The Detection Of P–Waves In Ambulatory ECG Recordings By The Use Of Wavelet Transforms

STEPHEN EDWARD D'AGUIAR

Master Of Philosophy

Aston University

December 2007

Aston University

The Detection Of P–Waves In Ambulatory ECG Recordings By The Use Of Wavelet Transforms

Stephen Edward d'Aguiar

Master Of Philosophy, 2007

Abstract

This research concerns the automated detection of P–waves in ambulatory ECG recordings using wavelet transforms. The P–wave is the electrical signal produced during the depolarisation of the atria. The interest in the detection of P–waves lies in their usefulness in the detection of arrhythmias of the heart.

Huntleigh Healthcare have already developed a commercial automatic ambulatory ECG analysis system, which classifies heartbeats using the information obtained from the ventricular activity. Improvements to this software can be obtained by augmenting this information with that obtained from the atrial activity.

Initially the Haar, Coiflet and Daubechies wavelets were compared for denoising ECG recordings. The Coiflet and Daubechies wavelets proved to be very similar in their ability to denoise signals, with both being superior to the Haar wavelet.

The algorithme à trous was identified as a suitable algorithm for implementing the detection of P–waves using wavelet transforms due to its time invariance. The Haar and quadratic spline wavelets were compared using this algorithm and both were able to accurately detect P–waves in real ECG recordings. Work on artificial ECG records however, suggested that the quadratic spline wavelet was superior to the Haar wavelet if a large amount of noise was present in the signal.

Acknowledgements

The author would like to acknowledge the help, advice and encouragement provided by David Evans, James Pardey and Svetlana Jouravleva in producing this research.

The author also thanks Vicky Bond for her help with matters of administration and Alex Brulo for computational assistance.

The author would also like to thank both the Engineering and Physical Sciences Research Council and Huntleigh Healthcare for their financial support, which has enabled this research to be carried out.

Figure 1.2 is reproduced by kind permission of Huntleigh Healthcare.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Overview

Cardiac diseases are a major cause of mortality in developed countries, accounting for over 4.5 million deaths in Europe in 1998 [World Health Organization 1999] and for about 22% of the total deaths in England and Wales [Office for National Statistics 2004, 2005 & 2006]. Consequently there is great interest in the diagnosis of these diseases and the Electrocardiogram (ECG) is the main technique used for this purpose. Because some symptoms of heart disease occur only infrequently there is a need for the continuous monitoring of the heart for periods of 24 hours or more. In order to carry out this monitoring, small portable devices, initially developed by Holter in the late 1950's [Holter 1961], are employed. Because a 24 hour ambulatory recording typically contains over 100,000 beats, the time it would take a cardiologist to complete the analysis becomes prohibitive and it is necessary to automate the analysis using computational techniques.

The ECG signal consists of a series of waves which correspond to the electrical activity of the heart; this electrical activity being related to the mechanical activity of the heart. These waves comprise the P–wave, denoting the depolarisation of the atria, QRS complex, corresponding to the depolarisation of the ventricles, and the T–wave, signifying the repolarisation of the ventricles. No signal for the repolarisation of the atria is detected in the ECG due to the fact that this event occurs during the depolarisation of the ventricles.

The automatic detection of QRS complexes is well documented, and has proved to be relatively easy to achieve due to the large signal to noise ratio of the R-wave. The automatic detection of P–waves has, however, proved to be more difficult. This is due to the relatively small signal to noise ratio and the fact that during arrhythmias the P–wave can become dissociated from the QRS complex and become hidden either within the QRS complex or the T-wave.

Huntleigh Healthcare, a company involved in diagnostic health care, have already developed an automated ECG analysis system (Medilog Darwin) which is commercially available. However there

is room for improvement in the software, particularly with regard to the detection of P–waves. The importance in the accurate determination of the P–waves lies in their usefulness in the diagnosis of arrhythmias of the heart. The aim of this research is to develop an accurate, automatic, P–wave detection algorithm, which can then be incorporated into the existing Medilog Darwin software.

1.2 The Heart

The physical action of the heart and its electrical activity are described. This is followed by a discussion of the diseases affecting the heart.

1.2.1 Physical Action

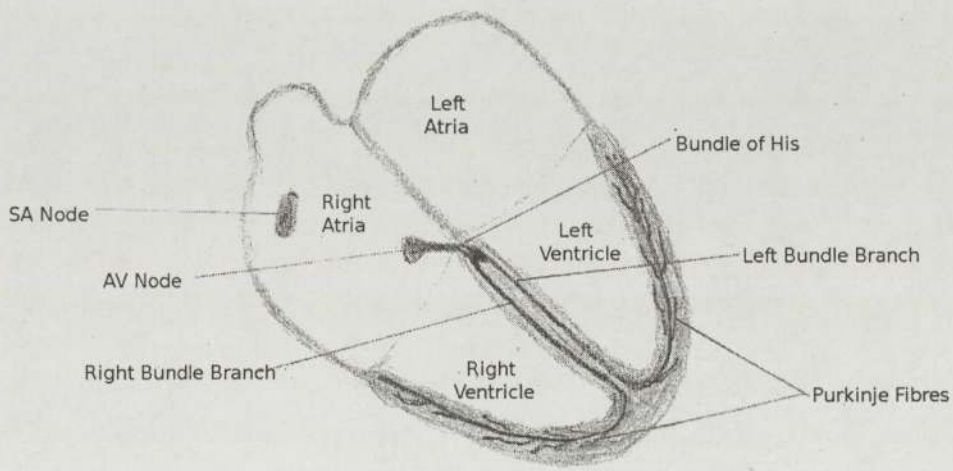

Figure 1.1: Cross Section of the Heart

The heart is the organ which pumps blood around the body, transporting oxygen and other essential molecules to, and removing carbon dioxide and other waste products from, the other organs of the body. It is divided into anatomically separate left and right sides and consists of four chambers, the right and left atria, situated at the top of the heart, and the right and left ventricles, which are beneath them (figure 1.1).

The cardiac cycle begins with deoxygenated blood from the inferior vena cava, the superior vena

cava, and the coronary sinus, entering the right atrium. The contraction of the right atrium forces this blood into the right ventricle, which in turn contracts to pump the blood into the lungs via the pulmonary artery. Oxygenated blood returns from the lungs through the pulmonary veins to the left atria, from where it is pumped to the left ventricle and then from there through the aorta to the rest of the body. Typically this cycle will be repeated around 100,000 times over a twenty four hour period.

1.2.2 Electrical Activity

The contractions of the heart are promoted by electrical changes which occur within the heart. In the normal functioning of the heart these changes are initiated by the sino–atrial (SA) node (figure 1.1), located in the upper wall of the right atrium near its junction with the superior vena cava, which contains the heart's main pacemaking cells. The impulse from the SA node spreads over both atria, prompting them to contract, forcing blood into the ventricles. This impulse is picked up by the atrio–ventricular (AV) node, located in the lower right atrium next to the interatrial septum, from where it is transmitted to the ventricles via the Bundle of His and the left and right bundle branches, thus promoting the contraction of the ventricles. After depolarisation both the atria and ventricles repolarise ready for the next heartbeat.

If the cells in the SA node fail to promote the necessary contractions in the heart, other groups of pacemaking cells, operating at slower rates, can take over the role of initiating these contractions. These cells are present, in decreasing order of discharge rate, in the AV node, the Bundle of His, and the ventricles.

1.2.3 Heart Disease

As stated in section 1.1 cardiac disease affects a large number of people every year. Cardiac disease comprises a number of different conditions, including myocardial infarction, strokes, arrhythmias and hereditary heart ailments. The ECG is a most useful tool for the detection of the last two conditions. Arrhythmias may affect the atria or ventricles (or both), and may be either faster (tachyarrhythmia) or slower (bradyarrhythmia) than the normal resting heart rate of between 60 to 100 beats per minute (bpm).

The most common arrhythmia which is encountered by doctors is atrial fibrillation, which is estimated to affect around half a million people in the United Kingdom [Arrhythmia Alliance 2005]. Although not a life–threatening condition in itself, it can cause the formation of blood clots which can then lead to strokes.

Other forms of tachyarrhythmia include supraventricular tachycardia, atrioventricular re–entry tachycardia, atrial flutter, ventricular tachycardia, and ventricular flutter. Bradyarrhythmias, which can cause fainting (syncope), are caused by atrioventricular block. This may be either first degree, type 1 second degree, type 2 second degree, or third degree.

The hereditary heart conditions comprise Wolff–Parkinson–White syndrome, long Q–T syndrome, and Brugada syndrome.

1.3 Electrocardiography

The various types of electrocardiography are introduced, along with the technical details of their recording. The relationship between the ECG and the activity of the heart is described.

1.3.1 Types Of Electrocardiography

Electrocardiography is an important diagnostic tool for the detection and classification of diseases of the heart, especially surface ECG which is a non-invasive technique (as opposed to catheter ECG, where the electrodes are inserted into the heart). There are three main types of surface ECG: standard 12–lead, exercise test, and ambulatory.

The standard 12–lead ECG is performed when a patient is admitted to a Coronary Care Unit (CCU). The purpose of this is to have a record for comparison against future ECG recordings obtained during the monitoring of the patient's heart during their stay in the CCU. The patient is in a supine position during the collection of these recordings.

Exercise stress testing is usually performed with the patient on a treadmill. The purpose of stress testing is to provoke myocardial ischemia, record the changes in the ECG, and to correlate these changes with the patient's symptoms. The test can be used for screening individuals who have high risk factors for coronary artery disease.

Both of the above methods of measuring the ECGs of patient's are only carried out for relatively short times (the standard 12–lead for less than a minute and the exercise stress test usually over 5, three minute, periods) and this means that they can miss important arrhythmias which occur only infrequently. In order to capture these episodes it is necessary to monitor the heart over longer periods of time using a Holter monitor.

Ambulatory ECG recording is carried out to enable the analysis of the patient's heart for periods of 24 hours or more, in order to gain information about arrhythmias which cannot be obtained during routine 12–lead ECG recording. The patient keeps a diary of their activities and symptoms during

the monitoring period and this is correlated to any irregularities in the ECG trace. Because a 24 hour ambulatory recording typically contains over 100,000 beats, it is unrealistic for a cardiologist to complete the analysis, and it becomes essential to automate the analysis using computational techniques.

1.3.2 Technical Details

The surface ECG works by measuring the electrical activity of the heart using several electrodes attached to the skin of the patient. For the standard 12–lead ECG there are three limb and six chest electrodes (the term lead, when applied to electrocardiography, referring to the signal obtained between two electrodes, and not the physical lead connecting an electrode to the recorder). One limb electrode is attached to the right arm, one to the left arm, and the third to the left leg. Together with the six chest electrodes they give twelve different channels of data, giving different electrical views of the heart's activity.

The recording of the ambulatory ECG uses fewer electrodes, the number of which is dependent on the number of data channels. The Medilog AR12 recorder used for the collection of data for this research has seven electrodes, producing three channels of data. These electrodes are positioned as in figure 1.2. The black (negative) and red (positive) electrodes produce the signal for channel one, the white (negative) and brown (positive) electrodes produce the signal for channel two, and the blue (negative) and orange (positive) electrodes produce the signal for channel three.

1.3.3 Relationship Of The ECG Recording To The Activity Of The Heart

For a normal heartbeat the initial feature of an ECG record is the P–wave, which is a representation of the depolarisation of the atria. In good quality recordings a notch can be detected at the peak of the P–wave, indicating the depolarisation of the right atrium occurring slightly before that of the left atrium due to the position of the SA node in the right atrium.

The next significant feature is the QRS complex, consisting of the Q–wave, R–wave, and S–wave, indicating the depolarisation of the ventricles. Of these three waves, the R–wave has by far the largest magnitude (with the other two waves being deflections in the opposite direction), making it an important fiducial point in ECG recordings.

The next peak is the T–wave, which denotes the repolarisation of the ventricles. This is sometimes followed by a very small peak, the U–wave, whose relationship to the electrical activity of the heart is not, as yet, fully understood.

It will have been noted that none of the above features relates to the repolarisation of the atria. This

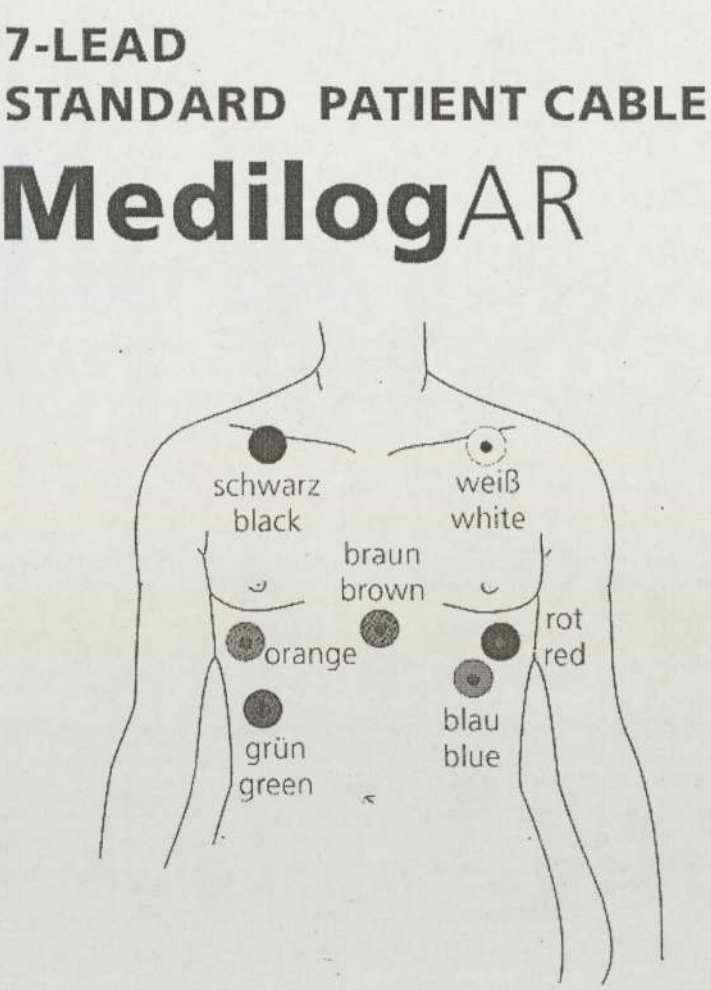

Figure 1.2: Electrode Placements for the Medilog AR12 Recorder

is due to fact that this stage in the cycle of the heartbeat occurs at the same time as the depolarisation of the ventricles, the far larger signal of which obscures any evidence of atrial repolarisation.

The shape of these waves and the intervals between them (figure 1.3) convey important information regarding the condition of the heart.

1.4 Research Problem

As has been previously mentioned, the typical length of an ambulatory ECG recording makes the automation of its analysis essential. Whilst the magnitude of the R–wave makes its detection easy, even in noisy recordings, the smaller magnitude of the P–wave causes problems due to the lower signal to noise ratio. Noise occurs as a result of several factors, the interference caused by the movement of the other muscles in the body, the contact between the electrodes and the skin becoming more insecure during the recording period, and external electrical interference. The most important of these is the

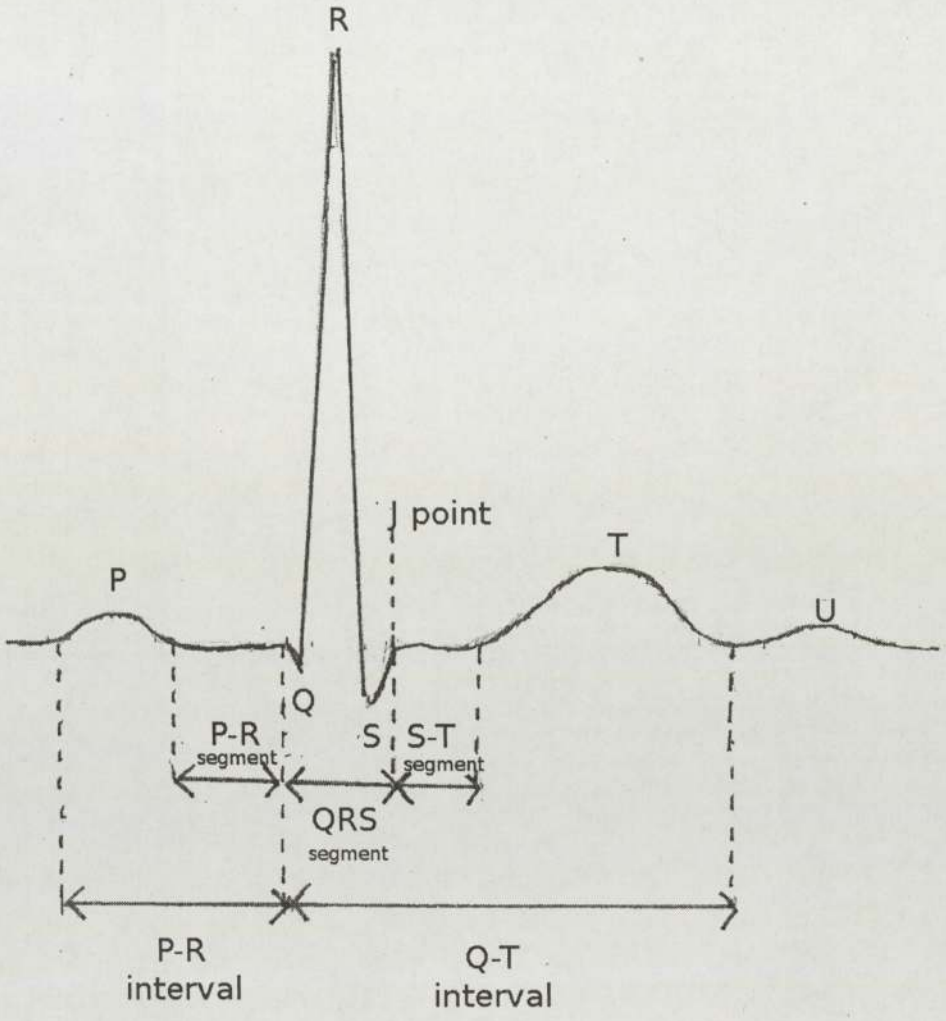

Figure 1.3: Principal Features of the ECG

influence of other muscles in the body, especially during exercise. Because accurate detection of the P–wave is required for the correct diagnosis of certain arrhythmias of the heart, the importance of a reliable automated P–wave detection system becomes apparent. To enable the correct diagnosis the position of the P–wave has to be accurately determined, in particular the onset of the P–wave which allows the accurate determination of the P–R interval (figure 1.3).

The focus of this research is to develop an accurate and reliable algorithm for the detection of P–waves in long–term ambulatory ECG recordings. This algorithm is then to be incorporated into the existing Medilog Darwin software to enable the accurate detection of P–waves in a commercially available software package. Of the various techniques used to attempt to solve this problem which were uncovered during a review of the recent literature, the wavelet transform was identified as the most promising and it is this approach which will be used in the current research.

