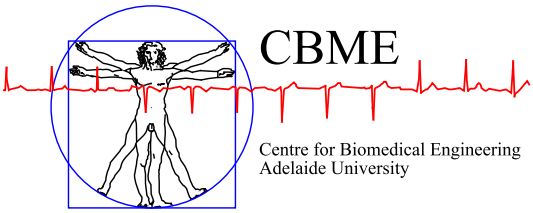

Signal Processing Techniques for Phonocardiogram De-noising and Analysis

by

Sheila R. Messer

B.S., University of the Pacific, Stockton, California, USA

Thesis submitted for the degree of

Master of Engineering Science

Adelaide University
Adelaide, South Australia

Department of Electrical and Electronic Engineering
Faculty of Engineering, Computer and Mathematical Sciences

July 2001

Contents

Abstract

The focus of this thesis is the de-noising and representation of phonocardiograms for subsequent analysis. The PCG has been proven to be a clinically significant diagnostic tool while being inexpensive, non-invasive, reliable and cheap. However, the PCG is corrupted by noise from a number of sources including thoracic muscular noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), peristaltic intestine noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), respiratory noises, foetal heartbeat noise if the subject is pregnant, noise caused by contact with the instrumentation and ambient noise. Thus, there is a need to de-noise the PCG signal. Because it is a complex, non-stationary signal, traditional methods of de-noising are not appropriate. Phonocardiogram de-noising techniques, which are explored, include wavelet de-noising, optimised wavelet de-noising, wavelet packet de-noising, the matching pursuit technique, and averaging. The time-frequency and time-scale de-noising methods performed roughly equally while removing significant amounts of noise from the signal. However, optimised wavelet de-noising performed slightly better than the other methods; thus, optimised wavelet de-noising in conjunction with averaging is recommended to be used in appropriate cases. Once the PCG has been de-noised, different methods of extracting features from the PCG and classifying the PCG according to this information were explored. The use of phase space diagrams, HT diagrams, instantaneous signal parameter extraction, and phase synchronisation between the ECG and PCG were investigated, but these investigations were limited by the quantity and quality of data available. The results presented are only indicative results, but they demonstrate that further work to investigate the use of these techniques with larger amounts of data would be worthwhile. Recommendations for future research in the area of phonocardiogram de-noising and classification are provided.

Statement of Originality

I hereby declare that this work contains no material which has been accepted for the award of any degree or diploma in any university or other tertiary institution and to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being available for loan and photocopying.

Sheila Renee Messer

Acknowledgements

This thesis could not have been completed without the help and support of a number of people. I would like to take this opportunity to thank them.

First, I would like to thank my supervisor Dr. Derek Abbott whose many ideas, guidance and enthusiasm have been a source of motivation to me. I also would like to thank my co-supervisor Dr. John Agzarian for his input and assistance in procuring the PCG/ECG recordings.

This research has been carried out in the Electrical and Electronic Engineering Department of Adelaide University. I would like to acknowledge everyone who made available their assistance, support and advice. Particularly, I wish to thank Prof. R. E. Bogner and Associate Prof. B. Davis for their willingness to answer questions at various stages. Prof. R.E. Bogner, Prof. J. Mazumdar, Associate Prof. C. Bertram, and Bradley Ferguson have been very helpful for kindly proof reading the manuscript and offering helpful corrections and suggestions. I would also like to thank my colleagues at the Centre for Biomedical Engineering, Stefan Enderling, Bradley Ferguson, Leonard Hall, Greg Harmer, Sam Mickan and Joseph Ng, for their general moral support throughout my time here, willingness to answer questions, and many good Friday night pub sessions. I would like to express appreciation to Leonard Hall, Jarrad Maple, and Mohammad Ali Tinati for their previous work on this project and allowing me to use their equipment.

I would like to thank the Rotary Foundation for their financial support in providing me with a scholarship. I would also like to acknowledge the support provided by local Rotary members especially those of the Adelaide South Rotary Club.

Last but not least, I would like to thank my family for their support throughout my academic career, especially my mother who always has encouraged me to learn beginning an early age.

Publications

Messer, S., Abbott, D. & Agzarian, J. (2000), Optimal wavelet denoising for smart biomonitor systems, *in* 'Proceedings of SPIE: Smart Electronics and MEMS, Vol. 4236', SPIE, pp. 66–78.

Messer, S., Abbott, D. & Agzarian, J. (2001*a*), Comparison of automatic de-noising methods for phonocardiograms with extraction of signal parameters via the hilbert transform, *in* 'Proceedings of the SPIE: Photonics West, Vol. 4304', SPIE, San Jose, California. In Press.

Messer, S. M., Abbott, D. & Agzarian, J. (2001*b*), 'Optimal wavelet denoising for phonocardiograms', *Microelectronics (Elsevier)* . Accepted.

List of Figures

List of Tables

Glossary

A/D	Analogue to Digital
ALE	Adaptive Line Enhanced
AR	Auto-regressive
CBME	Centre for Biomedical Engineering
CWT	Continuous Wavelet Transform
dB	Decibel
DWT	Discrete Wavelet Transform
ECG	Electrocardiogram
EHG	Electrophysterography
FFT	Fast Fourier Transform
FT	Fourier Transform
HRV	Heart Rate Variability
HT	Hilbert Transform
IDWT	Inverse Discrete Wavelet Transform
MP	Matching Pursuit
NRMSE	Normalised Root-mean-square Error
PC	Personal Computer
PCB	Printed Circuit Board
PCG	Phonocardiogram
QRS	QRS Complex–Waves on the ECG
S1	First Heart Sound
S2	Second Heart Sound
S3	Third Heart Sound
S4	Fourth Heart Sound
SA	Sinoatrial
SNR	Signal-to-noise-ratio
SPL	Sound Pressure Level
STFT	Short Time Fourier Transform
WD	Wigner Distribution
WP	Wavelet Packet
WT	Wavelet Transform

Table 1 Glossary

Chapter 1

Introduction

"Research is the process of going up alleys to see if they are blind."

Marston Bates (1906-1974)

AMERICAN ZOOLOGIST

1.1 Introduction

While the ancient Greeks believed that the human heart was the source of intelligence and others were convinced that it was the origin of emotions, it is now known that these theories are false (Marieb 1991). The diagnosis and treatment of heart conditions has greatly benefited from the technological leaps that have been made in the last century such as digital radiography, computed tomography, magnetic resonance imaging and ultrasound techniques (Macovski 1983, Rangayyan & Lehner 1988). Older and fundamental techniques such as auscultation (listening to sounds produced by the heart) have been pushed out of the spotlight with the advances made in these recent, flashier technologies (Rangayyan & Lehner 1988). Auscultation possesses many unsolved problems including controversy over the genesis of heart sounds (Chen, Durand & Lee 1997, Cloutier, Grenier, Guardo & Durand 1987, Durand & Pibarot 1995, Ewing, Mazumdar, Goldblatt & Vollenhoven 1985, Ewing, Mazumdar, Vojdani, Goldblatt & Vollenhoven 1986, Luisada 1965, Luisada 1980, Rangayyan & Lehner 1988), lack of quantitative analysis techniques (Rangayyan & Lehner 1988) and noise corruption (Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley 1990*a*, Groom, Herring, Francis & Shealy 1956, Hall 1999, Jiménez, Ortiz, Peña, Charleston, Aljama & González 1999, Karpman, Cage, Hill, Forbes, Karpman & Cohn 1975, Maple, Hall, Agzarian & Abbott 1999, Sava, Pibarot & Durand 1998, Zhang, Durand, Senhadji, Lee & Coatrieux 1998). Although the knowledge of the heart has increased, there remains much to be learned about the heart with many common heart conditions and diseases that continue to afflict a large percentage of the population. Phonocardiography remains in an evolutionary phase of development with a need to overcome the previously mentioned problems in order to facilitate its use as a routine medical diagnostic tool (Rangayyan & Lehner 1988).

Auscultation is of great importance because many pathological conditions of the heart produce murmurs and abnormal heart sounds before other diagnostic tools, such as an electrocardiogram (ECG), might detect them (Rangayyan & Lehner 1988). Thus, auscultation is the standard test used by physicians to assess the condition of the heart (Rangayyan & Lehner 1988). Auscultation also has advantages over other tests in that it is very easy to perform, non-invasive (Durand & Pibarot 1995), cheap and requires little

time. However, because the genesis of heart sounds is not well understood (Chen, Durand & Lee 1997, Cloutier et al. 1987, Durand & Pibarot 1995, Ewing et al. 1985, Ewing et al. 1986, Luisada 1965, Luisada 1980, Rangayyan & Lehner 1988), conclusions of a pathological condition should not solely be based on auscultation. The human ear possesses limitations which make it a non-ideal instrument for auscultation (Brown, Smallwood, Barber, Lawford & Hose 1999, Luisada 1965, Luisada 1980, Rangayyan & Lehner 1988, Zhang, Durand, Senhadji, Lee & Coatrieux 1998).

The limitations of the human ear were realized and devices designed to overcome this deficiency were introduced including various versions of the stethoscope and phonocardiography which is actually a recording of the heart sound (Bell, Long, Langham, Kos & Parten 1998, Cheitlin, Sokolow & McIlroy 1993, Grenier, Gagnon, Genest, Durand & Durand 1998, Guo, Moulder, Durand & Loew 1998, Lehner & Rangayyan 1987, Liang, Lukkarinen & Hartimo 1998, Luisada 1965, Lukkarinen, Korhonen, Angerla, Nopanen, Sikio & Sepponen 1997, Lukkarinen, Nopanen, Sikio & Angerla 1997, Lukkarinen, Sikio, Nopanen, Angerla & Sepponen 1997, Maple et al. 1999, Nandagopal, Mazumdar, Karolyi & Hearn 1981, Selig 1993, Shen & Sun 1997, Shino, Yoshida, Mizuta & Yana 1997). Phonocardiography has many advantages over traditional auscultation in that it can be replayed, analysed for spectral and timing information, and processed in different ways which may reveal hidden information.

Phonocardiography has not progressed as quickly as other diagnostic methods because lack of standards in equipment and recording locations and poor understanding of the heart sound mechanisms and the complexity of PCGs, but signal processing has aided in the development of phonocardiography (Wood & Barry 1995). An important use of signal processing in the area of phonocardiography is noise removal. There are many sources of noise which may pollute the PCG including thoracic muscular noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), peristaltic intestine noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), respiratory noises, foetal heart sound noise if the subject is pregnant, noise caused by contact with the instrumentation and ambient noise. Thus, the problem of noise removal in phonocardiography is very important for the development of phonocardiography as a widely used clinical tool. In the current study, various methods of de-noising phonocardiograms are explored and studied. This study also explores various

methods of extracting notable features from phonocardiograms.

The primary aims of this chapter are to supply the reader with the knowledge of the heart necessary to comprehend this thesis, provide justification for the research being conducted and perform a literature review of research in this area. In Section 1.2, a brief description of the human heart is given. The next two Sections 1.3 and 1.4 describe heart sounds and electrical activities of the heart. In Section 1.5, a literature review of the recent research in the area of phonocardiography being studied is presented. The final section briefly summarises the scope of the research presented in this thesis and provides motivation for the work.

In order to clarify terminology used in this thesis, a couple of words will be defined. For our purpose, a heartbeat is defined as a "single complete pulsation of the heart" (*The American Heritage Dictionary of the English Language* 2000), and thus, the term will be used to refer to a complete cycle of the heart. Heart sounds are an acoustical phenomenon which are caused by actions of the heart and are explained in Section 1.3.

1.2 Brief Description of the Heart

This section presents a brief description of the heart so that the reader may gain a basic understanding of the organ. The reader is referred to Passamani (2000) for a more detailed description. The heart is a hollow, muscular organ which is primarily composed of muscle tissue that contracts rhythmically to distribute blood throughout the body. The circulatory system consists of the heart, blood, and blood vessels. The purpose of the circulatory system is to distribute oxygen and nutrients to the body and to collect waste products. The heart is the most important organ in the body because if the heart stops and oxygen is not delivered to the brain, death will result within minutes.

The shape of the heart may be compared to an upside-down pear; it is situated slightly to the left in the chest cavity.

The heart has four chambers. The left and right atria, the upper chambers, receive inflow of blood from the veins. The right and left ventricles, the lower chambers, pump blood into arteries which distribute the it throughout the body. The right and left side of the

heart are separated from each other by a wall of tissue.

Figure 1.1 depicts the systemic and pulmonary circuits. Oxygen deficient, waste carrying blood flows into the heart and before it leaves the heart, the waste is removed and the blood is re-oxygenated. Oxygen depleted blood containing carbon dioxide enters the heart from the superior vena cava and inferior vena cava into the right atrium. It then passes through the right-antrioventricular valve, (a valve prevents blood from flowing backward in the heart and opens easily in the direction of blood flow but closes when blood attempts to flow the other way) known as the tricuspid valve, into the right ventricle. The right ventricle pumps the blood through the right semilunar valve, also known as the pulmonary valve, to the pulmonary artery which connects to the lungs. The lungs remove waste from the blood and replenish the oxygen supply. The blood is then conducted through the pulmonary veins to the left atrium where it passes through the left atrioventricular valve, known as the bicuspid or mitral valve, into the left ventricle. Finally, the left ventricle pumps the oxygen rich blood through the left semilunar valve (aortic valve), to the aorta, the largest artery. The blood is then distributed throughout the body by the network of blood vessels.

Both sides of the heart contract simultaneously, producing a heartbeat. The cardiac cycle or events of a single heartbeat has two phases: diastole, where the chambers of the heart are relaxed, and systole in which the chambers contract to pump blood. During the systole period, the atria contract before the ventricles resulting in efficient circulation of the blood.

The atria and ventricles relax and the atrioventricular valves open during diastole which allows blood to enter from the veins into the atria and then to the ventricles. At the beginning of systole, the atria contract to completely fill the ventricles. Then the ventricles contract pushing blood out the semilunar valves into the arteries. Meanwhile, the atrioventricular valves close so that blood does not flow backwards into the atria. When pressure has risen enough in the arteries, the semilunar valves close to prevent blood from flowing backwards into the ventricles. After the heart muscle relaxes, diastole begins again.

The sinoatrial (SA) node controls the contraction of the heart. The heart muscle is unique

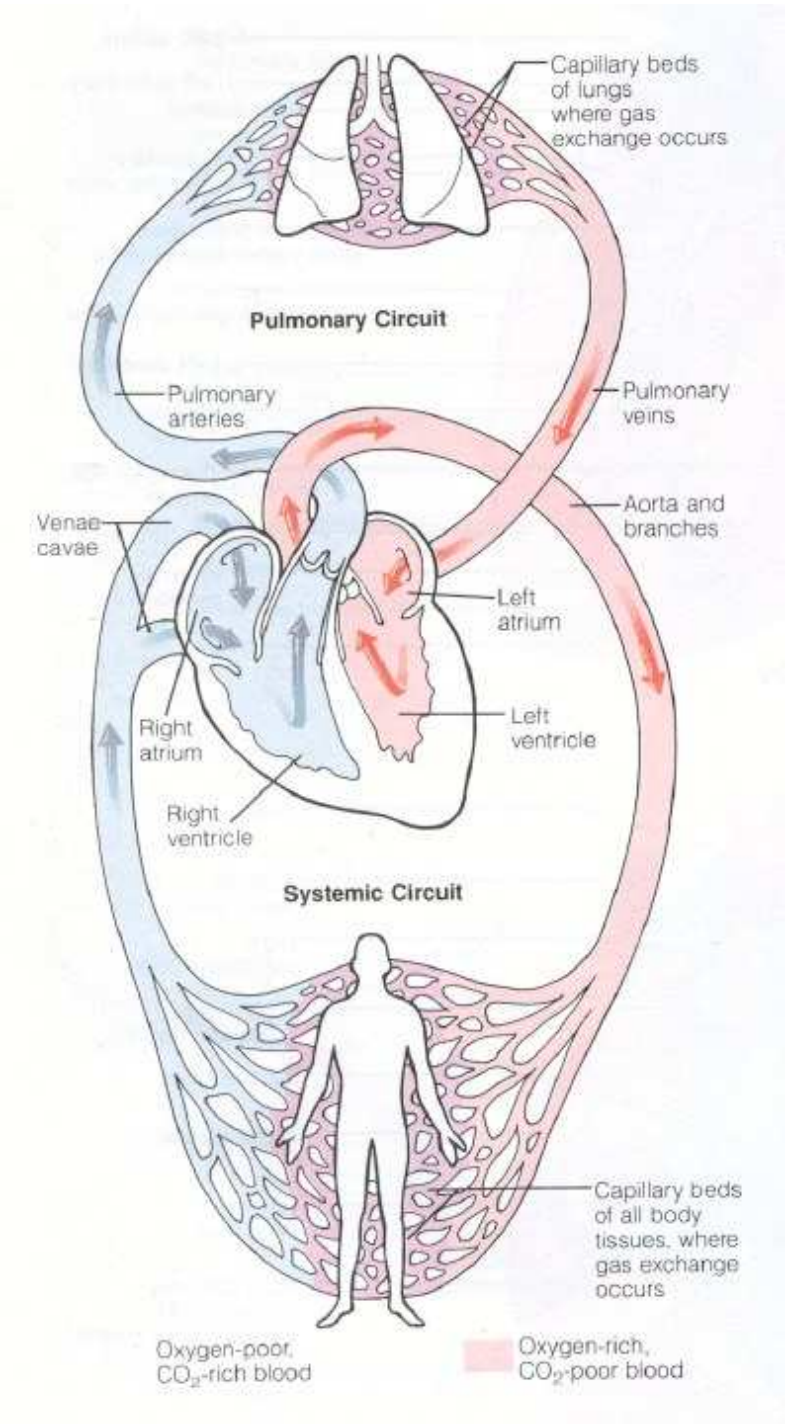

Figure 1.1 This diagram shows the systemic and pulmonary circuits where the left side of the heart acts as a pump for the entire system and the right side of the heart is the pump for the pulmonary system (Marieb 1991).

because it can contract on its own and does not need outside nerve impulses. Certain heart muscles may contract spontaneously, and these cells generate electrical signals which are conducted throughout the heart to produce a regular beat. The SA node is located in the upper right-hand corner of the right atrium and generates electrical signals faster than cells anywhere else in the heart, thus it sets the pace for the muscle to contract.

1.3 Heart Sounds

There are four heart sounds: S1, S2, S3 and S4. The oldest theory, valvular theory, states that heart sounds are caused by the vibrations which are produced during the opening and closing of the heart valves (Chen, Durand & Lee 1997, Leatham 1975). The cardiohemic theory was coined by Rushmer (1952) and further developed by Luisada, MacCanon, Coleman & Feigen (1971) and Luisada (1972). The cardiohemic theory asserts that heart sounds are caused by vibrations of the whole cardiac structure after the speed of the intracardiac blood flow is changed as a result of the opening and closing of the cardiac valves (Chen, Durand & Lee 1997). A new theory, which combined the older two theories, was presented in the 1980's by Durand & Guardo (1982), Durand, Genest & Guardo (1985), and Durand & Pibarot (1995). They believed that the intracardiac PCG is a summation of the vibrations produced by different cardiac structures, and the thoracic PCG is influenced by time-varying properties of the heart-thorax acoustic system (Chen, Durand & Lee 1997). Table 1.1 presents some basic characteristics of the four heart sounds.

Heart Sound	**Duration in secs.**	**Frequency Range (Hz)**
First Heart Sound	0.1-0.12	20-150
Second Heart Sound	0.08-0.14	50-600
Third Heart Sound	0.04-0.05	20-50
Fourth Heart Sound	0.04-0.05	Less than 25

Table 1.1 Some basic heart sound characteristics (Ewing 1989)

1.3.1 The First Heart Sound

The first heart sound occurs at the beginning of ventricular contraction when the mitral and tricuspid valves are closing (Durand & Pibarot 1995). It indicates the beginning of ventricular systole.

According to valvular theory, S1 is composed of two major high-frequency components which may be heard (Durand & Pibarot 1995). The high-frequency sounds, M1 and T1, correspond to the halting of the mitral and tricuspid valves when they are closing.

According to cardiohemic theory, the heart chambers and arterial vessels may not vibrate without forcing the enclosed blood to vibrate as well (Durand & Pibarot 1995). This theory is based on the assumption that the heart chambers, valves and blood vessels compose an interdependent system that vibrates as a whole. According the cardiohemic theory, S1 begins with a low-frequency component occurring when the first myocardial contractions begin after a rise in ventricular pressure (Durand & Pibarot 1995). The second component is comprised of a higher frequency caused by tension of the left ventricular structures, the myocardium contracting and the deceleration of the blood. The third component occurs when the aortic valve opens and is caused by the sudden acceleration of blood into the aorta. The fourth component occurs due to turbulence in the blood flowing quickly through the aorta. The right heart does not significantly contribute to the generation of the first heart sound.

1.3.2 The Second Heart Sound

The second heart sound (S2) occurs at the end of ventricular systole and the beginning of ventricular relaxation, after the closure of the aortic and the pulmonary valves (Durand & Pibarot 1995). S2 usually possesses higher frequency components than S1 and is generally shorter in duration than S1 (Tahmasbi 1994).

According to valvular theory, S2 is composed of two high-frequency components, A2 and P2, which may be caused when the aortic and pulmonary valve leaflets are stopped at the end of closure (Durand & Pibarot 1995). The cardiohemic theory states that S2 is caused when blood slows down and the flow is reversed into the aorta and pulmonary arteries

(Durand & Pibarot 1995).

1.3.3 The Third and Fourth Heart Sounds

Compared to the first two heart sounds, the third heart sound is quite weak, occurring during ventricular filling in early diastole once the mitral and tricuspid valves have opened (Durand & Pibarot 1995). The third heart sound is a low-pitched, localised sound occurring approximately 0.12-0.16 seconds after S2 (Cheitlin et al. 1993, Ewing 1989). The third heart sound (S3) is not normally heard except in young patients because of the thinness of their skin on the chest and it may be a pathological indicator of ventricular failure if heard in patients older than 40 years old (Durand & Pibarot 1995, Tahmasbi 1994).

The fourth heart sound (S4) is produced in late diastole when the ventricle fills as a result of atrial contraction (Durand & Pibarot 1995). If the fourth heart sound is heard, it is considered a clinical indicator of left ventricular hypertrophy and coronary artery disease (Durand & Pibarot 1995) A study done by Jordan, Taylor, Nyhuis & Tavel (1987) also concluded that an audible S4 is evidence of cardiac abnormality.

According to valvular theory, S3 and S4 are generated in the left ventricle when the mitral valve leaflets are halted at the end of opening (Durand & Pibarot 1995). Cardiohemic theory states that the deceleration of blood into the ventricles during the early and late diastolic stages, when the ventricles are being filled, causes the sounds to be generated (Durand & Pibarot 1995).

1.4 Electrical Activity of the Heart

The cardiac muscle has the intrinsic ability to depolarise and contract and is independent of extrinsic nerve impulses as is demonstrated by transplanted hearts which continue to beat rhythmically even after all nerve connections are severed (Marieb 1991).

The cardiac conduction system is composed of cardiac cells which do not contract but are specialised to initiate and distribute electrical impulses throughout the heart (Marieb

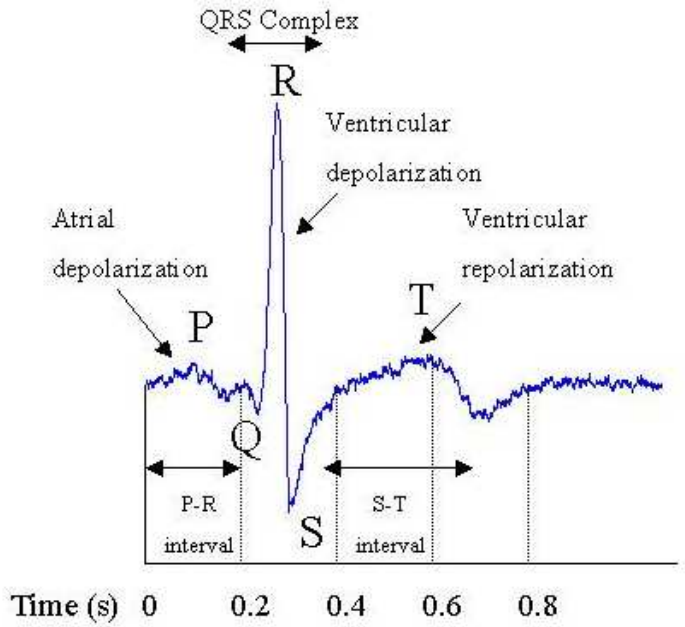

Figure 1.2 An electrocardiogram trace showing the three deflection waves and the intervals, modified from Marieb (1991).

1991). This allows the myocardium to depolarise and contract in an orderly fashion from the atria to the ventricles (Marieb 1991).

In Marieb (1991), the three deflection waves contained in a normal ECG are explained (See Figure 1.2). The first wave is small and is known as the P wave. The P wave lasts about 0.08 seconds and is caused by the movement of the depolarisation wave from the sinoatrial node to the atria. About 0.1 seconds after the P wave starts, the atrium contracts. The next wave, the QRS complex, is typically the most obvious. The QRS complex is a result of ventricular depolarisation and begins prior to ventricular contractions. The shape of the QRS complex indicates the size difference of the two ventricles and the time needed for each ventricle to depolarise. The QRS complex usually lasts about 0.08 seconds. The configuration and length of the QRS complex are not affected by the force of contraction or heart rate variation (Cheitlin et al. 1993). The final wave, the T wave lasting about 0.16 seconds, is a result of ventricular re-polarisation. Because re-polarisation is slower than depolarisation, the T wave is more distributed than the QRS wave and thus has a lower amplitude. Atrial re-polarisation normally occurs simultaneously with ventricular

excitation.

The P-R or P-Q interval has a duration of about 0.16 seconds and lasts from the beginning of the atrial excitation to the beginning of ventricular excitation (Marieb 1991). The Q-T interval (about 0.36 seconds) represents the time from the beginning of ventricular depolarisation until re-polarisation during which the ventricles contract (Marieb 1991). The duration of re-polarisation of the heart (S-T interval) becomes shorter as the heart rate increases (Cheitlin et al. 1993).

1.5 Literature Review

The aim of this section is to give the reader a background in the areas of phonocardiography dealt with in this study. These areas include de-noising of phonocardiograms and classification and feature extraction of PCGs in relation to their use as a diagnostic tool. Excellent reviews of research performed in the area of phonocardiography containing hundreds of references are given by Rangayyan & Lehner (1988) and Durand & Pibarot (1995).

1.5.1 Time-Frequency and Time-Scale Decomposition Based De-noising

Wavelets, wavelet packets, the matching pursuit algorithm, and other time-frequency and time-scale decomposition based de-noising algorithms are useful tools in the area of de-noising complex non-stationary signals which often are biomedical signals. This section reviews a few of these de-noising studies.

A signal is decomposed using the wavelet or wavelet packet transform and because of their structure, coherent signals can be represented in a few, large coefficients. The noise is disorderly and scattered throughout the signal, and therefore is represented by a large number of small coefficients which are thresholded and the signal then reconstructed.

Another method which may be used to effectively de-noise highly non-stationary signals is the matching pursuit method. The signal is expanded into waveforms which are called

atoms that are selected from a large dictionary of functions to match the local signal structures (Akay 1997). The decomposition is stopped at a certain number of time-frequency atoms when there are no remaining coherent structures that may be extracted from the signal. The remaining parts of the signal are incoherent structures and are assumed to be noise. The coherent structures are recomposed to form the signal leaving the noise discarded.

Matalgah & Knopp (1994) present multi-resolution filters based on wavelets with an analytic solution for the scale parameter that provides high stability relative to the level of additive noise. The filter performance is evaluated for detecting monochromatic signals embedded in white noise and may be shown to be comparable with matched filters.

Bertrand, Bohorquez & Pernier (1994) employ the wavelet transform to filter and analyse brain-evoked potentials. Generalised Wiener filtering in the wavelet domain is performed to de-noise the signal.

The wavelet transform is applied to electroencephalograms (EEGs) in order to decompose the signals so that they very closely resemble long ensemble averages (which effectively eliminate noise) at certain scales (Lim, Akay & Daubenspeck 1995). Thus, the need for long trial recordings may be successfully eliminated.

Whitmal, Rutledge & Cohen (1996) used a wavelet based method for extracting speech from background noise in digital hearing aids, and the result is an improvement over many of the previous methods.

Prochaázka, Mudrová & Storek (1998) demonstrated that wavelet de-noising is effective in removing additive noise from air pollution signals.

Carré, Leman, Fernandez & Marque (1998) employed the wavelet transform to de-noise uterine electrophysterographies (EHGs) which are corrupted by electronic and electro-magnetic noise that overlaps with the desired signal spectra making classical de-noising unsuitable. Two algorithms were used in this study. The first method uses the algorithm "á trou" with non-symmetrical filters that require little computation time and the results are comparable with classical de-noising techniques. However, some artefactual peaks are created by the first de-noising method. The second method which uses orthogonal wavelets

is an improvement upon the first method. After thresholding the decomposed signal, the results correspond to the "average of all circulant shifts de-noised by a decimated wavelet transform." The results of the second method are better than those of the first because the averaging performed in the method eliminates many of the peaks generated by the first method.

Wavelet transform domain filtering was used as an adaptive de-noising tool to remove heart sound noise from lung sound recordings (Hadjileontiadis & Panas 1998). Results of the study indicate that the WT domain filtering applied to lung sound recordings remove most of the superimposed heart sound noise producing a relatively noise free signal.

The discrete wavelet transform was shown to dramatically reduce the noise in heart sounds (Hall 1999, Maple et al. 1999, Maple 1999). Of the orthogonal wavelets employed in this study, the authors suggested using Daubechies wavelets of order 6 or 7 with a decomposition level between 5 and 7.

Ferguson & Abbott (2000) reported that wavelets were successfully used to de-noise terahertz pulses in a bio-sensor system.

Wavelet-based noise removal techniques were used to remove noise from biomechanical acceleration signals which were obtained from numerical differentiation of displacement data (Wachowiak, Rash, Quesada & Desoky 2000). The performance of wavelet-based de-noising was compared with the performance of four traditional de-noising techniques used in biomechanics. The study concluded that the wavelet-based noise removal techniques were very effective in removing noise from these signals and even more effective, in some cases, than the conventional methods used.

Zhang, Durand, Senhadji, Lee & Coatrieux (1998) reported that the matching pursuit method is quite effective in removing Gaussian noise from PCGs. Good results were obtained with quite high noise levels present. The normalised root-mean-square error (NRMSE) was computed between the original and reconstructed signals and the results showed that the NRMSE was around 2.2% for various PCGs tested.

Krishnan & Rangayyan (2000) used a number of signal decomposition techniques including wavelet, wavelet packets, and the matching pursuit method to de-noise knee joint vibration

signals, which are non-stationary and multi-componential in nature. Thus, there is no way of knowing *a priori* the noise content. Because noise exists in the desired signal spectrum, traditional filtering would not be very effective. The authors concluded that the matching pursuit method worked the best of the three methods followed by wavelet packet de-noising.

Jiménez, Ortiz, Peña, Charleston, Aljama & González (1999) evaluated the influence of de-noising via mutual WPs to detect the first and second heart sounds in foetal phonocardiograms. The envelope of the signal was then calculated using the Hilbert Transform. S1 and S2 were identified using the envelope and the results were compared with and without the de-noising. It was concluded that the de-noising improves the SNR and detection rate of the heart sounds.

A study was conducted by Jiménez, Ortiz, Peña, Charleston, Aljama & González (2000) which evaluated the influence of the wavelet and level of decomposition on two adaptive sub-band schemes, for processing foetal PCGs, which use Mutual Information and Wavelet Packets. Simulated fetal phonocardiograms containing different amounts of noise were filtered with combinations of the adaptive sub-band schemes where the level of decomposition and wavelets were varied. Two indices defined by the authors quantified the performance of the de-noising method. The optimal level of decomposition was found to be three and the Daubechies wavelet of the sixth order appeared to perform the best. The Mutual Information method gave better results than the Wavelet Packet Algorithm.

1.5.2 Other De-noising Methods

Because the desired signal, the PCG, and the noise spectrum overlapped, a modified adaptive line enhanced (ALE) filter was used to remove the background noise from heart sounds (Tinati, Bouzerdoum & Mazumdar 1996, Tinati 1998). Feedback was used to convert the original adaptive line enhancement filter into a pole-zero system. The basic concept of the modification to the ALE filter is the addition of poles into the pass-band region of the ALE filter. This scaled and delayed the output which was added to the input generating poles inside the unit circle. The modified ALE filter was tested using a sinusoid with additive white noise. The modified filter proved to perform better. Tests

also confirmed an increase in performance of the modified ALE filter over the original ALE filter when applied to heart sounds. The parameter values of the modified ALE filter like the number of filter taps, decorrelation delay, the feedback delay parameter, and the number of taps in the feedback loop were chosen on the basis that poles were to be placed close to the unit circle in the band-pass region. By positioning these poles in this manner, the frequency response curve of the filter in the pass-band filter was improved while the sidelobes were attenuated resulting in better performance.

1.5.3 Time-Frequency and Time-Scale Analysis

The Fourier Transform (FT) is not a very good tool to analyse the PCG because important time events and features are lost and the spectral resolution is not that good (El-Asir, Khadra, Al-Abbasi & Mohammed 1996, Ewing et al. 1985, Ewing et al. 1986, Matalgah, Knopp & Mawagdeh 1998, Sawada, Ohtomo, Tanaka, Tanaka, Yamakoski, Terachi, Shimamoto, Nakagawa, Satoh, Kuroda & Limura 1997). This is essentially due to the fact that with Fourier analysis it does not matter *when* frequency components appear in a signal, because the signal is integrated over *all* time. However, some studies were performed on heart sounds using the FFT to obtain information about their frequency content (Yoganathan, Gupta, Corcoran, Udwadia, Sarma & Bing 1976, Yoganathan, Gupta, Udwadia, Miller, Corcoran, Sarma, Johnson & Bing 1976) and Yoganathan (1976) provides several recommendations on how the FFT may be used and computed in order to avoid common pitfalls encountered when using the technique.

Using FFT analysis and linear predictive coding techniques, Hearn, Gopal, Ghista, Robinson, Tihal, Mazumdar & Bogner (1983), Nandagopal, Bogner & Mazumdar (1980), Nandagopal et al. (1981), Nandagopal & Mazumdar (1981), Nandagopal (1984), and Nandagopal, Mazumdar & Bogner (1984) investigated the spectral analysis of heart sounds and vibration analysis of the mitral and aortic heart valves and ventricular chamber. The aim of this work was to determine the valvular and myocardial properties and their pathologies with a non-invasive test. The heart sound signal analysis and the corresponding structural vibration analyses were correlated to develop this technique. The linear predictive coding procedure was found to have a higher frequency resolution for

analysing heart sounds than the FFT method.

The Short-Time Fourier Transform (STFT) is an improvement over the FT, but normally, the STFT does not provide adequate time-frequency resolution of the PCG compared to modern methods such as the wavelet transform. Spectrograms of phonocardiograms are reported as early as the 1940's and may trace their origins to the spectral analysis of speech (Geckeler, Likoff, Mason, Riesz & Wirth 1954, Potter 1947, Riesz 1949). Winder, Perry & Caceres (1965) provided an improvement to the clarity of the spectrogram by using contour displays. Spectrograms have proven to be of clinical use in identifying murmurs, and aortic ball variance in patients with valve prostheses (Geckeler et al. 1954, Hylen, Kloster, Herr, Starr & Griswold 1969). Wood & Barry (1994) concluded that the first heart sound consisted primarily of stationary or almost stationary components and superimposed impulses; thus, quasi-stationary methods like the STFT may be sufficient.

Improved time-frequency results may be obtained using the STFT if certain adjustments are made. Djebbari & Reguig (2000) demonstrated that, by adjusting the size of the sliding time window using a Hamming window, an acceptable result may be reached. The duration of the heartbeat, spectral content and time-frequency views may be seen. Furthermore, Jamous, Durand, Langlois, Lanthier, Pibarot & Carioto (1992) reported that a time-window duration of 16 to 32 ms for a spectrogram of a PCG in dogs provided adequate resolution.

Another alternative to the FFT as a tool to analyse spectral content is the maximum entropy method. Ewing et al. (1985) and Ewing et al. (1986) used the maximum entropy method instead of the FFT to analyse the third heart sound in children because the FFT did not provide adequate resolution.

Debiais, Durand, Pibarot & Guardo (1997) and Debiais, Durand, Guo & Guardo (1997) compared various time-frequency representation techniques, including the spectrogram, Choi-Williams distribution and Bessel distribution, in their ability to represent six different heart murmurs. The basic characteristics of several types of murmurs were determined and then were simulated in order to test the previously mentioned time-frequency representation techniques. It was found that no single technique was optimal for all 6 murmurs, but that a spectrogram using a Hamming window of 30 ms provided the best compromise

of the methods tested.

Another study compared several time-frequency methods, including the spectrograph, STFT and Binomial time-frequency transform, for their capacity to represent the first heart sound (Wood, Buda & Barry 1992). It was found that the Binomial representation provided the best resolution of the time-frequency details. The spectrogram proved to be better than the spectrograph because the spectrograph resulted in large amounts of smearing and the Fourier transforms have better resolution frequency estimates than bandpass filters which are used for the spectrograph.

Einstein, Kunzelman, Reinhall, Tapia, Thomas, Rothnie & Cochran (1999) reported that time-frequency analysis can be used as a tool to obtain a "signature" of radiated vibrations of the mitral valve which correspond with numerical models.

The wavelet transform (WT) provides much better resolution of heart sounds than the STFT. Khadra, Matalgah, El-Asir & Mawagdeh (1991) showed that the WT was useful in obtaining qualitative and quantitative measurements of the time-frequency characteristics of the PCG and demonstrated this fact with an example by calculating the time interval between A2 and P2 (components of the second heart sound) which may change if a pathological condition is present. A study comparing the STFT, Wigner Distribution (WD), and WT for time-frequency PCG representation concluded that the WT has clear advantages over the other two methods (Khadra et al. 1991). The WT also allows for clear measurement of the time gap between the A2 and P2 components of S2 (Khadra et al. 1991). Obiadat & Matalgah (1992) also concluded that the WT outperforms the STFT in time-frequency resolution and again points out the fact that the WT can detect A2 and P2 of the second heart sound. El-Asir et al. (1996) demonstrated that the WT was useful for separating normal and abnormal PCGs especially in classifying heart murmurs. Using the wavelet transform and neural networks, systolic murmurs have been classified with a high success rate (Shino et al. 1997). Matalgah et al. (1998) conducted a study where the FT, STFT, pseudo-WD, and WT were applied to normal and abnormal PCGs, and it was concluded that the WT was the most appropriate technique of the previously mentioned methods to analyse the PCG because it possesses good time resolution for high-frequency components, and allowed for the exact measurement of the time gap between the A2 and

P2 components. Kim, Lee, Yeo, Han & Hong (1999) also came to similar conclusions as Matalgah et al. (1998) in a study that compared the STFT, WD and WT. In the study done by Matalgah et al. (1998), an interesting time-frequency algorithm is presented that uses a combination of the WT and STFT.

An interesting three-part study first simulated the first heart sound, secondly compared the performance of five time-frequency representations of the simulated S1 while varying parameters including noise, and then used the best time-frequency representation to investigate the intracardiac and the thoracic first heart sounds in dogs (Chen, Durand & Lee 1997, Chen, Durand, Guo & Lee 1997, Chen, Durand, Lee & Wieting 1997). The five time-frequency representation methods were the spectrogram, time-varying autoregressive (AR) modelling, binomial reduced interference distribution, Bessel distribution and cone-kernel distribution. Of the five, the cone-kernel distribution was found to be the best technique for the time-frequency analysis of signals such as the simulated S1.

In recent years, the matching pursuit method has become an increasingly popular method to use in time-frequency representations of complex signals. Akay & Szeto (1995) performed a study which used the matching pursuit method to analyse the effects of morphine on highly non-stationary foetal breathing rates. The MP method was chosen over the WT method because it may not accurately represent signals whose FTs have narrow frequency support. The MP method was found to be superior to the STFT and WT in finding multiple periodicities in the foetal breathing signals. In a study, which applied the matching pursuit method to 11 PCGs containing heart sounds and different murmurs, it was found that the time-frequency representation of these PCGs was suitable and effective (Zhang, Durrand, Senhadji, Lee & Coatrieux 1998). Also, the effectiveness of the MP method is tied to the optimisation of the parameters used for the decomposition of the PCGs.

1.5.4 Classification and Feature Extraction

Classification and feature extraction have proven to be useful tools in extracting hidden information from the complex PCG signal that is not revealed to the naked eye. If this information is extracted and sorted based upon a classification system, various pathological conditions may be efficiently diagnosed. There are many types of classifications which

have been performed based upon temporal features, spectral content and time-frequency features. Some of these studies are reviewed below.

An early study that began in 1960 was aimed at developing a fast and reliable test for the early detection of cardiovascular disease (Gerbarg, Holcomb, Hofler, Bading, Schultz & Sears 1962). Large numbers of PCGs were recorded and were analysed to automatically classify the PCGs into different groups. Most heart sounds were generally classified correctly using spectral content.

Karpman et al. (1975) presented an interesting study which examined the use of synchronous averaging of PCG envelopes to make the system more robust to noise and then used the signal to access 80 patients having one of six common systolic murmurs. This technique resulted in a correct diagnosis rate of 89%.

Linear prediction analysis was used to classify various patterns (such as different types of murmurs) in phonocardiograms (Iwata, Suzumura & Ikegaya 1977). The system used a single beat recorded from the apex or portion 3L through an H-type filter. The characteristics for each class were shown using spectral contours and the classification rate was quite acceptable. Iwata et al. (1977) mentioned different PCG classification systems which have previously been reported that make use of PCGs recorded at several points on the chest of the subject along with the ECG and then analysed the temporal features such as zero crossings and amplitude. His system presented an improvement on such systems because only one PCG recording was needed.

Iwata, Boedeker, Dudeck, Pabst & Suzumura (1983) presented several procedures which may be used for classifying PCGs. A distance measure, which evaluates the similarity between two signals and is derived from linear predictive analysis, is used to classify PCGs of normal patients and subject with different types of murmurs. The total correct rate of classification was reported to be over 70%.

Baranek, Lee, Cloutier & Durand (1989) write of an interesting method for automatically detecting sounds and murmurs in patients who have an Ionescu-Shiley aortic bioprostheses, as automatic detection of murmurs and events is much more efficient than performing the process manually. The algorithm uses *a priori* knowledge of the spectral and temporal characteristics of the first and second heart sounds and other features. The PCG

envelope and noise levels are estimated iteratively in order to identify certain events in the PCG. The study concluded that this method is very effective in detecting the second heart sound and the aortic component of the second heart sound and perhaps may be useful in identifying murmurs.

The abilities of two different pattern recognition methods (or classifiers) to detect valvular degeneration were tested in patients with a porcine bioprosthetical heart valve (Durand, Blanchard, Cloutier, Sabbah & Stein 1990). The first method was based on the Gaussian-Bayes model and the second used the "nearest neighbour" algorithm with three distance measurements. Both methods performed very well with the best record (98%) being obtained using the Bayes classifier and two patterns of six features each.

A study was performed which employed the wavelet transform to describe and recognise certain cardiac events from the ECG (Senhadji, Carrault, Bellanger & Passariello 1995). The choice of wavelet family, analysing function and levels of decomposition were discussed. Different types of descriptions have been examined which are an energy-based representation and the extrema distribution estimated at every decomposition level. The capability of these methods has been tested using principal component analysis. Their ability to discriminate between normal and various pathological heart sounds has been analysed using linear descriminent analysis. The results need to be confirmed on a larger study population. It would be interesting to apply a similar technique to PCGs.

Using a technique which relied on the wavelet transform and the analytical signal, the distributions of the heart rate variation (HRV) in healthy subjects may be described by a single function while patients with a cardiopulmonary instability caused by sleep apnoea could not be, thus producing a method to discriminate between the two groups (Ivanov, Rosenblum, Peng, Mietus, Havlin, Stanley & Goldberger 1996). A new method was presented for exact classification of patients with different levels of coronary artery diseases and different ejection fraction (Tkacz, Kostka & Komorowski 1998). The wavelet transform and adaptive filters were applied to HRV signals resulting in a reasonably accurate and non-invasive method. HRV was studied using scale specific variance and a scaling exponent to differentiate between healthy and pathological subjects using the WT as a tool among other methods (Ashkenazy, Lewkowicz, Levitan, Havlin, Saermark, Moelgaard,

Thomsen, Moler, Hintze & Huikuri 1999, Ashkenazy, Lewkowicz, Levitan, Havlin, Saermark, Moelgaard & Thomsen 2000). The variance measure at certain scales separates healthy subjects from those with cardiac conditions. The scaling exponents perform better than the variance measure for cumulative survival probability. Significant amounts of research has been performed on the topic of HRV and the reader may refer these studies for further information (Cavalcanti 2000, Dickhaus, Maier, Khadra & Maayah 1998, Qader, Khadra & Dickhaus 1999, Raymond 1999, Sawada et al. 1997, Vinson, Khadra, Maayah & Dickhaus 1995).

Shen & Sun (1997) investigated the use of a non-Gaussian AR model and a parametric estimation in analysing normal and abnormal PCGs. The non-Gaussian AR model of PCGs is applied to uncover quadratic non-linear interactions and to classify two classes of phonocardiograms using the parametric bispectral estimate (which is shown to be highly immune to background noise). Many actual PCGs were analysed using these techniques revealing that quadratic non-linearities exist in normal and pathological PCGs. It was also shown that parametric bispectral techniques are useful for analysing the PCG.

A new method for graphically representing heart sounds and murmurs is presented in Tovalr-Corona & Torry (1997) which uses stylised diagrams. These diagrams can later be used as a standard for classifying murmurs because current methods are very subjective and often are just a written description of the murmur. Frequency domain analysis, which uses AR modelling, performs much better than time domain analysis for identification of heart sounds and murmurs. Two different diagrams give information in the time and frequency domains. Individual heartbeats are separated. Then the main heart sounds and other features such as murmurs are segmented with spectral estimation. Finally, two diagrams giving the shape in time and spectral content are drawn of the cycle with the loudest murmur.

Tovar-Corona & Torry (1998) present another study dealing with murmurs in which a time-frequency representation is found to give interesting information on the characteristics of murmurs. The Continuous Wavelet Transform is applied to the PCGs of 10 patients, who are diagnosed with Aortic Stenosis and have systolic murmurs, to obtain three dimensional time-frequency representations. The dominant frequencies are highlighted which

provides a very readable representation that reveals the variation of spectral content and intensity of the murmur. Some of the murmurs demonstrate characteristic rising and falling tones which suggests varying degrees of severity.

Multi-resolution wavelet analysis of heartbeat intervals was used to discriminate between healthy patients and those with cardiac pathologies (Thurner, Feurstein & Teich 1998, Thurner, Feurstein, Lowen & Teich 1998). Applying the wavelet transform with a scale window of between 16 to 32 beats, produced widths of the R-R wavelet coefficients that fall into separate sets for normal and heart-failure patients resulting in 100% classification.

Dickhaus, Khadra, Lipp & Schweizer (1992) and Dickhaus & Heinrich (1998) examined the use of time-frequency representations such as the STFT, WT and Wigner distribution as a clinical tool for analysing late potentials in the ECG. Overall, it was found that the WT performed the best for the task because the STFT provided poor resolution while the WD produced smearing. It was demonstrated that the signal energy of a certain time-frequency region could correctly discriminate between 90% of patients and control subjects which were surveyed.

It was demonstrated that differences between normal and abnormal heart sounds may be seen in the time-frequency distribution computed using the cochlear wavelet transform (Jandre & Souza 1997). These differences may be seen in the scalograms which have morphological differences in the duration and spectral content of heart sounds. In the future, it might be possible to design an automatic diagnostic system using the cochlear wavelet based scalogram.

In Kovacs & Torok (1998), an improved method for acoustical foetal heart rate monitoring is presented that reduces the number and length of the faulty sections of the foetal heart rate diagram when there are large interferences due to noise. The algorithm makes use of an adaptive correlation on the abstracted form of the detected signal power which is measured on two frequency channels. If this method is implemented using a portable electronic circuit, foetal heart rate may be passively monitored for long periods of time.

Bentley, Grant & McDonnell (1998) compared time-frequency and time-scale techniques for classifying native and bioprosthetic heart valves. Morphological features, which were obtained using the Choi-William distributions, achieved 96% correct classification for

native valve populations and 61% for those with the bioprosthetic heart valves. The WT feature set obtained better results with 100% correct classification for native valve populations and 87% for those with the implanted valves.

A heart sound analysis and acquisition system is described by Zehan, Shiyong, Li, Yuli & Shouzhong (1998). The system records the PCG and calculates the amplitude of various heart sound components, and the duration of each component of the heart sound. The system then analyses the heart sound based upon these factors and classifies the heart sounds as normal or pathological. Of 50 abnormal samples analysed, the coincidence rate compared to the opinion of a cardiologist was 86%.

A technique, called the phonocardiogram exercise test, for detecting cardiac reserve in healthy and diseased patients using the PCG is presented (Xiao, Xiao, Cao, Zhou & Pei 1999). Changes in the cardiac state may be seen in the PCG. Comparing the PCG before and after exercise, allows these changes to be examined because changes in the amplitude of S1 are closely linked with the maximum rate or rise of left ventricular pressure which measures cardiac contractility. A portable system on a laptop computer has been implemented which is relatively cheap, easy to use, and reliable.

1.6 Scope of Thesis and Justification of Research

Phonocardiograms are an important tool in the field of cardiology as changes in the heart may be reflected by PCGs much sooner than by other diagnostic tools such as ECGs. They are also non-invasive involving no trauma to the patient, quick, inexpensive and may be stored for future use. However, phonocardiograms are not as widely used as they could possibly be for a number of reasons including lack of standardised equipment and procedures and problems with noise.

Noise is produced from a number of sources including thoracic muscular noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), peristaltic intestine noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), respiratory noises, foetal heartbeat noise if the subject is pregnant, noise caused by contact with the instrumentation and ambient noise. Thus, the problem of noise removal in phonocardiography is very important for the development

of phonocardiography as a widely used clinical tool.

Although attempts have been made to model heart sounds (Baykal, Ider & Köymen 1991, Chen, Durand & Lee 1997, Cloutier et al. 1987, Debiais, Durand, Pibarot & Guardo 1997, Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley 1990*a*), the genesis of heart sounds is not well understood and it is not completely known *a priori* what the PCG signal content is resulting in difficulty using traditional de-noising methods which typically rely on *a priori* knowledge of the signal. Overlapping signal and noise spectra also present another problem. Several time-frequency and time-scale decomposition based algorithms for de-noising including wavelet, optimised wavelet, wavelet packets and matching pursuit de-noising algorithms, are examined to see which method performs better for de-noising PCGs. Averaging is also used as a de-noising method. This research is part of an ongoing effort to produce a cheap, reliable, and easy to use phonocardiogram system.

Once a clean PCG is obtained, it is useful to extract features and distinguishing characteristics from the signal in an attempt to differentiate healthy and pathological cases. Various methods were explored in this area including use of the phase-space diagram, Hilbert Transform, instantaneous signal parameters and the cardiosynchrogram.

Chapter 1 introduces the subject of the thesis and gives some background information. Chapter 2 gives a background and history of phonocardiography and provides information about the instrumentation and procedures used to obtain the data. Chapter 3 provides information on the theory of the de-noising methods used. Once the reader has the necessary background to understand the de-noising study, Chapter 4 gives the results and recommendations of the de-noising study. After the PCG has been de-noised, Chapter 5 presents the different methods that were explored in relation to feature extraction and classification of the PCG. Chapter 6 summarises the results and conclusions reached and provides recommendations for future research.

Chapter 2

Equipment and Data Acquisition

"Science lies in the intellect, not in the instruments"

Abraham Flexner (1866-1959)

AMERICAN EDUCATOR

2.1 Introduction

Auscultation of the heart may be defined as listening to the sounds and vibrations produced by heart muscles contracting, blood moving through veins, arteries, and the heart, and the motion of the heart valves (Tinati 1998). Auscultation is the standard test performed when evaluating the condition of the cardiac system because it is easy to perform, requires little time, is inexpensive and non-invasive. Auscultation is normally accomplished using a stethoscope to aid the imperfect human ear in detection of the heart sounds. Graphical recordings of heart sounds are called phonocardiograms, and a microphone picks up heart sounds which are then recorded by some device.

The purpose of this chapter is to demonstrate how the PCG and ECG recordings were obtained and to describe the equipment which was used in this process. In Section 2.2, a history of phonocardiography and auscultation is given to show how the technique was developed as a clinical tool (an extensive history and development of phonocardiography up to the 1950's is given in McKusick (1958)). Then, in Section 2.3, the signal acquisition process is illustrated and the equipment used is described. In the final section of this chapter, information is presented regarding the patients and the ECG and PCG recordings which were obtained.

2.2 History of Phonocardiography and Auscultation

2.2.1 Limitations of the Human Ear

Listening to the human heart (known as auscultation) is a very old and useful technique, which dates back to the days of Hippocrates around 400 BC, for the diagnosis of heart diseases and conditions (Selig 1993, Luisada 1965). The goal of auscultation is to predict as far in advance as possible the development of a cardiac pathological condition (Selig 1993).

Important factors to obtain the maximum benefit of cardiac auscultation as a clinical diagnostic tool include the ability to hear a wide range of frequencies with adequate sensitivity,

to selectively listen to different sounds and experience in recognising and classifying heart sounds and murmurs (Selig 1993). PCGs and electronic stethoscopes attempted to aid the human ear in overcoming certain difficulties in hearing, but they were introduced too early and lacked standardisation making them unreliable and difficult to use (Selig 1993).

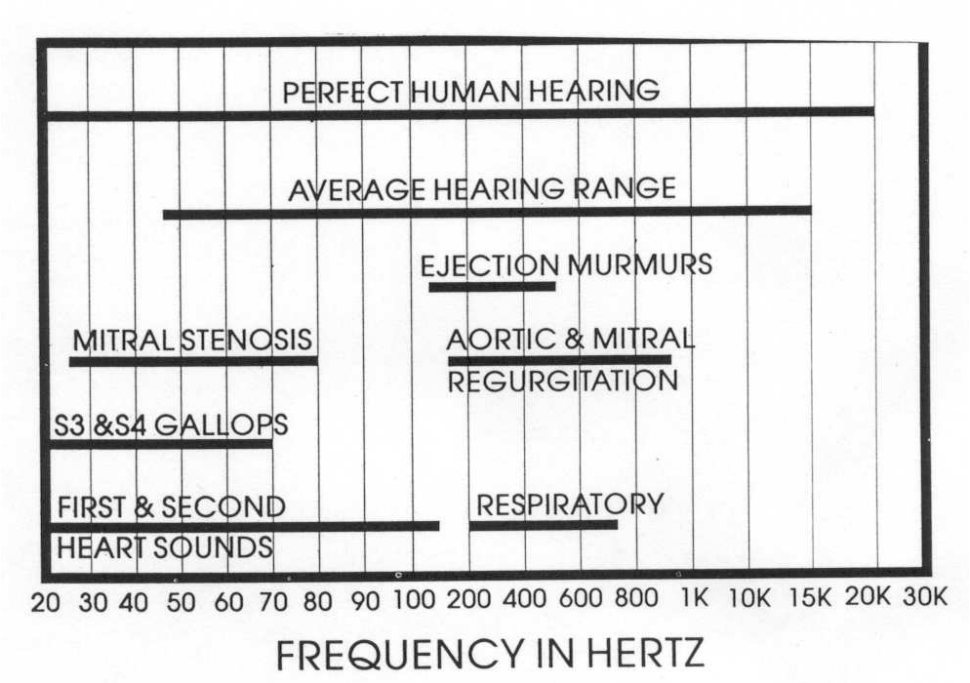

Figure 2.1 Relative frequency ranges from Selig (1993)

Due to the design and limitations of the human ear, it has difficulties discerning certain cardiac events which is demonstrated in Figures 2.1. The average adult hearing range is about 50-12,000 Hz with the ear performing best and possessing the best frequency discrimination between 1-2 kHz (Selig 1993). The human ear has difficulty discriminating between different frequencies of sound equally as the volume is varied often resulting in difficulties diagnosing heart murmurs and sounds; a remedy to this problem is to amplify the low-frequency sounds because high volume enables low-frequency sounds to be heard (Selig 1993). Ambient noise may also mask cardiac sounds resulting in difficulties discerning the actual heart sounds especially in hospital rooms or examining wards which may have up to 70 dBs of ambient noise while the first heart sound heard at the apex is only 20-60 dBs and the volume of the second heart sound at the aortic area is 30-70 dBs (Selig 1993, Groom et al. 1956). The third and fourth heart sounds are often missed because they have very low frequency content and are quite faint. There is also a time

delay of up to 1 second, known as temporal summation, associated with hearing the full intensity of a sound and the human ear cannot differentiate between time intervals smaller than 20 milliseconds (Selig 1993).

