Microelectronics Journal 32 (2001) 931–941

Microelectronics Journal

www.elsevier.com/locate/mejo

Optimal wavelet denoising for phonocardiograms

Sheila R. Messer[a], John Agzarian[b], Derek Abbott[a,*]

[a]Department of Electrical and Electronic Engineering, Centre for Biomedical Engineering (CBME), Adelaide University, Adelaide, SA 5005, Australia
[b]Hampstead Medical Clinic, Hampstead Gardens, SA 5086, Australia

Revised 15 June 2001; accepted 21 August 2001

Abstract

Phonocardiograms (PCGs), recordings of heart sounds, have many advantages over traditional auscultation in that they may be replayed and analysed for spectral and frequency information. PCG is not a widely used diagnostic tool as it could be. One of the major problems with PCG is noise corruption. Many sources of noise may pollute a PCG including foetal breath sounds if the subject is pregnant, lung and breath sounds, environmental noise and noise from contact between the recording device and the skin. An electronic stethoscope is used to record heart sounds and the problem of extracting noise from the signal is addressed via the use of wavelets and averaging. Using the discrete wavelet transform, the signal is decomposed. Due to the efficient decomposition of heart signals, their wavelet coefficients tend to be much larger than those due to noise. Thus, coefficients below a certain level are regarded as noise and are thresholded out. The signal can then be reconstructed without significant loss of information in the signal content. The questions that this study attempts to answer are which wavelet families, levels of decomposition, and thresholding techniques best remove the noise in a PCG. The use of averaging in combination with wavelet denoising is also addressed. Possible applications of the Hilbert transform to heart sound analysis are discussed. © 2001 Published by Elsevier Science Ltd.

Keywords: Wavelets; Denoising; Phonocardiogram; Heart sound analysis; Heartbeat analysis; Hilbert transform

1. Introduction

The stethoscope and human ear have their limitations in diagnosing heart defects and conditions. Modern technology has developed new tools, which are capable of revealing information that traditional tools such as the stethoscope alone cannot. For example, digital stethoscopes have been developed, which have the capacity to record and replay the heartbeat sound recordings, otherwise known as phonocardiograms (PCGs). The PCG is a particularly useful diagnostic tool because the graphic recordings show timings and relative intensities of heartbeat sounds and may reveal information that the human ear cannot [1–3]. With the aid of computers, the PCG data may be stored, managed, and manipulated for frequency and spectral content.

Electrocardiograms (ECGs), which reveal the electrical activity of the heart, and echocardiography, which uses ultrasound waves to create an image of the heart, are other methods that are used to assess the condition the heart. All of the mentioned techniques are non-invasive, but each method presents a separate set of issues and challenges.

With the advent of echocardiography, auscultation and phonocardiography have become less important. However, with the aid of computers and digital signal processing techniques, PCGs may reveal important information [1].

The ECG is not normally used unless a problem has been previously detected by auscultation (listening to the heart sounds), because the time required to set up an ECG recording is longer and hence not used as a standard test. However, PCGs are easily obtained by placing the stethoscope against the skin. The current problem with many PCG systems is noise from sources such as breath sounds, contact of the stethoscope with the skin, foetal heart sounds if the subject is pregnant, and ambient noise that may corrupt the heart sound signals.

The PCG would be a much more useful diagnostic tool if unwanted noise was removed revealing the heartbeat sound clearly. The current study examines methods of removing the noise from the PCG using namely wavelet analysis and averaging. Currently, there is no way of knowing a priori what the particular noise component is, or of determining the noise component once the measurement has been recorded. The previously mentioned sources of noise are by no means a complete list. In every case and situation, the noise will be different.

Although what produces each of the four heart sounds

* Corresponding author. Tel.: +61-8-8303-5748; fax: +61-8-8303-4360.
E-mail addresses: smesser@eleceng.adelaide.edu.au (S.R. Messer), dabbott@eleceng.adelaide.edu.au (D. Abbott).

0026-2692/01/$ - see front matter © 2001 Published by Elsevier Science Ltd.
PII: S0026-2692(01)00095-7

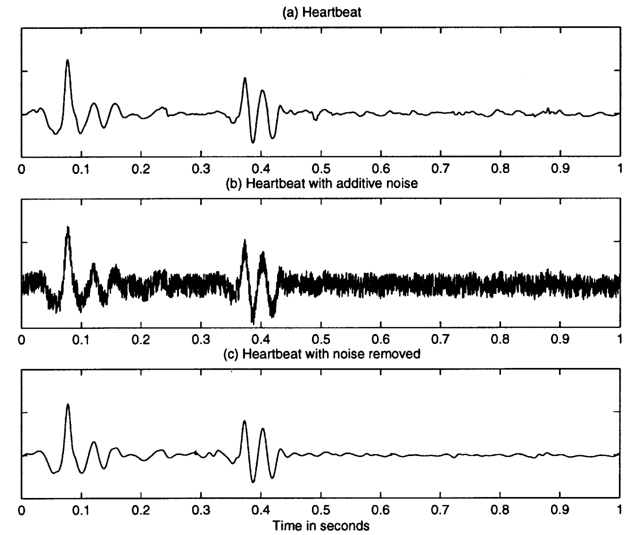

Fig. 1. This figure shows the principle of denoising a heartbeat. (a) Is the characteristic heartbeat signal (b) the heartbeat with 1 dB of additive white noise (c) the heartbeat with noise removed using wavelet analysis.

heard through a stethoscope is not exactly known, they are likely to be produced by a number of sources including the opening and closing of valves, vibrations of the cardiac structure, and acceleration and deceleration of blood [1–4]. The first and second heart sounds ('lub-dub') are the two that are generally heard by the human ear, and are usually the most visible on the PCG. They may be seen in Fig. 1(a) with the first heart sound, S1, occurring at about 0.1 s and the second heart sound, S2, happening at about 0.4 s.