1.5 Signal Processing

Any integrable stationary signal can be represented as a combination of sine and cosine waves, a fact which was first discovered in the early nineteenth century by Jean Baptiste Joseph Fourier whilst he was investigating the flow of heat [Fourier 1955]. A stationary signal is one which has constant mean and variance and which continues for an infinite length of time. In practice this means a periodic signal which has constant frequency over a long period of time as all signals have a beginning and an end. Using the Fourier transform, a stationary signal can be transformed from the time to the frequency domain.

Although the ECG is a quasi–periodic signal it is not stationary, varying both in amplitude and frequency over the course of a twenty–four hour recording. Therefore the Fourier transform is not well suited to the analysis of ECG signals. One approach to overcome the problem of analysing non–stationary signals is to use a windowed Fourier transform. With this method a small window of the signal is analysed at a time. The window is then moved along to the next section of the signal and the analysis repeated. The problem with this approach is that although small windows are good for the analysis of sudden changes in the signal they fail to pick up the lower frequencies. Conversely, larger windows are able to analyse the lower frequencies but are less able to locate the sudden changes in the signal. Another approach to this problem, which overcomes the difficulty of one window size not being able to adequately analyse all of the frequencies contained within the signal, is the wavelet transform.

1.6 Wavelets

The wavelet transform is a method of analysing signals by converting information regarding the time domain into the frequency domain, in the manner of a Fourier transform. The signal is analysed using different window lengths so that all of the frequency information is retained without losing its temporal context. This approach is termed multiresolution analysis and wavelets have been described as a "mathematical microscope" [Burke Hubbard 1998]. This means that the wavelet transform is suited to the analysis of non–stationary signals, such as ECG recordings.

A wavelet (ψ) is a function of zero average:

$$\int_{-\infty}^{+\infty} \psi(t)dt = 0 \tag{1.1}$$

which is dilated using a scaling parameter s and translated by u:

$$\psi_{u,s}(t) = \frac{1}{\sqrt{s}}\psi\left(\frac{t-u}{s}\right) \tag{1.2}$$

Wavelet transforms can be explained in terms of filters. Basically the signal is divided into two parts using a highpass and a lowpass filter. The lowpass filter is a moving average filter, and smooths the signal. The highpass filter is a moving difference filter, and detects the details of the signal.

For the Haar wavelet [Haar 2006], which is the simplest wavelet, the highpass filter is:

$$\frac{1}{\sqrt{2}}x(i) - \frac{1}{\sqrt{2}}x(i-1) \tag{1.3}$$

Whilst the lowpass filter is:

$$\frac{1}{\sqrt{2}}x(i) + \frac{1}{\sqrt{2}}x(i-1) \tag{1.4}$$

The outputs of both filters are then downsampled by removing all of the odd indexed coefficients. This can be done because all of the information regarding the signal is contained in the even indexed coefficients. From equation 1.3 it can be seen that each coefficient obtained from the highpass filter contains the scaled difference between the signal at that coefficient's time point and signal at the previous time point. Similarly, equation 1.4 shows that each coefficient obtained from the lowpass filter contains the scaled sum of the signal at that coefficient's time point together with the signal at the previous time point. The reason for downsampling the filter outputs is to reduce the computational load by keeping the vector length constant.

The downsampled output of the lowpass filter is then filtered again at the next scale in the same way, producing details at a lower frequency (the output of the highpass filter), and an even smoother approximation of the signal (the output of the lowpass filter). This procedure can be repeated to obtain information about even lower frequency components of the signal. If the signal length is a power of 2 then this procedure can be continued until only two data points remain. Due to the downsampling, all of the information concerning the signal can be contained in a vector of the same size as that of the original signal. All of the signal information is contained in the detail coefficients obtained at each stage of filtering, together with the approximation coefficients obtained from the final filtering stage. Therefore the signal may be reconstructed by reversing the process and recombining all of these coefficients using the reconstruction filters (which for the Haar wavelet are equations 1.5 and 1.6 for the highpass and lowpass filters respectively).

$$\frac{1}{\sqrt{2}}x(i) - \frac{1}{\sqrt{2}}x(i+1) \tag{1.5}$$

$$\frac{1}{\sqrt{2}}x(i) + \frac{1}{\sqrt{2}}x(i+1) \tag{1.6}$$

1.7 Summary Of Thesis

This chapter introduces the research topic and explains its relevance.

Chapter 2 contains a review of the recently published literature concerning the research problem.

Chapter 3 provides a description of the data used for this research, including the processing that was required and the production of data subsets.

Chapter 4 compares the ability of three types of wavelet transform (Coiflet, Daubechies and Haar) to denoise ECG recordings.

Chapter 5 describes the algorithme à trous and the software developed for its implementation using MATLAB. A comparison between the effectiveness of the Haar and quadratic spline wavelet transforms for detecting P–waves using this software is effected on an artificially produced P–wave.

Chapter 6 continues the comparison between the Haar and quadratic spline wavelets of their P–wave detection capabilities; this time on real data.

Chapter 7 concludes the thesis by providing a summary of the research carried out together with the results obtained, and suggests further research which could be undertaken to attain the goal of automatic P–wave detection in ambulatory ECG recordings.

Chapter 2

Literature Survey

2.1 Databases

A literature search was conducted to discover what previous, recent, research had been carried out on the detection of P–waves in ECG recordings. Twelve databases were searched using the keywords *ECG* and *P–wave*. Because only the most recent research was considered the databases were searched from the 1st of January 2002 until the 18th of February 2005. The databases searched were:–

Engineering Village 2 (including both the Compendex & INSPEC databases)
ISI Web of Science
IEEE Xplore
Blackwell Synergy
Science Direct
Wiley Interscience
Emerald
EBSCO
Ingenta
MetaPress
Swetswise
PubMed

Because so many databases were searched there were a large number of duplicate results found. Also some of the databases contained a large number of medical journals. The searches of these databases (ISI Web of Science, Ingenta, MetaPress and PubMed) were refined using the additional keywords, *detection*, *algorithm*, *automat** and *comput**.

The titles and abstracts of the papers obtained by the search were examined to determine which were likely to be of interest. Those papers which related to Holter or ambulatory ECG recording, algorithms, automatic P–wave analysis or detection, (artificial) neural networks, or arrhythmia were considered worthy of further study. A total of 13 papers were identified as being of interest from 2004 to date, whilst for the calendar years 2002 and 2003 respectively there were 6 and 5 papers identified.

When the papers had been read 15 were deemed to be of direct relevance (7 from 2004 up until the 18th of February 2005, and 4 each from 2002 and 2003) [Cañizares *et al.* 2003; Carrault *et al.* 2003; Clavier *et al.* 2002; Farrell, Xue, and Young 2003; Goutas *et al.* 2005; Hoffman *et al.* 2002; Kikawa and Oguri 2004; Marenco *et al.* 2003; Martínez *et al.* 2004; Senhadji *et al.* 2002; Sovilj, Jeras, and Magjarevic 2004; Sternickel 2002; Stridh *et al.* 2004; Tu *et al.* 2004; Tu, Zeng, and Yang 2004] with the others containing useful background information.

2.2 Summaries of Papers

In this summary, papers are divided into those which involve wavelet transforms, and those using alternative techniques. The papers involving the use of wavelet transforms are further subdivided into those which involve the analysis of P–waves, and those which analyse other features of the ECG. The papers using other methods are subdivided into those of direct interest to the current research, and those of incidental interest.

2.2.1 Wavelet Transforms — P–Waves

Martínez *et al.* [2004] initially distinguished two main groups of algorithms for the automatic delineation of ECG signals, those involving QRS detection and those involving wave delineation (where delineation is defined as the detection of the important parts of the wave, ie. start, end, and peak). Their concern in this paper was the delineation of ECG signals, and they provided a brief summary of the different techniques used by other researchers. They commented that whilst most QRS detection methods were validated using standard ECG databases, this approach was rarely encountered in papers using wave delineation algorithms, making comparison between different delineation methods difficult.

In this paper they evaluated a quadratic spline wavelet transform delineator, based on the work of Li, Zheng, and Tai [1995], on four standard ECG databases (Massachusetts Institute of Technology/Beth Israel Hospital (MIT/BIH), QT, European ST–T [Goldberger *et al.* 2000] and the Common Standards of Electrocardiography (CSE) multilead). They first offered a description of wavelet trans-

forms together with their usage before discussing the quadratic spline wavelet used. They stated that the wavelet–based approach can be considered as a differentiator filter–bank approach (see figure 2.1), the filter responses being

$$Q_k(e^{j\omega}) = \begin{cases} G(e^{j\omega}), & k = 1 \\ G(e^{j2^{k-1}\omega}), \ \prod_{l=0}^{k-2} H(e^{j2^l\omega}) & k \geq 2 \end{cases} \tag{2.1}$$

where G is the highpass filter, H is the lowpass filter, and k is the scale parameter of the wavelet transform.

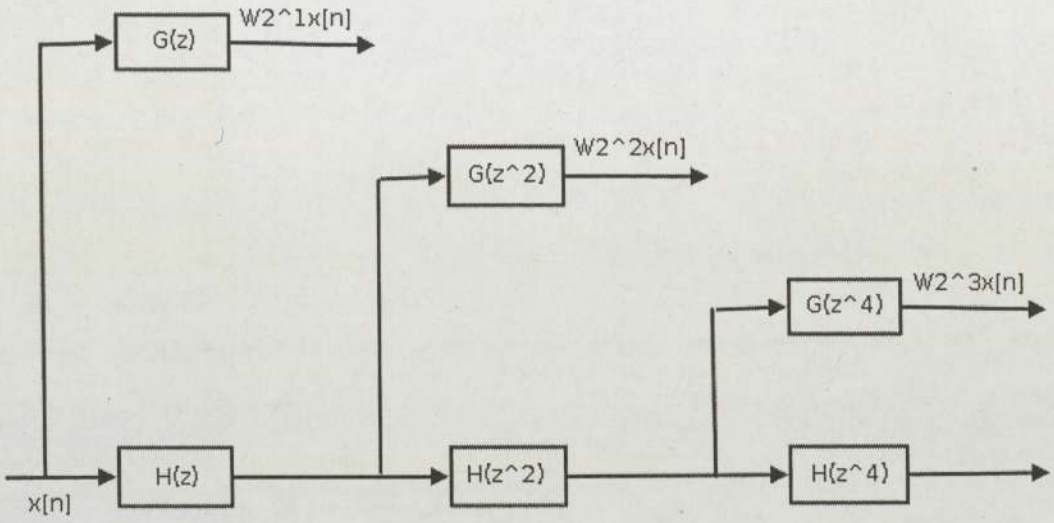

Figure 2.1: Filter–Bank Implementation of Algorithme À Trous

They commented on the fact that there has to be some adjustment made to the filters in order to handle signals sampled at frequencies other than 250 Hz: their suggestion being that the best solution is to compute a new set of filters with equivalent analogue frequency responses due to the time that it would take to resample the signal.

Most of the energy of ECG signals occurs within scales 2^1 to 2^5, with that attributable to QRS complexes occurring mainly in scales 2^1 to 2^4, whilst that attributable to P and T–waves occurs mostly in scales 2^4 and 2^5. However there is a problem with the influence of baseline wander becoming apparent in scale 2^5, and with high frequency noise affecting scales 2^1 and 2^2.

The algorithms described in this paper used the local maxima, minima and zero crossings at different scales to identify significant points in the ECG signal in four stages. Firstly the QRS complex was detected. Then the individual waves in the QRS complex together with the boundaries of the QRS complex were identified. Next the T–wave was detected and delineated and finally the P–wave was also detected and delineated.

The QRS complex was detected by setting a threshold value for each of the first four scales (2^1 to 2^4) and searching each scale for maxima above these values, the zero crossing between the maximum

positive and minimum negative pair at scale 2^1 being considered to mark a QRS complex.

The individual waves in the QRS complex were detected by searching scale 2^2 for local maxima before and after the marked position of the QRS complex which exceeded a threshold value. The zero crossings at scale 2^1 were then marked as the positions of the individual waves. The beginning of the QRS complex was then determined by searching scale 2^2 before the first wave in the complex for either a point below a threshold value or a local minima. A similar process was employed to detect the end of the QRS complex, this time by searching after the last wave in the complex.

The T–wave was detected by firstly defining a search window after the QRS complex, dependent on the RR interval, and then searching scale 2^4 for local maxima exceeding a threshold value. If none were found the process was repeated at scale 2^5. When a T–wave was identified the zero crossing in scale 2^3 was allocated as the T–wave. The onset and offset of the T–wave were determined in a similar manner to those of the QRS complex.

Finally the P–wave was detected and delineated in a similar fashion to the T–wave, this time a suitable search window being applied before the QRS complex.

– The results obtained for P–waves were comparable with those obtained using a low–pass–differentiator–based method on the same databases, the sensitivities for P–wave detection being respectively 98.87% and 97.7%, whilst those for positive predictivity were 91.03% and 91.17% respectively.

Clavier *et al.* [2002] used a wavelet transform, together with a hidden Markov model (HMM), for the automatic detection of P–waves in the signal obtained from standard lead II of a 12–lead ECG with the subject at rest.

Standard lead II was used because this is the lead in which the P–wave is most visible [Schamroth 1990]. They created their own database of signals, each sampled at 1 kHz for 1 minute, because although other databases were available they were not devoted to atrial fibrillation. Records from 145 patients were obtained, of which 82 had suffered from atrial fibrillation with the remaining 63, who had no such history, constituting the control group.

The signals were first subjected to a multiresolution analysis using Haar wavelets. These were chosen because they gave the best results out of all of the wavelets which were tried, possibly because of their ability to best approximate the first derivative of the ECG signal. The other types of wavelet tried were not mentioned in the paper.

The beats were then separated by detection of the QRS complex using the wavelet coefficients, with the signals being divided 350ms before the R peak.

The HMM used for P–wave delineation was based on a Bakis HMM (one in which each state has a transition to itself, its successor, or the successor of its successor) with 10 states, two each for the

P and T–waves and the QRS complex and four isoelectric states, but this was modified to allow the model to go backwards from some states and no transitions from a state to itself were allowed (figure 2.2). The estimation of the parameter set was accomplished from a training set of 240 beats taken from the control group which included all configurations encountered in the database. The most likely state sequence was estimated using the Viterbi algorithm [Viterbi 1967], a dynamic programming algorithm for finding the most likely sequence of hidden states that result in a sequence of observed events.

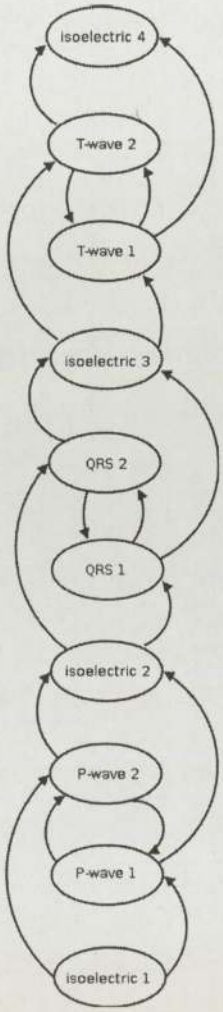

Figure 2.2: Permitted Transitions of the Modified Bakis HMM

After the P–waves had been isolated (ie. onset and offset detected), time, shape, spectral, and wavelet entropy parameters were measured to enable classification by discriminant analysis into two groups (history of atrial fibrillation and no history of atrial fibrillation). The parameters measured were P–wave duration, 5 different shape parameters, the energy, mean and variance computed in 3 different time and 6 different spectral bands from a Morlet continuous wavelet analysis (with values being obtained using both the whole P–wave and with QRS complex suppression, as at low frequencies

the Morlet wavelet extended to include the QRS complex), and 39 entropy parameters. This gave a vector of 117 parameters. Fisher discriminant analysis was used to reduce these parameters by removing those which were not significant. Two models were considered, one with only 3 features, and the other with 10, with both comprising the data for the wavelet analysis without a suppressed QRS complex.

The classifications obtained were poor, with both low sensitivity and specificity being obtained, the values being between 65% and 70%.

Sternickel [2002] claimed to have solved the problem of automatic P–wave detection in Holter recordings. Feature extraction was carried out by using three different types of discrete wavelet transforms (Daubechies, Coiflet and Harr), with the Coiflet 6 wavelet producing the best results. It was postulated that this was because the structure of the Coiflet 6 wavelet adapts best to most P–wave patterns.

The wavelet coefficients of the lower frequency bands contain the information regarding the general shape of the P–waves, while those of higher frequency bands contain information about the details of the P–wave shape. Therefore it is claimed that the exclusion of the highest one or two frequency bands will lead to better generalisation, and hence better classification results, when used as inputs of a neural network.

Classification into two classes (P–wave, or not P–wave) was then performed with a three layer feed forward neural network, with each coefficient linked to its own input neuron and, as it is meaningful that different frequency bands should be weighted equally, all input neurons which were members of the same frequency band linked to the same neuron in the hidden layer.

In spite of the claim that this method has solved the problem of P–wave detection in Holter recordings, and although 99.9% of P–waves were correctly identified, it produced one false positive result for every four true P–waves detected, these results being obtained when tested on an example dataset of a Holter recording of a healthy subject, the sampling rate being 200 Hz.

Sovilj, Jeras, and Magjarevic [2004] developed a real time P–wave detector using wavelet analysis. Their motivation for doing so was that real time measurements are necessary for applications such as P–wave synchronised pacing [Wagner 2001; Andreoli *et al.* 1983] (where heart stimulation is synchronised with the P–wave).

Their system was comprised of two different wave detectors, the first using adaptive thresholds for ECG segment detection, and the second using a more reliable detector which was based on a signal backwards searching algorithm. The second detector also processed data on the shape of the P and T–waves using statistical methods. This was done in order to provide information regarding the current

P and T–wave parameters for use in the prediction of expected P–wave parameters. The calculated expected P and T–wave parameters were then used by the first detector for making decisions based on the P–wave appearance, with P–wave parameters being used to make positive P–wave predictions and the T–wave parameters being used to prevent incorrect P–wave predictions.

The sampling rate used was 250 Hz, with decomposition of the signal being accomplished using a discrete wavelet transform with calculations made using the Mallat scheme [Mallat 1989]. The third and fourth wavelet transform scales were used for detection of the various components of the ECG signal because all of the components were clearly visible at these scales. Although all of the components were also visible at the fifth scale this scale was not used for two reasons, firstly because the computational time required was too long for a real time detection system, and secondly the influence of baseline drift became apparent at this scale.

In the detection of the various components of the ECG signal, the first to be detected was the QRS complex, followed by the P–wave, and finally the T–wave.

Selectivity was deemed to be more important than sensitivity in the backwards searching detector because its main purpose was the determination of P–wave parameters, and therefore it was considered better to fail to detect some real P–waves rather than to input parameters from false P–waves.