2.2.2 Development of the Art of Auscultation and the Stethoscope

Since the time of Hippocrates, doctors have applied their ears to the chest of the patient listening for sounds emanating from the heart (Selig 1993). However, it was not until 1628 that Harvey became the first person to record his belief that the heart caused an audible pulsation (Luisada 1965). In 1715, James Douglas heard and described aortic regurgitation sounds (Selig 1993). In 1764, William Hunter described arteriovenous fistula (Selig 1993). The idea of using sounds coming from internal organs to detect disease originated with Jean Nicholas Corvisart in 1806 (Selig 1993). Gaspard Laurent Bayle taught René Théophile Hyacinthe Laennec direct auscultation (ear applied to the chest of the patient), but Laennec thought this was quite inconvenient and often embarrassing for the patient (Selig 1993). So, Laennec rolled a sheet of paper into a cylinder and put one end on the patient's chest and the other to his ear which worked better than direct auscultation, thus inventing the stethoscope (from Greek "to view the chest") in 1818 (Selig 1993, Luisada 1965). Laennec described the first and second heart sounds and also described murmurs but they were misinterpreted due to incomplete knowledge of physiology (Luisada 1965).

In the second half of the twentieth century Laennec's stethoscope underwent modifications and improvements (See Figure 2.2) described by (Selig 1993). Monoaural flexible stethoscopes were designed by Nicholas Comins in 1829. In 1855, George Cammann invented a binaural small, rigid chest piece which was connected to two flexible tubes that went to the ears of the auscultator. The diaphragm and membrane chest piece were contrived by R. L. M. Bowls in 1894. In 1926, Howard B. Sprague introduced the combination chest piece.

Before the invention of the stethoscope, cardiac diagnosis was based on the patient's

account of their illness, the patient's skin colour, tongue condition and pulse (Selig 1993). In the 1820s, with the invention of the stethoscope, members of the medical fraternity challenged auscultation but autopsies proved it to be correct. However, the role of the stethoscope declined with the introduction of phonocardiography in 1908.

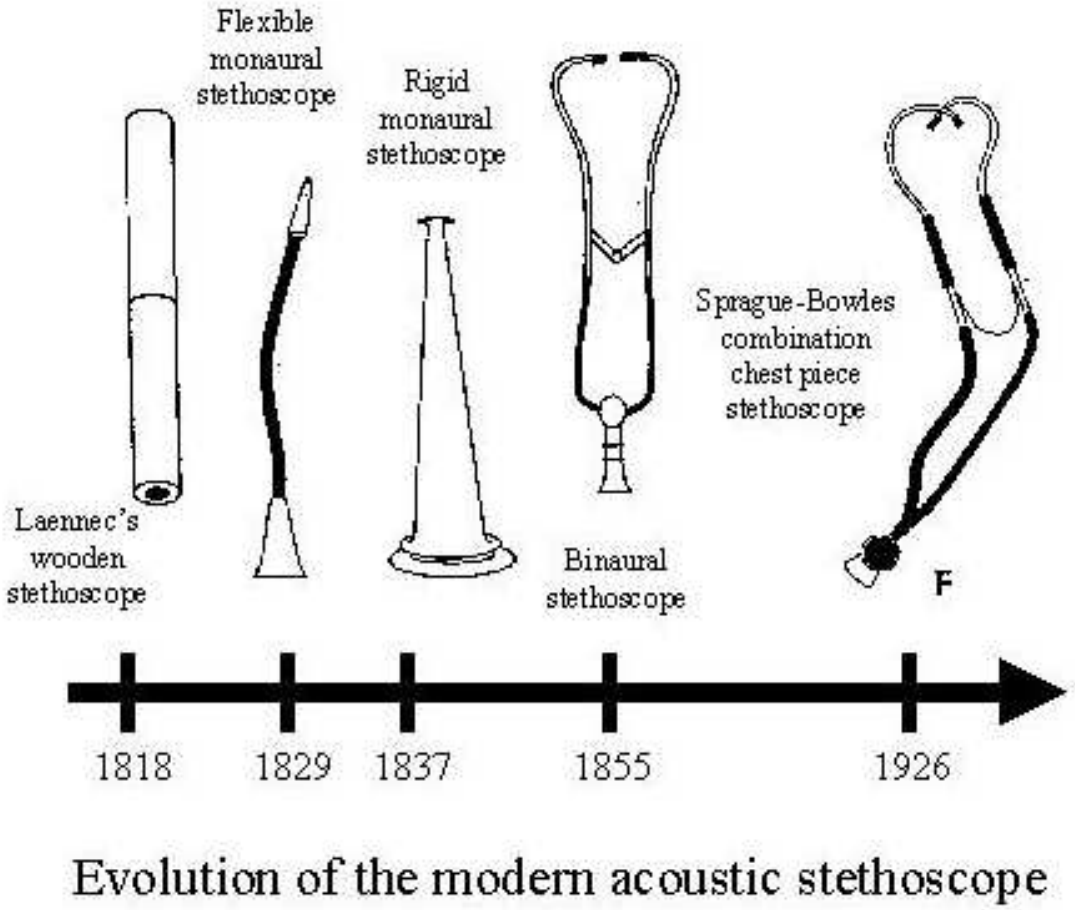

Figure 2.2 Timeline of the evolution of the acoustic stethoscope. Modified from Selig (1993).

2.2.2.1 From the Acoustic Stethoscope to the Electronic Stethoscope

Previous discussion has pointed out that the human ear, while having a wide hearing range, is not perfectly suited for auscultation. Electronic stethoscopes have evolved over the past 80 years but they were too early to the marketplace, were unreliable and difficult to use (Selig 1993). Even today, electronic stethoscopes are not preferred over acoustic

stethoscopes, as was proved in a study performed on many doctors and nurses where three acoustic stethoscopes and electronic stethoscopes were examined, acoustic stethoscopes were favoured 71% of the time and 29% of the time electronic stethoscopes were preferred (Grenier et al. 1998). The first electronic stethoscopes for teaching with multiple listener capability were introduced in 1923 (Selig 1993). In 1952, Maico Company developed a device that could provide amplification of +15 to -100 dB where the peak sensitivity was at 500 Hz (Selig 1993).

Limitations of the acoustic stethoscope include: the lack of amplification of heart sounds, attenuation of high frequency sounds, and the high pressure which is applied to the ears by some stethoscope models in order to isolate the heart sounds from the ambient noise (Grenier et al. 1998). Electronic stethoscopes overcome one of the main limitations of acoustic stethoscopes because they amplify heart sounds but they also introduce limitations including the generation of electronic noise, their sensitivity to impact and environmental noise, poor ergonomic design, and their use of electronic filters, which has no practical meaning to medical practitioners, instead of the standard bell and diaphragm (Grenier et al. 1998). An ideal electronic stethoscope would have the benefits of both the acoustic and electronic stethoscopes without having their limitations such as electronic and ambient noise, and they should be ergonomic and easy to use (Grenier et al. 1998). Thus, techniques for noise removal, are of crucial importance for electronic stethoscopes.

2.2.3 The Introduction of Phonocardiography

Various developments led to the PCG being favoured over the stethoscope by 1908 (Selig 1993). In 1893, K. Huerthle of Breslau described a method of inscribing heart sounds (Luisada 1965, Selig 1993). He connected the output of a microphone through an inductor whose secondary coil was connected to the nerve of a frog muscle preparation which marked on a smoked drum using an attached lever. The following year, the first PCGs similar to those in use today were recorded by William Einthoven with Geluc in Leyden using Lippman's capillary electrometer which was used to record the first ECGs (Selig 1993). In 1904, precordial vibrations with optical amplifications (called the Frank segment capsule) were recorded by Otto Frank in Munich. In 1909, Carl J. Wiggers modified

Frank's PCG and graphically displayed recordings; however, the method suffered from limitations in the frequency response of the membrane which was used in the frequency capsule. In 1907, Einthoven used a string galvanometer to record heart sounds, but the problem with this method is that there is an enormous range of intensities of heart sounds often resulting in damage to the string.

For the methods of indirect phonocardiography or phonocardiography using a string galvanometer, there was no way of applying electrical filtration so acoustic filtration was used (Selig 1993). This was accomplished by using an air leak in the tube connecting the chest piece with the Frank capsule or microphone.

In 1942, Maurice Rappaport and Howard Sprague outlined methods for use of electronic amplification and galvanometers (of a different type) in phonocardiography which became the predominant technique used in phonocardiography (Selig 1993). They reported that if low frequencies were filtered considerably then the PCG resembled auscultation.

Until the time that Bell Telephone Laboratories method of spectroscopy was used to study heart sounds and murmurs by Geckeler who named it cardiospectrography and by McKusick who called it spectral phonocardiography, all phonocardiography was oscillographic phonocardiography (Selig 1993). Developed in the 1950s, cardiospectrography or spectral phonocardiography possessed the capacity to spread out the frequency spectrum by using the Heterodyne electronic filter which could focus on specified frequency ranges. Huggins developed phase filtering which allowed heart sounds to be represented more accurately. With the arrival of the wide bandpass filter, transients were now able to be displayed. A problem existed with establishing normal values for PCGs because the range of sounds is enormous ranging from a faint diastolic murmur which would be represented by a mark less than 2 millimetres long to a large marking representing a systolic murmur. To correct this problem, voltage limitation was introduced; so that if one was searching for a faint diastolic murmur, the voltage could be limited to a small range, thus making the diastolic murmur more evident. To make the identification of heart sounds easier, reference signals such as the ECG were introduced. Frequency shifting is also now possible for PCG recordings making sub-audible frequencies audible.

Phonocardiography has not progressed as quickly as other diagnostic methods because of

lack of standards in equipment and recording locations, poor understanding of the heart sound mechanisms and the complexity of PCGs, but signal processing has aided in the development of phonocardiography (Wood & Barry 1995).

An important use of signal processing in the area of phonocardiography is noise removal. There are many sources of noise which may pollute the PCG including thoracic muscular noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), peristaltic intestine noise (Zhang, Durand, Senhadji, Lee & Coatrieux 1998), respiratory noises, foetal heart sound noise if the subject is pregnant, noise caused by contact with the instrumentation and ambient noise. The noise present in each case depends on the state of the subject, the instrumentation used, and the environment which the PCG was recorded in. Groom et al. (1956) found that the background sound level in hospitals and clinics was quite high being on the order of 60 to 70 dBs. A sound-proof room was constructed which reduced noise levels on the order of 35 dBs. The effect of noise on the auscultatory performance of forty doctors was measured under simulated stethoscope examinations. The average results of the group demonstrated that the same murmur heard in the sound proof room had to be increased 12 times in intensity to be heard under normal noise conditions in clinics and hospitals. Reducing hospital noise levels to the order of 35 dBs would not be economical except in the case of a sound-proof examining room which is not practical. Faint heart murmurs, which are very often of the most importance in diagnosing early heart disease, may be easily masked by ambient noise. Thus, commonly found levels of background noise seriously impair the ability of the medical practitioner to discern heart murmurs through a conventional stethoscope. Recording the heart sounds through phonocardiograms and applying noise removal techniques may reveal faint heart murmurs which previously would have gone undetected. Modern signal processing techniques may also have applications in the area of complex signal analysis of the PCG to reveal hidden information contained in the PCG.

2.2.4 Some Modern Phonocardiography Systems

The aim of this section is to provide a brief review of some of the more modern phonocardiography systems which have been developed that make use of the personal computer.

A simple system for phonocardiographic analysis is described in Nandagopal et al. (1981). The ECG and PCG are recorded via standard limb electrodes and appropriate microphone respectively and after some processing are recorded on a cassette tape. Frequency analysis is then applied to the PCG which could be used to determine the peak frequencies of the first and second heart sounds.

Lehner & Rangayyan (1987) presented a three-channel microcomputer system for phonocardiography. It records the PCG, ECG and carotid pulses. Using the ECG and carotid pulse as references, the PCG is broken down into systolic and diastolic parts. Parameters which represent the time and frequency domains of the signal are computed. Using these parameters, murmurs may be detected and classified.

Tavel, Brown & Shander (1994) reported that a system has been developed which has the capabilities to graphically display the PCG on a hand-held unit, but lacks the capabilities to re-play the recorded sound.

An electronic stethoscope has been developed which may be connected to a computer using a special module (Durand & Pibarot 1995). The PCG and ECG may be simultaneously recorded and re-played.

A more advanced phonocardiographic recording and processing system is described by Lukkarinen, Korhonen, Angerla, Nopanen, Sikio & Sepponen (1997), Lukkarinen, Nopanen, Sikio & Angerla (1997), and Lukkarinen, Sikio, Nopanen, Angerla & Sepponen (1997). A hand-held electronic stethoscope which contains an electret microphone and a bell-type chest piece is connected to a multimedia personal computer. The system has the capacity to record heart sounds using all the standard multimedia sound formats and then replay the PCG. The PCG and the STFT may be simultaneously displayed on the screen. User-definable digital filters may also be used to obtain better results. They report that the system is useful in diagnosing murmurs in children.

A non-invasive telemetric heart rate monitoring system is described by Torres-Pereira, Ruivo, Torres-Periera & Couto (1997). The system uses a piezoelectric sensor to detect the heart sounds. It has an advantage over electrocardiography in heart rate monitoring because it only uses a single probe instead of multiple wire connections. The most prominent heart sound, S1, is detected and used to determine the heart rate.

Another example of remote patient monitoring is given in Bell et al. (1998). The system acquires the ECG, PCG and patient temperature, and these are transmitted to a personal computer. Commercially available software may be used for displaying and analysing the data.

Yet another system has been developed which will simultaneously record the ECG and PCG (Guo et al. 1998). The system has the capacity to perform spectral and time-frequency analysis of the stored signals which may be useful for detecting murmurs and other conditions. This system was implemented using LabVIEW software and is relatively inexpensive and reliable.

2.3 Signal (ECG/PCG) Acquisition Process

2.3.1 Overview of the PCG-ECG System

The system for obtaining the PCG and ECG from a patient is represented by the block diagram in Figure 2.3. The ECG signal, a summation of voltages from different areas of the body, is amplified by a system described later and the PCG signal, a recording of heart sounds, is connected to the A/D converter. The A/D converter outputs a digital signal which may be stored on the computer for further manipulation. The signals are sampled simultaneously, so that both may be used at once. Further information on the PCG-ECG system including a circuit diagram and PCB layout may be found in Appendix C and the following sections.

2.3.2 Recording the PCG

2.3.2.1 Pick-up devices

Heart sounds are acoustic waves. Thus, a microphone is needed to pick up the sound waves. There are many types of microphones which may be used including magnetic microphones, condenser microphones, piezoelectric sensors, and electret microphones. There still exists a lack of standardisation for PCG pick-up devices although attempts have been

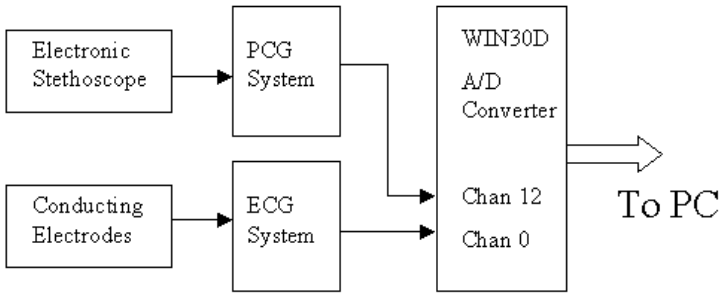

Figure 2.3 ECG-PCG system block diagram, modified from Maple (1999).

made to standardise microphones (Groom 1970, Mannheimer 1957, Takagi & Yoshimura 1964). The advantages and disadvantages of different microphone types are discussed below.

Magnetic microphones are not well suited for picking up heart sounds because they have a low frequency cutoff of about 50 Hz, while heart sounds have several frequency components less than 50 Hz (Tinati 1998), and they possess low impedance and sensitivity.

Condenser microphones are used for acoustical recordings and make use of the capacitance change caused by the movement of the diaphragm (Tinati 1998). They possess a cutoff frequency of around 20 Hz which makes them more suited for picking up low frequency components of heart sounds than magnetic microphones whose cutoff frequency is much higher (Tinati 1998). However, due to possessing a higher noise figure resulting from their higher impedance and sensitivity, they are unsuitable for phonocardiography (Tinati 1998). Another drawback is that they require a relatively high (60-120 V) DC voltage to polarise their plates (Borwick 1990).

The piezoelectric sensor is a device that is used in many heart sound studies (Beyar, Levkovitz, Braun & Palti 1984, Baranek et al. 1989, Cloutier et al. 1987, Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley 1990*a*, Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley 1990*b*, Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley 1990*c*, Ewing et al. 1986, Karpman et al. 1975, Jandre & Souza 1997, Lehner & Rangayyan 1987, Nandagopal et al. 1981, Tinati 1998, Tovalr-Corona & Torry 1997). Tinati (1998) describes how they function. Wafers

Figure 2.4 The Escope from Cardionics is an electronic stethoscope that is used to record heart sounds.

or slabs of crystalline or ceramic materials demonstrate electrical polarisation that varies with mechanical deformations. When the wafers are exposed to torsional or bending stresses, a potential difference is created between opposite faces of the slabs which is coined the piezoelectric effect. A normal piezoelectric microphone is composed of layers of oppositely polarised wafers all connected with metal electrodes attached to them. The piezo sensor is not being used in this study because the piezoelectric microphone is never used or readily available in a doctors' surgery and the emphasis in this project has been to obtain a practical system that would be used in a medical office. The piezo sensor may be losing information in the PCG because it only picks up the movement of the heart as it contracts (Maple 1999).

An electronic stethoscope, the Escope from Cardionics, is used to obtain an analogue signal which is a recording of the heart sounds. It uses an electret microphone which requires very little current and voltage as opposed to the condenser microphone due to a simplification in construction (Borwick 1990). Physicians possess extensive training in the use of the conventional stethoscope, and because the Escope closely resembles the traditional stethoscope just with additional functionality, it may easily be used by the doctor. A photo of the Escope may be seen in Figure 2.4 and the Escope specifications are given in Appendix A.

All of the types of microphones discussed pick up noise which can be a problem when

recording heart sounds. Microphones may be designed so that they aid in noise suppression . Hök (1991) presents a microphone design for suppressing noise and artifacts. Briefly, a microphone arrangement made up of a half-bridge of electret elements connected using polymer tubing and dedicated sensing heads is found to suppress noise and artifacts. The microphone system is made of widely available and relatively inexpensive components. The area of noise suppression microphones represents an area open for further research.

More information on types of microphones discussed and the development and study of microphones in the area of phonocardiography may be found in Blashkin & Yakovlev (1975), Borwick (1990), Dranetz & Orlacchio (1976), Sukimura & Funada (1971), Luisada & Zalter (1960), McKusick (1958), Obata, Yoshimura, Ide & Mike (1971), Padmanabhan, Fischer, Semmlow & Welkowitz (1989), Suzumura & Ikegaya (1977), Van Vollenhoven, Wallenburg, Van Rotterdam & Van Straaten (1968), Van Vollenhoven & Wallenburg (1970), Van Vollenhoven (1975), and Vermarient & Van Vollenhoven (1983).

2.3.2.2 Areas of the Chest for PCG Recordings

Typically, there are six areas of the human chest where phonocardiographic recordings and auscultation are performed. These areas have been derived from studies examining PCGs and comparing auscultation to the results of catheterisation and autopsy and surgical results. These areas are summarised in this section and are shown in Figure 2.5 (Luisada 1965, Luisada 1980). A study comparing phonocardiographic monitoring locations concluded that from among several recording sites there was no difference in S1 from a signal analysis perspective (Rice & Doyle 1995). Thus, the choice of the site should be based on practical reasons.

2.3.2.2.1 Left Ventricle Area (LVA) The LVA is the best place to hear any abnormalities in heart sounds which are caused by stenoses in the mitral valve and abnormalities in the left ventricle and the left atrium. Murmurs of subaortic stenosis and aortic insufficiency may be heard at this location which is located around the apex.

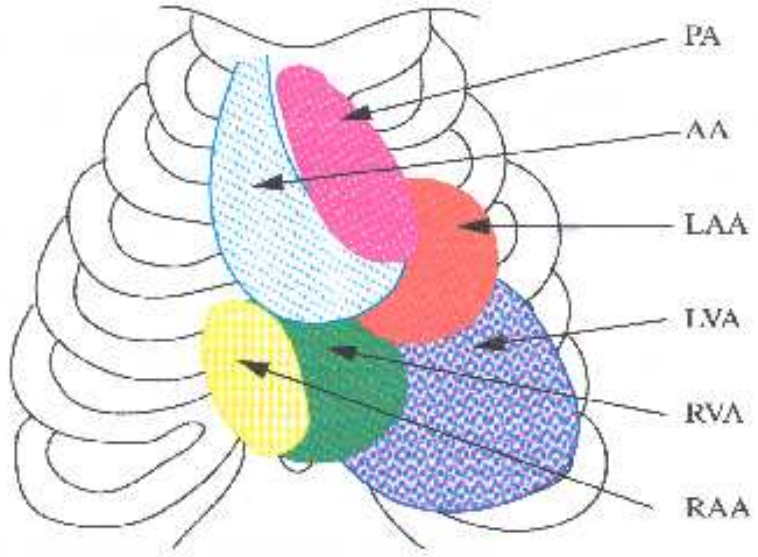

Figure 2.5 Areas of the chest for PCG recordings (Luisada 1980, Tinati 1998).

2.3.2.2.2 Right Ventricular Area (RVA) This is the best location to listen to murmurs of tricuspid stenosis or insufficiency, right ventricular or atrial "gallop", murmurs of pulmonary insufficiency and of ventricular septal defect. This area encompasses the lower part of the sternum and the third and fourth intercostal spaces on both sides of the sternum.

2.3.2.2.3 Left Atrial Area (LAA) When a mitral insufficiency is present, this area is the best location for recording the characteristic murmur that is produced. The area above and to the left of the apex is considered the LVA.

2.3.2.2.4 Right Atrial Area (RAA) This area is the optimal location for listening to and recording the murmur of tricuspid insufficiency. It is typically located 1 to 2 cm to the right of the sternum in the fourth and fifth intercostal spaces.

2.3.2.2.5 Aortic Area (AA) The sounds which are best heard at this location are those caused by aortic abnormalities and irregularities of a carotid or subclavian arteries. This area begins at the first intercostal space and covers up to the third intercostal space.

2.3.2.2.6 Pulmonary Area (PA) The PA is the best area to hear and record any abnormalities present in the pulmonary arteries. This area is located between the second and third left intercostal spaces close to the sternum.

2.3.2.3 The Recording Process

The PCG recording is made by placing the stethoscope normally on the aortic area. The patient is told to be still and quiet, lean forward a bit and hold their breath if possible. The PCG is recorded for about 30-60 seconds using the recording system described. A jack is plugged into the Escope. The output of the jack is connected to channel 12 of the A/D converter and sampled at 2500 Hz well above the Nyquist frequency for heart sounds (about 1000 Hz is about the highest frequency we see in heart sounds from Figure 2.1 (Selig 1993)). The PCG signal is passed through an anti-aliasing filter whose frequency response may be seen in Appendix C.

2.3.3 Recording the ECG

The purpose of the ECG system is to detect the electrical pulse generated which causes the heart to contract. The system for obtaining the electrocardiogram recording was inherited from work which Leonard Hall and Jarrad Maple completed. The system is described in Maple (1999) and Hall (1999), but because these reports are unpublished, the system will be described for the sake of completeness here.

In the work of Hall and Maple, a different stethoscope, the Analyst, was used, and references to it may be found. However, this particular stethoscope was found to have problems with the software and power supply, and the use of it was discontinued. For our study the ECG system is identical, but the Analyst is replaced by the Escope.

The specifications for the ECG system are as follows.

- The output is sampled by the A/D converter.
- The system voltage is 12 V.
- The system must have low output impedance.

- The maximum output voltage of the system should be 5 V.
- The system is designed for a sampling frequency of 5kHz and uses third order butterworth filter at 2 kHz.
- The system is mounted on the same PCB as the PCG.
- All cables in the system must be robust and shielded as protection from noise.
- The system must have high input impedance because the skin has high output impedance.

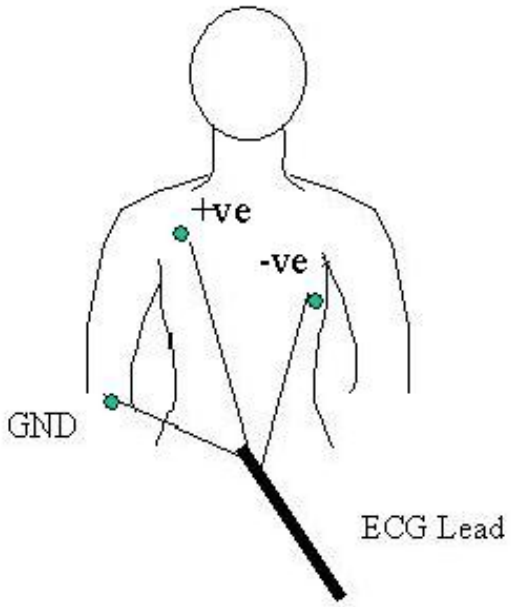

Figure 2.6 Positioning of the ECG electrodes. The +ve input is located just below the right collarbone of the patient, the -ve input is on the sixth rib under the left arm, and the Ground electrode is located on the right wrist. Modified from (Maple 1999)

To obtain the ECG, the voltage difference between two electrodes, in contact with the skin, is measured. The electrodes use a conductive gel and a pad with a fairly large surface area to give a highly conductive contact between the skin and the ECG leads. The electrodes may be positioned any number of ways. A simple configuration is used here as may be seen in Figure 2.6:

- +ve slightly below the right collarbone.
- -ve on the sixth rib under the left arm.
- Ground lead on the right wrist.

The body has a high impedance which could cause problems. The problem of high impedance of the body was resolved by connecting the electrodes directly to the +ve inputs of the operational amplifier. Next, the two signals are passed through a difference amplifier. A high pass filter is used to remove any DC component present. At the last stage of the signal processing, there is an anti-aliasing filter present at 2 kHz. The output impedance is 10 kΩ. The circuit diagram of the system and the frequency response of the filters may be seen in Appendix C. An oversight was made when setting the sampling frequency of the ECG. Although the ECG system was designed to sample at 5 kHz, we only sampled the ECG at 2.5 kHz. However, this is not important in this study because the only function of the ECG was to serve as a gating reference for the PCG. The sampling rate was correctly set for the PCG.

The original ECG-PCG system which used the Analyst was implemented on a printed circuit board (PCB), which is shown in Appendix C, because it was more reliable and noise free than a circuit built by hand on a breadboard. The PCB has inputs for a piezo sensor which is no longer being used and ECG electrodes. The output is passed to the WIN30D A/D card in the PC. The tracks on the PCB were designed to have a width of 30 mm for the ease of soldering and to avoid current overload. The PCB has dimensions of 10.16 cm by 7.62 cm so it may easily fit into a metal case which provides protection against interference.

2.3.4 The WIN-30D Analog to Digital Converter

The A/D converter is used to convert the analog signals from the ECG system and the Escope to digital signals which can be stored on the computer. See Table 2.1 for more information about the WIN-30D A/D converter.

A/D resolution	12 Bits
Nonlinearity	Less than +/- 1 LSB
A/D full scale input ranges	Unipolar range: 0 to +5VBipolar range: -5 to +5V
Number of A/D Inputs	16 single ended
A/D throughput rate	1 MHz

Table 2.1 WIN-30D Characteristics. Information from Hall (1999)

2.4 Data Records

Table 2.2 shows information about the subject and any cardiac conditions that might be present. The PCG and ECG were recorded simultaneously at the Hampstead Medical Clinic by Dr. John Agzarian for about 30 seconds as it is difficult to get a longer recording.

In this study, the PCG is analysed qualitatively rather than quantitatively because more work is needed in the area of quantitative phonocardiography before it can be completely standardised and widely used. Quantitative phonocardiography is limited by problems of standardisation including nonstandard nomenclature, equipment and recording techniques (Wood & Barry 1995). The quantitative aspect of phonocardiography has been largely ignored except for some studies which were done (Lehner & Rangayyan 1985, Luisada & Gamma 1954, Luisada & Zalter 1960, Luisada & Bernstein 1976, Mannheimer 1957, McKusick, Talbot, Webb & Battersby 1962, Obiadat & Matalgah 1992, Vermarien & van Vollenhoven 1984, Wood & Barry 1995). Vermarien & van Vollenhoven (1984) attempted to standardise microphones to aid in the development of quantitative phonocardiography analysis. The WT enables physicians to obtain qualitative and quantitative measurements of T-F characteristics of PCG signals (Obiadat & Matalgah 1992). Lehner & Rangayyan (1985) used the energy curve and power spectrum of the systole and diastole PCG segments which were computed and quantified using the concept of an energy distribution coefficient. By making use of this coefficient, various types of heart murmurs may be classified.

Patient Number	**Age**	**Sex**	**Pathological Condition**	**Comments**
1	87	F	Hypertension and heart murmur	Unusable
2	75	M	Mitral valve prosthesis	ECG leads reversed
3	79	M	Hypertension and heart murmur	ECG leads reversed
4	67	F	Hypertension and aortic stenosis	ECG leads reversed
5	70	M	Pharmacologically treated hypertension	
6	40	F	None	Patient talking during recording, rendered recording unusable
7	60	M	Aortic stenosis	
8	57	M	Angioplasty and systolic murmur	
9	84	F	Atrial fibrillation	
10	23	M	None	ECG leads reversed
11	43	F	Hypertension	ECG leads reversed
12	45	M	None	
13	42	M	None	Quality of recording very poor–unusable; strong family history of heart disease
14	8	M	Benign Murmur	ECG leads reversed
15	24	F	None	

Table 2.2 Patient Information

2.5 Chapter Summary

In this chapter, the human ear has been shown to have limitations in listening to heart sounds. Thus, additional tools are needed to listen to heart sounds. The development of the modern stethoscope and PCG were followed. Next, the specific equipment which was used in this study was presented. The final section discussed the recordings collected. In the next chapter, the theory and background needed to understand the results of the de-noising study will be introduced.

Chapter 3

Theory of De-Noising Methods

"I'd be glad to settle without the theory if I could even understand what this thing is–or what it's supposed to do."

Arthur C. Clarke (1917-)

ENGLISH SCIENCE FICTION AUTHOR

3.1 Introduction

The objective of this chapter is to provide the background theory needed to comprehend the de-noising methods used in the remainder of this thesis. Section 3.2 presents a basic overview of wavelet theory starting from the framework of the Fourier Transform and its application to signal de-noising. Section 3.3 gives a summary of wavelet packet theory and wavelet packet de-noising. Section 3.4 explains how the matching pursuit method may be applied to de-noising. In Section 3.5, some information regarding averaging in the context of de-noising is given.

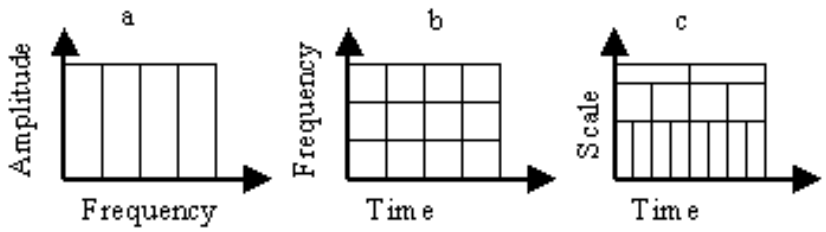

Figure 3.1 Comparison of a signal represented in different domains with (a) corresponding to the Fourier transform representation, (b) representing the short time Fourier transform, and (c) the wavelet transform

3.2 The Wavelet Transform and De-noising

Wavelet theory dates back to the work of Joseph Fourier, but most of the advances in the field have been made since the 1980s. This section gives a review of basic theory needed to understand wavelet de-noising.

3.2.1 Fourier Analysis

In 1822, Joseph Fourier discovered that any periodic function could be represented as an infinite sum of periodic complex exponential functions (Polikar 2000). The inclusive property of only periodic functions was later extended to any discrete time function. The Fourier Transform (FT) converts a signal expressed in the time domain to a signal

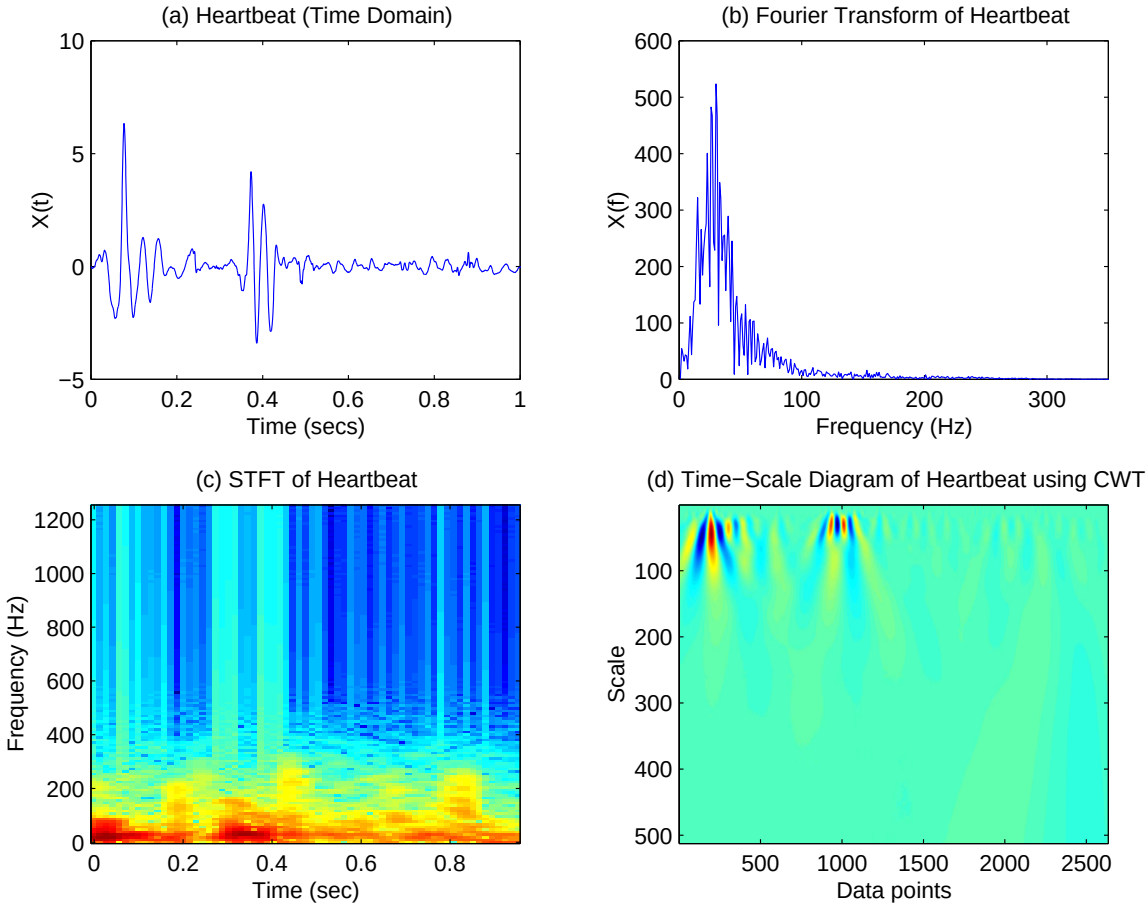

Figure 3.2 This figure shows a heartbeat in different representations with (a) giving the time domain representation of a heartbeat (b) showing the Fourier Transform representation (frequency magnitude) of the heartbeat (c) showing the Short Time Fourier Transform or a time-frequency domain representation of the heartbeat and (d) showing a time-scale diagram of the heartbeat which was obtained using the Continuous Wavelet Transform.

expressed in the frequency domain. The FT representation of a signal may be seen in Figure 3.1(a) and the FT representation of a heartbeat may be seen in Figure 3.2(b). The FT is widely used and usually implemented in the form of the Fast FT algorithm. The mathematical definition of the FT is given below

$$X(f) = \int_{\infty}^{-\infty} x(t)e^{-j2\pi ft}dt \tag{3.1}$$

The time domain signal $x(t)$ is multiplied by a complex exponential at a frequency f and integrated over all time. In other words, any discrete time signal may be represented by a sum of sines and cosines which are shifted and are multiplied by a coefficient that changes their amplitude. $X(f)$ are the Fourier coefficients which are large when a signal contains a frequency component around the frequency f. The peaks in a plot of the FT of a signal correspond to dominant frequency components of the signal.

The Fast Fourier Transform (FFT) is widely used, perhaps even too widely used. Yves Meyer states (Hubbard 1996),"Because the FFT is very effective, people have used it in problems where it is not useful–the way Americans use cars to go half a block..."

Fourier analysis is simply not effective when used on non-stationary signals because it does not provide frequency content information localised in time. Most real world signals exhibit non-stationary characteristics (such as heart sound signals), thus Fourier analysis is not appropriate.

3.2.2 Short Time Fourier Transform (STFT)

The problem with Fourier analysis is the fact that it does not matter *when* frequency components appear in a signal because the signal is integrated over all time in Equation 3.1. Thus, the frequency content of the signal is known, but its location in time is not known.

In an effort to address this problem, the STFT was developed in 1946 by Denis Gabor (Misiti et al. 1996). The STFT analyses a small section of the signal at a time which is known as windowing. The STFT is a compromise between the time and frequency representation of a signal providing information about the spectral content and when it occurs. The tradeoff is between rather imprecise time and frequency resolution, which

is determined by the window size. The STFT representation of a signal may be seen in 3.1(b) and the STFT representation of a heartbeat may be seen in Figure 3.2(c). The mathematical representation of the STFT is

$$\mathrm{STFT}_{\mathrm{x}}^{(\mathrm{w})}(\mathrm{t}', \mathrm{x}) = \int_{\mathrm{t}} [\mathrm{x}(\mathrm{t})\mathrm{w}^*(\mathrm{t} - \mathrm{t}')]\mathrm{e}^{-\mathrm{j}2\pi\mathrm{ft}}\mathrm{dt} \tag{3.2}$$

where $x(t)$ is the signal and $w(t)$ is the windowing function which is translated by a certain amount denoted as t'. The windowing process translates the complex conjugate of the window function along the length of the signal while multiplying the signal and windowing function at different points in time. The function of the exponential component in Equation 3.2 is to convert the product of multiplication of the signal and windowing function from the time domain to the frequency domain.

The problem with the STFT is a compromise in resolution. The smaller the window used, the better quickly changing components are picked up, but slowly changing details are not detected very well. If a larger window is used, lower frequencies may be detected, but localisation in time becomes worse.

3.2.3 The Wavelet Transform (WT)

The Wavelet Transform was developed as a method to obtain simultaneous, high resolution time and frequency information about a signal. The term "wavelet" was first mentioned in 1909 in a thesis by Alfred Haar (Misiti et al. 1996), although the progress in the field of wavelets has been relatively slow until the 1980s when scientists and engineers from different fields realized they were working on the same concept and began collaborating (Hubbard 1996). In the past decade, much has been written about wavelet theory and their diverse range of applications in the processing of biomedical signals, speech signal processing, physics, image processing and statistics (Akay 1997, Antoine 1999, Buckheit & Donoho 1995, Burrus, Gopinath & Guo 1998, Coifman & Wickerhauser 1998, Daubechies 1988, Daubechies 1992, Daubechies 1996, Hubbard 1996, Ivanov, Goldberger, Havlin, Peng, Rosenblum & Stanley 1999, Lankhorst & van der Laan 1995, Mallat 1989, Matalgah & Knopp 1994, Meyer & Ryan 1993, Misiti et al. 1996, Pan, Zhang, Dai & Zhang 1999, Polikar 2000, Strang & Nguyen 1996, Sun & Sclabassi 1998, Unser &

Aldroubi 1996, Wickerhauser 1994, Yang, Qiu & Koh 1994). For an interesting personal history about the origin of wavelets, Daubechies (1996) may be consulted.

The WT presents an improvement over the STFT because it obtains good time and frequency resolution simultaneously by using a variable sized window region (the wavelet) instead of a constant window size. Because the wavelet may be dilated or compressed as is seen in Figure 3.1(c), different features of the signal are extracted. While a narrow wavelet extracts high frequency components, a stretched wavelet picks up on the lower frequency components of the signal. An example of a heartbeat (a single pulsation of the heart which was recorded) represented in the time-scale domain by a wavelet transform is given in Figure 3.2 (d). Note the drastic improvement made in the representation of the heart cycle by the WT as opposed to the STFT. The resolution appears much better.

A wavelet $\psi(t)$ is defined as a square integrable function for which the following admissibility condition holds true:

$$\int_{-\infty}^{\infty} \frac{|\Psi(\omega)|^2}{|\omega|} d\omega < \infty \tag{3.3}$$

where $\Psi(\omega)$ is the Fourier transform of $\psi(t)$. Essentially, a wavelet is a signal of limited duration that has an average value of zero. If continuity is assumed at $\omega = 0$ the admissibility condition implies that $\Psi(0) = 0$. Thus, the function must change sign.

Examples of wavelets used in this study may be seen in Figure 3.3.

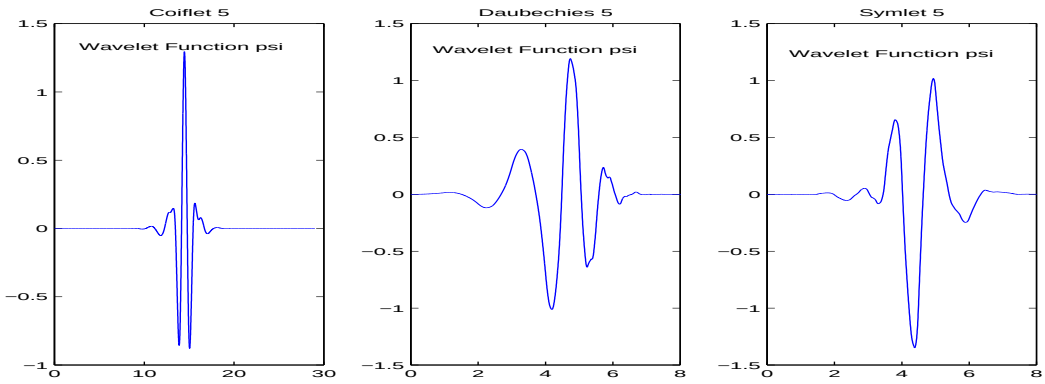

Figure 3.3 Examples of wavelets used in this study

A family of equally shaped functions can be extracted by translating and dilating this

wavelet:

$$\psi_{a,b}(t) = \frac{1}{\sqrt{|a|}}\psi(\frac{t-b}{a})\ a \neq 0, b \in \mathbb{R} \tag{3.4}$$

A signal may be decomposed into members of this wavelet family by using the wavelet transform:

$$(W_\psi f)(a,b) = \int_{-\infty}^{\infty} f(t)\overline{\psi_{a,b}(t)}dt. \tag{3.5}$$

The original signal may be recomposed from the wavelet transform by using the inverse wavelet transform:

$$f(t) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(W_\psi f)(a,b)\psi_{a,b}(t)\frac{\mathrm{d}a\,\mathrm{d}b}{a^2}. \tag{3.6}$$

Because a one-dimensional signal is transformed into a two-dimensional time-scale domain by the wavelet transform, there is redundancy in the signal representation. It is possible to recompose the original signal from a subset of the wavelet coefficients if we restrict the class of wavelets. If the family of scaled and shifted copies forms an orthonormal basis in the space of the square integrable functions, it is possible to recover the exact original signal. Daubechies (1988) discovered a set of these wavelets $\psi_{a,b}$, where

$$\begin{array}{l} a = 2^{-j} \\ b = 2^{-j}k \end{array} \quad (j,k \in \mathbb{Z}). \tag{3.7}$$

The recomposition of the signal from the wavelet-coefficients is quite simple because the decomposition reduces to an orthogonal projection

$$f(t) = \sum_{j,k} c_{jk}\,\psi_{j,k}(t) \tag{3.8}$$

where

$$c_{jk} = \int_{-\infty}^{\infty} f(t)\psi_{j,k}(t)\,\mathrm{d}t. \tag{3.9}$$

The mathematical description of the Continuous Wavelet Transform (CWT) is given by

$$CWT_f^\Psi(b,a) = \Psi_f^\psi(b,a) = \frac{1}{\sqrt{|a|}}\int f(t)\psi^*(\frac{t-b}{a})dt. \tag{3.10}$$

The scale, a, of the wavelet may conceptually be considered the inverse of the frequency. As seen in Figure 3.1 (c), the wavelet is compressed if the scale is low and dilated if the scale is high. Because the WT is computed in terms of scale instead of frequency, plots of the WT of a signal are displayed as time versus scale.

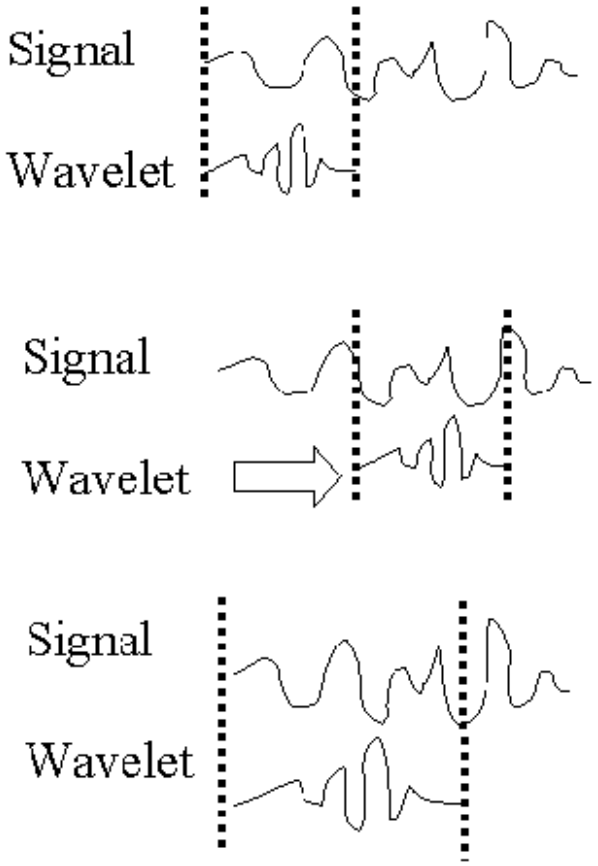

Figure 3.4 This figure illustrates how the CWT is calculated. At the top, the wavelet is compared to section at start of signal. A number is calculated (the wavelet coefficient) showing the degree of correlation between the wavelet and signal section. In the middle, the wavelet is shifted right and the coefficient calculated. This process is repeated for the whole signal. The final diagram, shows the wavelet being scaled and the process is repeated and done for all scales. Modified from Misiti et al. (1996)

The process of computing the CWT is very similar to that of the STFT. The wavelet is compared to a section at the beginning of a signal. A number is calculated showing the degree of correlation between the wavelet and signal section. The wavelet is shifted right and and the process is repeated until the whole signal is covered. The wavelet is scaled and the previous process is repeated for all scales. This process may be seen in Figure 3.4.

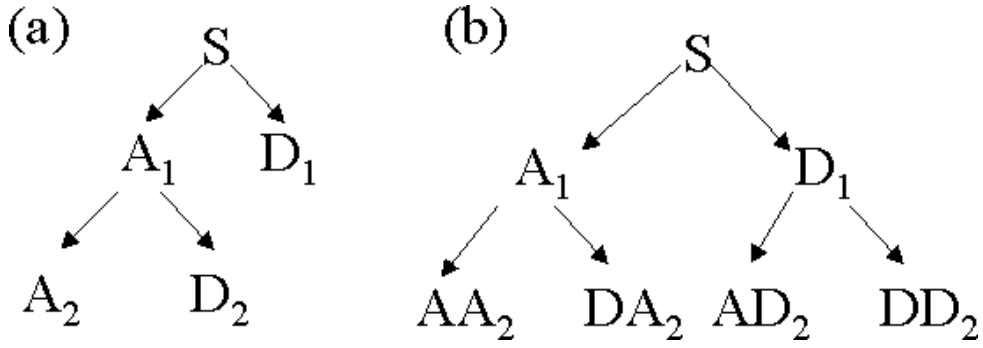

Figure 3.5 This figure illustrates how (a) the discrete wavelet transform decomposes the signal into details and approximations iteratively decomposing the approximations where in (b) wavelet packets iteratively decompose both the approximations and details. S, A, and D represent the signal, approximation and details respectively.

The CWT reveals much detail about a signal, but because all scales are used to compute the WT, the computation time required can be enormous. Therefore, the Discrete Wavelet Transform (DWT) is normally used. The DWT calculates the wavelet coefficients at discrete intervals of time and scale instead of at all scales. The DWT requires much less computation time than the CWT without much loss in detail. With the DWT, a fast algorithm is possible which possesses the same accuracy as other methods. The algorithm makes use of the fact that if scales and positions are chosen based on powers of two (dyadic scales and positions) the analysis is very efficient. Because the algorithm possesses the same accuracy as other methods, this method is often used and is used in the current study. An efficient way to implement this algorithm, using quadrature mirror filters, was developed in 1988 by Mallat and is known as a two-channel sub-band coder (Mallat 1989).

For a single level of decomposition, this algorithm passes the signal through two complementary filters (high-pass which is determined by the wavelet function ψ and low-pass which is determined by the scaling function ϕ) resulting in approximations which are high-scale, low-frequency components of the signal, and details, which are low-scale,

high-frequency components of the signal. This results in twice as many data-points so the data is down-sampled. For further levels of decomposition, successive approximations may be iteratively broken down into details and approximations as shown in Figure 3.5. Then, the signal may be reconstructed by up-sampling, passing the approximations and details through the appropriate reconstruction filters and combining the results. The de-construction and associated reconstruction filters are known as quadrature mirror filters (Meyer & Ryan 1993, Misiti et al. 1996). Because the DWT dyadic algorithm possesses the same accuracy as other methods, this method is often used and is used in the current study.

3.2.3.1 Wavelet Families and Properties

In order to effectively remove noise from a signal, the wavelet decomposition must approximate the signal with the smallest number of non-zero wavelet coefficients possible. It follows that the wavelet family ψ should be chosen so that the function is represented with a few large wavelet coefficients. The main properties that affect this are regularity, number of vanishing moments and the compactness of its support (Hubbard 1996, Mallat 1999).

A wavelet ψ has p vanishing moments if

$$\int_{-\infty}^{+\infty} t^k \psi(t) dt = 0 \qquad \text{for } 0 \leq k < p. \tag{3.11}$$

Therefore, ψ is orthogonal to any polynomial of degree $p-1$. If f is regular and may be approximated over a short period using a Taylor polynomial of degree k and if $k < p$, then the wavelet is orthogonal to this polynomial. The wavelet coefficients will be small for small scales; so a wavelet with a larger number of vanishing moments will represent a smooth function with smaller number of large coefficients. Vanishing moments influence what signal content is picked up by the wavelet transform (Hubbard 1996). With one vanishing moment, linear functions are not seen, and with two vanishing moments, quadratics are not picked up.

The size of support of a wavelet ψ is a measure of the temporal localisation of the wavelet meaning that the size of support is defined over the range which ψ has non-zero values.

Increasing numbers of large amplitude coefficients are generated by peaks in the input signal varying with the width of support of the wavelet. This may be a problem if the signal has many isolated peaks.

If the number of large amplitude coefficients needs to be reduced, the length of support must be shortened and the number of vanishing moments of ψ must be increased. If ψ has p vanishing moments, its support is at least $2p-1$ (Mallat 1999). Thus a compromise must be made. Daubechies wavelets represent the best tradeoff because they have minimum support for a given number of vanishing moments (Daubechies 1992). For Daubechies wavelet of order N, the support length of ψ and ϕ is $N-1$ and the vanishing moment of ψ is N (Misiti et al. 1996).

The order of regularity of a wavelet is the number of continuous derivatives which it possesses (Hubbard 1996). Poor regularity may introduce artifacts (Hubbard 1996). Regularity may be increased by increasing the length of support (Hubbard 1996) which increases with N.

Other properties which wavelets possess are symmetry and orthogonality. All of the wavelets used in this study are orthogonal due to their ability to perfectly reconstruct a signal and the availability of fast algorithms to perform the computation. Table 3.1 gives a summary of the properties of the wavelet families used in this study.

3.2.4 The Wavelet De-Noising Procedure

Wavelet analysis has been recognised to be useful in de-noising non-stationary biomedical signals. Krishnan & Rangayyan (2000) used the wavelet transform to de-noise knee joint vibration signals, Carré et al. (1998) applied the wavelet transform to de-noising uterine electrophysterographies, and Bertrand et al. (1994) employed the wavelet transform to filter and analyse brain-evoked potentials.

The general de-noising process may be outlined in three steps shown below (Misiti et al. 1996).

1. Decompose the signal

Property	**haar**	**dbN**	**symN**	**coifN**
Compactly supported orthogonal	•	•	•	•
Symmetry	•			
Assymmetry		•		
Near symmetry			•	•
Arbitrary number of vanishing moments		•	•	•
Vanishing moments for ϕ				•
Arbitrary regularity		•	•	•
Existence of ϕ	•	•	•	•
Orthogonal analysis	•	•	•	•
Biorthogonal analysis	•	•	•	•
Exact reconstruction	•	•	•	•
FIR filters	•	•	•	•
Continuous transform	•	•	•	•
Discrete transform	•	•	•	•
Fast algorithm	•	•	•	•
Explicit expression	•			

Table 3.1 Summary of the properties of various wavelet families (Misiti et al. 1996). The four wavelet groups from left to right are the haar wavelet (or Daubechies order 1), Daubechies wavelet family, symlets wavelet family, and the coiflet wavelet family.

Choose a wavelet and decomposition level N. Calculate the wavelet decomposition of the signal s at level N.

2. Threshold the detail coefficients

   For each level from 1 to N, select a threshold and threshold the detail coefficients.

3. Reconstruct the signal

   Compute wavelet reconstruction using the original approximation coefficients of level N and the thresholded detail coefficients of all levels from 1 to N.

There are many sources concerning wavelet de-noising theory and possible variation to the method described below (Burrus et al. 1998, Coifman & Donoho 1995, Coifman & Wickerhauser 1998, Matalgah & Knopp 1994, Pan et al. 1999, Prochaázka et al. 1998) but the method (Misiti et al. 1996) described here is the one used in the current study.

Wavelet analysis has the ability to reveal sharp discontinuities in a signal which other techniques such as the Fourier analysis miss. This fact is explained by comparing the poor time resolution possessed by Fourier analysis to the excellent time resolution properties of

the wavelet transform. When the signal is decomposed using Fourier analysis, the signal is decomposed into coefficients that are well localised in frequency but not in time. The rapid changes in the signal are usually represented by a very large number of sinusoids (as in Fourier analysis), but if the correct wavelet is chosen, it may be represented by a smaller number of wavelet coefficients. Because of the efficient decomposition of heart signals by the WT, their wavelet coefficients tend to be much larger than those due to noise which is disordered and scattered throughout the signal. Thus, coefficients below a certain level are regarded as noise and thresholded out. The signal is then reconstructed using the Inverse Discrete Wavelet Transform (IDWT) without significant loss of information. This principle can be applied to almost any ordered signal because the ordered signal will have most of its energy concentrated in a small number of wavelet coefficients whereas the noise will be disorderly and scattered throughout the signal being represented by a large number of small coefficients. Thus, even if the spectrum of the signal and noise overlap, this process may still be used.

