

Microelectronics Journal 31 (2000) 583-592

Microelectronics Journal

www.elsevier.com/locate/mejo

Sensor system for heart sound biomonitor

L.T. Hall^a, J.L. Maple^a, J. Agzarian^b, D. Abbott^{a,*}

^aCentre for Biomedical Engineering (CBME), Department of Electrical and Electronic Engineering, University of Adelaide, Adelaide, SA 5005, Australia ^bHampstead Medical Clinic, Hampstead Gardens, Adelaide, SA 5086, Australia

Received 15 December 1999; accepted 30 December 1999

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 per second 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. Also explored is the use of a phase space representation of the heart sound to exploit different properties of the signal. © 2000 Elsevier Science Ltd. All rights reserved.

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 heartbeat [3]. 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 occurs [2]. This 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

^{*} Corresponding author. Tel.: + 61-8-8303-5748; fax: + 61-8-8303-4360.

E-mail address: dabbott@eleceng.adelaide.edu.au (D. Abbott).

^{0026-2692/00/}\$ - see front matter © 2000 Elsevier Science Ltd. All rights reserved. PII: S0026-2692(00)00033-1

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 Heart-Sounds[®] in the USA. It comprises of a normal stethoscope diaphragm linked to a 'noise cancellation' microphone sampled at 2250 samples per second 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 µV/dyne force and 40 mV/micron displacement, capacitance of 20 nF, and frequency response between 0.02 and 2 kHz. The amplification and filtering circuit is a second order butterworth filter with cut-off 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 1 MHz, and full scale input ranges of 0 to + 5 V (for Unipolar range) and -5 to + 5 V (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 functions [1]. 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

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-j2\pi f t} dt.$$
 (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 (i.e. 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

The problem with the Fourier Transform arises from the fact that the signal is integrated over *all* time as in Eq. (1),



Fig. 1. Commonly used wavelets.

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) =
$$\int_{t} [x(t)w^*(t - t')] e^{-j2\pi f t} dt.$$
 (2)

The windowing technique involves translating (denoted in Eq. (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

To overcome resolution problems that make it impossible





to analyse the signal simultaneously in both the time and frequency domain, the wavelet transform (WT) 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 Fig. 1.

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 Fig. 2.

The mathematical equation describing the continuous wavelet transform (CWT) is

$$\operatorname{CWT}_{x}^{\psi}(\tau, s) = \Psi_{x}^{\psi}(\tau, x) = \frac{1}{\sqrt{|s|}} \int x(t) \psi^{*}\left(\frac{t-\tau}{s}\right) \mathrm{d}t \qquad (3)$$

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 Fig. 2. Therefore, plots of the WT 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 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





Fig. 3. Signal to noise ratio (SNR) of the denoised signals.

wavelet transform (DWT) is usually used. The DWT computes the wavelet coefficients at discrete intervals of time and scale, compared to the CWT that 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 Fig. 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 Mallat [6] 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



Fig. 4. Heart sound from analyst (healthy patient).

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. Testing of wavelet's and level of decomposition

Selection of wavelet and the level of decomposition were done on two criteria. The first test was automatically denoising a clean heart signal. The denoised clean heart sound is then tested against the original signal to get a signal to noise ratio that represents the information lost due to denoising. The second test we carried out was designed to determine how effectively white noise could be removed from a heart sound; 10 dB of white noise is added to the clean heart sound and then the result is automatically denoised using rigorous thresholding. In this way we get two values, one representing the information lost and the other representing the noise removed.

The results are shown in Fig. 3. From these results we can see db10 at a decomposition level of 5 performs very well. It can be seen that the advantages of increasing the level of decomposition above 5 are very limited. In practice db6 or db7 prove to be just as effective as db10. Provided the decomposition is above 5, all of the aforementioned wavelets perform quiet well.

4.4. Phase space representation

Apart from time-scale and wavelet analysis, the use of another representation in the form of a phase space diagram is also explored. The concept of a phase space representation of a signal is presented in Use of the Nonlinear Dynamical System Theory to Study Cycle to Cycle Variations from Spark Ignition Engine Pressure Data [7]. Theoretically, the phase space diagram of a signal is a plot of the rate of change (time derivative) of the signal against the original signal. If a pure sinusoid were to be plotted in this fashion, a circular shape would result with the size of the circle proportional to the amplitude of the sinusoid. The plot may be thought of as a superposition of larger circular shapes and lower amplitude impulsive shapes corresponding to higher frequency signals within the heart sound. The impulsive components are rapidly changing (hence a large value of d/dt) compared to a small amplitude component. This should therefore produce large vertical movements on the plot compared to smaller horizontal movements for this



Fig. 5. Properties of the db5 wavelet.

type of noisy signal. The plot should follow itself when plotting a clean periodically repeatable waveform.

4.5. Experimental data

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

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



Fig. 6. Discrete wavelet coefficients of the heart sound.



Fig. 7. Thresholding of the signal's 'details'.

present. Information about the db5 wavelet used, including the filters used by MATLAB to perform deconstruction and reconstruction of the signal is shown in Fig. 5.

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

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



Fig. 8. Denoised heart sound and thresholded coefficents.



Fig. 9. Original and denoised phase space plots.

the signal at the highest level, and the 'details' of all levels of the signal that are thresholded. The thresholding levels applied to the signal 'details' are shown in Fig. 7 below.

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. In practice the thresholding of coefficients is done automatically using fixed form soft thresholding.

Finally, the denoised heart sound in the time domain is shown in Fig. 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.

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



Fig. 10. Heart sound with a presystolic murmur.



Fig. 11. Denoised healthy heart sound originally corrupted with sound from the lungs.

shown in Fig. 9 superimposed over the original noisy heart sound.

As can be seen in these examples, the difference in the phase space diagram, before and after denoising, is obvious. The circles representing the individual beats become substantially smoother and much closer together. This is particularly obvious when observing the second heart murmur. The original phase space diagram is completely indistinguishable. After denoising, the individual circles are very obvious. Denoising also has another substantial advantage. In extremely quiet recordings the difference between the maximum and minimum may be as small as thirty samples. After denoising, these recordings become much clearer due to the increase in resolution. This is shown very clearly in Figs. 8 and 9.

The above heart sound was denoised over seven 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 Fig. 10. 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 Fig. 11. This time a 'Daubechies order 7 (db7)' wavelet over 7 scales was used for the denoising of the signal. Again, note that Fig. 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.

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 and 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 and Electronic Engineering, The University of Adelaide) for his ongoing interest in our research, and Mohammad Ali Tinati (Department of Electrical and 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. PhD thesis, Department of Electrical and Electronic Engineering, The University of Adelaide, SA, 1998.
- [3] M. Akay, Time Frequency and Wavelets in Biomedical Signal Processing, Series in Biomedical Engineering, IEEE Press, New York, 1997 pp. 271–301..
- [4] Mathworks Inc., Wavelet Toolbox user's guide, 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, Englewood Cliffs, NJ, 1998.
- [6] S. Mallat, A theory for multiresolution signal decomposition: the wavelet representation, IEEE Pattern Anal. Machine Intell. 11 (7) (1989) 674–693.
- [7] C. Letellier, S. Meunier-Guttin-Cluzel, G. Gouesbet, F. Neveu, T. Duverger, B. Cousyn, Use of the nonlinear dynamical system theory to study cycle to cycle variations from spark ignition engine pressure data, SAE Technical Paper Series, 971640, 1997, pp. 36–46.