise. The SBI was obtained using the maximum and the mean frequency envelopes extracted from the sonogram. A qualitative improvement in both the appearance of the AR sonograms and the shape of the individual AR spectra was noticeable. The AR approach consistently produced narrower spectra than the FFT and the shapes of the frequency envelopes derived from the AR sonogram and the FFT sonogram were also rather different. Despite these differences a strong correlation was observed between the value of the FFT-based SBI and the AR-based SBI. The mean value of the FFT-SBI is larger than that of the AR-SBI and the variance of the FFT-SBI is smaller than that of the AR-SBI based on a set of at least 20 sequentially recorded heartbeats. It was established that, for all cases, a statistically significant value for SBI could be obtained using 20 heartbeats if five spectra around the peak systole were used to estimate the SBI of an individual heartbeat.

The results obtained using the simulated data and real clinical data showed a close correlation in the comparison of the relative magnitudes of the FFT-SBI and the AR-SBI.

Chapter 8 gives an introduction to wavelets and illustrates their application with two examples. The wavelet transform provides the analyst with a scaleable time-frequency representation of the signal, which may uncover details not evidenced by conventional signal processing techniques. The signals used in this research were Doppler ultrasound recordings of blood flow velocity taken from the internal carotid artery and the femoral artery. Firstly, it was shown how wavelets could be used as an

alternative signal processing tool to the short-time Fourier transform for the extraction of the time-frequency distribution of Doppler ultrasound signals. Secondly, wavelet-based adaptive filtering is implemented for the extraction of maximum blood velocity envelopes in the post processing of Doppler signals.

9.1 Future work

The real-time system was developed on a DSP board that has been around since 1988. There is scope to improve the existing algorithms by replacing the 'C' programs with programs written in 'DSP32C assembly'. Alternatively the system could be implemented using a faster chip in order to allow complete real-time analysis of the Doppler signal across the range of expected Doppler shifts.

In chapter 5 an automatic heartbeat separation algorithm was described, the results of the algorithm showed limited success for very noisy signals and therefore more work is needed to produce a totally automatic and robust algorithm.

The results in this research for the comparison of clinical data used data from recordings with a small sample volume. It has been noted in several papers that mild stenosis do not have a dramatic affect on the whole velocity profile and only affect the flow patterns close to the vessel wall. It would therefore be interesting to compare the results of the small sample volume recordings with those taken using a large sample volume.

The programs developed during the course of this research have all been MS DOS based applications. The implementation of a Windows based interface would improve the graphical user interface.

This research has also investigated the possibility of applying wavelets to diagnostic vascular ultrasound. Some potential applications have been suggested in this research which warrant further research to investigate the benefits of the wavelet transform. In particular the investigation of the compression of the Doppler sonogram.

Some systems only keep a hard copy of the examination which allows a limited number of heartbeats to be recorded. A few commercial systems store the results onto disk for future examination. Wavelets have been found to be particularly good at compressing one dimensional signals and two dimensional images and therefore efficient storage of recorded sonograms could be used to help provide a manageable patient database for progressive assessment of arterial disease.

References

Akaike H. A new look at the statistical model identification. IEEE Trans. Automatic Control. 1974; 19:716-723.

AT&T Microelectronics. WE® DSP32C Digital Signal Processor - Information Manual, 1990.

AT&T Microelectronics. WE® DSP32 and DSP32C Application Software Library - Reference Manual, 1991.

Beach KW, Phillips DJ. Sensitivity and precision of fast Fourier transform spectral waveform analysis in mild carotid atherosclerotic disease. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. Edward Arnold, London. 1992; 6:57-68.

Blackshear WM, Phillips DJ, Thiele BD, *et al.* Detection of carotid and occlusive disease by ultrasonic imaging and pulsed Doppler spectrum analysis. Surgery. 1979; 86:5:698-706.

Born GVR. Pathogenesis and properties of atherosclerotic lesions as exemplified by carotid plaques. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. Edward Arnold, London. 1992; 3:23-32.

Brown PM, Johnston KW, Kassam M, Cobbold RSC. A critical study of ultrasound Doppler spectral analysis for detecting carotid disease. Ultrasound in Med. & Biol. 1982; 8:5:515-523.

Cardoso JCS, Ruano MG, Fish PJ. Nonstationarity broadening reduction in pulsed Doppler spectrum measurements using time-frequency estimators. IEEE Trans. Biomed. Eng. 1996; 43:12:1176-1186.

Carola R, Harley JP, Noback CR. Human anatomy & Physiology, International Edition. Second Edition. McGraw-Hill, Inc. New York. 1992.

David JY, Jones SA, Giddens DP. Modern spectral analysis techniques for blood flow velocity and spectral measurements with pulsed Doppler ultrasound. IEEE Trans. Biomed. Eng. 1991; 38:6:589-596.

Douville Y, Johnston KW, Kassam M. Determination of the hemodynamic factors which influence the carotid Doppler spectral broadening. Ultrasound in Med. & Biol. 1985; 11:3:417-423.

Evans DH. A pulse-foot seeking algorithm for Doppler ultrasound waveforms. Clin. Phys. Physiol. Meas. 1988; 9:3:267-271.

[1]Evans DH, McDicken, WN, Skidmore, R and Woodcock, JP. Doppler Ultrasound - Physics, Instrumentation, and Clinical Applications. John Wiley & Sons Ltd., Chichester, 1989.

[2]Evans DH, Schlindwein FS, Levene MI. An automatic system for capturing and processing ultrasonic Doppler signals and blood pressure signals. Clin. Phys. Physiol. Meas. 1989; 10:3:241-251.

[1]Fan L, Evans DH. Extracting instantaneous mean frequency information from Doppler signals using the Wigner distribution. Ultrasound in Med. & Biol. 1994; 20:5:429-443.

[2]Fan L, Evans DH. A real-time and fine resolution analyser used to estimate the instantaneous energy distribution of Doppler signals. Ultrasound in Med. & Biol. 1994; 20:5:445-454.

Forsberg F. On the usefulness of singular value decomposition—ARMA models in Doppler ultrasound. IEEE Trans. Ultrason. Ferroelec. Freq. Contr. 1991; 38:5:418-428.

Gibbons DT, Evans DH, Barrie WW, Cosgriff PS. Real-time calculation of ultrasonic pulsatility index. Med. & Biol. Eng. & Comput. 1981; 19:28-34.

Gosling RG, King DH. Continuous wave ultrasound as an alternative and complement to X-rays in vascular examinations. In: Reneman RS, ed. Cardiovascular applications of ultrasound. North-Holland Publishing Co. - Amsterdam - London. 1974; 266-285.

Gramatikov B, Georgiev I. Wavelets as alternative to short-time Fourier transform in signal-averaged electrocardiography. Med. & Biol. Eng. & Comput. 1995; 33:482-487.

Graps A. An introduction to wavelets. IEEE. Computational Science & Eng. 1995; 50-61.

Guo Z, Durand LG, Lee HC. Comparison of time-frequency distribution techniques for analysis of simulated Doppler ultrasound signals of the femoral artery. IEEE Trans. Biomed. Eng. 1994; 41:4:332-342.

Harris FJ. Windows, harmonic analysis and the discrete Fourier transform. Naval Undersea Center, San Diego, CA 92132. 1976.

[1]Harward TRS, Bernstein EF, Fronek A. Continuous-wave versus range-gated pulsed Doppler power frequency spectrum analysis in the detection of carotid arterial occlusive disease. Ann. Surg. 1986; 204:1:32-37.

[2]Harward TRS, Bernstein EF, Fronek A. Range-gated pulsed Doppler power frequency spectrum analysis for the diagnosis of carotid arterial occlusive disease. Stroke. 1986; 17:5:924-928.

Hedrick WR, Hykes DL, Starchman DE. Ultrasound Physics and Instrumentation. Third Edition, Mosby, 1995.

Hughes A, Graiwoig D. Statistics: A foundation for analysis. Addison-Wesley, London, 1971.

Hussey M. Diagnostic ultrasound. An introduction to the interactions between ultrasound and biological tissues. Blackie & Son Ltd., Glasgow, 1975.

Jones SA. Fundamental sources of error and spectral broadening in Doppler ultrasound signals. Critical Reviews in Biomed. Eng. 1993; 21:5:399-483.

Kaluzynski K. Analysis of application possibilities of autoregressive modelling to Doppler blood flow signal spectral analysis. Med. & Biol. & Comput. 1987; 25:373-376.

Kaluzynski K. Order selection in Doppler blood flow signal spectral analysis using autoregressive modelling. Med. & Biol. & Comput. 1989; 27:89-92.

Kaluzynski K, Palko, T. Effect of method and parameters of spectral analysis on selected indices of simulated Doppler spectra. Med. & Biol. Eng. & Comput. 1993; 31:249-256.

Kay SM. The effects of noise on the autoregressive spectral estimator. IEEE Trans. Acoust., Speech, Signal Processing. 1979; 27:5:478-485.

Kay SM, Marple Jr. SL. Spectrum analysis - a modern perspective, Proc. IEEE. 1981; 69:1380-1419.

[1]Keeton PIJ, Schlindwein FS. Application of wavelets in Doppler ultrasound. Sensor Review. 1997; 17:1:38-45.

[2]Keeton PIJ, Schlindwein FS, Evans DH. A study of the spectral broadening of simulated Doppler signals using FFT and AR modelling. Ultrasound in Med. & Biol. 1997; **in press.**

Kitney RI, Giddens DP. Linear estimation of blood flow waveforms measured by Doppler ultrasound, In: Salamon R, Blum B, Jørgensen M, eds. MEDINFO 86. North Holland: Elsevier Science Publishers B.V. 1986; 5:672-677.

Labs KH, Fitzgerald DE. Quantification of pulsed Doppler spectra for the diagnosis of minor to moderate atherosclerotic lesions: experience from in vitro and in vivo models. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. Edward Arnold, London. 1992; 13:126-141.

Loughborough Sound Images Ltd. DSP32C PC System Board - User Manual. Issue 2.02, 1991.

Marple Jr. SL. Resolution of conventional Fourier, autoregressive, and ARMA methods of spectrum analysis. Proceedings of the IEEE international conference on ASSP. Hartford, Conn. IEEE. 1977; 74-77.

Marple Jr. SL. Digital spectral analysis with applications. Prentice-Hall, Englewood Cliffs, N.J., 1987.

McDicken WN, Diagnostic ultrasonics. Principles and use of instruments. Third Edition. Longman Group UK Ltd., 1991.

Nichols WW, O'Rourke MF. McDonald's Blood flow in arteries, theoretical, experimental and clinical principles, Third Edition. Edward Arnold, London, 1990.

Oppenheim AV, Schafer RW. Digital signal processing. Prentice-Hall, Englewood Cliffs, N.J., 1975.

Oppenheim AV, Schafer RW. Discrete-time signal processing. Prentice-Hall, Englewood Cliffs, N.J., 1989.

Pourcelot L. Diagnostic ultrasound for cerebral vascular diseases. Present and future of diagnostic ultrasound. 1976; 141-147.

Qian S, Chen D. Joint Time-Frequency Analysis. Methods and Applications. Prentice-Hall, N.J., 1996.

Reneman RS, Van Merode T, Hoeks APG. Flow patterns and arterial wall dynamics in normal and diseased arteries. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. Edward Arnold, London, 1992; 5:42-54.

Rittgers SE, Thornhill BM, Barnes RW. Quantitative analysis of carotid artery Doppler spectral waveforms: Diagnostic value of parameters. Ultrasound in Med. & Biol. 1983; 9:3:255-264.

Satomura S, Study of the flow patterns in peripheral arteries by ultrasonics. J. Acous. Soc, Japan. 1959; 15:151-158.

Schlindwein FS, Evans DH. A real-time autoregressive spectrum analyzer for Doppler ultrasound signals. Ultrasound in Med. & Biol. 1989; 15:3:263-272.

Schlindwein FS, Evans DH. Selection of the order of autoregressive models for spectral analysis of Doppler ultrasound signals. Ultrasound in Med. & Biol. 1990; 16:1:81-91.

Schlindwein FS, Vieira MH, Vasconcelos CFM, Simpson DM. Real-time digital signal processing of Doppler ultrasound signals and calculation of flow parameters. Medical Progress through Technology. 1994; 20:81-89.

Schlindwein FS, Keeton PIJ, Dryden DJ. A simple modulator and mixer for directional Doppler signals. Revista Brasileira de Engenharia, RBE - Caderno de Engenharia Biomédica. 1996; 12:103-111.

Scott GC, Bogen DK, Korostoff E. FFT performance in the presence of noise. IEEE Trans. Biomed. Eng. 1987; 34:6:424-429.

Smith JL, Evans DH, Fan L, Thrush AJ, Naylor AR. Processing Doppler ultrasound signals from blood-borne emboli. Ultrasound in Med. & Biol. 1994; 20:5:455-462.

Strandness Jr. DE. The gold standard in the diagnosis of vascular disease. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. Edward Arnold, London, 1992; 1:3-11.

Strang G, Nguyen T. Wavelets and filter banks. Wellesley-Cambridge Press, 1996.

Thakor NV, Sherman D. Wavelet (Time-Scale) analysis in biomedical signal processing. In: Bronzino JD, ed. The Biomedical Engineering handbook. CRC Press, Inc. 1995; 886-906.

Vaitkus PJ, Cobbold RSC, Johnston KW. A comparative study and assessment of Doppler ultrasound spectral estimation techniques part II: Methods and results. Ultrasound in Med. & Biol. 1988; 14:8:673-688.

Welch PD. 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. 1967; 15:70-73.

Appendix 1

COMPUTER-BASED SPECTRUM ANALYSER

The following section describes the different options that were incorporated into the program to perform different types of analysis:

- *Input and output filenames* - The operator inputs the name of the input file and the type of data that is stored in the file (integers or floating-point numbers) and specifies the output filename to which the results are going to be stored.
- *Data frame length* - The number of samples used in each frame to estimate the frequency spectrum.
- *Spectral estimation algorithm* - Four algorithms were implemented these are described in more detail below:

Fourier-based spectral estimation

1. Fast Fourier transform
    - ◊ *FFT size* - Select size of FFT this can be larger than the *data frame size* to allow for zero padding.

AR-based spectral estimation

(I) Determination of AR coefficients

(II) Estimation of PSD using coefficients in (I)

2. (I) Yule-Walker (II) PSD obtained using equation [2.12].
3. (I) Burg (II) PSD obtained using equation [2.12].
4. (I) Yule-Walker (II) PSD obtained using equation [2.21].
    - ◊ *Model order* - Select model order of AR process.
    - ◊ *Extrapolation* - Number of lags that the ACS should be extrapolated up to before mirroring.

- *Anti-leakage window* - Either a rectangular or a Hamming window can be chosen.
- *Averaging* - Number of data frames that should be averaged before displaying the output.
- *Overlap* - The number of samples that the data frames should be overlapped by.

  '0' - The consecutive data frames are side by side and have no overlap.

  'Positive' - The consecutive data frames have a gap between them.

  'Negative' - The consecutive data frames overlap.

  $\Rightarrow$ *Increment = overlap + data frame size.*

Display results - Each spectrum that is computed can be displayed on the screen along with the data frame and the zero padded signal or in the case of the AR functions the extrapolated ACS can be viewed.

Format of the input and output files:

Input files:

- Binary input files - data can be either 16-bit integers or 32-bit floating-point numbers.

  Optional 4-byte header contains:

  16-bit unsigned integer = Sampling frequency

  16-bit integer = Heterodyne carrier frequency

- Text files - No header.

NB. Due to file manipulation text files cannot be used if overlapping data frames are required.

Output files:

Data is stored in a binary file as 32-bit floating-point numbers.

A 6-byte header is generated automatically if a header was included with the input file, the format of the header is as follows:

16-bit integer - representing the DSP function used:

0 : Fast Fourier transform

1 : Yule-Walker, PSD obtained using equation [2.12]

2 : Burg, PSD obtained using equation [2.12]

3 : Yule-Walker, PSD obtained using equation [2.21]

16-bit unsigned integer = Sampling frequency

16-bit integer = Heterodyne carrier frequency

Appendix 2

A SIMPLE MODULATOR AND MIXER FOR DIRECTIONAL DOPPLER SIGNALS

F.S. Schlindwein, P.I.J. Keeton and D.J. Dryden

(Published in Revista Brasileira de Engenharia, RBE - Caderno de Engenharia Biomédica, 1996; 12:103-111)

A SIMPLE MODULATOR AND MIXER FOR DIRECTIONAL DOPPLER SIGNALS

F.S. Schlindwein[1], P.I.J. Keeton[1] and D.J. Dryden[1]

ABSTRACT -- This technical note describes the design of a simple instrument that conveniently modulates and mixes the forward and reverse components outputs from a Doppler velocimeter. The output is a single signal around a programmable carrier frequency. A graphical explanation of the way the circuit works is provided, as well as the complete circuit diagram. The instrument is extremely simple to build and is a very useful tool in a vascular laboratory.

Key words: Doppler velocimeters, Ultrasound, Directional Doppler, Blood flow, Instrumentation, Modulator

INTRODUCTION

The radio frequency signal received by a Doppler velocimeter contains information of velocities towards the probe and away from the probe. Without sacrificing generality this kind of signal can be written as

$$S(t) = A_0 \cos(\omega_0 t + \phi_0) + A_f \cos(\omega_0 t + \omega_f t + \phi_f) + A_r \cos(\omega_0 t - \omega_r t + \phi_r) \quad (1)$$

where the subscripts refer to the transmitted base frequency (0), the forward component (f) and the reverse component (r). The transmitted frequency $f_0 = \omega_0 / 2\pi$ is of the order of 2 to 20 MHz. Let us consider a signal with forward and reverse components whose amplitude spectral density is that shown in figure 1, where the lower hump corresponds to the components ω_r.

[1] University of Leicester, Department of Engineering, University Road, Leicester LE1 7RH, England, UK, e-mail: fss1@le.ac.uk

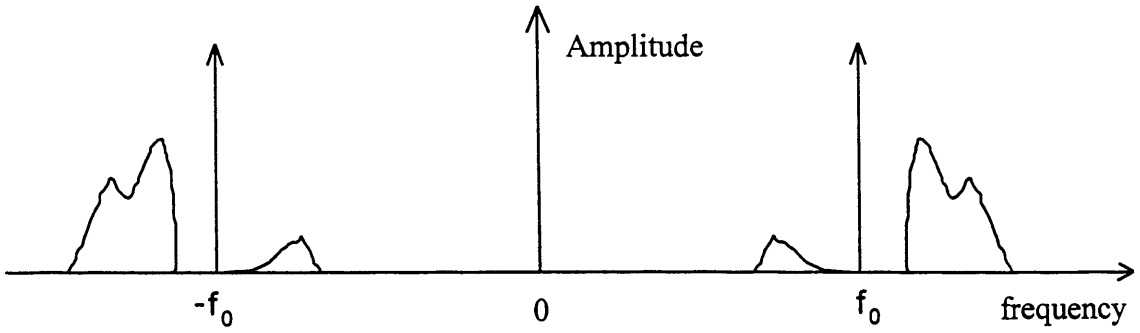

Figure 1. Radio frequency signal received by a CW Doppler velocimeter.

In the past, a simple coherent demodulator using the same angular frequency ω_0 from the master oscillator was applied to shift the signal S(t) down in frequency and produce an audio signal output, but this does not differentiate between movements towards and away from the transducer assembly because upper and lower sidebands of the forward and reverse components are brought to the same base audio band as seen in figure 2 (Wells, 1969).

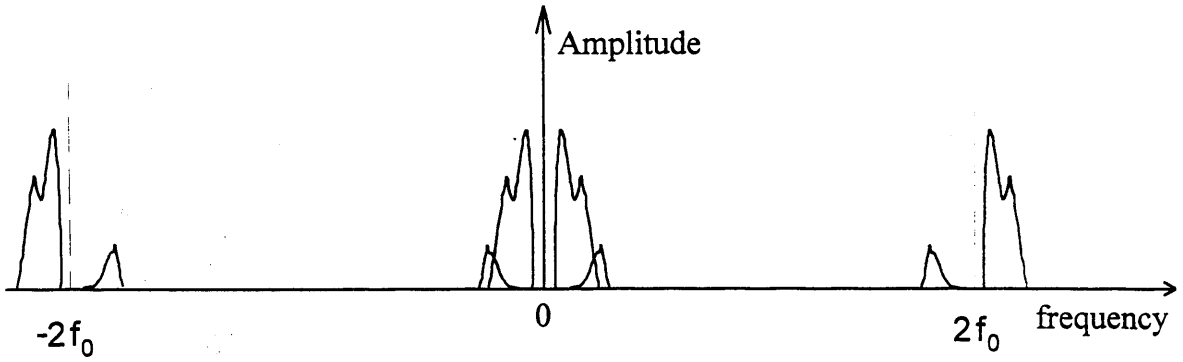

Figure 2. A simple coherent demodulator does not produce directional information because the upper and lower side bands are brought to the same base frequency band.

Nowadays most Doppler velocimeters are directional, that is, they provide separate information about forward flow and reverse flow. The most widely used technique of directional demodulation is quadrature phase detection, which results in two output audio signals, D(t)' and Q(t)' that contain the directional information in the relative phase between them. For a situation where flow is only towards the probe the direct signal D(t)' lags the quadrature signal Q(t)' by $\pi/2$, while for flow that is only receding from the probe it is the quadrature signal that lags the direct signal by $\pi/2$ (Evans et al., 1989). Quadrature phase

detection may be followed by phase domain processing (still in the Doppler velocimeter) resulting in the complete separation of the forward and reverse components into two separate signals F(t) and R(t) as shown in figure 3.

Directional Doppler velocimeters produce outputs either in phase quadrature form (as the Parks 806A does) or completely separated as a forward and a reverse component (as the Dopplex II does). Although these two kinds of output are mathematically equivalent and one can be transformed into the other, the best kind of output to handle is the one that uses separate forward and reverse signals because it can be recorded onto magnetic tape directly. This is not straightforward with quadrature signals because the process of recording and playing back the signal introduces different phase shifts to the two channels, corrupting or even completely destroying the directional information contained in quadrature signals.

The digital processing of the Doppler signals either in quadrature form or in the forward and reverse form requires two analogue to digital channels capable of sampling the signals at frequencies up to around 40 kHz. In the case of quadrature signals the traditional multiplex-sample/hold-A/D converter architecture of most commercial A/D boards for personal computers is not convenient. Simultaneous sampling is much preferred because the alternative procedure of dealing with the known phase shift introduced by the multiplexed sampling in the post-processing stage is quite awkward. Analogue to digital converter boards that allow simultaneous sampling are much more expensive than the traditional ones.