There remain two points which must be addressed: how to choose the threshold and how to perform the thresholding.

3.2.4.1 Soft or Hard Thresholding

There are two major methods for thresholding a signal, soft thresholding and hard thresholding. Hard thresholding is defined as

$$\begin{array}{rcll} y & = & x & \text{for } |x| > t \\ y & = & 0 & \text{for } |x| \leq t \end{array} \tag{3.12}$$

and soft thresholding as

$$\begin{array}{rcll} y & = & \text{sign}(x)(|x| - t) & \text{for } |x| > t \\ y & = & 0 & \text{for } |x| \leq t \end{array} \tag{3.13}$$

where $sign(x)$ is defined as

$$sign(x) = \begin{cases} +1 & \text{if } x > 0; \\ -1 & \text{if } x < 0. \end{cases}$$

and x is the original signal, y is the thresholded signal, and t is the threshold. Hard thresholding tends to create discontinuities at $x = \pm t$ because any values of the signal less than the threshold are immediately set to zero. With soft thresholding, the thresholded values are shrunk towards zero without creating the discontinuities.

3.2.4.2 Threshold Selection Rules

There are four threshold selection rules that are available to use with the MATLAB Wavelet Toolbox (Misiti et al. 1996) for optimised wavelet de-noising and are listed in Table 3.2. These threshold selection rules use statistical regression of the noisy coefficients over time to obtain a non-parametric estimation of the reconstructed signal without noise.

Rule Name	Description
rigrsure	selection using the principle of Stein's Unbiased Risk Estimate (SURE)
sqtwolog	fixed form threshold equal to the square root of two times the logarithm of the length of the signal
heursure	selection using a mixture of the first two options mentioned
minimaxi	threshold selection using the minimax principle

Table 3.2 Threshold selection rules

For the soft threshold estimator in the first method, a threshold selection rule which is based on Stein's Unbiased Estimate of Risk (SURE) detailed in Donoho & Jonstone (1992*a*), is used. An estimation of risk for a certain threshold value x_0 is obtained, and then by minimising the risks in x_0, a selection of the threshold value is obtained. The second method uses a fixed form threshold which results in minimax performance multiplied by a factor proportional to logarithm of the length of the signal. The third method is a combination of the first and second methods. If the signal-to-noise ratio is very small (for the third method), the SURE estimate is very noisy. If the signal-to-noise ratio is very small and the SURE estimate is very noisy, then the fixed form threshold is used. The fourth method uses a fixed threshold which is chosen to give minimax performance for mean square error. The minimax principle is used in the field of statistics to achieve the "minimum of the maximum mean square error." Figure 3.6 shows the results of each of the four threshold selection rules applied to a noisy signal.

The previously mentioned threshold selection methods were for white noise. Because other types of noise are quite often present, other threshold estimation models have been developed in Donoho & Jonstone (1992*b*) and Donoho (1995).

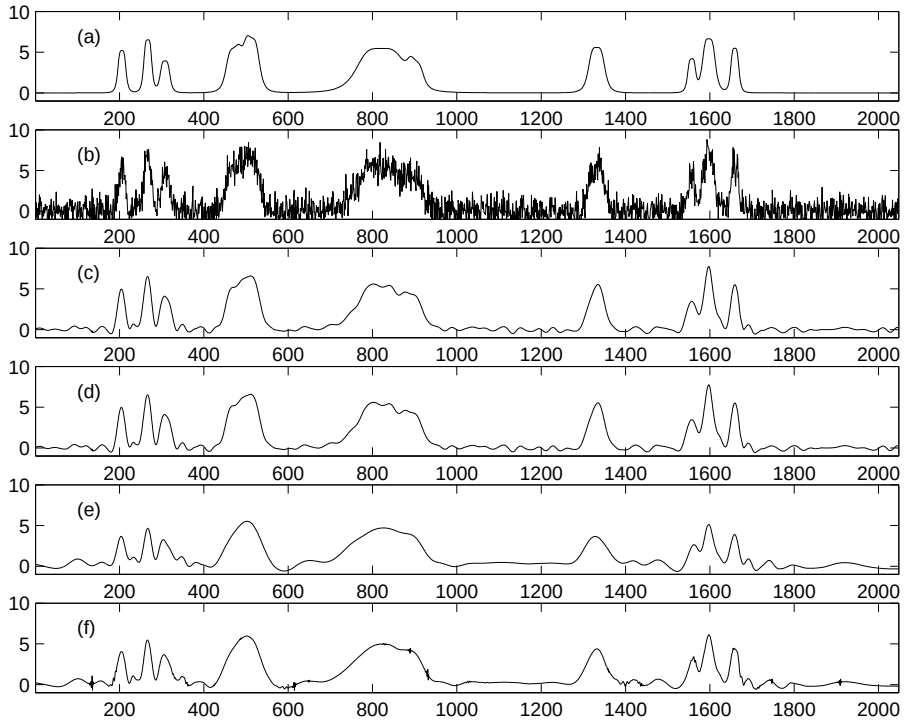

Figure 3.6 Example of applying optimised wavelet de-noising to a signal while varying the threshold selection rules (a) Original signal (b) Noisy signal - Signal to noise ratio = 2 dBs (c) De-Noised signal -heuristic SURE (d) De-Noised signal - SURE (e) De-Noised signal - Fixed form threshold (f) De-Noised signal - Minimax (All are de-noised using a Daubechies 10 wavelet with 10 levels of decomposition)

3.2.4.3 Threshold Rescaling Methods

There are three threshold rescaling methods which are available in the MATLAB Wavelet Toolbox for optimised wavelet de-noising (Misiti et al. 1996). They are listed in Table 3.3 and discussed below.

The underlying model for the noisy signal is basically of the following form

$$s(n) = f(n) + \sigma\, e(n) \tag{3.14}$$

where s is the complete signal, f is the signal without noise, e is the noise, σ is the strength of the noise, and time n is equally spaced (Misiti et al. 1996). The objective of the de-noising process is to suppress the noisy part of the signal s and recover f, which is the signal without noise.

In the simplest model, $e(n)$ is defined as Gaussian white noise N(0,1) and the noise level σ should be equal to 1. This model works well for families of functions f that are represented

with only a few non-zero wavelet coefficients. An example of this would be a function which is smooth in most places with few singularities. In the field of statistics, it would be said that the noise model is a regression model over time and the method may be seen as non-parametric estimation of the function f using an orthogonal basis.

Option	**Corresponding Model**
'one'	Basic model
'sln'	Basic model with unscaled noise
'mln'	Basic model with non-white noise

Table 3.3 Noise model options and corresponding models

Thresholding option *'one'* corresponds to the basic model. Normally, the noise level can be ignored, and it needs to be estimated. The detail coefficients at the finest scale may be considered noise coefficients with standard deviation equal to σ. The median absolute deviation of the coefficients is a good assessment of the noise level. Using a robust estimate of σ is very important because if the first level coefficients contain f details, then these details are represented by few coefficients, if the function f is regular enough and the robust estimate also helps eliminate signal end effects.

The option *'sln'* uses a single estimation of level noise based on the first-level coefficients only to perform threshold rescaling.

Threshold rescaling option *'mln'* performs threshold rescaling using level-dependent estimation of the level noise. When non-white noise e is believed to be present, the thresholds need to be rescaled using level-dependent estimation of the noise. Thus, the *'mln'* option is suitable for non-white noise.

3.3 Wavelet Packets (WP) and De-Noising

Wavelet packets were originally constructed by Coifman and Meyer (Coifman, Meyer & Wickerhauser 1992). Wavelet packet de-noising is very similar to wavelet de-noising, but it offers a wider range of possibilities for signal analysis. For n-levels of decomposition the approximations *and* details are broken down into a further level of details and approximations (as shown in Figure 3.5) resulting in 2^n possible ways to encode the signal (Misiti

et al. 1996). There are more ways of decomposing a signal using WP analysis compared to wavelet analysis because wavelet packet atoms are waveforms which are indexed by 3 parameters, position and scale which corresponds to the wavelet decomposition and frequency, instead of 2 as in the wavelet transform. The analysing window size, frequency and position can each be varied separately. For each orthogonal wavelet function, a library of WP bases is generated which can represent the signal in many combinations. With so many ways to represent the signal, a method must be used to select the best decomposition of the signal. An entropy-based search is performed using the adaptive filtering algorithm which is based on work by Coifman and Wickerhauser. If the reader wishes to know more about wavelet packets and their applications, there are a number of sources which may be consulted (Coifman, Meyer, Quake & Wickerhauser 1992, Coifman & Wickerhauser 1998, Durka & Blinowska 1998, Hubbard 1996, Jiménez, Ortiz, Peña, Charleston, Aljama & González 1999, Krishnan & Rangayyan 2000, Coifman, Meyer & Wickerhauser 1992, Wickerhauser 1994). The following Sections 3.3.1, 3.3.2, 3.3.3 and 3.3.4 provide further information about the methods implemented in MATLAB to generate wavelet packets and for finding the optimal representation of the signal using the many options presented by wavelet packet analysis (Misiti et al. 1996).

3.3.1 Wavelet Packet Generation

Obtaining wavelet packets is not very difficult when using orthogonal wavelets. If we use two filters of length $2N$ and label them $h(n)$, the reversed version of the low-pass decomposition filter divided by $\sqrt{2}$, and $g(n)$, the reversed version of the high-pass decomposition filter divided by $\sqrt{2}$, which correspond to the wavelet. By induction, the following functions $(W_n(x), n = 0, 1, 2, ...)$ may be defined

$$W_{2n}(x) = 2\sum_{k=0}^{2N-1} h(k)W_n(2x-k)$$

$$W_{2n+1}(x) = 2\sum_{k=0}^{2N-1} g(k)W_n(2x-k)$$

where $W_0(x) = \phi(x)$ is defined as the scaling function and $W_1(x) = \psi(x)$ is defined as the wavelet function.

3.3.2 Wavelet Packet Atoms

To obtain wavelet packet atoms, we begin with the functions $(W_n(x), n \epsilon N)$ and examine the family of analysing functions which are given below:

$$W_{j,n,k}(x) = 2^{-j/2} W_n(2^{-j}x - k),\ n \epsilon N,\ (j,k) \epsilon Z^2 \tag{3.15}$$

where k and j respectively correspond to time and scale indices as in wavelet analysis. n is not as easily explained. It can be shown that $W_n(x)$ oscillates about n times. Thus, for fixed values of j and k, $W_{j,n,k}$ analyses the fluctuations of the signal around $2^j \cdot k$ at the scale 2^{-j} and at different frequencies for the allowed values of n.

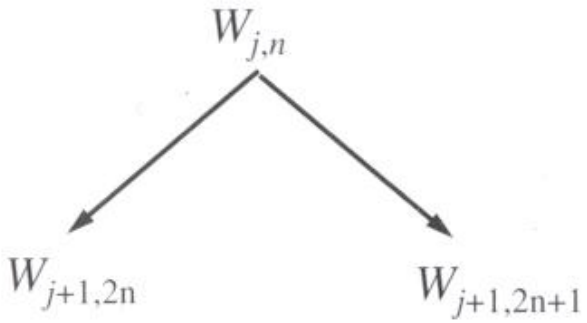

Figure 3.7 Wavelet packet tree (Misiti et al. 1996)

3.3.3 Organising Wavelet Packets in Trees

In the current implementation, WPs are organised in binary trees (Misiti et al. 1996). For the set of functions, $W_{j,n} = (W_{j,n,k}(x), k \epsilon Z)$, is the (j,n) wavelet packet where j and n must be positive integers to be represented in trees. In Figure 3.7, it may be seen that the notation $W_{j,n}$, where j is the scaling factor and n the frequency parameter, is used in the tree. It may be shown that the leaves of every connected binary subtree of the WP tree correspond to an orthogonal basis of the original space. Thus, for a finite energy signal, any WP basis allows for perfect reconstruction, and there are different ways available to represent the signal by making use of the information allocation in frequency scale sub-bands.

3.3.4 Choosing the Best Decomposition

There are numerous ways of decomposing a signal using WP analysis. A signal of length $N = 2^L$ may be expanded in up to 2^N combinations in a binary subtree representation of depth L (Misiti et al. 1996). Because so many possibilities exist for representing a signal, there must be a method to find the optimal representation of a signal for a given criterion efficiently. We are searching for a minimum of the criterion. Entropy-based criteria are very well suited for this task. In this study, we use four different entropy criteria which are defined below (Misiti et al. 1996). In the following formulas s is the signal, and $(s_i)_i$ are the coefficients of s in an orthonormal basis. The entropy E must be an additive cost function so that $E(0) = 0$ and $E(s) = \Sigma_i E(s_i)$. The criterion are:

- The (non-normalised) Shannon entropy is defined as $E1(s_i) = -{s_i}^2 \log(s_i^2)$ so $E1(s) = -\sum_i s_i^2 \log(s_i^2)$ with the convention of log(0)=0.
- The concentration of l^p norm with $1 \leq p < 2$. $E2(s_i) = |s_i|^p$ so $E2(s) = \sum_i |s_i|^p = \|s\|_p^p$.
- The logarithm of the "energy" entropy $E3(s_i) = \log(s_i^2)$ so $E3(s) = \sum_i \log(s_i^2)$ where log(0)=0.
- The threshold entropy $E4(s_i) = 1$ if $|s_i| > \varepsilon$ and 0 elsewhere so $E4(s)$ is such that $|s_i| > \varepsilon$ is the number of time instants when the signal is greater than a threshold ε.

3.3.5 De-Noising with Wavelet Packets

Wavelet packet de-noising is very similar to wavelet de-noising. There are many more decomposition options resulting in greater complexity but more flexibility. Basically, the signal is decomposed into a tree structure using the specified wavelet and level of decomposition. The best decomposition is selected using the criterion specified in Section 3.3.4. The decomposed signal is either soft or hard thresholded removing the noise and then re-composed.

3.4 Use of the Matching Pursuit Method to De-noise Signals

The matching pursuit algorithm was first proposed by Mallat & Zhang (1993) and may be used to de-noise signals (Krishnan & Rangayyan 2000, Zhang, Durand, Senhadji, Lee & Coatrieux 1998). The matching pursuit algorithm decomposes a given signal into a linear expansion of waveforms that are chosen to optimally match the signal structure from a redundant dictionary of functions (Gabor functions in our case). These basis functions have excellent time-frequency properties. The decomposition vectors are selected from a dictionary of waveforms and are chosen based upon the signal properties. The signal $x(t)$ is projected onto a dictionary of time-frequency atoms computed by scaling, translating and modulating a window function $g(t)$ (Krishnan & Rangayyan 2000):

$$x(t) = \sum_{n=0}^{\infty} a_n g_{\gamma_n}(t) \tag{3.16}$$

where

$$g_{\gamma_n}(t) = \frac{1}{\sqrt{s_n}}\, g\left(\frac{t-p_n}{s_n}\right)\, \exp[j(2\pi f_n t + \phi_n)] \tag{3.17}$$

and a_n are the expansion coefficients. The scale factor, s_n, controls the width of the windowing function and p_n adjusts the temporal placement. $\frac{1}{\sqrt{s_n}}$ is used as a normalising factor to restrict the norm of g_{γ_n} to 1. The term, f_n, is the frequency of the exponential function and ϕ_n is the phase of the exponential function. γ_n represents the set of parameters (s_n, p_n, f_n, ϕ_n).

The time frequency atoms are called Gabor atoms because the windowing function is defined as $g(t) = 2^{\frac{1}{4}} \exp(-\pi t^2)$, which is a Gaussian function although we may use many different types of windowing functions (Krishnan & Rangayyan 2000). With Gaussian functions optimal time-frequency resolution is obtained, and the equality criteria of the uncertainty principle is met (Cohen 1989).

The algorithm basically works as described below. The signal is iteratively projected onto the Gabor function dictionary and is decomposed being represented by a vector of a function chosen from the Gabor dictionary. So after the first iteration, the signal is

decomposed into two parts (Krishnan & Rangayyan 2000)

$$x(t) = \langle x, g_{\gamma_0} \rangle \, g_{\gamma_0}(t) + R^1 x(t) \tag{3.18}$$

where $\langle x, g_{\gamma_0} \rangle$ is the projection or mathematically speaking the inner product of $x(t)$ and the first time-frequency atom, $g_{\gamma_0}(t)$. $R^1 x(t)$ is what remains after the first decomposition by projecting $x(t)$ onto the Gabor dictionary and is known as the residue. The process is continued iteratively by projecting the residue onto the appropriate functions chosen from the dictionary. After M iterations, we may express the signal as (Krishnan & Rangayyan 2000)

$$x(t) = \sum_{n=0}^{M-1} \langle R^n x, g_{\gamma_n} \rangle \, g_{\gamma_n}(t) + R^M x(t) \tag{3.19}$$

where $R^0 x(t) = x(t)$. There are two ways of stopping the iterative decomposition of the signal: M may be pre-defined and limited or a check can be performed against the energy of the residue atoms, $R^M x(t)$, and if it small enough the process will be ceased. If M is set to be very large and a zero value for the residue is specified, the signal will be decomposed completely using more computing power for each decomposition performed.

In our case, the decomposition is terminated after extracting the first M coherent structures of the signal. By using a decay parameter, $\lambda(m)$, the first M coherent functions may be determined (Mallat & Zhang 1993):

$$\lambda(m) = \sqrt{1 - \frac{\|R^m x\|^2}{\|R^{m-1} x\|^2}} \tag{3.20}$$

where $\|R^m x\|^2$ is the residual energy at the mth level of decomposition. The decomposition should be continued until the decay parameter does not become any smaller because at this point coherent structures cannot be extracted and the remaining residues which are incoherent structures, may be assumed to be random noise.

The signal is reconstructed using M coherent structures as shown below (Krishnan & Rangayyan 2000)

$$x(t) = \sum_{n=0}^{M-1} \langle R^n x, g_{\gamma_n} \rangle \, g_{\gamma_n}(t) \tag{3.21}$$

3.4.1 Numerical Implementation of the Matching Pursuit with Gabor Dictionaries

This section describes the numerical implementation of the matching pursuit algorithm using the Gabor dictionary where the explanation is given in Mallat & Zhang (1993). For a more complete derivation, please consult Mallat & Zhang (1993). The method described was used in software written by Z. Zhang and may be downloaded by anonymous ftp at the address cs.nyu.edu from the /directory/pub/wave/software (Mallat & Zhang 1993). Our study used this software.

It may be said that for any $\gamma = (s, p, \frac{2\pi k}{N})$ and $\phi \in [0, 2\pi]$, real discrete time-frequency atoms are linked to complex atoms by the relationship

$$g_{(\gamma,\phi)} = \frac{K_{(\gamma,\phi)}}{2}(e^{i\phi} g_\gamma + e^{-i\phi} g_{\gamma^-}). \tag{3.22}$$

where the normalisation constant is defined as

$$K_{(\gamma,\phi)} = \frac{\sqrt{2}}{\sqrt{1 + \Re(e^{i2\phi} \langle g_\gamma, g_{\gamma^-} \rangle)}}, \tag{3.23}$$

with $\Re(z)$ being the real part of the complex number z. For any residue $R^n x$,

$$|\langle R^n x, g_{(\gamma,\phi)} \rangle| = K_{(\gamma,\phi)} |\Re(e^{-i\phi} \langle R^n x, g_\gamma \rangle)|. \tag{3.24}$$

If ϕ is chosen so that it is equal to the complex phase ϕ_γ of $\langle R^n x, g_\gamma \rangle$, then

$$|\Re(e^{-i\phi_\gamma} \langle R^n x, g_\gamma \rangle)| = |< R^n x, g_\gamma >|. \tag{3.25}$$

A search is performed for an index $\tilde{\gamma}_n$ that maximises $|\langle R^n x, g_\gamma \rangle|$ for γ in the subset $\mathbf{\Gamma}_\alpha$ of $\mathbf{\Gamma}$. Using a Newton algorithm, the area around $\tilde{\gamma}_n$ in $\mathbf{\Gamma}$ is examined for an index $\gamma_n = (s_n, p_n, \frac{2\pi k_n}{N}) \in \mathbf{\Gamma}$, where $|\langle R^n x, g_\gamma \rangle|$ is at a local maxima. It may be shown that there exists an $\alpha > 0$ such that

$$|\langle R^n x, g_{(\gamma_n,\phi_{\gamma_n})} \rangle| \geq \alpha \sup_{(\gamma,\phi) \in \mathbf{\Gamma}_\alpha \times [0,2\pi]} |\langle R^n x, g_{\gamma,\phi} \rangle|. \tag{3.26}$$

Because,

$$R^{n+1} x = R^n x - \langle R^n x, g_{(\gamma_n,\phi_{\gamma_n})} \rangle \, g_{(\gamma_n,\phi_{\gamma_n})}, \tag{3.27}$$

for any $\gamma \in \mathbf{\Gamma}_\alpha$,

$$\langle R^{n+1} x, g_\gamma \rangle = \langle R^n x, g_\gamma \rangle - \langle R^n x, g_{(\gamma_n,\phi_{\gamma_n})} \rangle \langle g_{(\gamma_n,\phi_{\gamma_n})}, g_\gamma \rangle. \tag{3.28}$$

must be computed in the next iteration.

So it is estimated that

$$\left\langle g_{(\gamma_n,\phi_{\gamma n})}, g_\gamma \right\rangle = \frac{K_{(\gamma_n,\phi_{\gamma n})}}{2}(e^{i\phi_{\gamma n}} \left\langle g_{\gamma_n}, g_\gamma \right\rangle + e^{-i\phi_{\gamma n}} \left\langle g_{\gamma_n^-}, g_\gamma \right\rangle). \tag{3.29}$$

In order to quickly calculate this inner product, an analytical formula that gives the inner product of two discrete complex Gabor signals is used. For $\gamma_1 = (s_1, p_1, \frac{2\pi k_1}{N})$ and $\gamma_2 = (s_2, p_2, \frac{2\pi k_2}{N})$ and $g(t) = 2^{\frac{1}{4}} e^{-\pi t^2}$, the inner product of two discrete Gabor signals is

$$\begin{aligned} \langle g_{\gamma_1}, g_{\gamma_2} \rangle &= K_{s_1} K_{s_2} \sqrt{\frac{2 s_1 s_2}{s_1^2 + s_2^2}} \exp(-ip_2 \frac{2\pi(k_2 - k_1)}{N}) \times \\ \sum_{m=-\infty}^{+\infty} \sum_{q=-\infty}^{+\infty} \Bigg(& \exp(-\pi \frac{(p_2 - p_1 + mN)^2}{s_1^2 + s_2^2}) \ \exp(-\pi \frac{(k_2 - k_1 + qN)^2}{N^2(s_1^{-2} + s_2^{-2})}) \times \\ & \exp(i \frac{s_2^2}{s_1^2 + s_2^2} \frac{2\pi}{N}(k_2 - k_1 + qN)(p_2 - p_1 + mN)) \Bigg) . \end{aligned} \tag{3.30}$$

If g_{γ_1} or g_{γ_2} is a discrete Dirac or a discrete complex exponential, other formulas must be used.

The numerical complexity for a single matching pursuit iteration is $O(N \log N)$. Comparatively, each iteration needs approximatively about the same CPU time as a Fast Fourier Transform on a signal of N samples.

3.5 De-noising Using Averaging

Averaging is a commonly used noise reduction method (Baykal et al. 1991, Brown et al. 1999, Cozic, Durand & Guardo 1998, Tinati 1998), and is known to reduce white noise because it is randomly distributed throughout the signal. Averaging may be used to produce a "characteristic heartbeat" which is an averaged heart cycle from a series of recorded heart sound cycles (Tinati 1998). Over short periods of time, heartbeats have the same statistical properties. Thus, the signal may be considered quasi-stationary over a short period of time (Tinati 1998).

According to basic probability theory (Beyar et al. 1984), the intensity of a random signal averaging of n cycles is attenuated by $\sqrt{n}$. Thus, if 20 cycles were averaged, random

signals in the recording would be attenuated by a factor of $\sqrt{20} \approx 4.5$ or if 50 cycles were averaged, the attenuation factor would be about $\sqrt{50} \approx 7$.

An important factor to consider in the use of averaging to de-noise PCGs is the type of signal sought. The mechanical activity of the heart can be classified into two categories: "deterministic" and "nondeterministic" (Beyar et al. 1984). In our case, any process that repeats itself exactly for each cycle may be considered deterministic. For non-deterministic events (eg. some heart murmurs), averaging the signal will tend to reduce our ability to discriminate these from deterministic characteristics in the heart sound signal – this is where wavelet de-noising can offer an advantage over simple averaging.

There are two further drawbacks in using averaging (Karpman et al. 1975). When heart sounds are averaged, variations in timing may lead to cancellation of part of the signal. However, with low-frequency signals, this effect is usually negligible. Also, with timing variations, there may be a small overlap of cardiac events which are quite close in time. The duration of these events may appear falsely widened.

3.5.1 Heartbeat Segmentation Algorithms

In order to average the heartbeats, the PCG must be segmented beat by beat. There are many algorithms to do this. For example, Liang et al. (1998) uses the spectrogram to calculate the envelogram of the heart sound signals. The envelogram is thresholded at certain values in order to segment the heartbeats. In an early study, performed by Gerbarg et al. (1962), timing and spectral content is used to automatically identify the first and second heart sounds. Iwata, Ishii, Suzumura & Ikegaya (1980) and Iwata et al. (1983) describe an algorithm which detects the temporal positions of the first and second heart sounds using spectral tracking by linear predictive analysis. Sava & Durand (1997) and Sava et al. (1998) present a method which detects the first and second heart sounds using the matching pursuit method and is relatively robust to background noise. Another method makes use of the ECG as a timing reference and the wavelet transform to treat the PCG (Zhou & Wang 1998). Rajan, Doraiswami, Stevenson & Watrous (1998) used a wavelet based bank of correlators to detect components of the PCG using a Morlet wavelet. Although the segmentation methods presented in Iwata et al. (1983), Liang

et al. (1998), Sava & Durand (1997), and Sava et al. (1998) do not require the ECG to also be recorded simultaneously, the algorithm used for PCG segmentation in the current study uses the ECG as a gating signal as they are both recorded simultaneously. This method is one of the simplest and most common segmentation methods to implement and the hardware and software for this process was readily available for use. The QRS complex of the ECG signals the beginning of the cycle and is used to separate each heartbeat. A description of the complete algorithm is given by Tinati (1998).

3.6 Chapter Summary

In the current chapter, we have introduced the theory and background of the de-noising techniques, wavelet de-noising, optimised wavelet de-noising, wavelet packet de-noising, the matching pursuit technique, and averaging, which were applied to the PCGs in this study in order that the reader may comprehend the results of the de-noising study presented in the next chapter.

Chapter 4

PCG De-noising Study

"Where is the information we have lost in the data?"

Hiroshi Inose (1927-2000)/John Robinson Pierce (1910-)

JAPANESE ENGINEERING PROFESSOR AND AUTHOR/AMERICAN ELECTRICAL ENGINEER, PROFESSOR AND AUTHOR

4.1 Introduction

The PCG is a useful diagnostic tool, but many sources of noise may pollute a PCG making it difficult to extract the relevant information contained in the signal.

There are many techniques used in digital signal processing to remove noise. Many of them rely on actually knowing which part of the signal contains the noise. However, in a PCG, the part of the signal containing the noise is not known *a priori*. Averaging is a technique which has been used to remove noise from PCGs. The WT, WP and the MP methods have also been used to remove noise in non-stationary signals such as the PCG. There has not been a study which has compared the use of the these techniques to obtain an algorithm on how to best de-noise a PCG. Some of the questions investigated by this study are:

- Which method/methods overall is/are more suited for PCG de-noising?
- How many heartbeat cycles need to be averaged in order to obtain the best de-noising result?
- Which wavelets are better suited for de-noising PCGs?
- How many levels of decomposition when using the wavelet and wavelet packet transform de-noising techniques yield optimal de-noising results?
- Which thresholding methods used by the optimal wavelet transform method are best suited for removing noise from of PCGs?
- How many time-frequency atoms should be used to recompose the signal when using the MP method?

4.2 Estimation of Noise in Recorded PCGs

In order to know of what type of noise is contained in the PCG, noise estimates were performed using methods described by Durand, Langlois, Lanthier, Chiarella, Coppens, Carioto & Bertrand-Bradley (1990*a*). In a study done by Durand, Langlois, Lanthier,

Chiarella, Coppens, Carioto & Bertrand-Bradley (1990*a*) on dogs, two tests were performed to estimate the background noise of the recording instrumentation and the physiological background noise which is inherent to the heart/thorax acoustic system. The test which estimates instrumental noise involved exposing the recording microphones to ambient air during quiet periods and recording signals, with efforts made to minimise interference caused by ambient acoustic vibrations. The method used to estimate the background physiological noise in combination with the instrumentation background noise was to compute the power spectrum of 60 milliseconds of the diastolic period of the cardiac cycle (when minimal cardiac activity occurred). The study concluded that both the diastolic and instrumental noise spectrum showed a behaviour approaching white noise. The intensity of the thoracic recording instrumentation noise was significantly less than the combination of the physiological and instrumentation background noise resulting in the conclusion that the instrumentation noise of the thoracic recording was insignificant compared with the physiological background noise.

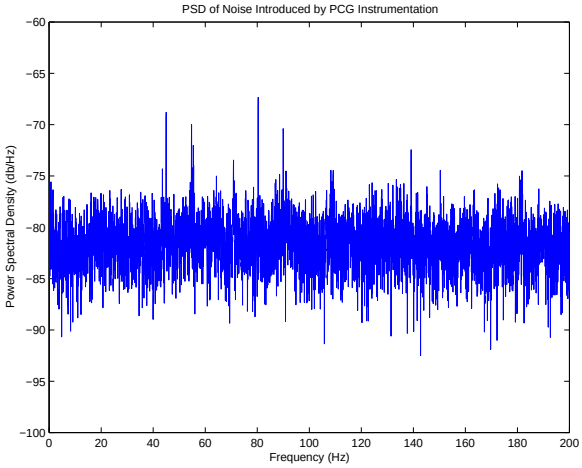

Figure 4.1 Power spectrum (in decibels per Hertz) of the instrumental background noise estimate.

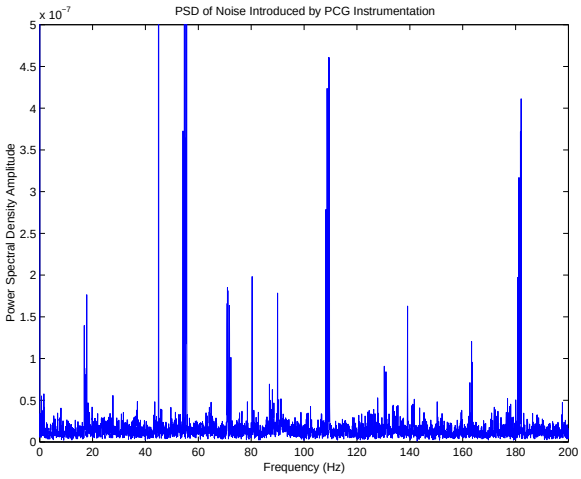

Figure 4.2 Amplitude of power spectrum of the instrumental background noise estimate.

We have used the same two tests to estimate the noise power spectrums. In order to obtain the power spectrum estimate of background instrumentation noise, several recordings were made by holding the microphone up to quiet air. The power spectral density of each recording was calculated, and then they were averaged and may be seen in Figures 4.1 and 4.2. The power spectral densities are shown as relative magnitude (Figures 4.2, 4.4

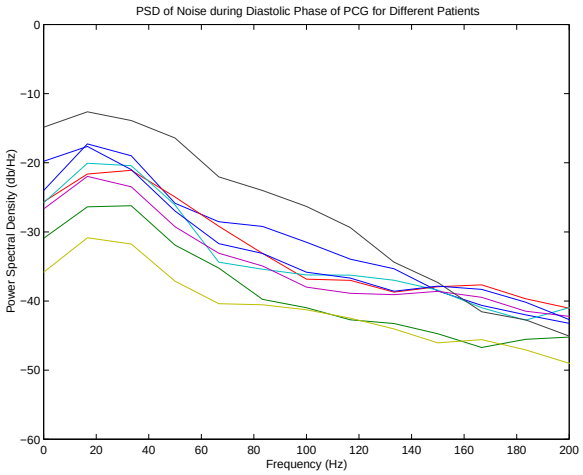

Figure 4.3 Power spectrum (in decibels per Hertz) of the noise estimate for the instrumental and physiological background noise taken during the diastolic phase of the PCG for several patients.

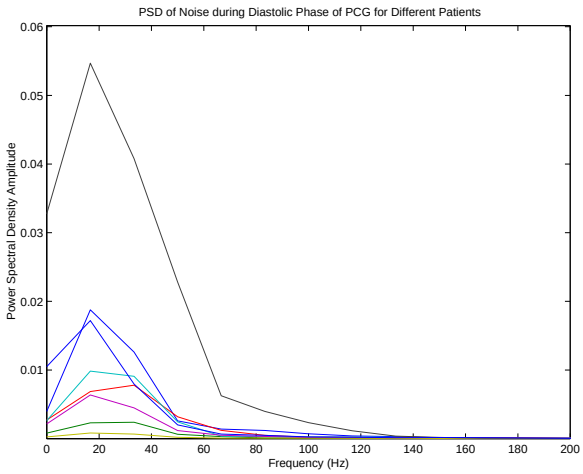

Figure 4.4 Amplitude of power spectrum of the noise estimate for the instrumental and physiological background noise taken during the diastolic phase of the PCG for several patients

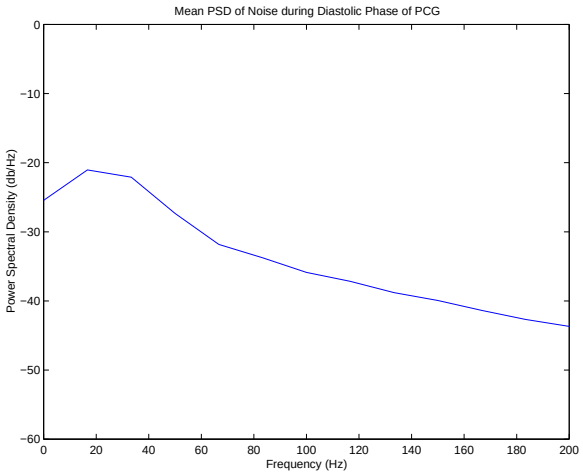

Figure 4.5 The mean power spectrum (in decibels per Hertz) of the noise estimate for the instrumental and physiological background noise taken during the diastolic phase of the PCG for several patients.

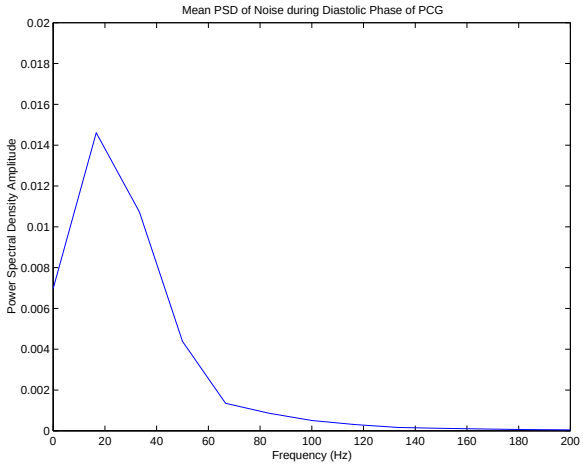

Figure 4.6 The mean amplitude of power spectrum of the noise estimate for the instrumental and physiological background noise taken during the diastolic phase of the PCG for several patients

4.6) and dB/Hz (Figures 4.1, 4.3 4.5) forms are for the ease of the reader. The algorithm used to compute the power spectral density is Welch's method which uses the FFT and time averaged, overlapping periodograms (Rabiner & Gold 1975, Welch 1967). It may be observed that the instrument background noise magnitude is about -80 to -85 dB/Hz and from the spectrum, it appears to be white noise.

To obtain the power spectrum estimate of the background instrumentation and physiological noise, care was taken to select 60 milliseconds of the diastolic period (with minimal cardiac action) of several patients who were a mixture of healthy and unhealthy. The power spectral density using the same method was calculated for the 60 milliseconds diastolic time period for each cycle and then was averaged for all heart cycles. The power spectrum for each patient may be seen in Figures 4.3 and 4.4. The power spectrum of the noise estimates, for each different patient, are all very similar. The power spectrum was then averaged for all patients and may be seen in Figures 4.5 and 4.6. The average power spectrum for the physiological and instrumentation background noise ranges from -20 to -40 dB/Hz which is significantly stronger than that of the instrumentation alone. The power spectrum of both types of background noise is not nearly as flat as that of the instrumentation background noise alone and as a whole resembles $\frac{1}{f}$ noise. However, the majority of the noise is under 40-60 Hz as shown in Figure 4.6, and in this range, a behaviour resembling that of white noise may be noted.

4.3 Measurement of Noise Removal from PCGs

Signal-to-noise-ratio (SNR) is a traditional method of measuring the amount of noise present in a signal. SNR is defined as $10 * \log_{10}(\text{Power}_{\text{signal}}/\text{Power}_{\text{noise}})$ measured in decibels. Two tests are performed using the SNR to measure the performance of optimised wavelet, WT, WP, and MP de-noising (Maple et al. 1999). Because there is currently no known method to calculate which combination of parameters best de-noise a signal when using the optimised wavelet, WT, WP or MP methods, tests must be performed to evaluate the de-noising capabilities of these different combinations. A known amount of noise was added to a "clean" heart sound recording. ("Clean" refers to the fact that although attempts were made to eliminate all environmental noise during the recording, there is

still some noise present in small amounts.) Using various parameters, the de-noising techniques were applied to the heart sound recordings which had white noise added. Then, the SNR of the de-noised signal and the original signal was calculated. The higher the SNR, the less noise there is present.

The procedure performed for the second test is to apply the de-noising techniques to clean PCG recordings and compute the SNR of the resultant signal and the original signal. This test determines how much information from the original signal is lost by the de-noising process. In other words, the more of the original, clean signal that remains after applying the de-noising technique, the better.

The concept of adding a known amount of noise to a clean PCG, then de-noising the signal, and seeing how much noise remains is also employed to measure how well averaging performs as a de-noising technique.

The data presented in most of the following figures represents "typical results" meaning that it is just an example of results normally obtained from conducting the de-noising performance measurements.

4.4 Optimised Wavelet De-noising

Optimised wavelet de-noising differs from ordinary wavelet de-noising in that it has been optimised with special thresholding options including threshold selection rules and threshold rescaling methods. Examining Table 4.1 and Figure 4.7 reveal that no one wavelet seems to give much better results than another. However, the wavelets Daubechies orders 1 and 2 , Coiflets order 1 and Symlets orders 1 and 2 perform worse than the other wavelets. This fact can most likely be explained by examining Figure 4.8 and by the properties of the wavelets themselves. Note that these lower order wavelets appear to lose more information than their higher order counterparts (except the higher order Daubechies wavelets). It is believed that the properties of the wavelets including the support length, regularity, and the number of vanishing moments influence this. For Daubechies wavelet of order N, the support length of ψ and ϕ is $N-1$ and the vanishing moment of ψ is N (Misiti et al. 1996). The order of regularity of a wavelet is the number of continuous derivatives which

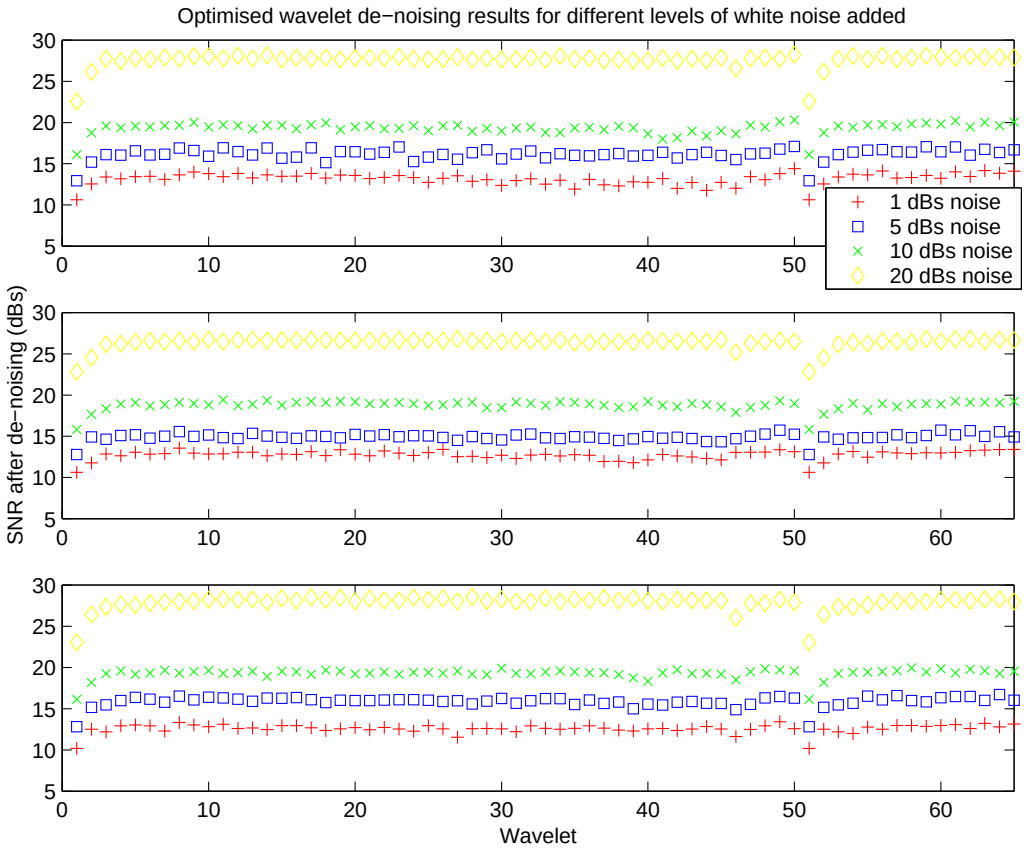

Figure 4.7 This figure shows wavelet de-noising results (as an SNR in dBs) while varying the wavelet used for different levels of white noise added to a three PCG samples (top–patient 15, middle–patient 10, bottom–patient 12). The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15

SNR (in dBs)	**Wavelet**	**Decomp level**	**Wavelet**	**Decomp level**	**Wavelet**	**Decomp level**
	Trial 1		**Trial 2**		**Trial 3**	
1	sym11	10	sym10	7	sym12	10
2	sym6	7	db16	8	coif2	10
3	sym11	10	db6	9	db10	9
4	db11	8	db7	7	sym14	10
5	sym12	10	sym9	8	db6	5
6	sym6	8	sym7	6	db6	10
7	coif5	10	db23	9	coif4	9
8	coif5	6	db31	9	coif4	10
9	coif5	8	db16	5	sym11	9
10	coif5	10	db40	5	sym14	4
11	sym1	10	db40	5	sym12	10
12	coif5	9	sym14	9	db23	9
13	coif5	10	db9	10	sym6	5
14	db14	6	db19	6	sym11	5
15	coif5	5	db10	5	coif4	6
16	sym14	7	sym6	5	sym15	8
17	sym9	4	sym14	5	sym13	9
18	sym9	5	db28	5	coif4	10
19	sym11	6	sym14	4	coif5	4
20	sym9	5	db10	5	db34	6

Table 4.1 Wavelet and decomposition level which obtained best de-noising results after adding a known amount of noise to three different characteristic heartbeats and then applying optimised wavelet de-noising. (coif = Coiflet, db = Daubechies, sym= Symlet)

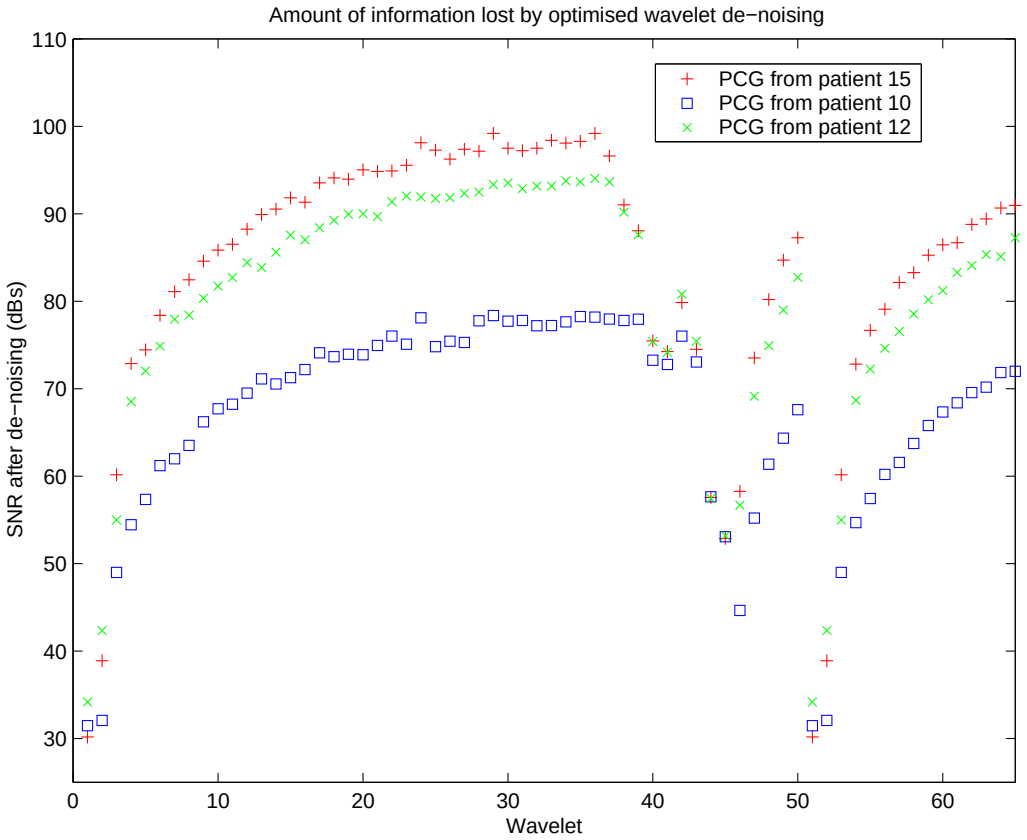

Figure 4.8 This figure shows how much information (as an SNR in dBs) was lost when applying optimised wavelet de-noising to three clean PCG samples while varying the wavelet. The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15 with 4 levels of decomposition used in each case

it possesses (Hubbard 1996). Poor regularity may introduce artifacts (Hubbard 1996). Regularity may be increased by increasing the length of support (Hubbard 1996) which increases with N. Vanishing moments influence what signal content is picked up by the wavelet transform (Hubbard 1996). With 1 vanishing moment, linear functions are not seen, and with 2 vanishing moments, quadratics are not picked up. Thus, by increasing the number of vanishing moments, the lower order components of the signal may be seen. Examining Figure 4.9, reveals that the level of decomposition does not appear

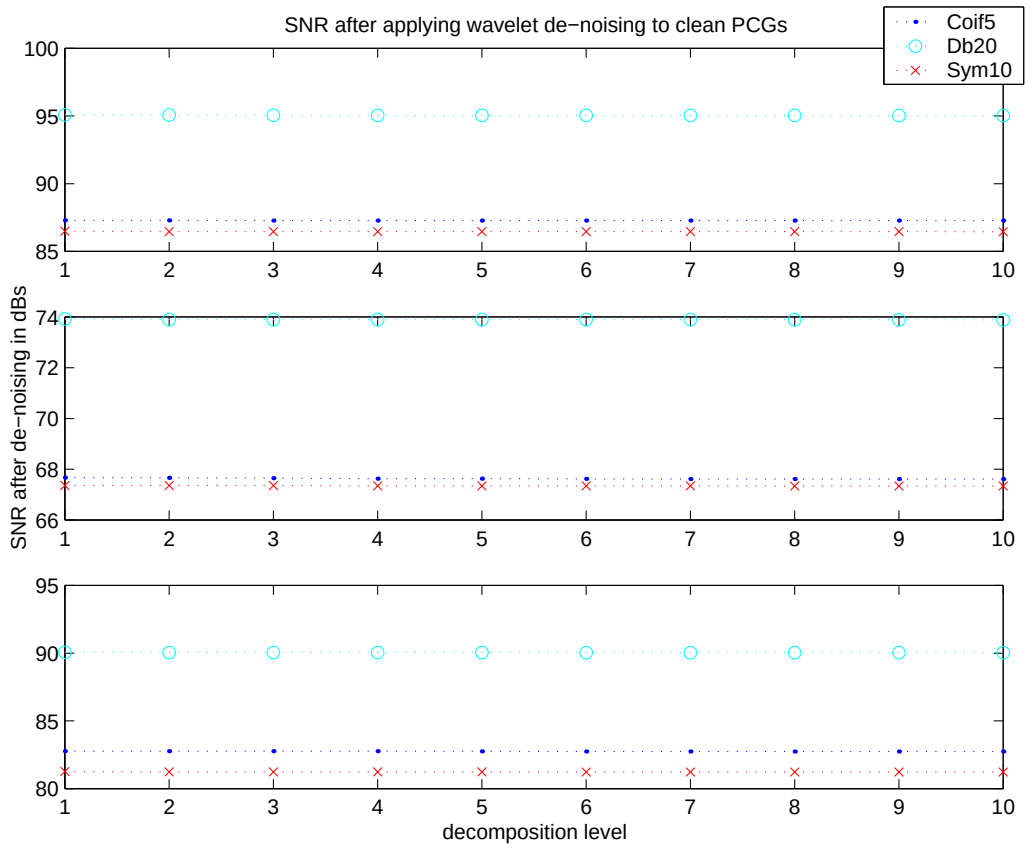

Figure 4.9 This figure shows how much information was lost from optimised wavelet de-noising results (as an SNR in dBs) applied to a three clean PCG samples (top–patient 15, middle–patient 10, bottom–patient 12) while varying the level of decomposition.

to influence the amount of information lost by the optimised wavelet de-noising process. Figures 4.10, 4.11, and 4.12 demonstrate the effect of varying the level of decomposition for the optimised wavelet de-noising procedure with different amounts of noise. It would seem that no additional gains are made by continuing the decomposition level beyond the fourth or fifth level. Additional levels of decomposition increase computation time; thus, we wish to minimise the level of decomposition used if no additional advantage can be

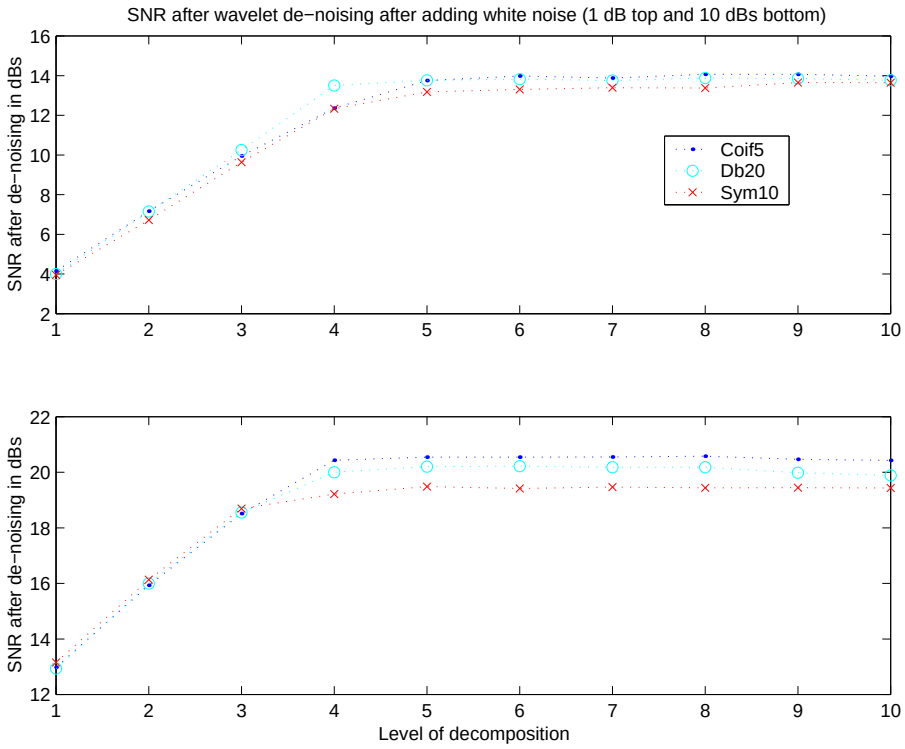

Figure 4.10 The effect of varying the level of decomposition for optimised wavelet de-noising applied to the heart sound recording of patient 15 for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

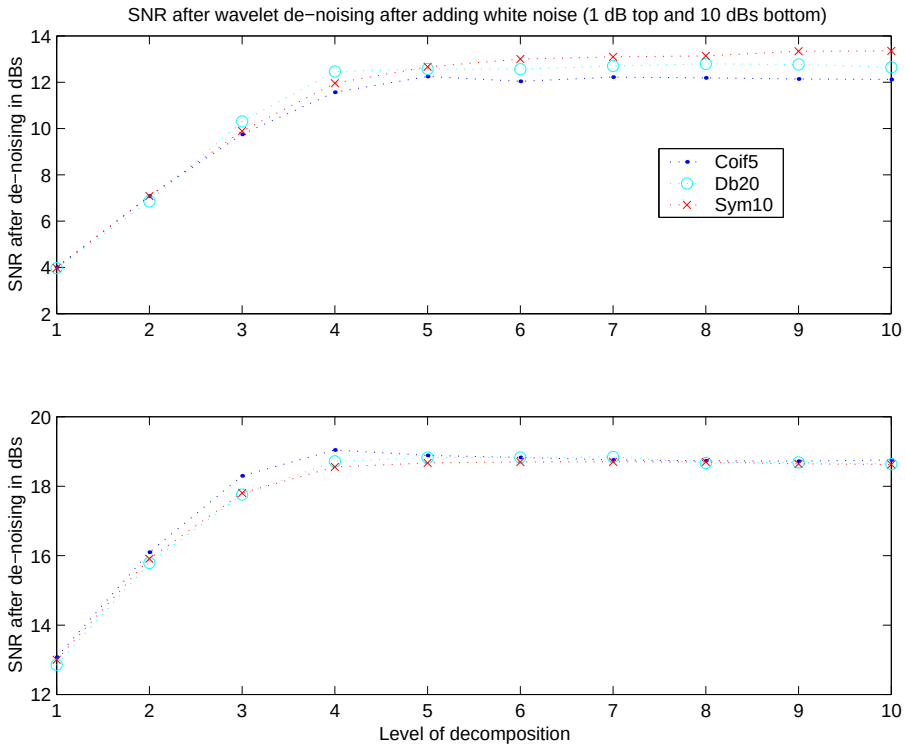

Figure 4.11 The effect of varying the level of decomposition for optimised wavelet de-noising applied to the heart sound recording of patient 10 for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

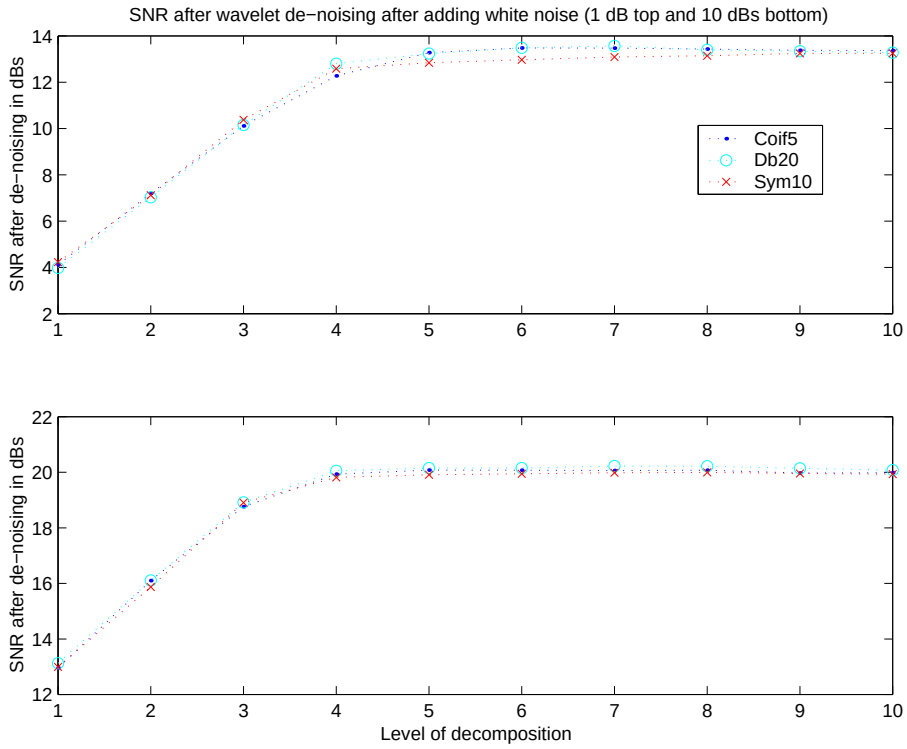

Figure 4.12 The effect of varying the level of decomposition for optimised wavelet de-noising applied to the PCG of patient 12 for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

obtained by further decomposing the signal.

