

Sensor System for Heart Sound Biomonitor

Jarrad L. Maple^a, Leonard T. Hall^a, John Agzarian^b, and Derek Abbott^a

^aCentre for Biomedical Engineering (CBME)

Department of Electrical and Electronic Engineering, University of Adelaide, SA 5005, Australia

^bHampstead Medical Clinic, Hampstead Gardens, SA 5086, Australia

ABSTRACT

Heart sounds can be utilised more efficiently by medical doctors when they are displayed visually, rather than through a conventional stethoscope. A system whereby a digital stethoscope interfaces directly to a PC will be described along with signal processing algorithms adopted. The sensor is based on a noise cancellation microphone, with a 450 Hz bandwidth and is sampled at 2250 samples/sec with 12-bit resolution. Further to this, we discuss for comparison a piezo-based sensor with a 1 kHz bandwidth. A major problem is that the recording of the heart sound into these devices is subject to unwanted background noise which can override the heart sound and results in a poor visual representation. This noise originates from various sources such as skin contact with the stethoscope diaphragm, lung sounds (patient breathing), and other surrounding sounds such as speech. Furthermore we demonstrate a solution using 'wavelet denoising.' The wavelet transform is used because of the similarity between the shape of wavelets and the time-domain shape of a heartbeat sound. Thus coding of the waveform into the wavelet domain is achieved with relatively few wavelet coefficients, in contrast to the many Fourier components that would result from conventional decomposition. We show that the background noise can be dramatically reduced by a thresholding operation in the wavelet domain. The principle is that the background noise codes into many small broadband wavelet coefficients that can be removed without significant degradation of the signal of interest.

Keywords: biomonitor sensors, wavelets, noise suppression

1. INTRODUCTION

The analysis of biological sounds within the human body or auscultation, as it is sometimes called, by the use of a stethoscope is a common practice of medical practitioners all over the world. Although this may appear to be a quite sufficient method of diagnosing heart defects, problems can easily go undetected due to limitations of the human ear in being able to distinguish defects from within the sound of a heartbeat³. There are far more effective methods such as the examination of a phonocardiogram (PCG) which is a visual display of heart sounds, or an electrocardiogram (ECG) which displays the electrical characteristics of the heart. These techniques are both non-invasive (i.e. they do not require surgery) but have their own share of issues and problems.

An ECG is not normally used unless a problem has been detected initially by auscultation. This is because an ECG requires a great deal of setting up and is therefore not practical for a general practitioner to use as a standard test. Phonocardiograms obtained from transducers positioned on the skin can easily be swamped by noise from surrounding sources such as the lungs, frictional skin contact, or speech and therefore are generally useless if conditions around the patient are not perfect.

The limitations of a PCG described above could be dramatically reduced by the development of a system that removes the unwanted noise but leaves important information about the characteristics of the heart being examined. This short paper is concerned with the study of a relatively new signal processing technique called wavelet analysis to achieve this goal.

Heart sounds, as heard through a normal stethoscope are produced by a biological membrane in the heart when an event such as the opening or closing of a valve, vibration of the cardiac structure, or acceleration or deceleration of blood occurs². The sound generally comprises of two parts, a first heart sound and a second heart sound, but sometimes aberrations and murmurs caused by valve imperfections and other obstructions to blood flow are also present in the sound. A 'click' sound also may be present if a patient has an artificial heart valve. It is important that the noise removal technique does not dispose of such information so as to become an ineffective analysis tool.

A digital stethoscope is used to transduce the heart sounds into an electrical signal. The heart sound data may then be manipulated with various signal processing tools available in MATLAB as an attempt to break down the signal into wavelet coefficients, thus allowing the removal of unwanted sounds by thresholding. It is intended that this paper introduce the reader to basic wavelet theory as it applies to the denoising of heart sounds, and briefly demonstrate some preliminary results that may inspire the reader to pursue this field further.