In the abstract of this paper a sensitivity of 98.5% was claimed for the detection of P–waves, although no quantitative results were presented within the main body of the paper.

Senhadji *et al.* [2002] decided to use a dyadic wavelet transform to perform adaptive QRS–T cancellation in order to detect hidden P–waves. Based on their previous research [Senhadji *et al.* 1995] they had already identified wavelets as a suitable method of analysing ECG signals. Their use of the dyadic wavelet transform was due to its good temporal localisation properties, as well as the speed of the calculations involved.

Two ECG channels were chosen so that one of them exhibited small amplitude P–waves. The signals from both channels were decomposed using wavelet transforms of scales 1 to 9 inclusive, with only the local extrema (exceeding a given threshold dependent on the scale) being retained to represent the signal. The extrema obtained from the ECG signal with the small amplitude P–waves were then subtracted from those obtained from the other ECG channel and the new values so obtained used to reconstruct the QRS–T cancelled signal. At this stage baseline drift and noise due to muscle activity was also removed.

This method was evaluated on records from the MIT/BIH database to determine how the signal to noise ratio of the P–waves was altered, with boxplots being used to illustrate the effect. The signal to noise ratio was significantly improved after adaptive QRS–T cancellation, and was greater than

that obtained by the other methods with which it was compared.

Carrault *et al.* [2003] used an adaptive ventricular activity cancellation module based on wavelet decomposition [Senhadji *et al.* 2002] and an artificial neural network (ANN) classifier for P–wave detection. They wished to distinguish between normal and abnormal beats using sinus rhythm, ventricular bigeminy, Left Branch Bundle Block (LBBB), Premature Ventricular Contraction (PVC) and Mobitz type II records taken from the MIT/BIH database.

There were two steps to this process, offline, where a set of high level symbolic characterisations of cardiac arrhythmias was produced using a Prolog type syntax to specify the characteristics of each class of arrhythmia, and online, where the signals were analysed and arrhythmias identified by matching the symbolic representation of the signal to prestored characterisations.

The offline learning process used Inductive Logic Programming (ILP) with the learning examples being a mixture of real and simulated signals, the simulated signals being obtained from the CARMEN Cardiac Model. High level characterisation of the cardiac arrhythmias was considered desirable because this was easily understandable by cardiology experts.

A multi layer feed forward ANN with one hidden layer with 10 nodes was used for the P–wave detection. Candidate P–waves were first classified as either P–waves or not P–waves, and then those that were classified as P–waves were classed as either normal beats or one of the arrhythmias. This was done by matching their symbolic representation with prestored characterisations.

Although normal beats, Mobitz type II, and ventricular bigeminy were in general correctly classified, a large number of PVC and LBBB records were given incorrect classifications.

2.2.2 Wavelet Transforms — Other ECG Features

Shyu *et al.* [2004] used the quadratic spline wavelet transform of Li, Zheng, and Tai [1995] for both QRS complex and Ventricular Premature Contraction (VPC) detection from Holter ECG traces. The features extracted provided the inputs for a fuzzy neural network. This method had difficulty in differentiating between LBBB and VPC, but otherwise the classifications obtained were reasonably accurate, the types of beat examined being normal, VPC, LBBB, Right Branch Bundle Block (RBBB) and Atrial Premature Contraction (APC).

Tasoulis *et al.* [2004] used principal component analysis together, with noise reduction using Daubechies–type wavelets, for ischemia episode detection. Five components were identified which contained 98% of the information required for the reconstruction of the ST–T segment. The European ST–T database (containing ambulatory recordings from patients who had either been diagnosed with,

or were suspected to suffer from, myocardial ischemia) was used for evaluating five different neural network training methods. In spite of the claim that the percentage of beats correctly classified was good, the numerical results for correct classification presented in the paper appear poor, with the best performance being reported for record E111 (ca 92% correct classification) and some other records only having correct classifications of 70% or less.

2.2.3 Wavelet Transforms — Summary

Wavelet transforms have become a popular choice for the analysis of ECG data and P–waves in particular, as can be seen from the fact that 1/3 of the papers found during the literature survey used them, with 3/4 of those papers using them for P–wave analysis. This popularity can be attributed to their suitability for the analysis of non–stationary signals.

Whilst a large proportion of researchers in this field use wavelet transforms there appears to be no general consensus as to which wavelets are the best to use, with the choice seemingly being influenced by the data being analysed and which sections of the ECG are of most interest. Haar [Clavier *et al.* 2002], Coiflet [Sternickel 2002], quadratic spline [Martínez *et al.* 2004], Morlet [Clavier *et al.* 2002], and Daubechies [Senhadji *et al.* 1995] wavelets have all been used with some success.

The third and fourth scales of the wavelet transform, regardless of the wavelet used, are those most favoured for P–wave analysis, being least susceptible to the effects of both noise and baseline wander. Scales greater than the fourth scale also use too much computer time for real time analysis [Sovilj, Jeras, and Magjarevic 2004].

2.2.4 Other Methods — Relevant To Research

Cañizares *et al.* [2003] presented algorithms for the identification of P–waves, the end of the T–wave, and the end of the QRS complex. P–waves from ECGs obtained by the authors were correctly identified in 97% of cases, as defined by two experienced cardiologists, with 2% being undetected and 1% being incorrectly identified. P–waves were only searched for in the TQ interval.

Farrell, Xue, and Young [2003] described the improvement made to the commercially available MAC–RHYTHM atrial analysis program of the GE Medical Systems 12SL resting analysis system. The new version of the program was then compared with the existing version, and also with the classifications obtained from cardiologists. The new version of the program had a discordance rate of 4.1% with the cardiologists classification, compared to 6.9% for the old version.

Goutas *et al.* [2005] used a digital fractional order differentiation based algorithm [Ferdi *et al.* 2003] to detect and delineate P and T–waves. The P–waves were detected by first detecting the R wave and then defining temporal search windows 225ms to 70ms before the peak of the R–wave. Fractional order differentiation is based on differentials of the form $\frac{d^{1/2}}{dx^{1/2}}$, $\frac{d^{1/3}}{dx^{1/3}}$ etc. (see Oldham and Spanier [1974] for further details). The concept of fractional order differentiation is not easy to understand because differentials of this form have no obvious physical interpretation. The algorithm used combined a lowpass filter for electromyographic noise (electronic noise caused by the muscles) removal with a highpass filter for eliminating baseline deviations and slope enhancement (equivalent to the application of a notch filter).

They claim that this method detects P and T–waves even in noisy signals from the MIT/BIH database, although they admit that their algorithm has yet to be validated against experienced cardiologists. However when Ferdi *et al.* [2003] compared their technique with other methods they found that it was inferior to an algorithm based on wavelet transforms proposed by Li, Zheng, and Tai [1995]. It is also not explained how the use of fractional differentials is supposed to improve the extraction of features from noisy signals.

Hoffman *et al.* [2002] based their work on a Multiple Model Adaptive Estimator (which is a bank of parallel Kalman filters, each with a different filter model and an algorithm to test for the adequacy of the assumed model in each filter) together with a threshold filter–switching algorithm to mimic a human analyst, firstly by searching an ECG trace for usable temporal marks and then branching out to find other key points. They assumed that the normal ECG wave sequence applied (P–wave followed by QRS complex and then a T–wave). The algorithm was tested on the same MIT ECG database recordings that had been used for development of the algorithm.

Kikawa and Oguri [2004] tackled the problem of non–heartbeats (false positive identifications) being incorrectly classified as heartbeats in the automatic analysis of Holter recordings using a Support Vector Machine (SVM). A SVM works by using a maximum–margin hyperplane to split training examples labelled "yes" or "no" into two groups such that the distance from those examples which are closest to the hyperplane is maximised. The only problem addressed in this paper was the incorrect detection of R–waves.

Marenco *et al.* [2003] developed an automatic T–wave subtraction algorithm for uncovering hidden P–waves in standard 12 lead ECG traces. A trial was then carried out to see if hidden P–waves which had been artificially induced by atrial pacing could be reliably extracted from within T–waves. One group (21 subjects) were patients who had undergone diagnostic electrophysiologic

study, whilst the other group (10 subjects) had undergone radiofrequency catheter ablation. P–waves were successfully uncovered from all patients in both groups.

Stridh *et al.* [2004] were concerned with the characterisation of atrial arrhythmias using time–frequency analysis, the hypothesis being that different types of atrial signal can be categorised by their fundamental frequency together with a number of harmonics. They discovered that slower fibrillation rhythms were associated with waves having a sharper shape, whilst faster rhythms were associated with waves having a more sinusoidal shape.

Tu, Zeng, and Yang [2004] firstly detected the QRS–T complex using histograms and genetic algorithms, subtracted this from the ECG signal and then used fractal dimension and Shannon entropy methods for the analysis of P–waves, or f–waves if atrial fibrillation was present, where the fractal dimension is a dimension which does not take an integer value. It was found that the mean value of the fractal dimension of P–waves was significantly greater than that of f–waves. The mean value of the Shannon entropy of P–waves was found to be smaller than that of f -waves, although no measure of variability was given. The data used in this paper came from ECGs recorded by the authors and also from the MIT/BIH database.

Tu *et al.* [2004] used the same QRS–T detection and subtraction method as above, but then used three different methods for differentiating between P–waves and f–waves. The techniques used were the power spectrum, the auto–correlation function (which describes the degree of correlation of values of a signal at different times), and the cross–correlation function (which describes the correlation of two different signals sampled at different times). ECG traces were obtained from the MIT/BIH database.

2.2.5 Other Methods — Not Directly Relevant To Research

Aytemir *et al.* [2005] used a commercial ECG system (Kardiosis ARS–LP) for Signal–Averaged Electrocardiography (SAECG) for the recording and analysis of P–waves. SAECG is a technique used to reduce the noise in the ECG signal by aligning and averaging the signals obtained from orthogonal X, Y and Z leads [Marriott and Conover 1998], with the averaging being carried out over at least 250 beats. They compared two groups of hypertensive patients, one group of 44 subjects with a history of paroxysmal atrial fibrillation (PAF) and the other group of 50 subjects who had no such history. It was found that the SAECG P–wave duration was a significant predictor of PAF, with the duration being significantly greater in those patients who had experienced PAF.

Choi *et al.* [2002] firstly used the QRS–T subtraction algorithm developed by Bard Electrophysiology, as used by Marenco *et al.* [2003], and then performed pulmonary vein pacemapping to obtain an individual P–wave catalogue for each vein in each patient. The use of pulmonary vein pacemapping is not relevant to the current research question.

Ho and Yeh [2005] developed a hybrid architecture combining correlation dimension (which can be simply defined as a numerical measure of the degree of chaos), together with feature extraction techniques into a case–based reasoning technique (which uses a library of previously solved problems to solve new problems) for ECG analysis. The turning points of the ECG signals were identified using a Min–Max Turning Points Selection algorithm and these were then used as the features to identify the most promising cases from a specific k means cluster. They used records from the MIT/BIH database for training and evaluation, and obtained sensitivities and specificities of greater than 99% for PVC, LBBB, and normal beats.

Madias [2004] compared the detection of P–waves recorded in three ways (standard ECG, "Lewis lead" [Goldman 1986; Dunn and Lipman 1989], and intracardiac ECG using a saline–filled central venous catheter). Although only 28 patients were studied, he concluded that there was no significant difference between the standard ECG and the "Lewis lead", but that the intracardiac ECG gave recordings with significantly larger P–waves, not surprisingly given Wilson's proximity effect, which states that the strength of the electrical signal from the heart diminishes algebraically with the distance from the heart, but that beyond a distance of 15cm the decrease becomes negligible.

Papaloukas *et al.* [2002] developed an ischemia detection method based on artificial neural networks. The QRS complex was detected and used to identify the J point (end of the QRS complex and start of the S–T segment), which then allowed an ischemic window to be identified. Principal Component Analysis was used to reduce dimensionality. The results obtained were better than those from a rule based approach, with sensitivity and selectivity of 90%, although their method couldn't be used to provide any interpretation for the decisions obtained.

Ros *et al.* [2004] concerned themselves with the development of an algorithm to detect which subjects will develop paroxysmal atrial fibrillation from ECG traces in which no atrial fibrillation is present. The data examined was that used in the 2001 Computers in Cardiology Challenge[1]. Parameters from the detected P–waves were used to attempt to differentiate between patients who had experienced atrial fibrillation and those who had not.

[1] URL:http://www.physionet.org/physiobank/database/afpdb/

Ryan *et al.* [2004] presented a method of representing the electrical depolarisation and repolarisation of cardiac cells using real–time volumetric animation techniques. Whilst this work is interesting it has no direct relevance to the detection of P–waves.

2.3 Summary

Various techniques have been utilised for the analysis of P–waves, including support vector machines [Kikawa and Oguri 2004], time–frequency analysis [Stridh *et al.* 2004; Tu *et al.* 2004], genetic algorithms [Tu *et al.* 2004], fractional order differentiation [Goutas *et al.* 2005] and entropy [Tu, Zeng, and Yang 2004]. However the most popular method used by researchers has been the wavelet transform.

Whilst many researchers have used the wavelet transform, and there is a general consensus that the P–waves are best analysed using the level four wavelet coefficients, there is no such consensus regarding the most suitable wavelet to use, with various wavelets being preferred by different researchers. These researchers have also used different ECG databases for the development and testing of their algorithms; a fact commented on by Martínez *et al.* [2004], who suggested that this makes the comparison between these algorithms problematic.

Some researchers have confined themselves to the detection of P–waves, whilst others were interested in the classification of atrial disorders, both as distinct from normal sinus rhythm and into different disorders.

In the light of the above it would be sensible to conduct the comparison of algorithms on standard databases, as recommended by Martínez *et al.* [2004]. Consideration should also be given to which wavelet will best generalise to the detection of P–waves, rather than using different wavelets for different problems and databases.

Chapter 3

Ambulatory ECG Recordings

3.1 Data Collection

The data used for this research was supplied by the Cardiology Products Division of Huntleigh Healthcare Limited, and consisted of six ambulatory recordings of approximately 24 hours each. These recordings were obtained using a Medilog AR12 digital Holter recorder (manufactured by TOM Medical) at 250 Hz and 16 ADC bit rate. These recordings were all of good quality. All of these recordings had been anonymised and no information regarding the age, gender, or health status of the subjects was available.

The recordings were supplied in both the Darwin and PhysioNet formats. The Darwin format was produced by the Huntleigh Healthcare Medilog Darwin software, whilst the PhysioNet format is that used by PhysioNet [Goldberger *et al.* 2000], an online research resource of biomedical signals. Neither of these formats are compatible with the software which was to be used for the research (MATLAB). Therefore the recordings in the PhysioNet format were used because a software package (WaveForm DataBase [WFDB]) was readily available to convert them into a format suitable for use with MATLAB.

For each recording there were three channels of data. The Medilog Darwin software produced three data outputs for each channel; one containing the raw data, another containing the bandpass filtered data, and a third containing the first derivative of the bandpass filtered data. The Medilog Darwin software also classified each beat according to its QRS morphology.

3.2 Data Processing — Preprocessing

Initially the data had to be converted into a format suitable for use with MATLAB. This was accomplished using the previously mentioned WFDB software package available from PhysioNet[1]. Struc-

[1] URL:http://www.physionet.org/physiotools/wfdb.shtml

tural arrays were produced containing the filename; the data from each of the three channels, raw, first derivative of the lowpass filtered, and bandpass filtered; the beat annotations; the time of the beat annotations, expressed as a data point number; the sampling frequency; the time in seconds; and the beat annotation times in seconds.

3.3 Data Processing — Preparation Of Data Subsets

Various data subsets were produced: single beats, isolated P-wave sections, and a continuous record of half an hour from one of the records ((aston2)). Records (aston4) and (aston6) were excluded from these data subsets due to the fact that they both contained a substantial number of non–normal beats, as designated by the Medilog Darwin software. Development and test sets of single beats, and the isolated P–waves from these beats, were produced from the other four records, each of which contained an equal number of beats from each record.

3.3.1 Single Beats

A random selection of 100 beats was taken from records *aston1, aston2, aston3*, and *aston5* (25 beats from each record). These records were chosen because they were almost totally composed of supraventricular (normal) beats with coupled P–waves (see section 3.4). The random selection was obtained by producing a random permutation of the integers from 1 to the length of the vector containing the beat annotations. The first 25 numbers of this new vector were taken as the random beat numbers, and their corresponding beat annotations checked to ensure that they corresponded to normal beats (i.e. that the beat annotation was 1). All of the beats selected in this way were found to be designated as normal by the Medilog Darwin software.

The process of extracting the beats from the records is described below. The time of the beat annotation was taken to be the time of the beat, with the duration of the beat taken to be from the one hundredth data point before, to the one hundred and fifty fifth data point after, this time, giving a total beat length of 256 data points. The reason for choosing the beat length to be 256 data points was because this length was the smallest power of 2 which would include all of the data from a complete beat (assuming a normal resting heart rate of about 70bpm), and the wavelet transform analysis to be performed was dependent on the data size being in powers of 2. The beat annotation time was the time of the maximum value of the first derivative of the lowpass filtered data as provided by the Medilog Darwin software. This point usually occurred on the initial upward slope of the R–wave, but was sometimes on the downward slope after the peak of the R–wave. Because a beat length of 256

data points was taken, portions of adjacent beats were included when the heart rate was greater than 60bpm. The disparity between the length of the isolated sections before and after the beat marker is explained by the differing length of time of the heartbeat before and after the QRS complex (figure 1.3).

For the denoising analysis which was to be performed, normalised versions of these datasets were produced. This was because a thresholding method was to be used for this purpose, and the noise level in the signals was unknown. Normalisation was achieved by the use of two separate methods. Firstly, the default method of the WaveLab software was employed. This method calculates the first level detail coefficients by the use of the Quadrature Mirror Filters (QMF) of the Coiflet 3 wavelet, next determines the absolute median value of these coefficients, and then divides the signal by this value. Secondly, normalisation was achieved by determining the mean and standard deviation of the signal, subtracting the mean value from the signal and then dividing by the standard deviation.

In order to analyse single beats using wavelet transforms without also including portions of adjacent beats, it is necessary to take only the data points corresponding to the single beat, and then to pad out the vector of data to 256 data points. This was accomplished by adapting the above method for beat isolation by taking into account the heart rate using the beat times for the preceding and succeeding beats. The start of the beat was defined as the time point 40% of the time between the beat and the preceding beat before the beat annotation, and the end of the beat as 60% of the time between the beat and the succeeding beat after the beat annotation. However, if these times were greater than one hundred data points before, in the case of the start of the beat, or one hundred and fifty five data points after, in the case of the end of the beat, the beat was isolated as described in the original algorithm in order to preserve the 256 data point length for the wavelet analysis.

The individual isolated beats could have been padded with zeros at each end to produce a vector of 256 data points. However this approach would have led to spurious discontinuities being detected in the signal, due to the end values of the isolated beats not being equal to zero. The approach actually implemented therefore, was to pad the beginning of the signal with the first value of the isolated beat, and the end of the signal with the last value of the isolated beat.