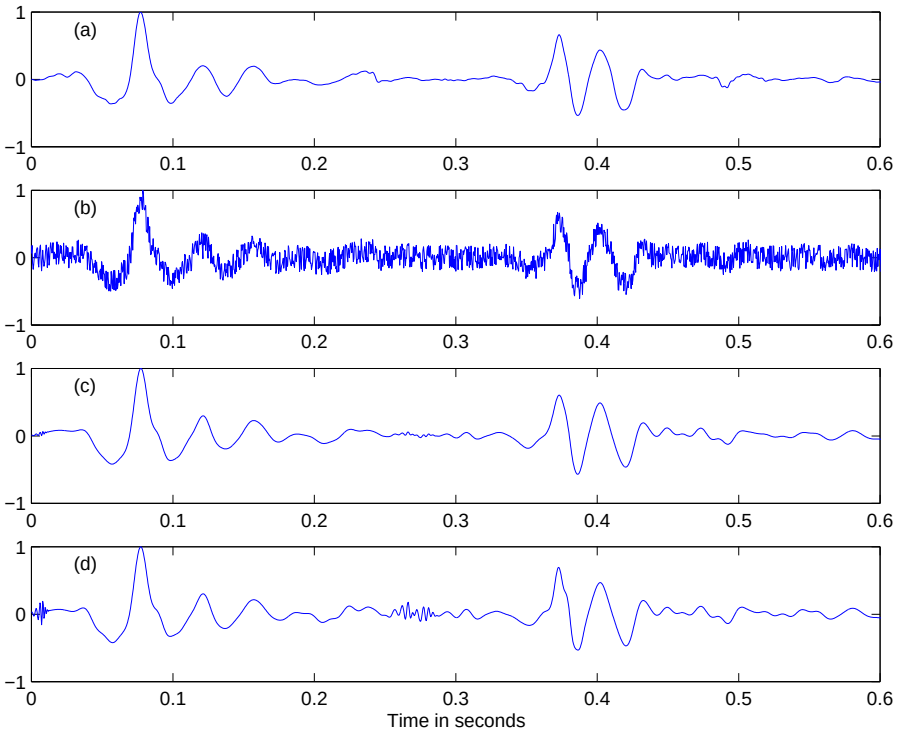

Figure 4.13 This figure demonstrates that hard thresholding can cause discontinuities in a signal. (a) Is the characteristic heartbeat signal, (b) the characteristic heartbeat with white noise added resulting in an SNR of 1 dB, (c) the characteristic heartbeat de-noised using soft thresholding, and (d) the characteristic heartbeat de-noised with hard thresholding. All values have been normalised using the maximum absolute value in (a).

Threshold Selection Rule	**Trial 1**	**Trial 2**	**Trial 3**
"Rigrsure"	13.65	13.78	13.70
"Heursure"	13.27	13.24	12.72
"Minimaxi"	9.30	8.09	9.84
"Sqtwolog"	7.26	8.09	7.37

Table 4.2 Typical SNR results after optimised wavelet de-noising using four threshold selection rules. White noise was added to three different characteristic heartbeats resulting in an SNR of 1 dB. Optimised wavelet de-noising was applied to the characteristic heartbeats using the same parameters (Symlet 14 wavelet with 10 levels of decomposition, soft thresholding, and used the basic noise model with a single estimation level of noise) except for varied threshold selection rules. The SNR was then calculated in dBs. The "rigrsure" outperforms all other methods in every case.

Soft thresholding definitely outperformed hard thresholding in the threshold selection category. Soft thresholding almost always gave a better SNR after de-noising than hard

thresholding. Figure 4.13 shows that hard thresholding may cause discontinuities in a signal.

Of the four threshold selection rules, the minimax and SURE ("rigrsure") threshold selection schemes are more conservative than the others, and therefore should be used when small details of the signal lie in the noise range. The two other schemes remove the noise more aggressively. Because small details of the PCG signal are located in the noise range, it was believed that the "rigrsure" method would be more effective. Our expectations proved to be true. The rigorous SURE method almost always produced better results than the other methods with the heuristic SURE coming in second. Table 4.2 shows typical results for each of the four threshold selection rules.

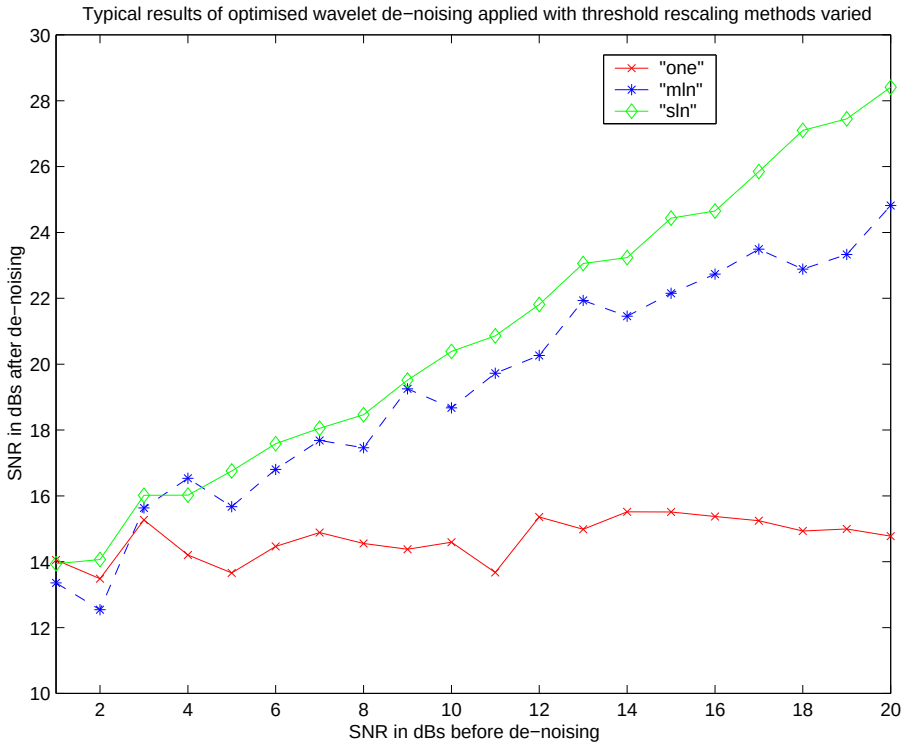

Figure 4.14 This figure is a comparison of the different threshold rescaling methods used by optimised wavelet de-noising. Noise in increments of 1 dB from 1 to 20 dBs is added to a clean characteristic heartbeat. The signal is then de-noised using optimised wavelet de-noising (Daubechies 14 wavelet with 10 levels of decomposition using soft thresholding, and the rigorous SURE threshold selection rule) varying the threshold rescaling methods. Overall, the "sln" method performs the best.

The "sln" method performed the best of three threshold rescaling methods available. For example, in Figure 4.14, with large amounts of noise present, the methods all perform roughly equally with the "sln" method producing slightly better results than the others.

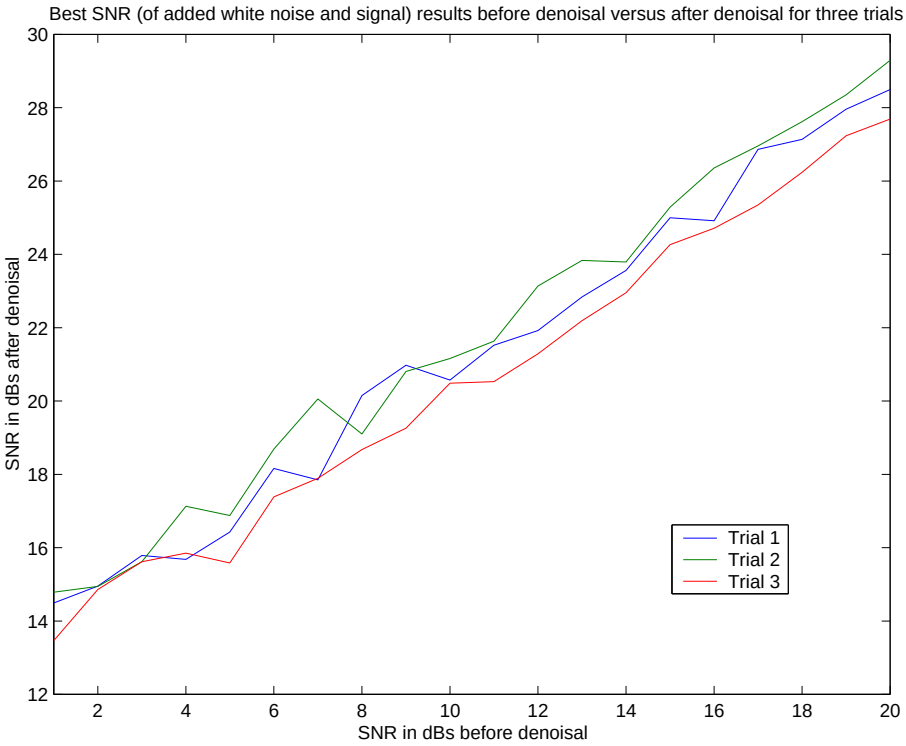

Figure 4.15 Best SNR results before de-noising versus after optimised wavelet de-noising for three trials. Each trial is a different, clean PCG. A known amount of white noise in decibels is added to a characteristic heartbeat. Different combinations of wavelets, thresholding techniques, and levels of decomposition are tried. The highest SNR after de-noising if displayed for each trial. It is interesting to note that the best results for each trial are very similar.

However, as less noise is present, the "sln" method appears to be, by far, the best. The "mln" method is a close second. The "one" method gives about the same SNR after de-noising no matter what amount of noise was present in the signal.

Figure 4.15 shows the best results for adding known amounts of noise to three different PCGs and then applying optimised wavelet de-noising with varying parameters. The best results are very similar proving the results of optimised wavelet de-noising are easily reproducible for heart sounds. Also, the SNRs after optimised wavelet de-noising appear to be relatively linear corresponding to the linear SNR between the signal and the addition of white noise.

Although there was no evidence that a single wavelet was the best suited for de-noising heart sounds, there were some wavelets which were slightly better than others and certain wavelets which would not be recommended for this purpose such as the very low order wavelets. We reached the conclusion that a decomposition level of 5 produced reasonable results while decomposing the signal further often produced marginal benefits and increased computation time. Soft threshold definitively outperformed hard thresholding. Of the four threshold selection rules, the "rigrsure" rule performed the best, and the best choice of the threshold rescaling methods proved to be the "sln" method.

4.5 Wavelet De-noising

Figure 4.16 shows typical results of wavelet de-noising with the wavelet being varied for three different PCGs which had different amounts of noise added. Observing the results of Figure 4.16, it may be seen that the results are similar to those of the optimised wavelet de-noising. No one wavelet produces significant improvements over another with the lower order wavelets again slightly under-performing their higher order counterparts due to insufficient numbers of vanishing moments. Wavelet de-noising appears not perform as well as optimised wavelet de-noising. It is evident from Figure 4.17 that the lower order wavelets again lose more information than the higher order wavelets of the same family. More information is lost by wavelet de-noising than by the optimised wavelet de-noising process.

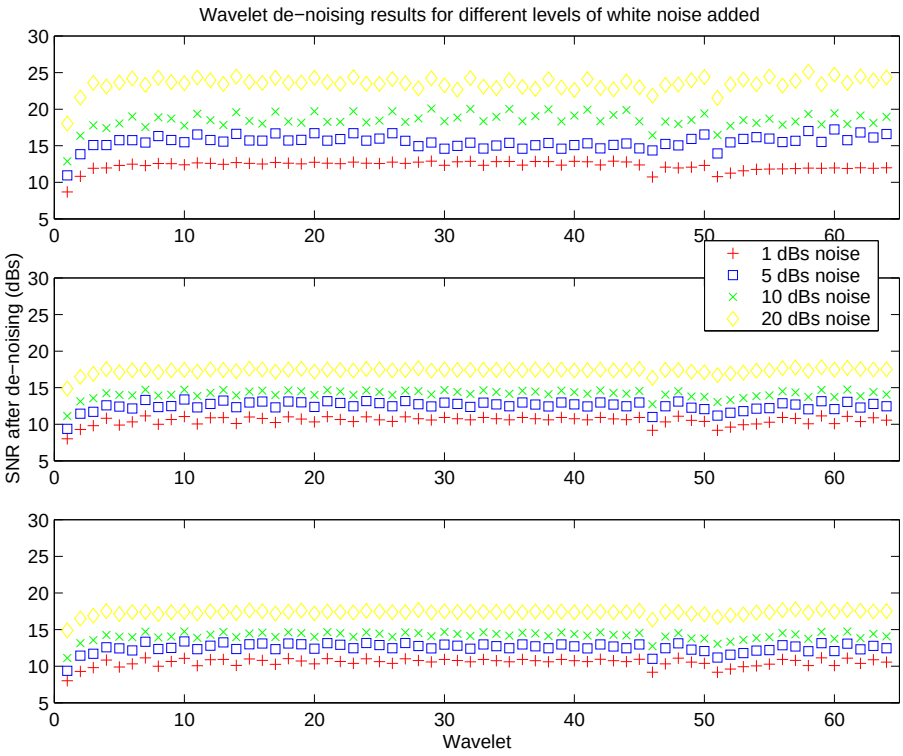

Figure 4.16 Wavelet de-noising results (as an SNR in dBs) for different levels of white noise added to 3 different PCGs. The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15 with 4 levels of decomposition used in each case.

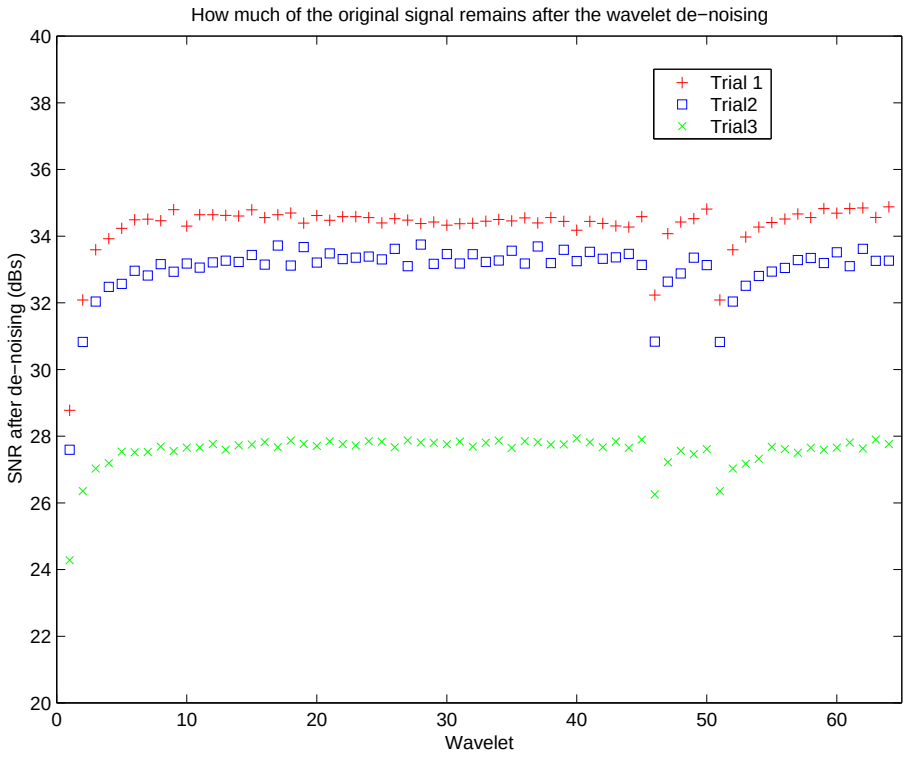

Figure 4.17 How much of the original signal content remains (expressed as an SNR in dBs) after wavelet de-noising is applied to 3 "clean" PCGs. The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15 with 4 levels of decomposition used in each case.

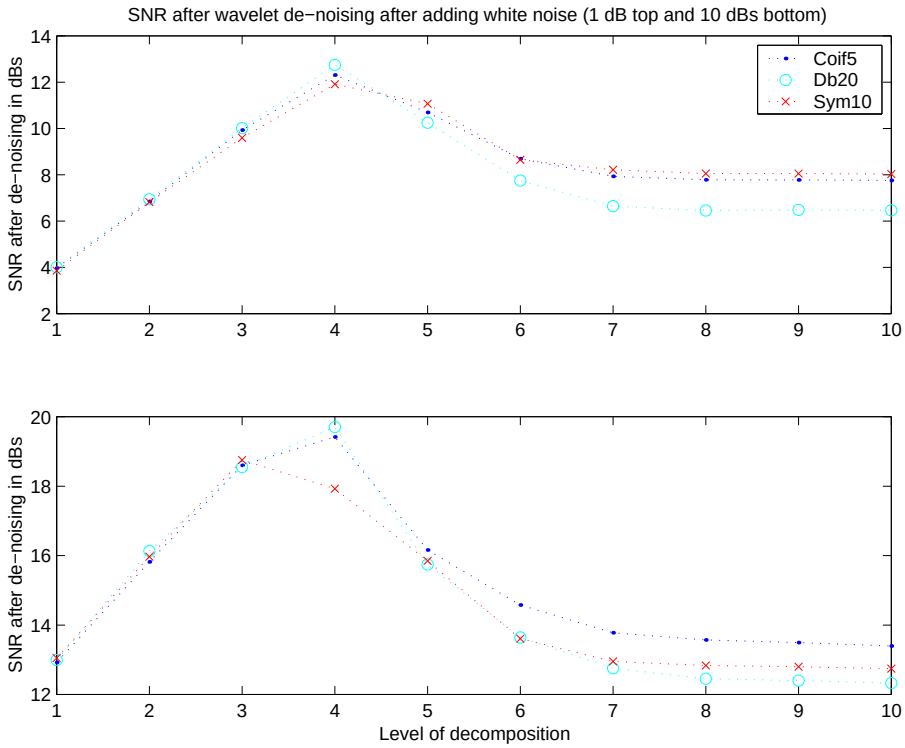

Figure 4.18 Effect of varying the level of decomposition for wavelet de-noising of a PCG (Trial 1) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

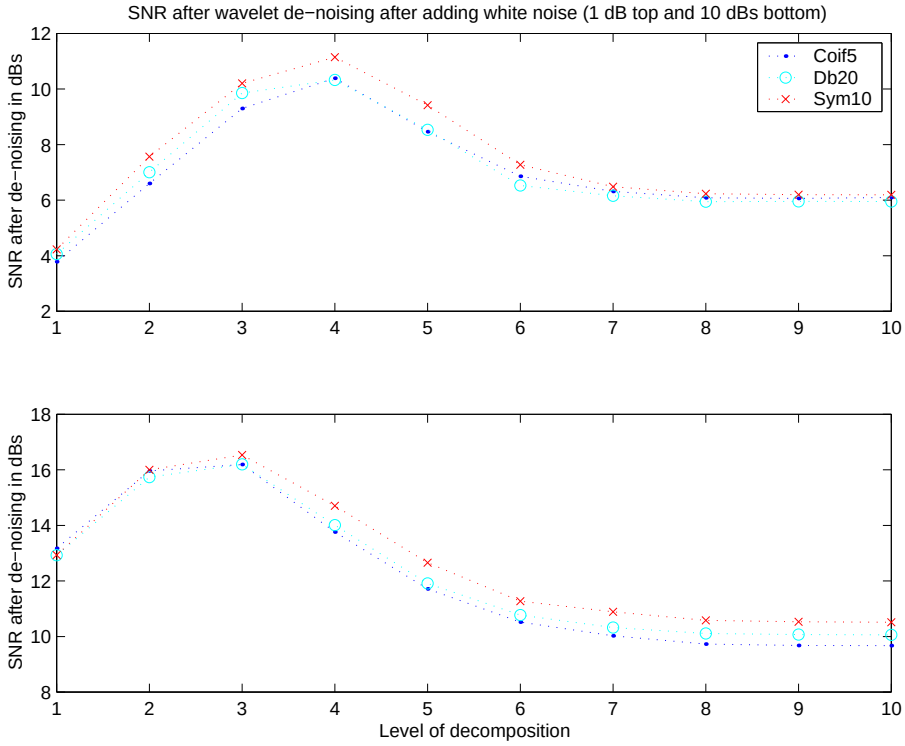

Figure 4.19 Effect of varying the level of decomposition for wavelet de-noising of a PCG (Trial 2) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

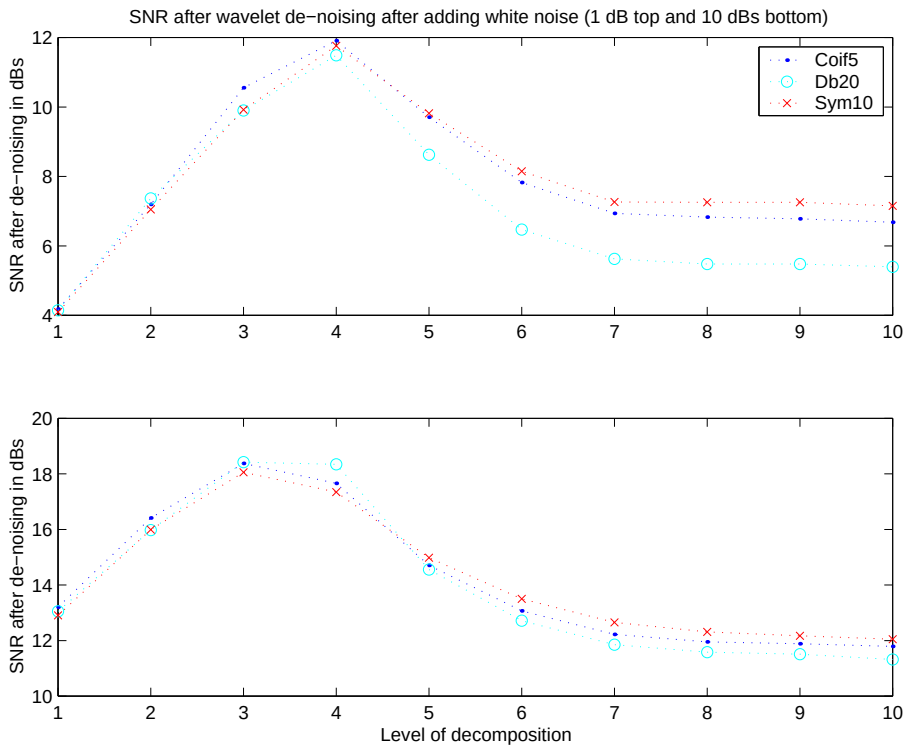

Figure 4.20 Effect of varying the level of decomposition for wavelet de-noising of a PCG (Trial 3) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

To illustrate the effects of varying the number of levels of decomposition used in the de-noising process, we must consult Figures 4.18, 4.19 and 4.20. By surveying the figures, it is observed that 3 or 4 levels of decomposition are optimal with results becoming comparatively quite poorer for successive levels of decomposition. This may be explained by examining Figure 4.21 in which we see that much of the original signal content is lost by successive levels of decomposition. Thus, there is a tradeoff between the information lost as a result of successive levels of decomposition and the gains made in de-noising using further levels of decomposition and is optimised at about 3 or 4 levels of decomposition.

Table 4.3 shows typical results for three different PCGs with various amounts of white noise added after wavelet de-noising.

Overall, most wavelets performed roughly equally except for some of the lower order wavelets for reasons previously explained. Decomposition levels of 3-5 were found to perform the best.

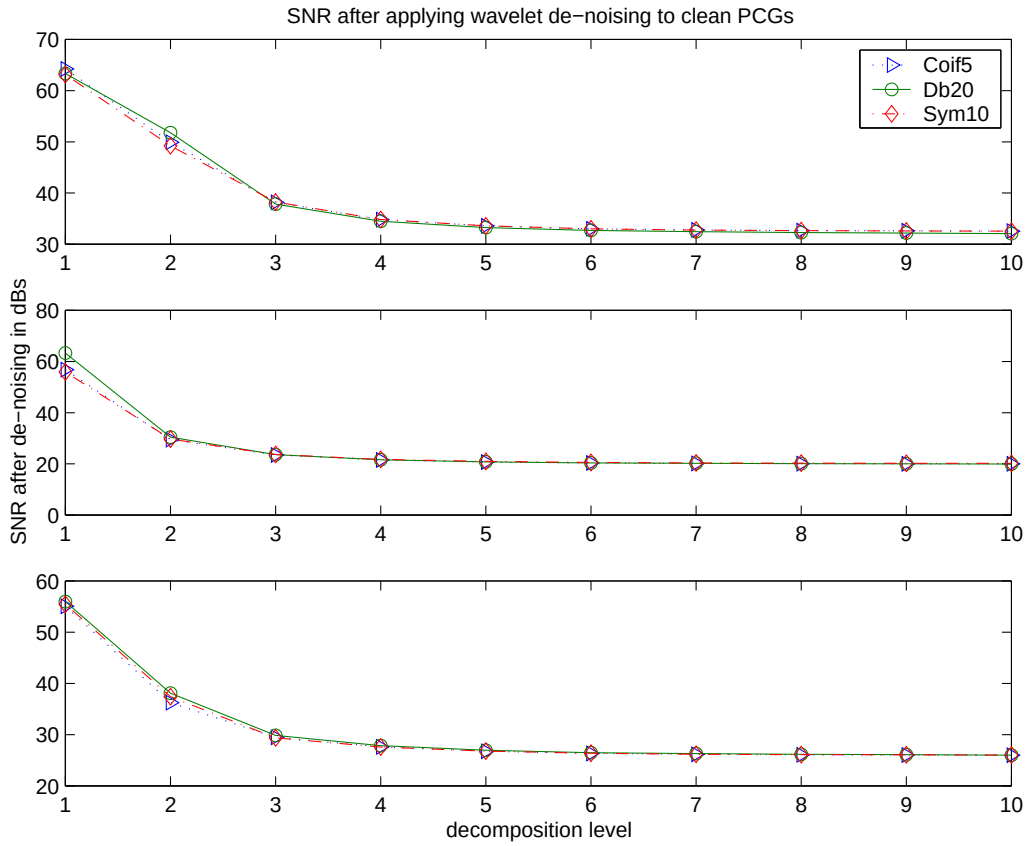

Figure 4.21 Degree of information loss from the wavelet de-noising process (SNR in dBs) when it was applied to a three clean PCG samples while varying the level of decomposition.

Amount of White Noise Added	Trial 1	Trial 2	Trial 3
1 dB	12.89	11.14	11.91
5 dBs	17.24	13.54	14.96
10 dBs	20.07	16.78	18.64
20 dBs	28.22	24.80	25.87

Table 4.3 This table lists the best results (using SNR measured in dB) of all the combinations tried for wavelet de-noising with varying amounts of white noise added. Trial 1 is a 24 year old, healthy, female. Trial 2 is a 43 year old female with hypertension, Trial 3 is an 84 year old female with atrial fibrillation.

4.6 Wavelet Packet De-noising

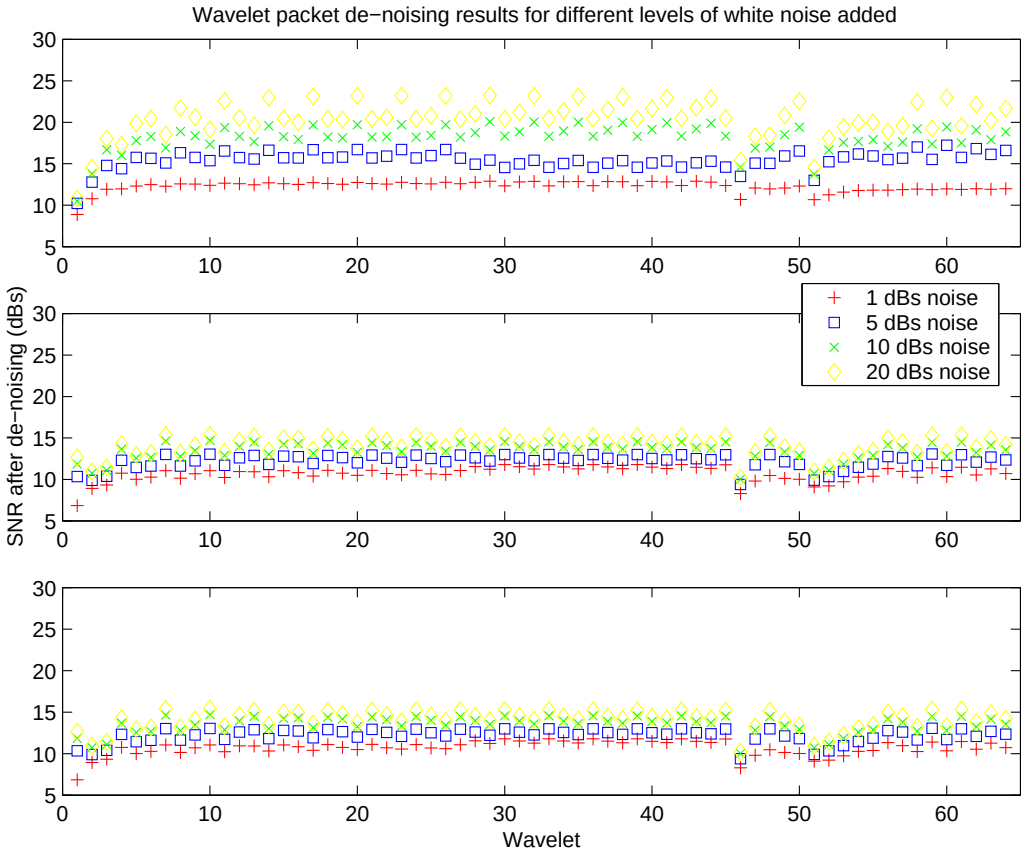

Figure 4.22 Wavelet packet de-noising results (as an SNR in dBs) for different levels of white noise added to a 3 different PCGs. The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15 with 4 levels of decomposition used in each case.

Figure 4.22 shows typical results of wavelet packet de-noising with the wavelet being varied for three different PCGs which have had different amounts of noise added. Observing the results of Figure 4.22, it may be seen that the results are fairly similar to those of the optimised wavelet and wavelet de-noising except that optimised wavelet de-noising again performs better. No one wavelet produces significant improvements over another except that the lower order wavelets again perform poorer than their higher order counterparts due insufficient numbers of vanishing moments. It is evident from Figure 4.23 that the lower order wavelets again lose more information than the higher order wavelets of the same family. More information is lost by WP de-noising than by the optimised wavelet and the WT de-noising process.

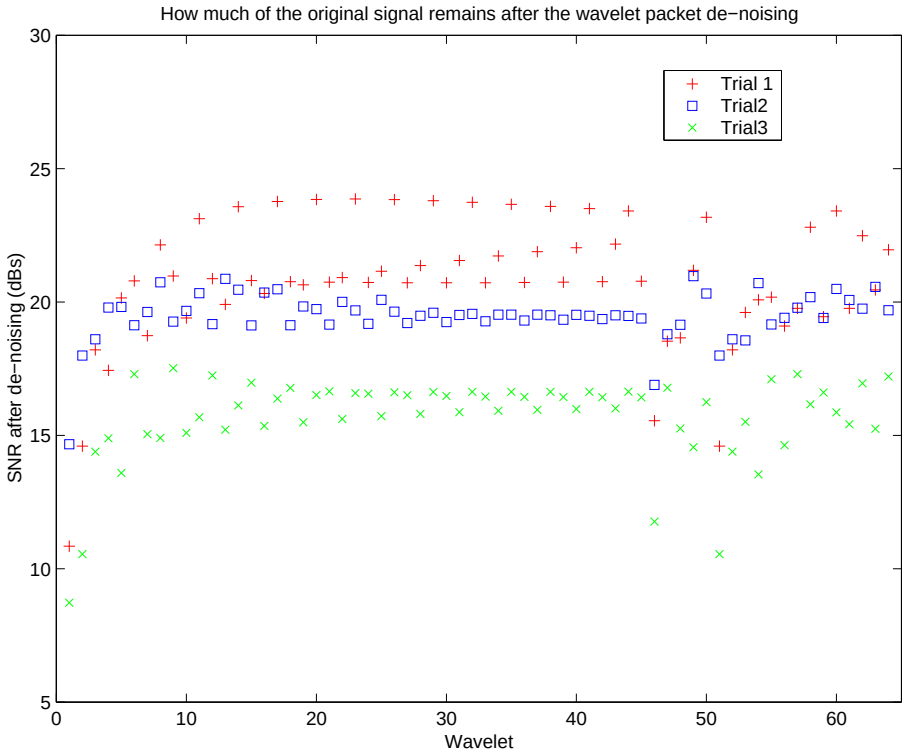

Figure 4.23 This figure shows how much of the original signal content remains (expressed as an SNR in dBs) after wavelet packet de-noising is applied to 3 "clean" PCGs. The x-axis represents the different wavelets respectively: Daubechies Orders 1-45, Coiflet Orders 1-5, Symlets Orders 1-15 with 4 levels of decomposition used in each case.

Amount of White Noise Added	Trial 1	Trial 2	Trial 3
1 dB	13.55	11.79	11.90
5 dBs	17.23	13.15	14.65
10 dBs	20.07	16.90	18.37
20 dBs	28.17	24.97	26.09

Table 4.4 This table lists the best results (using SNR measured in dB) of all the combinations tried for wavelet packet de-noising with varying amounts of white noise added. Trial 1 is a 24 year old, healthy, female. Trial 2 is a 43 year old female with hypertension, Trial 3 is an 84 year old female with atrial fibrillation.

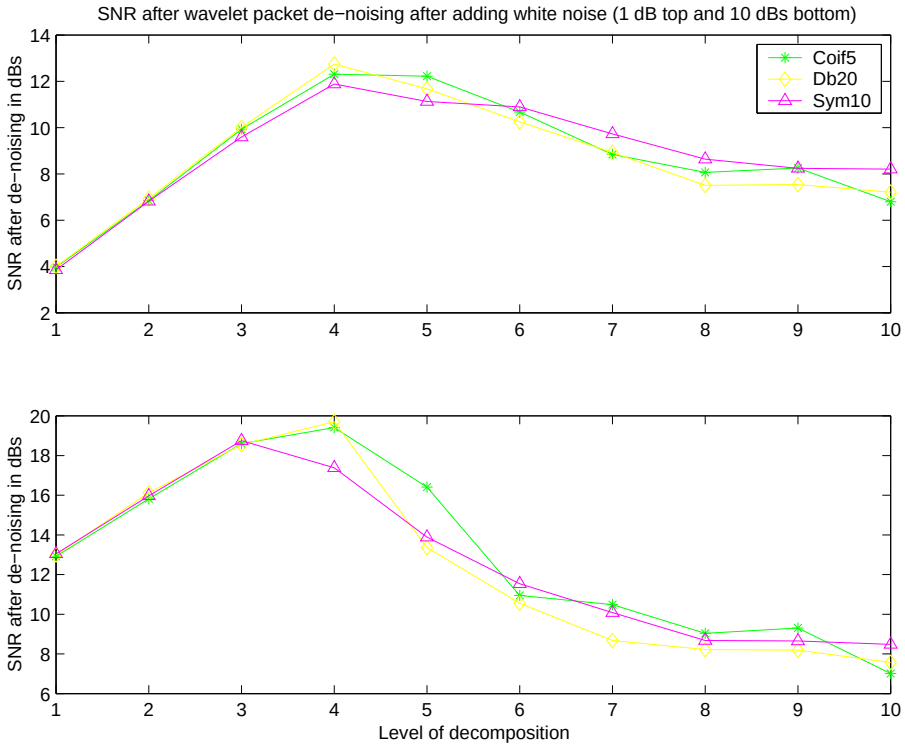

Figure 4.24 Effect of varying the level of decomposition for wavelet packet de-noising of a PCG (Trial 1) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

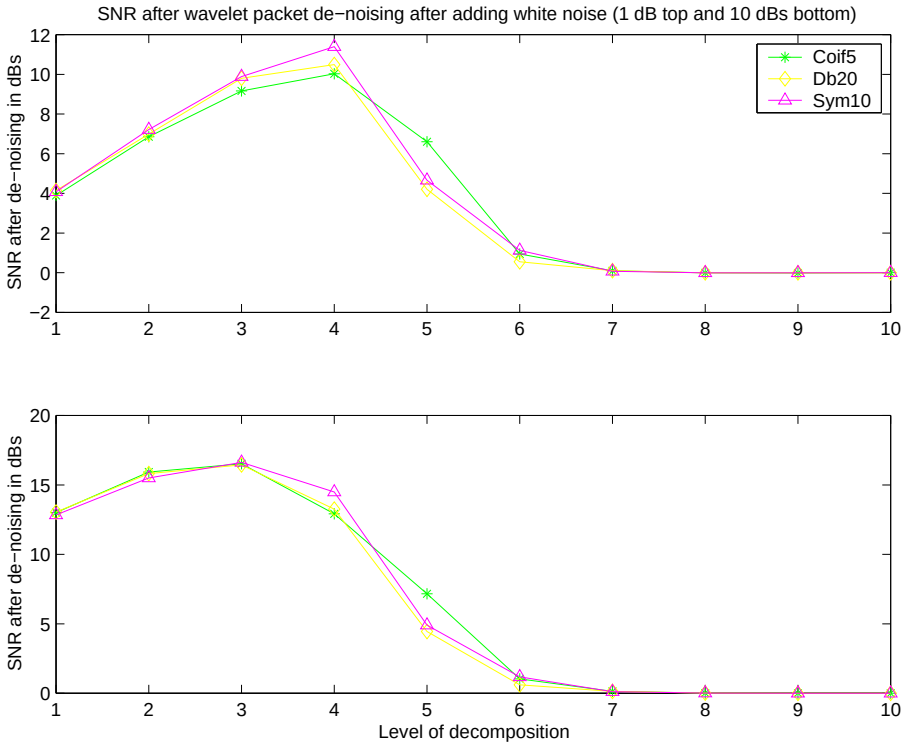

Figure 4.25 Effect of varying the level of decomposition for wavelet packet de-noising of a PCG (Trial 2) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

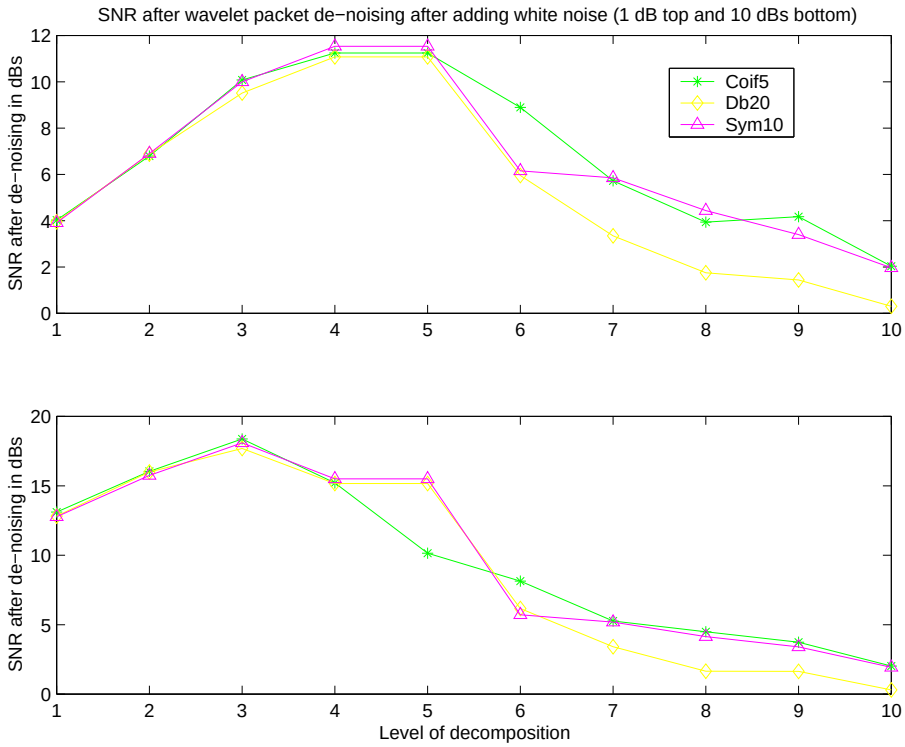

Figure 4.26 Effect of varying the level of decomposition for wavelet packet de-noising of a PCG (Trial 3) for various wavelets with additive white noise at levels of 1 dB and 10 dBs.

Table 4.4 shows typical results for three different PCGs with various amounts of white noise added after wavelet packet de-noising.

To illustrate the effects of varying the number of levels of decomposition used in the de-noising process, we must consult Figures 4.24, 4.25 and 4.26. For large amounts of noise added (as in SNRs of 1 dB), it may be seen that the optimal level of decomposition is approximately 4 or 5. For lesser amounts of noise, it appears that fewer levels of decomposition are needed because with a SNR of 10 dBs the optimal level or decomposition to use is about 3. Examining Figure 4.27, it may be seen that much of the original signal content is lost by successive levels of decomposition. The amount of information lost by the wavelet packet de-noising process increases steadily with additional levels of decomposition until about the fifth or sixth level of decomposition. Thus, there is a tradeoff between the information lost as a result of successive levels of decomposition and the gains made in de-noising using further levels of decomposition.

Overall, most wavelets performed roughly equally except for some of the lower order

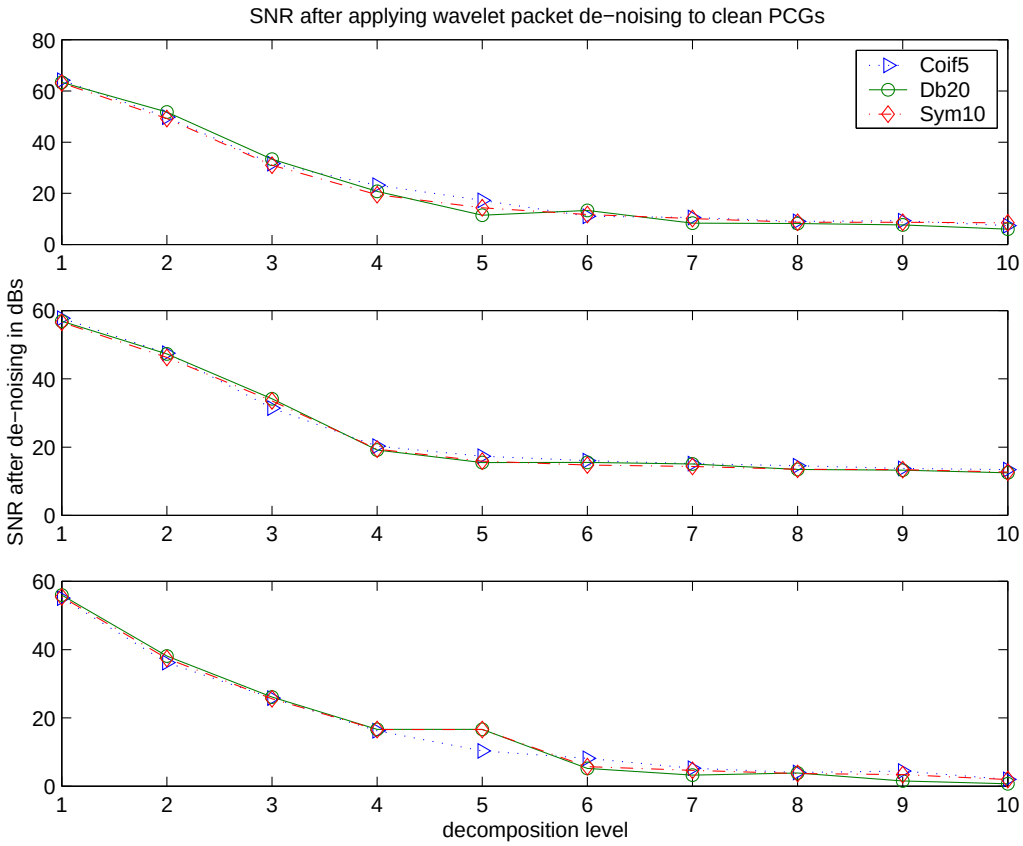

Figure 4.27 The degree of information lost when the wavelet packet de-noising process (measured in SNR in dBs) was applied to three clean PCG samples while varying the level of decomposition.

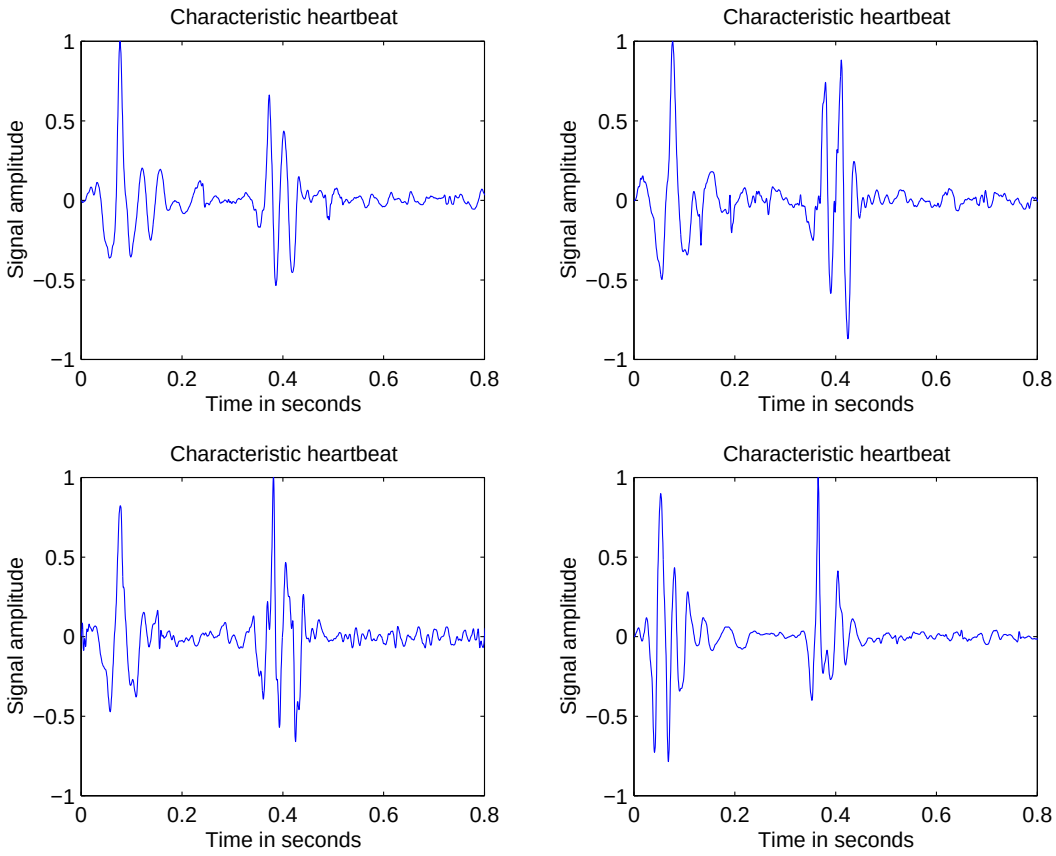

Figure 4.28 Four similar characteristic heartbeats. These four characteristic heartbeats are recorded from a single subject with a normal heart at four different times. Each PCG contained approximately 30 heartbeats and was then averaged. Although, the characteristic heartbeats are not exactly the same, if examined closely, it can be seen that they are very similar. This demonstrates that the recording of a PCG for a person under similar conditions should be easily reproducible, and heart sound recording series over short periods of time are quasi-stationary and thus may be averaged.

wavelets. Decomposition levels of 3-5 were found to perform the best.

4.7 Averaging

Figure 4.28 demonstrates that the recording of a PCG for a person under similar conditions is easily reproducible, and heart sound series over short periods of time are quasi-stationary and thus may be averaged to reliably reduce noise. However, it is important to remember that there are certain drawbacks to averaging (Karpman et al. 1975). Variations in the timing of individual cycles may lead to cancellation of part of the signal, but, with low-frequency signals, this effect is usually negligible. Also, with timing variations,

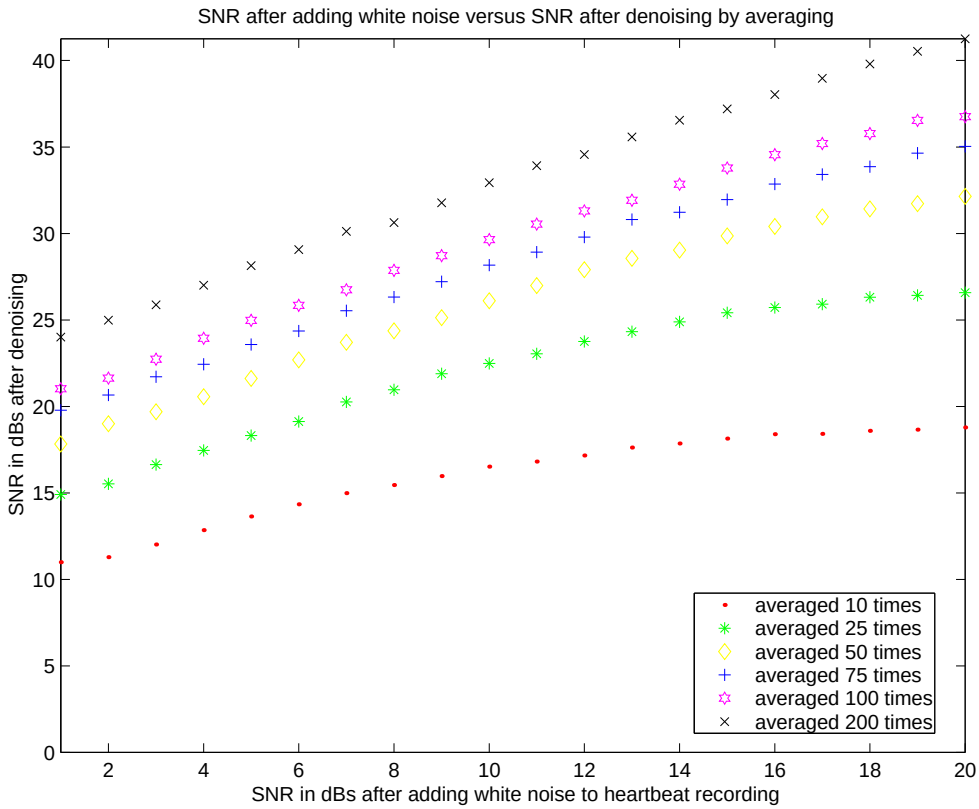

Figure 4.29 The SNR after adding white noise to a series of heart sound cycles versus the SNR after averaging these series of heart cycles to obtain a characteristic heartbeat and reduce noise. There seems to be marginal improvement in SNR when little noise is present in the signal and the signal is not averaged a fair number of times. For example, with an SNR of 1 dB, after averaging the signal 10 times, we see the noise levels decrease as the SNR approaches 11 dBs, but with an initial SNR of 20 dBs and averaging 10 heart cycles, the SNR is still about 20 dBs meaning the noise remains. Averaging a series of 50-75 cycles seems to give the best results here in terms of recording and computation time tradeoff.

there may be a small overlap of cardiac events which are quite close in time resulting in falsely widening the duration of these events.

Averaging seemed to produce significant improvements especially if there was a large amount of noise present in the signal. Figure 4.29 shows that averaging a series of 50-75 heart sound cycles seems to give the best result in terms of recording and computation time tradeoff as is difficult to obtain a long, clean recording and also increases the computation time required.

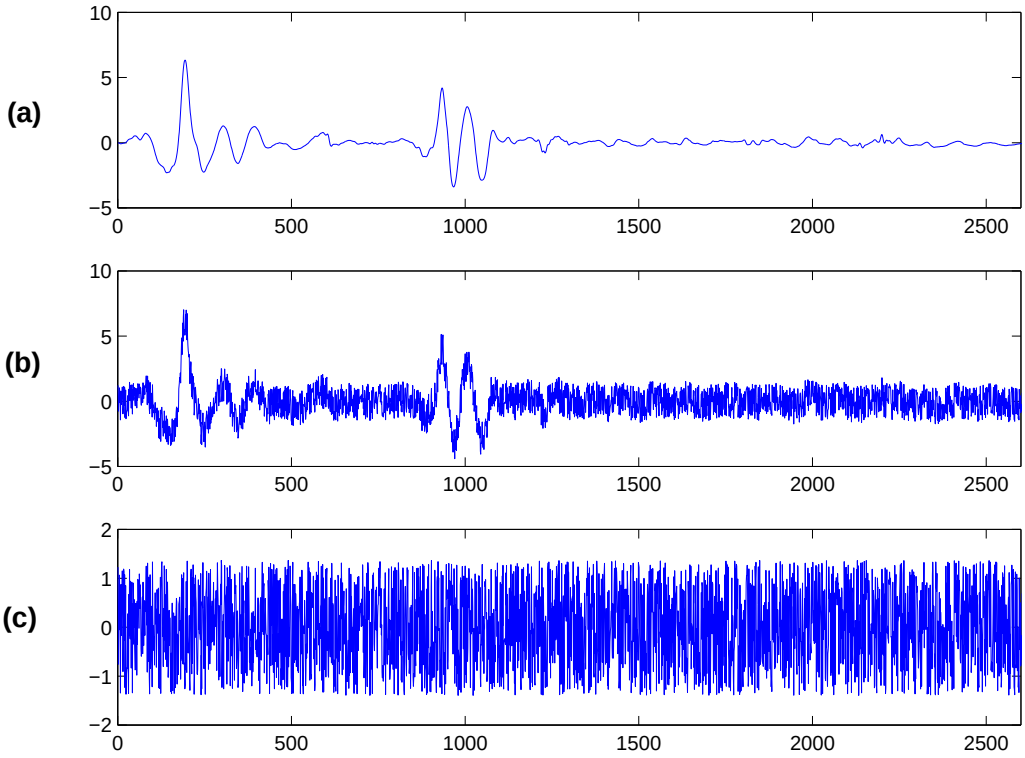

Figure 4.30 This figure shows (a) a clean heart sound cycle (from patient 15), (b) the heart sound cycle with 1 dB additive white noise, and (c) the additive white noise.

4.8 Matching Pursuit

The matching pursuit method (described in Section 3.4) was used to de-noise PCGs. The results of applying the two de-noising tests previously described in Section 4.3 are presented in this section. Figures 4.30 and 4.31 show the effect of varying the number of time-frequency atoms used to reconstruct the signal. Figure 4.30 shows a typical heart sound cycle, then the cycle with additive white noise and finally the noise itself. In Figure 4.31, we can see the effect of varying the number of time-frequency atoms used in the MP de-noising process. Using only a single time-frequency atom does not yield much detail and leaves out the second heartsound. Increasing the number of time-frequency atoms to three, the second heartsound becomes visible. Using seven time-frequency atoms, reveals further detail. More detail is obtained by increasing the number of time-frequency atoms further. There is a point when increasing the number of time-frequency atoms used does not really yield much more information about the signal and only increases computation time.