A single channel output is particularly well suited to spectral analysis since only one A/D channel and one spectrum analyser are required to deal with bi-directional flow. The instrument described here achieves exactly this. The design goal is: We want to mix the forward and reverse components around a programmable carrier of angular frequency ω_c, subtracting the undesired side bands, to produce a single analogue signal that contains the forward spectrum at frequencies above the carrier, say at $(\omega_c + \omega_f)$ for an individual component of angular frequency ω_f and the reverse spectrum at frequencies below the carrier, say at $(\omega_c - \omega_r)$ for a reverse component of angular frequency ω_r, as explained below.

DESIGN CHARACTERISTICS - A GRAPHICAL APPROACH

The description of the circuit can be entirely derived from a very simple graphical analysis of the frequency representation of the input signals and the required output. It follows.

The forward and reverse signals can be represented as in figure 3 while the required output is as in figure 4. The procedure to obtain that is easily followed referring to figures 5 and 6. First, both the forward component signal, F(t), and the reverse signal, R(t), are modulated by the carrier, $H(t) = \cos(\omega_c t)$. The complete circuit of the crystal controlled sine wave generator that produces the carrier signal is shown in figure 9. The resulting signal from

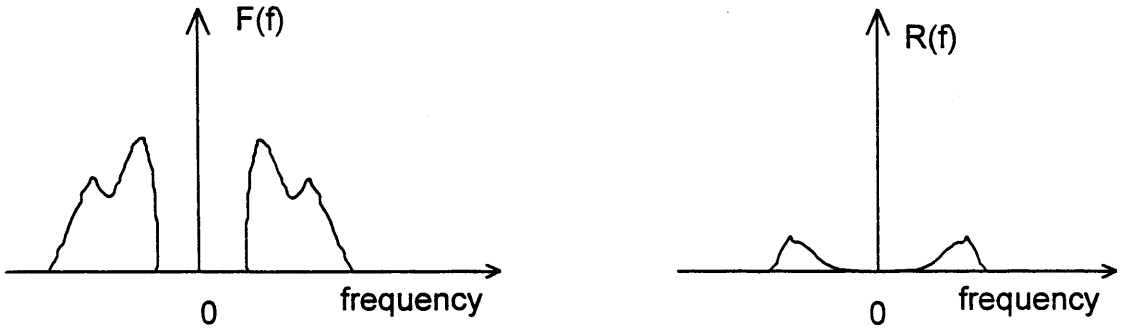

Figure 3. The frequency representation of the forward and reverse spectra.

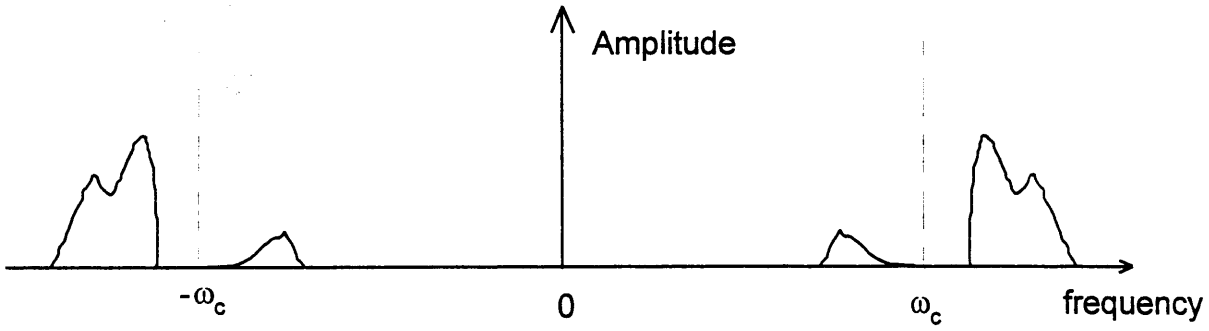

Figure 4. The required output from the instrument: Forward and reverse components modulated about a programmable carrier frequency ω_c and mixed to produce a single signal.

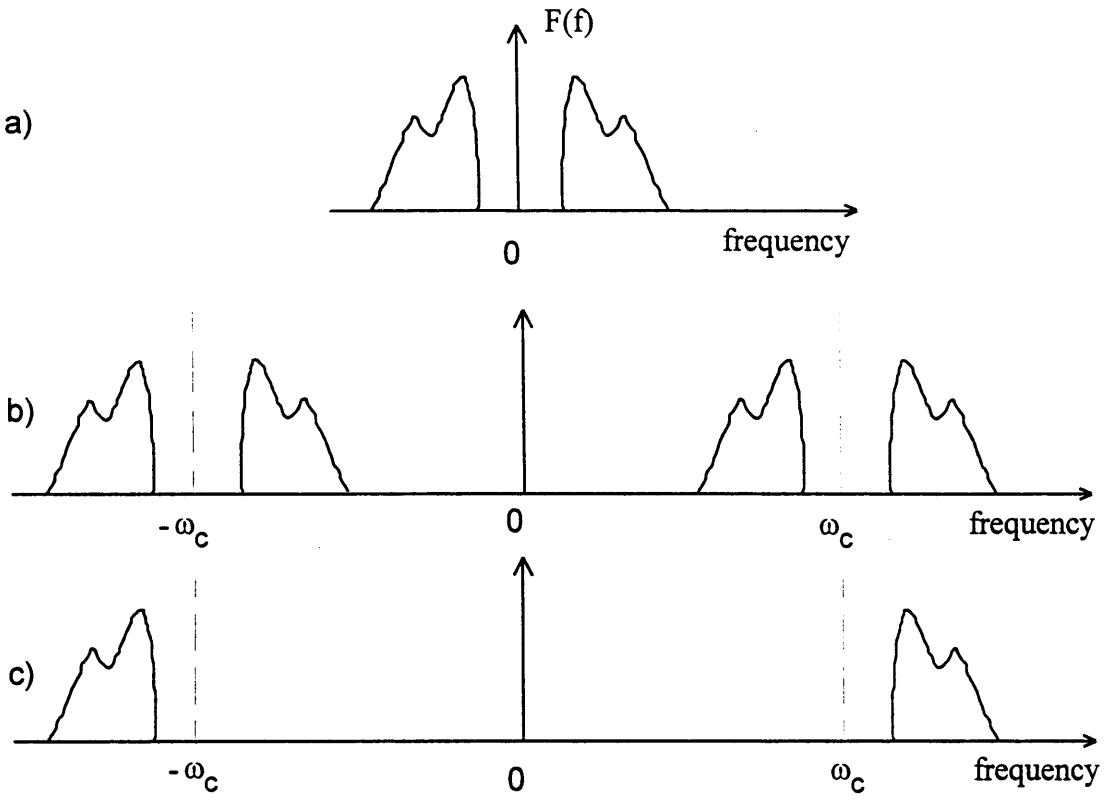

Figure 5. The forward component (a) is modulated (b) and then high-pass filtered (c).

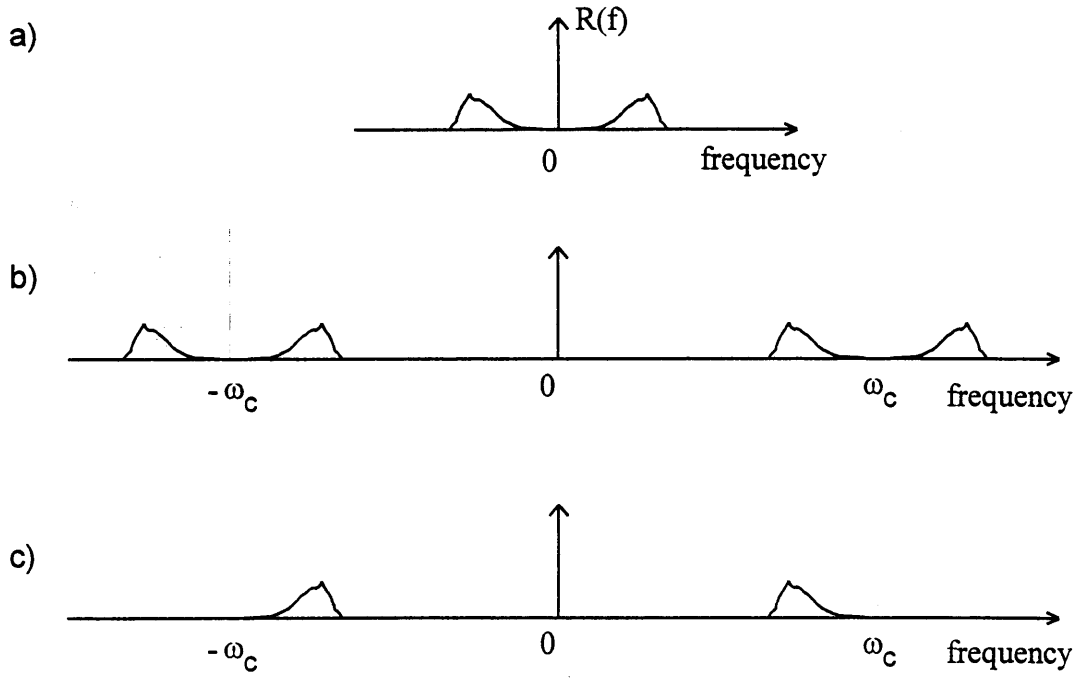

Figure 6. The reverse component (a) is modulated (b) and then low-pass filtered (c).

the modulation of the forward component is high-pass filtered to remove the lower side band, while the signal resulting from the modulation of the reverse component is low-pass filtered

to remove the upper side band. Finally the filtered signals are added producing the required output. The diagram with the modulators, filters and adder is shown in figure 7 and the complete circuit in figure 10.

The result of adding the signals represented in figure 5c and 6c is obviously the required signal of figure 4.

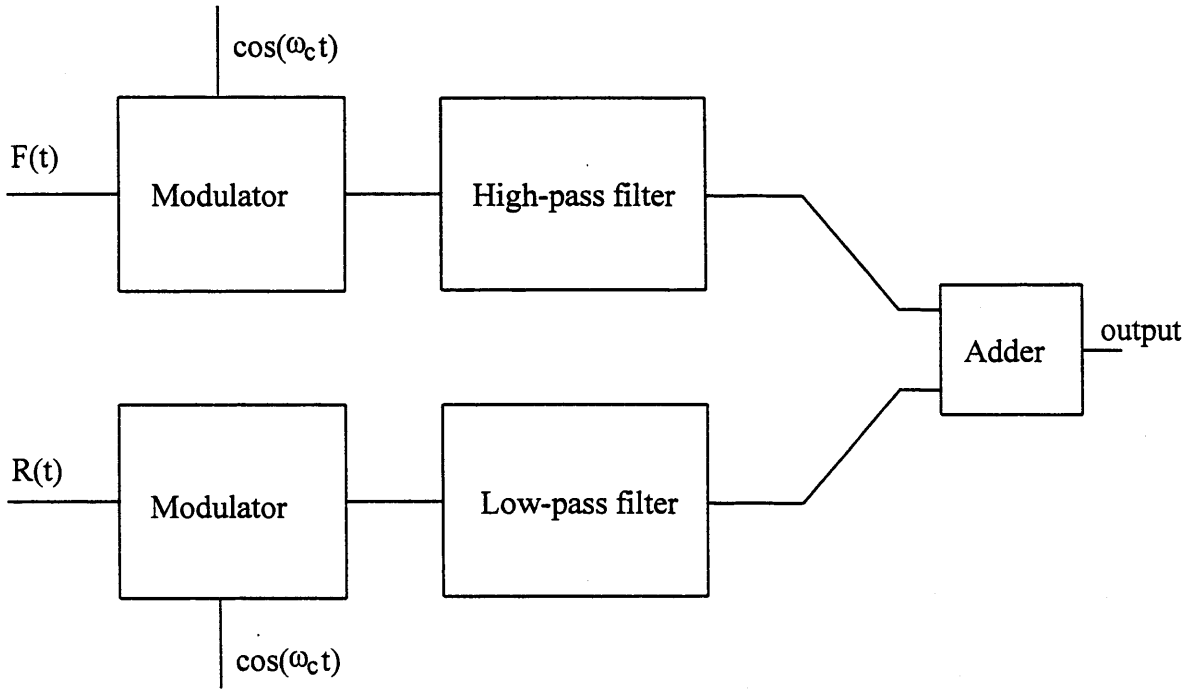

Figure 7. Block diagram of the independent modulator and mixer.

The frequency of the carrier signal is programmable via a 3-bit BCD switch to be 500 Hz, 1, 2, 4, or 8 kHz. The setting of this BCD switch is read by an IBM-compatible personal computer using the parallel printer port programmed as an input. The frequency selection switch, represented by T1, T3 and T4 in figure 9, produces the six input bits labeled 'FREQ' after the EPROM 2732. These are used as input codes to the filters. Both the high-pass and the low-pass filters are implemented using switched capacitor filter IC chips - two 7th order Elliptical modules in cascade for each, the low-pass implemented using the American Megatrends Incorporated S3528 and the high-pass, the S3529 as shown in figure 10. For the IC filters the specified pass band ripple is only 0.05 (typ.) per stage, the stop band attenuation is better than 51 dB per stage (102 dB total) for frequencies higher (for the low-pass) than 1.3 times the cut-off frequency. The phase is non-linear around the cut-off frequency - and

therefore the group delay is not flat, but this is of little consequence for the kind of processing we perform on the Doppler signals: We are interested in the amplitude spectrum only; phase distortion is not relevant. The choice of the carrier frequency automatically reprograms the cut-off frequencies of both filters. It is left to the operator to ensure that the carrier frequency is larger than the maximum frequency component contained in the reverse signal.

CONCLUSION

The circuit described performs its function well and produces a clean signal as shown in the amplitude sonogram of figure 8. It has been in use in our research laboratory for over two years. We feel that the graphical description given here is of value for teaching the subject and might even help designers with similar problems. The full diagram of the circuit is given in figures 9 and 10. A critical part of the adjustments for the circuit is the carrier suppression. This was performed using a spectrum analyser, measuring the output of the modulators IC2 and IC9 (Motorola MC1496) and carefully adjusting the multi-turn potentiometers VR1 and VR2 of the circuit of figure 10.

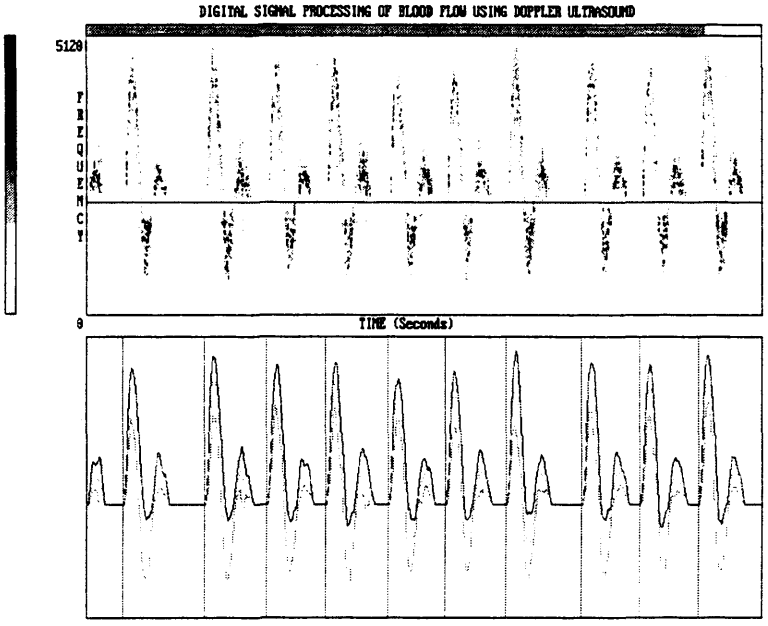

Figure 8. The sonogram of a bi-directional Doppler signal processed by the instrument together with the results of the detection of the individual heart beats. The grey-scale for the amplitude sonogram is linear.

REFERENCES

Wells, P.N.T. (1969), *Physical Principles of Ultrasonic Diagnosis*, Academic Press, London, UK.

Evans, D.H., McDicken, W.N., Skidmore, R. and Woodcock, J.P. (1989), *Doppler Ultrasound*, John Wiley & Sons, Chichester, UK.

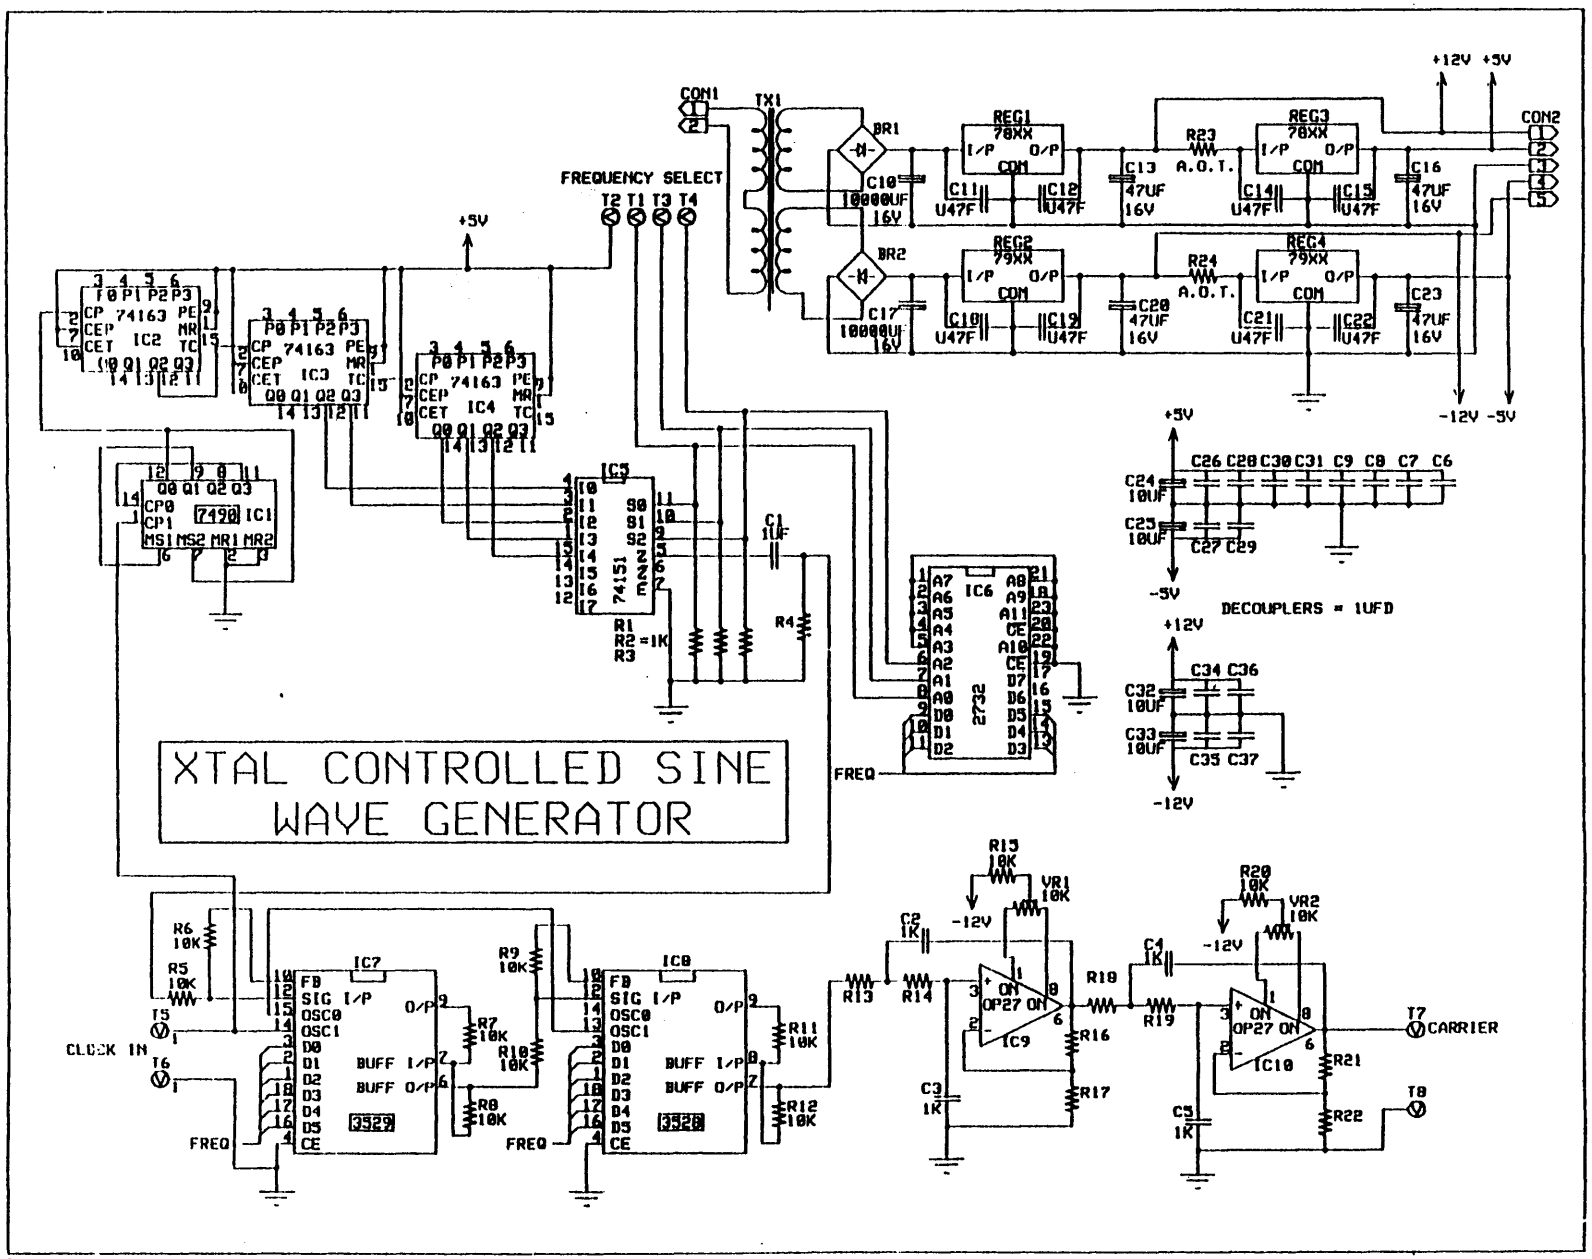

Figure 9. The crystal controlled carrier generator. Only the fundamental frequency goes through the band pass filter to produce the carrier signal.