A set of data was also obtained by isolating the 100 random beats together with their preceding and succeeding beats. The purpose of this data set was to show the context of the random beats within their ECG records.

A data set containing a further 100 random beats, taken from the four records previously used and checked to ensure that no beats from the first 100 random beats were replicated, was produced as a test set. Again, a data set containing the preceding and succeeding beats was obtained.

3.3.2 Preparation Of Data Subsets — Isolated P–Waves

Because the initial objective of the research was the detection and classification of coupled P–waves the analysis of complete beats was not necessary. As the coupled P–waves were associated with a QRS complex, the position of which had already been determined by the Medilog Darwin software, the section of the beat containing only the P–wave could be isolated. This was accomplished by adapting the software previously written for producing the isolated single beats. Software was also written to produce the same isolated P–wave sections for the raw data from channels 2 and 3, and for the bandpass filtered data for all three channels.

Normalised versions of each of these data sets was also produced using the same two methods previously used for normalising the isolated whole beats. This was in order to enable a numerical comparison between the wavelet coefficients obtained from different signals.

The P–wave sections were visually examined to determine if the P–waves had been successfully isolated from the rest of the beats. This was carried out by comparing the corresponding isolated single beat with the isolated P–wave section, and seeing if the P–wave in the isolated single beat was present in the isolated P–wave section, whilst other components of the ECG signal (QRS complex and T–wave) were absent.

The data set containing the second set of 100 random beats was also processed in the same manner to produce a data set containing only the isolated P–wave sections.

3.3.3 Preparation Of Data Subsets — Record *aston2*

The aim of this research is to be able to detect P–waves from a continuous ambulatory ECG recording. Therefore record *aston2* was randomly chosen, and the first half hour of data, containing 1,895 beats, extracted for analysis. All of the beats in this portion of the record were classified as being normal according to the Medilog Darwin software, except for the first beat, which, on examination, turned out to be an artifact caused by the recorder and not a beat at all. The variables contained in this data file were the same as those in the data file of the random selection of 100 beats and the signals were un–normalised.

3.4 Data Description

A visual examination of the raw data was undertaken using MATLAB. This revealed that the P, R and T–waves were all positive in data channels 1 and 2 for all records, as usual in ECG recordings of the normal heart in most leads. The polarity of the R and T–waves in data channel 3 differed

Record	Total Beats	Ventricular Beats	Percentage Ventricular
aston1	113769	115	0.10
aston2	94537	2	less than 0.01
aston3	92447	351	0.38
aston4	100738	6380	6.33
aston5	110016	42	0.04
aston6	80419	19619	24.40

Table 3.1: Ventricular Beats

according to the record, although the P–waves were positive in all six records. For records *aston1* and *aston4* both the R and T–waves were negative, whilst in records *aston3, aston5* and *aston6* they were both positive. Record *aston2* exhibited a negative R–wave but a positive T–wave. These differences seem a little strange as it was to be expected that the electrode positions, as shown in figure 1.2, would have been the same for each recording, and that therefore the direction of each wave would have been the same for each channel 3 recording.

The visual examination also detected a difference in the amplitude of the signals from record *aston3* as compared to the other five records, with a comparison of the channel 1 raw data revealing that the maximum R–wave value was around 100 times greater than that in the other five records. This difference in the intensity of the *aston3* signal as compared to the other records is possibly explained by the recording settings used, although no information regarding these was supplied with the recordings.

Huntleigh Healthcare were contacted regarding both the differences in the polarity of the waves in different recordings and the increased amplitude of the *aston3* recording, but were unable to provide answers to these questions.

The data was examined to determine the number of ventricular beats in each record (QRS complexes without associated P–waves) as determined by the Medilog Darwin software. From the results in table 3.1 it can be seen that records *aston4* and *aston6* contained substantial numbers of QRS complexes unassociated with a P–wave.

Because the wavelet analysis was dependent on the individual beats being successfully isolated, the lengths of the individual beats were consequently of interest. Therefore an investigation was carried out into the distribution of the individual, isolated, beat lengths of the six records. The beat lengths were determined using the method detailed in the preprocessing section for the isolation of individual beats. The number of beats in each record with a length of over one second (table 3.2), together with relevant statistics (mean, mode, median, and standard deviation: table 3.3) were calculated using a specially written script. Histograms of the distribution of the beat lengths for each record were also produced (figure 3.1).

Record	Length (h:m:s)	Total Beats	Beats Over 1 sec	Percentage Over 1 sec
aston1	23:07:36	113769	800	0.70
aston2	22:36:06	94537	25260	26.72
aston3	22:14:12	92447	16999	18.39
aston4	22:09:39	100738	5803	5.76
aston5	21:25:59	110016	516	0.45
aston6	21:14:43	80419	32230	40.08

Table 3.2: Beats Longer Than 1 Second

Record	Mean	Mode	Median	Standard Deviation	Rate (bpm)
aston1	0.73s	0.68s	0.71s	0.092s	82
aston2	0.86s	0.67s	0.83s	0.202s	70
aston3	0.87s	0.82s	0.81s	0.272s	69
aston4	0.79s	0.91s	0.79s	0.152s	76
aston5	0.70s	0.61s	0.67s	0.129s	86
aston6	0.95s	0.78s	0.89s	0.282s	63

Table 3.3: Beat Length Statistics

The large number of beats with lengths greater than 1 second indicate that there is likely to be a problem with the correct isolation of complete individual beats using the current algorithm, with portions of the beats being excluded.

3.4.1 Isolated P–Waves

Visual examination of the isolated P–wave data for both the development and test sets revealed that a few of the P–waves had not been correctly isolated. In the development set it appeared that the isolated section for beat number 56 did not contain a P–wave (figure 3.2), whilst the isolated sections for beat numbers 58, 67, 71, 72, and 73 contained only part of the P–wave for that beat (figure 3.2). In the development set the isolated sections for beat numbers 66, 68, and 74 only contained part of the P–wave of those beats.

All of the isolated P–wave sections in both data sets in which the P–wave was not correctly isolated were from record *aston3*. As mentioned earlier in this section the signals from this record had an amplitude approximately 100 times greater than those from the other five records.

An examination of the first derivative of the lowpass filtered data of all of the records revealed that whilst those of the other five records all exhibited one large minimum/maximum pair, indicating the position of the R–wave, those from record *aston3* had several minimum/maximum pairs of equal magnitude around the time of the R–wave (figure 3.4). Due to the fact that the isolation of the P–wave sections is based on the location of the R–wave, as determined by the maximum value of the first derivative of the lowpass filtered data, the failure of the algorithm to correctly isolate some

P–waves in record *aston3* may be attributed to the uncertainty regarding the exact position of the R–wave thus caused.

3.5 Summary

Several differences were found between the six recordings. Records (aston4) and (aston6) both contained a high number of ventricular beats, with nearly 1 in 4 of the beats in record (aston6) being so designated.

All recordings contained some beats whose length was greater than one second, with records (aston2), (aston3) and (aston6) containing a particularly high proportion of such beats. This indicates a potential problem with the isolation of single beats, and also with the correct isolation of P–wave sections, as the algorithm is based on the heart rate not being less than 60bpm. This potential problem could, however, be solved by increasing the window length when the heart rate is lower than 60bpm.

The polarity of the R and T–waves differed in the channel 3 recordings, although the P–waves exhibited the same polarity in all records for all three channels.

Record (aston3) consisted of signals with an amplitude 100 times greater than that of the other five recordings. Some P–waves had not been correctly isolated in record (aston3), with the search window determined by the algorithm containing only part of (in five cases), and in one case (beat number 56) no part of, the P–wave. An examination of the first derivatives of the recordings revealed that whilst those of the other five records exhibited one large minimum/maximum pair, denoting the position of the R–wave, those of record (aston3) had several large minimum/maximum pairings. This had the effect of making the determination of the position of the R–wave unreliable for this recording.

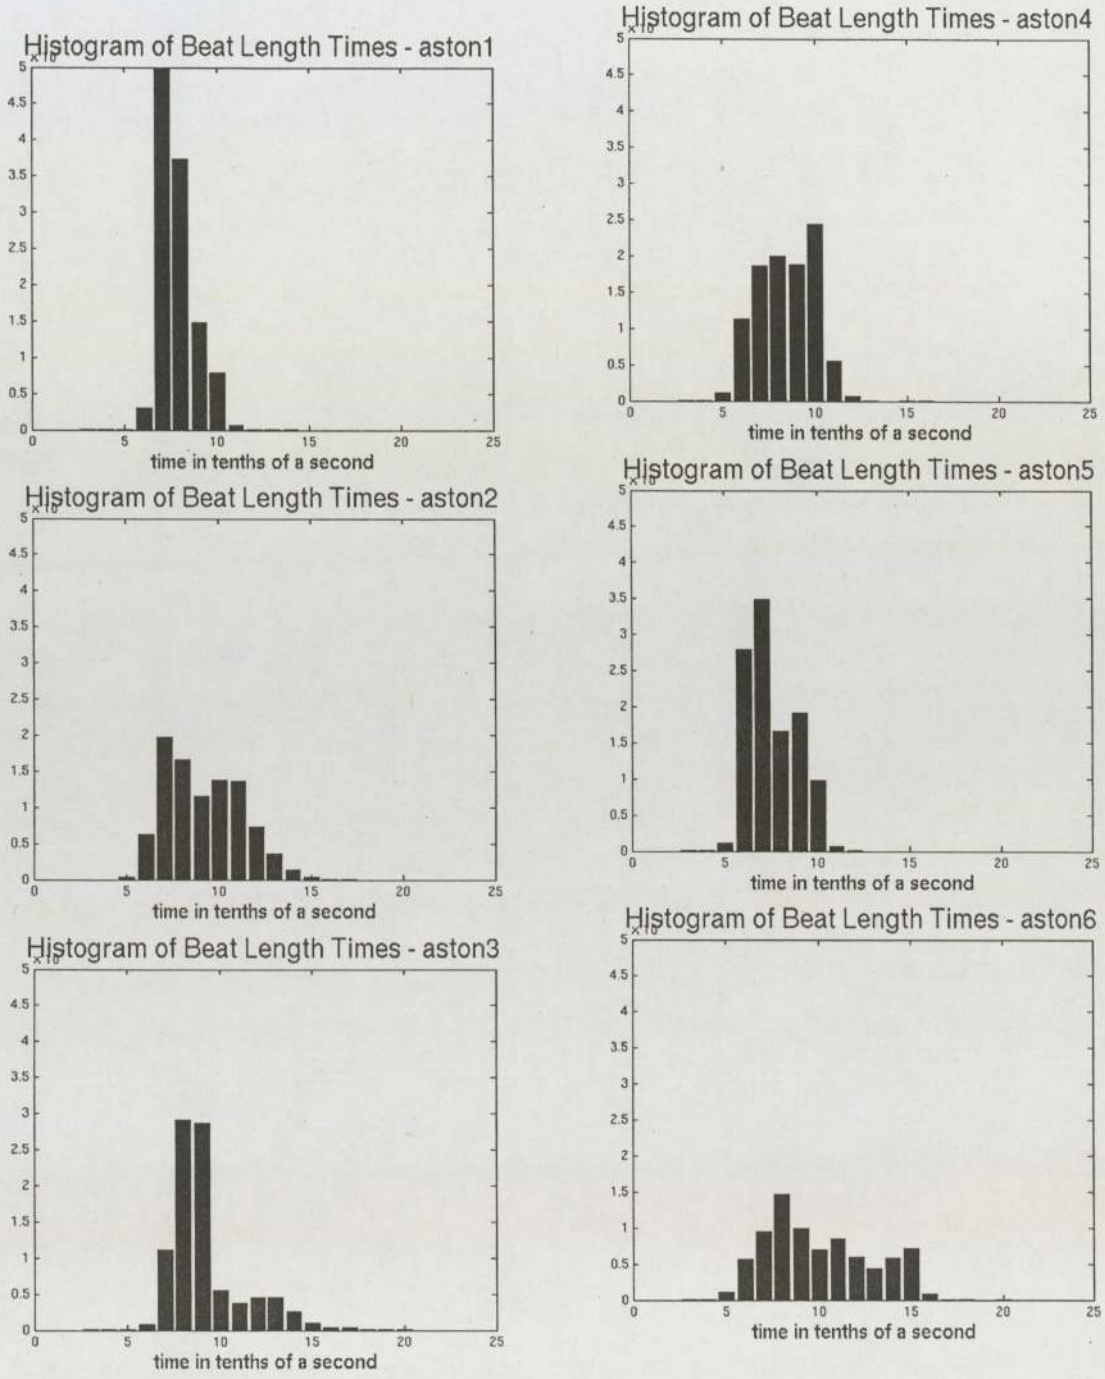

Figure 3.1: Distributions of Beat Lengths for Records *aston1* to *aston6*

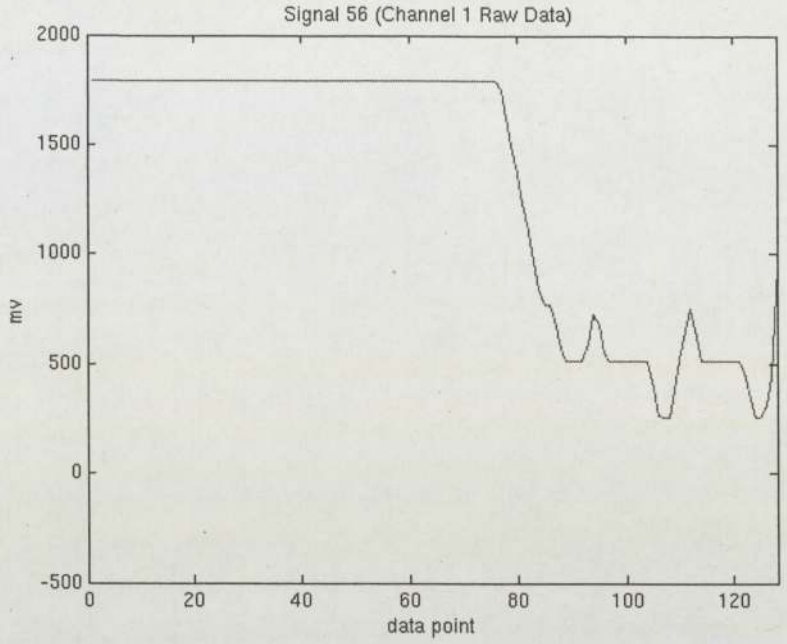

Figure 3.2: P–Wave Missing From Isolated Section

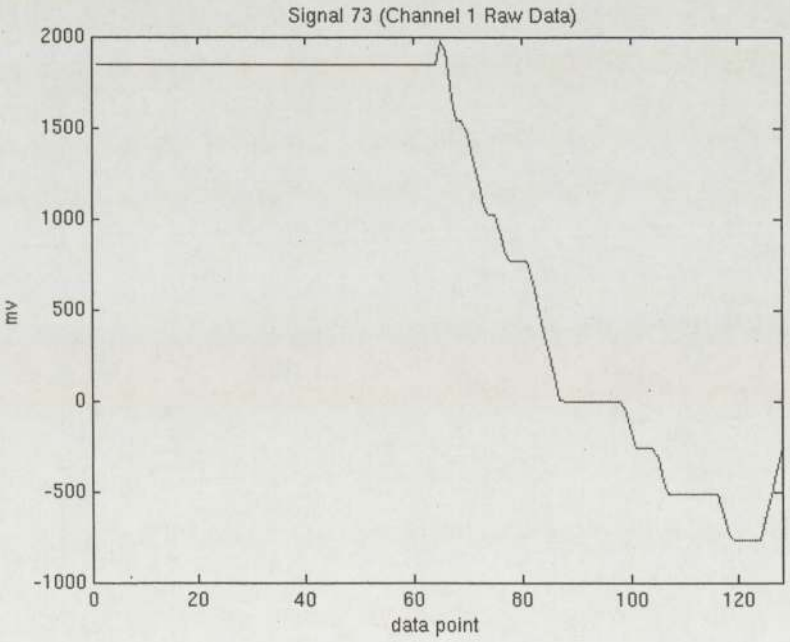

Figure 3.3: Only Part Of P–Wave Isolated

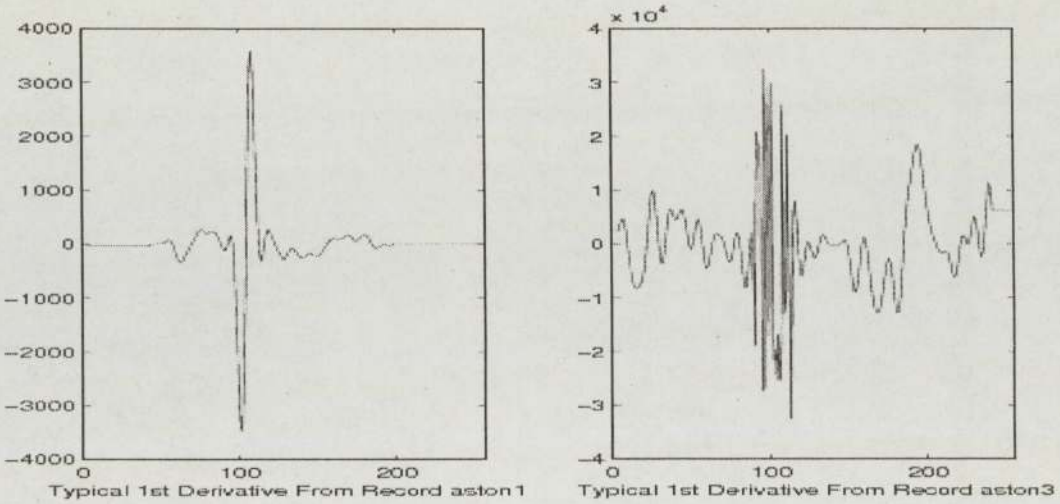

Figure 3.4: Comparison of First Derivatives

Chapter 4

Denoising ECG Signals

4.1 Introduction

Three wavelets, Haar, Daubechies and Coiflet, were used to perform denoising of the initial random selection of 100 beats described in section 3.3. The Haar wavelet was chosen because it is the simplest of all the wavelets, whilst the Coiflet and Daubechies wavelets had both been previously successfully used for ECG analysis by other researchers [Sternickel 2002; Cuesta-Frau *et al.* 2000; Sánchez *et al.* 2002; Senhadji *et al.* 1995; Tasoulis *et al.* 2004; Li *et al.* 1995]. Denoising was performed using the WaveLab software package. The aim of this denoising analysis was to provide a comparison between the three wavelets and to discover which wavelet performed most successfully.

4.2 Software

The WaveLab software package is a collection of MATLAB functions which implement a variety of algorithms related to wavelet analysis. It was developed by, and is freely available from, the Department of Statistics at Stanford University[1]. The most recently available version, WaveLab802, was used, and was installed according to the instructions provided.