An electronic stethoscope is used to record heart sounds. Various digital signal processing tools namely wavelets are employed to remove noise from signal. The remainder of this article will introduce the reader to basic wavelet theory in relation to heart sounds and the methods of denoising heart sounds, which were investigated.

Wavelets may be used to denoise the PCG as shown in Fig. 1. The signal is decomposed by a discrete wavelet transform. Because of the efficient decomposition of heart signals, their wavelet coefficients tend to be much larger than those due to noise. Thus, coefficients below a certain level are regarded as noise and thresholded out. The signal is then reconstructed without significant loss of information. The questions that this study attempts to answer are, which wavelet families, levels of decomposition, and thresholding techniques best remove the noise in a PCG. This study is an extension of the work described by Maple et al. [5].

Averaging may also be used to produce a characteristic heartbeat [4]. Although heartbeats are considered non-stationary signals, they are periodic in the sense that heartbeats regularly repeat. 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. This study examines the use of averaging in combination with wavelet analysis to remove noise from a PCG.

2. Equipment

An electronic stethoscope (the Escope from Cardionics), using an electret microphone outputs the heart sound as an analogue signal. This analogue signal is converted to a digital signal (sampled at 2500 samples/s with 12-bit resolution) and stored on the computer for further use. The electrical activity of the heart is also simultaneously recorded to serve as a reference signal. MATLAB software is then used to analyse the signal and perform the signal processing. For more detail on the equipment and recording procedure, refer to our previous study [6].

3. Wavelet theory

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

3.1. Fourier analysis

In 1822, Joseph Fourier discovered that any periodic function could be represented as an infinite sum of periodic complex exponential functions [7]. 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 expressed in the frequency domain. The FT representation of a signal may be seen in Fig. 2(a). 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

$$X(f) = \int_{+\infty}^{-\infty} x(t)\mathrm{e}^{-j2\pi ft}\mathrm{d}t \tag{1}$$

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.

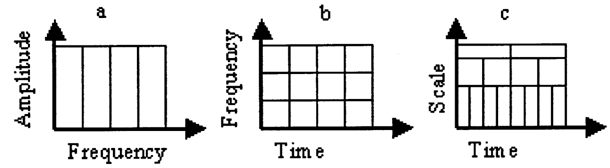

Fig. 2. Comparison of a signal represented in different domains with (a) corresponding to the Fourier transform representation, (b) representing the short time Fourier transform, and (c) the wavelet transform.

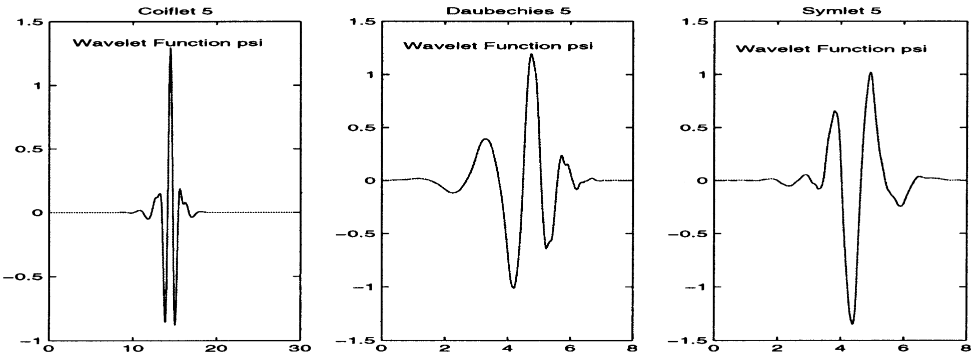

Fig. 3. Examples of wavelets used in this study.

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 [8], "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 localized in time. Most real world signals exhibit non-stationary characteristics (such as heart sound signals). Thus, Fourier analysis is not adequate.

3.2. Short time Fourier transform

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 Eq. (1). Thus, the frequency content of the signal is known but its location in time is not known.

In an effort to answer this problem, the STFT was developed in 1946 by Denis Gabor [9]. 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 on the frequency content, when it occurs. The trade-off 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 Fig. 2(b). The mathematical representation of STFT is

$$\mathrm{STFT}_x^{(w)}(t',x) = \int_t [x(t)w^*(t-t')]\mathrm{e}^{-j2\pi ft}\mathrm{d}t \tag{2}$$

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 Eq. (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 localization in time becomes worse.

3.3. The wavelet transform

The Wavelet transform (WT) 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 [9], although the progress in the field of wavelets has been relatively slow until the 1980's when scientists and engineers from different fields realized they were working on the same concept and began collaborating [8].

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 Fig. 2(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.

A wavelet is a signal of limited duration that has an average value of zero. Examples of wavelets used in this study may be seen in Fig. 3.

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

$$\mathrm{CWT}_x^{\Psi}(\tau,s) = \Psi_x^{\psi}(\tau,x) = \frac{1}{\sqrt{|s|}}\int x(t)\psi^*\left(\frac{t-\tau}{s}\right)\mathrm{d}t \tag{3}$$

The scale, s, of the wavelet may conceptually be considered the inverse of frequency. As seen in Fig. 2(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.