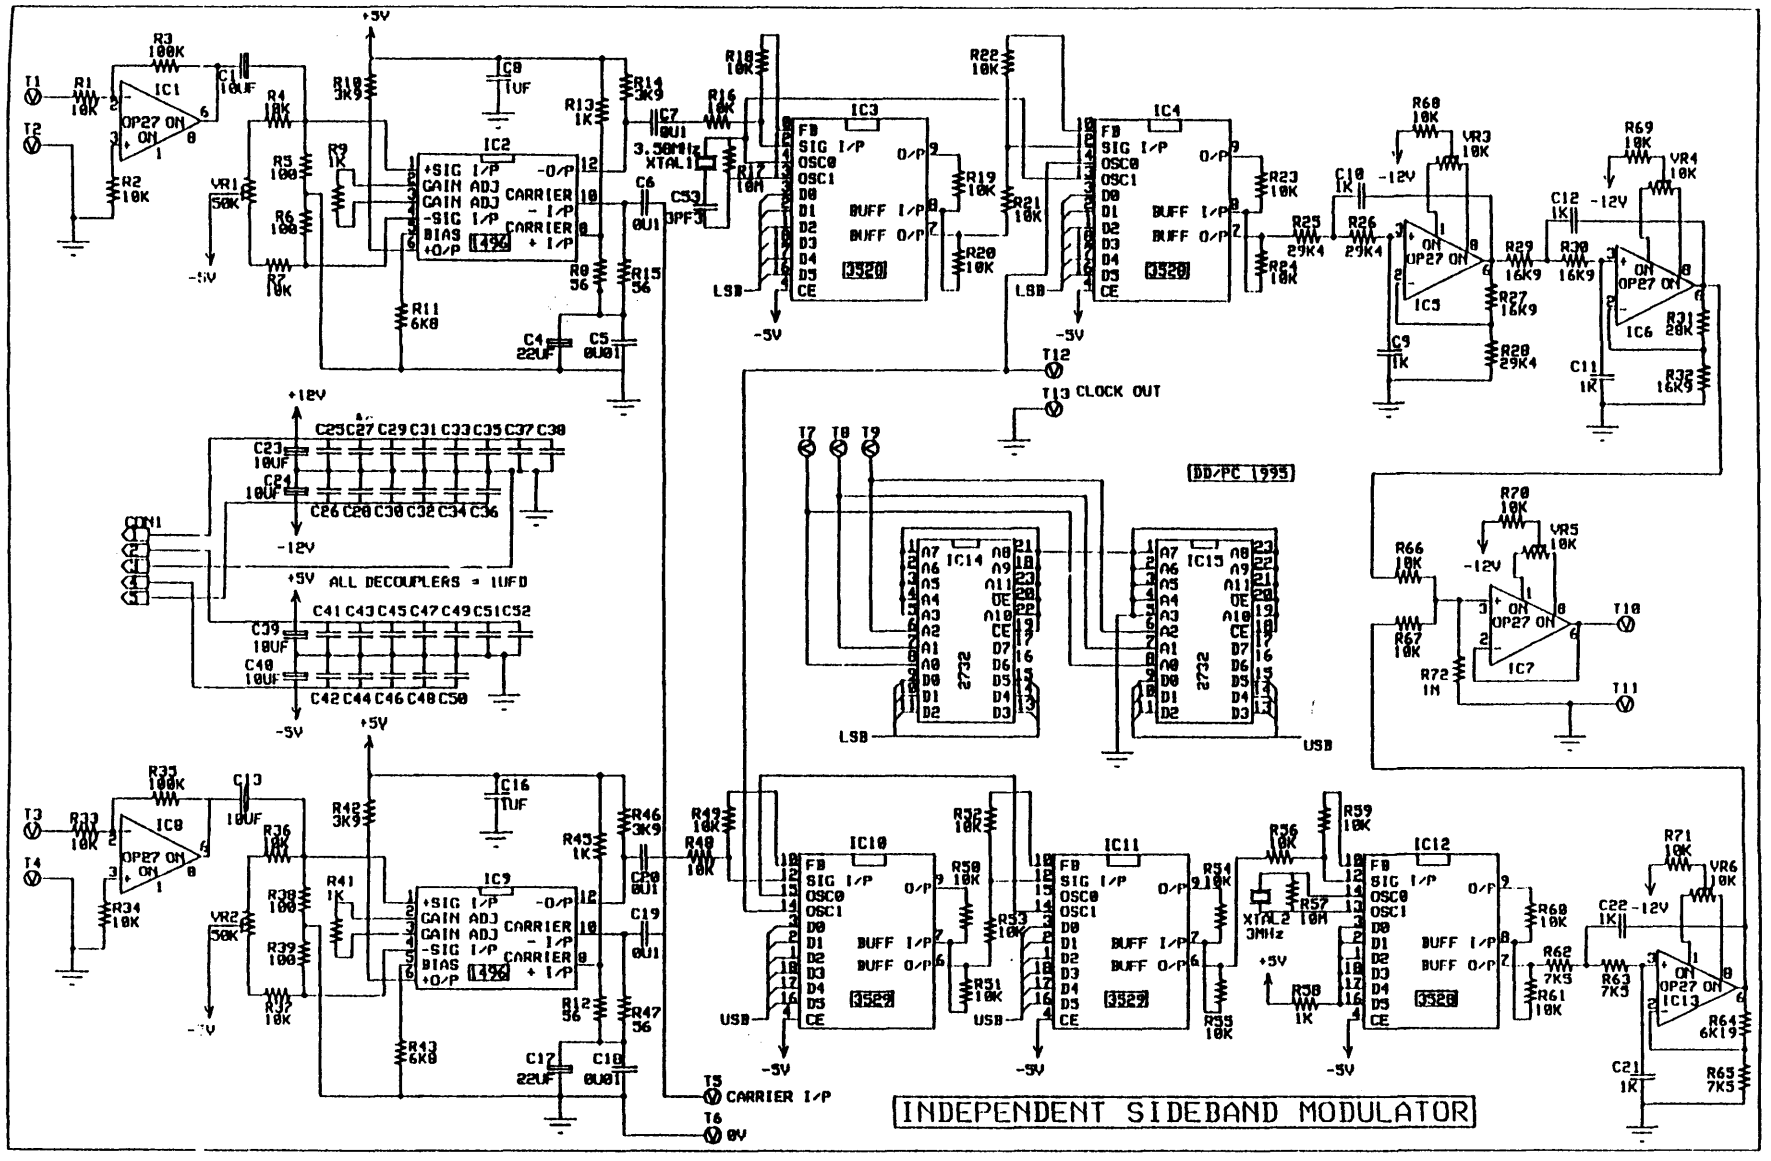

Figure 10. Complete circuit diagram for the modulator, filters and adder.

Appendix 3

A STUDY OF THE SPECTRAL BROADENING OF SIMULATED DOPPLER SIGNALS USING FFT AND AR MODELLING

P.I.J. Keeton, F.S. Schlindwein

(In Press: Ultrasound in Med. & Biol. 1997)

A STUDY OF THE SPECTRAL BROADENING OF SIMULATED DOPPLER SIGNALS USING FFT AND AR MODELLING

Abstract — Doppler ultrasound is used clinically to detect stenosis in the carotid artery. The presence of stenosis may be identified by disturbed flow patterns distal to the stenosis which cause spectral broadening in the spectrum of the Doppler signal around peak systole. This paper investigates the behaviour of the spectral broadening index (SBI) derived from wide band spectra obtained using autoregressive modelling (AR) compared to the SBI based on the fast Fourier transform (FFT) spectra. Simulated Doppler signals were created using white noise and shaped filters to analyse spectra typically found around the systolic peak and to assess the magnitude and variance of AR and FFT-SBI for a range of signal-to-noise ratios. The results of the analysis show a strong correlation between the indices calculated using the FFT and AR algorithms. Despite the qualitative improvement of the AR spectra over the FFT the estimation of SBI for short data frames is not significantly improved using AR.

keywords: Doppler ultrasound, carotid stenosis, autoregressive model, spectral analysis, spectral broadening.

INTRODUCTION

The 'fast Fourier transform' (FFT) is the most commonly used digital signal processing tool for estimating the spectrum of Doppler ultrasound signals. This is due to the computational efficiency and widespread availability of FFT algorithms. It has been suggested however that the FFT is not necessarily the best tool for analysing Doppler blood flow signals. Kitney and Giddens (1986) stressed the better performance on spectral tracking and spectral resolution of autoregressive (AR) spectral estimation when short frames are used. Kaluzynski (1987) reported the advantages of using AR modelling for the analysis of pulsed Doppler signals especially for short data lengths. Vaitkus, Cobbold and Johnston (1988) addressed the good spectral matching ability of the AR modelling approach using a simulated stochastic stationary Doppler signal with a known theoretical spectrum as a reference test sequence. They also added white noise in order to test the performance of six alternative spectral estimation methods with signals at various signal-to-noise ratios (SNRs).

Accurate extraction of Doppler shifts using current signal processing algorithms is extremely difficult since Doppler ultrasound signals are highly non-stationary, this is due to the pulsatile flow through the vessel. The stationarity of the signal is further reduced if the flow pattern is disturbed as a result of an obstructed vessel. If the blood flow over the cardiac cycle is to be observed, it is necessary to use time frames that are no longer than the length of time that the signal can be considered stationary. If longer time frames are used then the frequency spectra will be smeared and the consecutive frames will not provide a detailed indication of how the velocities within the vessel are changing with respect to time. The data frames are typically shorter than 20 ms (Evans et al. 1989, Nichols and O'Rourke 1990) and this limits the signal duration that can be used to estimate the frequency spectrum.

For the analysis of stochastic signals the FFT approach is known to have an inherent inability to produce stable spectra. Vaitkus, Cobbold and Johnson (1988) results demonstrated that the FFT approach has the largest variance for all frequency

values when compared to five other spectral estimation techniques at an SNR of 20 dB using 20 realisations of 256 samples of the simulated stationary Doppler spectrum. There is also a trade-off between frame acquisition rate and frequency resolution. For very short time segments the FFT spectra will have only a few frequency bins. Zero padding can be used to interpolate between the sparse points but this offers no improvement in the underlying spectral resolution. Another problem with the Fourier based spectral estimation approach is spectral leakage. This can be improved with anti-leakage windows, but only at the expense of further reduction of spectral resolution. The autoregressive spectral estimation produces more stable spectra from short segments of data and also produces spectra that have good spectral matching ability, "closely approximating the theoretical spectrum with good statistical consistency" (Vaitkus, Cobbold and Johnston 1988). The AR approach does not assume periodicity, offering an interesting alternative to FFT-based techniques for the analysis of Doppler signals of disturbed blood flow signals such as the kind of signals collected at post-stenotic sites.

This study is divided into two parts, firstly the relationship between the spectra obtained using FFT and AR was studied using a diagnostic parameter for simulated symmetrical wide band signals using data frames containing 64 and 256 samples. Secondly, three filters were used to simulate Doppler spectra typically found around the systolic peak which represented signals taken from a healthy carotid artery and signals that represented diseased carotid arteries. Gaussian noise was added to the simulated Doppler signals over a range of SNRs. The accuracy and robustness of the parameter used to measure spectral broadening for both the FFT and AR spectra were compared for a range of signal to noise ratios.

Spectral broadening index (SBI)

The spectral broadening index (SBI) has been used as the parameter traditionally associated with the measurement of flow disturbances that occur with stenosis (Harward et al. 1986, Labs and Fitzgerald 1992). Kaluzynski and Palko (1993) studied the behaviour of SBI as well as other indices with different conditions of spectral analysis for simulated signals. In this study SBI was calculated from both

the FFT and the AR-based spectra. There are several definitions of SBI, the one used in this work is given by:

$$SBI = \frac{f_{\max} - f_{mean}}{f_{\max}} \qquad [1]$$

where f_{max} = the maximum frequency component in the spectrum above a threshold used to reject low level noise components and f_{mean} = the average mean frequency of the spectral components that exceed this threshold.

The SBI is calculated by selecting a threshold to reject low level noise components as can be seen in figure 1. It has been observed by Schlindwein et al. (1994) that if a simple threshold is implemented then there is no obvious plateau on the SBI *vs.* threshold plot where the SBI can be considered constant. Figure 2 shows the SBI and the resistance index (RI) (Pourcelot, 1976) for a real Doppler signal. The Doppler ultrasound signal was taken from an internal carotid artery. The site of the scan was 2 to 3 diameters distal to the stenosis. A continuous wave Doppler velocimeter, with a 4 MHz focused 1.5mm beam width, was used for data collection. The beam transfixed the vessel axis at an angle of around 50°. It can be seen from figure 2 that there is a range of thresholds over which the value of RI is constant but the SBI is changing (RI has been included here to illustrate the window where the envelopes used to determine the diagnostic parameters fit the signal, i.e. the range of thresholds where noise is rejected with minimal loss of the true Doppler signal). The reason why there is no plateau for SBI is due to the trends in the mean and maximum frequency envelopes: As the threshold is increased, f_{max} decreases and f_{mean} increases. The uncertainty caused by the positioning of the threshold results in an inconclusive SBI value. This problem can be avoided if a threshold is defined that is set at a pre-determined level with respect to the individual spectra. It is proposed that the threshold should be set 6dB below the maximum component of each of the individual spectra as suggested by Labs and Fitzgerald (1992). This significantly reduces the sensitivity of the parameter to gain and also produces a single estimate of SBI for a particular recording.

If a high grade stenosis is affecting the flow pattern, then it is possible to have simultaneous forward and reverse flow components as a result of turbulence. The presence of reverse flow is in itself a strong indication of disturbed flow. In this situation the definition of SBI as described in [1] is not adequate.

Estimation of the frequency spectrum is further complicated by the presence of noise in the signal, it has been reported that typical clinical SNR levels are in the range of 0-20dB (Forsberg, 1991). If the signal processing algorithm is not stable in the presence of noise then the variance of the spectral estimate is going to be large. Parameters derived from these spectra will have a large variance making it clinically infeasible to produce a statistically significant estimate since the number of heartbeats that need to be sequentially recorded for averaging the SBI will be too large. Therefore one of the aims of this paper is to establish the magnitude and variance of SBI in the presence of noise at SNRs present in a clinical environment.

METHODS & DISCUSSION

The behaviour of the SBI index obtained from spectral estimates based on the FFT approach was compared to that obtained from spectra produced by autoregressive modelling. The first algorithm was a straightforward 256 point FFT with no overlap and no anti-leakage window. The second algorithm was identical to the first except that the data was multiplied by an 256 point Hamming window. The AR algorithm was based on an implementation of the Yule-Walker equations using the recursive Levinson-Durbin approach with model order 12 as described by Schlindwein and Evans (1989). Both techniques produced spectral estimates with 128 frequency bins.

The signals used in this study were generated using filtered Gaussian white noise which has a uniform theoretical spectrum from zero to half the sampling frequency. The shape of the filter was adapted to generate a series of wide band signals with varying bandwidths and frequency response characteristics. The ideal shape of the filter was approximated using an 8-order IIR Yule-Walker digital filter whose coefficients were determined using a least squares fit to the desired filter shape.

The design was performed using MATLAB. A Gaussian white noise sequence of N samples was filtered and the first 256 filtered samples were discarded so that transients caused by the filter would not be processed. The impulse response of the filter was used to determine the number of samples that should be rejected to ensure that the initial transients had decayed sufficiently within the first 256 samples. Figure 3 shows the frequency response of the symmetrical band-pass filters that were used to create the wide band signals.

Comparison of FFT-SBI & AR-SBI

Preliminary tests were carried out to establish the relationship between FFT-SBI (with and without an anti-leakage window) and AR-SBI by comparing the magnitude and variance of the estimates. The SBI was calculated from 64 separate spectral estimates obtained using a 64 point data frame and a 256 point data frame using each of the algorithms described above. The results obtained were plotted against one another and are shown in figure 4.

Table 1 shows the exact value of the SBI determined from the frequency response of each filter and the mean and variance of the SBIs determined from the spectral estimates using each of the algorithms described above. It can be seen that for the 256 point data frame the mean of the AR-SBI is closer to the true value than the FFT estimate and generally the variance of the SBI estimate is smaller for AR than FFT. When a data frame of only 64 points is used the variance of all the algorithms increases and the average SBI is not so close to the true value.

It can be seen from figure 4 that there is a linear relationship between the SBIs calculated using the three separate algorithms. The correlations between the indices are shown in table 2 and table 3 for the 64 and the 256 point data frames respectively. The correlation coefficients for each of the comparisons exceed the 0.1% value, therefore the results are significant at the 0.1% level ($p < 0.001$). The correlation between the FFT (no anti-leakage window) and the AR algorithm is the strongest. The correlation between the FFT (anti-leakage window) and the other algorithms gives the worst result. From these results it is apparent that the application of an anti-leakage

window does not improve the estimation of SBI, in fact the variance of the SBI when using FFT and an anti-leakage window is larger than that of the SBI obtained for FFT spectra with no anti-leakage window.

Effect of AR model order on SBI

When implementing the AR algorithm there is a certain degree of choice in the selection of model order for the AR model. There are various techniques for establishing the 'best' order for a particular process. For Doppler signals Schlindwein and Evans (1990) suggested that the AR model order (p) should be around 12. The effect of using different model orders was tested using the same symmetrical wide band signals that were used to compare FFT-SBI and AR-SBI. Figure 5 shows the result of model order on the estimation of AR-SBI. If the model order is less than about 8 then there are not enough poles to adequately model the signal and therefore the spectral estimates are not representative of the true spectrum (Kaluzynski, 1989). If the model order is too high then the filter attempts to model the local fluctuations in noise as well as the underlying signal producing a biased estimate (Akaike, 1974). In addition the computational time to calculate the autoregressive coefficients using the Levinson-Durbin recursive algorithm is proportional to p^2, therefore in real-time applications of the AR algorithm it is sensible to avoid unnecessary computation. For AR model orders between 10 and 16 the variance of AR-SBI is smaller than that of FFT-SBI for all cases.

Estimation of SBI from simulated Doppler signals

Having established that the FFT and AR algorithms produce comparable indices for symmetrical signals the second stage of the study was to look at the ability of the algorithms to produce estimates of SBI in the presence of noise and with asymmetric spectra, which are more comparable to Doppler signals. Simulated Doppler signals were created using filters modelled on the typical frequency spectra obtained around the systolic peak of the cardiac cycle. Three types of filter were used to simulate different degrees of spectral broadening. Figure 6 shows the frequency responses of the filters that were used to generate the simulated Doppler spectra. The first filter has only slight spectral broadening and is characteristic of a healthy carotid

artery, where plug flow is present. The second two spectra simulate the effect of a stenosis which would lead to spectral broadening as the flow would be disturbed. A second Gaussian white noise signal was used to add noise to the simulated Doppler signals to generate noisy signals with SNR from +10dB down to -10dB.

The signals were processed using the FFT and AR algorithms using data frames of 64 and 256 samples. Figure 7 shows examples of typical spectra obtained using each of the algorithms with a data frame of 256 samples, it can clearly be seen that the AR spectral estimate is closer to the shape of the filter (refer to figure 6(b)) than the FFT spectra. All of the spectral estimators produce spectra with a varying degree of ripple in the pass band, figure 8 shows the variance in the spectral estimates for the simulated Doppler spectra (figure 6(c)) using AR with a data frame of 256 points. The results show that although there is a high variance in the pass band, the bandwidths of the estimated spectra are stable. The average SBI and the variance was estimated from 128 and 32 frames for each signal using the 64 and 256 sample data frames respectively. The same 6dB threshold as for the first part of the paper was used here to determine SBI from the simulated Doppler spectra. The results are shown in figure 9 and figure 10.

It can be seen that for high SNRs the average SBI is constant for both FFT and AR. Figure 9 also shows the true SBI for each of the filters, all of the algorithms underestimate the actual value, this is particularly true for the 64-point data frame. When the 256-point data frame is used the estimate of SBI is better, with AR-SBI producing the most accurate result. The variance of the SBIs for the short data frames is approximately the same for both the FFT and AR, however the AR performs marginally better than the FFT for a 256 point data frame. The ability of the FFT to perform as well as AR particularly for the short data frames can be attributed to the fact that the SBI is solely dependent upon the rise and fall of the wide band frequency spectral estimate, the variance across the pass band does not affect SBI as long as the ripple is uniform. Therefore even though the individual spectral estimates for FFT are qualitatively worse than AR, when the results are averaged over a number of frames

the FFT is able to closely match the shape of the filter of figure 6(c). Figure 11 shows the averaged spectra for all the algorithms using 64 and 256 samples.

When the level of noise approaches the threshold the SBI starts to increase. The SNR at which the SBI starts to drift is dependent upon the bandwidth of the simulated Doppler signal. For signals with a high degree of spectral broadening the noise will start to cause the SBI to drift at higher signal-to-noise ratios than signals with a lesser degree of spectral broadening. This is due to the larger relative amplitude of the noise for a broad band signal compared to a narrow band signal. It can be seen from figures 9 and 10 that the magnitude and variance of the SBI are stable when the SNR is higher than 3dB for the 256-point data frame (7dB for the 64-point data frame). Therefore in a normal clinical environment where the SNR is around 10-20dB the results of SBI can be considered stable. In a very noisy environment where the SNR is approaching 0dB the results will be affected by the noise.

CONCLUSIONS

The results of this study have shown that there is a strong correlation between the estimates of SBI using both the FFT and AR algorithms. The actual value of SBI is more accurate with AR for a 256 point data frame although for both the data frames all the algorithms underestimate the true value. The use of an anti-leakage window on the FFT algorithm does not reduce the variance of the spectral broadening.

The average SBI for all algorithms is constant in low-level noise, it is therefore possible to estimate the SBI in this situation. The variance of the estimates for both FFT and AR at high signal-to-noise ratios remains fairly constant until the level of noise reaches the threshold where the SBI is measured. When the level of noise is comparable to the threshold the variance of the FFT and AR SBIs increase and the average SBI starts to increase as a result of noise contamination, as expected. The SNR at which the estimate of SBI tends to drift is dependent upon the degree of spectral broadening present in the signal. It should be noted that the magnitude of SBI tends to approximately 0.5 at low SNR.

For AR model orders greater than 10 the average value of AR-SBI is stable. For AR model orders between 10 and 16 the variance of AR-SBI is smaller than that of FFT-SBI.