WaveLab provides a Graphic User Interface (GUI) for signal analysis, invoked by the command `WLBrowser`, and it was this feature that was used for the initial analysis. There were some initial problems with the use of the GUI, with some programs needing editing and other programs being missing (see Appendix A).

[1] `URL:http:/www-stat.stanford.edu/~wavelab/`

4.3 Denoising Using The WaveLab GUI

The WaveLab GUI was used for denoising of the channel 1 raw data of the random selection of 100 beats (CinC_absdat) described in section 3.3. Three different wavelets were compared, the Haar, Daubechies, and Coiflet.

The denoising analysis was carried out using the following procedure. Initially the data file containing the channel 1 raw data from the 100 random beats (CinC_absdat) was loaded into MATLAB. Then each isolated beat was subjected to the following procedure in turn. The beat was saved as a variable x, which was then saved in a file named kitlab, which was the default WaveLab GUI filename for user defined data. This was to enable the signal to be loaded correctly. This file was then loaded into WaveLab by first selecting Data and then User Defined in the GUI. Next the beat was Normalised by selecting Xforms and Normalize, this being the required WaveLab pre–processing for naturally occurring data. The way in which this method achieves the Normalisation of the signal is by calculating the first level detail coefficients by the use of the Quadrature Mirror Filters (QMF) of the Coiflet 3 wavelet, then determining the absolute median value of these coefficients, and finally dividing the signal by this value. The wavelet to be used for the denoising was chosen by selecting Params, then Wavelet, and then the desired wavelet name (either Haar, Daubechies, or Coiflet). The denoising was then carried out using a thresholding method by selecting Actions and WT DeNoise. To obtain the denoised signal in the MATLAB workspace, the command `load recon` was entered. This loaded the denoised signal into the workspace as the variable xhat, which was then saved in a data file. The data files of the denoised signals for the the Haar, Coiflet,and Daubechies wavelets were named xhat_haar, xhat_coif, and xhat_daub, respectively.

The root mean squared error between the denoised and original signals for each one of the 100 beats using each of the three wavelets was calculated, together with the combined mean squared error for each wavelet.

4.3.1 Results

The combined mean squared error for the Coiflet and Daubechies wavelet were very similar, being 2.0245 and 2.0490 respectively. The error using the Haar wavelet was greater than the other two wavelets used for denoising, having a value of 2.3703. The root mean squared errors (RMSEs) of the three wavelets were compared using a paired t–test and significant differences were found between each pair of wavelets at the 5% significance level, with the difference between the Coiflet and Daubechies wavelets being the least, having a p–value of 0.0079. A comparison between the wavelets for the beats from each record was undertaken. Significant differences between the Coiflet and Haar wavelets,

and Daubechies and Haar wavelets, were found for all of the four records, whilst only record *aston1* showed a significant difference between the Coiflet and Daubechies wavelets, with the other three records showing no significant difference between these two wavelets. These differences can clearly be seen from the boxplots (figure 4.1). The root mean squared error (RMSE) for each individual beat was also generally greater for the Haar wavelet reconstruction than for that of the other two wavelets, as can been seen from figure 4.2.

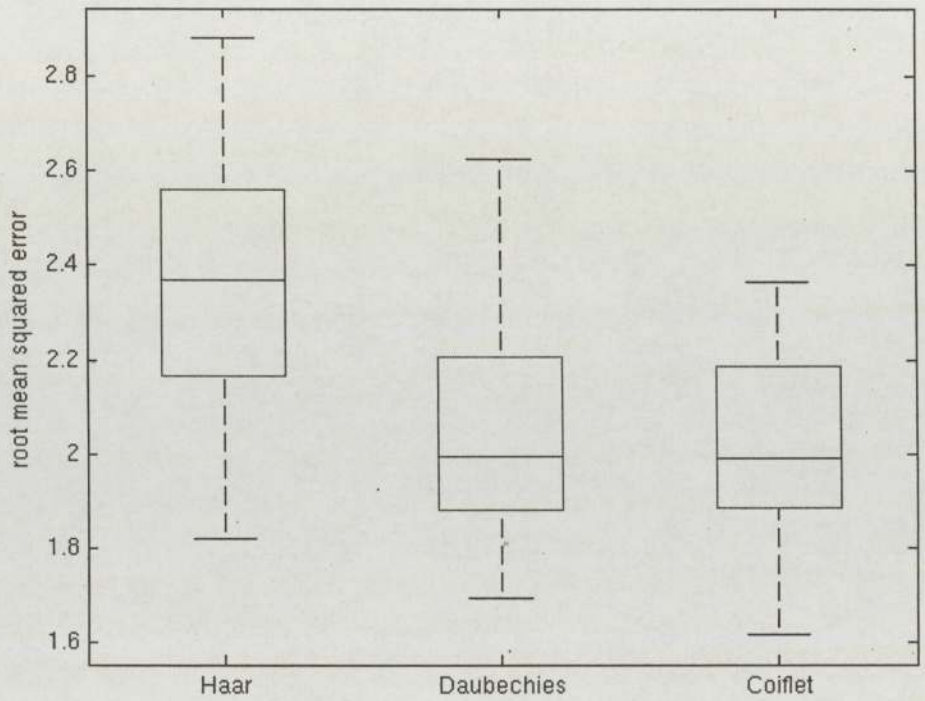

Figure 4.1: Wavelet Denoising Comparison (Boxplots)

A visual examination of the reconstructed, denoised, signals, with that of the original signals showed that the Haar wavelet had given a less smooth reconstruction than the other two wavelets (as illustrated in figure 4.3). This can be attributed to the fact that the Haar wavelet is the least smooth of the three wavelets compared.

4.4 Replication of Denoising Results

The reason for attempting to replicate the previous results without using the WaveLab GUI is that loading each individual beat separately and performing the wavelet analysis is very time consuming. If a large amount of data has to be processed, as is the case in ambulatory ECG recordings, then a

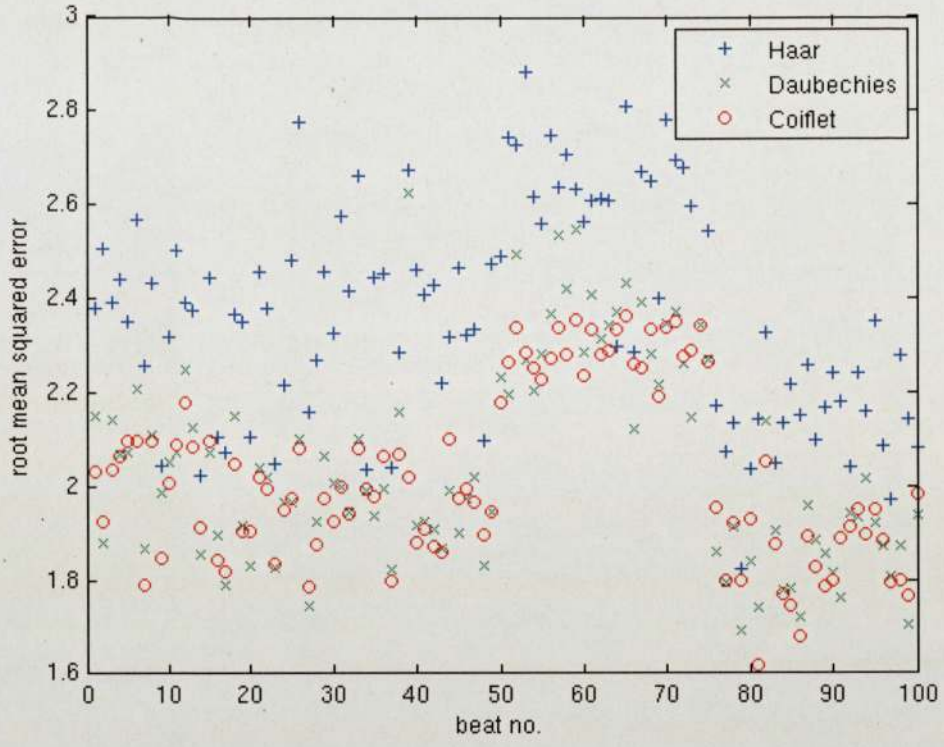

Figure 4.2: Wavelet Denoising Comparison (Individual Beats)

quick method of processing the data in batches is required.

The way in which the WaveLab GUI implemented the wavelet analysis was examined, with particular reference to the relationship between the various programs (figures 4.4 and 4.5).

4.4.1 Normalisation of Signals

A program was written to load CinC_absdat and to normalise all of the beats in the data file using the WaveLab default method for Normalising signals.

4.4.2 Denoising of Signals

The WaveLab default method for denoising signals is slightly different when the Daubechies wavelet is used, as compared with the Coiflet and Haar wavelets. This difference comes about because the default method for the Daubechies wavelet uses boundary filters developed by Cohen *et al.* [1993], with filter coefficients calculated by Cohen, Daubechies, and Vial [1993], for calculating the wavelet coefficients at either end of the signal. Boundary filters are used in an attempt to resolve the problems arising from the analysis of signals with finite lengths, the more usual approach being to wrap the

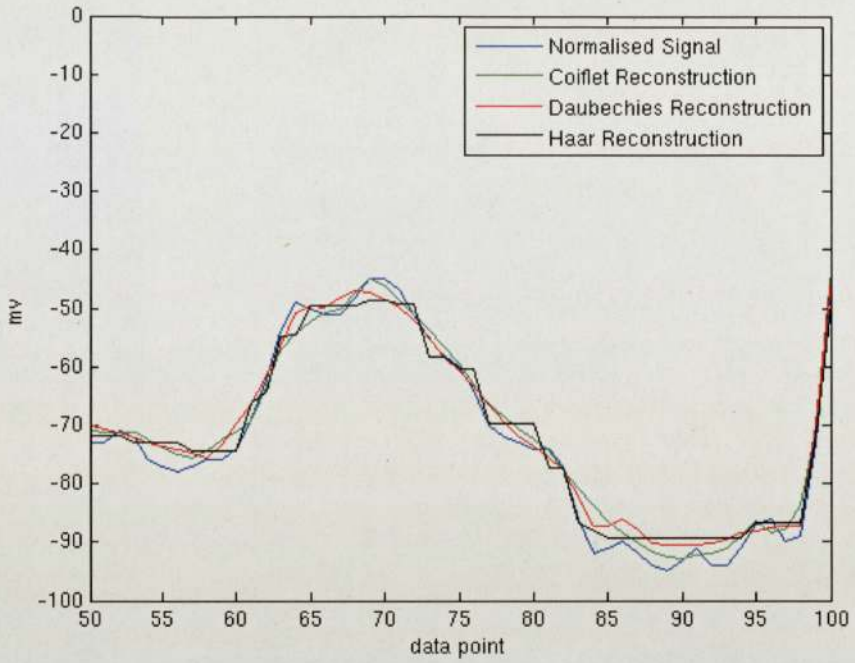

Figure 4.3: Signal Reconstruction Comparison: P–Wave Detail

signal back around itself.

Selecting Params and then Wavelet in the WaveLab GUI allows the user to choose one of five wavelets for the wavelet analysis, Haar, Daubechies, Average–Interpolating, Coiflet, or Symlet. This defines the QMF to be produced for the denoising when Actions, and then WT DeNoise are selected. When WT DeNoise is selected a program named **do_shrinkage.m** is called which produces the appropriate QMF, carries out the forward wavelet transform, performs the denoising using the VisuShrink algorithm [Donoho and Johnstone 1994] and finally carries out the reverse wavelet transform on the denoised coefficients.

Coiflet and Haar Wavelets

As has been previously mentioned, the QMFs are produced in a different manner for the Daubechies wavelet because of the use of boundary filters. The QMFs for the Coiflet and Haar wavelets are produced using a program called **MakeONFilter.m** and the wavelet coefficients are produced by **FWT_PO.m**. The inverse transform of the denoised coefficients is then performed using **IWT_PO.m**.

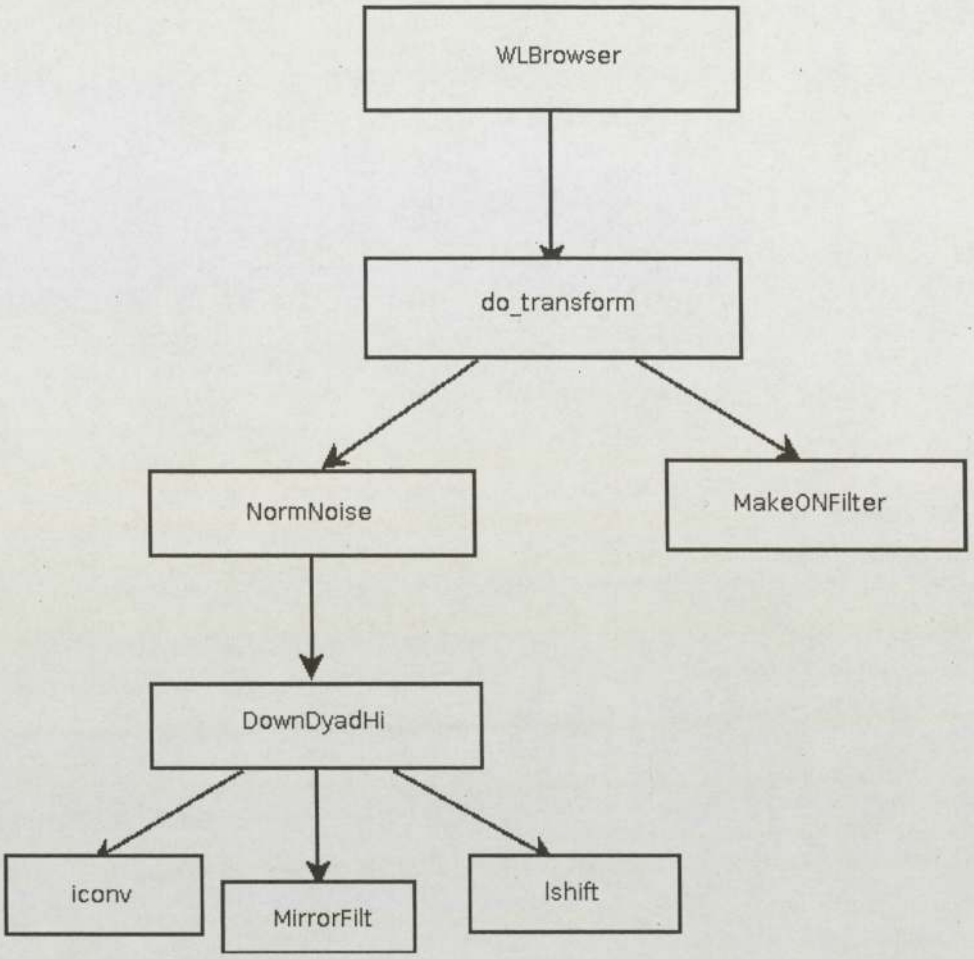

Figure 4.4: WaveLab Programs Used For Normalising Signals

Daubechies Wavelet

For the Daubechies wavelet both the QMF and the coefficients are produced by **FWT_CDJV.m**, with the inverse transform of the denoised coefficients being effected by **IWT_CDJV.m**.

4.4.3 Software

A program was written to enable the denoising of the ECG beats to be carried out quickly. This software loaded and normalised the data, and then carried out the denoising analysis for each of the three wavelets using the WaveLab software as described above.

4.4.4 Results

The results obtained from running this program were compared to those previously obtained by the use of the WaveLab GUI and found to be identical, indicating that the software which had been

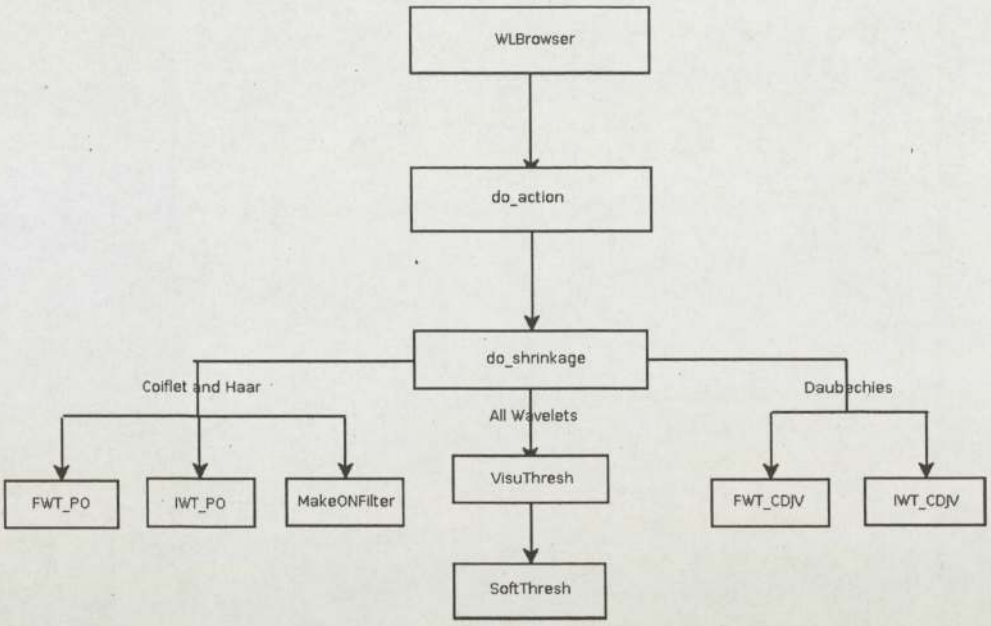

Figure 4.5: WaveLab Programs Used For Denoising Signals

written for the denoising process performed correctly.

4.4.5 Determination of Wavelet Coefficients

The rationale behind obtaining the results without using any WaveLab software is that the aim of the research is to incorporate the wavelet analysis detection algorithm into the Medilog Darwin software, and the WaveLab programs will not be available in Medilog Darwin.

As a first step to the implementation of the wavelet analysis without the use of WaveLab software, a program was written to produce the wavelet coefficients for any of the WaveLab supported wavelets which only involved the use of one WaveLab program (**MakeONFilter.m**). Although this program still made use of **MakeONFilter.m** in order to produce the wavelet coefficients, it avoids using the three WaveLab programs **do_action.m**, **do_shrinkage.m**, and **FWT_PO.m**.

The Coiflet and Haar wavelet coefficients produced using this program proved to be identical to those obtained using the WaveLab GUI, indicating that this software was suitable for the denoising of ECG signals using these wavelets. The Daubechies wavelet coefficients were produced without the use of the boundary filters devised by Cohen *et al.* [1993]. They were therefore not all identical to those produced by the WaveLab GUI. However, they were all identical to those obtained using the usual Daubechies QMF with no boundary filters, again indicating the suitability of the software for the denoising of ECG recordings.

This program was then used as a basis for producing programs to produce the Haar and Daubechies

wavelet coefficients without the use of WaveLab programs by directly creating the appropriate QMF.

4.5 Conclusions

The Coiflet and Daubechies wavelets are better suited to the denoising of ECG signals than the Haar wavelet. This would appear to be because they are both smoother wavelets than the Haar wavelet.