The process of computing the CWT is very similar to that of the STFT. The wavelet is compared to a section at the

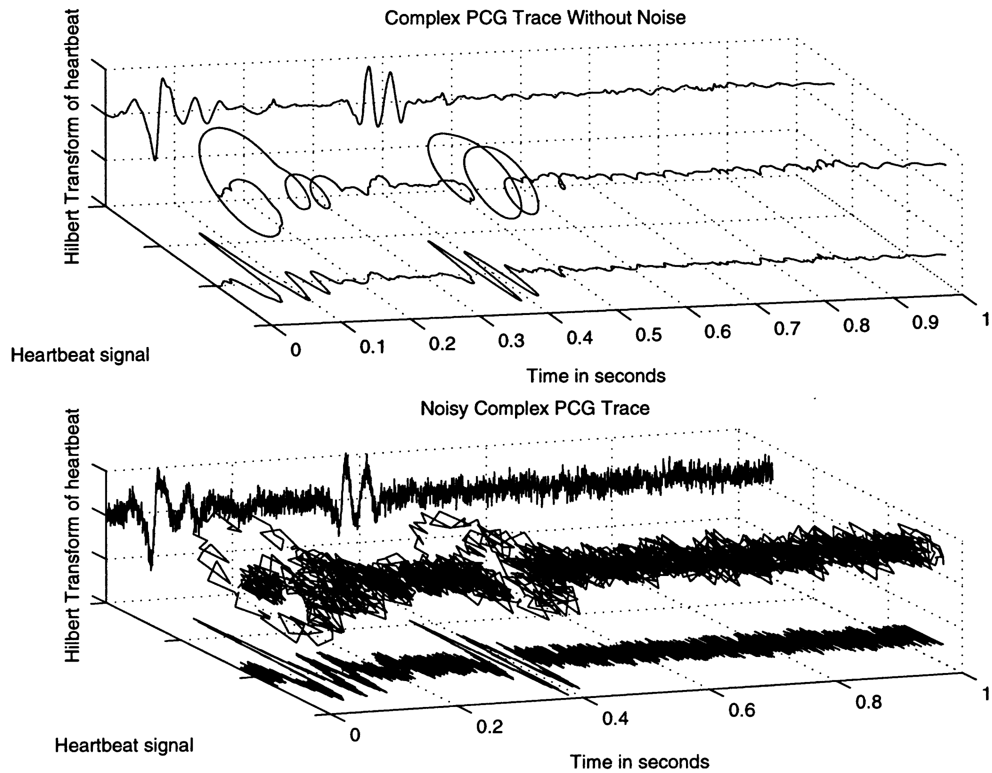

Fig. 4. This figure shows a complex PCG trace first with additive white noise and secondly without noise. The x-axis represents time, the y-axis represents the PCG signal and the z-axis represents the HT of the PCG.

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 the process is repeated until the whole signal is covered. The wavelet is scaled and the previous process is repeated for all scales.

The CWT reveals more 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.

4. Averaging theory

Averaging is known to reduce white noise because it is randomly distributed throughout the signal and may also be used to produce a 'characteristic heartbeat' which is an averaged heartbeat from a series of beats [4]. 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 [4].

According to basic probability theory [10], 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 heartbeat series 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 denoise heartbeats is the type of signal sought. The mechanical activity of the heart can be classified into two categories: 'deterministic' and 'nondeterministic' [10]. In our case, any process that repeats itself exactly for each beat may be considered deterministic. For non-deterministic events (e.g. such as 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 denoising can offer an advantage over simple averaging.

The algorithm for averaging the PCG signal uses the ECG as a gating signal because they are both recorded simultaneously. 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 [4].

5. Hilbert transform theory

The Hilbert transform (HT) may be used to calculate the instantaneous attributes of a signal. The mathematical definition of the Hilbert transform is [11]

$$y(t) = \pi^{-1} \int_{\infty}^{-\infty} \frac{x(\tau)}{t - \tau} \, \mathrm{d}\tau \tag{4}$$

The Hilbert transform (HT) can be considered a convolution between the signal and $(1/\pi)t$. 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 $(\pi/2)t$.

Fig. 4 borrows the concept of a complex trace from seismic data analysis [12]. The signal and its HT 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 which may be seen in Fig. 5. The instantaneous frequency is calculated through the analytic method using the HT. The instantaneous frequency is mathematically defined in Eq. (5) where s is the signal and $H(s)$ is the HT of the signal [13].

$$\Theta = \frac{1}{2\pi} \frac{d}{\mathrm{d}t} \left[\arctan\left(\frac{H[s(t)]}{s(t)} \right) \right] \tag{5}$$

The instantaneous frequency of a signal may be used to demonstrate the effectiveness of denoising [14].

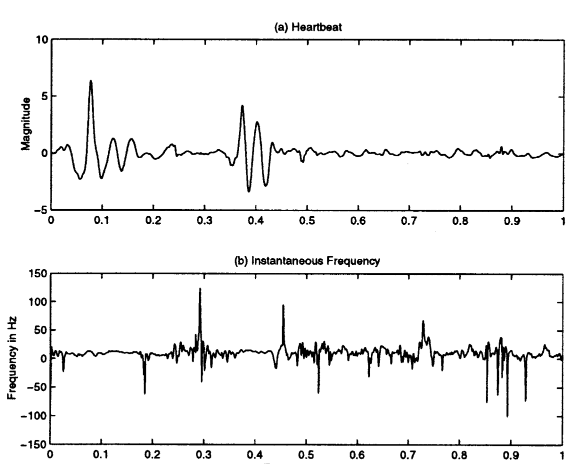

Fig. 5. This figure shows (a) a characteristic heartbeat as a reference (b) the instantaneous frequency of a heartbeat. It is interesting to note that at the first and second heart sound the dominant frequency appears to be about 10 Hz. There are artifacts present outside the principal areas of interest (S1 and S2) due to computational rounding errors.

6. Denoising of heartbeats

6.1. Signal decomposition and reconstruction

The MATLAB wavelet toolbox provides built in routines for using the DWT to decompose a signal into wavelet coefficients and then to reconstruct the signal using the inverse discrete wavelet transform (IDWT). Many wavelet families are available in the toolbox. However, in the current study only orthogonal wavelets are examined because they allow for perfect reconstruction of a signal. The process used by MATLAB in decomposing and reconstructing a signal applies a series of highpass and lowpass filters in succession. This procedure was developed by Mallat in 1988, and is a well known signal processing tool called the two-channel subband coder [15]. The signal is passed through highpass and lowpass filters and downsampled to keep the original number of datapoints. This procedure results in details, which are low scale, high frequency elements of the signal, and approximations, which are high scale, low frequency elements of the signal. This decomposition can be performed for many levels with the decomposition process being iterated for successive approximations [9,16].