Despite the fact that the results of this study have shown that there is a distinct qualitative improvement in the spectral estimates for AR compared to those of the FFT, the ability of the AR model to produce stable spectra does not result in a significant improvement in the determination of SBI over the FFT approach. For short data frames the FFT is able to produce similar SBI results to AR.

REFERENCES

Akaike H. A new look at the statistical model identification. IEEE Trans. Automatic Control. 1974; 19:716-723.

Evans DH, McDicken, WN, Skidmore, R and Woodcock, JP. Doppler Ultrasound - Physics, Instrumentation, and Clinical Applications. John Wiley & Sons Ltd, Chichester, 1989.

Forsberg F. On the usefulness of singular value decomposition—ARMA models in Doppler ultrasound. IEEE Trans. Ultrason. Ferroelec. Freq. Contr. 1991; 38:5:418-428.

Harward TRS, Bernstein EF, Fronek A. Continuous-wave versus range-gated pulsed Doppler power frequency spectrum analysis in the detection of carotid arterial occlusive disease. Ann. Surg. 1986; 204:1:32-37.

Kaluzynski, K. Analysis of application possibilities of autoregressive modelling to Doppler blood flow signal spectral analysis. Med. & Biol. Eng. & Comput. 1987; 25:373-376.

Kaluzynski K. Order selection in Doppler blood flow signal spectral analysis using autoregressive modelling. Med. Biol. Eng. & Comput. 1989; 27:89-92.

Kaluzynski K, Palko, T. Effect of method and parameters of spectral analysis on selected indices of simulated Doppler spectra. Med. & Biol. Eng. & Comput. 1993; 31:249-256.

Kitney RI, Giddens DP. Linear estimation of blood flow waveforms measured by Doppler ultrasound, In: Salamon R, Blum B, Jørgensen M, eds. MEDINFO 86. North Holland: Elsevier Science Publishers B.V. 1986; 5:672-677.

Labs KH, Fitzgerald DE. Quantification of pulsed Doppler spectra for the diagnosis of minor to moderate atherosclerotic lesions: experience from in vitro and in vivo models. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. London: Edward Arnold, 1992; 126-141.

Nichols WW, O'Rourke MF. McDonald's Blood flow in arteries, theoretical, experimental and clinical principles, Third Edition. London: Edward Arnold, 1990.

Pourcelot L. Diagnostic ultrasound for cerebral vascular diseases. In: Donald I, Levi S, eds. Present and future of diagnostic ultrasound. Rotterdam: Kooyker, 1976; 141-147.

Schlindwein FS, Evans DH. A real-time autoregressive spectrum analyzer for Doppler ultrasound signals. Ultrasound in Med. & Biol. 1989; 15:3:263-272.

Schlindwein FS, Evans DH. Selection of the order of autoregressive models for spectral analysis of Doppler ultrasound signals. Ultrasound in Med. & Biol. 1990; 16:1:81-91.

Schlindwein FS, Vieira MH, Vasconcelos CFM, Simpson DM. Real-time digital signal processing of Doppler ultrasound signals and calculation of flow parameters. Medical Progress through Technology. 1994;20:81-89.

Vaitkus PJ, Cobbold RSC, Johnston KW. A comparative study and assessment of Doppler ultrasound spectral estimation techniques part II: Methods and results. . Ultrasound in Med. & Biol. 1988; 14:8:673-688.

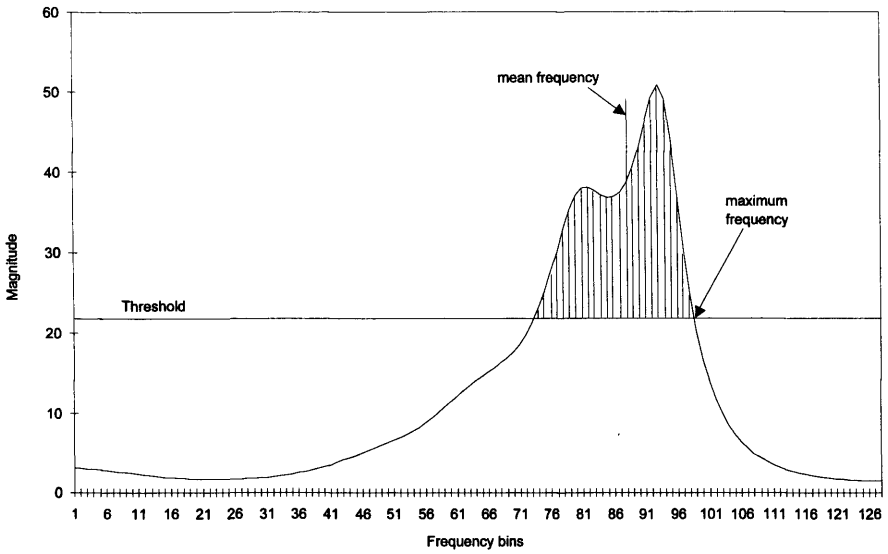

Figure 1 Determination of f_{max} and f_{mean} from frequency spectrum.

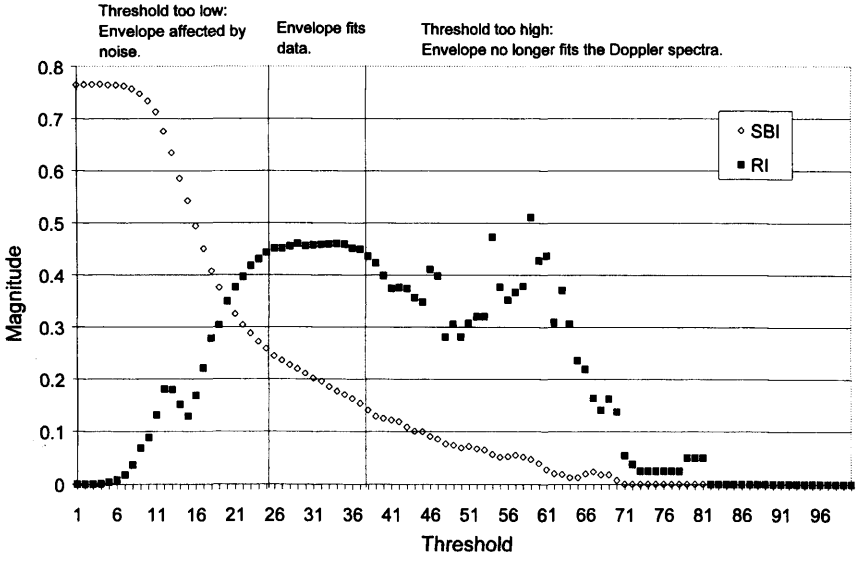

Figure 2 SBI & RI vs. threshold for an internal carotid artery.

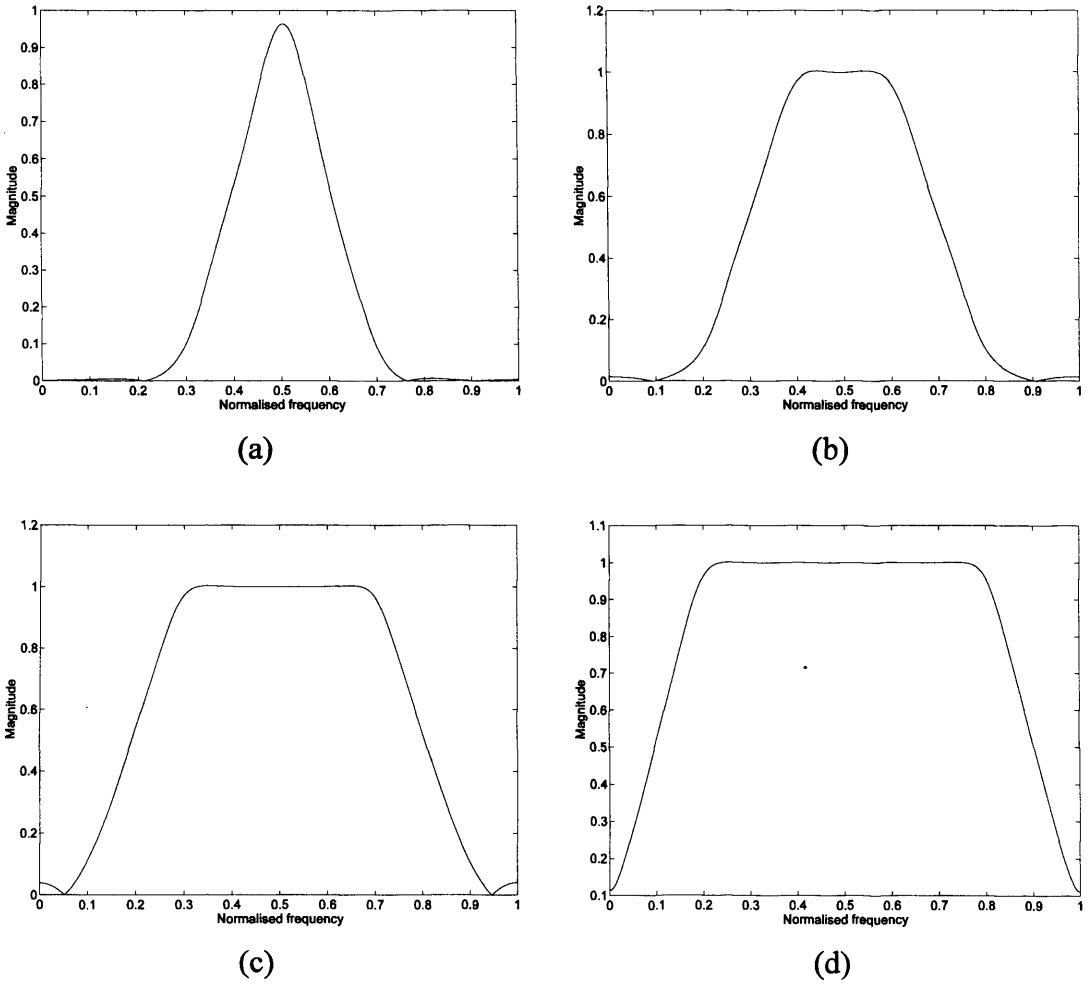

Figure 3 Frequency response of filters used to generate symmetrical wide band signals. Normalised 6 dB bandwidths for each filter are (a) 0.2 (b) 0.4 (c) 0.6 (d) 0.8.

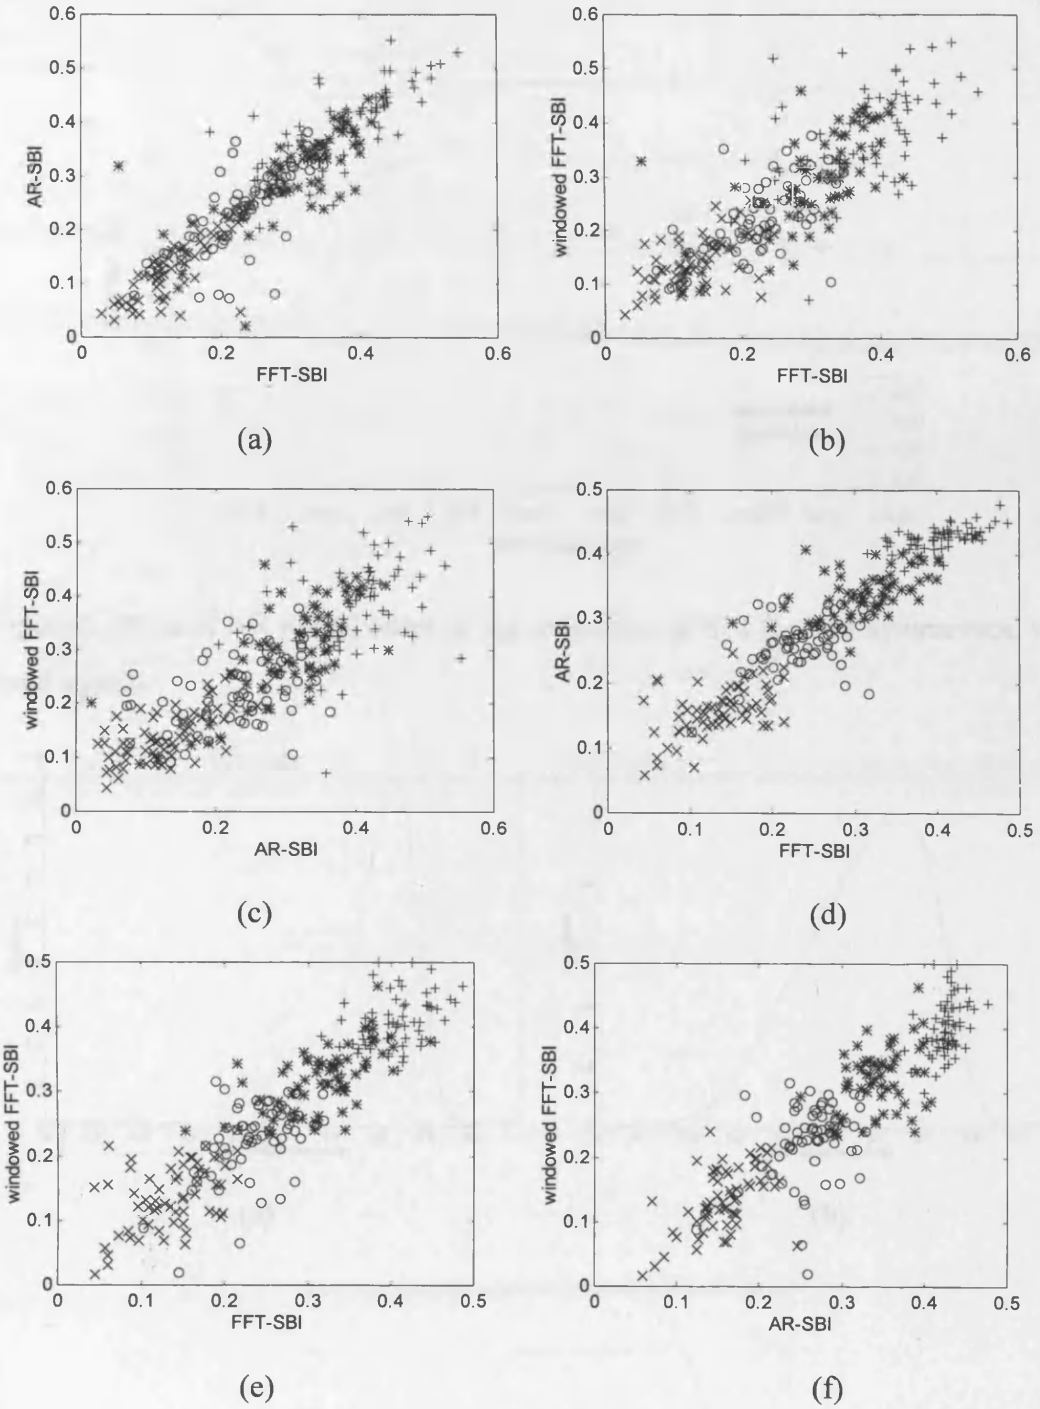

Figure 4 Comparison of FFT-SBI, AR-SBI and windowed FFT-SBI for four symmetrical wide band signals (a),(b),(c) 64-point data frame (d),(e),(f) 256-point data frame.

Normalised Bandwidth	
×	0.2
o	0.4
*	0.6
+	0.8

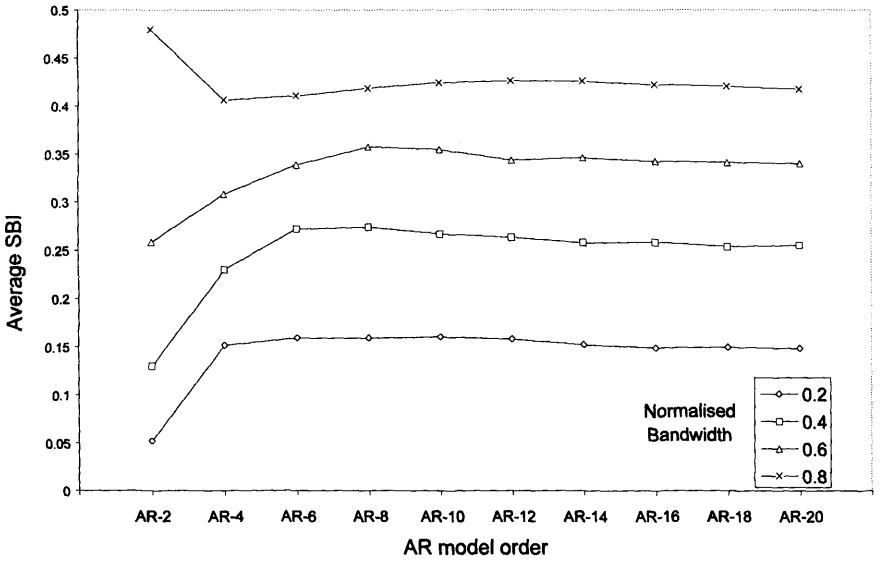

Figure 5 Effect of AR model order on the estimation of SBI for four symmetrical wide band signals.

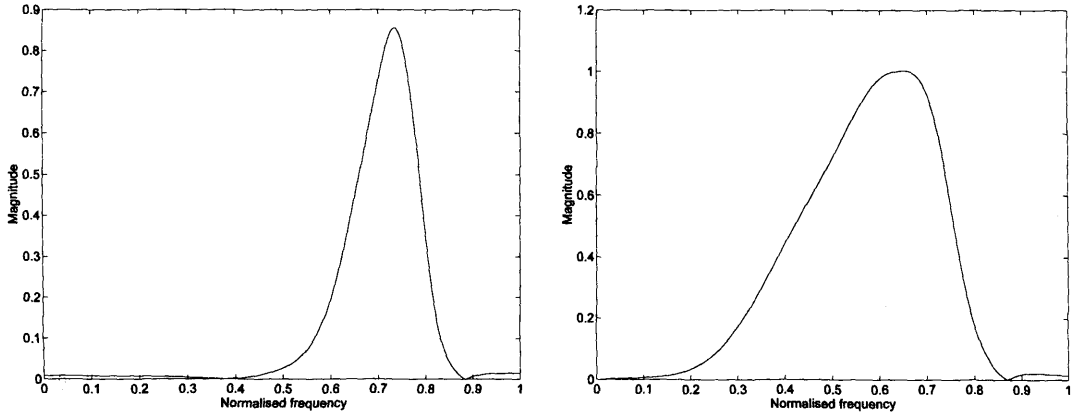

(a)

(b)

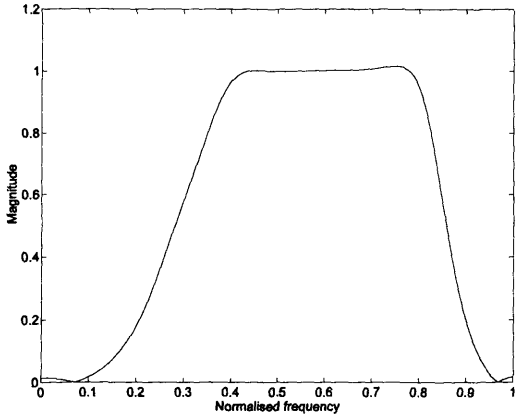

(c)

Figure 6 Frequency response of the filters used to simulate the Doppler spectra.

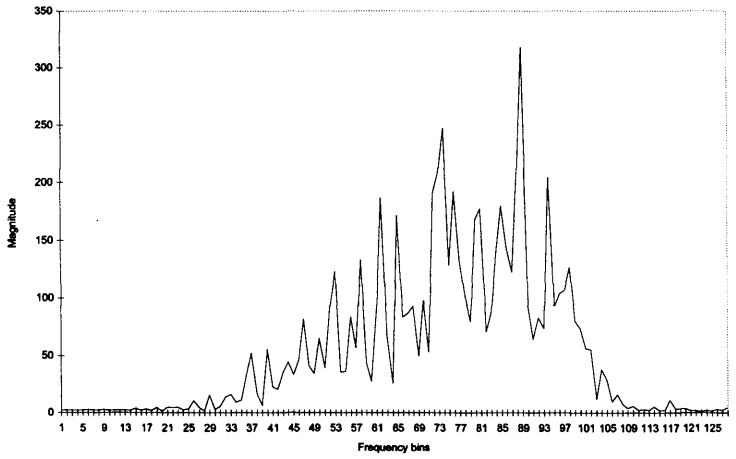

(a)

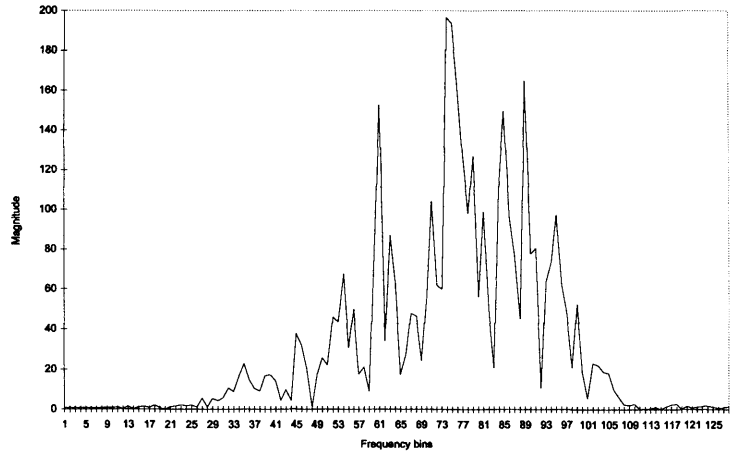

(b)

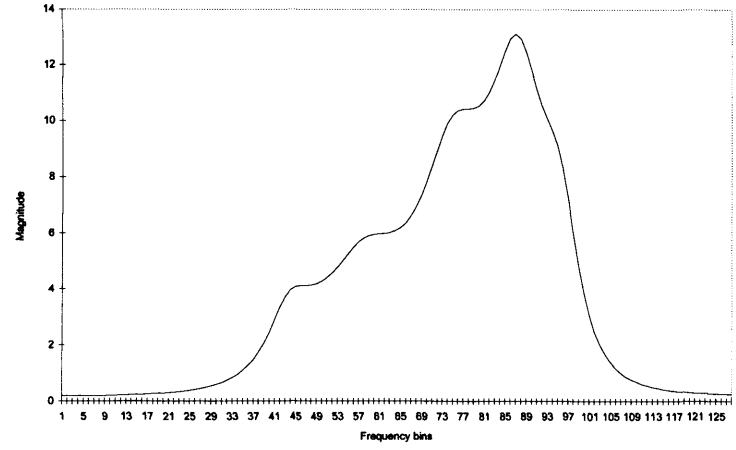

(c)

Figure 7 Frequency response for simulated Doppler spectrum. (a) FFT (b) FFT-Hamming window (c) AR model order 12.