2. SENSOR EQUIPMENT

2.1 The Analyst

The Analyst is a digital stethoscope developed by HeartSounds® in the USA. It comprises of a normal stethoscope diaphragm linked to a 'noise cancellation' microphone sampled at 2250 samples/sec with 12-bit resolution and 450 Hz bandwidth. Hardware such as amplifiers and memory are included internally within the system. The system links into a PC through the serial port and communicates with the Sound Management Software (SMS) installed on the PC. The Analyst has the ability to record and store up to three separate heart sounds on board, they then must be transferred to the PC for viewing on SMS. A later model of the Analyst, the ICON (Analyst XT), will be released later this year and will have improved noise cancellation, a higher sampling rate, and better digital resolution. The SMS software proved to be unreliable on a Windows NT system, therefore new software was written to dispose with unwanted effects experienced with using this system.

2.2 Piezoelectric Sensor System

The equipment used to compare the output from the Analyst system described above comprises a piezoelectric sensor linked to amplification and filtering circuits (to give sufficient signal gain and eliminate aliasing) and then to an analogue to digital converter card in a PC. The piezo sensor used is the Hewlett Packard HP 21050A, composed of crystalline material arranged in slabs and exhibits the piezoelectric effect by creating a changing electrical polarisation when it is deformed. The displacement of the crystal, caused by vibrations from the sound waves entering the system, cause a voltage to be developed across its output terminals proportional to the sound. The piezo crystal displays sensitivity of 10 $\mu\text{V}/\text{dyne}$ force and 40 mV/micron displacement, capacitance of 20nF, and frequency response between 0.02 Hz and 2 kHz. The amplification and filtering circuit is a second order butterworth filter with cutoff at 1 kHz and provides 40 dB of gain to the signal. The analogue to digital converter card is a WIN-30D from which the sampled analogue signal can be retrieved and processed by MATLAB. Features of the WIN-30D include 12 bit resolution, 16 single ended inputs, nonlinearity of less than +/- 1 least significant bit (LSB), throughput rate of 1MHz, and full scale input ranges of 0 to +5V (for Unipolar range) and -5 to +5V (for Bipolar range). This system provides accurate, high-resolution signals accompanying a higher bandwidth that can be compared to the lower resolution signals recorded from the Analyst.

3. BACKGROUND OF WAVELET ANALYSIS

Wavelet analysis has evolved through years of research dating back to the work of Joseph Fourier, but it is still a relatively new concept with most formal research on the topic being conducted in the past 10 years. In the following, we provide a brief background of the history behind several important signal processing techniques to give the reader a sufficient understanding of wavelet theory.

3.1 Fourier Analysis

In 1822, J. Fourier, a famous French Mathematician, discovered that any periodic function could be expressed as an infinite sum of periodic complex exponential functions¹. This property of periodic functions was later generalized to non-periodic functions and then to (both periodic and non-periodic) discrete time functions. The process of converting a signal to the frequency domain from the time domain is achieved with the Fourier Transform (FT). The Fourier Transform is still frequently used today (mainly in the form of the Fast Fourier Transform algorithm with the aid of computers) and is mathematically described below as

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft} dt \quad (1)$$

Here, the time domain signal $x(t)$ is multiplied by a complex exponential at some certain frequency f , and then integrated over *all* time. The results of this calculation are the Fourier Coefficients $X(f)$, which are large when the signal contains a

frequency component close to or equal to f . These large coefficients correspond to the peaks at various frequencies exhibited on a plot of the Fourier Transform of a signal.

Using the Fourier Transform has one main disadvantage in that it does not provide enough information when used on non-stationary signals (ie signals whose frequency is time varying). The Fourier Transform only determines the frequency components of a signal, but not their location in time. Since most signals of practical interest to engineers (such as heart sound signals) are of a non-stationary nature, the Fourier Transform is not particularly suited to their analysis.

3.2 The Short Time Fourier Transform (STFT)

The problem with the Fourier Transform arises from the fact that the signal is integrated over *all* time as in equation (1), therefore no matter where in time the component with frequency f appears, it will affect the result of the integration equally as well, hence giving an accurate representation of the frequencies present in the signal, but no information about the instant of time at which the specific frequency occurs.