Software has been written to produce wavelet coefficients identical to those obtained from WaveLab, demonstrating that suitable programs can be written which can perform wavelet analysis. These programs should then be suitable for incorporation into the Medilog Darwin software.

Chapter 5

Algorithme À Trous

5.1 Introduction

The usual implementation of wavelet analysis involves downsampling of the output of both the highpass and lowpass filters at each level in order to reduce the computational load. This leads, however, to the loss of some temporal information concerning the signal, which becomes more marked the lower the frequency being examined. This is of concern because the object of the research is to detect P–waves, and it is the time of these P–waves which is of importance. The algorithme à trous is an algorithm which preserves all of the temporal information of the signal, but at the cost of increasing the computational load.

5.2 Background

The algorithme à trous, literally "algorithm with holes", was developed by Holschneider *et al.* [1988, 1989] for computing the wavelet transform of an arbitrary signal in real–time. It is time–invariant, and is therefore suitable for extracting temporal information from a signal. The reason for this time–invariance is that instead of downsampling the filter outputs at each level and using the same filter, as is usually the case with wavelet analysis, the algorithme à trous retains the whole of the output from both the highpass and lowpass filters, whilst doubling the filter length at each level by inserting a zero at the front of the filter and between each of the filter coefficients. This manner in which this done is illustrated in figure 5.1 for an original filter with four coefficients (A1, A2, A3, and A4).

5.3 Implementation And Evaluation

A program (**FWT_ATrou.m**) was available in the WaveLab software package for the implementation of the algorithme à trous using a quadratic spline wavelet. This program was used as the basis for

Original Filter							A1	A2	A3	A4						
Level 1 Filter					0	A1	0	A2	0	A3	0	A4				
Level 2 Filter	0	0	0	A1	0	0	0	A2	0	0	0	A3	0	0	0	A4

Figure 5.1: Algorithme À Trous Filter

writing the software for this research. Programs were written to implement the algorithme à trous for both the Haar wavelet and the quadratic spline wavelet used by Martínez *et al.* [2004]. These programs were then tested by evaluating their performance on an artificially produced P–wave representative of the P–waves in the channel 1 raw data from record *aston1*. This synthetic P–wave (figure 5.2) was produced by the use of quadratic splines.

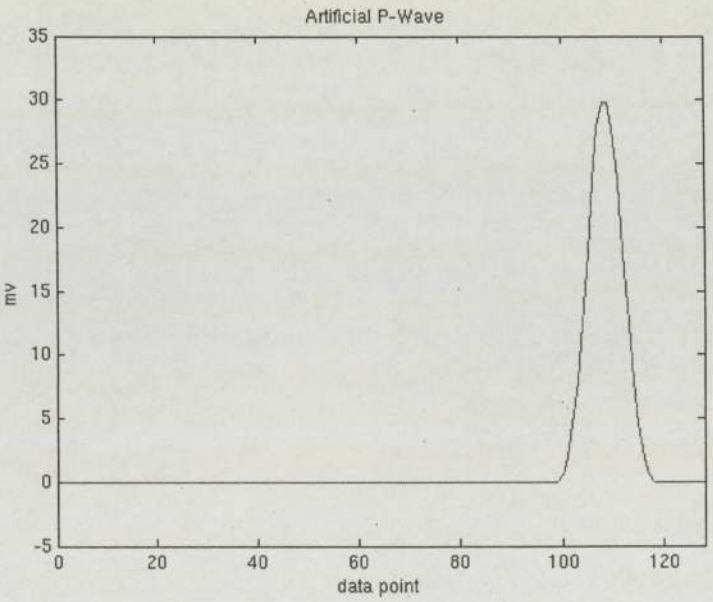

Figure 5.2: Synthetic P–Wave

When the two programs were tested on the artificial P–wave it was found that the point where the wavelet coefficients crossed zero did not exactly match up with the peak of the P–wave, but that it occurred slightly before the peak. The amount by which the time of the zero crossing differed from the peak of the P–wave increased with the level of the wavelet decomposition, with the zero crossing of the wavelet coefficients corresponding to the highest frequency being closest to the P–wave peak. Investigation of the software revealed that the signal was left shifted at each decomposition level, which caused the increasing divergence of the zero crossing from the P–wave peak. This was corrected by right shifting the coefficients by the appropriate amount, $2^{(d-1)}$, (where d is the level of

decomposition) after the analysis had been completed.

After amending the programs, they were again tested on the synthetic P–wave, and this time the zero crossing point coincided with the peak of the P–wave, at data point 108, for levels three, four, and five (figures 5.3 and 5.4).

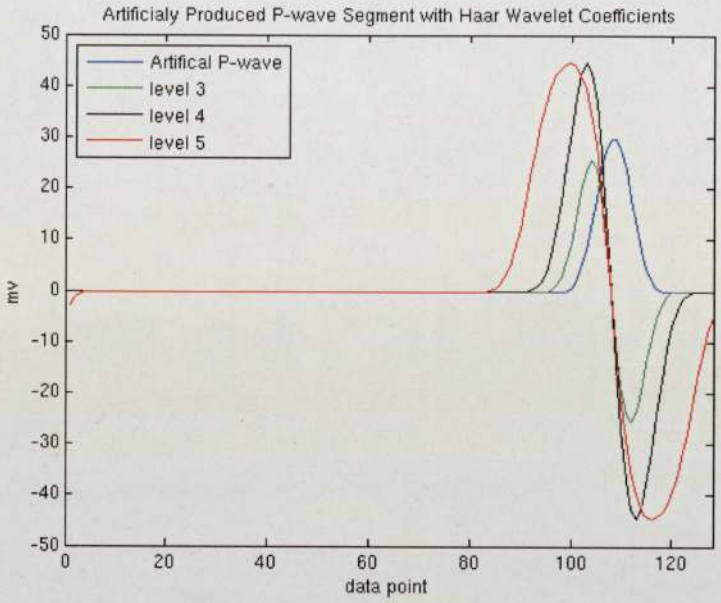

Figure 5.3: Third, Fourth and Fifth Level Haar Detail Coefficients of the Artificial P–Wave

5.3.1 Noisy Signals

The programs were also tested on the synthetic P–wave with artificial Gaussian noise added, initially with a mean of zero and a standard deviation of five (chosen because this gave a very noisy signal without completely obscuring the P–wave), to determine what effect this would have on the correspondence between the zero crossing and the peak of the P–wave. The level three, four, and five detail coefficients were again used for the comparison as these were the most informative regarding information concerning the P–wave (figures 5.5 and 5.6). The maximum value of the noisy synthetic P–wave occurred at data point 107. The zero crossings for the quadratic spline wavelet for levels three, four, and five were all closest to data point 108, the same data point as that of the underlying artificial P–wave. This was also true for the Haar wavelet, except for the zero crossing of the level three coefficients, which was closest to data point 107.

To determine what effect excessive noise would have on the wavelet analysis, Gaussian noise with a mean of zero and a standard deviation of fifty was added to the synthetic P–wave. This had the effect of totally obscuring the P–wave in the signal (figure 5.7).

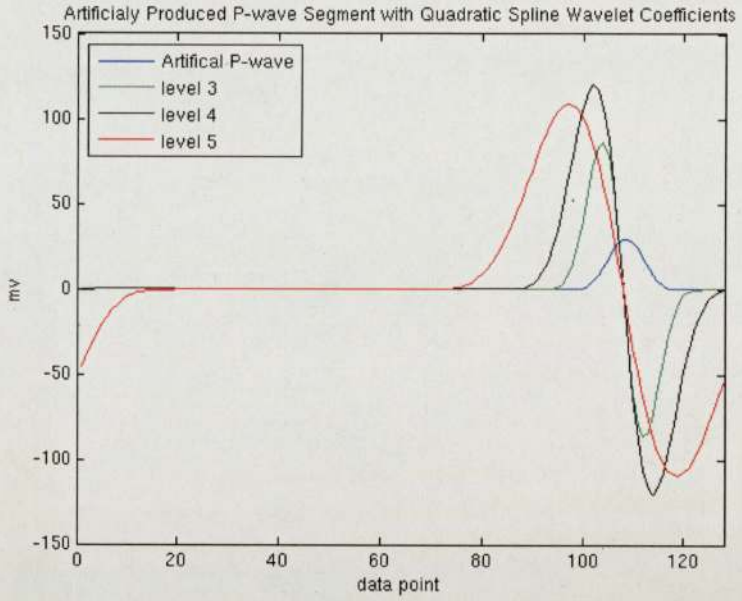

Figure 5.4: Third, Fourth and Fifth Level Quadratic Spline Detail Coefficients of the Artificial P–Wave

The results obtained using the Haar wavelet were very poor for this signal (figure 5.8), with no clear indication of the position of the P–wave peak at any of the three levels.

Results were better for the quadratic spline wavelet (figure 5.9), with the level four and five zero crossings indicating a maxima at data point 112, only four data points away from the underlying P–wave maxima at data point 108. The level three zero crossing, however was as poor as the Haar wavelet zero crossings at indicating the P–wave peak.

5.4 Conclusions

The algorithme à trous is a suitable algorithm for the detection of P–waves, as shown by its ability to accurately detect the peak of an artificially produced P–wave. This ability is not affected by the addition of a small amount of noise, but if the noise is so great as to make visual identification of the P–wave a problem then the Haar wavelet, and level three quadratic spline wavelet, cannot determine the location of the P–wave maxima. The results are better for the level four and five quadratic spline wavelets, differing by only 0.016 seconds from the actual time of the P–wave peak.

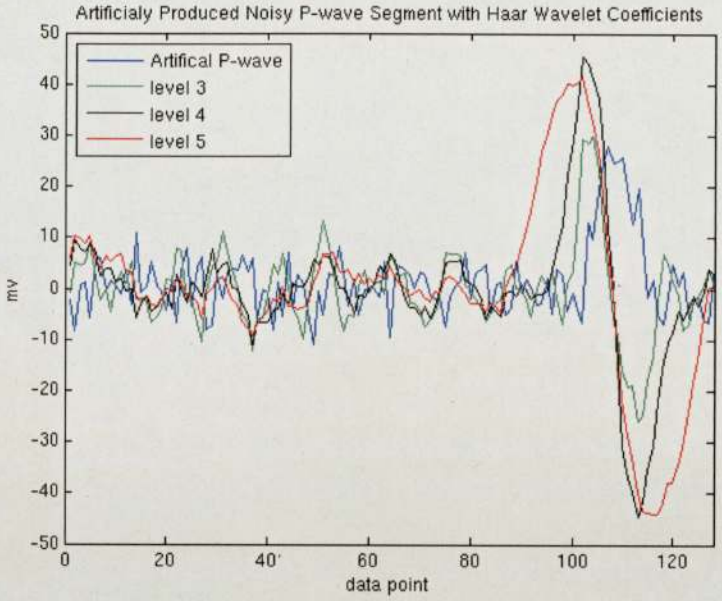

Figure 5.5: Third, Fourth and Fifth Level Haar Detail Coefficients of the Artificial P–Wave with Added Noise (sd 5)

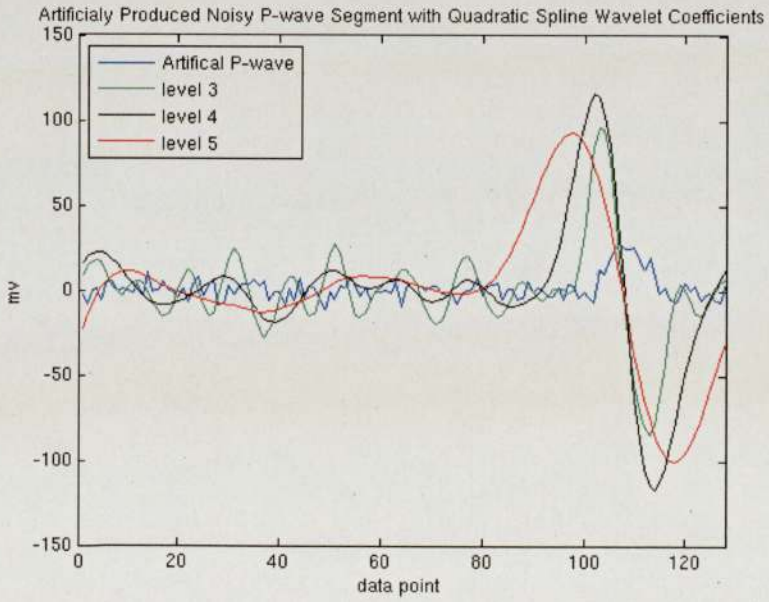

Figure 5.6: Third, Fourth and Fifth Level Quadratic Spline Detail Coefficients of the Artificial P–Wave with Added Noise (sd 5)

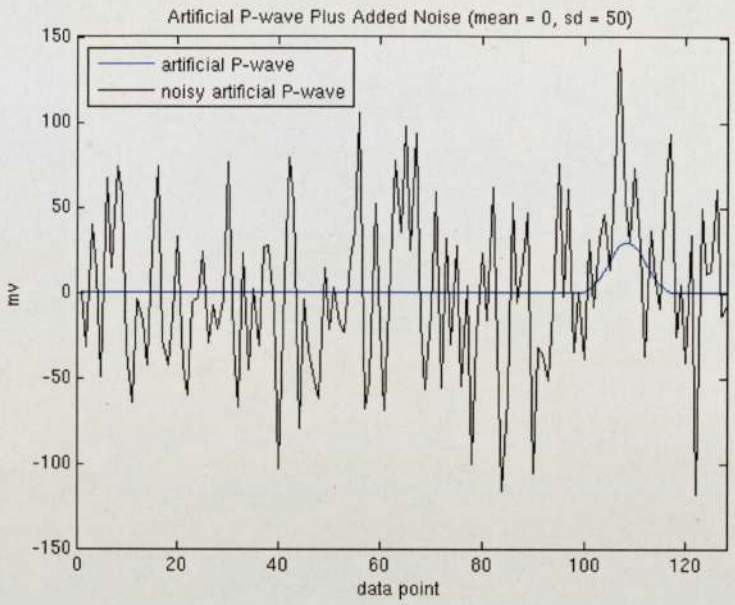

Figure 5.7: P–Wave and P–Wave with Added Noise

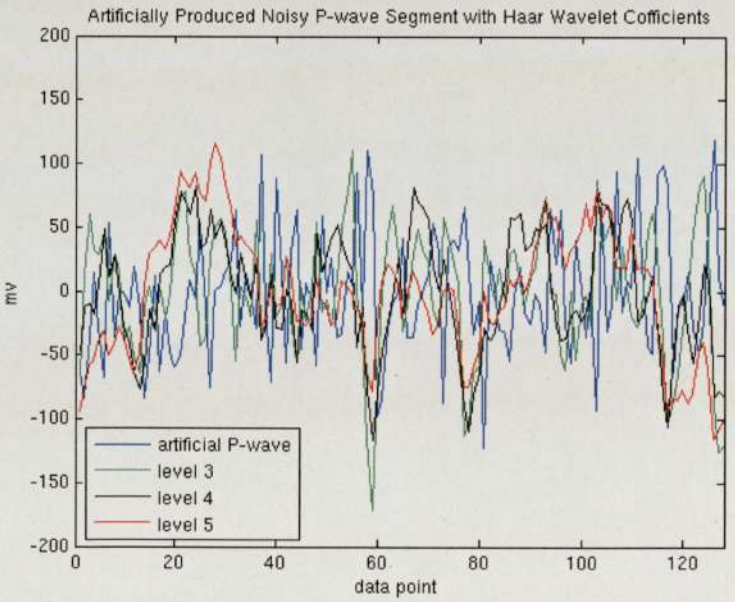

Figure 5.8: Third, Fourth and Fifth Level Haar Detail Coefficients of the Artificial P–Wave with Added Noise (sd 50)

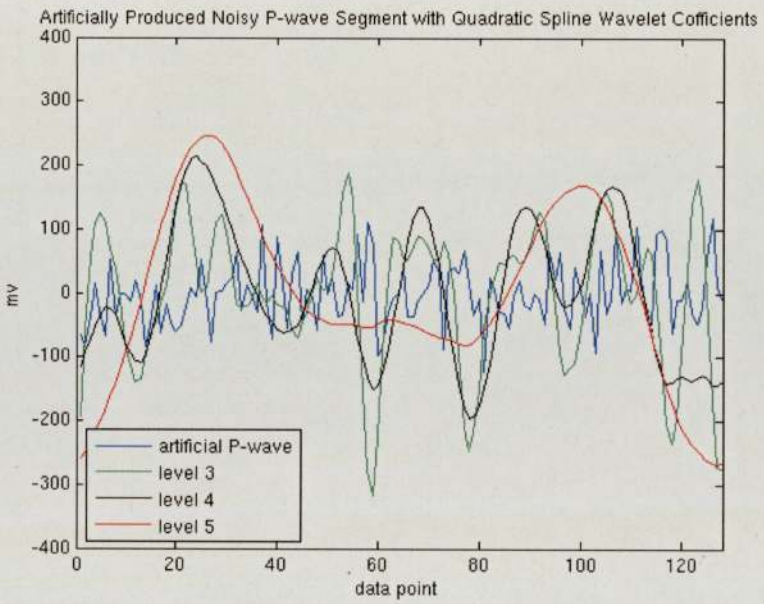

Figure 5.9: Third, Fourth and Fifth Level Quadratic Spline Detail Coefficients of the Artificial P–Wave with Added Noise (sd 50)

Chapter 6

P–Wave Detection

6.1 Zero Crossing Detection

Software was developed to detect the zero crossings of the wavelet coefficients and the results obtained from the implementation of this software are presented.

6.1.1 Program Development

As previously described in section 1.6 it is the detail wavelet coefficients which contain the information regarding the changes in the signal. The point at which the detail wavelet coefficients cross the y axis at any given scale indicates a maximum value of the signal. The scale imparts information concerning the gradient of the signal, with the smallest scales containing the most information concerning the steepest gradients. Therefore a program needed to be written in MATLAB to detect the point where the value of the wavelet coefficients crossed zero. This would then allow P–wave maxima to be detected when the algorithm was applied to the appropriate wavelet scale. The program development was carried out on the P–wave sections isolated from the first set of 100 random beats.