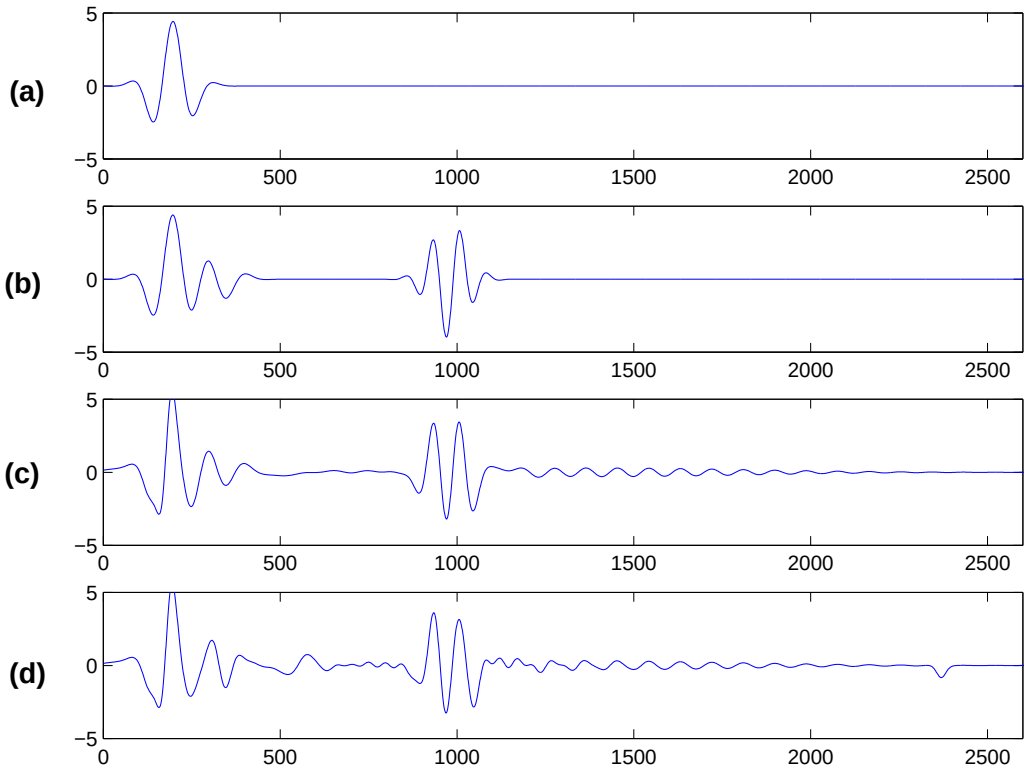

Figure 4.31 This figure shows the reconstruction of the heart sound cycle (from patient 15) shown in Figure 4.30 (b) after matching pursuit de-noising. In (a) the signal is reconstructed from a single time-frequency atom, (b) from 3 time-frequency atoms, (c) from 7 time-frequency atoms and (d) from 11 time-frequency atoms.

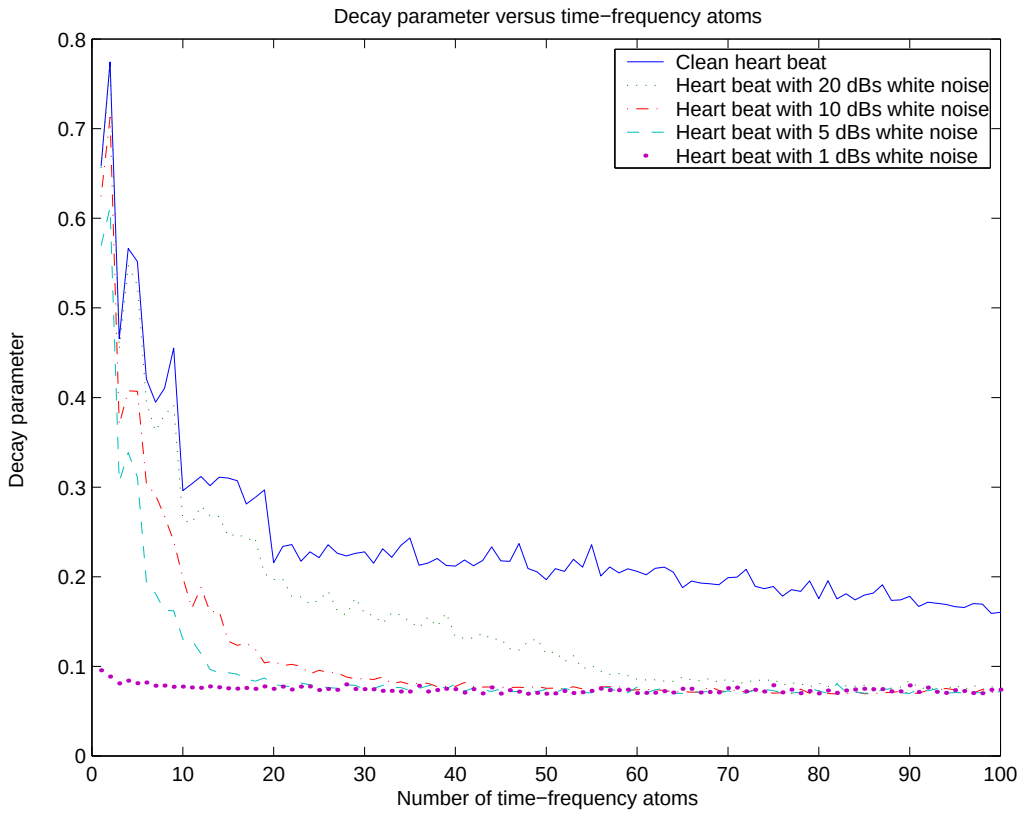

Figure 4.32 Plot of the decay parameter against the number of time-frequency atoms used by the MP method for the heart sound cycle (from patient 15) shown in Figure 4.30.

This point where decomposition should be stopped can be measured using a decay parameter. When the decay parameter levels out, no further significant benefit is obtained by using increasing numbers of time-frequency atoms as the structures left in the signal are incoherent which are unable to be extracted by the MP method and are assumed to be noise. Figures 4.32, 4.33, and 4.34 are a plot of the decay parameter against the number of time-frequency atoms used by the MP method which is applied to different heart sound cycles with various amounts of noise added. The point where the decay parameter levels off is where the decomposition should be stopped. Note that with increasing levels of noise the number of time-frequency atoms used becomes smaller. This means that the noise is interfering with the coherent, structured part of the signal making it possible for fewer coherent structures to be extracted from the signals. In Figures 4.35 and 4.36, it may be seen how the SNR in dBs between the original signal and de-noised signal and the decay parameter act as the number of time-frequency atoms used varies with different amounts of noise. As the decay parameter levels out, the best de-noising results are achieved. Further decomposition only uses more computational power.

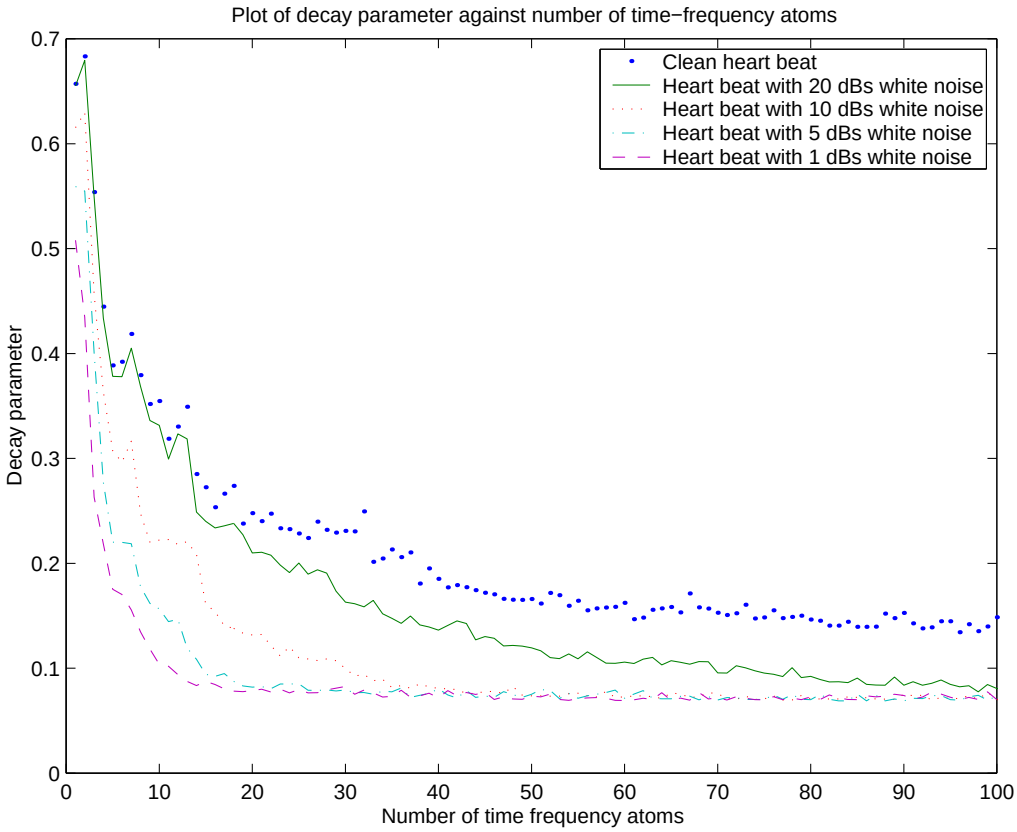

Figure 4.33 Plot of the decay parameter against the number of time-frequency atoms used by the MP method for another heart sound cycle (from patient 10).

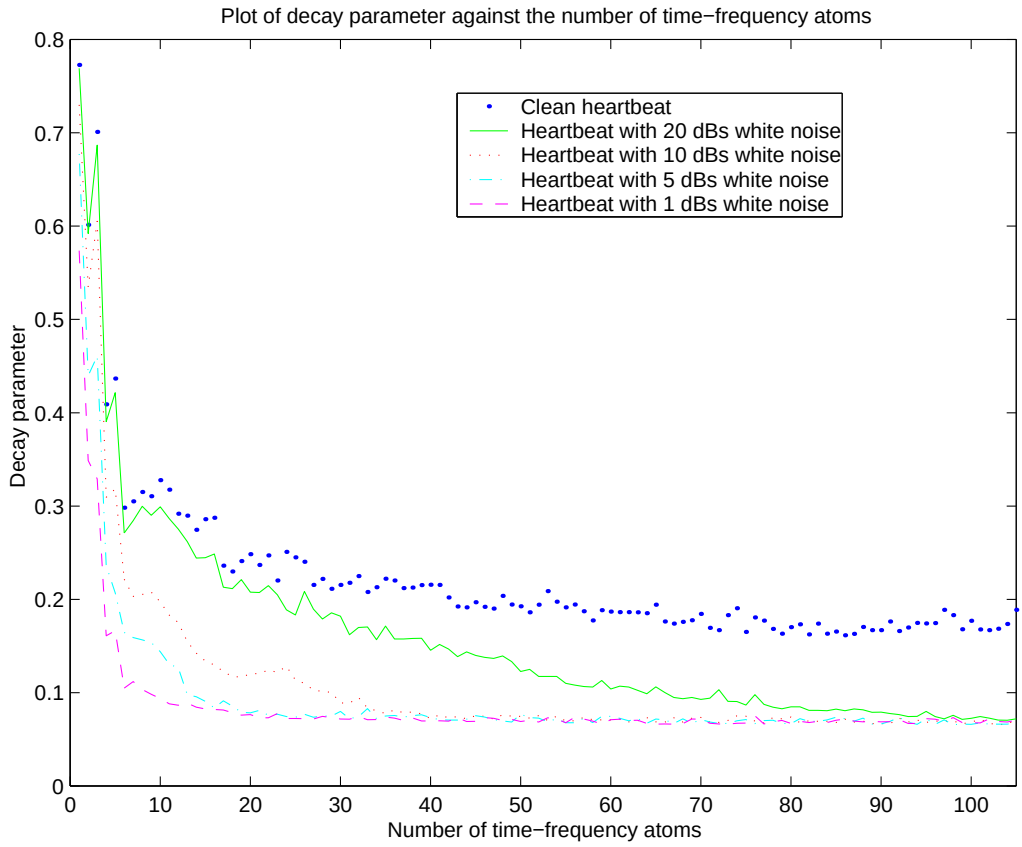

Figure 4.34 Plot of the decay parameter against the number of time-frequency atoms used by the MP method for the another heart sound cycle (from patient 12).

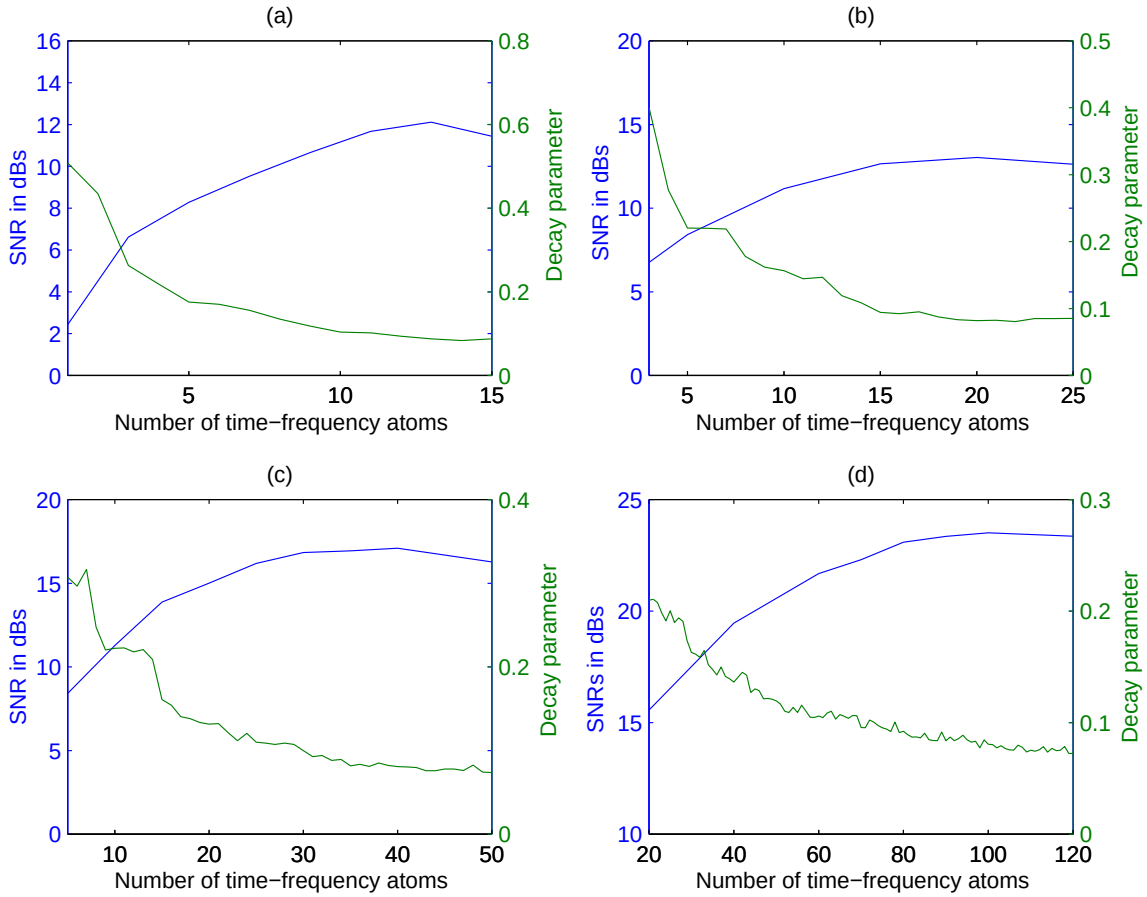

Figure 4.35 This figure shows (a) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to a heart sound cycle (patient 10) with a SNR of 1 dB, (b) the decay parameter and the SNR plotted against the number of time-frequency atoms after applying MP de-noising to the same recording with a SNR of 5 dBs, (c) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to the same recording with a SNR of 10 dBs and (d) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to the same recording with a SNR of 20 dBs additive white noise. The decay parameter is plotted in green and the SNR in dBs is plotted in blue. It may be observed that as the decay parameter levels out, the best de-noising results are achieved.

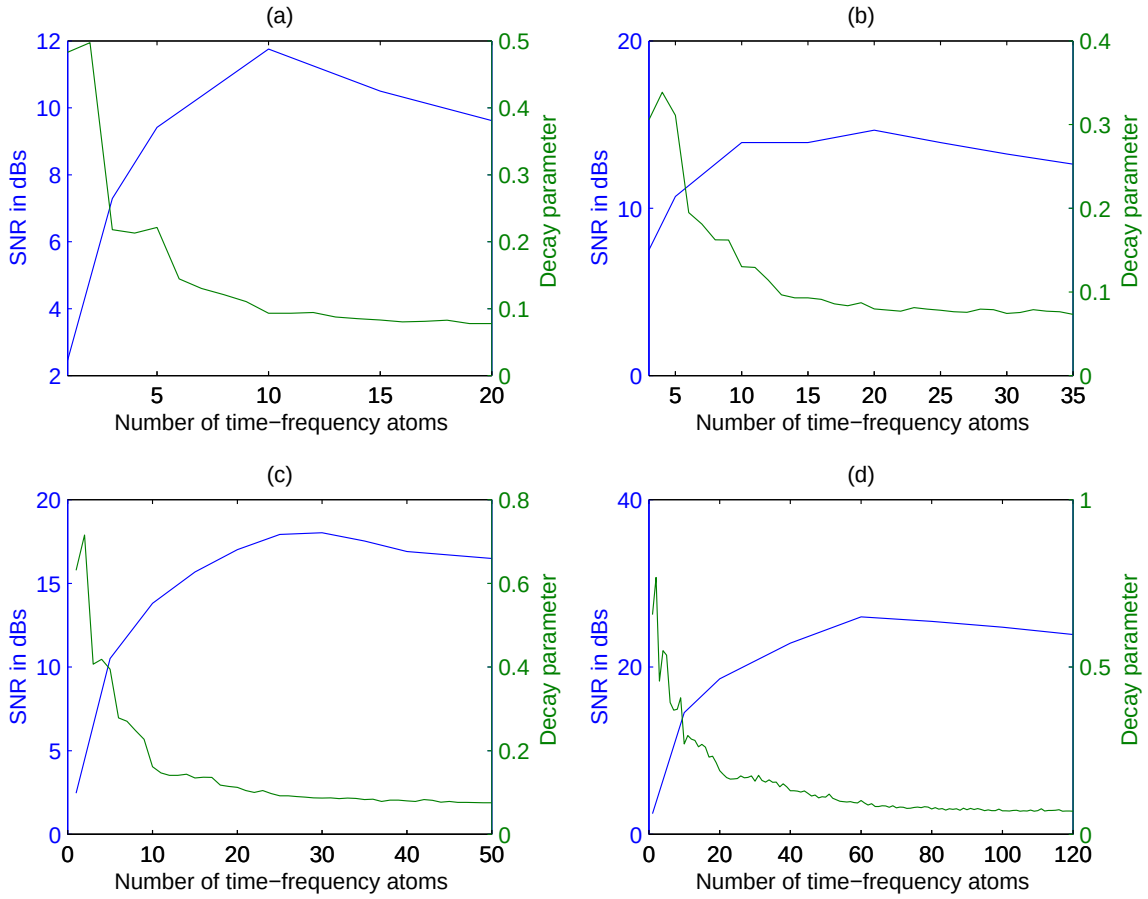

Figure 4.36 This figure shows (a) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to a heart sound cycle (from patient 15) with 1 dB additive white noise, (b) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to the same recording with 5 dBs additive white noise, (c) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to the same recording with 10 dBs additive white noise and (d) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to the same recording with 20 dBs additive white noise. The decay parameter is plotted in green and the SNR in dBs is plotted in blue. It may be observed that as the decay parameter levels out, the best de-noising results are achieved.

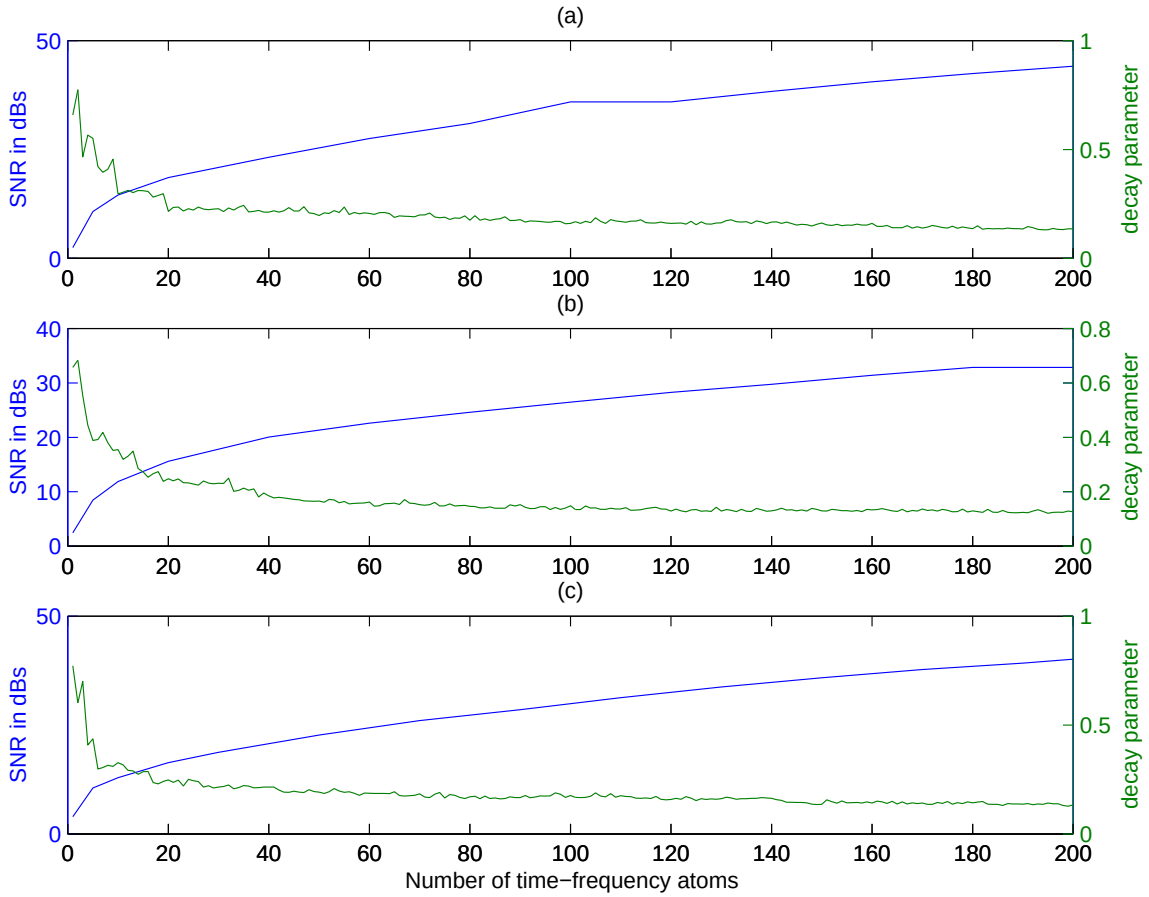

Figure 4.37 This figure shows (a) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to a heart sound cycle (from patient 15) (b) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising is applied to a heart sound cycle (from patient 10), and (c) the decay parameter and the SNR plotted against the number of time-frequency atoms after MP de-noising a heart sound cycle (from patient 12). The decay parameter is plotted in green and the SNR in dBs is plotted in blue.

In order to see how much information is lost by the MP de-noising process, the reader may refer to Figure 4.37. This figure shows the decay parameter and the SNR after applying the MP de-noising process to clean PCGs plotted against the number of time-frequency atoms. It may be seen that a fair amount of information is lost as the SNR between the original and "de-noised" PCGs is only approximately 40 dBs at best and decreases even further as more time-frequency atoms are used.

Overall, the MP method de-noises fairly well, but attention must be given to the number of time-frequency atoms used which may be determined by observing the decay parameter. Also, the amount of information lost by the de-noising process itself is of concern.

4.9 Results and Discussion

Patient	**Wavelet**	**Optimised Wavelet**	**Wavelet Packet**	**Matching Pursuit**
Patient 15	12.60	13.56	12.64	11.74
Patient 10	11.13	12.25	10.81	12.11
Patient 12	12.21	13.69	8.45	12.09

Table 4.5 Comparison of typical results for various de-noising methods for three different PCGs which had 1 dB of additive white noise. The results are given as the SNR in dBs of the original clean PCG and the de-noised version.

Patient	**Wavelet**	**Optimised Wavelet**	**Wavelet Packet**	**Matching Pursuit**
Patient 15	15.99	16.67	15.99	14.71
Patient 10	14.54	15.47	13.77	13.03
Patient 12	15.45	16.57	15.75	14.01

Table 4.6 Comparison of typical results for various de-noising methods for three different PCGs which had 5 dBs of additive white noise. The results are given as the SNR in dBs of the original clean PCG and the de-noised version.

The purpose of this section is to comparatively discuss the different de-noising methods which have been presented and point out any conclusions that may be drawn from this research. The purpose of Figure 4.38, which gives an example of a heart sound cycle which has had noise added to it and then the cycle after optimised wavelet, wavelet, WP, and MP de-noising have been applied, is to visually demonstrate how well the various de-noising

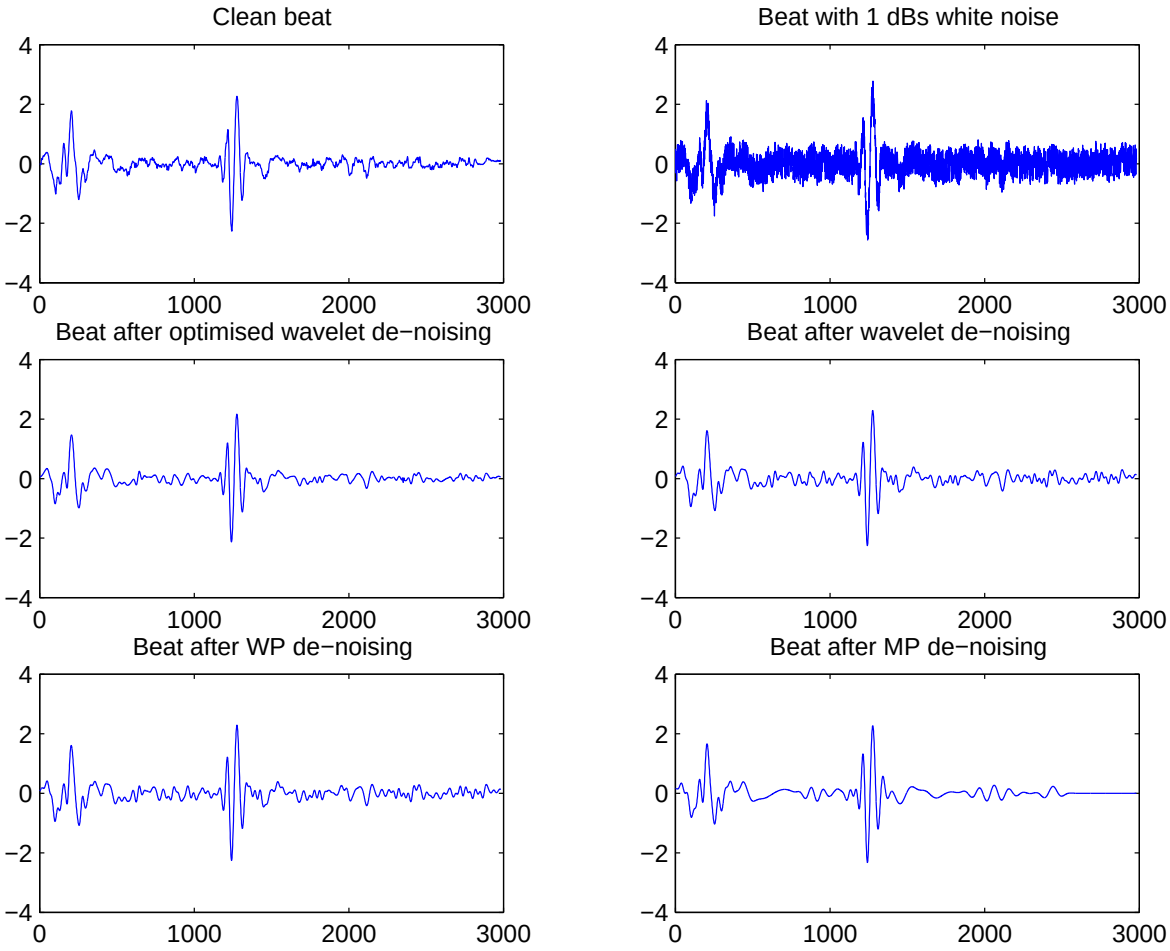

Figure 4.38 Here is an example of the various de-noising methods all applied to the same noisy heart sound recording. A clean heart sound cycle is shown, then the heart cycle with a SNR of 1 dB, then the cycle after optimised wavelet de-noising has been applied to the noisy cycle, then the cycle after wavelet de-noising, next the cycle after WP de-noising, and finally the cycle after MP de-noising. Note the numbers shown the on the x-axis are the number of samples and the y-axis represents the magnitude of the PCG.

Patient	**Wavelet**	**Optimised Wavelet**	**Wavelet Packet**	**Matching Pursuit**
Patient 15	19.58	20.73	19.51	17.70
Patient 10	18.01	18.70	18.02	17.09
Patient 12	19.18	20.20	19.18	17.50

Table 4.7 Comparison of typical results for various de-noising methods for three different PCGs which had 10 dBs of additive white noise. The results are given as the SNR in dBs of the original clean PCG and the de-noised version.

Patient	Wavelet	Optimised Wavelet	Wavelet Packet	Matching Pursuit
Patient 15	27.23	27.95	27.00	25.87
Patient 10	25.77	26.24	25.76	23.51
Patient 12	27.69	28.34	27.27	10.52

Table 4.8 Comparison of typical results for various de-noising methods for three different PCGs which had 20 dBs of additive white noise. The results are given as the SNR in dBs of the original clean PCG and the de-noised version.

methods perform. It is interesting to note that the MP de-noising method does not pick up some of the finer details from the original PCG. The low energy coherent structures have been destroyed by the high level of noise and are unable to be extracted by the MP method. Tables 4.5, 4.6, 4.7, and 4.8 show typical results for optimised wavelet de-noising, wavelet de-noising, wavelet packet de-noising and matching pursuit de-noising applied to different PCGs with various amounts of white noise added. Figure 4.39 summarises the results of the previously mentioned tables into bar charts. Examining these tables and the figure, it becomes apparent that the optimised wavelet de-noising seems to perform slightly better than the other methods for removing white noise from PCGs. The other methods appeared to perform about equally as well as each other.

A study comparing wavelet, WP and MP de-noising applied to knee-joint vibrations, which are complex, non-stationary signals, concluded that the MP method outperformed wavelet and WP de-noising with the WP de-noising performing better than the wavelet de-noising method (Krishnan & Rangayyan 2000). For a large amount of white noise (SNR=0 dB) added to the signal, the MP method only outperformed the others by a fairly small amount, but with smaller amounts of noise (SNR=10 dBs) added to the signal, the MP method outperformed the WP and wavelet de-noising methods by a much more significant amount. This would suggest that large amounts of noise destroy low-energy coherent structures within that signal making it impossible for them to be extracted from the noisy signal by the MP method. Given the results presented in Krishnan & Rangayyan (2000) and the similarity in the types of signals, both being complex, non-stationary signals, it was anticipated that the MP method would perform the best followed by the WP method. The MP method is a greedy algorithm which decomposes the given signal using a whole dictionary of functions. It was expected that WP de-noising would

perform better than wavelet de-noising because WP analysis adaptively chooses the best basis based upon an entropy search and has more decomposition structure combinations to choose from. De-noising results with wavelet and WP depend greatly on the selection of the threshold value for the coefficients (Krishnan & Rangayyan 2000). This fact may be used to partly explain why optimised wavelet de-noising outperformed the other methods because it has various options for the methods used to perform thresholding and noise modelling. Another reason optimised wavelet de-noising may have had a higher SNR after de-noising compared to the other methods is that it appears to lose less information by the de-noising process alone. This fact may be confirmed by examining Figures 4.8, 4.9, 4.17, 4.21, 4.23, 4.27, and 4.37.

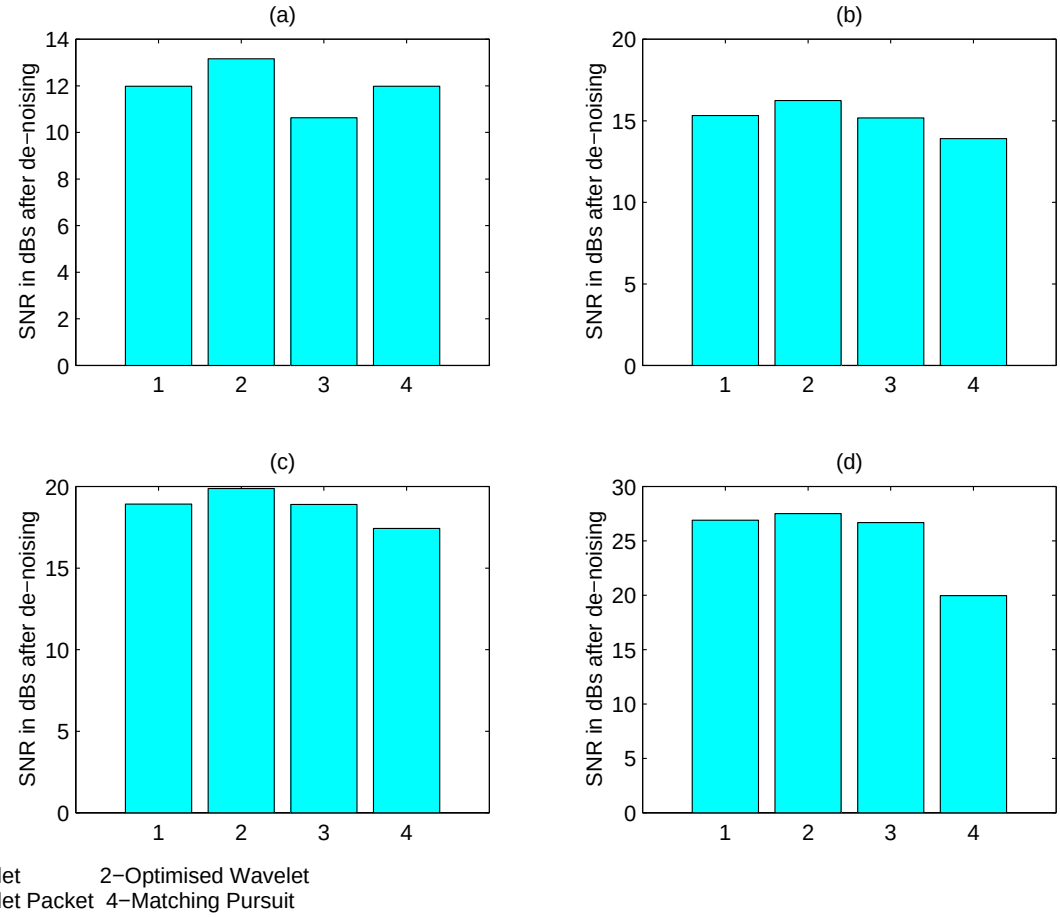

Figure 4.39 (a) The SNR after applying de-noising techniques to PCGs with 1 dB additive white noise, (b) the SNR after applying de-noising techniques to PCGs with 5 dBs additive white noise, (c) the SNR after applying de-noising techniques to PCGs with 10 dBs additive white noise, and (d) the SNR after applying de-noising techniques to PCGs with 20 dBs additive white noise where 1 represents wavelet de-noising, 2 is optimised wavelet de-noising, 3 represents wavelet packet de-noising and 4 is matching pursuit de-noising.

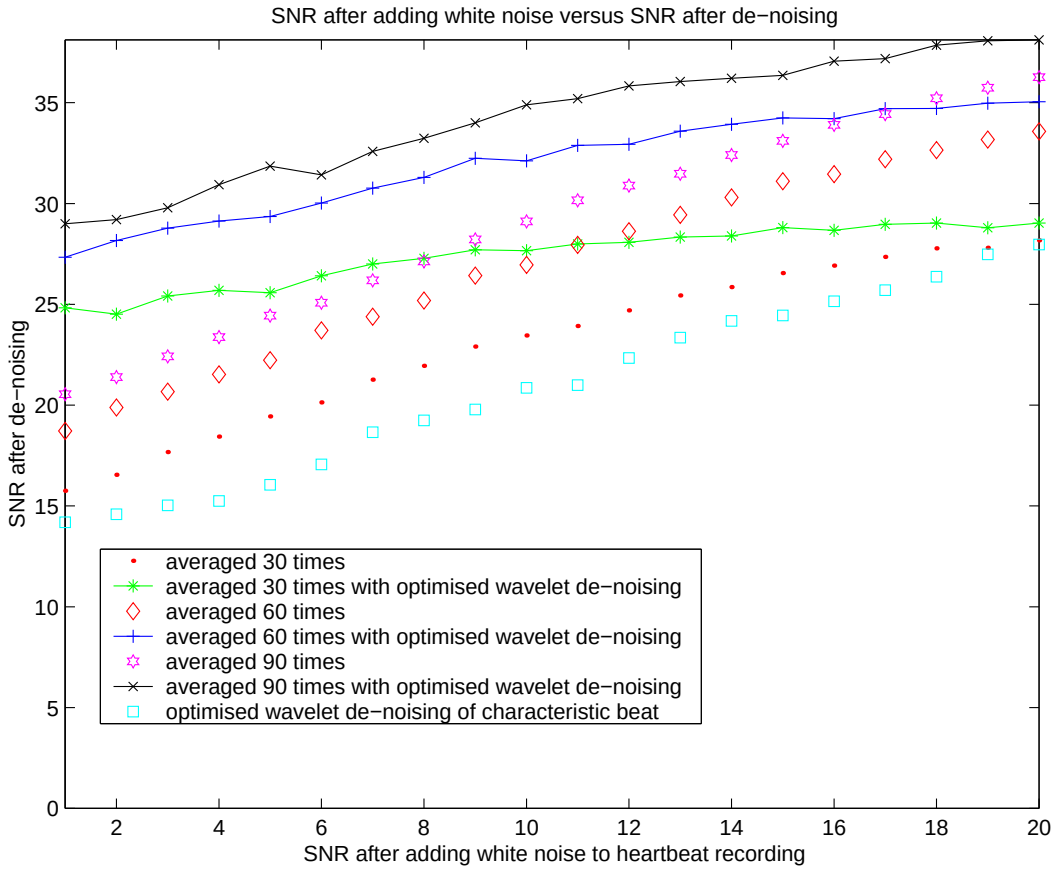

Figure 4.40 This graph shows the SNR after adding white noise to a PCG with a number of heart sound cycles versus the SNR after de-noising the signal. Various methods are tried: optimised wavelet de-noising only, averaging only, and optimised wavelet de-noising combined with averaging. The optimised wavelet de-noising combined with averaging was the most successful de-noising method. PCG signals of 30, 60 and 90 cycle were averaged. It appears that averaging at least 60 cycles would be recommended. There is a significant improvement of averaging 60 cycles over 30 cycles while averaging 90 cycles over 60 cycles does not show nearly the same improvement.

There was no evidence for optimised wavelet de-noising that any one wavelet was much better than another wavelet. However, some of the lower order wavelets did not perform very well. We reached the conclusion that a decomposition level of 5 produced reasonable results while decomposing the signal further often produced marginal benefits and increases computation time. Soft thresholding definitively outperformed hard thresholding. Of the four threshold selection rules, the "rigrsure" rule performed the best, and the best choice of the threshold rescaling methods proved to be the "sln" method.

For both wavelet and WP de-noising, most wavelets performed roughly equally except for some of the lower order wavelets for reasons previously explained. Decomposition levels

of 3-5 were found to perform the best. The WP de-noising process seems to lose more of the original signal content than the wavelet de-noising process with the wavelet de-noising process losing even more original signal content than optimised wavelet de-noising.

Overall, the MP method de-noises the PCG about as well as WT or WP de-noising, but attention must be given to the number of time-frequency atoms used which may be determined by observing the decay parameter.

Averaging seemed to produce significant improvements especially if there is a large amount of noise present in the signal. Averaging a series of 50-75 heart sound cycles seems to give the best result in terms of recording and computation time tradeoff.

Averaging may be used in combination with one of the other de-noising methods. Because optimised wavelet de-noising performed the best, this method was chosen to be used. Figure 4.40 shows a comparison of using wavelet de-noising only, averaging only and wavelet de-noising combined with averaging. It clearly demonstrates that combining the techniques is much more effective. It shows that given a choice between averaging 30, 60 or 90 cycles that 60 cycles provides a good compromise in terms of de-noising and recording and computation time. However, as previously discussed in Section 4.7 in some cases averaging is not appropriate.

4.10 Chapter Summary

This chapter has presented the results of the de-noising study of PCGs and made recommendations based upon these results. Now that the PCG is relatively free of noise, the signal can be analysed for data that may differentiate between healthy and unhealthy hearts. The next chapter explores various methods of doing just this.

Chapter 5

PCG Data Analysis

"Data! Data! Data!" he cried impatiently. "I can't make bricks without clay."

Sherlock Holmes

CHARACTER IN ARTHUR CONAN DOYLE'S THE COMPLETE SHERLOCK HOLMES, THE ADVENTURE OF THE COPPER BREECHES

5.1 Introduction

After the PCG signal is de-noised, the signal should be analysed for hidden information that may differentiate a subject with a normal heart from a patient with a pathological condition. There are many techniques for data analysis, but we have chosen to explore phase synchronisation between the ECG and PCG, the phase space diagram and the use of the Hilbert Transform (HT) to construct the HT diagram and to extract instantaneous signal features. These techniques were chosen because they produce easy to read graphs and diagrams that reveal information hidden in the PCG that may not be seen by a simple visual inspection of the signal itself. No statistical analysis techniques were chosen because statistical techniques normally require very long recordings, and PCG recordings are typically not very long due to the difficulties presented by the recording techniques used to obtain a clean recording. Also, our exploration of the use of these data analysis techniques was limited by a smaller than expected number of PCG recordings obtained from actual patients at the Hampstead Medical Clinic by Dr. J. Agzarian. We, therefore, present initial indicative results only. In the final chapter we will suggest directions for an in-depth study of this area.

5.2 Phase Space and Hilbert Transform Diagrams

5.2.1 Phase Space Diagrams

The phase space diagram is a plot of the signal itself versus the rate of change of a signal and is readily explained in Letellier, Meunier-Guttin-Cluzel, Gouesbet, Neveu, Duverger & Cousyn (1997). The phase space diagram has been applied to various cardiovascular signals in Dutt & Krishnan (1999), Hall (1999), Maple et al. (1999), and Maple (1999). Figure 5.1 aids in the explanation of the phase space diagram. The heart sound is basically a complex signal which is a sum of discrete time sinusoids and random impulses representing the noise (Hall 1999). One of the most fundamental signals is the sinusoid which is shown in Figure 5.1 (a), and the phase space diagram of the sinusoid in Figure 5.1 (a) is shown in Figure 5.1 (b). It is a oval plot. Now, we explore what an increase in

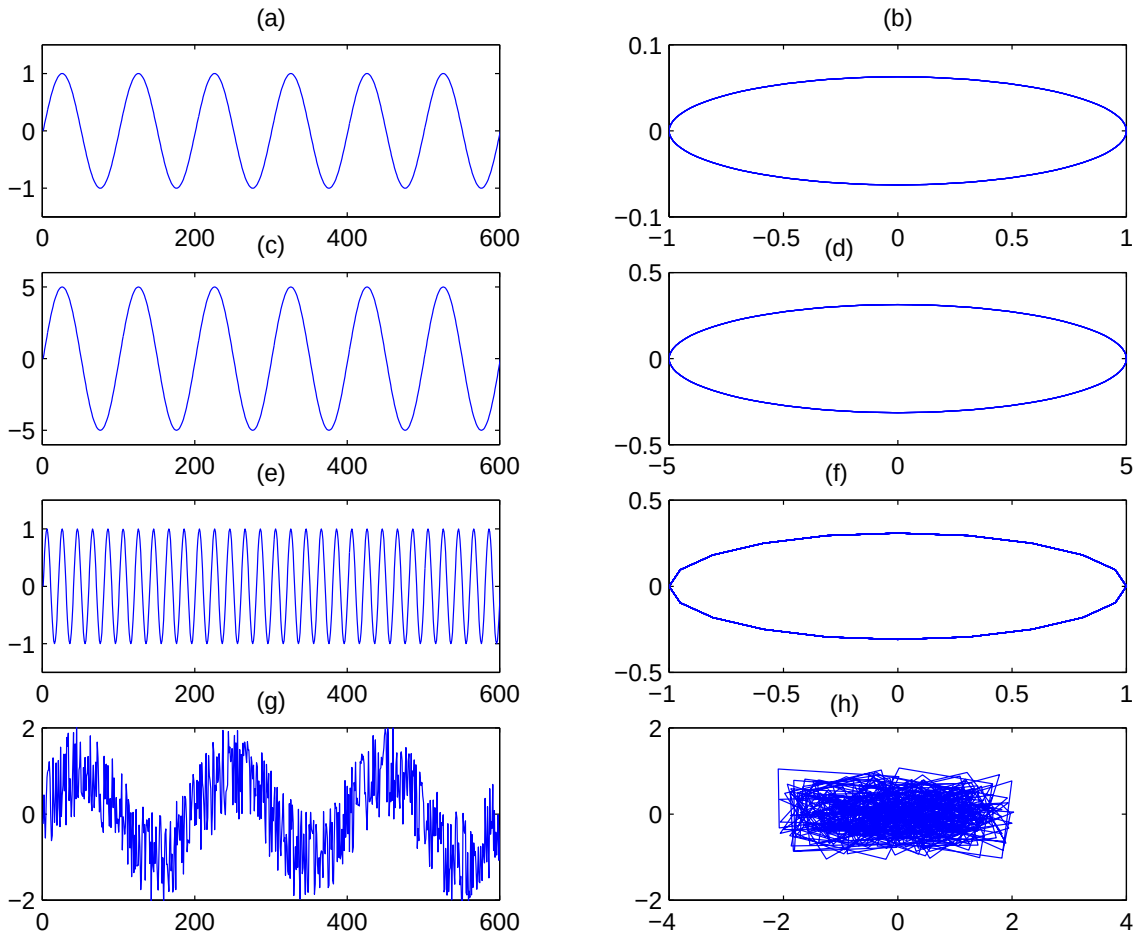

Figure 5.1 This figure is an aid for explaining phase space diagrams. (a) Is a sinusoid at 5 Hz, (b) the phase space diagram of (a), (c) is a sinusoid at 5 Hz with a magnitude 5 times larger than the sinusoid in (a), (d) the phase space diagram of (c), (e) is a sinusoid at 50 Hz, (f) is a phase space diagram of (e), (g) is the same sinusoid as in (a) but with 1 dB of additive white noise, and (h) is the phase space diagram of (g). The left hand column shows the signal amplitude versus time. The right hand column contains the corresponding phase space plots which are obtained by plotting the rate of change of the signal against the original signal. The units are arbitrary in all cases.

the magnitude of the sinusoid does to the phase space diagram by observing Figure 5.1 (c) and (d). The larger the magnitude of the sinusoid, the larger the oval which is created in the phase space diagram. Next, let us examine the effect of increasing the frequency of the sinusoid by referring to Figure 5.1 (e) and (f). It is obvious that as the frequency increases the the circle becomes larger in the rate of change direction. This is because higher frequency signals have a larger rate of change. Finally, the effect of noise on the sinusoid is shown in Figure 5.1 (g) and (h) where it is evident that the randomness of the noise produces irregular patterns in the phase space diagram, and the noise produces large deflections in the rate of change direction because the noise contains components that are of much higher frequency than the sinusoid. Thus, for PCGs, the phase space diagram should be composed of large circular shapes for the low frequency, high amplitude sinusoidal components of the major heart sounds where the low frequency components have a small derivative producing large horizontal movements with small vertical deflections. High frequency, low amplitude parts of the signal will be represented by small squiggles in the rate of change direction. The signal component caused by noise will produce large values in the rate of change direction for even small magnitudes of noise. If the PCG is quite noisy, it will be easily revealed by the phase space diagram. Also, if heart sound cycles are quite regular, they should follow fairly similar paths on the phase space diagram. The phase space diagram can, therefore, be used as a tool to reveal how well de-noising techniques perform as shown in Maple et al. (1999).

5.2.2 Hilbert Transform Diagram

The Hilbert Transform (HT) may be used to calculate the instantaneous attributes of a signal and to display the HT diagram. The mathematical definition of the Hilbert Transform is (Ersoy 1997)

$$y(t) = \pi^{-1} \int_{\infty}^{-\infty} \frac{x(\tau)}{t-\tau} d\tau. \tag{5.1}$$

The Hilbert Transform (HT) can be considered a convolution between the signal and $\frac{1}{\pi t}$. The HT can be realized by an ideal filter whose amplitude response is unity and whose phase response is a constant ninety degree lag. The HT is called the quadrature filter

because it shifts the phase of the spectral components by $\frac{\pi}{2}t$. More information about the theory of the HT and its applications may be found in a number of sources including Braun & Feldman (1997), Bogner (1999), Bolton (1983), Claerbout (1976), Dishan (1996), Ersoy (1997), Feldman (1994), *Instantaneous Frequency–A Seismic Attribute Useful in Structural and Stratigraphic Interpretation* (2000), Gao, Dong, Want, Li & Pan (1999), Jiménez, Charleston, Peña, Aljama & Ortiz (1999), Oppenheim, Schafer & Buck (1998), Panter (1965), Randall (1987), Taner, Koehler & Sheriff (1979), and Yang et al. (1994). The HT is calculated using MATLAB software with the method being further explained in MATLAB help files and documentation. Basically, the analytic signal of a signal has a one-sided FT meaning that negative frequencies are zero. To approximate the analytic signal, the FFT of the input sequence is calculated, the FFT coefficients that correspond to negative frequencies are replaced with zeros, and finally the inverse FFT of the result is calculated.

The Hilbert Transform diagram is simply the signal plotted against its Hilbert Transform. Figure 5.2 is an aid to explain the HT diagram. A sinusoidal wave results in a circular pattern in the HT diagram as may be seen in Figure 5.2 (a) and (b). From Figure 5.2 (c) and (d), we can see that increasing the magnitude of the sinusoid simply increases the size of the circular pattern created in the HT diagram. Increasing frequency without changing the magnitude of the sinusoid does not change the circular pattern as shown in Figure 5.2 (e) and (f). Random white noise destroys the neat pattern followed by a clean sinusoid as seen in Figure 5.2 (g) and (h).

5.2.3 Comparison of Phase Space and Hilbert Transform Diagrams

Because the phase space diagram is a plot of the derivative of a signal against itself, it picks up high frequency content or rapidly changing details of a signal and the low frequency or slowly changing details tend to be diminished. In other words, when a derivative of a signal is taken, the derivative will spike when there is a rapid change in the signal. On the other hand, the Hilbert Transform treats all frequencies equally.