In 1946, Denis Gabor ⁴ realised the need to exhibit signals in the time-frequency domain. He developed a technique involving 'windowing' the signal which maps the signal into a two dimensional space of time and frequency. This technique is known as the Short Time Fourier Transform (STFT), and the mathematical theory behind this technique is best described by

$$STFT_X^{(w)}(t', x) = \int_t [x(t) \cdot w^*(t - t')] \cdot e^{-j2\pi ft} dt. \quad (2)$$

The windowing technique involves translating (denoted in equation (2) by the translation parameter t') the complex conjugate of the window function $w(t)$ along the length of the signal $x(t)$, and multiplying the two functions $x(t)$ and $w^*(t)$ at the different instants of time. The exponential part converts the result of each multiplication to the frequency domain at that instant as done in the ordinary Fourier Transform.

3.3 The Wavelet Transform (WT)

To overcome resolution problems which make it impossible to analyse the signal simultaneously in both the time and frequency domain, the Wavelet Transform was developed. The first recorded mention of the term 'wavelet' was in 1909, in a thesis by Alfred Haar ⁴, but over the last 10 years a sharply increasing amount of research has been conducted into its possible applications.

The main difference between the STFT and the WT is that the WT uses a variable sized window region (or wavelet) to examine the signal which helps to reduce resolution problems significantly. A wavelet is essentially a waveform of limited duration that has an average value of zero. Some examples of commonly used wavelets used in signal processing are shown in Figure 1.




Figure 1 – Commonly used Wavelets

The theory behind this approach is that being able to dilate (stretch) or compress the wavelet, different features of the signal will be extracted. For example a narrow wavelet will show up higher frequency components, while a stretched wavelet will show up lower frequency components of the signal. A comparison between the constant window regions used in STFT analysis and the variable window regions used in WT analysis is exhibited in Figure 2.




Figure 2 - Window Regions of STFT and WT Analyses

The mathematical equation describing the Continuous Wavelet Transform (CWT) is

$$CWT_x^\psi(\tau, s) = \Psi_x^\psi(\tau, s) = \frac{1}{\sqrt{|s|}} \int x(t) \cdot \psi^*\left(\frac{t-\tau}{s}\right) dt. \quad (3)$$

The quantity s is referred to as the 'scale' of the wavelet, which can be thought of as the inverse of frequency. The wavelet is compressed if the scale is low and stretched if the scale is high, which is also evident in Figure 2. Therefore, plots of the Wavelet Transform of a signal represent scale against time rather than frequency against time. The process involved in creating the CWT is very much the same as that involved with the STFT, except that the wavelet traverses the signal many times (as indicated by the translation parameter τ), with each traversal computed with a different scale. The process is illustrated in Figure 3.




Figure 3 - Process involved with creating a CWT of a signal

The CWT is extremely good at displaying information about the signal in great detail, but due to limited computational power available with computer systems, the Discrete Wavelet Transform (DWT) is usually used. The DWT computes the wavelet coefficients at discrete intervals of time and scale, compared to the CWT which shifts the wavelet and changes the scale in a continuous nature in order to calculate the WT coefficients. The DWT is used to make the process more efficient without a significant loss of detail. It turns out that if scales and translations based on powers of 2 are used, there is virtually no detail lost in the analysis and the signal may be perfectly reconstructed. This power of 2 transformation is precisely the method used when a DWT of a signal is computed.

4. DENOISING WITH WAVELETS

4.1 Signal Decomposition and Reconstruction

The MATLAB Wavelet Toolbox provides standard routines for using the DWT to decompose a signal into wavelet coefficients and then reconstruct it with the Inverse Discrete Wavelet Transform (IDWT). A diverse selection of wavelets such as the 'Morlet' and 'Mexican Hat' wavelets shown in Figure 1 are available in this package. The procedure used by MATLAB to achieve this decomposition and recomposition of a signal involves the process of applying numerous highpass and lowpass filters in succession. This scheme was developed by Mallat⁶ in 1988, and is actually a widely used signal processing tool known as a two-channel subband coder. The process actually breaks the signal successively into 'approximations' and 'details', whereby approximations represent the slowly changing (low frequency) features of a signal, and conversely the details represent the rapidly changing (high frequency) features of the signal. A detailed description of this process is provided in the *Wavelet Toolbox Users Guide*⁴ and also in *Introduction to Wavelets and Wavelet Transforms: A Primer*⁵.