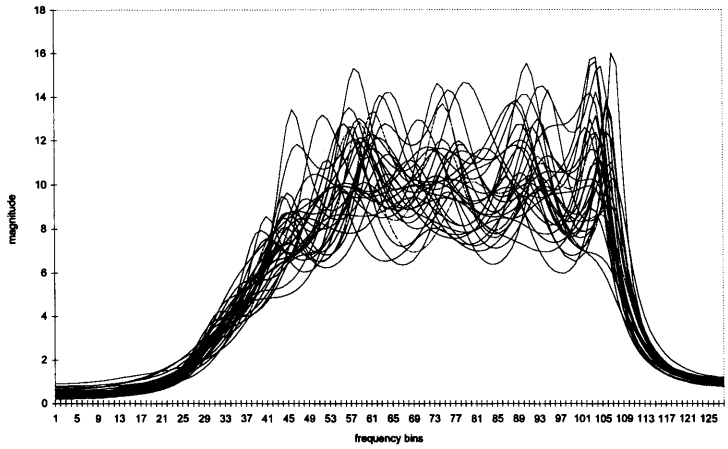

Figure 8 Spectral estimates for simulated Doppler spectra using AR (model order 12) on consecutive 256 point data frames.

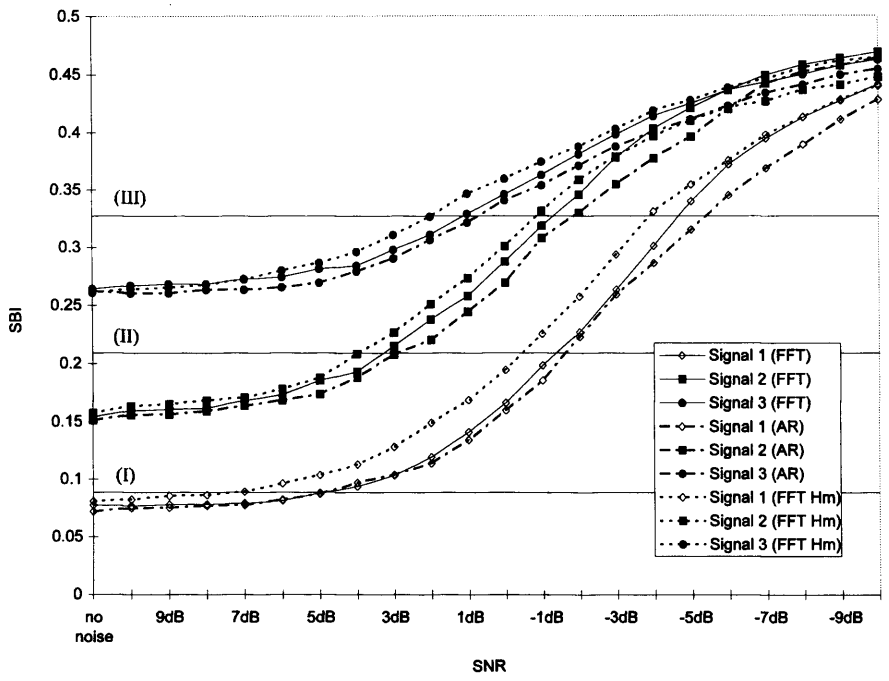

(a)

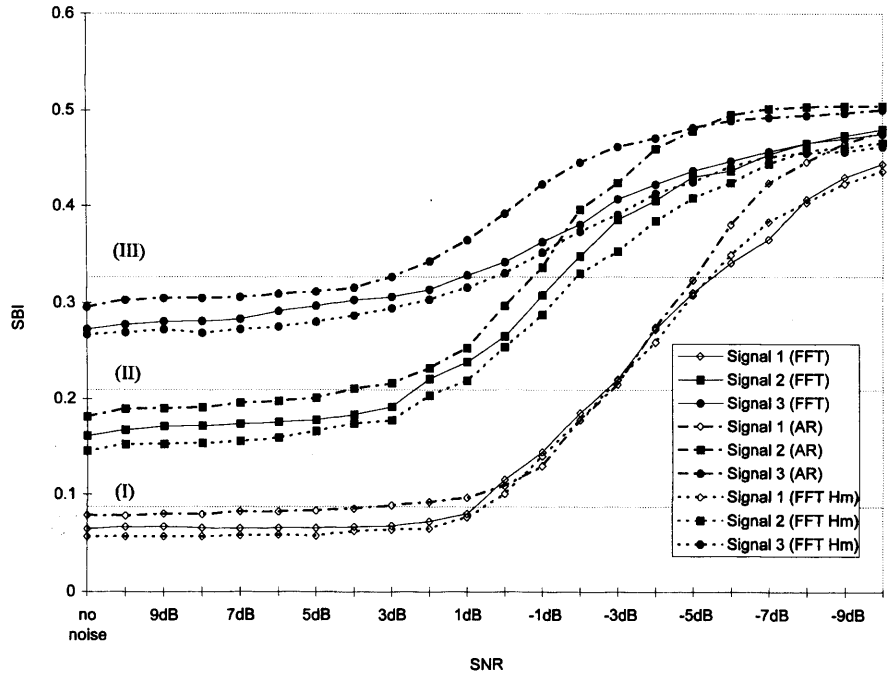

(b)

Figure 9 Average SBI for three simulated Doppler signals in noise (a) 64 point data frame (b) 256 point data frame. True SBI for simulated Doppler spectra indicated using the lines I, II & III for the three asymmetric filters in figure 6.4(a,b,c) respectively.

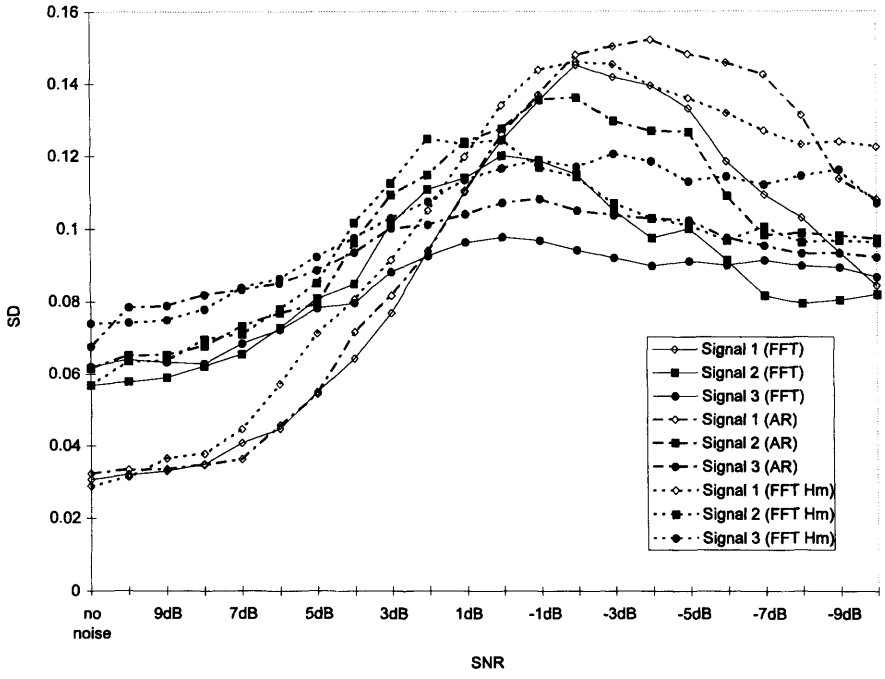

(a)

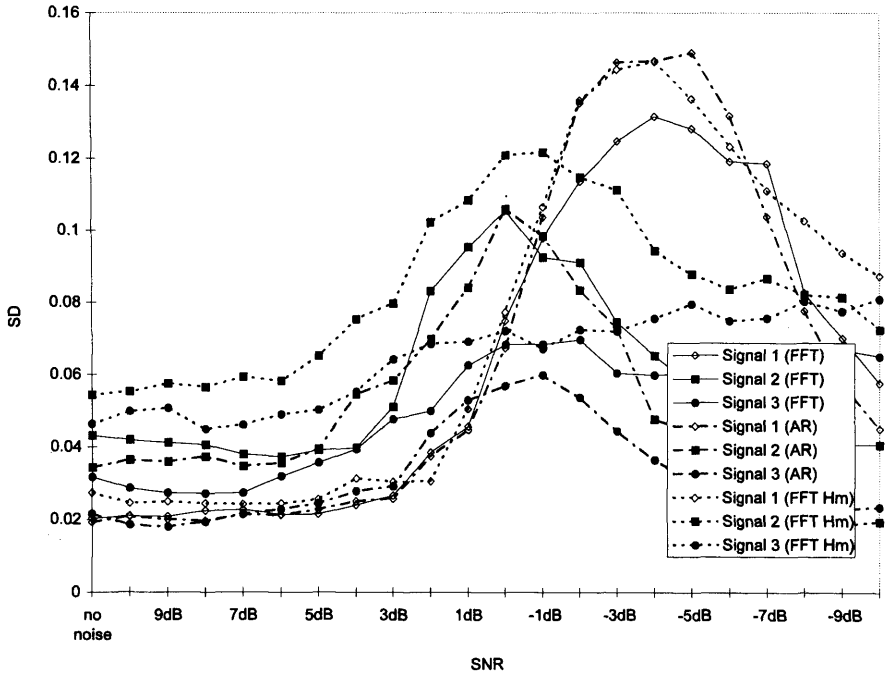

(b)

Figure 10 Standard deviation of SBI for three simulated Doppler signals in noise (a) 64 point data frame (b) 256 point data frame.

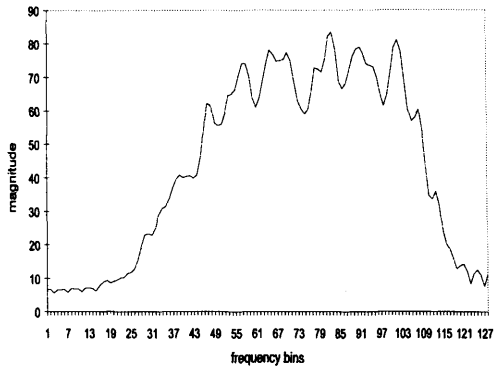

(a) FFT, 64 point data frame

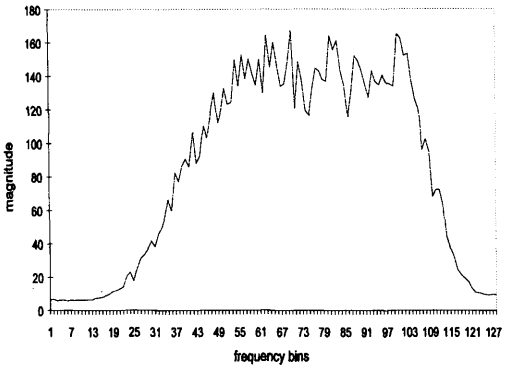

(b) FFT, 256 point data frame

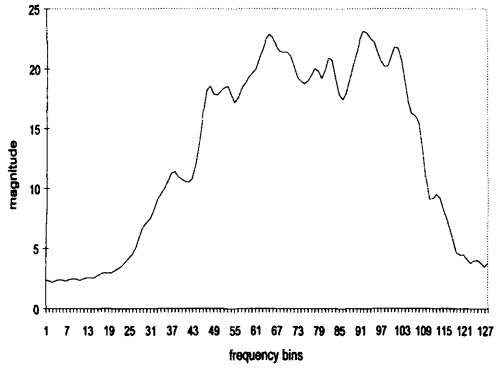

(c) FFT-Hm, 64 point data frame

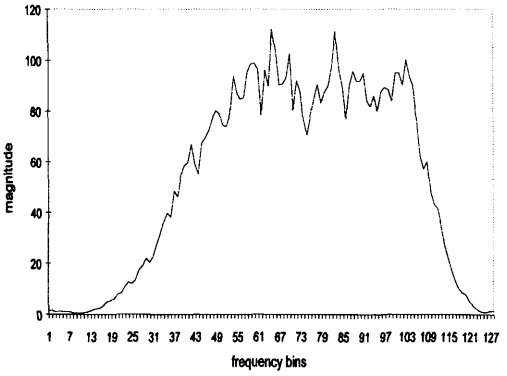

(d) FFT-Hm, 256 point data frame

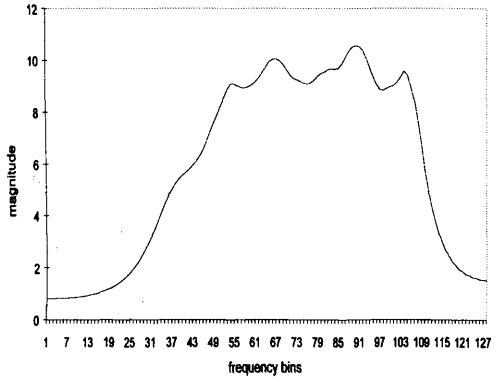

(e) AR order 12, 64 point data frame

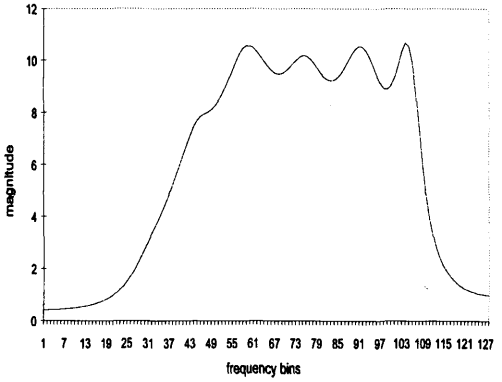

(f) AR order 12, 256 point data frame

Figure 11 Averaged frequency spectra for a stationary simulated Doppler signal.

Table 1 Comparison of true SBI for four wide band signals and the estimates of SBI determined from the FFT and AR frequency spectra.

Normalised Bandwidth	True SBI	FFT Av. SBI		FFT (Hamming) Av.SBI		AR Av. SBI	
		64	256	64	256	64	256
0.2	0.167	0.133 +/- 0.053	0.131 +/- 0.043	0.141 +/- 0.048	0.123 +/- 0.060	0.129 +/- 0.059	0.146 +/- 0.023
0.4	0.286	0.227 +/- 0.062	0.237 +/- 0.053	0.223 +/- 0.066	0.211 +/- 0.081	0.229 +/- 0.076	0.258 +/- 0.054
0.6	0.375	0.312 +/- 0.076	0.333 +/- 0.038	0.300 +/- 0.089	0.322 +/- 0.068	0.308 +/- 0.080	0.342 +/- 0.026
0.8	0.444	0.378 +/- 0.081	0.396 +/- 0.053	0.379 +/- 0.090	0.394 +/- 0.064	0.397 +/- 0.074	0.426 +/- 0.031

Table 2 Matrix of Pearson Product-Moment Correlation Coefficients for 64-point data frame.

	AR	FFT (Hamming)	FFT
FFT	0.90	0.84	1.0
FFT (Hamming)	0.83	1.0	
AR	1.0		

Table 3 Matrix of Pearson Product-Moment Correlation Coefficients for 256-point data frame.

	AR	FFT (Hamming)	FFT
FFT	0.93	0.91	1.0
FFT (Hamming)	0.90	1.0	
AR	1.0		

Appendix 4

A STUDY OF SPECTRAL BROADENING OF CLINICAL DOPPLER ULTRASOUND SIGNALS

P.I.J. Keeton, F.S. Schlindwein

(Submitted to European Journal of Ultrasound, 1997)

A STUDY OF SPECTRAL BROADENING OF CLINICAL DOPPLER SIGNALS USING FFT AND AUTOREGRESSIVE MODELLING

Abstract - This paper investigates the behaviour of the spectral broadening index (SBI) derived from spectra obtained using autoregressive (AR) modelling compared to that of SBI based on fast Fourier transform (FFT) analysis of clinical Doppler ultrasound scans. Doppler signals from internal carotid arteries of patients with normal and diseased vessels with up to 80% stenosis were analysed. A threshold at -6 dB of the maximum magnitude component of each individual spectrum was implemented to reject low-level noise. The SBI was obtained using the maximum and the mean frequency envelopes extracted from the sonogram. A qualitative improvement in both the appearance of the AR sonograms and the shape of the individual AR spectra was noticeable. The AR approach consistently produced narrower spectra than the FFT and the shapes of the frequency envelopes derived from the AR sonogram and the FFT sonogram were also rather different. Despite these differences a strong correlation was observed between the value of the FFT-based SBI and the AR-based SBI. The mean value of the FFT-SBI is larger than that of the AR-SBI and the variance of the FFT-SBI is smaller than that of the AR-SBI based on a set of at least 20 sequentially recorded heartbeats. It was established that, for all cases where significant stenosis was present, a statistically significant value for SBI could be obtained using 4 or more heartbeats if five spectra around the peak systole were used to estimate the SBI of each individual heartbeat.

key words: Doppler ultrasound, autoregressive model, spectral analysis, spectral broadening, carotid stenosis.

INTRODUCTION

Doppler ultrasound is used clinically to assess stenosis in the carotid arteries. The presence of stenosis may be indicated by disturbed flow distal to the site of stenosis and this causes a broadening of the spectrum of the Doppler signal around peak systole. Nowadays colour Doppler is used for the detection of stenosis, but quantification of its severity can be done using an index extracted from the Doppler spectrum such as the spectral broadening index (SBI).

The spectral estimation of Doppler ultrasound signals is normally performed using the 'short-time Fourier transform' (STFT). Various researchers however have demonstrated that some more modern spectral estimation techniques can produce better spectral estimates, especially when short data frames are required: Marple (1977) demonstrated the better spectral resolution of autoregressive (AR) and autoregressive moving average (ARMA) models compared to FFT; Kitney and Giddens (1986) reported the better performance of AR spectral estimation on spectral tracking and spectral resolution; Kaluzynski (1987) addressed the advantages of using AR spectral estimation for pulsed Doppler signals; Vaitkus *et al.* (1988) stressed the good spectral matching ability of the AR modelling approach under various signal-to-noise ratio (SNR) situations; David *et al.* (1991) tested three modern spectral estimation techniques, including AR, with the STFT approach and concluded that, provided that the model order is chosen appropriately, the modern techniques are superior to the traditional FFT-based approach.

Autoregressive spectral estimation produces more stable spectra from short segments of data and since the AR approach does not assume periodicity, it offers an interesting alternative to FFT-based techniques. This ability to produce spectral estimates from short segments of data is important for the analysis of Doppler signals of arteries with disturbed blood flow such as the kind of signals collected at post-stenotic sites.

Kaluzynski and Palko (1993) studied the behaviour of SBI and other indices under different conditions for the spectrum analysis of simulated signals and concluded that the instability of the spectral estimates (of simulated data) has only a limited effect on the indices derived from the spectrum. Keeton *et al.* (1997) also used simulated data and studied the robustness of Fourier-based and AR-based SBI in noise and the behaviour of AR-SBI with model order. They concluded that although AR had better spectral matching characteristics than the FFT approach there was no significant improvement in the estimation of the SBI by using the AR technique even in the presence of noise.

The main aim of this paper is to test whether the afore mentioned conclusions derived for simulated data are applicable to real clinical Doppler signals, i.e., how the AR spectral estimation technique compares with FFT-based spectral estimation for the quantification of spectral broadening for clinical Doppler signals.

METHOD

In this study real clinical data from both healthy and diseased patients were used as opposed to simulated signals. The SBI has been used as the parameter traditionally associated with the measurement of flow disturbances that occur with stenosis (Harward *et al.* 1986, Labs and Fitzgerald 1992). The SBI was calculated from both the FFT and the AR-based spectra at peak systole and also by averaging the SBI value of five spectra starting at peak systole. The correlation between both FFT-SBI and AR-SBI and the degree of stenosis is presented for 10 documented cases. In all cases more than 20 heart beats were processed and by measuring the SBI of 5 spectra per heart beat, more than 1000 estimations of the systolic spectrum and SBI were performed for each approach.

The spectral broadening index used in this study is defined as:

$$SBI = \frac{f_{max} - f_{mean}}{f_{max}} \qquad [1]$$

where f_{max} = the maximum frequency component in the spectrum above a -6dB threshold (measured from the maximum component of the individual frequency spectra) used to reject low level noise components and f_{mean} = the average mean frequency of the spectral components that exceed this threshold.

In this study forward and reverse flow are combined to produce a single signal with both positive and negative frequency shifts modulated around a programmable 'central' frequency (selected by the operator to be one of the following values: 500 Hz, 1, 2, 4, or 8 kHz). If forward and reverse flow components are present simultaneously then equation [1] is not adequate since it is only suitable for establishing the spectral broadening of a single wide band lobe. In order to overcome this problem the reverse flow was eliminated from the calculation of SBI and only the forward flow components were used to determine the f_{mean} envelope.

Spectrum Estimation

The FFT power spectral density estimation (PSDE) was obtained using N=256 real data points with a Hamming window and producing spectra with 128 frequency bins.

The AR power spectral density is given by either

$$PSD_{AR}(f) = \frac{T\sigma_w^2}{\left| 1 + \sum_{k=1}^{p} a_k \exp(-j2\pi fkT) \right|^2} \qquad [2]$$

where a_k are the autoregressive coefficients, σ_w^2 is the variance of the driving white noise input and T is the sampling period, or approximated by using

$$\hat{PSD}_{AR}(f) = T\sum_{n=-M}^{M-1} \hat{R}_{xx}(n)\exp(-j2\pi nkT) \qquad [3]$$

where $\hat{R}_{xx}(n)$ is an AR-based extrapolation of the biased estimate of the autocorrelation sequence (ACS) $R_{xx}(n)$ derived from the data sequence. The extrapolation is given by

$$\hat{R}_{xx}(n) = -\sum_{k=1}^{p} a_k R_{xx}(n-k). \qquad [4]$$

The AR PSDE can therefore be obtained if σ_w^2 and the coefficients a_k are known. In this paper the method used for estimating the AR parameters is based on the Yule-Walker equations (Kay and Marple, 1981). The results were compared to those obtained using Burg's algorithm (Kay and Marple, 1981) in which the coefficients are obtained directly from the data samples without estimating the ACS. Typical results of this comparison are illustrated in figure 1. The fact that the spectra produced using each of the three above mentioned AR approaches are very similar to one another can be appreciated from figure 1. In this study M = 128 and p = 12 (after Schlindwein and Evans, 1990) were used.

The algorithms for performing AR modelling are computationally more demanding than the FFT approach and a real-time AR system requires more computational power than that of a standard personal computer. Modern DSP boards combined with standard microcomputers allow a flexible software approach to the implementation of real-time algorithms for analysing a wide range of signals. One such system capable of real-time AR spectrum analysis of Doppler ultrasound signals was described by Schlindwein and Evans in 1989.