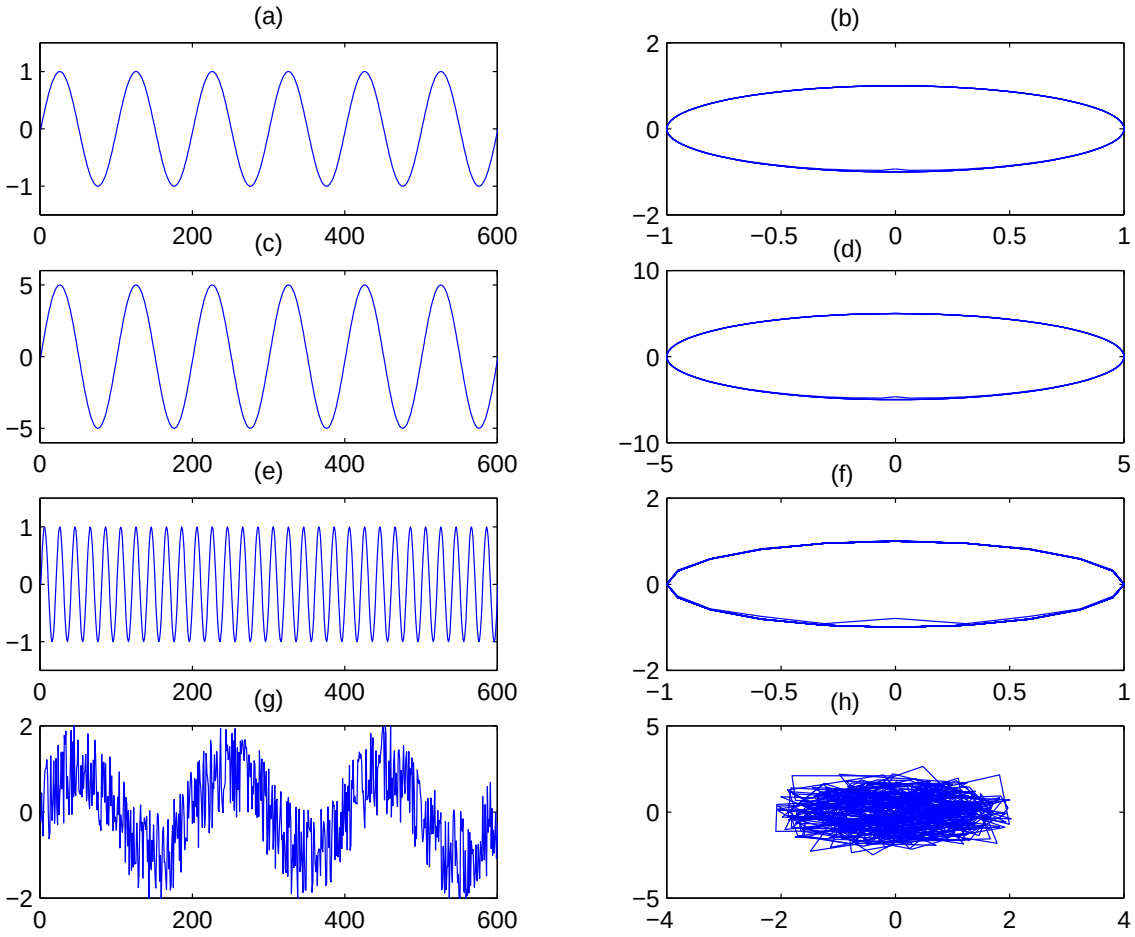

Figure 5.2 This figure is an aid for explaining Hilbert Transform diagrams. (a) Is a sinusoid at 5 Hz, (b) the HT diagram of (a), (c) is a sinusoid at 5 Hz with a magnitude 5 times larger than the sinusoid in (a), (d) the HT diagram of (c), (e) is a sinusoid at 50 Hz, (f) is the HT diagram of (e), (g) is the same sinusoid as in (a) but with 1 dB of additive white noise, and (h) is the HT diagram of (g). The left hand column shows the signal amplitude versus time. The right hand column contains the corresponding Hilbert Transform plots which are obtained by plotting the Hilbert Transform of the signal against the original signal. The units are arbitrary in all cases

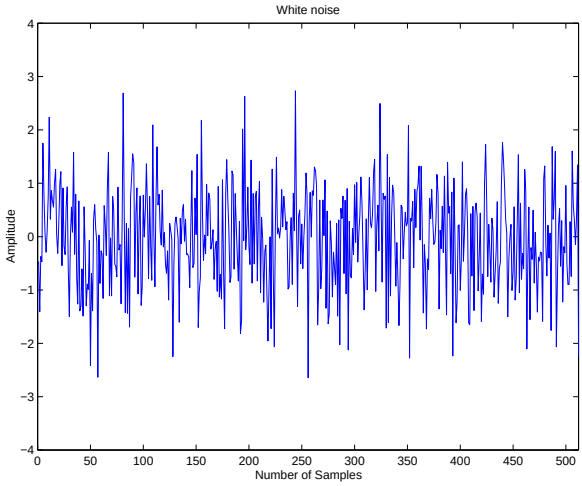

Figure 5.3 White Noise

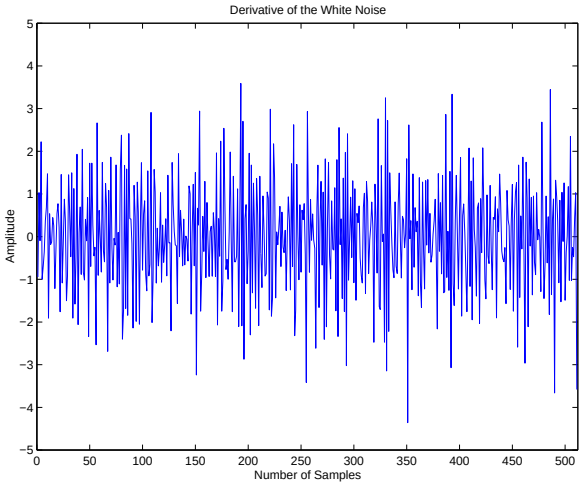

Figure 5.4 Derivative of the White Noise

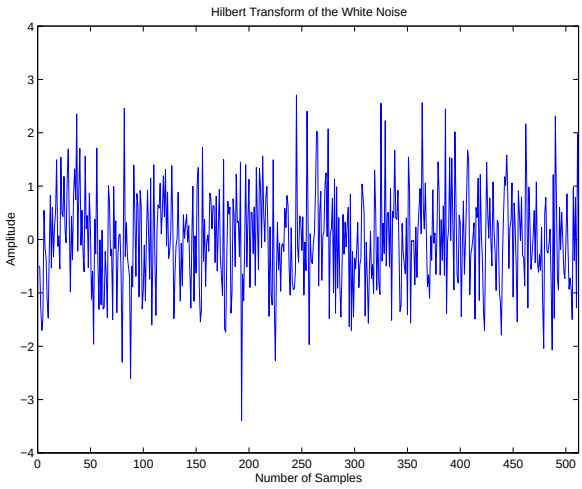

Figure 5.5 Hilbert Transform diagram of the White Noise

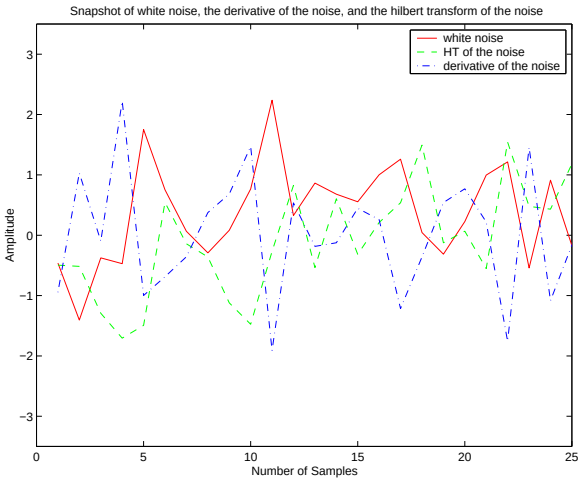

Figure 5.6 Snapshot of the White Noise, Derivative of the Noise, and Hilbert Transform of the Noise

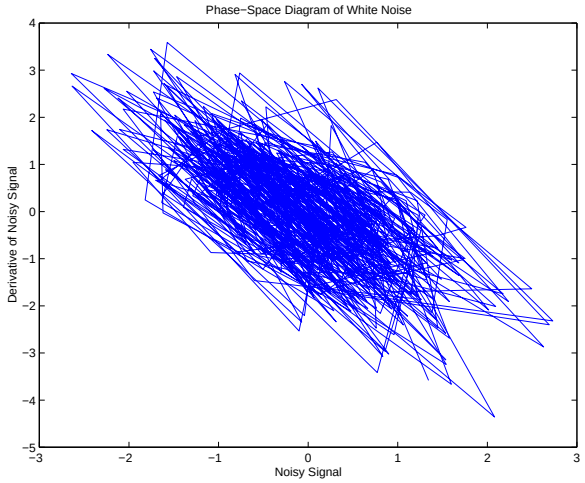

Figure 5.7 Phase Space Diagram of the White Noise

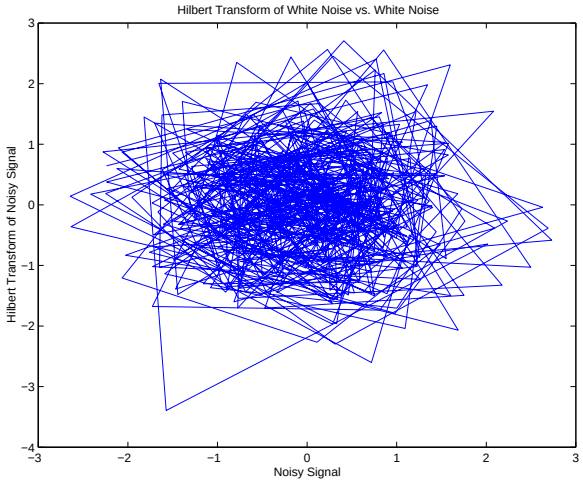

Figure 5.8 Hilbert Transform of the White Noise

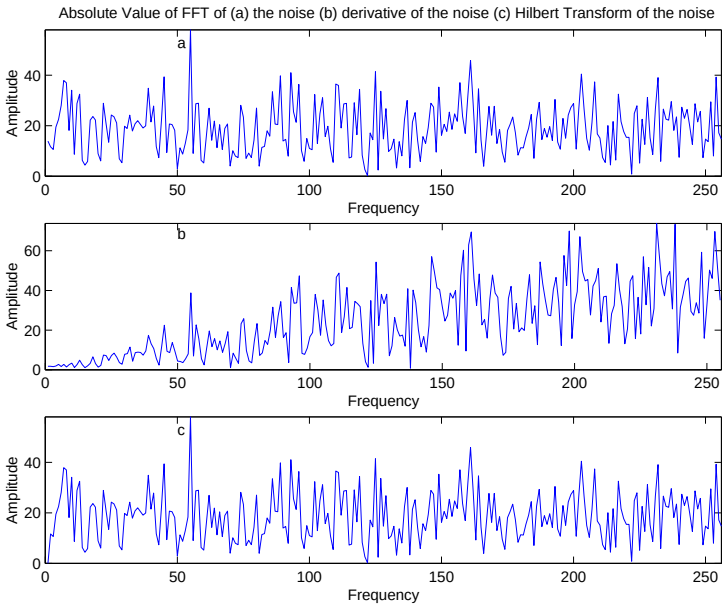

Figure 5.9 (a) FFT of the white noise, (b) FFT of the derivative of the white noise and (c) FFT of the Hilbert Transform of the noise

This concept is further explained by using an example where the HT and phase-space diagrams are applied to white noise. Figures 5.3, 5.4, and 5.5 show the white noise, the derivative of the noise and the Hilbert Transform of the white noise respectively. Figures 5.3, 5.4, and 5.5 do not appear very different. However, upon closer examination of a small portion of the signals in Figure 5.6, it may be seen that the signals are actually quite different. The derivative of the noise detects rapidly changing details of the noise. The Hilbert Transform of the noise has a $\pi/2$ phase lag behind the noise signal. Figure 5.8 that shows the noise plotted against its Hilbert Transform, is more evenly distributed (almost in circular fashion) than Figure 5.7 which shows the noise plotted against its derivative. This confirms that the Hilbert Transform treats all frequencies equally while the phase-space diagram accentuates higher frequencies thus giving the phase space diagram an oval planar distribution. Figure 5.9 illustrates the fact that the Hilbert Transform treats all frequencies equally while taking the derivative of a signal accentuates higher frequencies and virtually ignores lower frequencies.

It was hoped that the phase space and HT diagrams would reveal information that was not apparent by viewing the PCG because the HT has proven to be a useful too in analysing ECGs. For example, Bolton (1983) used the HT for representation and pattern recognition of ECGs. Due to lack of PCG data, we have not been able to test this theory extensively. However, from the limited data available we examined the difference between normal PCGs and PCGs from patients possessing murmurs. Figure 5.11 shows that the phase space diagrams of the normal PCGs and phase space diagrams of PCGs with murmurs are different. Murmurs normally contain much higher frequency content than normal PCGs. Phase space diagrams accentuate higher frequencies, thus we could expect phase space diagrams to emphasise the higher frequencies contained in the murmurs. Examining Figure 5.12, which shows the HT diagrams of the same PCGs, we cannot see a notable difference between the normal and pathological heart sound recordings.

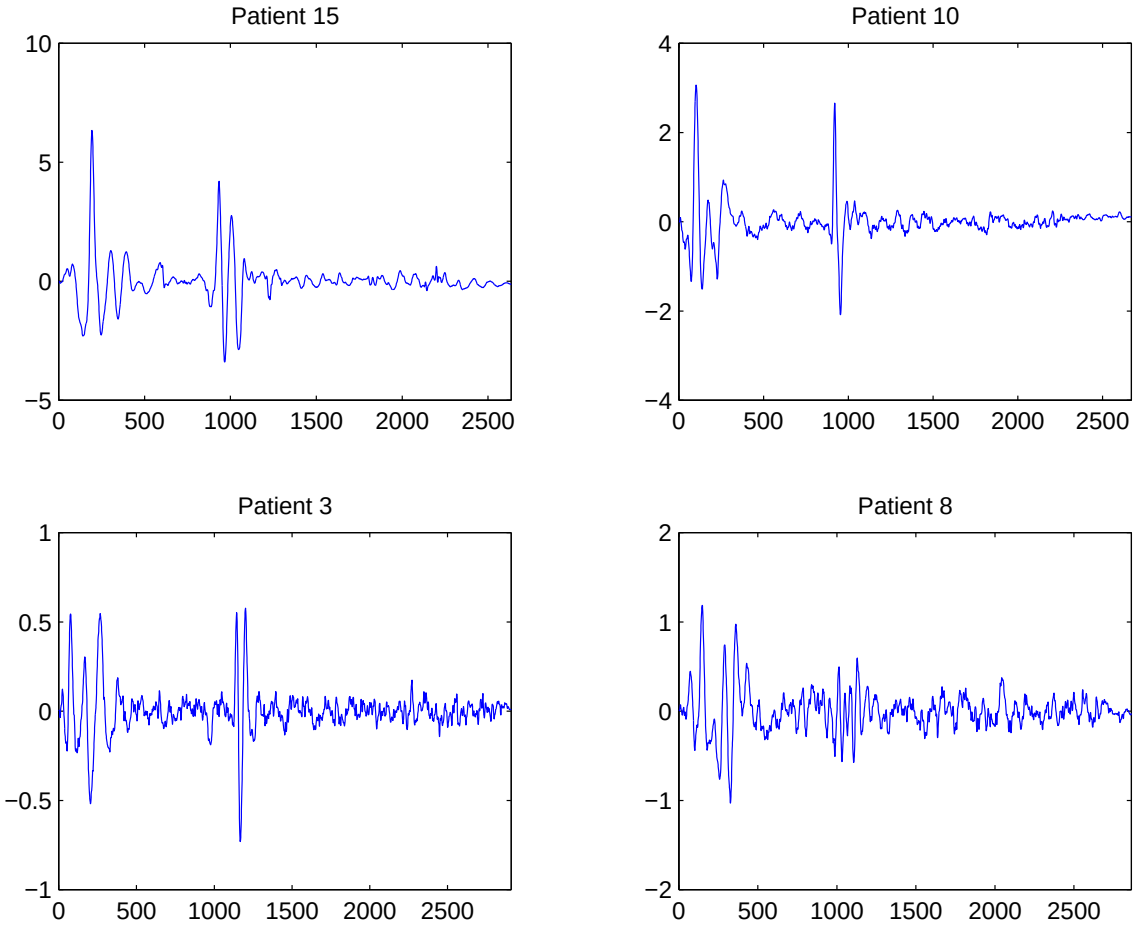

Figure 5.10 The characteristic heartbeat of four patients. Patients 10 and 15 are normal subjects whereas Patients 3 and 8 have heart murmurs.

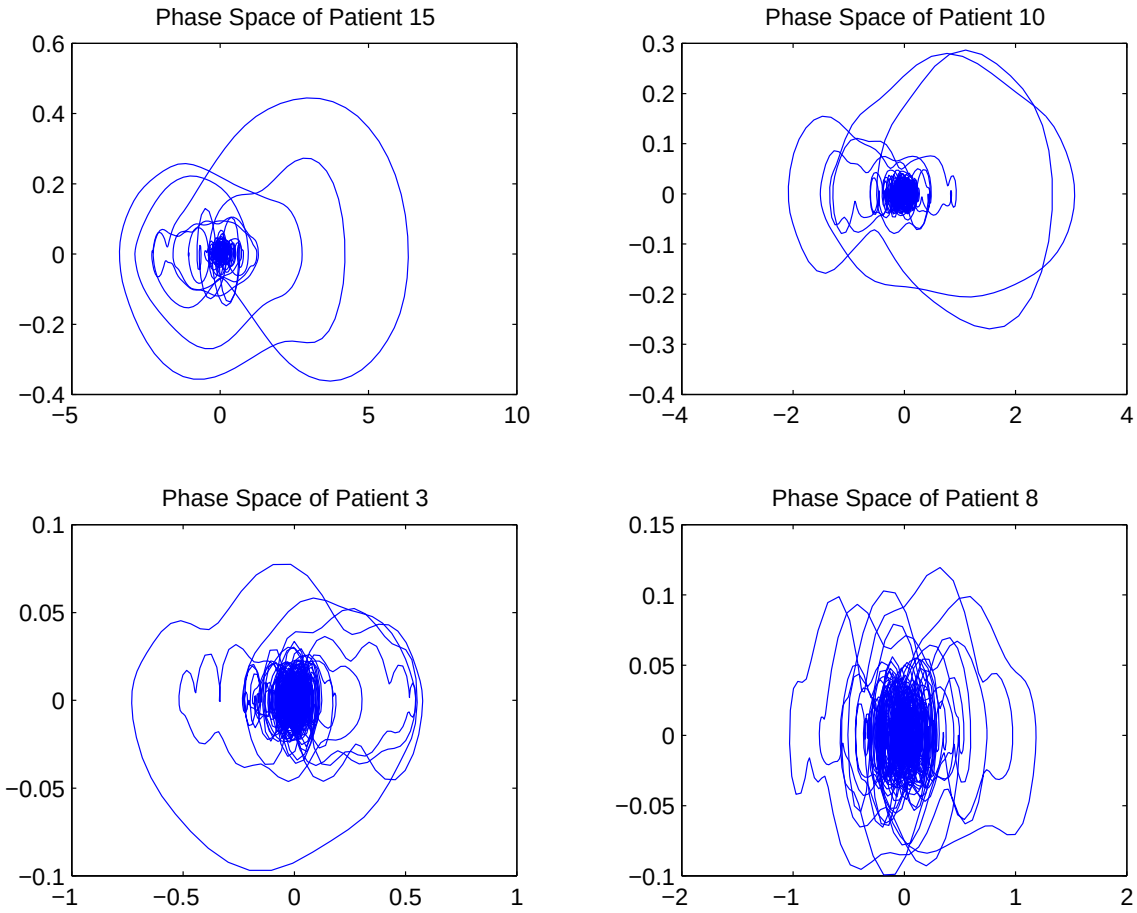

Figure 5.11 The phase space diagrams of four patients where the PCG is plotted against its' derivative. Patients 10 and 15 are normal subjects whereas Patients 3 and 8 have heart murmurs. The characteristic heartbeats of these patients are shown in Figure 5.10.

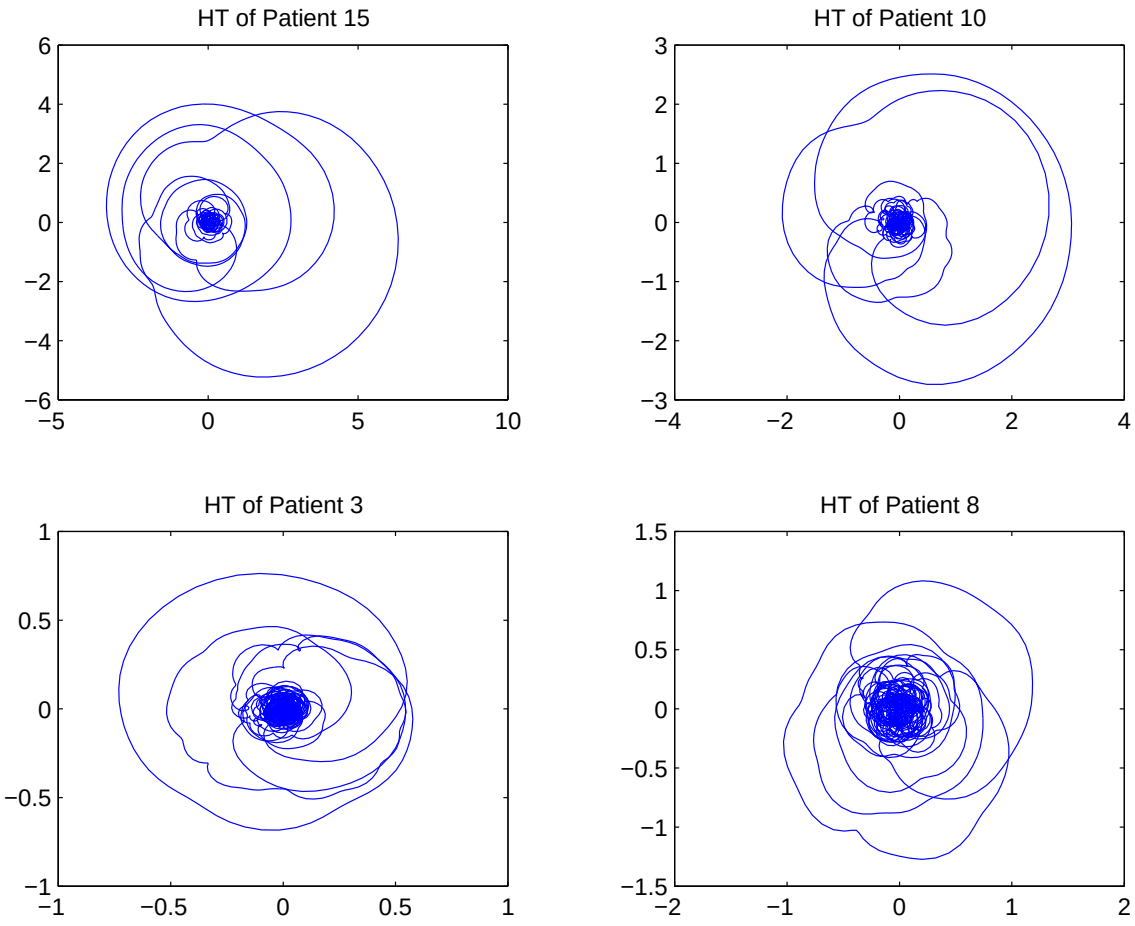

Figure 5.12 The Hilbert Transform diagrams of four patients where the PCG is plotted against its HT. Patients 10 and 15 are normal subjects whereas Patients 3 and 8 have heart murmurs. The characteristic heartbeats of these patients are shown in Figure 5.10.

5.3 Use of the HT to Calculate Instantaneous Signal Parameters of the PCG

The Hilbert Transform may be used to calculate the instantaneous signal parameters including the instantaneous amplitude, phase and frequency of the signal. The instantaneous parameters are defined as follows for s, the real signal, and $H(s)$, the Hilbert Transform of the signal (Bogner 1999, Gao et al. 1999):

$$e(t) = \sqrt{s^2(t) + H^2[s(t)]}, \tag{5.2}$$

$$\Theta(t) = \arctan(\frac{H[s(t)]}{s(t)}), \text{ and} \tag{5.3}$$

$$f(t) = \frac{1}{2\pi}\frac{d}{dt}[\arctan(\frac{H[s(t)]}{s(t)})] = \frac{s\frac{dH[s(t)]}{dt} - H[s(t)]\frac{ds}{dt}}{s^2 + H[s(t)]^2} \tag{5.4}$$

where $e(t)$, $\Theta(t)$, and $f(t)$ are the instantaneous amplitude, phase and frequency of s(t) respectively.

The instantaneous frequency of a signal is the derivative of the instantaneous phase and may be used to demonstrate the effectiveness of the de-noising technique (Carré et al. 1998). The instantaneous amplitude is the magnitude of the complex analytical signal found by using the HT and is sometimes called the envelope.

We may examine how well the PCG de-noising techniques performed by plotting the instantaneous frequency. Figure 5.13 shows a characteristic heartbeat, then the cycle with noise added, and finally the noisy cycle with noise removed both by the wavelet and wavelet packet de-noising processes. The corresponding instantaneous frequencies are also shown. We can clearly see from the instantaneous frequency plots when there are large amounts of noise present. It is also interesting to note that around S1 and S2, the instantaneous frequencies remain relatively constant at low-frequencies supporting the well known fact that S1 and S2 are composed of several low-frequency sinusoidal components.

It was thought that the instantaneous signal parameters might reveal information about the PCG that was not readily extracted by viewing the PCG itself or currently used techniques. We have attempted to begin an investigation into this area, but once again, due to limited data, we could not extensively explore the concept.

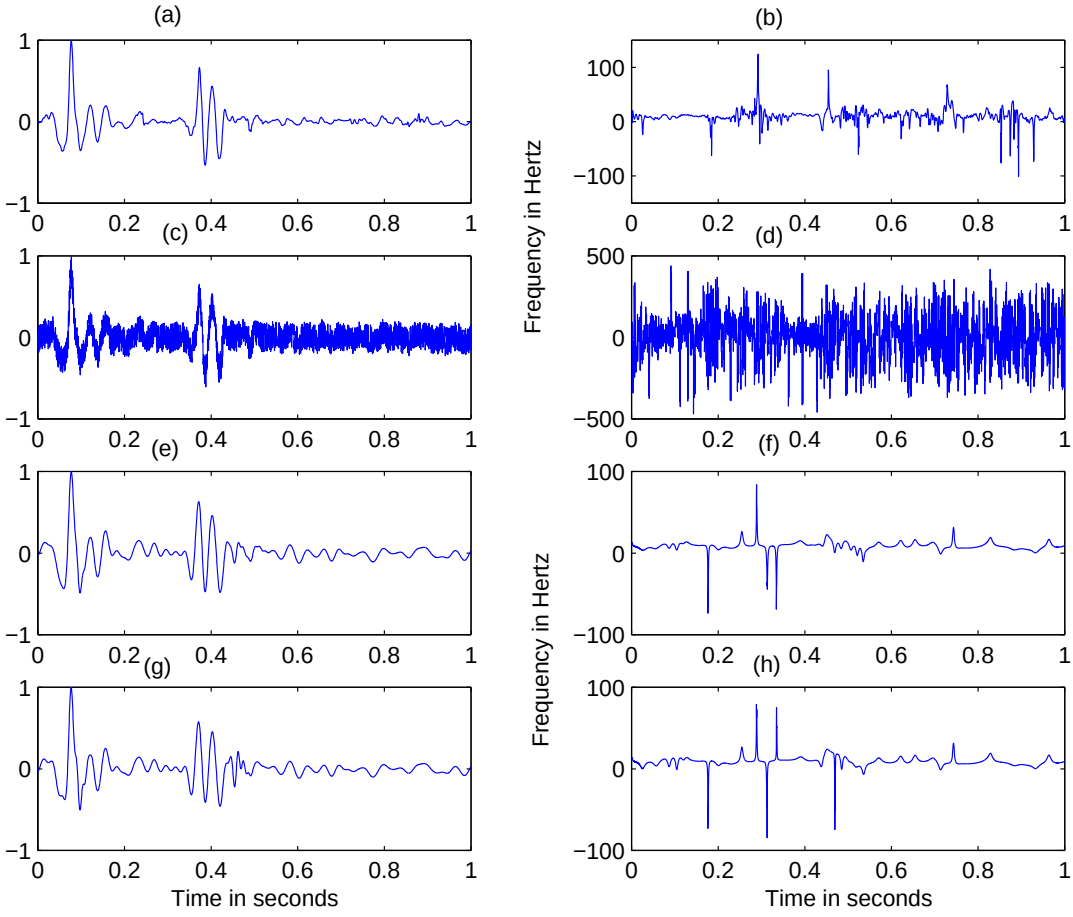

Figure 5.13 (a) Characteristic heartbeat, (b) instantaneous frequency of (a), (c) 1 dB of white noise added to (a), (d) instantaneous frequency of (c), (e) Noisy heartbeat de-noised by wavelet technique, (f) instantaneous frequency for (e), (g) Noisy characteristic heartbeat de-noised by wavelet packet technique, and (h) instantaneous frequency of (g)

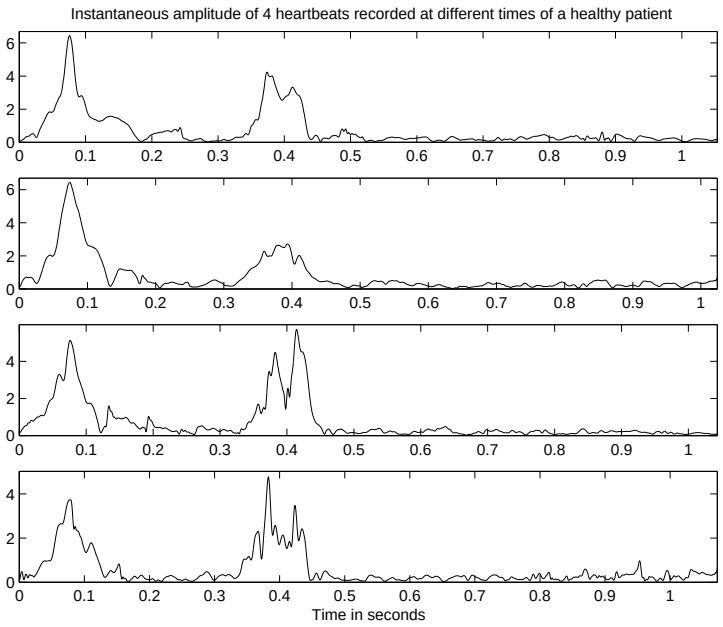

Figure 5.14 The instantaneous amplitude of 4 characteristic heartbeats recorded at different times from the same normal patient. They are all fairly similar demonstrating that this technique is reproducible.

The instantaneous amplitude is an alternative method of looking at the PCG data. Figure 5.14 demonstrates that recording the PCG of a patient is reproducible because plots of the instantaneous amplitude of a PCG of one individual recorded on four different occasions are very similar. Figure 5.15 shows the instantaneous amplitude of PCGs for patients with various pathological conditions and patients with normal hearts. We were limited by the number of PCG recordings available, but by examining this plot we may see that the healthy patients appear to have a well defined and compact S1 and S2 whereas some of the patients with pathological conditions do not.

Figure 5.16 shows the instantaneous frequency of the characteristic heartbeats of various patients with pathological heart conditions and normal hearts. These characteristic heartbeats correspond to the same ones used in Figure 5.15. It was hoped that the instantaneous frequency of the characteristic heartbeats might have obvious differentiating features between the pathological and normal characteristic beats. Examining Figure 5.16, does not reveal this to be the case except that the instantaneous frequency of the normal characteristic heartbeats appear to be "quieter" than the pathological cases which could produce extra signal content especially in the case of heart murmurs. There is a

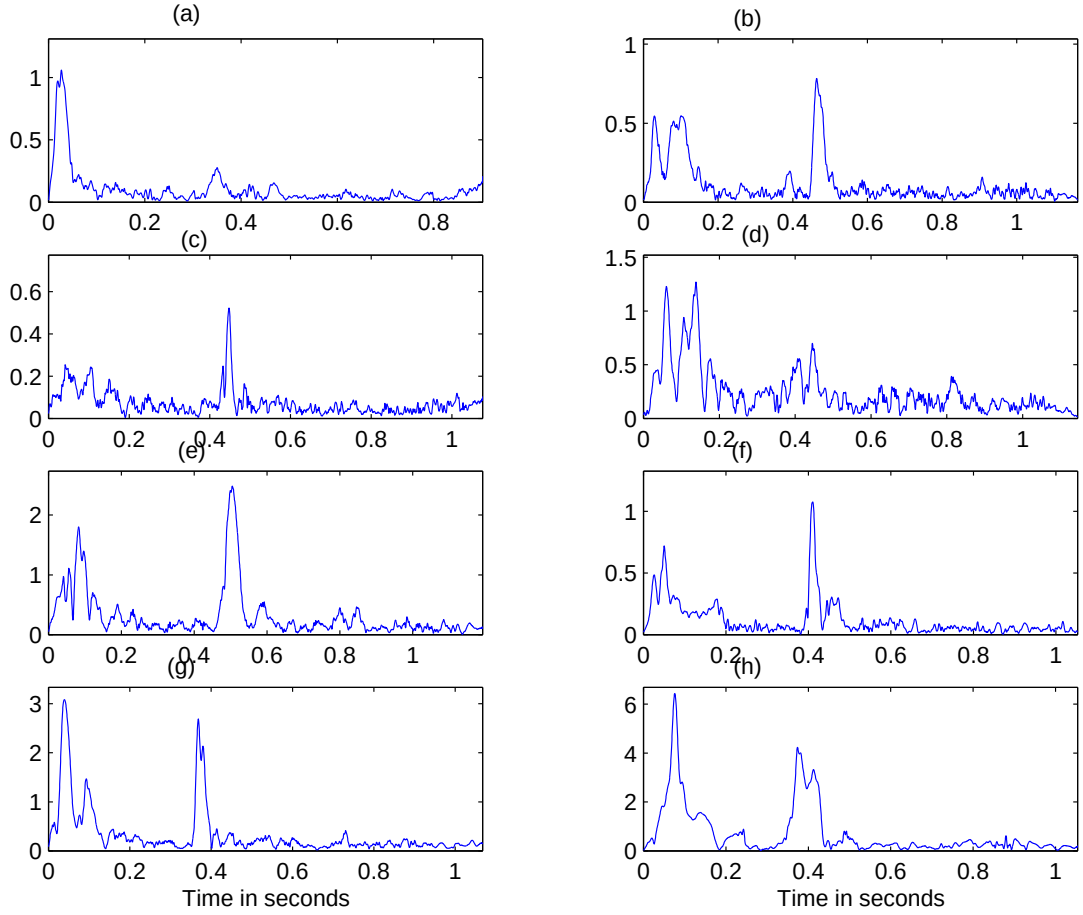

Figure 5.15 (a) Instantaneous amplitude of a characteristic heartbeat of a patient with a mitral valve prosthesis, (b) instantaneous amplitude of a characteristic heartbeat of a patient with a heart murmur and hypertension, (c) instantaneous amplitude of a characteristic heartbeat of a patient with a an aortic stenosis and hypertension, (d) instantaneous amplitude of a characteristic heartbeat of a patient with a systolic murmur and angioplasty, (e) instantaneous amplitude of a characteristic heartbeat of a patient with atrial fibrillation, (f) instantaneous amplitude of a characteristic heartbeat of a patient with hypertension, (g) instantaneous amplitude of a characteristic heartbeat of a normal patient, and (h) instantaneous amplitude of a characteristic heartbeat of a normal patient

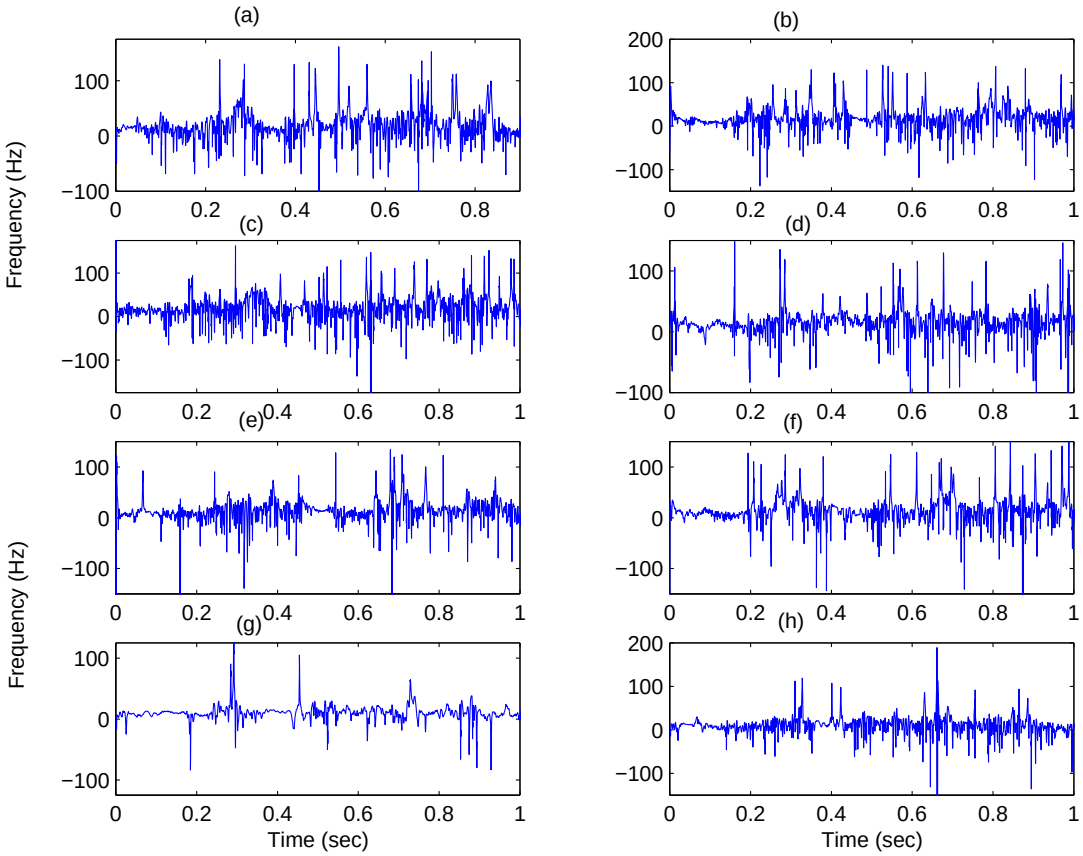

Figure 5.16 (a) Instantaneous frequency of a characteristic heartbeat of a patient with a mitral valve prosthesis, (b) instantaneous frequency of a characteristic heartbeat of a patient with a heart murmur and hypertension, (c) instantaneous frequency of a characteristic heartbeat of a patient with a an aortic stenosis and hypertension, (d) instantaneous frequency of a characteristic heartbeat of a patient with a systolic murmur and angioplasty, (e) instantaneous frequency of a characteristic heartbeat of a patient with atrial fibrillation, (f) instantaneous frequency of a characteristic heartbeat of a patient with hypertension, (g) instantaneous frequency of a characteristic heartbeat of a normal patient, and (h) instantaneous frequency of a characteristic heartbeat of a normal patient. PCGs correspond to those used in Figure 5.15.

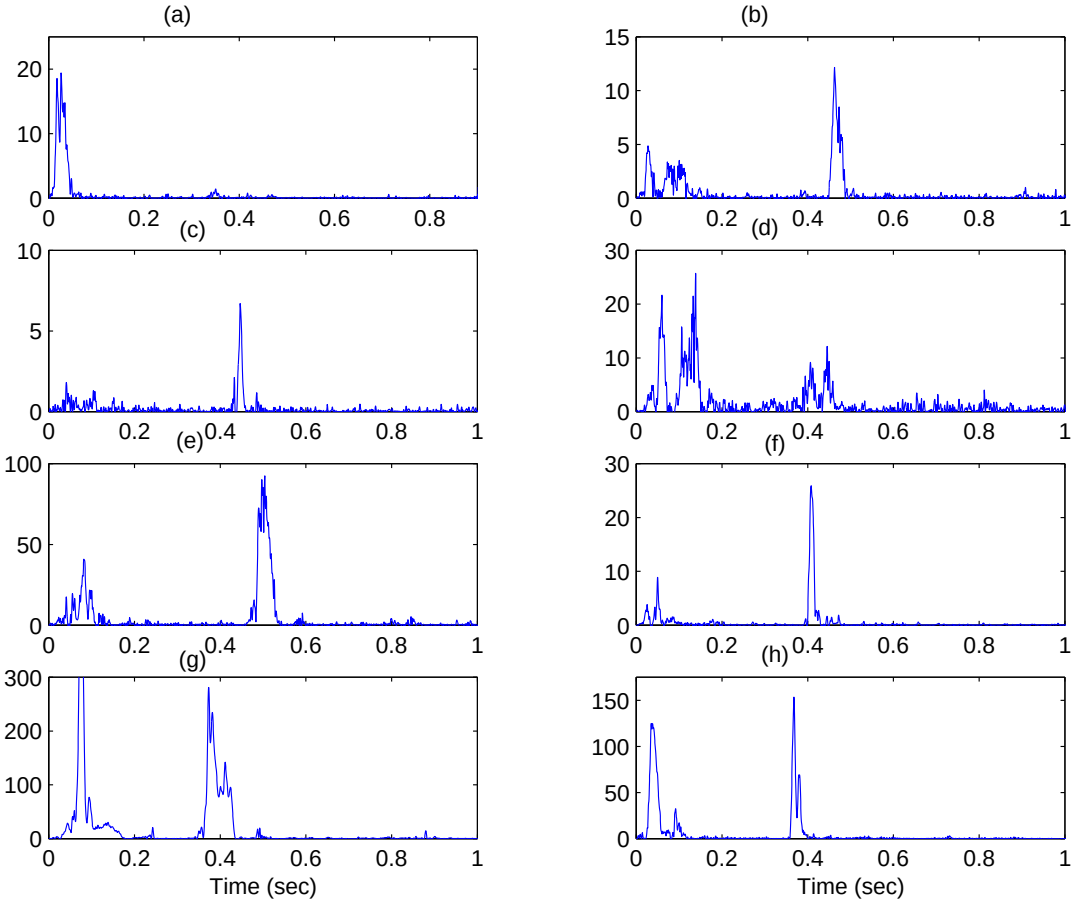

Figure 5.17 (a) Moment of velocity of a characteristic heartbeat of a patient with a mitral valve prosthesis, (b) moment of velocity of a characteristic heartbeat of a patient with a heart murmur and hypertension, (c) moment of velocity of a characteristic heartbeat of a patient with a an aortic stenosis and hypertension, (d) moment of velocity of a characteristic heartbeat of a patient with a systolic murmur and angioplasty, (e) moment of velocity of a characteristic heartbeat of a patient with atrial fibrillation, (f) moment of velocity of a characteristic heartbeat of a patient with hypertension, (g) moment of velocity of a characteristic heartbeat of a normal patient, and (h) moment of velocity of a characteristic heartbeat of a normal patient. Units are in volts squared per seconds. PCGs correspond to those used in Figure 5.15.

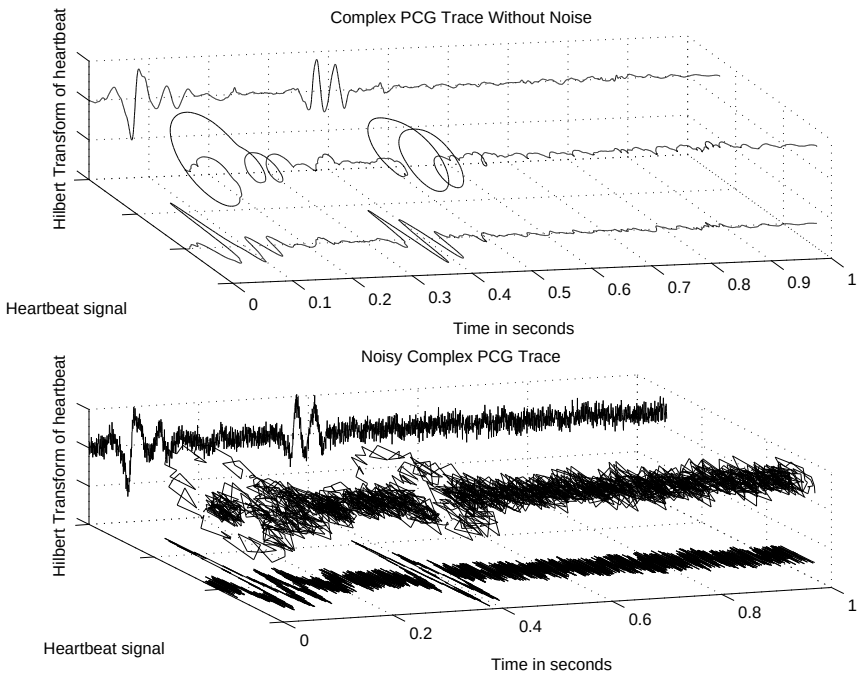

Figure 5.18 Complex PCG trace first with additive white noise and secondly without noise.

problem with using the instantaneous frequency as graphical tool to reveal distinguishing features between pathological and healthy PCGs because the instantaneous frequency plots are plagued with distracting spikes as may be seen in Figure 5.16. This is because in the formula for calculating the instantaneous frequency the numerator is divided by a denominator that is very small at certain times producing a huge spike in the instantaneous frequency. This is the true instantaneous frequency; however, when plotting the signal these huge spikes obscure the critical features of the signal. This problem is dealt with somewhat by smoothing in our software. Smoothing, however, only partially solves the problem and distorts the true value of the instantaneous frequency. Thus, we introduce the concept of the moment of velocity (Davis 2001) which is very similar to instantaneous frequency except that the denominator of the instantaneous frequency equation is removed thus avoiding the distracting spikes. The moment of velocity is defined as

$$\text{moment of velocity} = s\frac{dH[s(t)]}{dt} - H[s(t)]\frac{ds}{dt}. \tag{5.5}$$

Further information on the moment of velocity may be found in Appendix D. If we compare Figures 5.15 and 5.17, it may be seen that the envelope of the signal and the moment of velocity are actually very similar except that the rapidly changing components appear to be accentuated in the moment of velocity plot. This occurs because there are

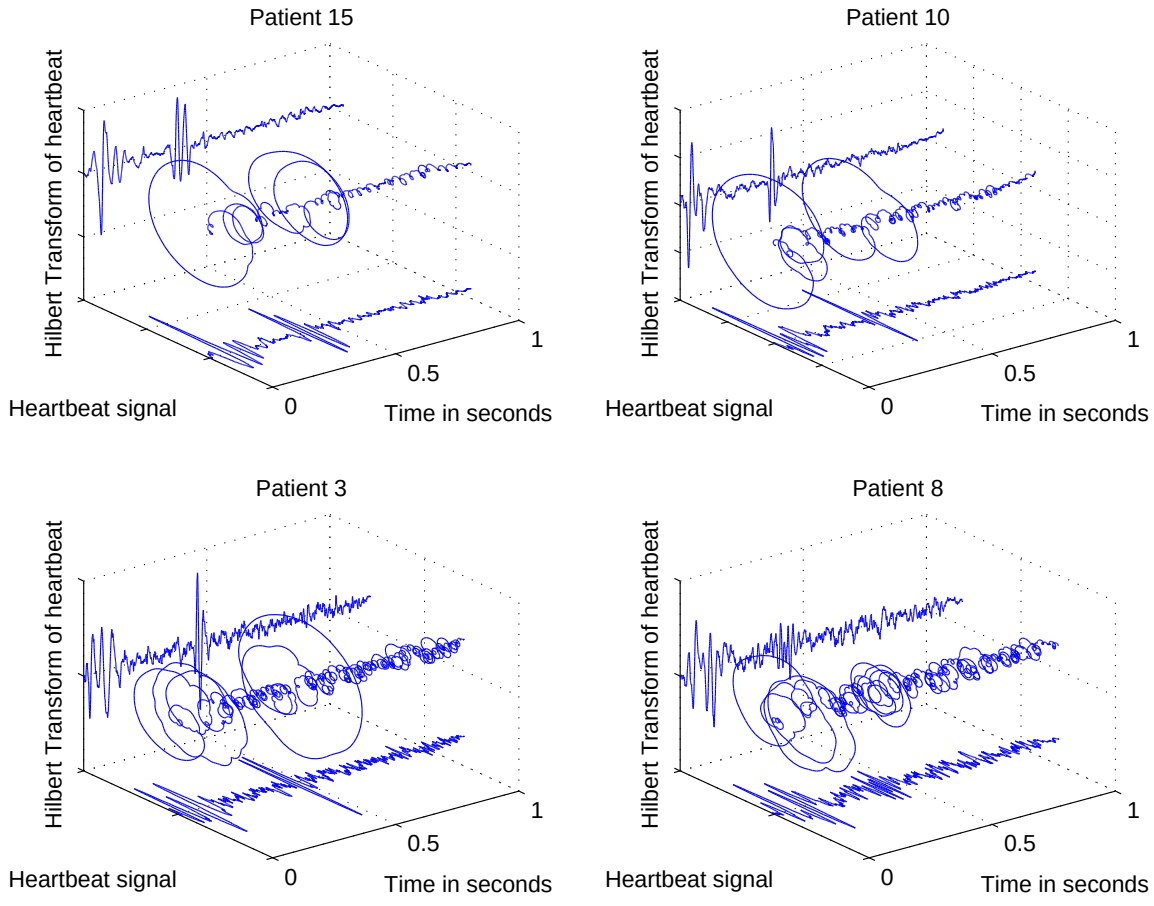

Figure 5.19 This figure shows a complex PCG trace of four different characteristic heartbeats. Patients 10 and 15 are normal subjects whereas Patients 3 and 8 have heart murmurs.

derivatives in the formula for the moment of velocity. However, as can be seen some detail is lost by the moment of velocity. When finer details of the signal are important, it would seem better to use the instantaneous amplitude. When focusing on higher frequency components (such as in the case of heart murmurs) or rapidly changing features, the moment of velocity would be more useful.

Figure 5.18 borrows the concept of a complex trace from seismic data analysis (Taner et al. 1979). The signal and its Hilbert Transform are projected on their prospective axes with the complex trace being a vector sum of the two. This view reveals many features of the signal. The length of the complex trace vector is the instantaneous amplitude. The orientation angle (usually measured relative to the positive axis of the plane where the real signal is projected) is the instantaneous phase. The time rate of change of the phase angle is the instantaneous frequency. It was hoped that the alternative views of the PCG presented by the complex PCG trace would reveal additional information about the PCG that is not readily seen by another technique. Due to limited data, we have not had the opportunity to extensively test this theory. However, Figure 5.19 presents a sample view of the complex PCG trace of two healthy patients and two patients with heart murmurs. There really are no drastic differences but a couple small differences were noticed. The sections outside the the two major heart sounds, which are represented by the big circles, appear to be less compact and well defined, than their non-pathological counterparts, in the complex PCG traces of the heart murmurs. This is believed to be caused by the additional signal content of the murmurs. Also, S1 and S2 are also less compact and well-defined than those seen in the complex PCG traces of the healthy patients.

5.4 Phase Synchronisation

Synchronisation is defined as the "adjustment of frequencies of periodic self-sustained oscillators due to weak interaction" and is also known as phase locking or frequency entrainment (Rosenblum, Pikovsky, Schafer, Tass & Kurths 1999). Phase synchronisation is simply where a relationship exists between the phases of interacting systems while no account is taken of the amplitudes.

Synchronisation is a data analysis technique commonly used in the field of biological experimental studies and the modelling of interaction between different physiological systems (Rosenblum et al. 1999). Examples of this synchronisation include phase locking of respiration with a mechanical ventilator or locomotory rhythms, phase locking of chicken embrion heart cells with external stimuli and interaction of the sinus node with ectopic pacemakers, and heart rate synchronisation with external audio or visual stimuli (Rosenblum et al. 1999).

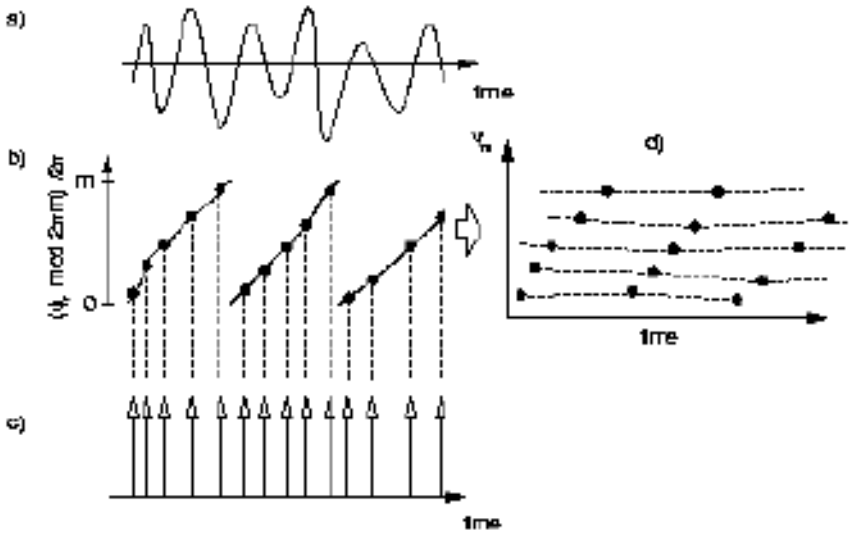

Figure 5.20 This figure demonstrates how the phase stroboscope known as a synchrogram functions. Points in the slow signal (a) are analysed according to a the phase of a fast signal (c). The phase of the slow signal is calculated at these points and wrapped modulo $2\pi m$ (b) and then plotted as shown in (d). In this case, m = 2. $n : m$ phase synchronisation may be recognised in the plot (d) by n horizontal lines. This figure is from Rosenblum et al. (1999)

Because phase synchronisation seems to be a fairly common occurrence between interacting biological systems, we wished to investigate if there was phase synchronisation occurring between the PCG and ECG and if it was present if the occurrence of phase synchronisation differentiated between healthy hearts and hearts with pathological conditions. In work done by Schafer, Rosenblum & Kurths (1998), Schafer, Rosenblum, Abel & Kurths (1999), and Rosenblum et al. (1999), respiration data and the ECG were analysed for phase synchronisation using a stroboscopic technique. The phase of one of the systems was observed at points in time when the phase of the other system had a certain value. The phase of the respiratory signal ψ_r was calculated at times t_k of the appearance of

the k-th R-peak of the ECG and this phase was plotted against t_k. If there was $n:1$ synchronisation and no noise present, there would be n separate values of the respiratory phase, so n horizontal lines would be visible. Noise may smear these lines and create some bands. In order to determine if $n:m$ phase locking is present the respiratory phase was wrapped onto the $[0, 2\pi m]$ and m oscillations were taken to be one cycle and plotted as

$$\psi_m(t_k) = \frac{1}{2\pi}(\phi_r(t_k) \bmod 2\pi m), \tag{5.6}$$

versus t_k. This plot is called a cardiorespiratory synchrogram (Rosenblum et al. 1999) and the concept is illustrated in Figure 5.20.

5.4.1 ECG-PCG Phase Sychronisation, The Cardiosynchrogram

It was decided to investigate if there was phase synchronisation occurring between the PCG and ECG and if the occurrence of phase synchronisation differentiated between normal hearts and hearts with pathological conditions. The same technique presented in the work of Rosenblum et al. (1999), Rosenblum, Pikovsky, Schafer, Tass & Kurths (2000), Schafer et al. (1998), Schafer et al. (1999) to produce the cardiorespiratory synchrogram was used here with minor modifications. The concept of the phase stroboscope was used again. The phase of the phonocardiogram signal ϕ_r at the times t_k of the appearance of the k-th R-peak of the ECG was observed and plotted versus t_k. The appearance of n horizontal lines indicates $n:m$ synchronisation. An example of this may be seen in Figure 5.21. There are four "cardiosynchrograms" with two of patients with pathological cardiac conditions (Figure 5.21 (a) and (b)) and two of healthy individuals (Figure 5.21 (c) and (d)). Note than in Figure 5.21 (b) there appears to be a 1:1 phase synchronisation. Longer PCG recordings and a larger variety of samples are needed though in order to investigate this phenomenon more thoroughly and accurately. For example, Rosenblum et al. (1999) used over 1,500 ECG cycles in their cardiorespiratory synchrograms. In contrast, our data was limited to 20-30 cycles.

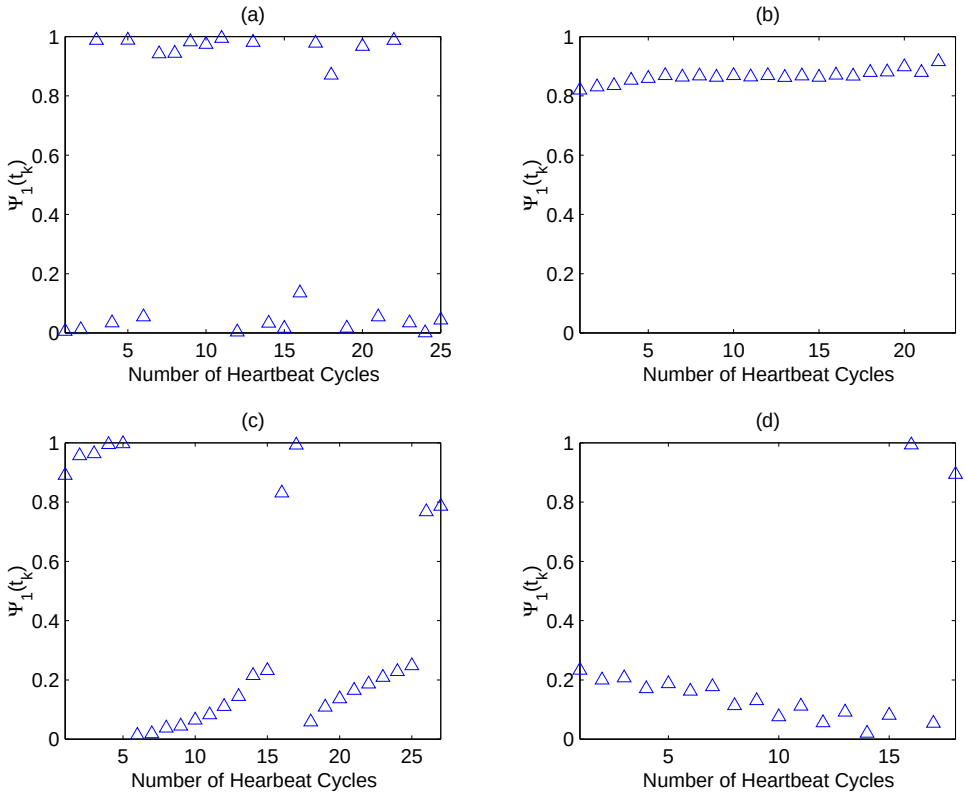

Figure 5.21 The four charts are cardiosynchrograms where the ECG R wave is used as a stroboscopic point to examine the phase of the PCG to see if there is any phase synchronisation occurring. (a) is the cardiosynchrogram of patient 4, who had hypertension and aortic stenosis, (b) is the cardiosynchrogram of patient 8 who had an angioplasty and systolic murmur, (c) is the cardiosynchrogram of Patient 15 who had a normal heart, and (d) is the cardiosynchrogram of Patient 13 who also had a normal heart. For all cases, m=1.

5.5 Chapter Summary

In this chapter, we have attempted to extract information that is not readily viewed by the eye or current data analysis methods from the PCG that would differentiate between healthy hearts and hearts with pathological conditions. We have explored the use of phase space diagrams, HT diagrams, instantaneous signal parameter extraction, and phase synchronisation between the ECG and PCG, but were limited by the amounts and quality of data available to us. These are only indicative results, but they show that further work to investigate the use of these techniques with larger amounts of data would be worthwhile. In the next and final chapter, we will summarise and discuss the work presented and then suggest future research avenues.

Chapter 6

Conclusion and Future Directions

"Stupidity consists in wanting to come to a conclusion."

Elemire Zolla (1926-)

ITALIAN PHILOSOPHER AND ESSAYIST

6.1 Introduction

This chapter will firstly present a summary of the information that was presented in this thesis, then will discuss the conclusions reached and finally will provide recommendations for future research.

6.2 Summary

Chapter 1 introduced the subject of the thesis and provided some necessary background information including a literature review of the work performed in this area.

Chapter 2 reviewed the background and history of phonocardiography and provided information about the instrumentation and procedures used to obtain the data. It was argued that tools in addition to the human ear were needed to clearly hear heartsounds due to limitations of the human ear. The development of the modern stethoscope and PCG were reviewed. Next, the specific equipment which was used in our study was presented. The final section discussed the recordings collected.

Chapter 3 surveyed the background and theory of the de-noising methods used in this study in order that the reader could comprehend the results of the de-noising study presented in the next chapter. We introduced the theory and background of the de-noising techniques, wavelet de-noising, optimised wavelet de-noising, wavelet packet de-noising, the matching pursuit technique, and averaging, which were applied to the de-noise the PCGs.

Once this necessary background to the de-noising study was covered, Chapter 4 examined the results and recommendations of the de-noising study.