The first attempt at writing a program found the minimum value of the wavelet coefficients in the search window, together with its index. Then the coefficients were searched backwards from this point to discover the maximum value of the wavelet coefficients, together with its index. The minimum absolute value of the wavelet coefficients between these two points, together with its index, was then determined. This information was then used to produce a cross marking the position of the zero crossing of the wavelet coefficients, and hence the presumed maxima of the P–wave. This algorithm was applied to the detail wavelet coefficients at levels three, four and five. The purpose of this was to confirm that the level four detail wavelet coefficients were the most accurate indicator of the P–wave maxima, as had been found in previous research [Anant *et al.* 1994; Li, Zheng, and Tai 1995; Bahoura, Hassani, and Hubin 1997; Martínez, Olmos, and Laguna 2000; Sovilj, Jeras, and Magjarevic

2004; Martínez *et al.* 2004]. Research conducted by Sahambi, Tandon, and Bhatt [1997], however, used the level three detail wavelet coefficients because they claimed that interference from baseline drift became apparent in the level four detail coefficients, something that other researchers claimed only became apparent in the level five detail coefficients. Their findings are perhaps a little surprising as they used the MIT/BIH database, as used by Li, Zheng, and Tai [1995]; Bahoura, Hassani, and Hubin [1997]; and Martínez *et al.* [2004], and the CSE database, also used by Martínez *et al.* [2004] for their research.

The theory behind this algorithm is that the P–wave should be the largest peak within the search window, and therefore that the maximum and minimum wavelet coefficients obtained are associated with the P–wave. The minimum absolute value of the wavelet coefficients is assumed to be the point closest to the zero crossing, and therefore to be the best indication of the position of the P–wave maxima.

The P–wave markers obtained from the the third, fourth and fifth level wavelet coefficients were visually compared with the isolated P–wave segments. From this comparison it was decided that the fourth level detail coefficients did indeed give the best indication of the P–wave maxima. Therefore the remainder of the research concentrated on the information obtained from the level four detail coefficients.

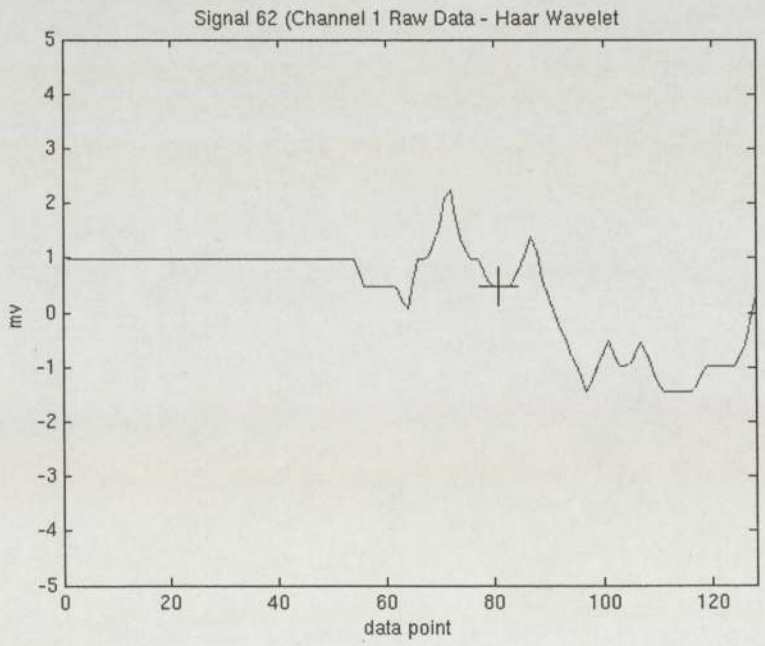

Figure 6.1: Incorrectly Detected P–Wave

When this program was assessed on the initial normalised isolated P–wave data (all three channels, both original and bandpass filtered) it was discovered that in some instances the marked zero crossings

were not close to the P–wave maxima as determined visually (figure 6.1). In other cases, where the P–wave had not been included in the search window (figure 6.2), the marker for the zero crossing was meaningless because it was positioned on the front end padded section. To correct these problems the program was adapted by searching backwards from the minimum wavelet coefficient value to find the point where the sign of the wavelet coefficients changed from negative to positive (or zero).

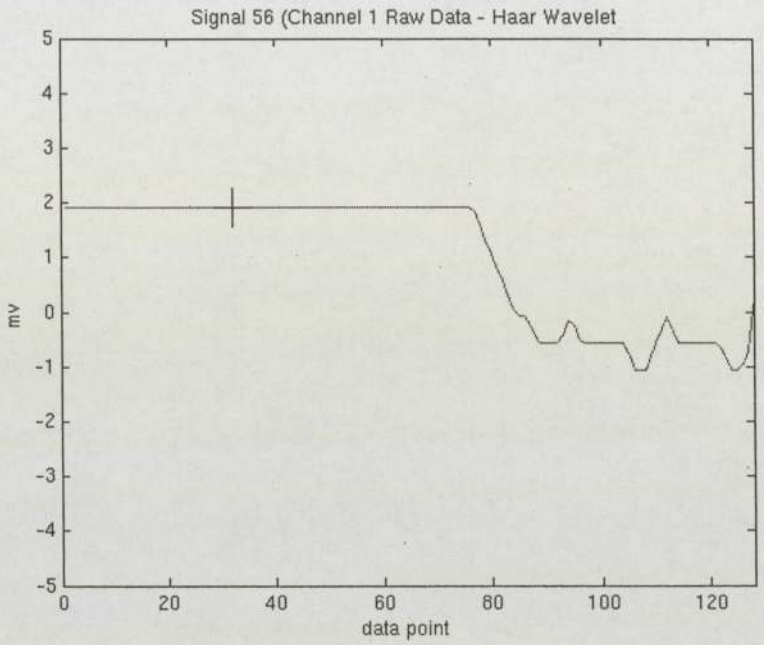

Figure 6.2: P–Wave Absent From Search Window

This algorithm still relies on the P–wave being the largest feature in the search window, and with the minimum wavelet coefficient being associated with it. It does, however, address the problems of the minimum absolute value of the wavelet coefficients between the minimum and maximum values not necessarily being associated with the P–wave maxima, and also of the index of the last value tested, if the values are identical, being the output of the program (the problem illustrated in figure 6.2).

When the efficiency of the new program was tested on the same data it was found to improve the efficiency of detecting the P–wave maxima. The program was further refined by comparing the absolute value of the two wavelet coefficients on either side of the sign change and taking the wavelet coefficient with the smallest absolute value, and its index, as the zero crossing.

6.1.2 Results

The effectiveness of this program was assessed on the un–normalised isolated P–wave sections of the development set of 100 random beats. An examination of the P–wave markers produced for this data set revealed that, out of the 100 beats, there were 18 in which the position of the P–wave was incorrectly marked in at least one of the twelve signal/wavelet combinations. Fifteen of these beats were from the 25 which had been randomly selected from record *aston3*. This record had already been noted as being different from the other records provided, and the production of the isolated P–wave sections for this record had also proved to be problematical (section 3.4). Of the other three beats which had at least one incorrectly marked P–wave, two, beat number 34 (record *aston2*) and beat number 92 (record *aston5*), were both noisy. The other beat to have at least one P–wave incorrectly detected was beat number 78 (record *aston5*).

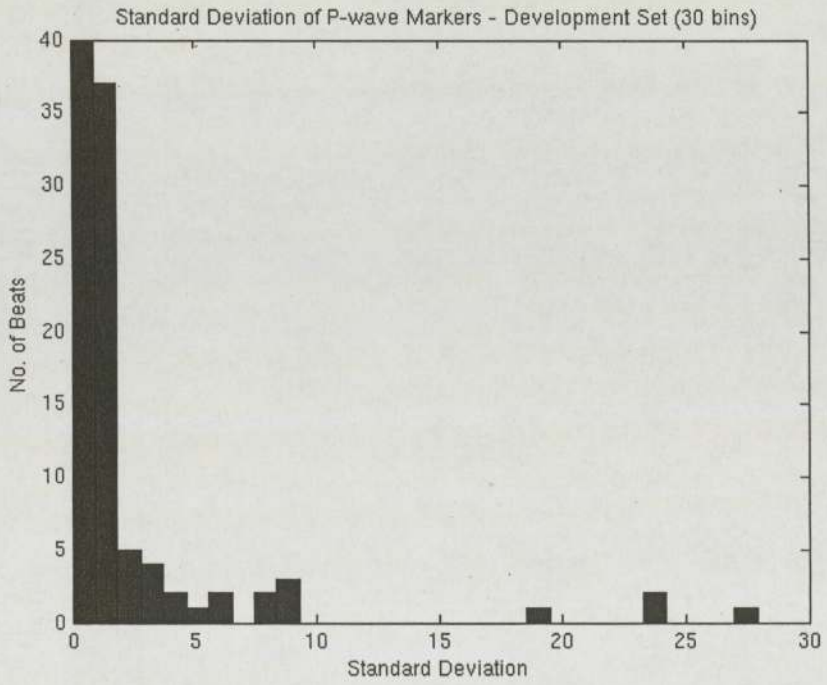

Figure 6.3: Histogram of the Standard Deviations of the P-wave Markers for the Development Set (30 bins)

For one third of the 18 beats which had incorrectly detected P–waves the zero crossing algorithm had failed to correctly detect the P–wave in any of the twelve signal/wavelet combinations. All of these beats were from record *aston3* (beat numbers 56, 58, 64, 67, 71, and 73). Of these 6 beats there had been problems with the isolation of the P–wave in five cases, beat number 56, where no P–wave was present in the isolated section, and beat numbers 58, 67, 71, and 73, which contained only part of the P–wave. The remaining beat for which no correct P–wave detections were observed (number

	no. of beats with sd > 5	no. of beats with sd < 5
Incorrectly marked P–wave	12	6
Correctly marked P–wave	0	82

Table 6.1: P–wave Markers and Standard Deviation Greater Than 5 (Development Set)

	no. of beats with sd > 2	no. of beats with sd < 2
Incorrectly marked P–wave	16	2
Correctly marked P–wave	6	76

Table 6.2: P–wave Markers and Standard Deviation Greater Than 2 (Development Set)

64) had the P–wave situated on the T–wave of the preceding beat.

The percentage of beats which had the P–wave correctly detected in all signal/wavelet combinations was 82%. If record *aston3* was discounted, due to it being atypical of the records supplied and the problems encountered with correct P–wave isolation, then this rose to 96%.

The standard deviations for the indices of the P–wave markers for each beat were calculated and a histogram plotted (figure 6.3). From this it can be seen that whilst the majority of the beats had P–wave markers that were in close agreement there were others where the standard deviation was high.

A comparison of beats where at least one of the P–wave markers was incorrect, was made with those which had a standard deviation of greater than five (table 6.1). Of the 18 beats having at least one incorrect P–wave marker, twelve also had P–wave marker standard deviations greater than five. All twelve of the beats having P–wave marker standard deviations greater than five were visually considered to have a P–wave marker incorrectly positioned in at least one of the signal/wavelet combinations.

When a similar comparison was made, but with beats having a standard deviation greater than two (table 6.2), sixteen of the beats were contained in both groups. Two beats (numbers 59 and 72) had at least one incorrect P–wave marker but standard deviations less than two, whilst six beats (numbers 2, 8, 12, 21, 44, and 65) had standard deviations greater than two although all of the P–wave markers appeared to be correct.

6.2 Program Testing

The revised zero crossing detection program was implemented on the test set of 100 isolated random P–wave sections to evaluate its performance. The initial evaluation being carried out visually.

6.2.1 Results

Out of the 100 beats in the test set there were 12 which had the P–wave incorrectly detected in at least one of the signal/wavelet combinations. All of these beats were from those randomly selected from record *aston3*.

Two of the beats (numbers 54 and 66) had no correctly detected P–waves in any of the twelve signal/wavelet combinations. The isolated section for beat number 66 only contained part of the P–wave. The algorithm detected the P–wave maxima in front of the actual maxima for all signal/wavelet combinations of beat number 54.

Of all of the beats in the test set, 88% had the P–waves correctly detected in all signal/wavelet combinations. If the beats from record *aston3* were discounted, due to the problems mentioned in section 6.1.2, then all of the P–waves were correctly identified, in all beats, in all signal/wavelet combinations.

The standard deviations for the indices of the P–wave markers for each beat were calculated and a histogram plotted (figure 6.4).

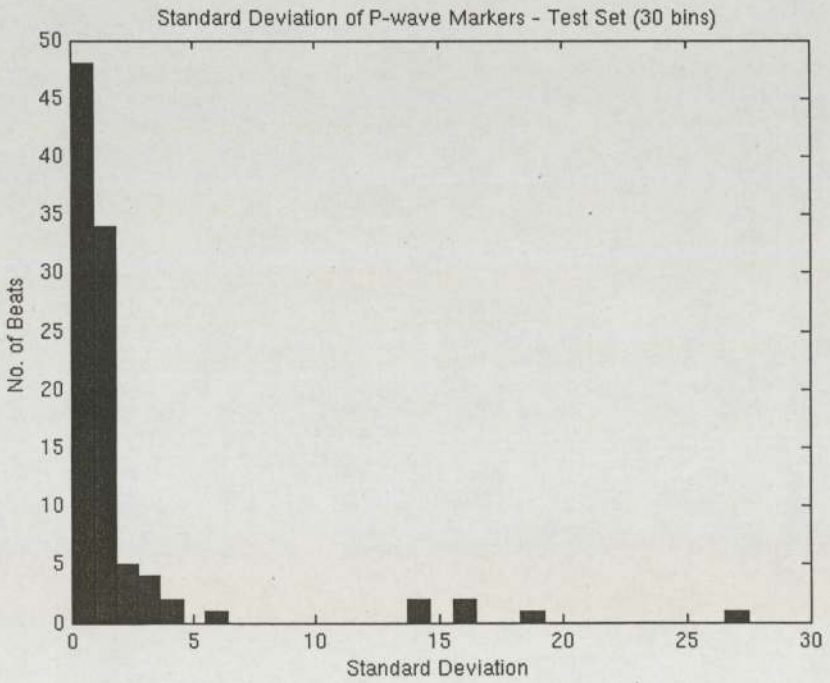

Figure 6.4: Histogram of the Standard Deviations of the P-wave Markers for the Test Set (30 bins)

A comparison of beats where at least one of the P–wave markers was incorrect, was made with those which had a standard deviation of greater than five (table 6.3). Of the 12 beats having at least one incorrect P–wave marker, seven also had P–wave marker standard deviations greater than five. All seven of the beats having P–wave marker standard deviations greater than five were visually considered

	no. of beats with sd > 5	no. of beats with sd < 5
Incorrectly marked P–wave	7	5
Correctly marked P–wave	0	88

Table 6.3: P–wave Markers and Standard Deviation Greater Than 5 (Test Set)

	no. of beats with sd > 2	no. of beats with sd < 2
Incorrectly marked P–wave	11	1
Correctly marked P–wave	4	84

Table 6.4: P–wave Markers and Standard Deviation Greater Than 2 (Test Set)

to have a P–wave marker incorrectly positioned in at least one of the signal/wavelet combinations.

When a similar comparison was made, but with beats having a standard deviation greater than two (table 6.4), eleven of the beats were contained in both groups. One beat (number 54) had at least one incorrect P–wave marker but a standard deviation less than two, whilst four beats (numbers 23, 35, 59, and 63) had standard deviations greater than two although all of the P–wave markers appeared to be correct.

6.3 Continuous Recording

As in section 6.1 the development of the software used is described, followed by a presentation of the results obtained from its use.

6.3.1 Program Development

Because the purpose of this research was to detect P–waves in ambulatory ECG recordings, the algorithms which had been developed had to be capable of operating on continuous recordings. Therefore one of the records containing predominantly normal beats (*aston2*) was chosen at random and the first half hour of the recording used for evaluation purposes.

On examination, this portion of the record contained 1,895 beats, of which all but the first were considered to be normal (according to the QRS morphology) by the Medilog Darwin software. On examination, the first annotated beat turned out to be not a genuine beat at all, but merely an artifact caused by the recorder.

Programs were adapted to analyse a continuous signal rather than the 100 isolated P–wave sections. Although now applied to a continuous signal, the algorithm still used a search window between the R–wave and the T–wave of the preceding beat for determining the position of the P–wave, as previously

described in section 3.3.1. The program for determining the zero crossings was adapted to calculate the zero crossing for each beat in the record. The signals for both the raw and bandpass filtered data, for each of the three channels, together with the P–wave marker for either the Haar or quadratic spline wavelet, were plotted as subplots in a figure. This gave a figure (6.5) containing twelve subplots, each with a window of 2000 data points, which could be scrolled to enable the viewing of any part of the record.

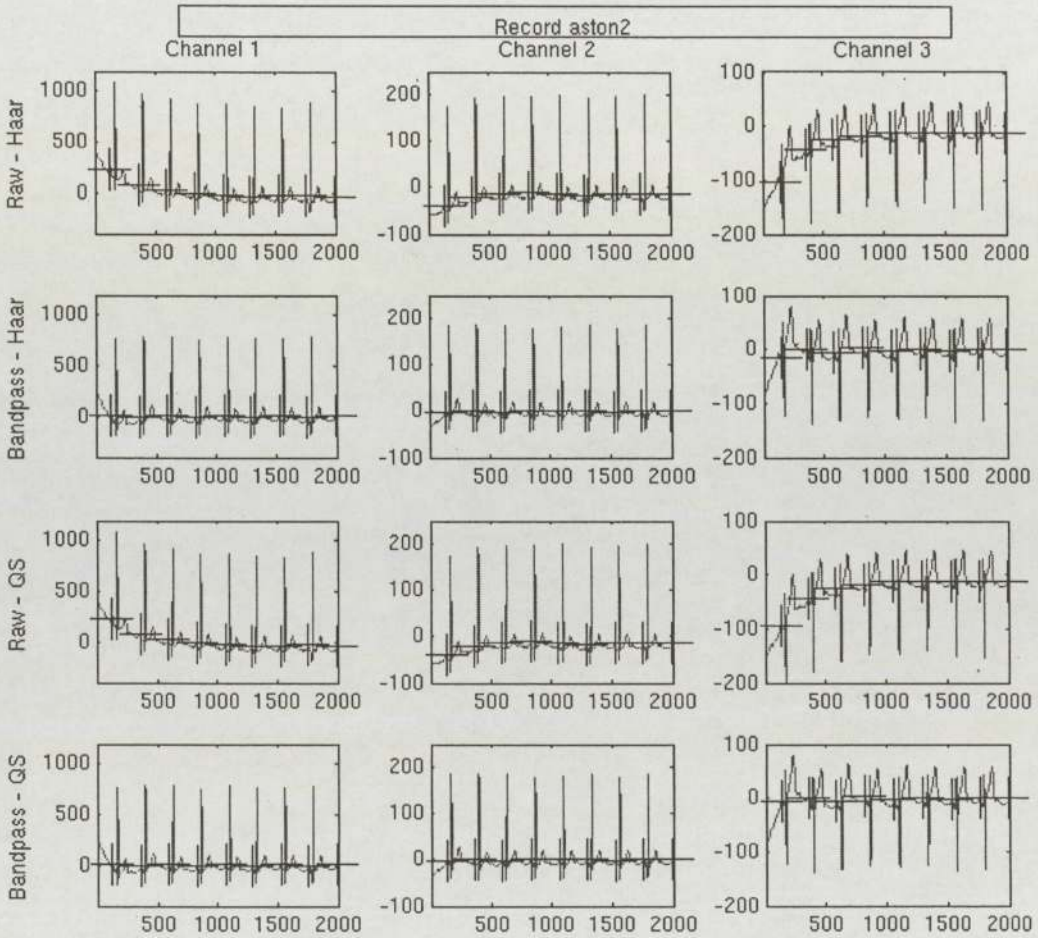

Figure 6.5: First 2000 Data Points Of Record *aston2*

6.3.2 Results

All of the P–wave markers for all twelve of the combinations of data channel, wavelet, and raw or bandpass data, were checked visually to determine whether they appeared to give an accurate indication of the position of the P–wave. In the instances where a P–wave marker did not appear

	no. of beats with sd > 5	no. of beats with sd < 5
Incorrectly marked P–wave	24	3
Correctly marked P–wave	1	1866

Table 6.5: P–wave Markers and Standard Deviation Greater Than 5 (Continuous Recording)

to be correct, the beat concerned was re–examined for all the data channels and both of the wavelet transforms. Out of the 1,894 beats in the record there were only 27 (1.4% of the total number of beats) in which the P–wave maxima appeared to be incorrectly detected in at least one of the twelve signal/wavelet combinations.