6.2. White noise removal by thresholding

A classic problem with most noise removal methods is actually knowing the noise content from a time series and how to classify the signal. With wavelet denoising, it is not necessary to know which part of the signal is white noise. David Donoho and Iain Johnstone proved mathematically that if a special kind of orthogonal basis existed, then it would do the best job at removing white noise from a signal. Then, in 1990, Donoho conversed with Domonique Picard and Gerard Kerkyachaian and realized that they had found the answer to the question of white noise removal using wavelets [8]. The WT is applied to the signal and all coefficients below a certain size are discarded [8]. This technique makes use of the fact that some of the decomposed wavelet coefficients correspond to signal averages and others are associated with details on the original signal. If the smaller details are eliminated from the signal decomposition, the original signal can be extracted from the remaining coefficients and the main signal characteristics will remain intact because an orthogonal wavelet transform compresses the 'energy' of the signal into a few large components and the white noise is very disordered and hence it is scattered throughout the transform in small coefficients [8,17].

6.3. Optimal parameter selection for wavelet denoising of PCGs

When using wavelets to denoise PCGs, there are many factors that must be considered. Examples of these choices are which wavelet, level of decomposition, and thresholding methods to use.

6.3.1. Choosing the wavelets to use

In general, the more a wavelet resembles the signal, the better it denoises the signal. However, because it is very difficult and time-consuming to actually create a useful wavelet, the wavelets which will be considered for use will be previously defined.

MATLAB provides several families of wavelets including the Morlet, Mexican hat, Meyer, Haar, Daubechies, Symlets, Coiflets and Spline biorthogonal wavelets and provides further documentation about these wavelet families [9]. In order to obtain perfect reconstruction results, only orthogonal wavelets will be considered. The orthogonal wavelet transform[1] has certain benefits. It is very concise, allows for perfect reconstruction of the original signal and is not very difficult to calculate [9]. The ease of calculation exists because each wavelet coefficient is calculated with only one scalar product of the signal and the wavelet [8]. Being computationally inexpensive, orthogonal wavelets allow a transform to be computed which contains the same number of points as the original signal [8]. Wavelets that have the properties of orthogonality and a scaling function ϕ allow the use of the fast algorithm. The wavelets which satisfy these criteria are the Haar, Daubechies, Coiflets, and Symlets.

In MATLAB, the Daubechies family of wavelets consists of 45 wavelets. The Haar wavelet is the first and most simple wavelet in this family. The Daubechies family of wavelets is not explicitly mathematically defined except for the Haar wavelet. Most Daubechies wavelets are not symmetrical.

The Symlet family of 45 wavelets is more symmetrical than the Daubechies but is still not exactly symmetric. Ingrid Daubechies modified her family of wavelets so that their symmetry would be increased while they remained simple. Only the first 15 wavelets in this family were examined because their increasing complexity requires much more computation time. For example, on the same computer under similar conditions, the denoising of the same sample with a Daubechies 15 wavelet took about.27 s and a Symlet 15 took about 1.27 s and using a Daubechies 25 wavelet it took about 0.50 s of computation time but using a Symlet 25 wavelet it increased to 54.38 s.

The Coiflet family of 5 wavelets was built by Daubechies at the request of Coifman. These wavelets are more symmetrical than the Daubechies wavelets.

For each decomposition level, the signal is passed through quadrature filter banks resulting in approximations and details. This decomposition can be performed for many levels with the decomposition process being iterated for successive approximations. For every wavelet, decomposition levels of N from 1 to 10 were tested.

6.3.2. Choosing threshold parameters

The two common methods of thresholding a signal are soft thresholding and hard thresholding which are used in the MATLAB wavelet toolbox [9]. The definitions of the two methods of thresholding are given below where x_0 represents the threshold and x denotes signal [9].

$$\text{Hard Thresholding: thresholded signal} = x \text{ for } |x| > x_0, \text{ and } 0 \text{ for } |x| \le x_0$$

$$\text{Soft Thresholding: thresholding signal} = \text{sign}(x)(|x| - x_0) \text{ for } |x| > x_0, \text{ and } 0 \text{ for } |x| \le x_0$$

In other words, hard thresholding is setting the elements whose absolute values are less than the threshold to zero. For soft thresholding, the elements whose absolute values are lower than the threshold are set to zero, and then the nonzero coefficients remaining are shrunk towards zero.

Although hard thresholding is the simplest method, soft thresholding can produce better results than hard thresholding. The reason for this is that hard thresholding may cause discontinuities at $\pm x_0$ because those values less than the threshold are set to zero [9].

There are four threshold selection rules that are available to use with the wavelet toolbox [9] and are listed in Table 1. 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. For the soft threshold estimator in the first method, a threshold selection rule which is based on Stein's unbiased estimate of risk (SURE), is used. An estimation of risk for a certain threshold value x_0 is obtained, and then by minimizing 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 (SNR) is very small (for the third method), the SURE estimate is very noisy. If the SNR 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.' Fig. 6 shows the results of each of the four threshold selection rules applied to a noisy signal.

Table 1
Threshold selection rules

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

[1] A wavelet basis is orthogonal if every wavelet is at ninety degrees to every other wavelet.