After the PCG was de-noised, different methods of extracting features from the PCG and classifying the PCG according to this information were explored in Chapter 5. We attempted to extract information that is not readily viewed by eye or current data analysis methods from the PCG that would differentiate between healthy hearts and pathological subjects. We explored the use of phase space diagrams, HT diagrams, instantaneous signal parameter extraction, and phase synchronisation between the ECG and PCG, but

we were limited by the quantity and quality of data available. The results presented are only indicative, but they demonstrated that further work to investigate the use of these techniques with larger amounts of data would be worthwhile.

6.3 Discussion and Conclusions

6.3.1 PCG De-noising

A study which compared wavelet, wavelet packet and matching pursuit de-noising applied to knee-joint vibrations, which are complex, non-stationary signals, concluded that the MP method outperformed wavelet and WP de-noising with the WP de-noising performing better than the wavelet de-noising method (Krishnan & Rangayyan 2000). Given the results presented in Krishnan & Rangayyan (2000) and the similarity in the types of signals, both being complex, non-stationary signals, it was expected that the MP method would perform the best followed by the WP method. We anticipated that WP de-noising would perform better than wavelet de-noising because WP analysis adaptively chooses the best basis based upon an entropy search and has more decomposition structure combinations to choose from. However, the time-frequency and time-scale de-noising methods performed roughly equally except for a slightly better performance by optimised wavelet de-noising in de-noising the PCGs as witnessed by the evidence provided in Chapter 4. Tables 4.5, 4.6, 4.7, and 4.8 show typical results for optimised wavelet de-noising, wavelet de-noising, wavelet packet de-noising and matching pursuit de-noising applied to different PCGs with various amounts of white noise added. Figure 4.39 summarises the results of the previously mentioned tables into bar charts. Examining these tables and the figure, it becomes apparent that optimised wavelet de-noising seems to perform slightly better than the other methods for removing white noise from PCGs. The other methods appeared to perform about equally as well as each other. De-noising results with wavelet and WP techniques depend greatly on the selection of the threshold value for the coefficients (Krishnan & Rangayyan 2000). This fact may be used to partly explain why optimised wavelet de-noising outperformed the other methods because it has various options for the methods used to perform thresholding and noise modelling. Another possible reason

optimised wavelet de-noising had a higher SNR after de-noising compared to the other methods is that it appears to lose less information by the de-noising process alone which is supported by Figures 4.8, 4.9, 4.17, 4.21, 4.23, 4.27, and 4.37.

Figure 4.38 visually demonstrated how well the various de-noising methods perform. Note that the MP de-noising method does not pick up some of the finer details from the original PCG because the low energy coherent structures have been destroyed by the high level of noise and are unable to be extracted by the MP method. Evidence of this fact also may be seen in the study done by Krishnan & Rangayyan (2000). For a large amount of white noise (SNR=0 dB) added to the signal, the MP method only outperformed the others by a fairly small amount, but with smaller amounts of noise (SNR=10 dBs) added to the signal, the MP method outperformed the WP and wavelet de-noising methods by a much more significant amount. This phenomenon indicates that large amounts of noise destroy low-energy coherent structures within that signal making it impossible for them to be extracted from the noisy signal by the MP method.

For optimised wavelet de-noising, it did not appear that any one wavelet performed much better than another wavelet. However, some of the lower order wavelets did not perform very well due insufficient numbers of vanishing moments. For a Daubechies wavelet of order N, the support length of ψ and ϕ is $N-1$ and the vanishing moment of ψ is N (Misiti et al. 1996). The order of regularity of a wavelet is the number of continuous derivatives which it possesses (Hubbard 1996). Poor regularity may introduce artifacts (Hubbard 1996). Regularity may be increased by increasing the length of support (Hubbard 1996), which increases with N. Vanishing moments influence what signal content is picked up by the wavelet transform (Hubbard 1996). With one vanishing moment, linear functions are not seen, and with two vanishing moments, quadratics are not picked up. Thus, by increasing the number of vanishing moments, the lower order components of the signal may be seen. We reached the conclusion that a decomposition level of 5 produced reasonable results, while decomposing the signal further often produced marginal benefits and increased computation time. Soft thresholding definitively outperformed hard thresholding. Of the four threshold selection rules, the "rigrsure" rule performed the best because small details of the PCG signal are located in the noise range, and the best choice of the threshold rescaling methods proved to be the "sln" method due to the presence of

white noise in the signal.

For both wavelet and WP de-noising, most wavelets performed roughly equally except for some of the lower order wavelets due to the properties of the wavelets especially insufficient numbers of vanishing moments. Decomposition levels of 3-5 were found to perform the best. The WP de-noising process seems to lose more of the original signal content than the wavelet de-noising process, while the wavelet de-noising process loses even more original signal content than optimised wavelet de-noising.

Overall, the MP method de-noises the PCG about as well as WT or WP de-noising, but attention must be given to the number of time-frequency atoms used which may be determined by observing the decay parameter as it levels out. Also, the amount of information lost by the de-noising process, if not enough time-frequency atoms are used, itself is of concern.

Averaging seemed to produce significant improvements especially if there is a large amount of noise present in the signal. Averaging a series of 50-75 cycles seemed to give the best result in terms of recording and computation time tradeoff.

Averaging may be used in combination with one of the other de-noising methods. Because optimised wavelet de-noising performed better than the other time-frequency and time-scale methods, this method is recommended to be used in conjunction with averaging in certain cases. Figure 4.40 showed a comparison of using wavelet de-noising only, averaging only and wavelet de-noising combined with averaging. It clearly demonstrated that combining the techniques is much more effective and given a choice between averaging 30, 60 or 90 cycles, averaging 60 cycles provides a good compromise in terms of de-noising and recording and computation time. However, in some cases, averaging is not appropriate and may have certain drawbacks (Karpman et al. 1975). If a non-deterministic component of the signal such as a murmur is averaged, that component will be distorted or possibly disappear. Variations in the timing of individual heart cycles may lead to cancellation of part of the signal, but, with low-frequency signals, this effect is usually negligible. With timing variations, there may be a small overlap of cardiac events which are quite close in time resulting in falsely widening the duration of these events.

6.3.2 PCG Data Analysis

After the PCG signal is de-noised, the signal was analysed for hidden information that may differentiate a subject with a healthy heart from a patient with a pathological condition. The techniques that were used are the (a) phase space diagram, (b) the use of the Hilbert Transform to construct the HT diagram and to extract instantaneous signal features, and (c) phase synchronisation between the ECG and PCG. These techniques were chosen because they produce easy to read graphs and diagrams that reveal information hidden in the PCG which may not be seen by a simple visual inspection of the signal itself. No statistical analysis techniques were chosen because statistical techniques normally require very long recordings, and PCG recordings are typically not very long due to the difficulties presented by the recording techniques used to obtain a clean recording. Also, we were limited in our investigation of these data analysis techniques by a smaller than expected number of PCG recordings. Thus, the results presented are initial indicative findings only.

It was thought that the phase space and HT diagrams would reveal information that was not apparent by viewing the PCG. From the limited data available we examined the difference between normal PCGs and PCGs from patients possessing murmurs. Figure 5.11 demonstrated that there appeared to be a difference between the normal PCGs and the PCGs of murmurs, confirming findings by Hall (1999) and Maple (1999). Murmurs normally contain much higher frequency content than normal PCGs. Phase space diagrams accentuate higher frequencies, thus we could expect phase space diagrams to emphasise the higher frequencies contained in the murmurs. Examining Figure 5.12, which shows the HT diagrams of the same PCGs, we can not see a notable difference between the healthy and pathological heart sound recordings. It would be worthwhile to further examine the use of phase space diagrams and HT diagrams to identify pathological cardiac conditions. It would be expected that they might be used to identify different characteristics as phase space diagrams accentuate higher frequencies while the HT diagrams treat all frequencies equally.

The instantaneous frequency was used to visually demonstrate how well the PCG de-noising techniques performed. Figure 5.13 showed a characteristic heart sound, then the heart sound with noise added, and finally the noisy heart sound with noise removed both

by the wavelet and wavelet packet de-noising processes. The corresponding instantaneous frequencies were also shown. It is clearly visible from the instantaneous frequency plots when there are large amounts of noise present.

It was thought that the instantaneous signal parameters could reveal information about that PCG that was not readily extracted by viewing the PCG itself or currently used techniques. We were limited by the number of recordings available. We demonstrated that using the instantaneous amplitude of the PCG of a patient is reproducible as shown in Figure 5.14. In Figure 5.15, which showed the instantaneous amplitude of PCGs for patients with various pathological conditions and patients with normal hearts, the healthy patients appeared to have a well defined and compact S1 and S2 whereas some of the patients with pathological conditions did not. It was hoped that the instantaneous frequency of the characteristic heart sounds might contain obvious differentiating features between the pathological and normal characteristic heart sounds. Examining Figure 5.16, did not reveal this to be the case except that the instantaneous frequency of the normal characteristic heart sounds appear to be "quieter" than the pathological cases which could produce extra signal content especially in the case of heart murmurs. There is a problem with using the instantaneous frequency as graphical tool to reveal distinguishing features between pathological and healthy PCGs because the instantaneous frequency plots are plagued with distracting spikes. Thus, we introduced the concept of the moment of velocity (Davis 2001) which is very similar to instantaneous frequency except that the denominator of the instantaneous frequency equation is removed so that division by small numbers is avoided thus escaping the distracting spikes. The envelope of the signal and the moment of velocity were shown to be very similar except that the rapidly changing components appear to be accentuated by the moment of velocity. This occurs because there are derivatives in the formula for the moment of velocity. Some detail is lost by the moment of velocity although it could be seen better if plotted on a logarithm scale. When finer details of the signal are important, it would seem better to use the instantaneous amplitude. When focusing on higher frequency components (such as in the case of heart murmurs) or rapidly changing features, the moment of velocity would be more useful than the signal envelope.

We introduced the concept of the complex trace, as shown in Figure 5.18, which is a tech-

nique borrowed from seismic data analysis. It was anticipated that the alternative views of the PCG presented by the complex PCG trace would reveal additional information about the PCG that is not readily seen by another technique. We did not see any drastic differences between the complex PCG traces of healthy and diseased subjects (Figure 5.19), but a couple small differences were noticed. In the complex PCG traces of the heart murmurs, the sections other than the two major heartsounds, which are represented by the big circles, appear to be less compact and well defined, than their non-pathological counterparts, like there might be additional signal content, such as that of the murmurs, represented. Furthermore, S1 and S2 are also less compact and well-defined than those seen in the complex PCG traces of the healthy patients.

Because phase synchronisation between biological systems is a commonly occurring phenomenon (Rosenblum et al. 1999), we decided to test if there was phase synchronisation occurring between the PCG and ECG and if the occurrence of phase synchronisation differentiated between healthy hearts and hearts with pathological conditions. Of the four cardiosynchrograms shown in Figure 5.21, there appears to be a 1:1 phase synchronisation between the ECG and PCG in Figure 5.21 (b). Longer PCG recordings and a larger variety of samples are needed though in order to investigate this phenomenon more thoroughly and accurately. For example, Rosenblum et al. (1999) used over 1,500 ECG cycles in their cardiorespiratory synchrograms whereas we only had 20-30 cycles available for use.

6.4 Future Research Directions

The use of genetic algorithms to de-noise signals has not been widely explored. Implementing a genetic algorithm to de-noise heart sounds could be useful. Some work in the area of genetic algorithm de-noising has been done but has not been extremely promising. Lankhorst & van der Laan (1995) studied signal approximation by wavelet-like functions using genetic algorithms but found that existing methods were faster and provided comparable approximation quality but that this algorithm provided greater flexibility. Vertan, Vertan & Buzuloiu (1997) presented a reduced computational genetic algorithm for noise removal which used non-linear filters and used a model-free approach which would be

important for heart sounds. Although it was tested on two-dimensional images, it is most likely modifications could be made to the algorithm so that it could be used to process time domain signals. With recent increases in computational power and advances in genetics algorithms and wavelets, a new algorithm or modification of these algorithms could prove to be an effective method of de-noising highly non-stationary signals such as the PCG. Another method of obtaining a relatively noise-free PCG would be to design a special pick-up device that would suppress the noise. There has been some research performed in this area (Hök 1991), but much more work could be done.

Extracting information from the PCG has proven to be a valuable tool in diagnosing various cardiac diseases and conditions. Spectrograms of PCGs have proven to be of clinical use in identifying murmurs, and aortic ball variance in patients with valve prostheses in early work done by Geckeler et al. (1954) and Hylen et al. (1969). More recent work in the identification of murmurs from PCGs using time-frequency and time-scale techniques has been done by Debiais, Durand, Pibarot & Guardo (1997), Debiais, Durand, Guo & Guardo (1997), El-Asir et al. (1996), and Zhang, Durrand, Senhadji, Lee & Coatrieux (1998). Xiao et al. (1999) presented a technique, called the phonocardiogram exercise test, for detecting cardiac reserve in healthy and diseased patients using the PCG is presented. Changes in the cardiac state may be seen in the PCG. This technique made use of the first heart sound. Comparing the PCG before and after exercise, allows these changes to be examined because changes in the amplitude of S1 are closely linked with the maximum rate or rise of left ventricular pressure which measures cardiac contractility.The timing and intensity of the second heartsound is of importance in patients with pulmonary stenosis (Cheitlin et al. 1993). Also, if the time interval between the A2 and P2 components of the second heartsound change, a pathological condition may be present (Durand & Pibarot 1995) and this time interval has been successfully calculated by Khadra et al. (1991) using the WT. For patients possessing a heart valve implant, the power spectra of the second heartsound may be used to indicate the condition of the valve (Durand & Pibarot 1995). Spectral analysis of the second heartsound may also be used to obtain a non-invasive estimation of the pulmonary and systemic arterial pressures which are used to diagnose pulmonary hypertension (Durand & Pibarot 1995). If the third heartsound is identified in patients over 40 years old, it can be a pathological sign of ventricular failure

(Durand & Pibarot 1995). The fourth heartsound can be a sign of left ventricular hypertrophy and coronary artery disease (Durand & Pibarot 1995). Thus, we have highlighted that information of clinical importance can be extracted from the PCG. We believe that further research into the use of the PCG as a clinical diagnostic tool is warranted due to the fact that it has been proven to be a clinically significant diagnosis tool, is inexpensive, non-invasive, reliable and cheap. We acknowledge the need to develop standards for PCG equipment and recording techniques to make the process more universal and reliable. Below, we provide some suggestions for areas of future research.

Neural networks could prove to be a valuable tool in the area of data analysis and classification of PCGs. It has already been used in many studies involving ECG signals. Dickhaus & Heinrich (1996) discussed the idea of using wavelet networks to classify ECG signals. Khadra, Abdallah & Nashash (1998) studied life threatening ventricular arrythmias using neural wavelet analysis. The algorithm used by Khadra et al. (1998) is based on a linear-approximation distance-thresholding compression and back-propagation neural network and was used to analyse ECG data efficiently and reliably compared with other methods. Carranza & Andina (2001) used wavelet pre-processing for ECGs and then have planned to apply neural networks to the ECG data for classification. The use of neural networks have also been explored to some extent for PCGs. Akay, Akay, Welkowitz & Kostis (1994) investigated the use of a wavelet-based fuzzy neural network to non-invasively identify coronary artery disease based upon clinical examination variables and information extracted from the PCG. The results of this study demonstrated that wavelet-based fuzzy neural networks are possibly able to differentiate between healthy and pathological individuals.

We recommend that the data analysis techniques which were initially presented in the previous chapter be further explored by using more PCG recordings of various cardiac conditions as we were severely limited in the number of recordings available at our disposal. The phase space diagram, HT diagram, instantaneous amplitude or velocity, and the cardiosynchrogram are all techniques which warrant further investigations.

We did not perform statistical analysis on our PCG data. However, after consulting with Ivanov (2000), who has performed substantial research in the area of ECG statistical

analysis (Havlin, Amaral, Ashkenazy, Goldberger, Ivanov, Peng & Stanley 1999, Ivanov et al. 1996, Ivanov, Amaral, Goldberger & Stanley 1997, Ivanov, Amaral, Goldberger, Havlin, Rosenblum, Struzik & Stanley 1999, Ivanov, Bunde, Amaral, Havlin, Fritsch-Yelle, Baevsky, Stanley & Goldberger 1999, Ivanov, Amaral, Goldberger, Havlin, Rosenblum, Struzik & Stanley 2000), a few ideas were proposed to extract some statistical features from the PCG data. If the following statistical techniques are going to be used, it would be better if slightly longer recordings, than the 30 second recordings we had available to us, were used. The values of the peaks of the first and second heartsounds of each cycle could be extracted from the PCG. Then, the student t-test could be applied to these peaks to investigate if there exists statistically different values of the peaks. Also, the shape of the histograms of the peak values of S1 and S2 could be examined. Then, the same test could be applied to not only the peak points but to all points along the first and second heartsounds. Next, the shape of the histograms of the first and second heartsounds could be examined to investigate if they are different. The time period between the peaks of the first and second heartsound for each cycle could be calculated and then used to estimate the average and standard deviation of this time period. Correlations could also be calculated. These techniques could be performed for patients who are healthy and who have pathological cardiac conditions to see if there are any differences between the groups revealed by these techniques.

In summary, the PCG is a potentially rich source of clinical information and a number of possible engineering-based signal analysis approaches appear fruitful for future research.

Appendix A

Escope Specifications

Audio gain	27 dB @ 200 Hz
Frequency response	Heart sounds: 100 to 240 Hz (-3 dB) SPL/SPL and 45 to 900 Hz (-20 dB) Breath sounds: 125 to 350 Hz (-3 dB) and 50 to 2000 Hz (-20 dB)
Maximum output	124 dB SPL, undistorted
Microphone	Sound pressure type with electret microphone element located in chest piece
Speaker	Dynamic type with 8 ohm impedance located in control box
Sounds out	Audio output with capability to drive 30 ohm headphone or Cardionics stethoscope. Nominal 0.5 Vpp level suitable for recording on cassette tape.
Weight	255 grams
Length	87.6 cm from chest piece to binaural earpieces
Controls	VOLume/OFF: Control sound volume and turns power off when adjusted to fully CCW detent position. Internal power time shuts off power after 2 minutes. RESTART: Restarts power timer. HS/BS: Selects heart sounds or breath sounds frequency response range via slide switch on control module.
Power Source	9 Volt alkaline battery powering E-scope for approximately 150 hours

Table A.1 This gives the specifications for the Escope, the electronic stethoscope, used to record the phonocardiograms (Cardionics 1999).

Appendix B

Some Data From Patient Recordings

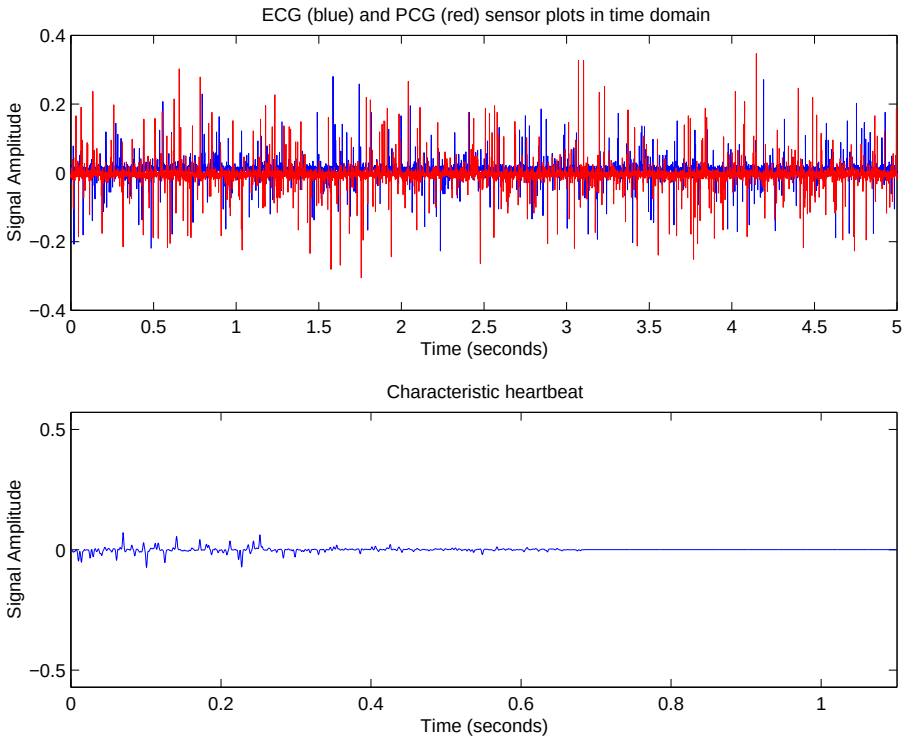

Figure B.1 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #1

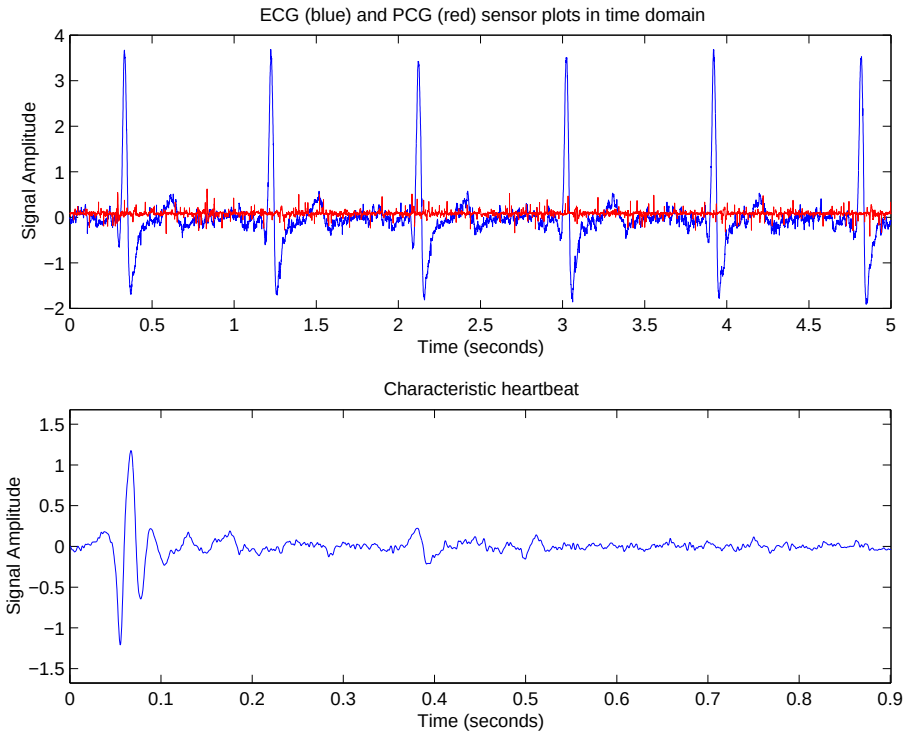

Figure B.2 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #2

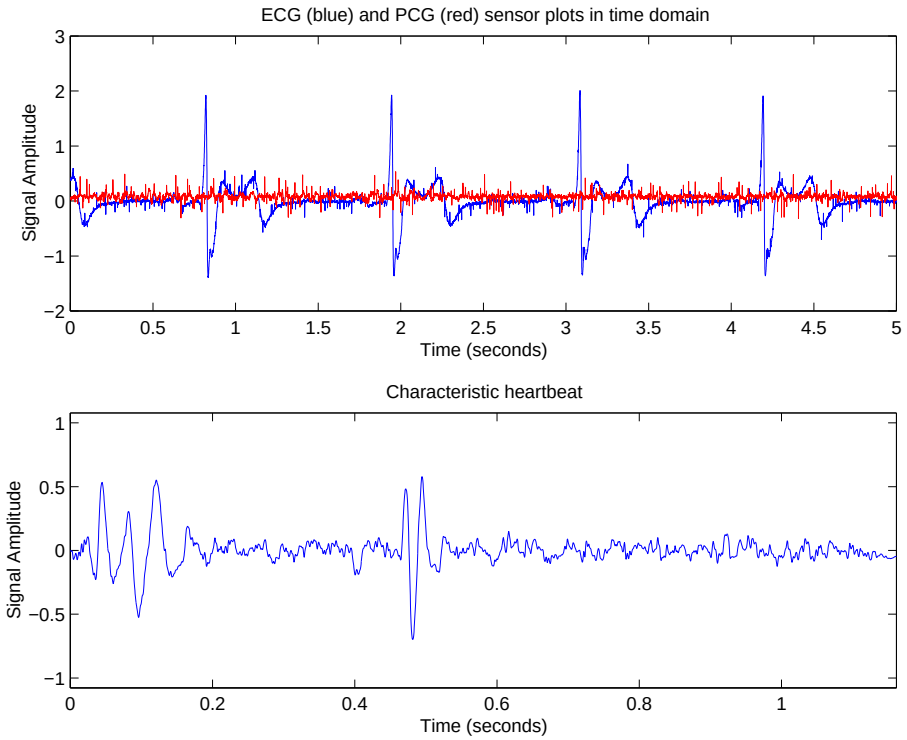

Figure B.3 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #3

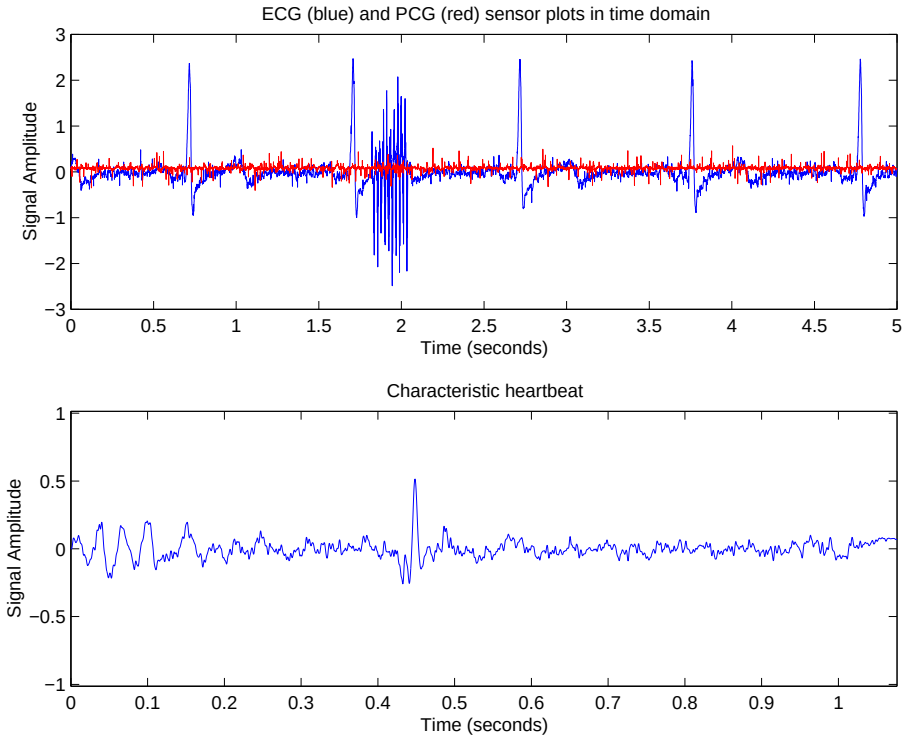

Figure B.4 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #4

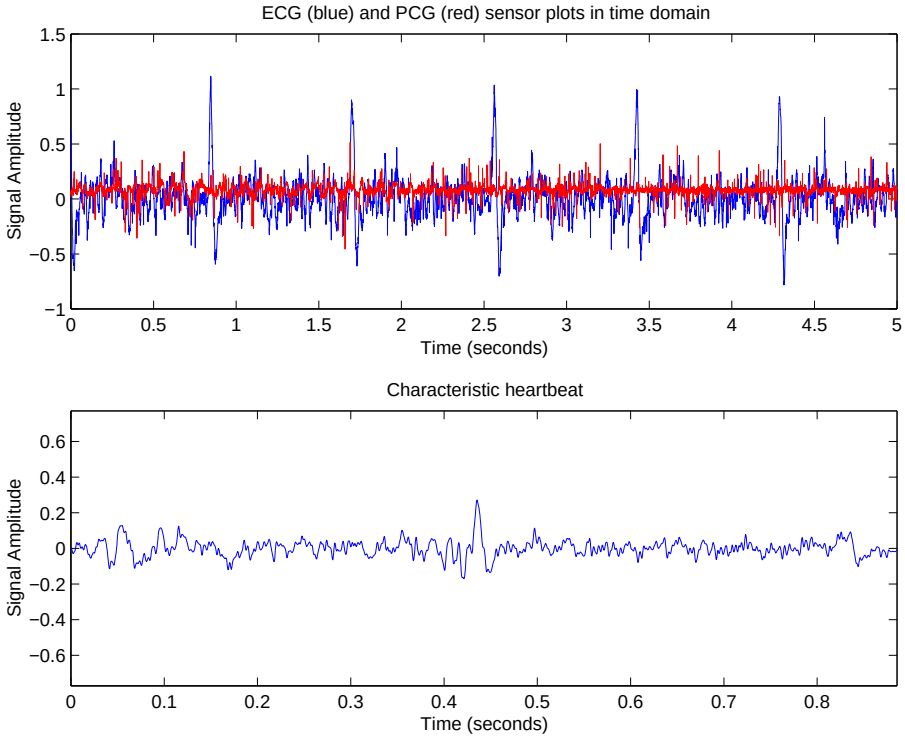

Figure B.5 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #5

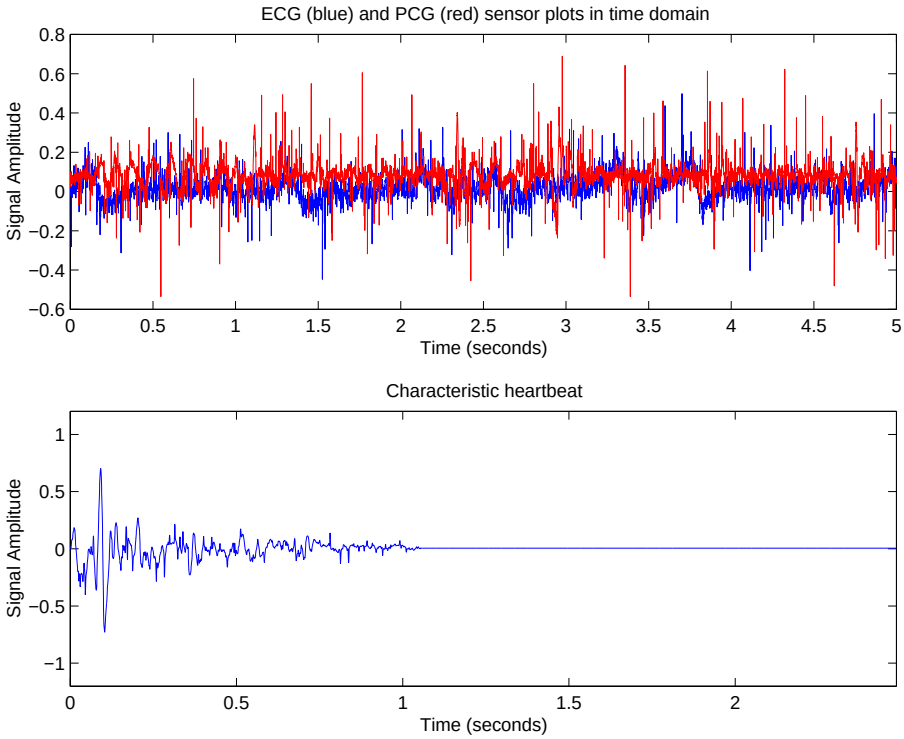

Figure B.6 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #6

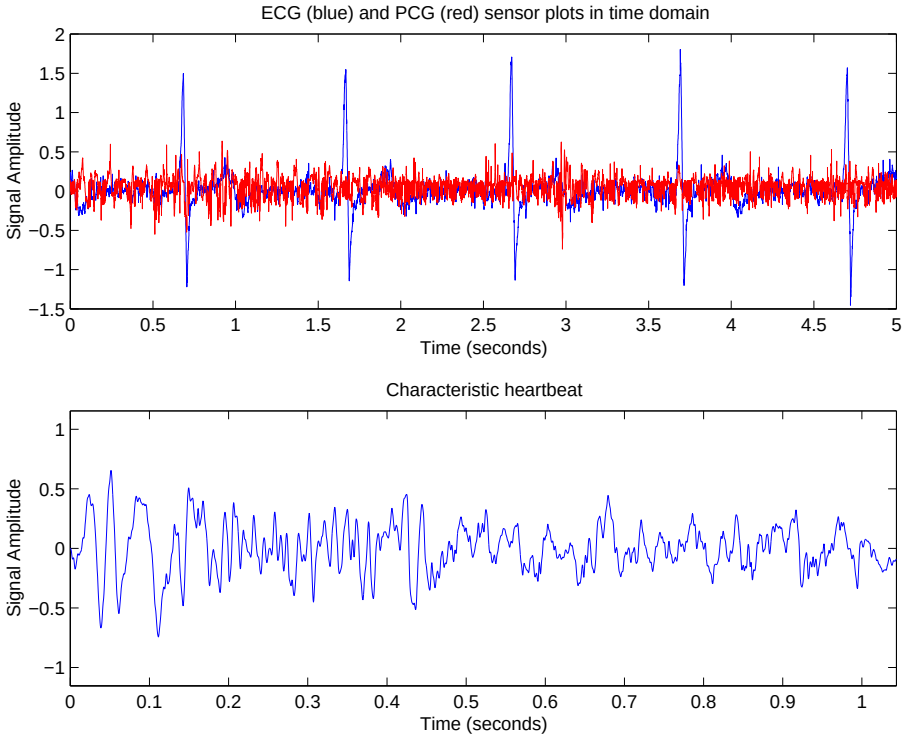

Figure B.7 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #7

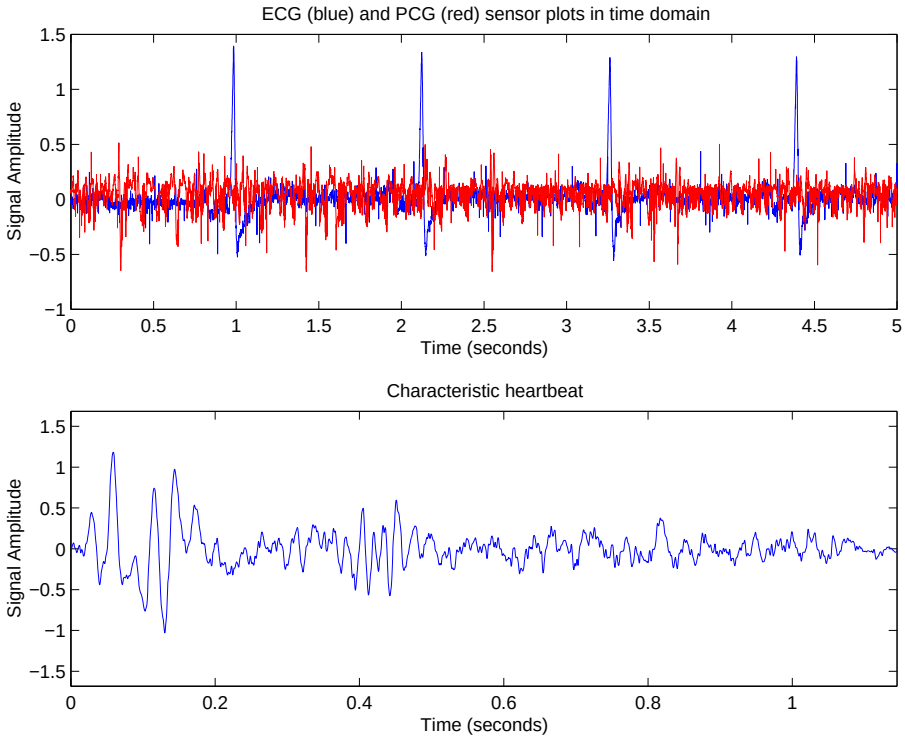

Figure B.8 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #8

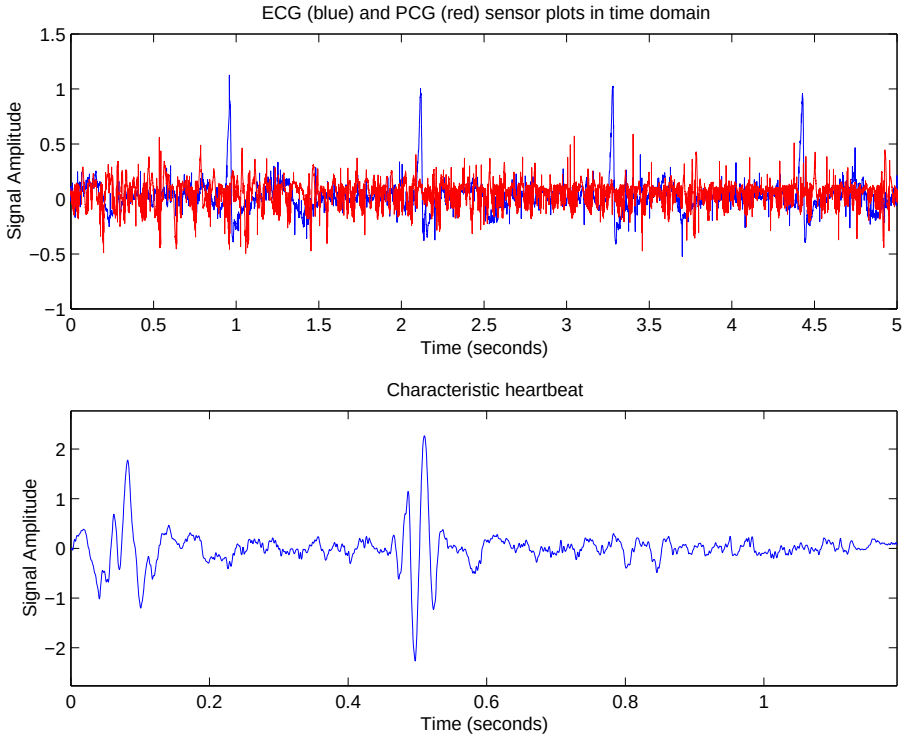

Figure B.9 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #9

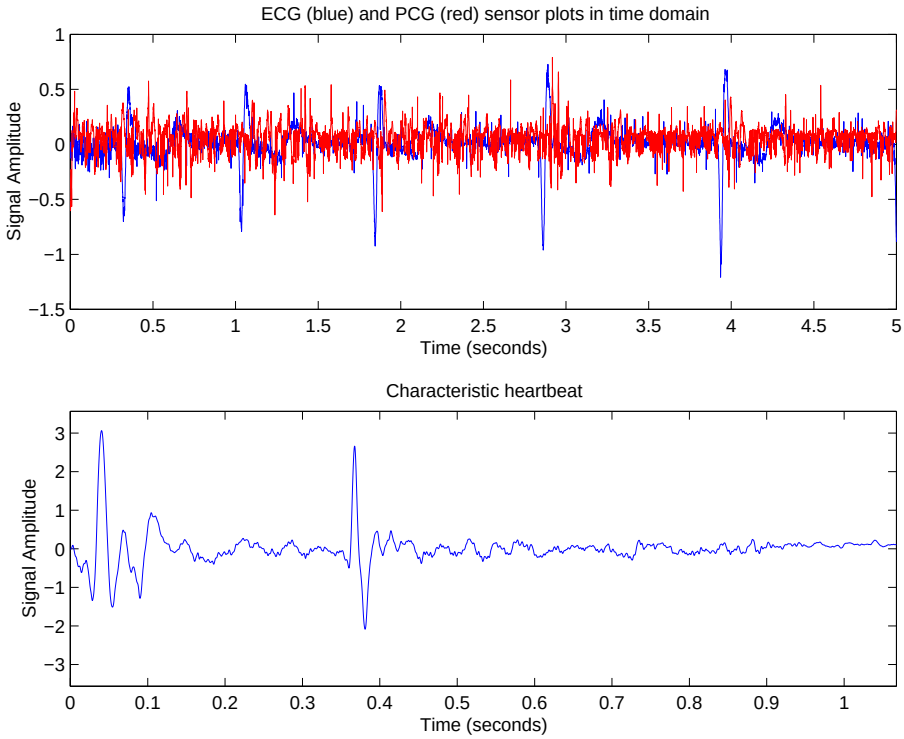

Figure B.10 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #10

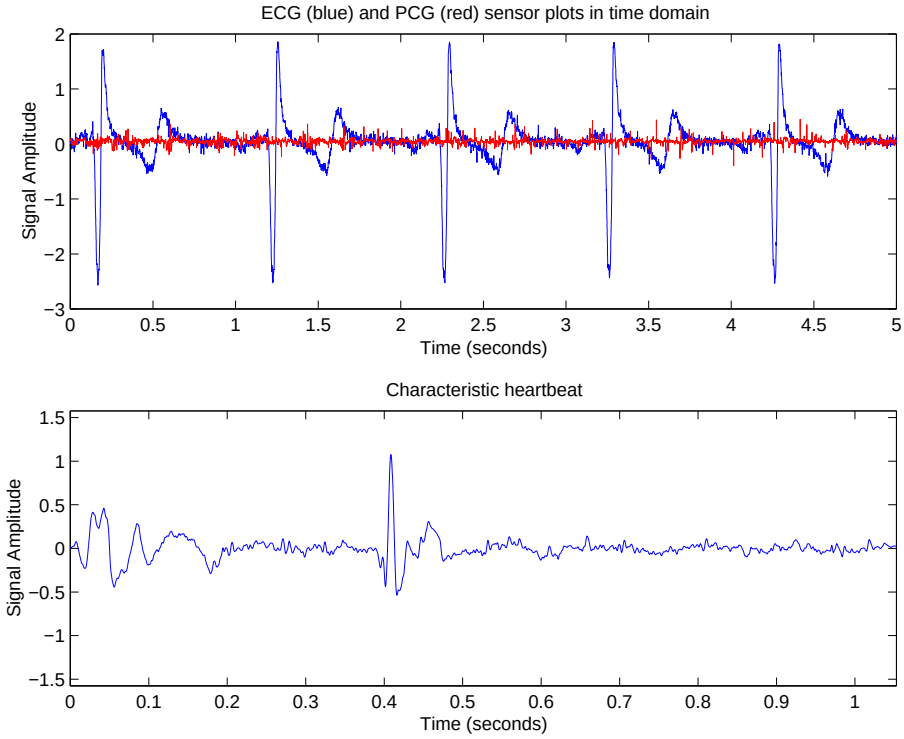

Figure B.11 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #11

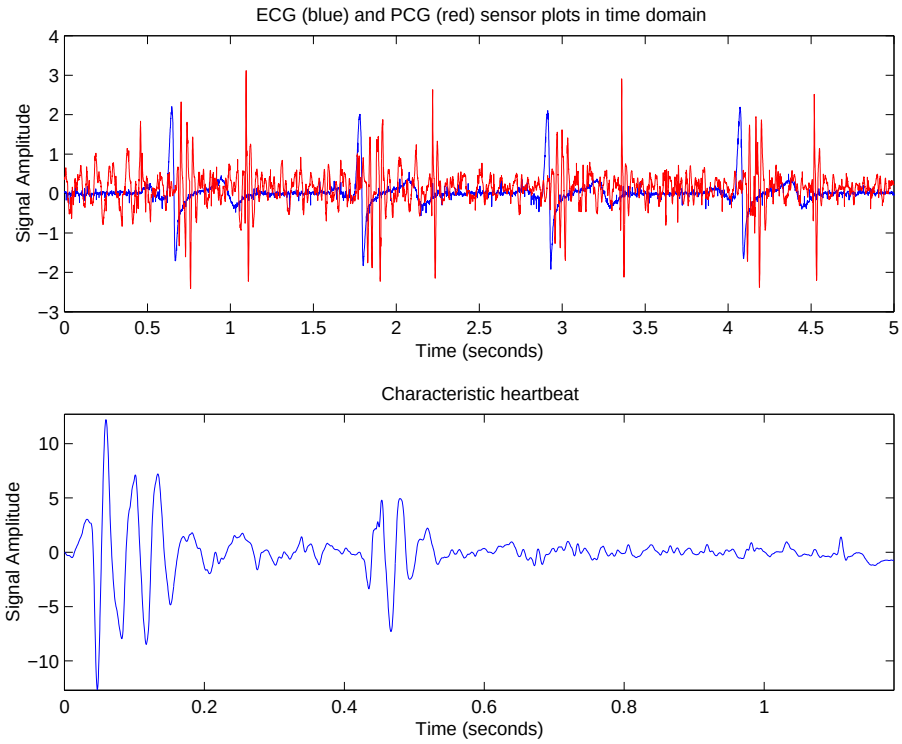

Figure B.12 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #12

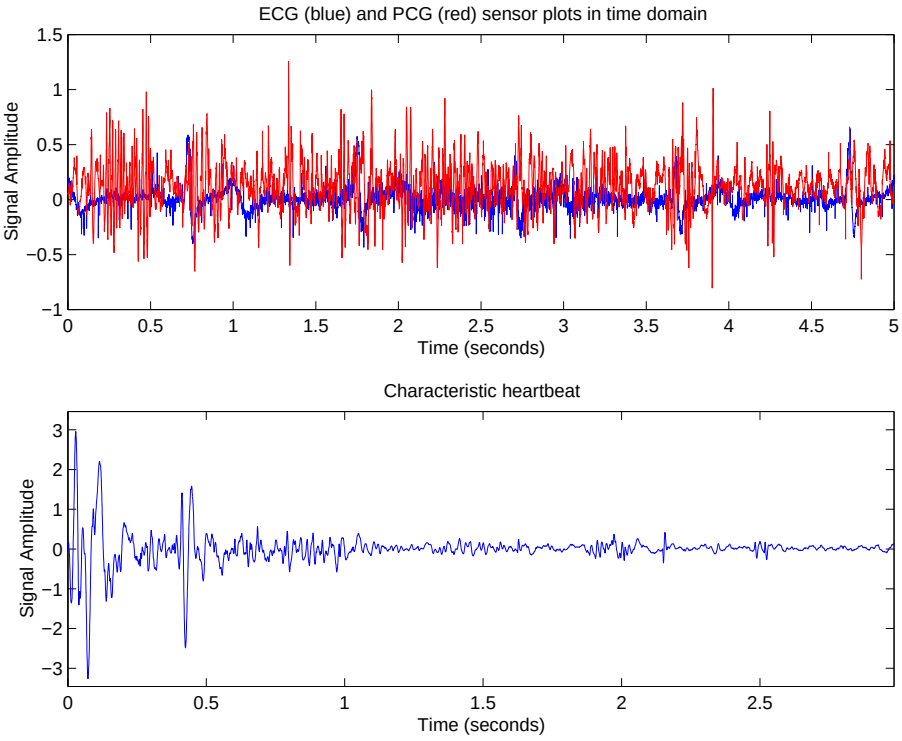

Figure B.13 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #13

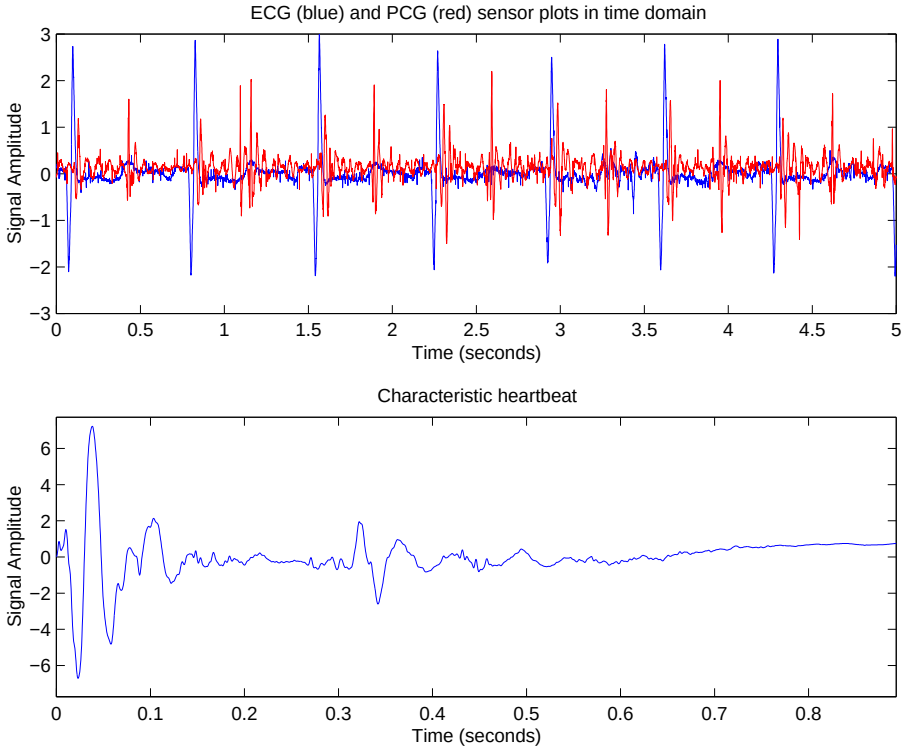

Figure B.14 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #14

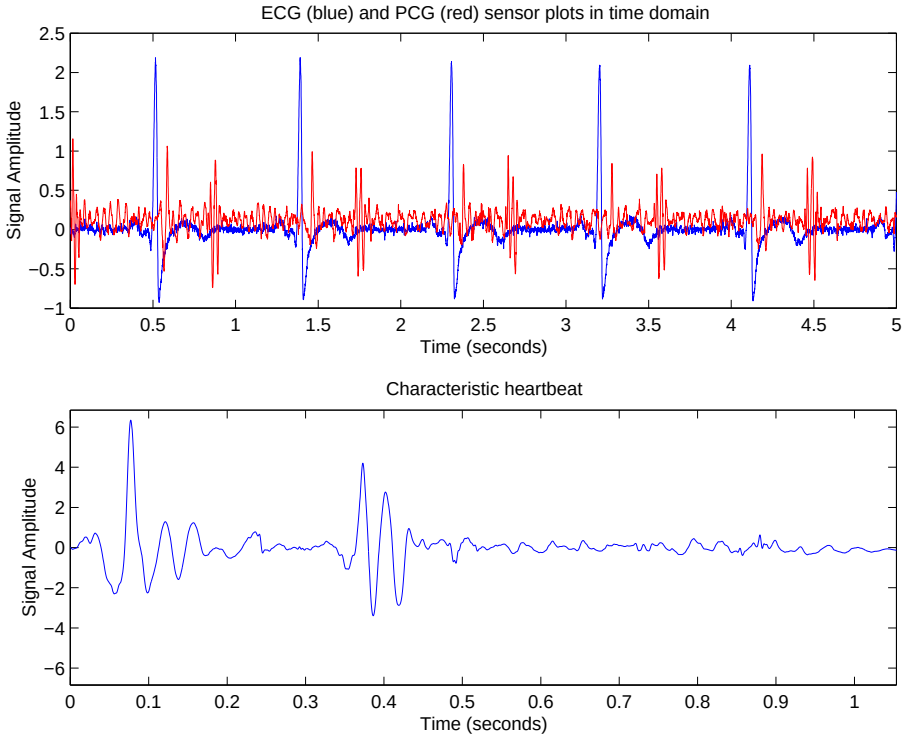

Figure B.15 5 Second Sample of ECG/PCG Recording and Characteristic Heartbeat From Patient #15

Appendix C

Information on the Design of the PCG/ECG System

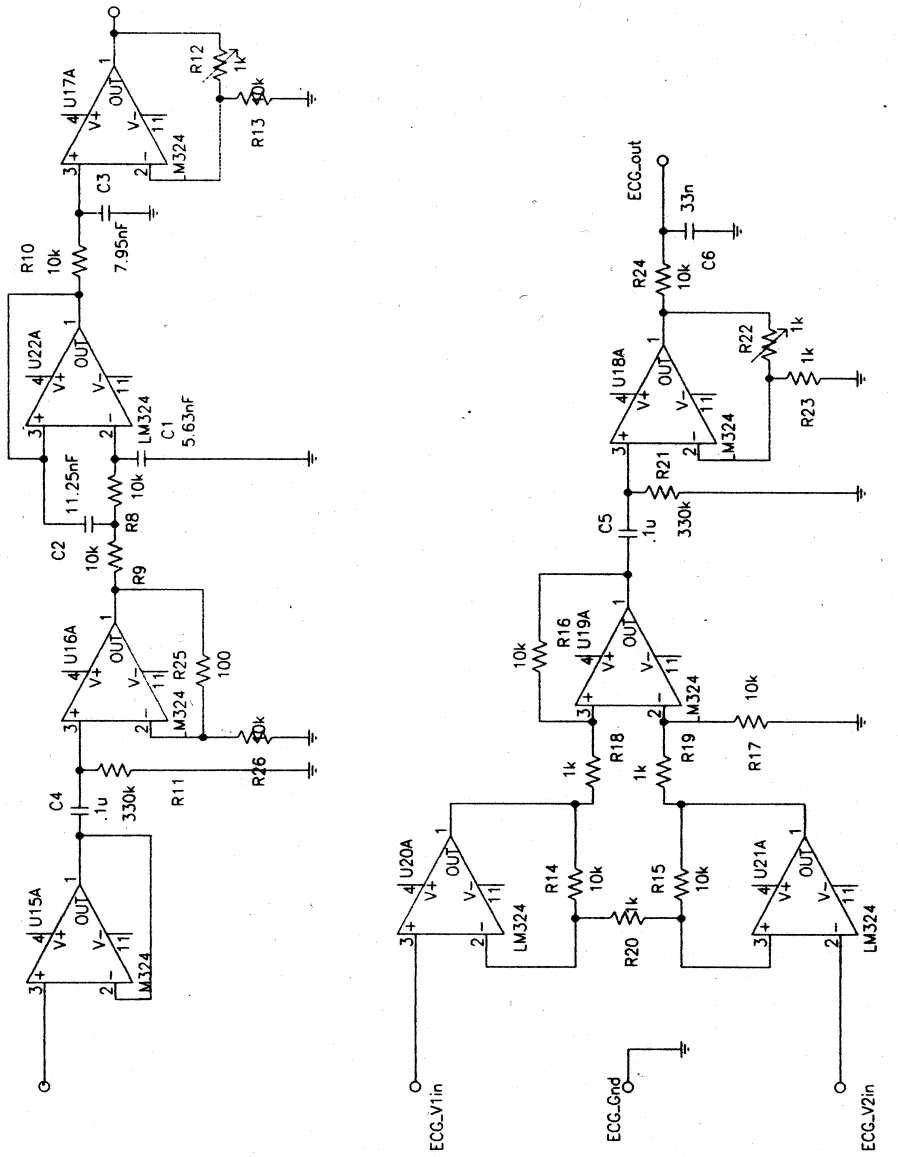

Figure C.1 PCG/ECG System Circuit Diagram (Hall 1999)

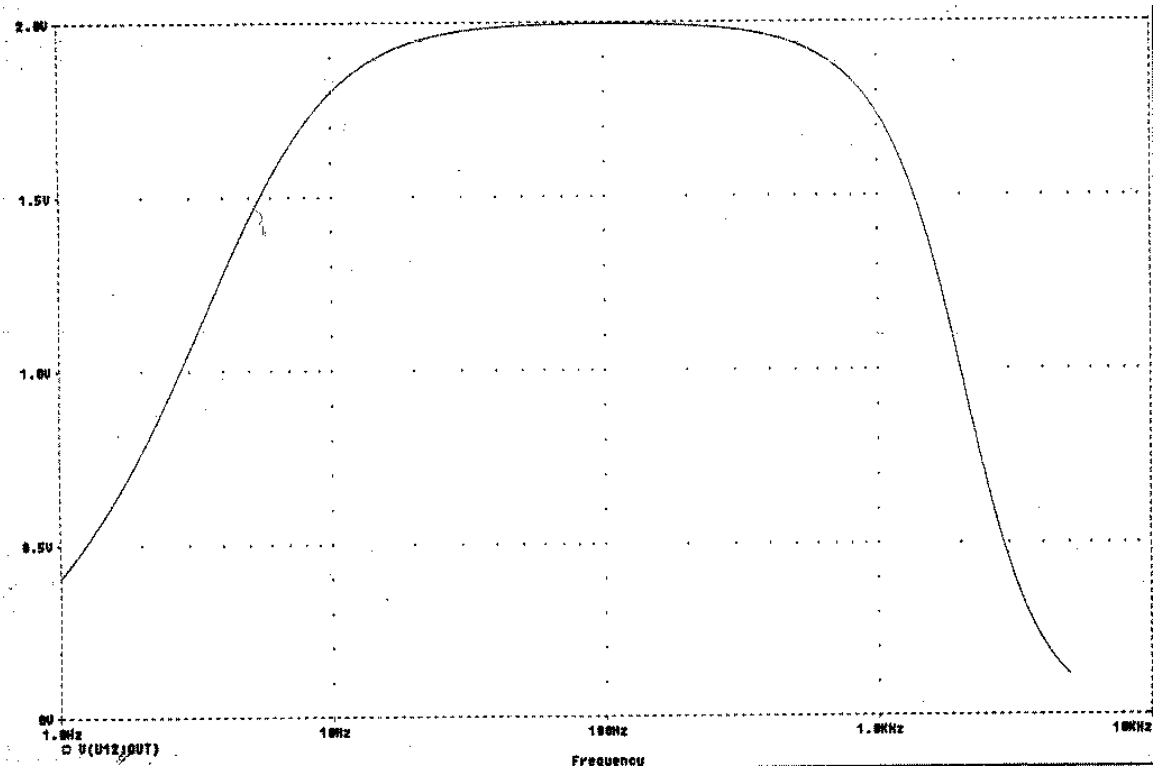

Figure C.2 PCG Filter Frequency Response (Hall 1999)