Real-time Doppler ultrasound

A system has been developed for real-time digital signal processing of the Doppler signal using a Digital Signal Processor (AT&T DSP32C capable of 12.5 MIPS) combined with an IBM PC compatible microcomputer - 486 DX2 66 MHz. The front end to the system is an analogue processor unit that combines the forward and reverse waves into a single signal around a modulation frequency chosen by the operator (Schlindwein *et al.*, 1996). The input is taken from pre-recorded continuous wave Doppler ultrasound scans that have been stored on Digital Audio Tapes, with the forward and reverse signals stored on separate channels. The system can also be used on-line with Doppler signals obtained directly from a Doppler velocimeter.

The spectral frame rate was set at 80 frames per second, corresponding to a data frame length of 12.5 ms irrespective of the sampling rate. All data is processed for sampling frequencies 5.12 kHz, 10.24 kHz or 20.48 kHz, the choice depending on the observed bandwidth. For each frequency estimate 128 frequency bins are computed. Table 1 shows the number of data points and the degree of zero padding, for the FFT approach, for each sample rate to maintain both a standard frame rate and the number of frequency bins.

Off-line analysis

The system was taken off-line so that the same digital data could be analysed using both the FFT and AR algorithms. Ten Doppler ultrasound scans were digitised and stored on disk, each scan comprised of at least 20 sequentially recorded heartbeats. Table 2 contains details of the recordings that were used in this study. All the recordings were taken from internal carotid arteries (ICA) of (i) patients with known diseased arteries ranging in severity of stenosis and (ii) healthy subjects. A

subjective estimate of the degree of stenosis was made at the time of the examination using colour Doppler images of the artery (both axial and cross-sectional views).

For patients with stenosis the site of the scan was around 2 diameters distal to the stenosis. For normal subjects the recordings were taken from a position distal from the carotid branch, which corresponded to the point where the data collection was made for the patients with stenosis. A continuous wave Doppler velocimeter, with a 4 MHz focused 1.5mm beam width, was used for data collection. The beam transfixed the vessel axis at an angle of around 50°.

RESULTS & DISCUSSION

Figure 1(a-d) shows the output for several heartbeats using each of the four methods described (FFT, Yule-Walker AR [2], AR using [3-4] and Burg AR) and figure 1(e-h) shows a typical spectrum obtained from around the systolic peak using each of the algorithms. It can be seen that there is no significant difference between the frequency spectra obtained using the different AR algorithms. There is a distinct qualitative improvement however in the frequency spectra obtained using the AR algorithms over the FFT which is in agreement with other published results (Vaitkus *et al.*, 1988; Kitney and Giddens, 1986; Kaluzynski, 1987; Keeton *et al.*, 1997).

A threshold was applied to the sonogram to reject low-level noise. The threshold was defined as 6dB below the maximum component of the individual frequency spectrum. Figure 2(a,b) shows the raw maximum and mean frequency envelopes extracted from the sonograms obtained using the FFT and AR algorithms respectively. The envelopes have been superimposed on the original sonograms to show the ability of the 6dB threshold to successfully reject low-level noise and retain the true Doppler signal. The two envelopes are then low-pass filtered using a 5-point moving-average filter. It can be seen from figure 2(c,d) that there are significant differences between the envelopes obtained using the FFT and AR algorithms. These differences affect the positioning of the systolic peak in each cardiac cycle. There is also a distinct difference between the relative spectral widths of the frequency spectra

obtained using the FFT and AR algorithms: The AR spectra are noticeably narrower (figure 2).

There are two approaches for producing an estimation of the average value of a parameter derived from the sonograms: The first is averaging spectra from different heart beats, and then calculating the parameter from the averaged spectrum; the second is calculating the parameter from the individual spectra, and only then averaging the values obtained. It was decided to use the second method here to estimate the SBI in order to avoid the inherent difficulties of time alignment of the first technique and the possibility of smearing the averaged spectrum if the alignment is not perfect.

The results of averaging between 1-5 estimates of SBI around each systolic peak (figure 3) were studied to assess whether the effects of spectral instability resulting from the frequency analysis of short data frames can be reduced. The mean value calculated for a particular heartbeat was then averaged over a number of heartbeats (always more than 20) in order to obtain a statistically valid SBI (Labs and Fitzgerald, 1992). The variance of the overall SBI for a particular patient was improved when the estimate of SBI for each heartbeat was obtained using more than one spectrum, as can be observed in figure 4. This was the case for both the FFT and the AR sonograms.

The correlation between the SBIs obtained using the FFT and AR algorithms for the 10 cases documented earlier in the paper is illustrated in figure 4, from where it can be seen that there is a strong correlation between the AR-SBI and the FFT-SBI. If the estimate of SBI is made from the average of several estimates taken from around the systolic peak then there is an improved correlation between AR-SBI and FFT-SBI. The correlation coefficients for using either a single estimate of SBI/heartbeat or 5 estimates of SBI/heartbeat are shown in figure 4. The correlation coefficients for each of the comparisons exceed the 0.1% value, therefore the results are significant at the 0.1% level ($p < 0.001$).

Having established that there is a strong correlation between the FFT-SBI and the AR-SBI, the magnitude and the standard deviation of the overall SBIs calculated using 5 estimates of SBI/heartbeat are compared in figure 5. The magnitude of the AR-SBI is significantly smaller than that of the FFT-SBI which is concurrent with the relative spectral width trends seen in figure 2(c,d). The standard deviation of the AR-SBI estimates are generally larger than the FFT-SBI. The smaller magnitude of the AR-SBI results in a larger coefficient of variation as can be seen in figure 5(c). The SBIs calculated using the FFT algorithms are therefore generally more stable. From figure 5(a) it is also apparent that the magnitude of the SBI in general increases with the severity of stenosis with a noticeable step at around 50% stenosis - between patients 5 and 6 in figure 5(a).

The minimum number of estimations needed to produce a statistically significant value for a parameter depends on the variance of the estimates, the mean value of the parameter and the error and level of confidence desired and is given by

$$n_{5\%,\,95\%} = \left\lceil \frac{1.96\sigma}{0.05\bar{x}} \right\rceil^2 \qquad [5]$$

for a 95% level of confidence with an error less than 5%, if one assumes a normal distribution (Hughes and Graiwoig 1971). The number of estimations needed to produce significant results was calculated and is illustrated in figure 6. For patients with a high degree of stenosis the overall SBI can be determined using fewer individual estimates of SBI: typically $n < 10$ for FFT and $n < 20$ for AR. For a lesser degree of stenosis the mean value of SBI will be lower resulting in a higher value for n. Since 5 estimates of SBI are being taken from each heartbeat and a minimum of 20 heartbeats are being used to estimate the overall SBI this means that the results for all cases where significant stenosis is present are statistically significant. In the case of minor to moderate stenosis care should be taken in interpreting the results due to the statistical instability resulting from the inability to record a sufficient number of sequential heartbeats, this is particularly true for the AR approach (figure 6).

CONCLUSIONS

Based upon the results obtained for the two techniques, it is clear that the AR approach can help improve the quality of the spectra of clinical Doppler signals (figs. 1(e-h)). This agrees with other works based on simulated data. The AR sonograms do not suffer from some of the intrinsic problems that affect the FFT based spectral estimation and hence there is a distinct qualitative improvement in the visualisation of the dynamic flow over the cardiac cycle (figs. 1(a-d)). With the sample cases examined it was found that the AR algorithms all produced similar sonograms, and all the AR algorithms investigated produced smoother sonograms than the FFT approach.

A strong correlation was observed between the value of the FFT-SBI and the AR-SBI. The mean value of the FFT-based SBI is larger than that of the AR-based SBI and the variance of the FFT-SBI is smaller than that of the AR-SBI.

If only one spectrum per heartbeat is used for the estimation then the number of heart beats required in order to produce a statistically significant value for SBI (with a 95% level of significance) is perhaps too large for clinical practice when the degree of stenosis is mild. If five spectra around the peak systole are used for the estimation then, for all cases where significant stenosis was present, a statistically significant value for SBI could be obtained using 4 or more heartbeats.

This research has also shown that despite the qualitative improvement in the individual frequency spectra, there is no quantitative advantage in using the AR approach over the FFT for the determination of SBI due to its poorer variance and the additional computational complexity.

Acknowledgements - The authors wish to acknowledge the collaboration with the Department of Medical Physics of the Leicester Royal Infirmary. In particular we are indebted to Prof. D.H. Evans for making the collection of the data possible, and to A.J. Thrush for her valuable help both in collecting the data and providing the

annotations with the degree of stenosis. We also wish to acknowledge the University of Leicester for the financial support and a scholarship for PIJK

REFERENCES

David JY, Jones SA, Giddens DP. Modern spectral analysis techniques for blood flow velocity and spectral measurements with pulsed Doppler ultrasound. IEEE Trans. Biomed. Eng. 1991; 38:6:589-596.

Evans DH, McDicken, WN, Skidmore, R and Woodcock, JP. Doppler Ultrasound - Physics, Instrumentation, and Clinical Applications. John Wiley & Sons Ltd, Chichester, 1989.

Harward TRS, Bernstein EF, Fronek A. Continuous-wave versus range-gated pulsed Doppler power frequency spectrum analysis in the detection of carotid arterial occlusive disease. Ann. Surg. 1986; 204:1:32-37.

Hughes A and Graiwoig, D. Statistics: A foundation for analysis. Addison-Wesley, London, 1971.

Kaluzynski K. Analysis of application possibilities of autoregressive modelling to Doppler blood flow signal spectral analysis. Med. & Biol. Eng. & Comput. 1987; 25:373-376.

Kaluzynski K, Palko, T. Effect of method and parameters of spectral analysis on selected indices of simulated Doppler spectra. Med. & Biol. Eng. & Comput. 1993; 31:249-256.

Kay SM, Marple Jr. SL. Spectrum analysis - a modern perspective, Proc. IEEE 1981; 69:1380-1419.

Keeton PIJ, Schlindwein FS, Evans DH. A study of the spectral broadening of simulated Doppler signals using FFT and AR modelling. Ultrasound in Med. & Biol. 1997; **in press.**

Kitney RI, Giddens DP. Linear estimation of blood flow waveforms measured by Doppler ultrasound, In: Salamon R, Blum B, Jørgensen M, eds. MEDINFO 86. North Holland: Elsevier Science Publishers B.V. 1986; 5:672-677.

Labs KH, Fitzgerald DE. Quantification of pulsed Doppler spectra for the diagnosis of minor to moderate atherosclerotic lesions: experience from in vitro and in vivo models. In: Labs KH, Jäger KA, Fitzgerald DE, Woodcock JP, Neuerburg-Heusler D, eds. Diagnostic Vascular Ultrasound. London: Edward Arnold, 1992; 126-141.

Marple SL. Resolution of conventional Fourier, autoregressive, and ARMA methods of spectrum analysis. Proceedings of the IEEE international conference on ASSP. Hartford, Conn. IEEE, 1977; 74-77.

Nichols WW, O'Rourke MF. McDonald's Blood flow in arteries, theoretical, experimental and clinical principles, Third Edition. London: Edward Arnold, 1990.

Schlindwein FS, Evans DH. A real-time spectrum analyzer for Doppler ultrasound signals. Ultrasound in Med. & Biol. 1989; 15:3:263-272.

Schlindwein FS, Evans DH. Selection of the order of autoregressive models for spectral analysis of Doppler ultrasound signals. Ultrasound in Med. & Biol. 1990; 16:1:81-91.

Schlindwein FS, Keeton PIJ, Dryden DJ. A simple modulator and mixer for directional Doppler signals. Revista Brasileira de Engenharia, RBE - Caderno de Engenharia Biomédica. 1996; 12:103-111.

Vaitkus PJ, Cobbold RSC, Johnston KW. A comparative study and assessment of Doppler ultrasound spectral estimation techniques part II: Methods and results. Ultrasound in Med. & Biol. 1988; 14:8:673-688.

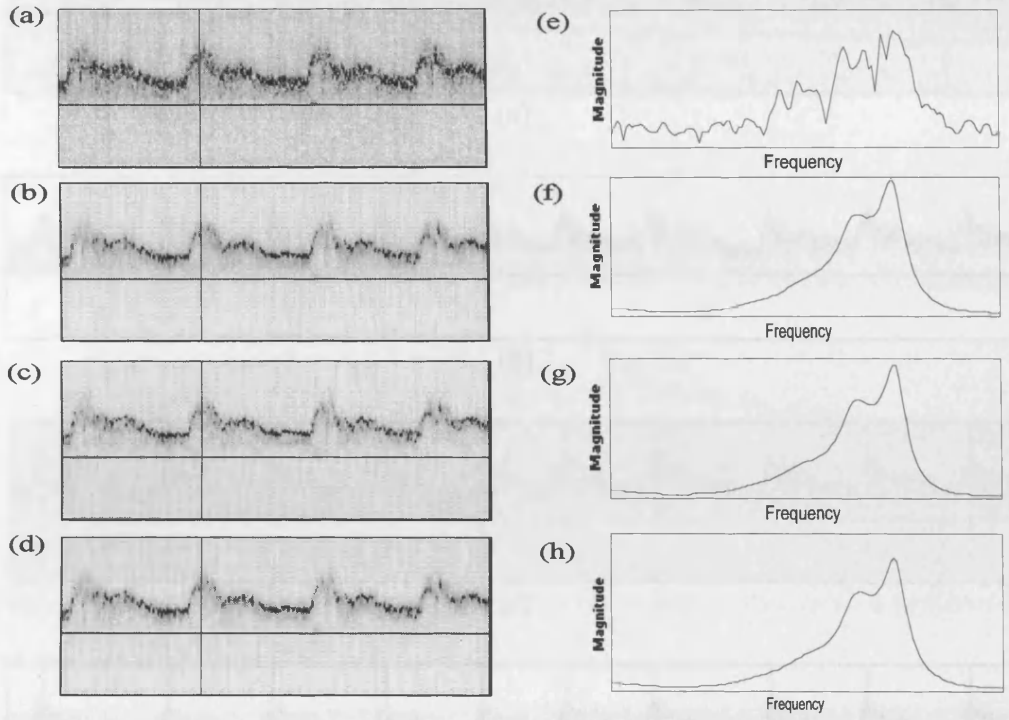

Figure 1 Sonograms of internal carotid arteries (ICA) using (a) FFT, (b) AR with Yule-Walker approach and eqn [3] for PSD, (c) AR with Yule-Walker approach and eqn [5] for PSD, (d) Burg and eqn [3] for PSD, (e)-(h) Frequency spectra taken from respective sonograms around the peak systole of the cardiac cycle (as marked with the vertical line on (a)-(d).

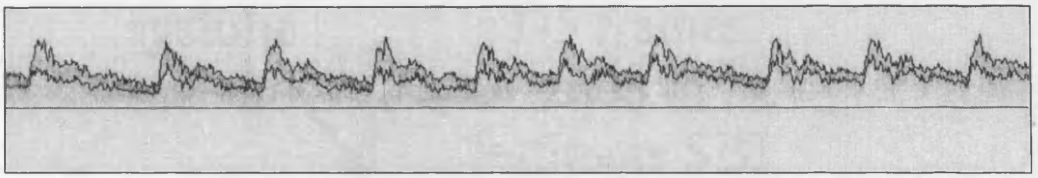

(a)

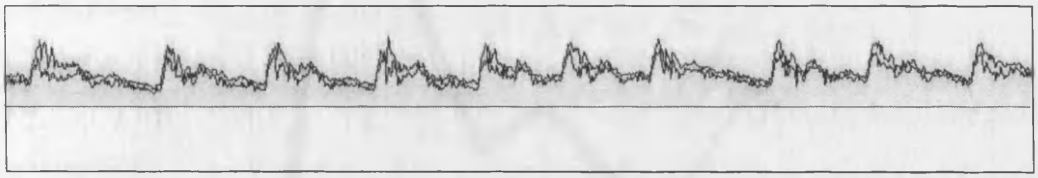

(b)

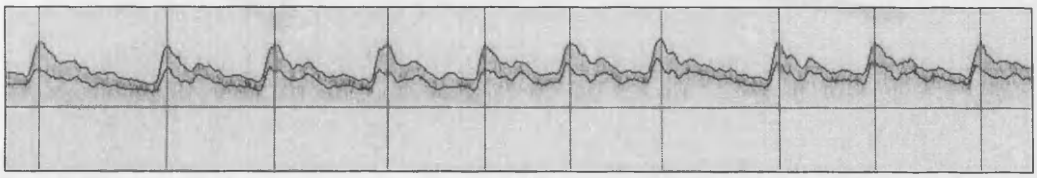

(c)

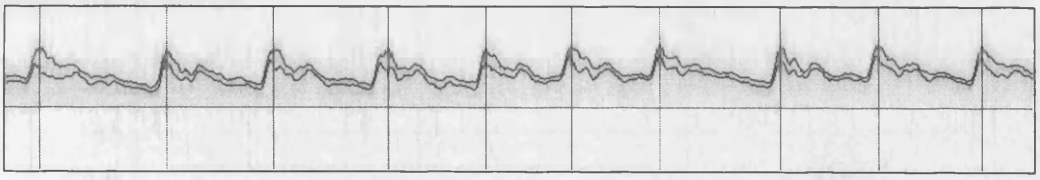

(d)

Figure 2 Raw maximum and mean frequency envelopes extracted from the (a) FFT (b) AR sonograms using the -6dB threshold. Figures (c) & (d) show the low-pass filtered envelopes of (a) & (b) respectively and the comparison between the location of the systolic peaks for each heartbeat based on the FFT and AR sonograms.

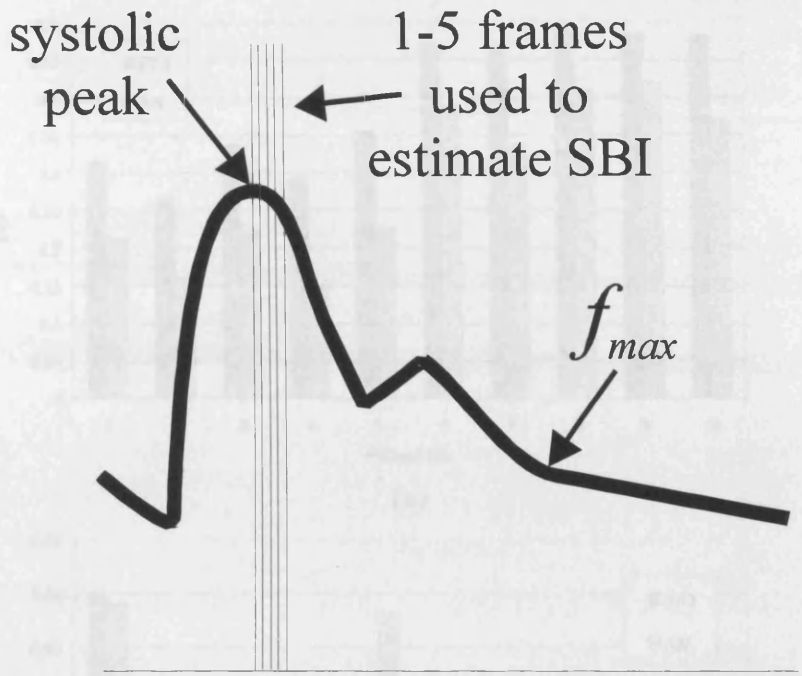

Figure 3 Illustrates the consecutive frames after the systolic peak used to produce an average SBI/heartbeat.

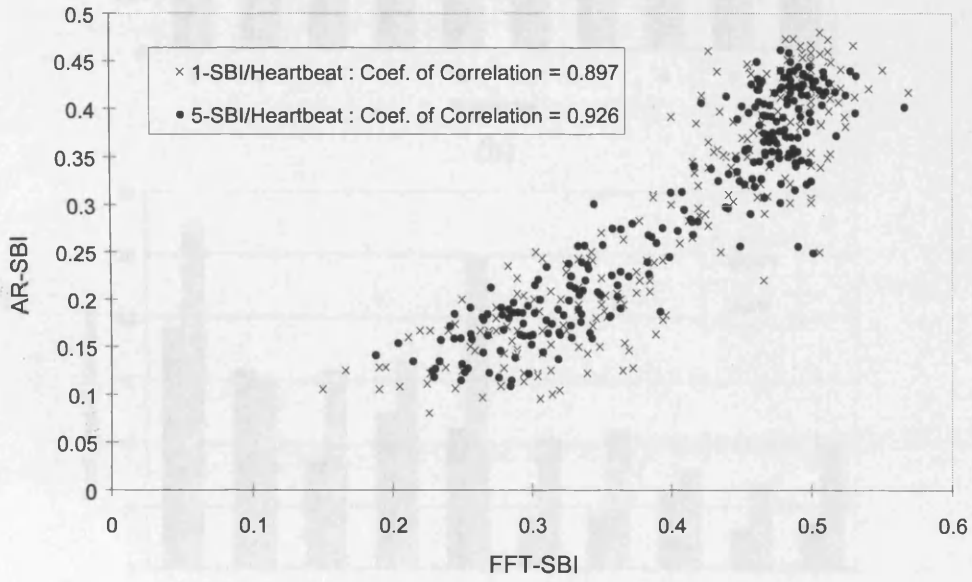

Figure 4 Scattergram of FFT-SBI versus AR-SBI for 1-SBI/Heartbeat and 5-SBI/Heartbeat.