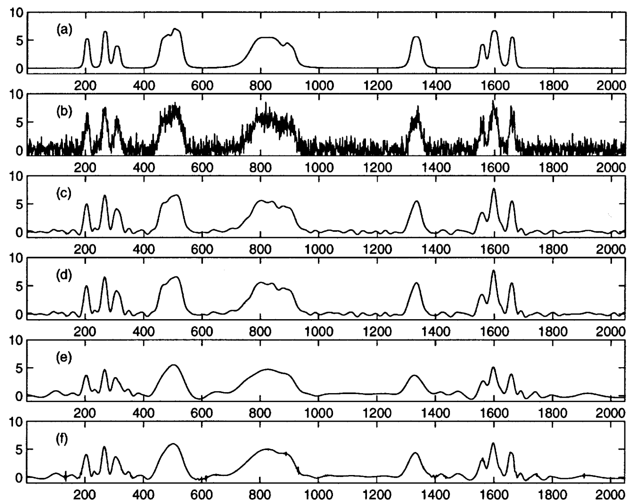

Fig. 6. (a) Original signal (b) Noisy signal — Signal to noise ratio = 2 (c) Denoised signal — heuristic SURE (d) Denoised signal — SURE (e) Denoised signal — fixed form threshold (f) Denoised signal — minimax (All are denoised using a Daubechies 10 wavelet with 10 levels of decomposition).

The basic noise model used for wavelet denoising in the MATLAB wavelet toolbox is: $s(n) = f(n) + \sigma e(n)$ 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 [9]. The objective of the denoising process is to suppress the noisy part of the signal s and recover f, which is the signal without noise.

There are three threshold rescaling methods available in MATLAB wavelet toolbox which are 'one', 'sln', and 'mln' [9]. The scheme 'one' follows the basic noise model. The option 'sln' corresponds to the basic noise model with unscaled noise and performs threshold rescaling using only a single estimation of level noise which is computed based on the first level coefficients of the decomposition. The 'mln' method corresponds to the basic noise model with non-white noise and performs threshold rescaling which is based on a level-dependent estimation of the noise at that decomposition level.

6.4. The use of averaging to denoise

Averaging in combination with wavelet denoising is examined in the context of how much improvement there is in the SNR and how many cycles should be averaged.

6.5. Use of HT

HT is used to demonstrate how effective de-noising is. The instantaneous frequency of a signal may be used to demonstrate how effective denoising is [14]. The use of the instantaneous signal parameters are also explored as a means of extracting additional information from the PCG.

6.6. Measuring the results of denoising

SNR is a traditional method of measuring the amount of noise present in a signal. SNR is defined as $10 \times \log_{10}(\text{Power}_{\text{signal}}/\text{Power}_{\text{noise}})$ measured in decibels. Two tests are performed using the SNR to measure the performance of wavelet denoising [5]. Because there is currently no known method to calculate which wavelet and thresholding parameters best denoise a signal, tests must be performed to evaluate the denoising capabilities of wavelets and thresholding parameter 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 wavelet denoising technique was applied to the heart sound recording which has noise added. Then the SNR was calculated for the denoised signal and the original signal. The higher the SNR, the less noise there is. Another test is to apply the wavelet denoising technique to a clean recording and compute the SNR of the resultant signal and the original signal. This test determines the loss in the information in the original signal by the denoising process. In other words, the more of the original clean signal that remains after applying the wavelet transform, the better.

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

6.7. Experimental results

While no one wavelet seemed to consistently give better results than another, there are some recommendations that can be made. The Coiflet 4 and 5, the Daubechies 11, 14, and 20 and the Symlet 9, 11, and 14 seemed to produce slightly better results in terms of noise removal than most of the wavelets tried, while the Daubechies 1, Coiflet 1, and Symlet 1 and 2 seemed to produce marginally worse results. The lower order wavelets did not perform as well as their higher order components due to their properties such as support length, regularity, and the number of vanishing moments [18].

It appeared that using 5 levels of decomposition usually proved adequate. For three trials (a trial being a clean heartbeat), varying amounts of noise was added (1, 5, and 10 dBs) and the levels of decomposition were varied from 1 to 10 with other denoising parameters remaining constant. In most cases, using a decomposition level of 5 achieved an SNR after denoising within 95% of the best denoising result for that set of data. However, in a test performed where noise was added to three trials in increments of 1 dB from 1 to 20 dBs, the best denoising results obtained in each case ranged from using 4 to 10 levels of decomposition.

The wavelet transform process of decomposition and

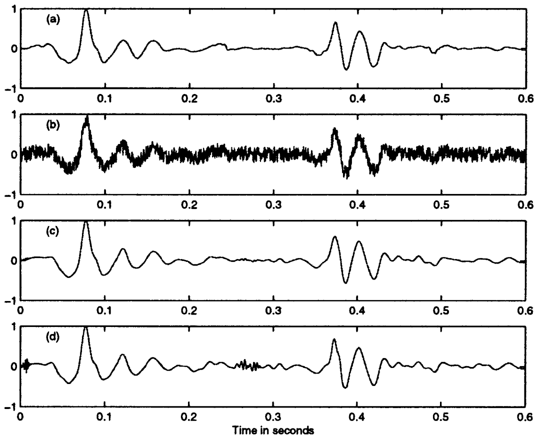

Fig. 7. This figure demonstrates that hard thresholding can cause discontinuities in a signal. (a) characteristic heartbeat signal (b) the heartbeat with white noise added resulting in an SNR of 1 dB (c) the heartbeat denoised using soft thresholding (d) the heartbeat denoised with hard thresholding.

recomposition was applied to several clean heartbeats. The SNR of the original signal and the signal after WT is applied are calculated. This number represents the loss in information in the wavelet analysis process. Any wavelet that consistently gave an SNR of under 60 dBs was considered to have lost too much information from the original signal. The wavelets which seemed to lose the most information were the Daubechies 1, 2, 3, 43, 44, 45 and the Symlet 1 and 2.