4.2 Noise Removal by Thresholding

The Wavelet Toolbox provides the facility to view and zoom in on the signal's decomposed wavelet coefficients. One important feature of wavelet analysis is that it is capable of exposing sharp discontinuities in a signal, which other signal processing techniques such as Fourier Analysis miss. This operation can be more efficient with certain types of wavelets than with others, and is a property that has been exploited in our analysis. Various discontinuities show up in the wavelet decomposition, therefore the introduction of noise is readily revealed and hence can be removed by thresholding certain components of the wavelet decomposition. Another technique, which can be exploited with the Wavelet Toolbox, is the comparison a noise free heart sound with the same heart sound combined with unwanted noise (such as that from the lungs). The wavelet decompositions of both signals can then be compared to determine which coefficients correspond to the noise source and therefore can be removed effectively.

The thresholding operation involves removing coefficients of the DWT which lie below a specified value, and then reconstructing the signal with the IDWT. This is a very powerful concept because signals with energy concentrated in a small number of wavelet dimensions will have coefficients that are relatively large compared to any other signal present that has its energy concentrated over a larger number of wavelet dimensions. Therefore applying the thresholding operation to the DWT will effectively remove any unwanted signal or noise, even if the instantaneous frequency spectra of the two signals overlap. The only problem left is to determine the most suitable wavelet basis function or 'mother wavelet' to apply to the signal to achieve this result.

4.3 Experimental Data

Consider the 6-second heart sound recording taken from a healthy patient with the Analyst shown in Figure 4.




Figure 4 – Heart Sound from Analyst (Healthy Patient)

You will notice that this recording contains a significant amount of background noise, despite the fact that it was taken in a controlled, relatively noise free environment. As a first demonstration, we will now show this signal broken down into its discrete wavelet coefficients by applying wavelets from the 'Daubechies' family of order 5 (db5) to the signal. Trial and error with various wavelet families has revealed this wavelet to be most suitable for analysing the healthy heart sound shown above. This is due to heart beat signal having most of its energy distributed over a small number of db5 wavelet dimensions (or scales), and therefore the coefficients corresponding to the heart beat signal will be large compared to any other noisy signal present. Information about the db5 wavelet used, including the filters used by MATLAB to perform deconstruction and reconstruction of the signal is shown in Figure 5.




Figure 5 – Properties of the db5 Wavelet

The discrete wavelet coefficients of the signal are shown in Figure 6. Note that the transform has been taken over seven levels of decomposition (or scale).




Figure 6 – Discrete Wavelet Coefficients of the Heart Sound

It is now useful to present how the coefficients have been thresholded at the various levels of decomposition. It would be unwise to threshold each coefficient in the wavelet domain since it would be impossible to determine which coefficients belong to a certain signal. To overcome this problem, the Wavelet Toolbox provides a useful feature allowing the user to be able to threshold the signal in the time domain. Actually, it is the 'approximations' of the signal at the highest level, and the 'details' of all levels of the signal that are thresholded. The optimum thresholding levels applied to the signal 'details' are shown in Figure 7 below.




Figure 7 – Thresholding of the signal's 'details'

In detail level d_7 , a low threshold is applied to retain the general shape of the signal whilst removing any unwanted low amplitude noise components. The detail levels d_6 through d_4 are thresholded just above the point where the heartbeat signal appears to protrude through the noise. It is at these detail levels (or scale) that the shape of the applied wavelet corresponds very closely with the shape of the heartbeat signal, and therefore it is quite easy to distinguish between the heart sound and the noise. Detail level d_3 shows up some low amplitude correlation of the wavelet with the heart sound, but it is quite difficult to distinguish between the heart sound and the noise at this level, and therefore the entire detail level is thresholded out. Detail levels d_2 and d_1 correspond closely to the noise component, and are thresholded accordingly.