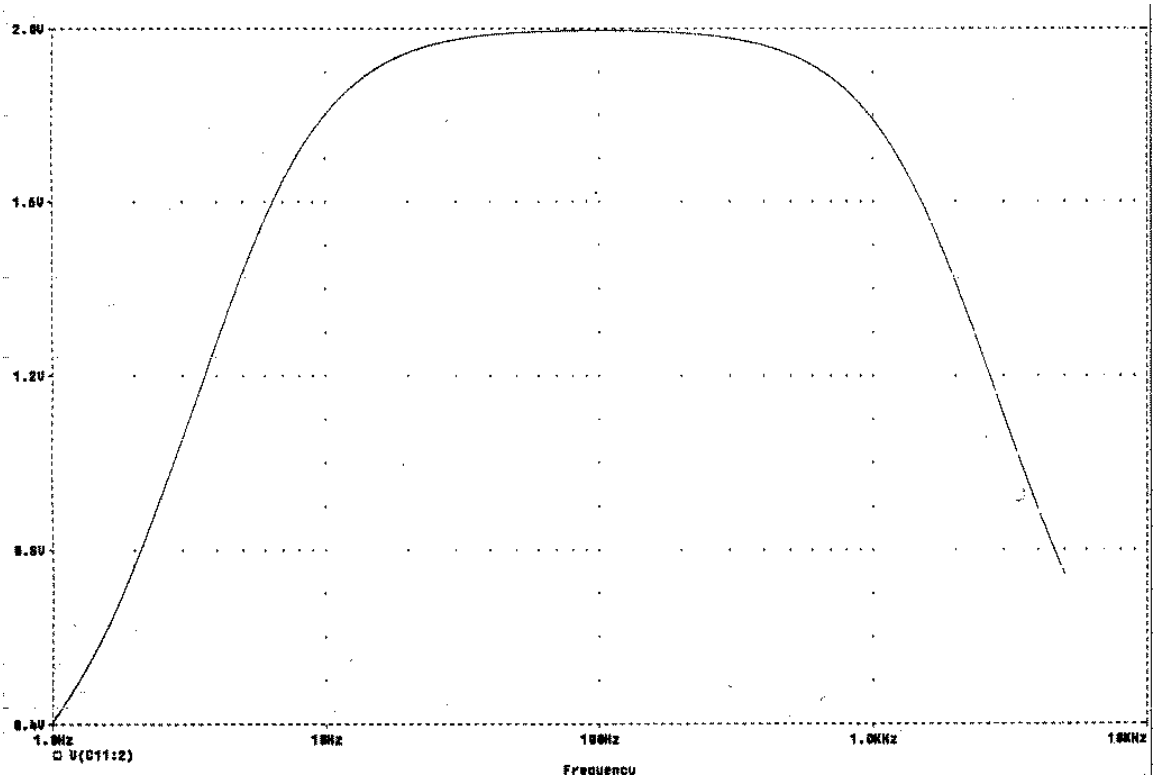

Figure C.3 ECG Filter Frequency Response (Hall 1999)

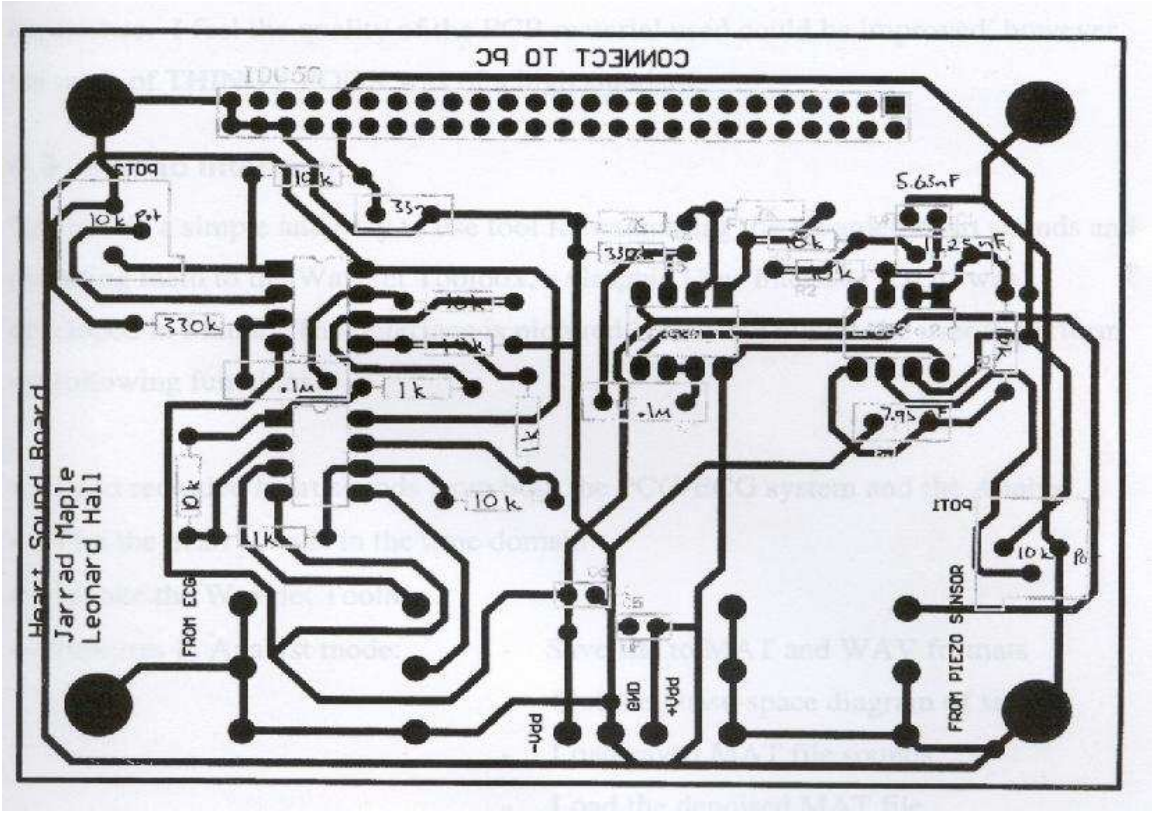

Figure C.4 ECG/PCG System PCB Layout (Hall 1999)

Appendix D

Moment of Velocity

We used the moment of velocity because we experienced problems with the frequent and distracting peaks in the instantaneous frequency plot when the denominator of the equation became very small. The moment of velocity is very similar to the instantaneous frequency and is defined as the numerator of the instantaneous frequency equation. It was thought that the moment of velocity would retain the main features portrayed by the instantaneous frequency while not having the large spikes caused by a small divisor. The

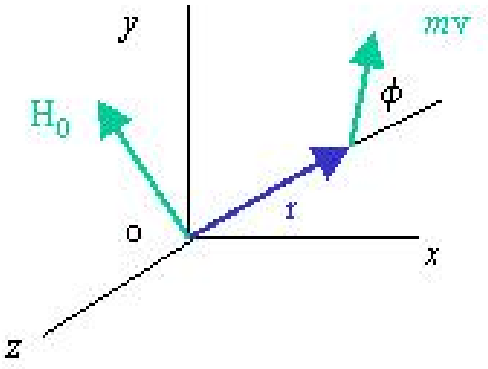

Figure D.1 Coordinate system used in context of defining angular momentum, modified from Beer & Johnston (1999)

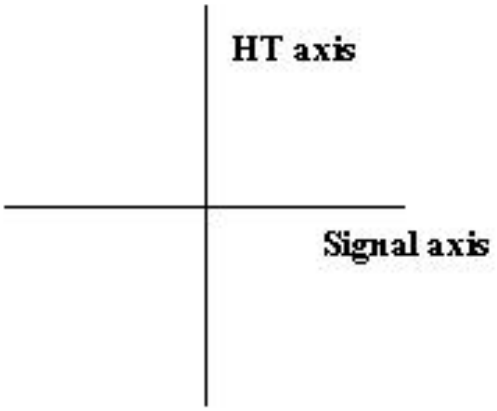

Figure D.2 Signal and Hilbert Transform analytic plane

moment of velocity is a concept borrowed from the the field of particle dynamics. The moment of velocity is derived from a commonly used quantity in particle physics called angular momentum or moment of momentum (further explained in Beer & Johnston (1999)) which is defined as

$$H_0 = \mathbf{r} \times m\mathbf{v} \tag{D.1}$$

where $\mathbf{r}$ is the position vector of $\mathbf{P}$ and H_0 is a vector perpendicular to the plane containing $\mathbf{r}$ and $m\mathbf{v}$ (See Figure D.1). If the vectors $\mathbf{r}$ and $m\mathbf{v}$ are resolved into components and the cross product formula is applied we obtain

$$\begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{j} \\ x & y & z \\ mv_x & mv_y & mv_z \end{vmatrix} \tag{D.2}$$

where the components of H_0 represent the moments of the linear momentum $m\mathbf{v}$ about

the coordinate axes. If the determinant is expanded we obtain

$$\begin{aligned} H_x &= m(yv_z - zv_y) \\ H_y &= m(zv_x - xv_z) \\ H_z &= m(xv_y - yv_x). \end{aligned} \tag{D.3}$$

So, if there exists a particle moving in the xy plane, $z = v_z = 0$ making H_x and H_z equal to zero. The angular momentum is perpendicular to the xy plane and is defined by the scalar

$$H_0 = H + z = m(xv_y - yv_x). \tag{D.4}$$

Because we are dealing with a signal and not a particle possessing mass, we use the moment of velocity which is simply the angular momentum without the mass included. The moment of velocity is defined as

$$\text{moment of velocity} = xv_y - yv_x \tag{D.5}$$

or in terms of the signal and the Hilbert Transform analytic plane (See Figure D.2) the moment of velocity is defined as

$$\text{moment of velocity} = s\frac{dH[s(t)]}{dt} - H[s(t)]\frac{ds}{dt}. \tag{D.6}$$

Bibliography

Akay, M. (1997), 'Wavelet applications in medicine', *IEEE Spectrum* **34**(5), 50–56.

Akay, M. & Szeto, K. H. (1995), 'Analyzing fetal breathing rates using matching pursuits', *IEEE Engineering in Medicine and Biology* **14**(2), 195–198.

Akay, Y. M., Akay, M., Welkowitz, W. & Kostis, J. (1994), 'Noninvasive detection of coronary artery disease', *IEEE Engineering in Medicine and Biology* pp. 761–764.

Antoine, J. P. (1999), *Wavelets in Physics*, Cambridge University Press, chapter Wavelet analysis: a new tool in physics, pp. 9–22.

Ashkenazy, Y. C., Lewkowicz, M., Levitan, J., Havlin, S., Saermark, K., Moelgaard, H. & Thomsen, P. B. (2000), 'Discrimination between healthy and sick cardiac autonomic nervous system by detrended heart rate variability analysis', http://xxx.lanl.gov/ps/chao-dyn/9810008.

Ashkenazy, Y., Lewkowicz, M., Levitan, J., Havlin, S., Saermark, K., Moelgaard, H., Thomsen, P. B., Moler, M., Hintze, U. & Huikuri, H. (1999), 'Scale specific and scale independent measures of heart rate variability as risk indicators', http://xxx.lanl.gov/ps/physics/9909029.

Baranek, H. L., Lee, H. C., Cloutier, G. & Durand, L. G. (1989), 'Automatic detection of sounds and murmurs in patients with Ionescu-Shiley aortic bioprostheses', *Medical and Biological Engineering and Computing* pp. 449–455.

Baykal, A., Ider, Y. Z. & Köymen, H. (1991), 'Use of signal averaging in analysis of the digital phonocardiograms', *Annual International Conference of the IEEE Engineering in Medicine and Biology Society* **13**(5), 2103–2104.

Beer, F. P. & Johnston, Jr., E. R. (1999), *Vector Mechanics for Engineers: Dynamics*, Third SI Metric edn, McGraw-Hill Ryerson.

Bell, T., Long, R., Langham, J., Kos, P. & Parten, M. (1998), Untethered patient monitoring, *in* 'Proceedings of the 11th IEEE Symposium on Computer-Based Medical Systems', IEEE, pp. 228–232.

Bentley, P., Grant, P. & McDonnell, J. (1998), 'Time-frequency and time-scale techniques for the classification of native and bioprosthetic heart valve sounds', *IEEE Transactions on Biomedical Engineering* **45**(1), 125–128.

Bertrand, O., Bohorquez, J. & Pernier, J. (1994), 'Time-frequency digital filtering based on an invertible wavelet transform: An application to evoked potentials', *IEEE Transactions on Biomedical Engineering* **41**(1), 77–88.

Beyar, R., Levkovitz, S., Braun, S. & Palti, Y. (1984), 'Heart-sound processing by average and variance calculation– physiologic basic and clinical implications', *IEEE Transactions on Biomedical Engineering* **31**(9), 591–596.

Blashkin, I. I. & Yakovlev, G. P. (1975), 'A contact microphone for phonocardiography', *Biomedical Engineering* **9**(3), 169–170.

Bogner, R. E. (1999), 'Hilbert Transforms', *Wiley Encyclopaedia of Electrical and Electronic Engineering* **9**, 83–94.

Bolton, R. J. (1983), Representation and pattern recognition of hilbert transformed electrocardiograms, Technical Report EE83/9, University of Queensland, Department of Electrical Engineering, Brisbane, Australia.

Borwick, J. (1990), *Microphones: technology and technique*, Focal Press, London and Boston.

Braun, S. & Feldman, M. (1997), 'Time-frequency characteristics of non-linear systems', *Mechanical Systems and Signal Processing* **11**(4), 611–620.

Brown, B. H., Smallwood, R. H., Barber, D. C., Lawford, P. V. & Hose, D. R. (1999), *Medical Physics and Biomedical Engineering*, Medical Science Series, Institute of Physics Publishing.

Buckheit, J. B. & Donoho, D. L. (1995), *Wavelets and Statistics*, Lecture Notes in Statistics, Spring-Verlag New York, Inc., chapter WaveLab and Reproducible Research, pp. 55–82.

Burrus, C. S., Gopinath, R. A. & Guo, H. (1998), *Introduction to Wavelets and Wavelet Transforms: A Primer*, Prentice-Hall, Inc., Upper Saddle River, New Jersey 07458.

Cardionics (1999), *Operator's Manual E-Scope Electronic Stethoscope, Part No. 718-7120*, Cardionics, Inc., 910 Bay Star Blvd., Webster, TX 77598, USA.

Carranza, R. & Andina, D. (2001), Medical wavelet-neural diagnosis in chagasic cardiopaties, *in* N. E. Mastorakis, ed., '2001 World Scientific Engineering and Engineering Society International Conferences: Neural Network and Applications', World Scientific and Engineering Society, pp. 3011–3014.

Carré, P., Leman, H., Fernandez, C. & Marque, C. (1998), 'Denoising of the uterine EHG by an undecimated wavelet transform', *IEEE Transactions on Biomedical Engineering* **45**(9), 1104–1113.

Cavalcanti, S. (2000), 'Arterial barareflex influence on heart rate variability: a mathematical model-based analysis', *Medical and Biological Engineering and Computing* **38**, 189–197.

Cheitlin, M. D., Sokolow, M. & McIlroy, M. B. (1993), *Clinical Cardiology*, sixth edn, Prentice-Hall International Inc.

Chen, D., Durand, L. G., Guo, Z. & Lee, H. C. (1997), 'Time-frequency analysis of the first heart sound. part 2: An appropriate time-frequency representation technique', *Medical and Biological Engineering and Computing* **35**, 311–323.

Chen, D., Durand, L. G. & Lee, H. C. (1997), 'Time-frequency analysis of the first heart sound. part 1: Simulation and analysis', *Medical and Biological Engineering and Computing* **35**, 306–310.

Chen, D., Durand, L. G., Lee, H. C. & Wieting, D. W. (1997), 'Time-frequency analysis of the first heartsound: Part 3: Application to dogs with varying cardiac contrac-

tility and to patients with mitral mechanical prosthetic heart valves', *Medical and Biological Engineering and Computing* **35**, 455–461.

Claerbout, J. F. (1976), *Fundamentals of geophysical data processing*, McGraw-Hill international series in earth and planetary sciences, McGraw-Hill, Inc.

Cloutier, G., Grenier, M., Guardo, R. & Durand, L. G. (1987), 'Spectral analysis of closing sounds produced by Ionescu-Shiley bioprosthetic aortic heart valves', *Medical and Biological Engineering and Computing* **25**, 492–496.

Cohen, L. (1989), 'Time-frequency distributions–a review', *Proceedings of the IEEE* **77**(7), 941–981.

Coifman, R. R. & Donoho, D. L. (1995), *Wavelets and Statistics*, Lecture Notes in Statistics, Spring-Verlag New York, Inc., chapter Translation-Invariant De-Noising, pp. 125–150.

Coifman, R. R., Meyer, Y., Quake, S. R. & Wickerhauser, M. V. (1992), Signal processing and compression with wavelet packets, *in* Y. Meyer & S. Roques, eds, 'Progress in Wavelet Analysis and Applications, Proc. Int. Conf. "Wavelets and Applications"', Editions Frontiéres, Fig-sure-Yvette, France, pp. 77–93.

Coifman, R. R., Meyer, Y. & Wickerhauser, V. (1992), *Wavelets and Their Applications*, Jones and Bartlett Publishers, Inc., Boston, MA, U.S.A, chapter Size Properties of Wavelet Packets, pp. 453–470.

Coifman, R. R. & Wickerhauser, M. V. (1998), Experiments with adapted wavelet de-noising for medical signals and images, *in* M. Akay, ed., 'Time frequency and wavelets in biomedical signal processing', IEEE, Inc., chapter 12, pp. 323–346.

Cozic, M., Durand, L.-G. & Guardo, R. (1998), 'Development of a cardiac acoustic mapping system', *Medical and Biological Engineering and Computing* **36**, 431–437.

Daubechies, I. (1988), 'Orthonormal bases of compactly supported wavelets', *Communications on Pure and Applied Mathematics* **41**, 909–996.

Daubechies, I. (1992), Ten lectures on wavelets, *in* 'CBMS-NSF regional conference series in applied mathematics', Philadelphia, U.S.A.

Daubechies, I. (1996), 'Where do wavelets come from?-a personal point of view', *Proceedings of the IEEE* **84**(4).

Davis, B. (2001), Personal email correspondence.

Debiais, F., Durand, L. G., Guo, Z. & Guardo, R. (1997), 'Time-frequency analysis of heart murmurs. Part II: Optimisation of time-frequency representations and performance evaluation', *Medical and Biological Engineering and Computing* **35**, 480–485.

Debiais, F., Durand, L. G., Pibarot, P. & Guardo, R. (1997), 'Time-frequency analysis of heart murmurs: Part 1: Parametric modelling and numerical simulations', *Medical and Biological Engineering and Computing* **35**, 474–479.

Dickhaus, H. & Heinrich, H. (1996), 'Classifying biosignals with wavelet networks (a method for noninvasive diagnosis)', *IEEE Engineering in Medicine and Biology* **15**(5), 103–111.

Dickhaus, H. & Heinrich, H. (1998), Analysis of ECG late potentials using time-frequency methods, *in* M. Akay, ed., 'Time frequency and wavelets in biomedical signal processing', IEEE, Inc., chapter 4, pp. 101–115.

Dickhaus, H., Khadra, L., Lipp, A. & Schweizer, M. (1992), Ventricular late potentials studied by nonstationary signal analysis, *in* 'Proceedings of the Annual International Conference of the IEEE: Vol. 2', Engineering in Medicine and Biology Society, pp. 490–491.

Dickhaus, H., Maier, C., Khadra, L. & Maayah, T. (1998), Time course of increasing nonlinear properties of heart rate variability after heart transplantation, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vol. 20, No 1', IEEE Engineering in Medicine and Biology Society. 334-337.

Dishan, H. (1996), 'A wavelet-based algorithm for the Hilbert Transform', *Mechanical Systems and Signal Processing* **10**(2), 125–134.

Djebbari, A. & Reguig, F. M. (2000), Short-time Fourier Transform analysis of the phonocardiogram signal, *in* 'The 7th IEEE International Conference on Electronics, Circuits and Systems', IEEE, pp. 844–847.

Donoho, D. L. (1995), 'De-noising by soft-thresholding', *IEEE Transactions of Information Theory* **41**(3), 613–627.

Donoho, D. L. & Jonstone, I. M. (1992*a*), Adapting to unknown smoothness by wavelet shrinkage, Technical report, Department of Statistics, Stanford University.

Donoho, D. L. & Jonstone, I. M. (1992*b*), Minimax estimation by wavelet shrinkage, Technical report, Department of Statistics, Stanford University.

Dranetz, A. I. & Orlacchio, A. W. (1976), *Shock and vibration handbook*, second edn, McGraw-Hill, New York, chapter Piezoelectric and piezoresistive pick-ups.

Durand, L. G., Blanchard, M., Cloutier, G., Sabbah, H. N. & Stein, P. D. (1990), 'Comparison of pattern recognition methods for computer-assisted classification of spectra of heart sounds in patients with a porcine bioprosthetic valve implanted in the mitral position', *IEEE Transactions on Biomedical Engineering* **37**(12), 1121–1129.

Durand, L. G., Genest, J. & Guardo, R. (1985), 'Modelling of the transfer function of the heart-thorax acoustic system in dogs', *IEEE Transactions on Biomedical Engineering* **32**, 592–601.

Durand, L. G. & Guardo, R. (1982), A model of the heart-thorax acoustic system, *in* Schwartz, ed., 'Applications of computers in medicine', IEEE Eng. Med. Biol. Soc., pp. 29–41.

Durand, L. G., Langlois, Y. E., Lanthier, T., Chiarella, R., Coppens, R., Carioto, S. & Bertrand-Bradley, S. (1990*a*), 'Spectral analysis and acoustic transmission of mitral and aortic valve closure sounds in dogs: Part 1 modelling the heart/thorax system', *Medical and Biological Engineering and Computing* **28**, 269–277.

Durand, L. G., Langlois, Y. E., Lanthier, T., Chiarella, R., Coppens, R., Carioto, S. & Bertrand-Bradley, S. (1990*b*), 'Spectral analysis and acoustic transmission of mitral

and aortic valve closure sounds in dogs: Part 3 effects of altering heart rate and p-r interval', *Medical and Biological Engineering and Computing* **28**, 431–438.

Durand, L. G., Langlois, Y. E., Lanthier, T., Chiarella, R., Coppens, R., Carioto, S. & Bertrand-Bradley, S. (1990*c*), 'Spectral analysis and acoustic transmission of mitral and aortic valve closure sounds in dogs: Part 4 effect of modulating cardiac inotropy', *Medical and Biological Engineering and Computing* **28**, 439–445.

Durand, L.-G. & Pibarot, P. (1995), 'Digital signal processing of the phonocardiogram: Review of the most recent advancements', *Critical Reviews in Biomedical Engineering* **23**(3/4), 163–219.

Durka, P. J. & Blinowska, K. J. (1998), In pursuit of time-frequency representation of brain signals, *in* M. Akay, ed., 'Time frequency and wavelets in biomedical signal processing', IEEE, Inc., chapter 15.

Dutt, D. N. & Krishnan, S. M. (1999), Application of phase space technique to the analysis of cardiovascular signals, *in* 'Proceedings of The First Joint BMES/EMBS Conference Serving Humanity, Advancing Technology', BMES/EMBS, p. 914.

Einstein, D., Kunzelman, K. S., Reinhall, P., Tapia, M., Thomas, R., Rothnie, C. & Cochran, R. P. (1999), Non-invasive determination of mitral valve acoustic properties: A proposed method to determine tissue alterations due to disease, *in* 'Proceedings of the First Joint BMES/EMBS Conference', BMES/EMBS, p. 183.

El-Asir, B., Khadra, L., Al-Abbasi, A. & Mohammed, M. M. J. (1996), Time-frequency analysis of heart sounds, *in* 'TENCON '96 Proceedings, Digital Processing Applications: Vol. 2', TENCON, pp. 553–558.

Ersoy, O. (1997), *Fourier-Related Transforms, Fast Algorithms and Applications*, Prentice-Hall, Inc., Upper Saddle River, New Jersey 07458.

Ewing, G. J. (1989), A new approach to the analysis of the third heart sound, Master's thesis, University of Adelaide.

Ewing, G., Mazumdar, J., Goldblatt, E. & Vollenhoven, E. V. (1985), Use of the maximum entropy method in the spectral analysis of the third heart sound in children, *in* 'Proc. 4th European Conference on Mechanocardiography'.

Ewing, G., Mazumdar, J., Vojdani, B., Goldblatt, E. & Vollenhoven, E. V. (1986), 'A comparative study of the maximum entropy method and the fast Fourier transform for the spectral analysis of the third heart sound in children', *Australian Physical and Engineering Sciences in Medicine* **9**(3), 117–126.

Feldman, M. (1994), 'Non-linear system vibration analysis using hilbert transform-i. free vibration analysis method 'freevib", *Mechanical Systems and Signal Processing* **8**(2), 119–127.

Ferguson, B. & Abbott, D. (2000), Signal processing for t-ray bio-sensor system, *in* 'Proceedings of SPIE: Smart Electronics and MEMS II, Vol. 4236', SPIE, pp. 157–169.

Gao, J., Dong, X., Want, W.-B., Li, Y. & Pan, C. (1999), 'Instantaneous parameters extraction via wavelet transform', *IEEE Transactions on Geoscience and Remote Sensing* **37**(2), 867–870.

Geckeler, G. D., Likoff, W., Mason, D., Riesz, R. R. & Wirth, C. H. (1954), 'Cardiospectrogram', *American Heart Journal* **48**, 189–196.

Gerbarg, D. S., Holcomb, Jr., F. W., Hofler, J. J., Bading, C. E., Schultz, G. L. & Sears, R. E. (1962), 'Analysis of phonocardiogram by a digital computer', *Circulation Research* **11**, 569–576.

Grenier, M., Gagnon, K., Genest, J., Durand, J. & Durand, L. (1998), 'Clinical comparison of acoustic and electronic stethoscopes and design of a new electronic stethoscope', *The American Journal of Cardiology* **81**, 653–656.

Groom, D. (1970), 'Standardisation of microphones for phonocardiography: the microphone pick-up', *Cardiology* **55**, 129–135.

Groom, D., Herring, O., Francis, W. & Shealy, G. (1956), 'The effect of background noise on cardiac auscultation', *American Heart Journal* **52**(5), 781–790.

Guo, Z., Moulder, C., Durand, L.-G. & Loew, M. (1998), Development of a virtual instrument for data acquisition and analysis of the phonocardiogram, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vol. 20, No. 1', IEEE Engineering in Medicine and Biology Society, pp. 436–439.

Hadjileontiadis, L. J. & Panas, S. M. (1998), 'A wavelet-based reduction of heart sound noise from lung noise', *International Journal of Medical Informatics* **52**(1-3), 183–190.

Hall, L. (1999), Private communication.

Havlin, S., Amaral, L., Ashkenazy, Y., Goldberger, A. L., Ivanov, P. C., Peng, C. K. & Stanley, H. (1999), 'Application of statistical physics to heartbeat diagnosis', *Physica A* **274**, 99–110.

Hearn, T. C., Gopal, D. N., Ghista, D. N., Robinson, J., Tihal, H., Mazumdar, J. & Bogner, R. (1983), Towards quantitative determination of valvular calcification and pressure drops across malfunctioning heart valves, *in* 'Proceedings of the 1983 Frontiers of Engineering and Computing in Health Care', IEEE, IEEE Publication, Columbus, Ohio, pp. 127–131.

Hök, B. (1991), 'Microphone design for bio-acoustic signals with suppression of noise and artifacts', *Sensors and Actuators A* **25**(527-533).

Hubbard, B. B. (1996), *The World According to Wavelets*, A K Peters Ltd., 289 Linden Street, Wellesley, MA 02181.

Hylen, J. C., Kloster, F. E., Herr, R. H., Starr, A. & Griswold, H. E. (1969), 'Sound spectrographic diagnosis of aortic ball variance', *Circulation* **39**, 849–858.

Instantaneous Frequency–A Seismic Attribute Useful in Structural and Stratigraphic Interpretation (2000), Technical report, Buereau of Economic Geology, The University of Texas at Austin, http://www.gri.org/.

Ivanov, P. (2000), Personal correspondence.

Ivanov, P. C., Amaral, L. A. N., Goldberger, A. L., Havlin, S., Rosenblum, M. G., Struzik, Z. & Stanley, H. E. (1999), Multifractality in human heartbeat dynamics.

Ivanov, P. C., Amaral, L. N., Goldberger, A. L., Havlin, S., Rosenblum, M. G., Struzik, Z. & Stanley, H. E. (2000), Beyond 1/f: Multifractality in human heartbeat dynamics, *in* D. Abbott & L. B. Kish, eds, 'Unsolved Problems of Noise and Fluctuations', American Institute of Physics, pp. 273–280.

Ivanov, P. C., Amaral, L. N., Goldberger, A. L. & Stanley, H. E. (1997), 'Stochastic feedback and the regulation of biological rhythms', http://xxx.lanl.gov/ps/cond-mat/9710325.

Ivanov, P. C., Bunde, A., Amaral, L. N., Havlin, S., Fritsch-Yelle, J., Baevsky, R. M., Stanley, H. E. & Goldberger, A. L. (1999), 'Sleep-wake differences in scaling behaviour of the human heartbeat: Analysis of terrestrial and long-term space flight data', *Europhysics Letters* .

Ivanov, P. C., Goldberger, A. L., Havlin, S., Peng, C. K., Rosenblum, M. G. & Stanley, H. E. (1999), *Wavelets in Physics*, Cambridge University Press, chapter Wavelets in medicine and physiology, pp. 391–419.

Ivanov, P. C., Rosenblum, M. G., Peng, C., Mietus, J., Havlin, S., Stanley, H. E. & Goldberger, A. L. (1996), 'Scaling behaviour of heartbeat intervals obtained by wavelet-based time-series analysis', *Nature* **383**(6598), 323–327.

Iwata, A., Boedeker, R. H., Dudeck, J., Pabst, W. & Suzumura, N. (1983), Computer aided analysis of phonocardiogram, *in* 'Medinfo 83 Proceedings of the Fourth World Conference on Medical Informatics', IFIP, pp. 569–572.

Iwata, A., Ishii, N., Suzumura, N. & Ikegaya, K. (1980), 'Algorithm for detecting the first and second heart sounds by spectral tracking', *Medical and Biological Engineering and Computing* **18**(1), 19–26.

Iwata, A., Suzumura, N. & Ikegaya, K. (1977), 'Pattern classification of the phonocardiogram using linear prediction analysis', *Medical and Biological Engineering and Computing* pp. 407–412.

Jamous, G., Durand, L., Langlois, Y., Lanthier, T., Pibarot, P. & Carioto, S. (1992), 'Optimal time-window duration for computing time/frequency representations of normal phonocardiograms in dogs', *Medical and Biological Engineering and Computing* **30**, 503–508.

Jandre, F. & Souza, M. (1997), Wavelet analysis of phonocardiograms: Differences between normal and abnormal heart sounds, *in* 'Proceedings of the 19th International Conference of IEEE/EMBS: Vol. 4', IEEE/EMBS, pp. 1642–1644.

Jiménez, A., Charleston, S., Peña, M., Aljama, T. & Ortiz, R. (1999), Performance of the Hilbert Transform in fetal phonocardiography for cardiotacogram generation, *in* 'Proceedings of the First Joint BMES/EMBS Conference', BMES/EMBS, p. 298.

Jiménez, A., Ortiz, M. R., Peña, M. A., Charleston, S., Aljama, A. T. & González, R. (1999), 'The use of wavelet packets to improve the detection of cardiac sounds from the fetal phonocardiogram', *Computers in Cardiology* **26**, 463–466.

Jiménez, A., Ortiz, M. R., Peña, M. A., Charleston, S., Aljama, A. T. & González, R. (2000), 'Performance of two adaptive subband filtering schemes for processing fetal phonocardiograms: Influence of the wavelet and the level of decomposition', *Computers in Cardiology* **27**, 427–430.

Jordan, M. D., Taylor, C. R., Nyhuis, A. W. & Tavel, M. E. (1987), 'Audibility of the fourth heart sound', *Archives of Internal Medicine* **147**, 721–726.

Karpman, L., Cage, J., Hill, C., Forbes, A. D., Karpman, V. & Cohn, K. (1975), 'Sound envelope averaging and the differential diagnosis of systolic murmurs', *American Heart Journal* **90**(5), 600–606.

Khadra, L., Abdallah, M. & Nashash, H. (1998), Neural wavelet analysis of life threatening ventricular arrhythmias, *in* 'Southeastcon '98 Proceedings IEEE', IEEE, pp. 228–229.

Khadra, L., Matalgah, M., El-Asir, B. & Mawagdeh, S. (1991), 'The wavelet transform and its applications to phonocardiogram signal analysis', *Med. Inform.* **16**(3), 271–277.

Kim, I. Y., Lee, S. M., Yeo, H. S., Han, W. T. & Hong, S. H. (1999), Feature extraction for heart sound recognition based on time- frequency analysis, *in* 'Proceedings of The First Joint BMES/EMBS Conference', BMES/EMBS, p. 960.

Kovacs, F. & Torok, M. (1998), An improved phonocardiographic method for fetal heart rate monitoring, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 20, No. 4', IEEE Engineering in Medicine and Biology Society, pp. 1719–1722.

Krishnan, S. & Rangayyan, R. M. (2000), 'Automatic de-noising of knee-joint vibration signals using adaptive time-frequency representations', *Medical and Biological Engineering and Computing* **38**, 2–8.

Lankhorst, M. M. & van der Laan, M. D. (1995), Wavelet-based signal approximation with genetic algorithms, *in* 'Proceedings of the Fourth Annual Conference on Evolutionary Programming', MIT Press, Cambridge, MA, pp. 237–255.

Leatham, A. (1975), *Auscultation of the heart and phonocardiography*, Churchill Livingstone.

Lehner, R. J. & Rangayyan, R. M. (1985), Microcomputer system for quantification of the phonocardiogram, *in* 'IEEE/Engineering in Medicine and Biology Society Annual Conference 1985', IEEE/EMBS, pp. 849–854.

Lehner, R. J. & Rangayyan, R. M. (1987), 'A three-channel microcomputer system for segmentation and characterization of the phonocardiogram', *IEEE Transactions on Biomedical Engineering* **34**(6), 485–489.

Letellier, C., Meunier-Guttin-Cluzel, S., Gouesbet, G., Neveu, F., Duverger, T. & Cousyn, B. (1997), Use of the nonlinear dynamical system theory to study cycle to cycle variations from spark ignition engine pressure data, *in* 'SAE Technical Paper Series 971640', pp. 36–46.

Liang, H., Lukkarinen, S. & Hartimo, I. (1998), 'A boundary modification method for heart sound segmentation algorithm', *Computers in Cardiology* **25**, 595–595.

Lim, L. A., Akay, M. & Daubenspeck, J. A. (1995), 'Identifying respiratory-related evoked potentials', *IEEE Engineering in Medicine and Biology* pp. 174–178.

Luisada, A. A. (1965), *From auscultation to phonocardiography*, The C.V. Mosby Company, St. Louis.

Luisada, A. A. (1972), *The sounds of the normal heart*, Warren H. Green, Inc.

Luisada, A. A. (1980), 'Sounds and pulses as aids to cardiac diagnosis', *Medical Clinics of North America* **64**, 3–32.

Luisada, A. A. & Bernstein, J. G. (1976), 'Better resolution and quantisation in clinical phonocardiography', *Cardiology* **47**, 113–126.

Luisada, A. A. & Gamma, G. (1954), 'Clinical calibration in phonocardiography', *American Heart Journal* **48**, 826–834.

Luisada, A. A., MacCanon, D. M., Coleman, B. & Feigen, L. P. (1971), 'New studies on the first heart sound', *American Journal of Cardiology* **28**, 140–149.

Luisada, A. A. & Zalter, R. (1960), 'A new standardized and calibrated phonocardiographic system', *IRE Transactions on Medical Electronics* **7**, 15–22.

Lukkarinen, S., Korhonen, P., Angerla, A., Nopanen, A.-L., Sikio, K. & Sepponen, R. (1997), Multimedia personal computer based phonocardiography, *in* 'Bridging Disciplines for Biomedicine, 18th Annual Conference of the IEEE: Vol. 5', 18th Annual Conference of the IEEE Engineering in Medicine and Biology Society, pp. 2303–2304.

Lukkarinen, S., Nopanen, A., Sikio, K. & Angerla, A. (1997), 'A new phonocardiographic recording system', *Computers in Cardiology* **27**, 117–120.

Lukkarinen, S., Sikio, K., Nopanen, A.-L., Angerla, A. & Sepponen, R. (1997), Novel software for real-time processing of phonocardiographic signal, *in* 'Proceedings of the 19th International Conference of the IEEE Engineering in Medicine and Biology Society', IEEE Engineering in Medicine and Biology Society, pp. 1455–1457.

Macovski, A. (1983), *Medical Imaging Systems*, Prentice Hall, Englewood Cliffs, N.J.

Mallat, S. (1989), 'A theory for multiresolution signal decomposition: the wavelet representation', *IEEE Pattern Analysis and Machine Intelligence* **11**(7), 674–693.

Mallat, S. (1999), *A wavelet tour of signal processing*, second edn, Academic Press, San Diego, CA, U.S.A.

Mallat, S. & Zhang, Z. (1993), 'Matching pursuits with time-frequency dictionaries', *IEEE Transactions on Signal Processing* **41**(12), 3397–3415.

Mannheimer, E. (1957), 'Standardization of phonocardiography', *American Heart Journal* **54**, 314–315.

Maple, J. (1999), Private communication.

Maple, J., Hall, L., Agzarian, J. & Abbott, D. (1999), 'Sensor system for heart sound biomonitor', *Proc. SPIE Electronics and Structures for MEMS* **3891**, 99–109.

Marieb, E. N. (1991), *Human Anatomy and Physiology*, second edn, The Benjamin/Cummings Publishing Company Inc.

Matalgah, M. & Knopp, J. (1994), Improved noise stability multiresolution filters based on a class of wavelets, *in* 'Proceedings of the 37th Midwest Symposium on Circuits and Systems: Vol. 2', IEEE, pp. 817–820.

Matalgah, M., Knopp, J. & Mawagdeh, S. (1998), Iterative processing method using Gabor wavelets and the wavelet transform for the analysis of phonocardiogram signals, *in* M. Akay, ed., 'Time frequency and wavelets in biomedical signal processing', IEEE, Inc., chapter 10, pp. 271–304.

McKusick, V. A. (1958), *Cardiovascular Sound in Health and Disease*, The Williams and Wilkins Company, Baltimore.

McKusick, V., Talbot, S. A., Webb, G. N. & Battersby, E. J. (1962), *Handbook of physiology: circulation*, Vol. I, American Society of Physiology, Washington, USA, chapter Technical aspects of the study of cardiovascular sound, pp. 681–694.

Meyer, Y. & Ryan, R. D. (1993), *Wavelets Algorithms and Applications*, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.

Misiti, M., Misiti, Y., Oppenheim, G. & Poggi, J.-M. (1996), *Wavelet Toolbox: For Use With Matlab*, The Math Works Inc.

Nandagopal, D. (1984), Phonocardiogram frequency analysis techniques and non-invasive determination of heart valve calcification, PhD thesis, Adelaide University.

Nandagopal, D., Bogner, R. E. & Mazumdar, J. (1980), Spectral analysis of second heart (SII) using linear predicition coding, *in* 'Proceedings of Second International Conference on Mechanics in Medicine and Biology', Osaka, Japan, pp. 102–103.

Nandagopal, D. & Mazumdar, J. (1981), 'A study of temporal variation in heart sound frequency spectra using Fast Fourier Transform', *Australasian Physical and Engineering Sciences in Medicine* **4**(2), 47–50.

Nandagopal, D., Mazumdar, J. & Bogner, R. E. (1984), 'Spectral analysis of second heart sound in normal children by selective linear prediction coding', *Medical and Biological Engineering and Computing* **22**, 229–239.

Nandagopal, D., Mazumdar, J., Karolyi, G. & Hearn, T. (1981), 'An instrumentation system and analysis procedure for phonocardiographic studies', *Australian Journal of Biomedical Engineering* pp. 16–20.

Obata, S., Yoshimura, S., Ide, H. & Mike, Y. (1971), 'Application of motional impedance to measurement of phonocardiography', *Japanese Journal of Medical Electronics and Biological Engineering* **9**(5), 344–350.

Obiadat, M. & Matalgah, M. (1992), Performance of the Short-Time Fourier Transform And Wavelet Transform to phonocardiogram signal analysis, *in* 'Proceedings of the 1992 ACM/SIGAPP Symposium on Applied Computing Volume II: Technological Challenges of the 1990's', ACM/SIGAPP, pp. 856–862.

Oppenheim, A. V., Schafer, R. W. & Buck, J. R. (1998), *Discrete Time Signal Processing*, Prentice Hall Signal Processing Series, second edn, Prentice-Hall, Inc., Upper Saddle River, New Jersey 07458.

Padmanabhan, V., Fischer, R., Semmlow, J. L. & Welkowitz, W. (1989), High sensitivity PCG transducer for extended frequency applications, *in* 'Images of the Twenty-First

Century. Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 1', IEEE/EMBS, IEEE, New York, USA, pp. 57–58.

Pan, Q., Zhang, L., Dai, G. & Zhang, H. (1999), 'Two denoising methods by wavelet transform', *IEEE Transactions on Signal Processing* **47**(12), 3401–3406.

Panter, P. F. (1965), *Modulation, Noise, and Spectral Analysis*, McGraw-Hill, Inc.

Passamani, E. (2000), '"heart" in microsoft encarta online encyclopedia 2000', http://encarta.msn.com.

Polikar, R. (2000), 'The wavelet tutorial', http://www.public.iastate.edu/ rpolikar/WAVELETS/waveletindex.html.

Potter, R. K. (1947), *Visible Speech*, Van Nostrand Company. Inc., New York.

Prochaázka, A., Mudrová, M. & Storek, M. (1998), *Signal Analysis and Prediction*, Applied and Numerical Harmonic Analysis, Birkhäuser, chapter Wavelet Use for Noise Rejection and Signal Modelling, pp. 215–226.

Qader, J. H. A., Khadra, L. M. & Dickhaus, H. (1999), Nonlinear dynamics in hrv signals after heart transplantations, *in* 'Fifth International Symposium on Signal Processing and its Applications', Signal Processing Research Centre, pp. 231–234.

Rabiner, L. R. & Gold, B. (1975), *Theory and Application of Digital Signal Processing*, Prentice Hall, Englewood Cliffs, NJ.

Rajan, S., Doraiswami, R., Stevenson, M. & Watrous, R. (1998), Wavelet based bank of correlators approach for phonocardiogram signal classification, *in* 'Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis', IEEE, pp. 77–80.

Randall, R. B. (1987), *Frequency Analysis*, third edn, Brüel and Kjaer, Denmark.

Rangayyan, R. M. & Lehner, R. J. (1988), 'Phonocardiogram signal analysis: A review', *CRC Critical Reviews in Biomedical Engineering* **15**(3), 211–236.

Raymond, B. (1999), Visualisation and pattern recognition of heart rate variability, PhD thesis, Adelaide University, Adelaide, South Australia.

Rice, M. L. & Doyle, D. J. (1995), Comparison of phonocardiographic monitoring locations, *in* 'IEEE Conference in Medicine and Biology 17th Annual Conference and 21 Canada Medical and Biological Engineering', IEEE, pp. 685–686.

Riesz, R. R. (1949), Low-frequency spectrography: Some applications in physiological research, *in* 'Proc A. I. E. E., I. R. E. Conference', New York, New York.

Rosenblum, M. G., Pikovsky, A. S., Schafer, C., Tass, P. & Kurths, J. (2000), Detection of phase synchronization from the data: Application to physiology, *in* D. Broomhead, E. Luchinskaya, P. McClintock & T. Mullin, eds, 'Stochastic and Chaotic Dynamics in the Lakes: STOCHAOS', American Institute of Physics. 154-161.

Rosenblum, M., Pikovsky, A., Schafer, C., Tass, P. A. & Kurths, J. (1999), Phase synchronization: from theory to data analysis.

Rushmer, R. H. (1952), *Cardiovascular dynamics*, W. B. Saunders Company.

Sava, H. & Durand, L.-G. (1997), Automatic detection of cardiac cycle based on an adaptive time-frequency analysis of the phonocardiogram, *in* 'Proceedings of the 19th International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 3', IEEE Engineering in Medicine and Biology Society, pp. 1316–1319.

Sava, H., Pibarot, P. & Durand, L. (1998), 'Application of the matching pursuit method for structural decomposition and averaging of phonocardiographic signals', *Medical and Biological Engineering and Computing* **36**, 302–308.

Sawada, Y., Ohtomo, N., Tanaka, V., Tanaka, G., Yamakoski, K., Terachi, S., Shimamoto, K., Nakagawa, M., Satoh, S., Kuroda, S. & Limura, O. (1997), 'New technique for time series analysis combining the maximum entropy method and non-linear least squares method: its value in heart rate variability analysis', *Medical and Biological Engineering and Computing* **35**, 318–322.

Schafer, C., Rosenblum, M. G., Abel, H.-H. & Kurths, J. (1999), 'Synchronization in the human cardiorespiratory system', *Physical Review E* **60**(1), 857–870.

Schafer, C., Rosenblum, M. G. & Kurths, J. (1998), 'Heartbeat synchronized with ventilation', *Nature* **392**(6673), 239–240.

Selig, M. B. (1993), 'Stethoscopic and phonoaudio devices: Historical and future perspectives', *American Heart Journal* **126**(1), 262–268.

Senhadji, L., Carrault, G., Bellanger, J. & Passariello, G. (1995), 'Comparing wavelet transforms for recognizing cardiac patterns', *IEEE Engineering in Medicine and Biology* **14**(2), 167–173.

Shen, M. & Sun, L. (1997), The analysis and classification of phonocardiogram based on higher-order spectra, *in* 'Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics', IEEE, pp. 29–33.

Shino, H., Yoshida, H., Mizuta, H. & Yana, K. (1997), Phonocardiogram classification using time-frequency representation, *in* 'Proceedings of the 19th International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 4', IEEE Engineering in Medicine and Biology Society, pp. 1636–1637.

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

Sukimura, A. K. I. N. & Funada, T. (1971), 'Absolute calibration of phonocardiographic microphones and measurement of chest wall vibration', *Medical and Biological Engineering* **9**, 683–692.

Sun, M. & Sclabassi, R. J. (1998), Wavelet feature extraction from neurophysiological signals, *in* M. Akay, ed., 'Time frequency and wavelets in biomedical signal processing', IEEE, Inc., chapter 11.

Suzumura, N. & Ikegaya, K. (1977), 'Characteristics of the air cavities of phonocardiographic microphones and the effects of vibration and room noise', *Medical and Biological Engineering and Computing* **15**(3), 240–247.

Tahmasbi, M. S. (1994), VLSI implementation of heart sounds maximum entropy spectral estimation, Master's thesis, University of Adelaide.

Takagi, S. & Yoshimura, S. (1964), 'Trial and success in the technical realization of requirements for phonocardiographic microphones', *Medical and Biological Engineering* **2**, 123–134.

Taner, M. T., Koehler, F. & Sheriff, R. E. (1979), 'Complex seismic trace analysis', *Geophysics* **44**(6), 1041–1063.

Tavel, M. E., Brown, D. D. & Shander, D. (1994), 'Enhanced auscultation with a new graphic display system', *Archives of Internal Medicine* **154**(8), 893.

The American Heritage Dictionary of the English Language (2000), fourth edn, Houghton Mittlin Company.

Thurner, S., Feurstein, M. C., Lowen, S. B. & Teich, M. C. (1998), 'Receiver-operating-characteristic analysis reveals superiority of scale-dependent wavelet and spectral measures for assessing cardiac dysfunction', *Physical Review Letters* **81**(25), 5688–5691.

Thurner, S., Feurstein, M. C. & Teich, M. C. (1998), 'Multiresolution wavelet analysis of heartbeat intervals discriminates healthy patients from those with cardiac pathology', *Physical Review Letters* **80**(7), 1544–1547.

Tinati, M. A. (1998), Time-Frequency and Time-Scale Analysis of Phonocardiograms with Coronary Artery Disease Before and After Angioplasty, PhD thesis, The University of Adelaide.

Tinati, M. A., Bouzerdoum, A. & Mazumdar, J. (1996), Modified adaptive line enhancement filter and application to heart sound noise cancellation, *in* 'Fourth International Symposium on Signal Processing and its Applications. ISSPA 96. Vol. 2', pp. 815–818.

Tkacz, E., Kostka, P. & Komorowski, D. (1998), An application of wavelet transform and adaptive filters for decomposition of the hrv signals in the case of patients with coronary artery disease, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 20, No. 6', IEEE Engineering in Medicine and Biology Society, pp. 3120–3122.

Torres-Pereira, L., Ruivo, P., Torres-Periera, C. & Couto, C. (1997), A non-invasive telemetric heart rate monitoring system based on phonocardiography, *in* 'Proceedings of the IEEE International Symposium on Industrial Electronics: Vol. 3', IEEE, pp. 856–859.

Tovalr-Corona, B. & Torry, J. N. (1997), 'Graphical representation of heart sounds and murmurs', *Computers in Cardiology* **24**, 101–104.

Tovar-Corona, B. & Torry, J. N. (1998), 'Time-frequency representation of systolic murmurs using wavelets', *Computers in Cardiology* **25**, 601–604.

Unser, M. & Aldroubi, A. (1996), 'A review of wavelets in biomedical applications', *Proceedings of the IEEE* **84**(4), 626–638.

Van Vollenhoven, E. (1975), Physical problems in calibration of microphones for phonocardiography, *in* 'International Conference on Biomedical Transducers', Electronic Industries Association, pp. 220–222.

Van Vollenhoven, E. & Wallenburg, J. (1970), 'Calibration of air microphones for phonocardiophy', *Medical and Biological Engineering* **8**(3), 309–313.

Van Vollenhoven, E., Wallenburg, J., Van Rotterdam, A. & Van Straaten, J. (1968), 'Calibration of contact microphones for phonocardiography', *Medical and Biological Engineering and Computing* **6**, 71–82.

Vermarien, H. & van Vollenhoven, E. (1984), 'The recording of heart vibrations: a problem of vibration measurement on soft tissue', *Medical and Biological Engineering and Computing* **22**, 168–178.

Vermarient, H. & Van Vollenhoven, E. (1983), 'Influence of design parameters of piezoelectric contact microphones on recording heart sounds', *Acta Cardiology* **38**, 290–292.

Vertan, C., Vertan, C. I. & Buzuloiu, V. (1997), Reduced computation genetic algorithm for noise removal, *in* 'IEE Conference Publication, n 443, pt 1', IEE, pp. 313–316.

Vinson, M., Khadra, L., Maayah, T. & Dickhaus, K. (1995), Detecting chaotic behaviour in hrv signals in a human chardiac transplant recipient, *in* 'International Confer-

ence on Acoustics, Speech and Signal Processing 1995: Volume 2', ICASSP IEEE, pp. 1364–1367.

Wachowiak, M. P., Rash, G. S., Quesada, P. M. & Desoky, A. H. (2000), 'Wavelet-based noise removal for biomechanical signals: a comparative study', *IEEE Transactions on Biomedical Engineering* **47**(3), 360–368.

Welch, P. D. (1967), 'The use of the Fast Fourier Transform for the estimation of power spectral: A method based on time averaging over short, modified periodograms', *IEEE Transactions on Audio and Electroacoustics* **15**, 70–73.

Whitmal, N. A., Rutledge, J. C. & Cohen, J. (1996), 'Reducing correlation noise in digital hearing aids', *IEEE Engineering in Medicine and Biology* pp. 88–96.

Wickerhauser, M. V. (1994), *Adapted wavelet analysis from theory to software*, A. K. Peters, Ltd.

Winder, D. E., Perry, L. W. & Caceres, C. A. (1965), 'Heart sound analysis: A three dimensional approach', *The American Journal of Cardiology* **16**, 547–551.

Wood, J. C. & Barry, D. T. (1994), 'Quantification of first heart sound frequency dynamics across the human chest wall', *Medical and Biological Engineering and Computing Electrocardiography, Myocardial Contraction and Blood Flow supplement* pp. S71–S78.

Wood, J. C. & Barry, D. T. (1995), 'Time-frequency analysis of the first heart sound', *IEEE Engineering in Medicine and Biology* pp. 144–151.

Wood, J. C., Buda, A. J. & Barry, D. T. (1992), 'Time-frequency transforms: A new approach to first heart sound frequency dynamics', *IEEE Transactions on Biomedical Engineering* **39**(7), 730–740.

Xiao, Y., Xiao, S., Cao, Z., Zhou, S. & Pei, J. (1999), 'The phonocardiogram exercise test', *IEEE Engineering in Medicine and Biology* **18**(4), 111–115.

Yang, H., Qiu, L. & Koh, S.-N. (1994), 'Application of instantaneous frequency estimation for fundamental frequency detection', *Nanyang Technological Univ Source: Proceed-*

ings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis pp. 616–619.

Yoganathan, A. P. (1976), 'Fast Fourier transform in the analysis of biomedical data', *Medical and Biological Engineering* pp. 239–244.

Yoganathan, A. P., Gupta, R., Corcoran, W. H., Udwadia, F. E., Sarma, R. & Bing, R. J. (1976), 'Use of the fast Fourier transform in the frequency analysis of the second heart sound in normal man', *Medical and Biological Engineering* **14**(4), 455–460.

Yoganathan, A. P., Gupta, R., Udwadia, F., Miller, J. W., Corcoran, W. H., Sarma, R., Johnson, J. L. & Bing, R. J. (1976), 'Use of the fast fourier transform for frequency analysis of the first heart sound in normal man', *Medical and Biological Engineering* **14**(1), 69–73.

Zehan, C., Shiyong, Z., Li, F., Yuli, P. & Shouzhong, X. (1998), The quantitative analysis approach for a heart sound information system, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Vol. 20, No. 3', IEEE Engineering in Medicine and Biology Society, pp. 1620–1621.

Zhang, X., Durand, L.-G., Senhadji, L., Lee, H. C. & Coatrieux, J.-L. (1998), 'Analysis-synthesis of the phonocardiogram based on the matching pursuit method', *IEEE Transactions on Biomedical Engineering* **45**(8), 962–971.

Zhang, X., Durrand, L.-G., Senhadji, L., Lee, H. C. & Coatrieux, J.-L. (1998), 'Time-frequency scaling transformation of the phonocardiogram based of the matching pursuit method', *IEEE Transactions on Biomedical Engineering* **45**(8), 972–979.

Zhou, P. & Wang, Z. (1998), A computer location algorithm for ECG, PCG and CAP, *in* 'Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vol. 20, No. 1', IEEE Engineering in Medicine and Biology Society, pp. 220–222.