Soft thresholding definitely outperformed hard thresholding in the threshold selection category. Soft thresholding almost always gave better SNR after denoising than hard thresholding. Fig. 7 shows how hard thresholding may cause discontinuities in a signal.

Of the four threshold selection rules, the minimax and

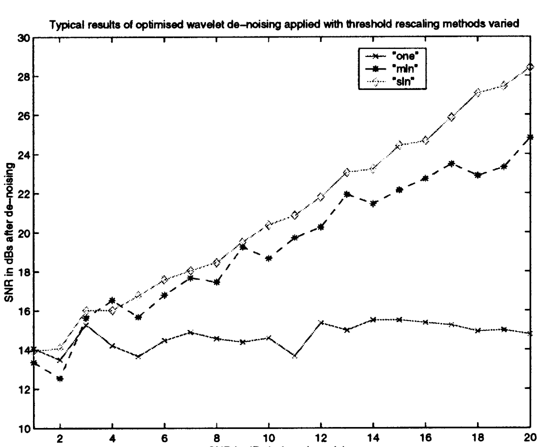

Fig. 8. This figure is a comparison of the different threshold rescaling methods. Noise in increments of 1 dB from 1 to 20 dBs is added to a clean characteristic heartbeat. The signal is then denoised using wavelet denoising (Daubechies 14 wavelet with 10 levels of decomposition using soft thresholding, and the rigorous SURE threshold selection rule) varying the threshold rescaling methods. Overall, the 'sln' method performs the best.

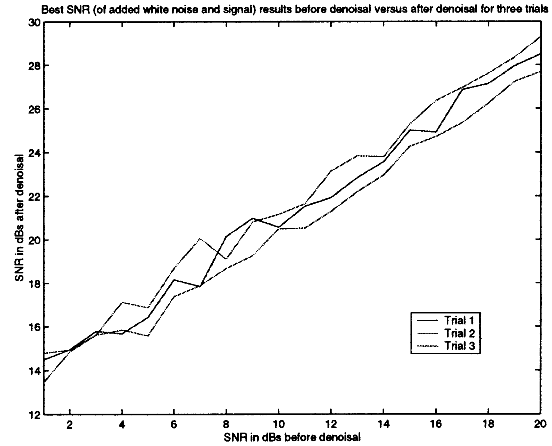

Fig. 9. Best SNR results before denoising versus after wavelet denoising for three trials. Each trial is a different, clean heartbeat. A known amount of white noise in decibels is added to a characteristic heartbeat. Different combinations of wavelets, thresholding techniques, and levels of decomposition are tried. The highest SNR after denoising is displayed for each trial. It is interesting to note that the best results for each trial are very similar.

SURE ('rigrsure') threshold selection schemes are more conservative than the others, and therefore should be used when small details of 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

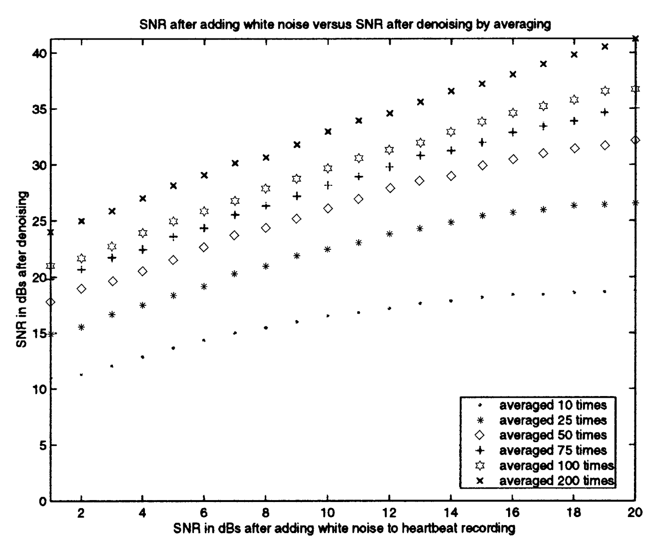

Fig. 10. This chart shows the SNR after adding white noise to a series of heartbeats and the SNR after averaging these series of heartbeats to obtain a characteristic heartbeat and reduce noise. There seems to be marginal improvement in SNR when little noise is present in the signal and the signal is 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 heartbeats, the SNR is still about 20 dBs meaning the noise remains. Averaging a series of 50–75 beats seems to give the best results here in terms of recording and computation time trade-off.

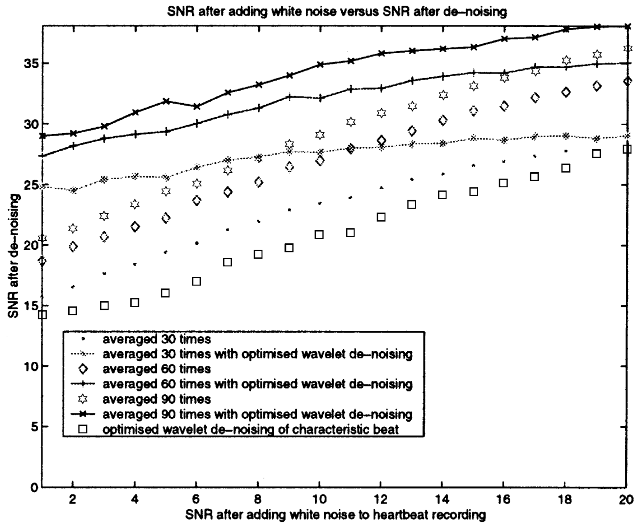

Fig. 11. This graph shows the SNR after adding white noise to a series of heartbeats versus the SNR after denoising the signal. Various methods are tried: wavelet denoising only, averaging only, and wavelet denoising combined with averaging. The wavelet denoising combined with averaging is the most successful denoising method. Heartbeat signals of lengths of 30, 60 and 90 were averaged. It appears that averaging at least 60 heartbeats would be recommended. There is a significant improvement of averaging 60 beats over 30 heartbeats while averaging 90 beats over 60 beats does not show nearly the same improvement.

rigorous SURE method almost always produced better results than the other methods with the heuristic SURE coming in second. Table 2 shows typical results for each of the four threshold selection rules.