Finally, the denoised heart sound in the time domain is shown in Figure 8 along with the thresholded wavelet coefficients of the signal. We notice in this figure that the 'clean' denoised heart sound (darker colour) is superimposed on top of the original 'noisy' heart sound (lighter colour) to emphasise the effectiveness of this noise removal operation.




Figure 8 – Denoised heart sound and thresholded coefficients

We now turn our attention to another noisy heart sound recorded with the Analyst which contains a heart defect known as a Presystolic Murmur ending in the first accentuated heart sound. The denoised heart sound is shown in Figure 9 superimposed over the original noisy heart sound.




Figure 9 – Heart sound with a Presystolic Murmur

The above heart sound was denoised over 7 scales using the technique previously discussed, but in this instance, the optimal wavelet bases was found to be the 'Daubechies order 9 (db9)' wavelet. A murmur is basically the effect of a heart valve allowing blood to rush through it producing a relatively high frequency sound compared to the normal heart beat. In this instance, the rushing of blood occurs before the heart contraction (or first heart sound), and this is shown clearly in the above figure.

Finally, let us consider a normal heart sound corrupted with noise from the lungs of a heavily breathing patient. Performing the same wavelet analysis again yields the results shown in Figure 10. This time a 'Daubechies order 7 (db7)' wavelet over 7 scales was used for the denoising of the signal. Again, note that Figure 10 shows the denoised heart sound superimposed over the original noisy signal. We can see from this figure that the noise removal is not perfect, but it is evident that the heart sound has been separated from the lung sound quite adequately.




Figure 10 – Denoised healthy heart sound originally corrupted with sound from the lungs

5. CONCLUSIONS

We have provided a brief introduction to the fundamental concepts of Wavelet Analysis and how they apply to the removal of noise from heart sound recordings. The characteristics the equipment used to acquire these recordings has been provided, along with the discussion of another system from which the recordings will be compared to. Several examples of wavelet denoising of various types heart sounds have been achieved with the use of MATLAB's Wavelet Toolbox. This has been provided as a demonstration of the results one can expect with this type of analysis.

ACKNOWLEDGEMENTS

We would like to thank Nariman Habili (Department of Electrical & Electronic Engineering, University of Adelaide) for referring us to Dr. Homer Nazeran (Department of Biomedical Engineering, Flinders University), who took the time to provide some useful information regarding this research area.

We would also like to take the time to thank Jeffrey Werner (Heartsounds USA Inc. Wayzata, Minnesota, USA) and Ed Goss (Product Manager, Heartsounds Inc. Toronto, Ontario Canada) for providing technical information on the Analyst stethoscope.

Finally, we would like to mention Sam Mickan (Department of Electrical & Electronic Engineering, The University of Adelaide) for his ongoing interest in our research, and Mohammad Ali Tinati (Department of Electrical & Electronic Engineering, The University of Adelaide) for allowing us to make use of his heart sound recording equipment.

REFERENCES

1. *Wavelet Tutorial* - www2.iastate.edu/~rpolikar/WAVELETS/waveletindex.html
2. M.A. Tinati, *Time-frequency and time-scale analysis of phonocardiograms with coronary artery disease before and after angioplasty. (Adelaide Thesis Ph.D.)*, Dept. of Electrical and Electronic Engineering, The University of Adelaide SA, 1998
3. M. Akay, *Time frequency and wavelets in biomedical signal processing*, pp. 271-301, IEEE Press series in biomedical engineering, 1997
4. Mathworks Inc., *Wavelet Toolbox User's Guide* - First printing , Chapter 1 – Wavelets: A new tool for signal analysis, March 1996
5. S. Burrus, R. A. Gopinath, H. Guo, *Introduction to Wavelets and Wavelet Transforms: A Primer*, Prentice Hall, New Jersey, 1998
6. Mallat, S. "A theory for multiresolution signal decomposition: the wavelet representation", IEEE Pattern Anal. and Machine Intell., vol. 11, no. 7, pp 674-693., 1989