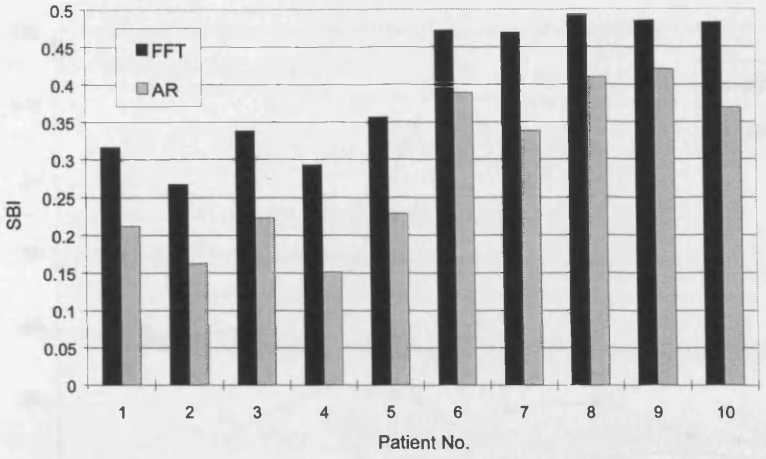

(a)

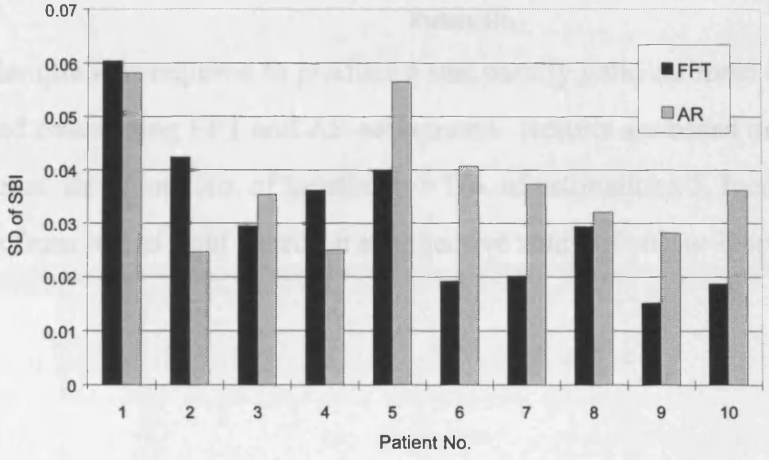

(b)

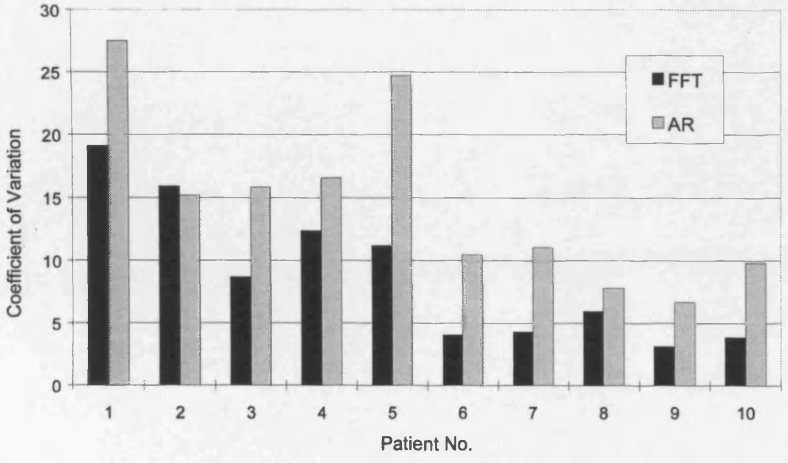

(c)

Figure 5 (a) Magnitude of SBI (b) Standard deviation (c) Coefficient of variation for 10 documented cases using FFT and AR sonograms based on 5 estimates of SBI/heartbeat. Increasing severity of stenosis from left to right based on a subjective study of colour Doppler ultrasound scans.

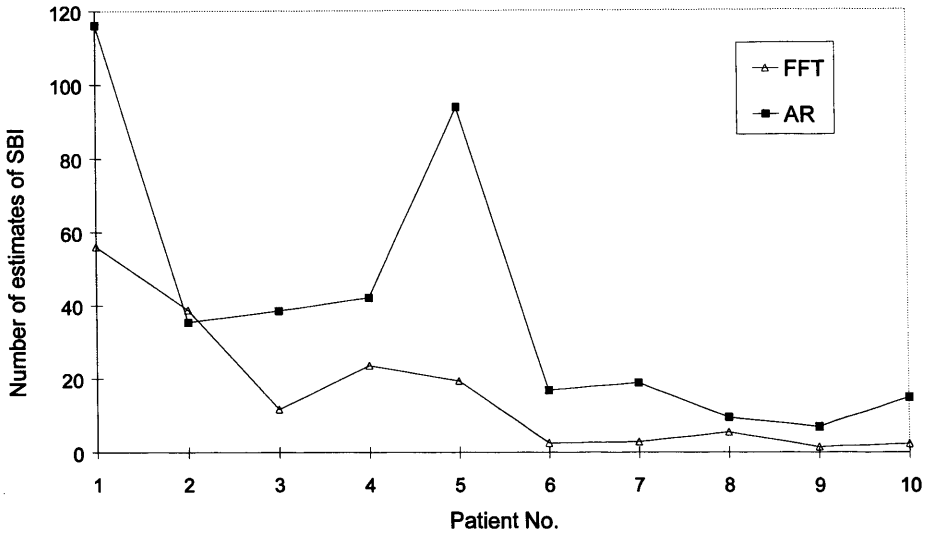

Figure 6 Sample size required to produce a statistically valid estimate of SBI for 10 documented cases using FFT and AR sonograms.. Results are based on 5 estimates of SBI/heartbeat, therefore, No. of heartbeats = No. of estimations/5. Increasing severity of stenosis from left to right based on a subjective study of colour Doppler ultrasound scans.

Table 1 - Zero padding is used to obtain frames of 256 data points and produce spectra with 128 frequency bins.

sampling frequency f_{sam} (Hz)	number of data points	zero padding
20480	256	0
10240	128	128
5120	64	192

Table 2 - Internal carotid artery (ICA) Doppler ultrasound data used in evaluation of AR against FFT.

Signal	Degree of Stenosis (Approx. % of lumen occluded)	Sample Frequency (kHz)
1	0	10.24
2	0 (Slight wall thickening)	10.24
3	20	10.24
4	25	10.24
5	40-50	10.24
6	50	10.24
7	65	5.12
8	>70	10.24
9	80	10.24
10	>80	10.24

Appendix 5

APPLICATION OF WAVELETS IN DOPPLER ULTRASOUND

P.I.J. Keeton, F.S. Schlindwein

(Published in Sensor Review, 1997; 17:1:38-45)

Research articles

Application of wavelets in Doppler ultrasound

P.I.J. Keeton and F.S. Schlindwein

The authors

P.I.J. Keeton is a research student and **F.S. Schlindwein** is a Lecturer in the Department of Engineering, University of Leicester, Leicester, UK. Tel: 0116 252 5321; Fax: 0116 252 2618.

Abstract

Provides an introduction into wavelets and illustrates their application with two examples. The wavelet transform provides the analyst with a scaleable time-frequency representation of the signal, which may uncover details not evidenced by conventional signal processing techniques. The signals used in this paper are Doppler ultrasound recordings of blood flow velocity taken from the internal carotid artery and the femoral artery. Shows how wavelets can be used as an alternative signal processing tool to the short time Fourier transform for the extraction of the time-frequency distribution of Doppler ultrasound signals. Implements wavelet-based adaptive filtering is for the extraction of maximum blood velocity envelopes in the post processing of Doppler signals.

Sensor Review
Volume 17 · Number 1 · 1997 · pp. 38–45
 · ISSN 0260-2288

Introduction

Doppler ultrasound is a non-invasive technique which is widely used in medicine for the assessment of blood flow in intact vessels. The technique has improved much since Satomura first demonstrated the application of the Doppler effect to the measurement of blood velocity in 1959. Doppler velocimeters work by emitting a focused ultrasound beam with a base frequency f into the body via a piezoelectric transducer and detecting the change in frequency that occurs after the beam is reflected or scattered by moving targets. This Doppler shift frequency Δf_D is proportional to the speed of the moving targets:

$$\Delta f_D = \frac{2vf\cos\theta}{c} \quad (1)$$

where v = magnitude of the velocity of target, Δf_D = Doppler shift, f = frequency of transmitted ultrasound, c = magnitude of the velocity of ultrasound in blood, and θ = angle between ultrasonic beam and direction of motion. If the ultrasound probe is positioned so that the beam insonates a blood vessel, the Doppler shift frequency is proportional to the velocity of the blood within the sample volume (1). Since there is a population of targets (mostly red blood cells) scattering back the ultrasound signal, there is a corresponding distribution of Doppler shift frequencies, in other words a wideband spectrum, which ideally corresponds to a histogram of the velocities of all the scattering particles. The band of Doppler shift frequencies normally falls within the audio range (Δf_D < 20kHz). Since the velocity components are proportional to the frequency shifts it is possible to track the velocity distribution by obtaining the power spectral density estimate (PSDE) of the demodulated signal.

For many years derivatives of the Fourier transform (FT)(2) have been the main tools used in signal processing for obtaining the PSDE.

$$X(f) = \int_{-\infty}^{\infty} x(t)\exp(-j2\pi ft)dt \quad (2)$$

The FT in its original form assumes that the signal exists for all time. This for practical purposes is not a realistic assumption and does not give any information about how the signal is changing with respect to time. This is not a problem when the signal being analysed is stationary, that is, the statistical properties of the signal are not changing with time. Most

useful signals, however, are non-stationary; for the analysis of these signals it is necessary to identify and locate the changing frequency characteristics.

The short-time Fourier transform (STFT)(3), also known as the time-dependent or the windowed-Fourier transform (Oppenheim and Schafer, 1989), attempts to analyse non-stationary signals by dividing the whole signal into shorter data frames.

$$X(k) = \sum_{n=0}^{N-1} x(n)w(n-n_0)\exp\left(\frac{-j2\pi nk}{N}\right) \quad (3)$$

The output of successive STFTs can provide a time-frequency representation of the signal. To accomplish this the signal is truncated into short data frames by multiplying it by a window so that the modified signal is zero outside the data frame. The frequency spectrum for the data frame is calculated using the fast Fourier transform (FFT). In order to analyse the whole signal the window is translated in time and then reapplied to the signal.

One of the limitations of the STFT is that the time frame for analysis of the signal is fixed. A more flexible approach would be to use a scalable window: a compressed window for analysing high frequency detail and a dilated window for uncovering low frequency trends within the signal.

Wavelets

The wavelet transform addresses the problem of fixed resolution by using base functions that can be scaled. These "wavelets" act in a similar way to the windowed complex exponentials that are used in the STFT, except that with the wavelet transform the length of signal being analysed is not fixed. Figure 1 illustrates some of the more common wavelets. At first glance it is apparent that wavelets are better suited to analysing transient signals, since they are well localized in time, whereas sinusoids extend over all time. Wavelets are not just well localized in time, they are also well localized in frequency, although not as well as sinusoids. The property of time and frequency localization is known as compact support and is one of the most attractive features of wavelet analysis.

Continuous wavelet transform

The continuous wavelet transform (CWT)(4) compresses or dilates and translates a "mother" wavelet and correlates it with the signal at all times and scales:

$$CWT(a,b) = -\int_{-\infty}^{\infty} x(t)\frac{1}{\sqrt{a}}\psi\left(\frac{t-b}{a}\right)dt \quad (4)$$

Figure 1 Examples of common wavelets

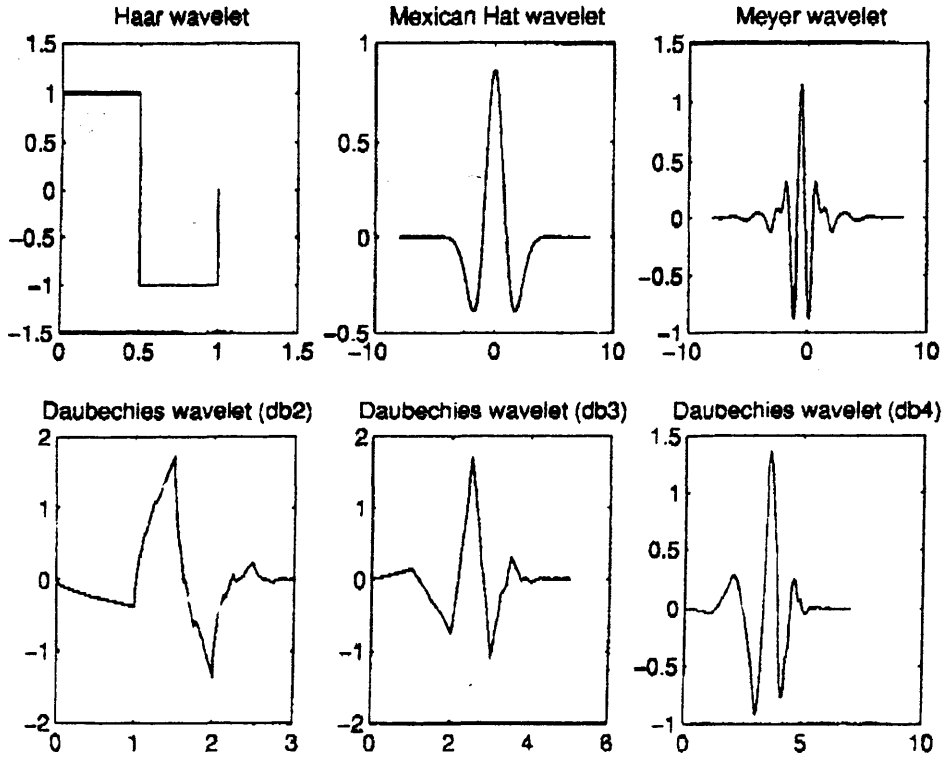

where $\psi(t)$ is the wavelet, b is the translation factor and a is the dilation factor. The output of the transform shows the correlation between the signal and wavelet as a function of time. If the signal and wavelet are a good match then the correlation between the signal and the wavelet is high resulting in a large coefficient. The choice of wavelet depends on the application. The concept of scale in the CWT is analogous to the inverse of frequency in the FT. When the wavelet is highly compressed it extracts the localized high frequency details of the signal. When the wavelet is fully dilated the length of the wavelet is more comparable to the length of the signal and therefore it extracts the low frequency trends of the signal.

Figure 2(a) shows a chirp signal with a sinusoid and an instantaneous spike superimposed. Figure 2(b-c) shows the results of analysing the signal using the STFT and the CWT respectively. The results of the Fourier analysis illustrate the trade-off between time and frequency resolution. If the sinusoid is going to be resolved with any degree of accuracy then the time resolution of the chirp and the localization of the instantaneous spike are compromised. If the time resolution is improved then the frequency resolution becomes inadequate to resolve the sinusoid and the chirp. The CWT pinpoints the instantaneous spike and shows the sinusoid and chirp. The CWT also shows the oscillatory behaviour in the time scale diagram, this corresponds to the relative phase difference between the wavelet and the signal. The CWT does not resolve sinusoidal components as well as Fourier analysis, but its ability to localize transient events makes it an attractive alternative to the STFT.

Wavelets have already been used in biomedical applications for the analysis of the EMG and ECG. In this paper the potential of using wavelets for the analysis of blood flow is explored. The Doppler signal from ultrasound analysis is highly non-stationary and it has been reported that the STFT is not necessarily the best tool for extracting the Doppler shifts from the received signal (Kaluzynski, 1987).

The CWT is computationally intensive and the amount of information that the transform provides is extremely large. A more efficient routine is the discrete wavelet transform (DWT) which looks at the signal at specific dilations of the mother wavelet. The efficiency of the DWT is better than that of the FFT. To understand the DWT and to see where wavelets come from it is necessary to introduce the concept of filter banks which are the discrete equivalent of wavelets.

Figure 2 (a) Signal composed of a chirp with a sinusoid and an impulse superimposed; (b) STFT analysis using a sliding 128-point FFT; (c) CWT decomposition. Scale is the inverse of frequency

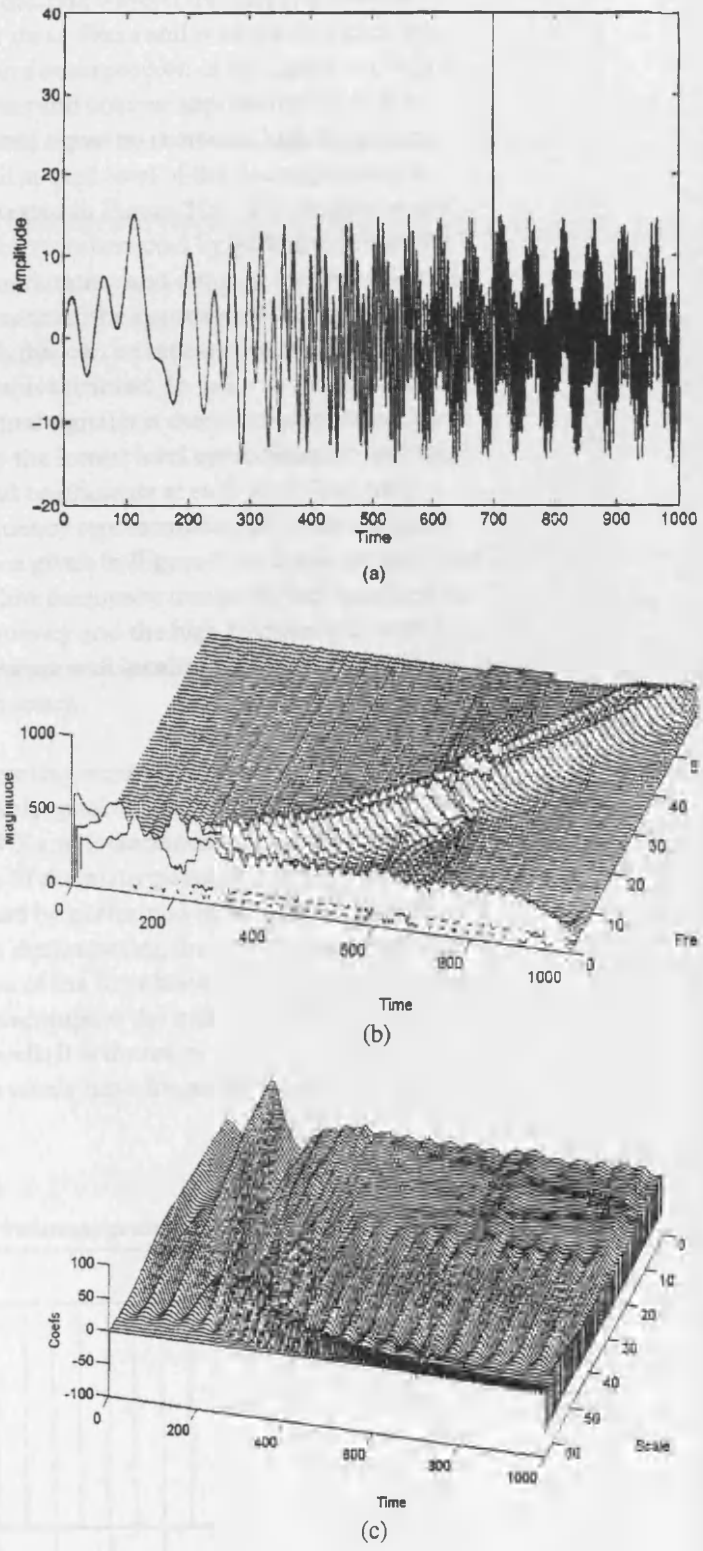

Filter banks and sub-band coding

A single filter has a particular frequency response and after passing a signal through the filter some of the information within the signal is lost. A single filter cannot be used to reconstruct the original signal from the filtered output since once the information has been lost it cannot be retrieved. If two filters are used, one which retains the low frequency information and the other which keeps the high frequency information, they can be designed in such a way that, together, they contain all the original information within the signal. The outputs from these two filters can be combined to reconstruct the original signal. The set of filters used to split the spectral information in such a way is known as a filter bank. It is possible to extend this concept to further decomposition of the signal into finer and finer frequency bands. In order to manage the data which are output from the filters it is necessary to downsample the outputs from the filters; otherwise, as the signal is passed through each level of the filter bank, the amount of data is doubled. Generally the idea is to represent the signal more efficiently. If the signal is decomposed n times it is not beneficial to have 2^n times more data than the original signal. For a set of filters to act as "perfect reconstruction filters" it is necessary that they have special characteristics, since real filters do not have perfect cut-off frequencies and therefore there is some overlap between adjacent filters in order to retain all the information. The filters that were developed in the early 1980s to perform this task are known as quadrature mirror filters (Strang and Nguyen, 1996).

Discrete wavelet transform

The discrete wavelet transform is derived from these filters and is based on a multiresolution decomposition of the signal to give a coarser and coarser approximation to the original signal by removing high frequency detail at each level of the decomposition as illustrated in Figure 3(a). The original signal can be reconstructed by adding together the approximation and detail at the lowest level to reconstruct the approximation at the higher level, this can be repeated until the original signal is retrieved. In order to reconstruct the original signal it is therefore only necessary to keep the lowest level approximation and the detail coefficients at each level. The time-frequency representation of the decomposition is given in Figure 3(b). It can be seen that the low frequency trends are well localized in frequency and the high frequency components are well localized in time but not in frequency.

Wavelet packet analysis

Wavelet packet analysis is an extension of the DWT and it turns out that the DWT is only one of the many possible decompositions that could be performed on the signal. Instead of just decomposing the low frequency component of the filter bank each time, it is possible to decompose the high frequency component as well. It is therefore possible to subdivide the whole time-frequency plane into different

Figure 3 (a) DWT decomposition of a time varying signal (b) Time-frequency representation using the DWT

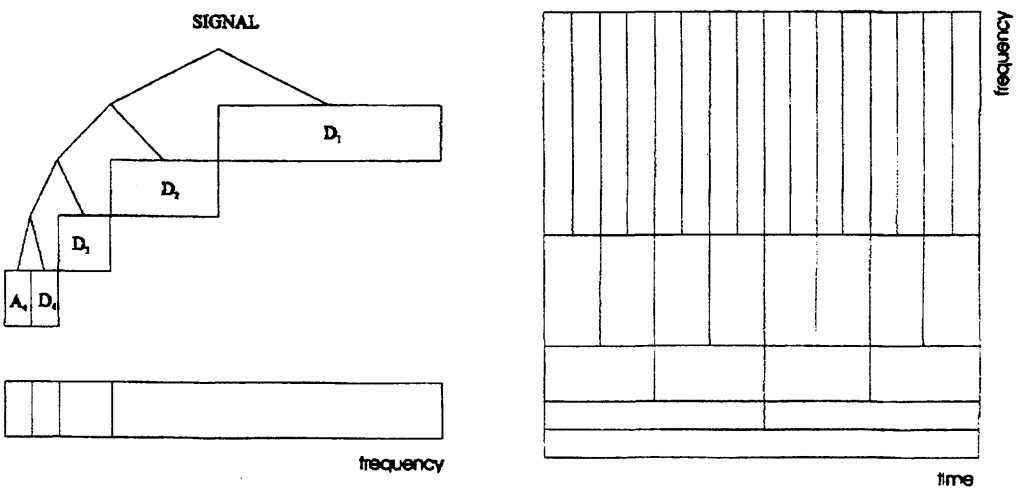

time-frequency tilings as can be seen from Figure 4.