The ‘sln’ method performed the best of three threshold rescaling methods available. For example, in Fig. 8, with large amounts of noise present, the methods all perform roughly equally with the ‘sln’ method producing slightly better results than the others. 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 denoising no matter what amount of noise was present in the signal.

Fig. 9 shows the best results for adding known amounts of noise to three different heartbeats and then applying wavelet denoising with varying parameters. The best results are very similar proving that the results of wavelet denoising are easily reproducible for heartbeat sounds. Also, the SNRs after wavelet denoising appear to be relatively linear corresponding to the linear SNR between the signal and the addition of white noise.

Averaging seemed to produce significant improvements especially if there is a large amount of noise present in the signal. Fig. 10 shows that averaging a series of 50–75 beats seems to give the best result in terms of recording and computation time tradeoff. It is difficult to obtain a long, clean recording which also increases the computation time required.

Fig. 11 shows a comparison of using wavelet denoising alone and wavelet denoising 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 beats that 60 beats provides a good compromise in terms of denoising and recording and computation time.

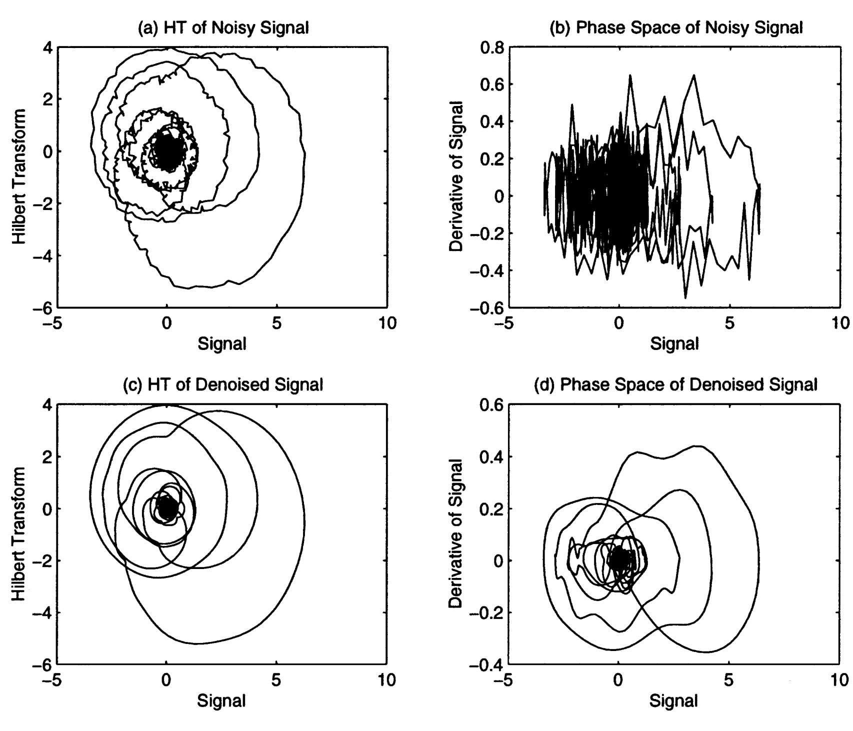

Fig. 12. This figure shows (a) HT of a characteristic heartbeat with 20 dBs of additive noise (b) the phase space diagram of the same noisy signal (c) the HT of the denoised heartbeat and (d) the phase space diagram of the denoised heartbeat.

Table 2
Typical SNR results after denoising using four threshold selection rules. White noise was added to three different characteristic heartbeats resulting in an SNR of 1 dB. Wavelet denoising was applied to the 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

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

The HT may be used to demonstrate how effective the denoising process performs but the instantaneous frequency and phase space diagram are better suited for this task. It may be seen that the signal is much cleaner after being denoised in Fig. 4, but this dramatic improvement is only because the signal has much additive noise. The HT of a signal does not accentuate the noise in a signal and thus does not show such dramatic improvement after denoising as is seen in phase space diagrams shown in [5]. The phase space diagrams take the derivative of a signal thus accentuating the noisy, high frequency, content. This fact is demonstrated in Fig. 12.

The instantaneous amplitude is an alternative method of looking at the PCG data. Fig. 13 demonstrates that recording the PCG of a patient is reproducible because plots of the instantaneous amplitude of a PCG recorded on 4 different occasions are very similar. Fig. 14 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.

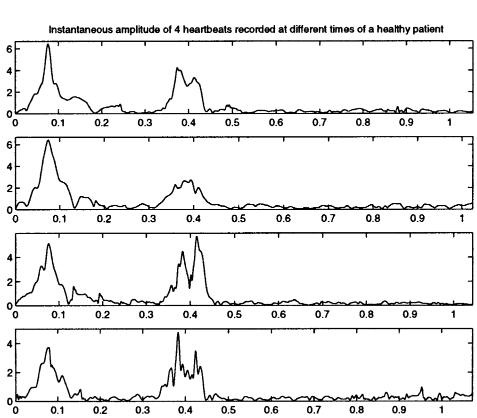

Fig. 13. This figure shows the instantaneous amplitude of 4 heartbeats recorded at different times from the same healthy patient. They are all fairly similar demonstrating that this technique is reproducible.

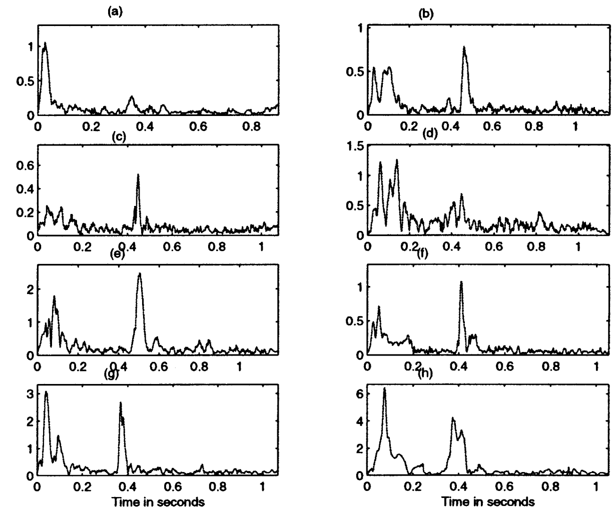

Fig. 14. (a) Instantaneous amplitude of a heartbeat of a patient with mitral valve prosthesis. (b) Instantaneous amplitude of a heartbeat of a patient with a heart murmur and hypertension. (c) Instantaneous amplitude of a heartbeat of a patient with an aortic stenosis and hypertension. (d) Instantaneous amplitude of a heartbeat of a patient with a systolic murmur and angioplasty. (e) Instantaneous amplitude of a heartbeat of a patient with atrial fibrillation. (f) Instantaneous amplitude of a heartbeat of a patient with hypertension. (g) Instantaneous amplitude of a heartbeat of a healthy patient. (h) Instantaneous amplitude of a heartbeat of a different healthy patient.

We have demonstrated that the HT has many interesting applications to PCGs. The complex PCG trace reveals much information about the signal. The HT may also be used to calculate the instantaneous parameters of a signal. The phase space diagram demonstrates denoising better visually, instead of plotting the HT and the signal, although instantaneous frequency calculated using the HT effectively demonstrates denoising.

7. Conclusions

We have demonstrated that wavelet denoising techniques and averaging are useful for removing white noise from heart sounds. Wavelet denoising in combination with averaging produced the best denoising results of the methods applied. However, there may be certain clinical cases, for example if a pathological condition was only present in some beats and not others, where wavelet denoising alone should be employed as averaging reduces a sequence of heartbeats to a single characteristic heartbeat.

Although there was no evidence that a single wavelet was the best suited for denoising heartbeats, there were some wavelets which were slightly better than others and certain wavelets which would not be recommended for this purpose. 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 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. Averaging heartbeats to produce a characteristic heartbeat significantly reduced noise. Averaging a sequence of about 60 heartbeats gave optimal denoising benefits in terms of computation time and difficulty in recording a clean PCG.

The use of the HT was explored in relation to the analysis of heartbeats. More information is displayed in the complex PCG trace. The instantaneous frequency of a PCG, as a tool of clinical use, which perhaps would reveal information about the condition of the heart, merits further research.

Another method of denoising the PCG which should be explored in the future would be the matching pursuit method [19]. Also, more research is needed to determine what sort of noise corrupts PCGs. Once, the sort of noise polluting the PCG is determined, a system could employ different denoising techniques based upon what sort of noise was present.

Acknowledgements

We would like to thank Leonard Hall, Jarrad Maple and Mohammed Ali Tinati (all of the Department of Electrical and Electronic Engineering, The University of Adelaide) for allowing us to make use of their heart sound recording equipment and software.

We would also like to thank colleagues at the Centre for Biomedical Engineering (Electrical and Electronic Engineering, The University of Adelaide) including Brad Ferguson, Leonard Hall, Greg Harmer, and Samuel Mickan, for providing assistance and inputs in various areas along the way.

References

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

[2] S. Lukkarinen, A. Nopanen, K. Sikio, A. Angerla, A new phonocardiographic recording system, Computers in Cardiology 27 (1997) 117–120.

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

[4] M.A. Tinati, Time-frequency and time-scale analysis of phonocardiograms with coronary artery disease before and after angioplasty. PhD thesis, The University of Adelaide, October 1998.

[5] J. Maple, L. Hall, J. Agzarian, D. Abbott, Sensor system for heart sound biomonitor, Proceedigs of SPIE Electronics and Structures for MEMS 3891 (1999) 99–109.

[6] S. Messer, D. Abbott, J. Agzarian, Optimal wavelet denoising for smart biomonitor systems, Proceedings of SPIE: Smart Electronics and MEMS 4236 (2000) 66–78 SPIE, December 2000.

[7] R. Polikar, The wavelet tutorial, http://www.public.iastate.edu/rpolikar/WAVELETS/waveletindex.html, August 2000.

[8] B.B. Hubbard, The World According to Wavelets, A K Peters Ltd, 289 Linden Street, Wellesley, MA 02181, 1996.

[9] M. Misiti, Y. Misiti, G. Oppenheim, J.-M. Poggi, Wavelet Toolbox: For Use With MATLAB, The Math Works Inc, 1996.

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

[11] O. Ersoy, Fourier-Related Transforms, Fast Algorithms and Applications, Prentice-Hall, Inc, Upper Saddle River, New Jersey 07458, 1997.

[12] M.T. Taner, F. Koehler, R.E. Sheriff, Complex seismic trace analysis, Geophysics 44 (1979) 1041–1063.

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

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

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

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

[17] M. Akay, Wavelet applications in medicine, IEEE Spectrum 34 (1997) 50–56.

[18] S. Messer, D. Abbott, J. Agzarian, Comparison of automatic de-noising methods for phonocardiograms with extraction of signal parameters via the hilbert transform, Proceedings of the SPIE 4304 (2001) SPIE, January 2001.

[19] S. Mallat, Z. Zhang, Matching pursuits with time–frequency dictionaries, IEEE Transactions on Signal Processing 41 (1993) 3397–3415.