The advantage of wavelet packet analysis is that it is possible to combine the different levels of decomposition in order to achieve the optimum time-frequency representation of the original signal.

Application of the wavelet transform to real signals

The concept of being able to decompose a signal totally and then perfectly reconstruct the signal again is nice, but it is not particularly useful by itself. In order to make use of this tool it is necessary to manipulate the wavelet coefficients to identify characteristics of the signal that were not apparent from the original time domain signal.

The next section of this paper outlines how the wavelet transform could be used to analyse Doppler ultrasound signals. First, the extraction of a time-frequency sonogram of the Doppler signal using wavelet analysis compared to the conventional FFT is considered. Second, wavelets are used in the post-processing of the sonogram to extract information that can be used for clinical diagnosis.

Time-frequency analysis of the Doppler signal

Conventional time-frequency analysis of the Doppler signal involves processing the signal using data frames of fixed duration. The maximum size of the data frame is dependent on the stationarity of the signal. The Doppler signal is highly non-stationary and therefore a time frame of <20ms is used for analysis. If a longer time frame is used then the transient behaviour of the blood flow becomes blurred. Figure 5 shows typical FFT sonograms for two Doppler signals; the first signal is taken from the internal carotid artery and the second from the femoral artery. The sonograms show the periodic heartbeats and within each beat it is possible to visualize the systolic and diastolic flow as the heart contracts and then relaxes. The same signals were analysed using wavelet packet analysis (db3 wavelet, see Figure 1); the results are shown in Figure 5.

Wavelet packet analysis produces a time-frequency sonogram that resembles the FFT result. The advantage of wavelet packet analysis over the FFT is the optimization of the time-frequency resolution.

Doppler ultrasound systems process the acquired data in real-time. This is done using

Figure 4 (a) Total decomposition of a time varying signal using wavelet packet analysis (b) Time-frequency representation for each level of the decomposition

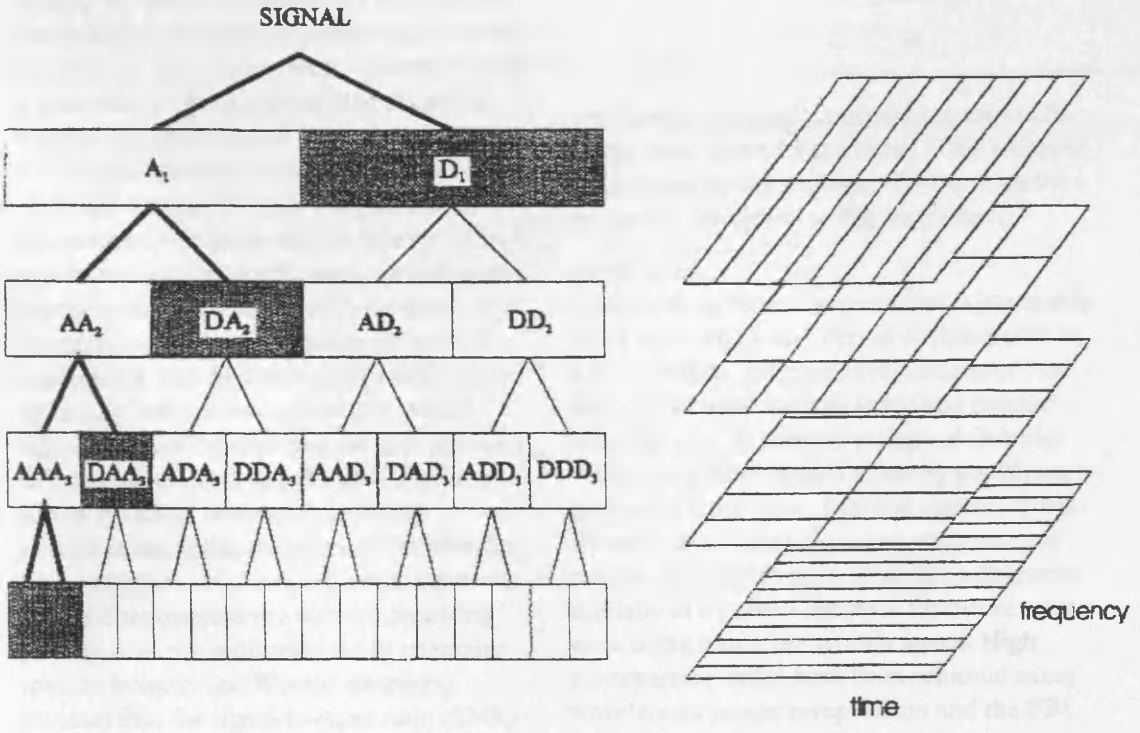

the FFT by processing the data in frames that comply with the stationarity of the signal. For real-time wavelet packet analysis it would be necessary to split the data into larger frames (≈ 1s). This is because wavelet packet analysis performs a total decomposition on a signal so that all the possible time-frequency representations are available. An algorithm is used to find the optimum combination of these levels to produce the sonogram. The computational requirements for producing a time-frequency sonogram using wavelet packet analysis are far greater than for the conventional FFT.

Figure 5 Sonograms obtained using a 256-point FFT for two Doppler ultra sound scans (a) internal carotid artery (b) femoral artery

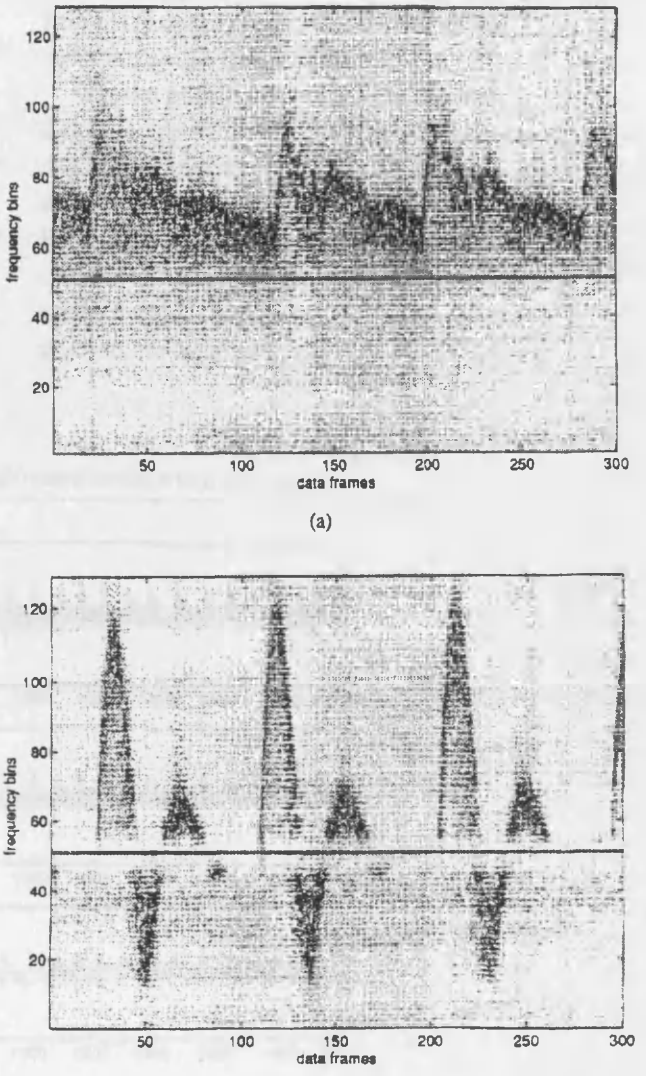

Post-processing of the time-frequency decomposition

Following time-frequency analysis of the Doppler signal a clinical diagnosis is made to determine whether or not the vessel is diseased. This decision is based on a subjective study of the sonogram and further quantitative analysis. One kind of quantitative analysis is performed using frequency envelopes extracted from the sonogram. The envelopes are normally obtained by selecting a threshold to reject low level noise components. Figure 6 shows the extracted frequency envelopes for the two examples over a number of cardiac cycles.

Denoising the frequency envelopes

The raw frequency envelopes extracted from a sonogram are usually corrupted by noise (Figure 6) due to statistical instability in the PSDE. To use the envelopes for quantitative analysis it is necessary to remove the noise to identify the underlying trends. Conventionally a five-point moving average filter is used to denoise the envelopes as illustrated in Figure 6. Discrete wavelet analysis can be used to clean the noisy envelopes. A signal can be represented efficiently using a relatively small number of wavelet coefficients, assuming an appropriate selection of wavelet is made. It is therefore possible to discard small wavelet coefficients that do not contain a significant amount of information about the overall signal. The coefficients that are kept are used to reconstruct the original signal. Figure 6 shows the effect of denoising the raw envelopes using this technique. This wavelet-based adaptive approach has advantages over a fixed filter because the wavelet denoising process does not indiscriminately attenuate specific frequencies. Wavelet denoising assumes that the signal-to-noise ratio (SNR) is significantly large, that is, that the noise coefficients are small compared to the coefficients used to model the signal. This seems to be the case for the removal of noise from the frequency envelopes as Figure 6 shows.

Compression

Computer storage of patient data is invaluable since data can be transferred electronically in a few seconds. Efficient compression of data ensures minimal storage space and quicker data transfer. At present storage of Doppler ultrasound recordings is made by producing a grey scale hard copy of several cardiac cycles. By using 2-D wavelet image compression it may be possible to store recorded sonograms digitally in a patient database for future reference using much less storage space. High compression ratios have been achieved using wavelets for image compression and the FBI are using wavelets to reduce storage space

required to store their database of fingerprints.

Conclusions

Wavelet analysis provides an interesting alternative to conventional Fourier methods. The compact support inherent in wavelets makes wavelets an attractive choice for time-frequency analysis. Fourier analysis has a fixed time-frequency resolution. This means that there is a trade-off between the resolution of transient events and underlying trends in the signal. Wavelets are better suited to resolving transient events since wavelet analysis incorporates the concept of scale into the equation, which gives the analyst the flexibility to look at the time domain signal at different resolutions: a compressed wavelet for analysing high frequency detail and a dilated wavelet for detecting underlying trends.

Wavelet packet analysis of the Doppler signal can provide a time-frequency decomposition similar to that of the FFT. The advantage of wavelet packet analysis over the FFT is the ability to optimize the time-frequency decomposition of the sonogram. The computational overheads for wavelet packet analysis are far greater than for FFT analysis

Figure 6 Envelopes extracted from the STFT-sonograms: (a) Internal carotid artery; (b) Femoral artery

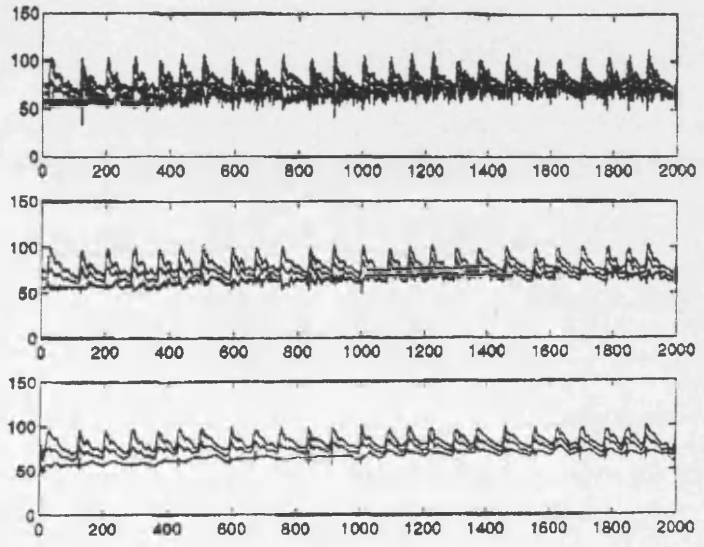

(a) **Top**: Raw envelopes. **Middle**: 5-point MA filtered envelopes. **Bottom**: Wavelet denoised envelopes.

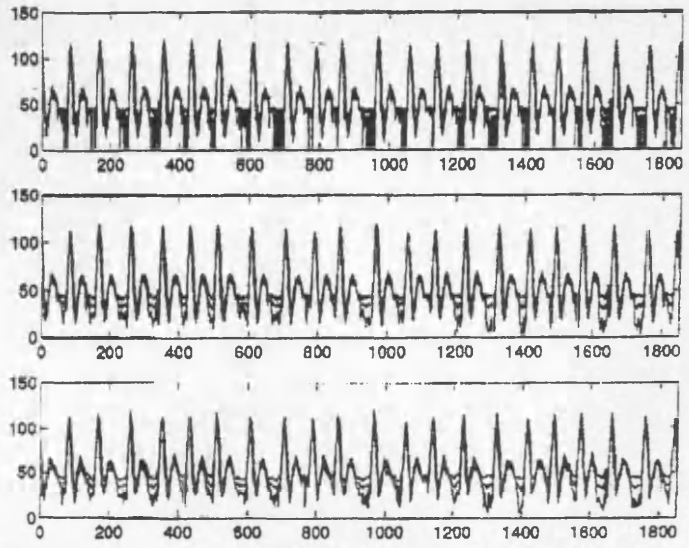

(b) **Top**: Raw envelopes. **Middle**: 5-point MA filtered envelopes. **Bottom**: Wavelet denoised envelopes.

of the Doppler signal making real-time implementation difficult.

The DWT can be used to denoise the frequency (blood velocity) envelopes extracted from the sonograms by manipulating the wavelet coefficients. Large wavelet coefficients contain significant information about the signal while small coefficients are expendable. If a signal is contaminated with noise and the SNR is high, then it is possible to remove the noise by implementing a threshold to remove small coefficients.

This paper has presented the fundamental ideas behind continuous and discrete wavelet analysis. The application of wavelets to Doppler ultrasound has been investigated and potential areas for further research have been outlined.

References

Kaluzynski, K. (1987), "Analysis of application possibilities of autoregressive modelling to Doppler blood flow signal spectral analysis", *Medical & Biological Engineering & Computing*, Vol. 25, pp. 373-6.

Oppenheim, A.V. and Schafer, R.W. (1989), *Discrete-Time Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ.

Strang, G. and Nguyen, T. (1996), *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Cambridge, MA.

Appendix 6

SONOGRAMS OF SEVERAL SEQUENTIALLY RECORDED HEARTBEATS MEASURED USING CLINICAL DOPPLER ULTRASOUND AND PROCESSED USING THE SHORT-TIME FOURIER TRANSFORM

This appendix contains examples of the ten clinical Doppler ultrasound signals recorded at the Leicester Royal Infirmary and used in the examination of spectral broadening in chapter 7. Both healthy subjects and patients with diseased internal carotid arteries with varying degrees of stenosis were used in the study. Figure 1 illustrates the sonograms of several sequentially recorded heartbeats obtained for each subject. The signals were analysed using the short-time Fourier transform with a data frame of 12.5 ms. For patients with stenosis the site of the scan was around 2 diameters distal to the stenosis. For normal subjects the recordings were taken from a position distal from the carotid branch, which corresponded to the point where the data collection was made for the patients with stenosis. A continuous wave Doppler velocimeter, with a 4 MHz focused 1.5 mm beam width, was used for data collection. The beam transfixed the vessel axis at an angle of around 50°. A subjective estimate of the degree of stenosis was made at the time of the examination using colour Doppler images of the artery (both axial and cross-sectional views).

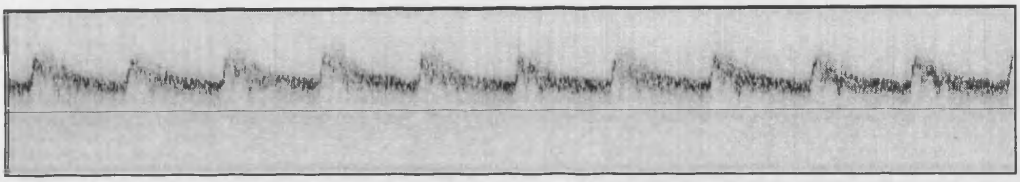

Patient 1 ≈ 0 % stenosis; f_{samp} = 10240 kHz
(slight wall thickening)

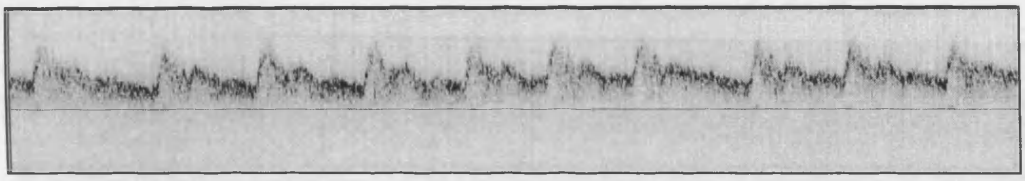

Patient 2 ≈ 0 % stenosis; f_{samp} = 10240 kHz

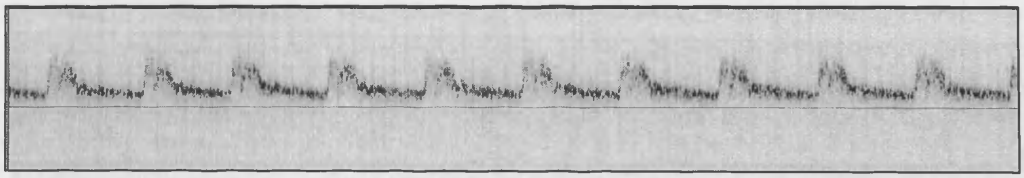

Patient 3 ≈ 20 % stenosis; f_{samp} = 10240 kHz

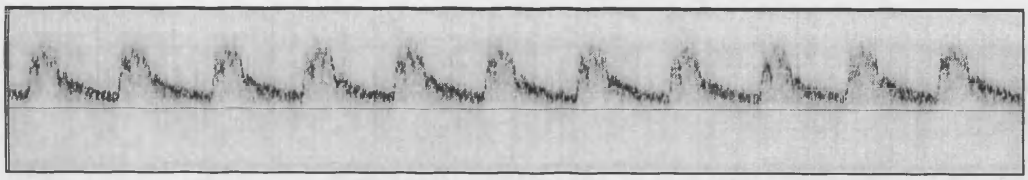

Patient 4 ≈ 25 % stenosis; f_{samp} = 10240 kHz

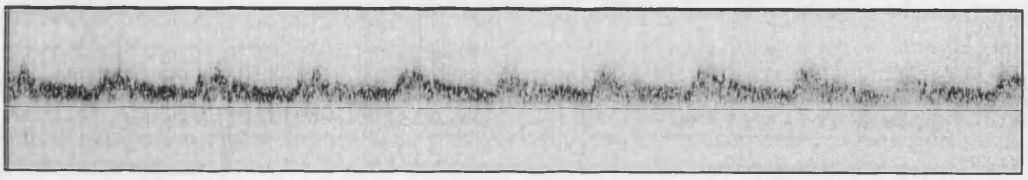

Patient 5 ≈ (40 - 50) % stenosis; f_{samp} = 10240 kHz

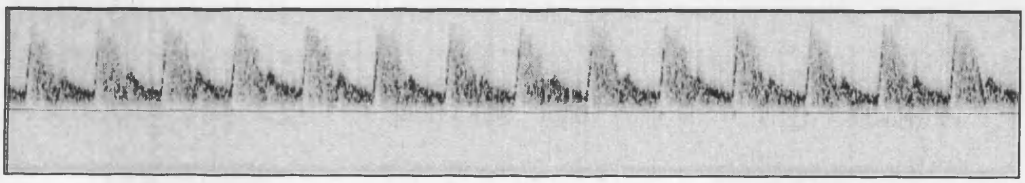

Patient 6 ≈ 50 % stenosis; f_{samp} = 10240 kHz

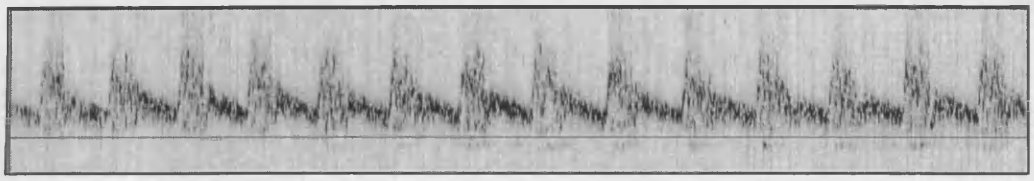

Patient 7 ≈ 65 % stenosis; f_{samp} = 5120 kHz

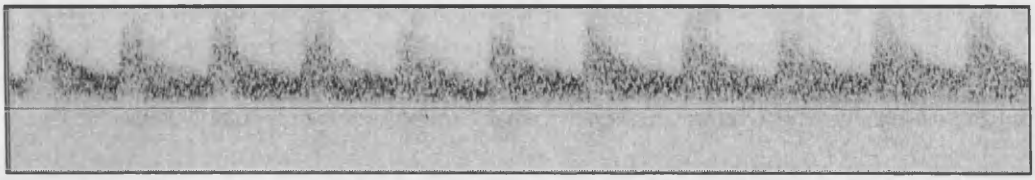

Patient 8 > 70 % stenosis; f_{samp} = 10240 kHz

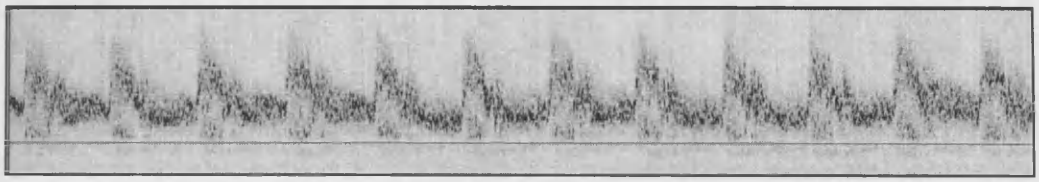

Patient 9 ≈ 80 % stenosis; f_{samp} = 10240 kHz

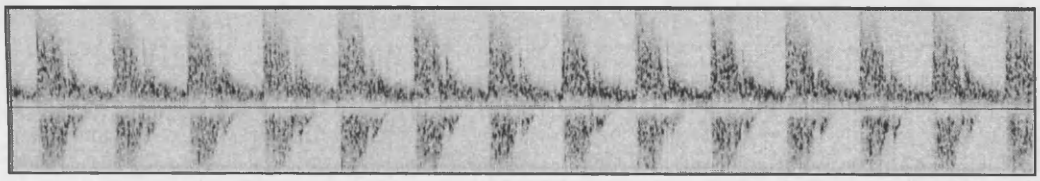

Patient 10 > 80 % stenosis; f_{samp} = 10240 kHz

Figure 1 FFT sonograms illustrating several heartbeats recorded from each of the 10 subjects used in chapter 7.