The standard deviation of the indices of the P–wave markers for each beat was calculated and histograms plotted. From these histograms it appeared that the overwhelming majority of the standard deviations were less than five (figures 6.6 and 6.7). Therefore it was decided to examine those beats which had standard deviations greater than five in order to determine how they corresponded with those P–wave markers which had been visually identified as incorrectly marking a P–wave maxima (table 6.5).

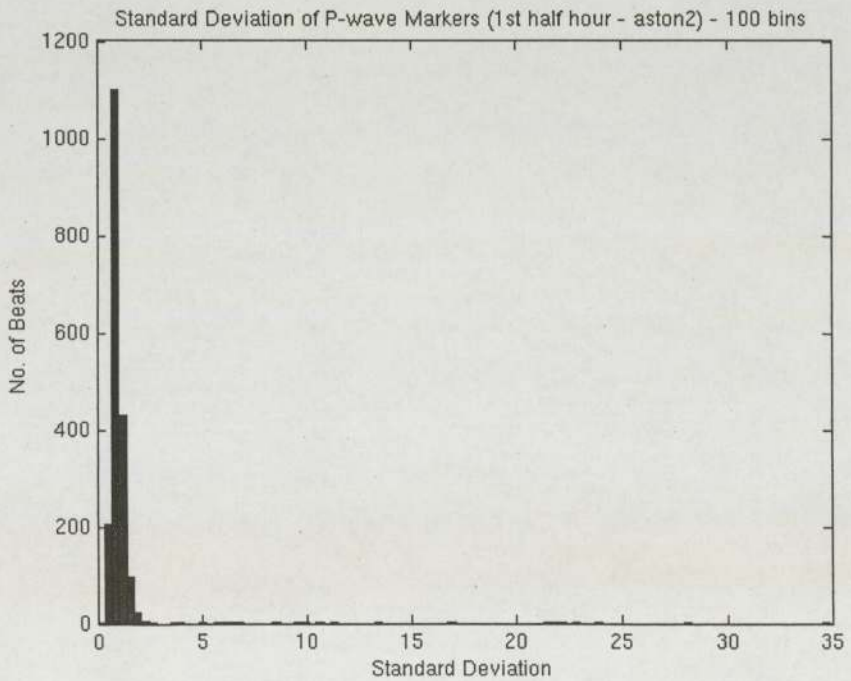

Figure 6.6: Histogram of the Standard Deviations of the P-wave Markers for the First Half Hour of Record *aston2* (100 bins)

Twenty–five of the beats had P–wave marker standard deviations greater than 5, and of those, twenty–four corresponded to beats which had been visually identified as having the P–wave maxima

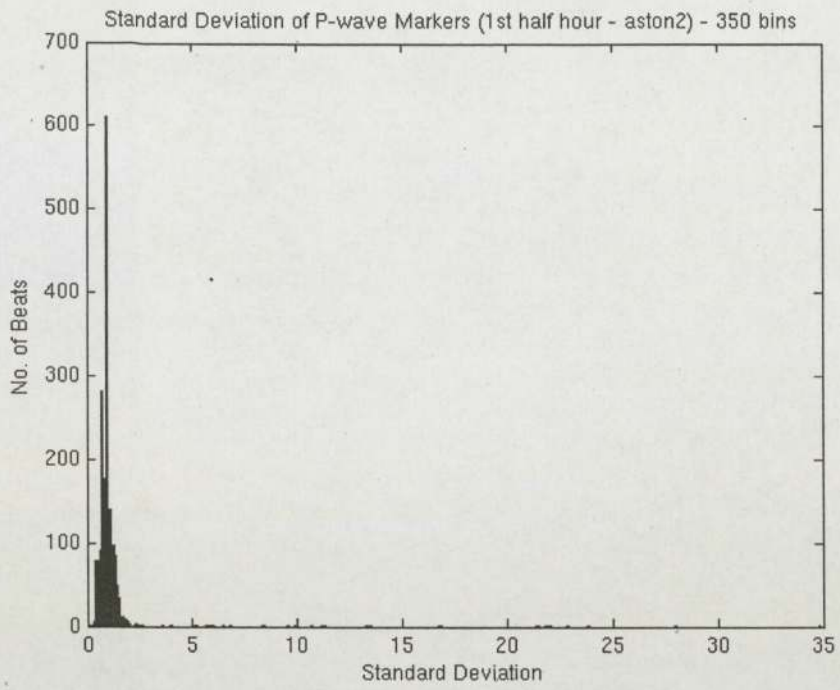

Figure 6.7: Histogram of the Standard Deviations of the P-wave Markers for the First Half Hour of Record *aston2* (350 bins)

incorrectly marked in at least one of the data channel/wavelet combinations. The one beat which had not been so identified was beat number 858 (figure 6.8), which only had a standard deviation slightly above this threshold (5.2049). As can be seen from figure 6.8, the channel 1 and 2 signals are very noisy and it is difficult to discern the P–wave in these channels. However the channel 3 signal is practically noise free, and the P–wave can easily be located in this channel.

Three beats, therefore, had standard deviations of five or less whilst having been visually identified as having at least one incorrect P–wave marker. These were beat numbers 1239, 1498, and 1682. Of these three, beat number 1498 had its P–wave obscured under the T–wave of the preceding beat (figure 6.9), whilst the signals for the other two beats were very noisy in the region where the P–wave was to be expected (figures 6.10 and 6.11).

From figure 6.9 it can be seen that the P–wave markers produced by using the Haar wavelet are situated around the peak of the T–wave from the preceding beat, whilst those produced by the use of the quadratic spline wavelet locate the peak of the P–wave before the apex of the T–wave from the preceding beat. This suggests that the Haar wavelet transform is unable to distinguish between P and T–waves, but that the quadratic spline wavelet has greater power to discriminate between the two different waves.

For beat numbers 1239 and 1682 the signals are very noisy for all three channels, especially so in

the case of beat number 1239, and a visual identification of the P–wave in these cases proved to be impossible.

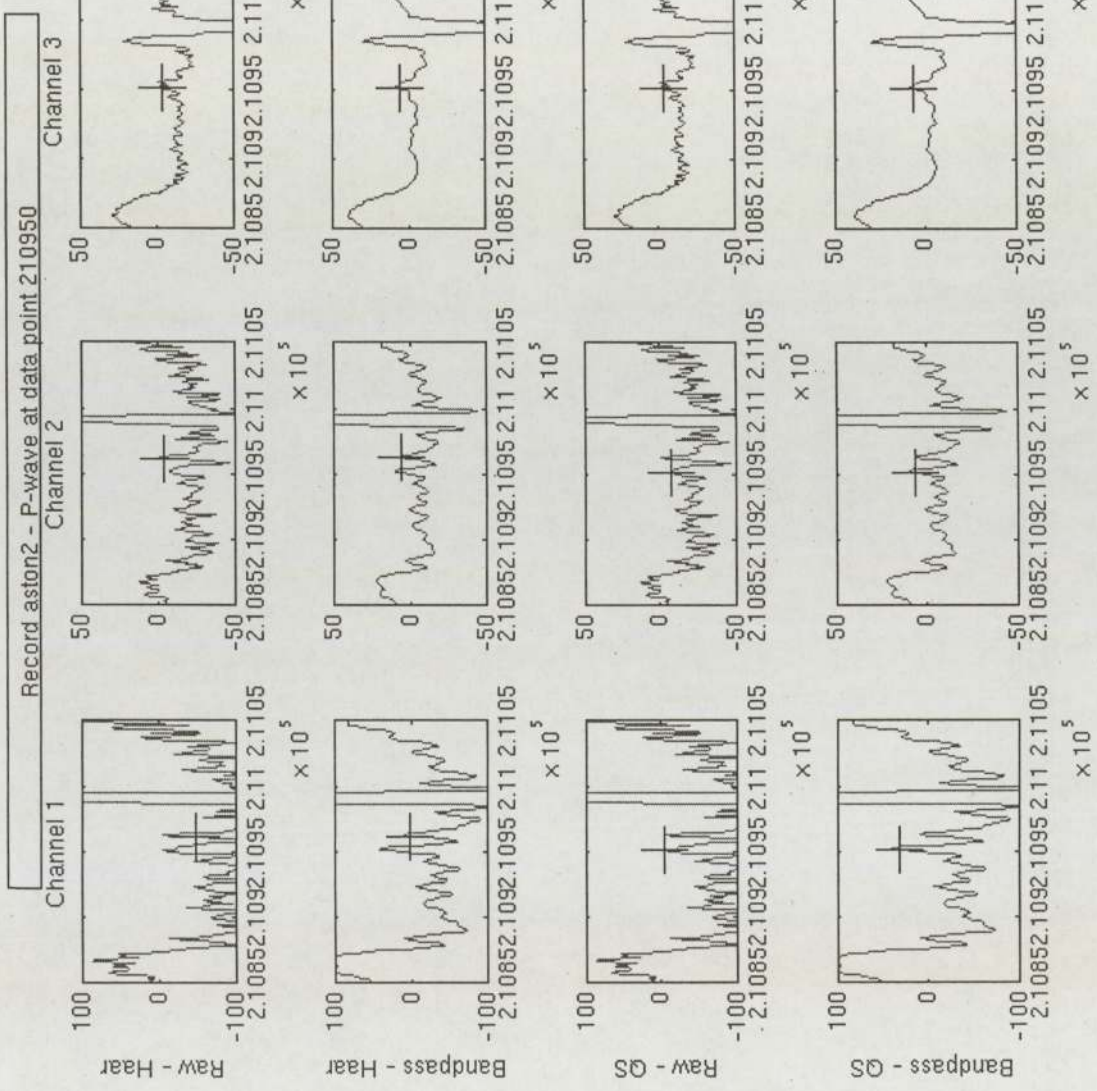

Figure 6.8: Beat Number 858

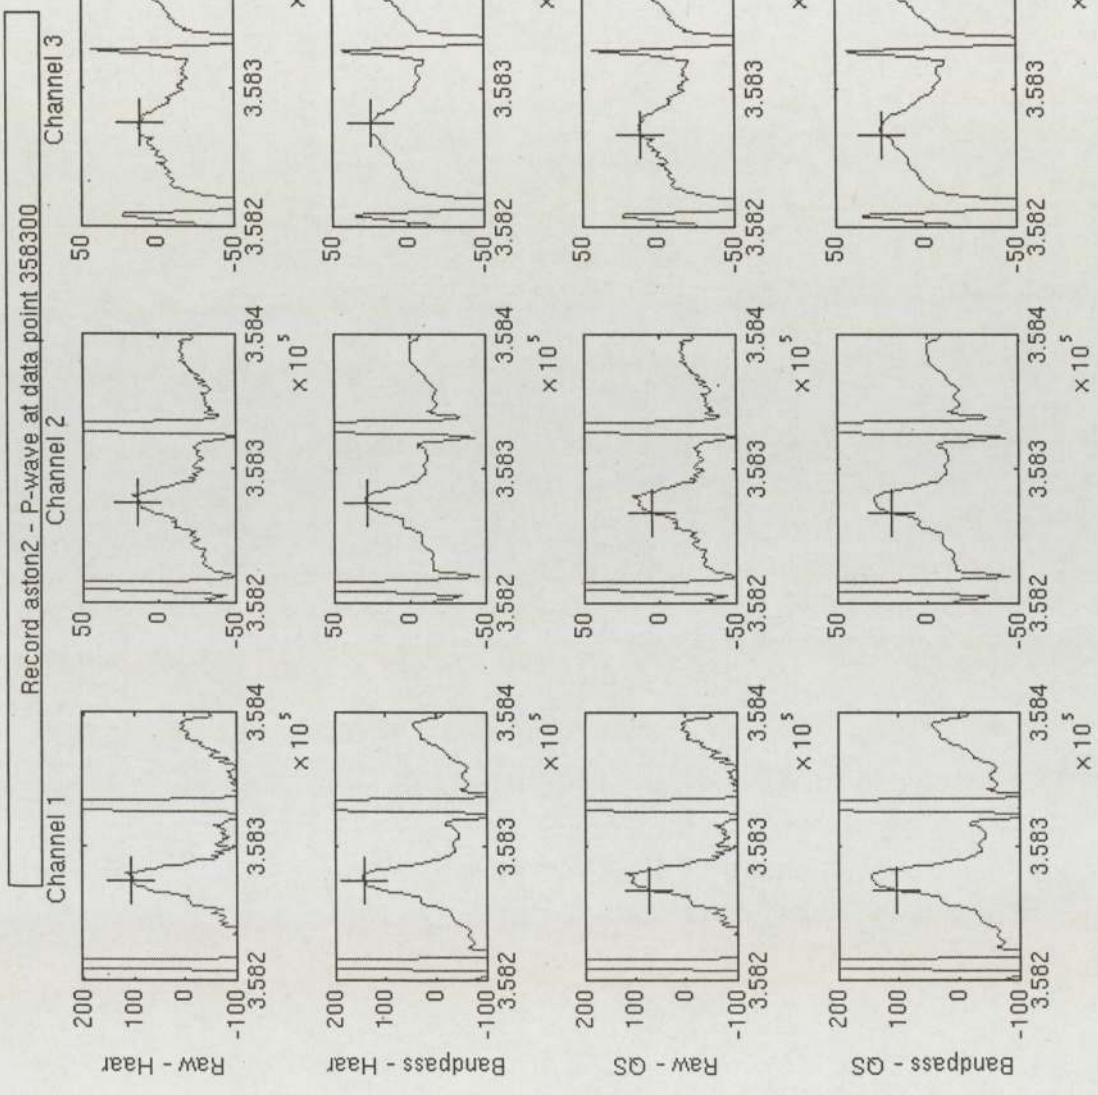

Figure 6.9: Beat Number 1498

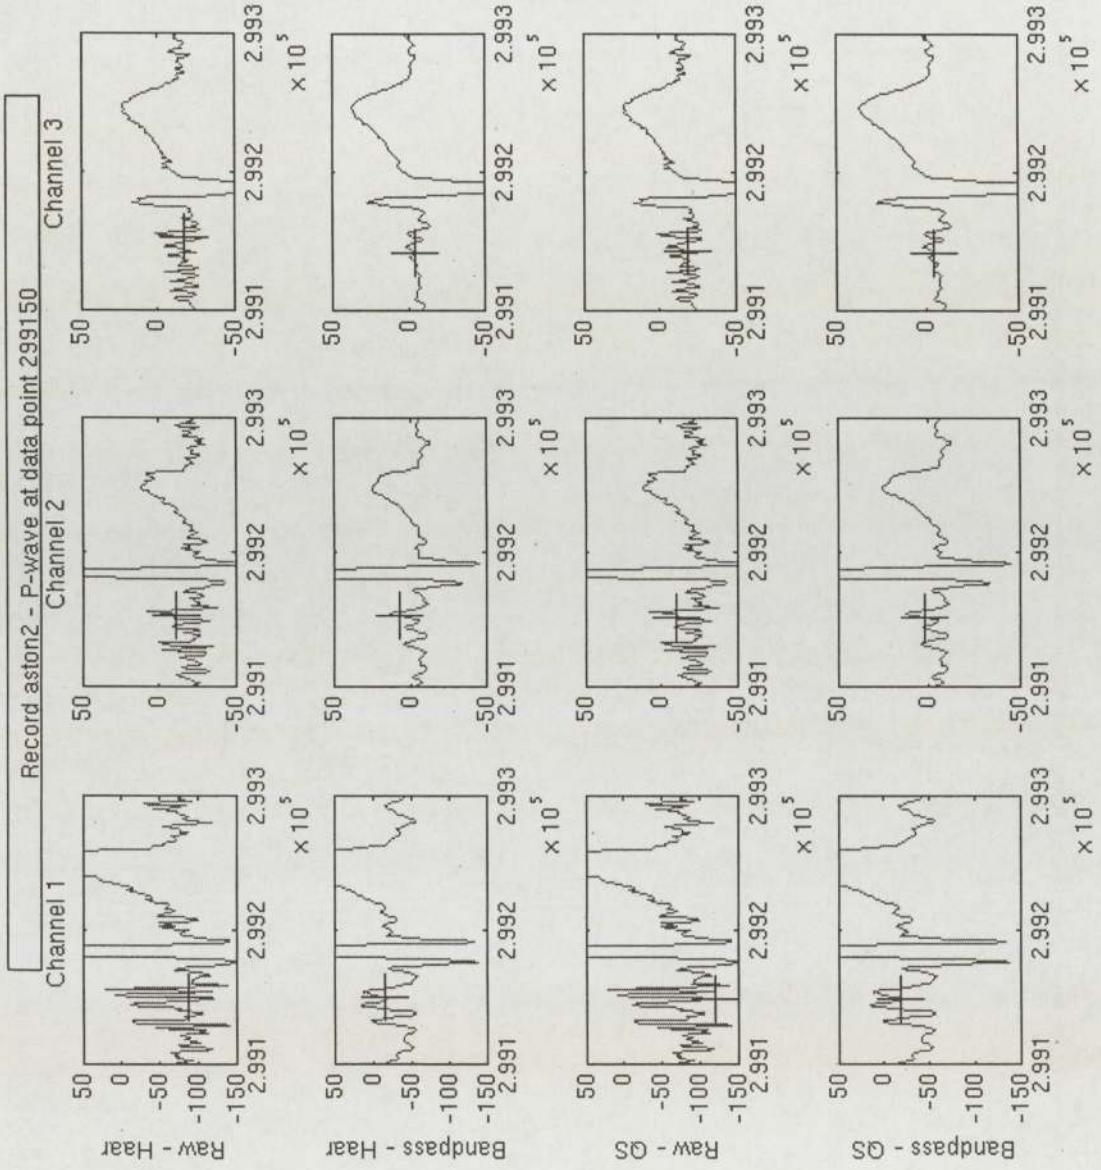

Figure 6.10: Beat Number 1239

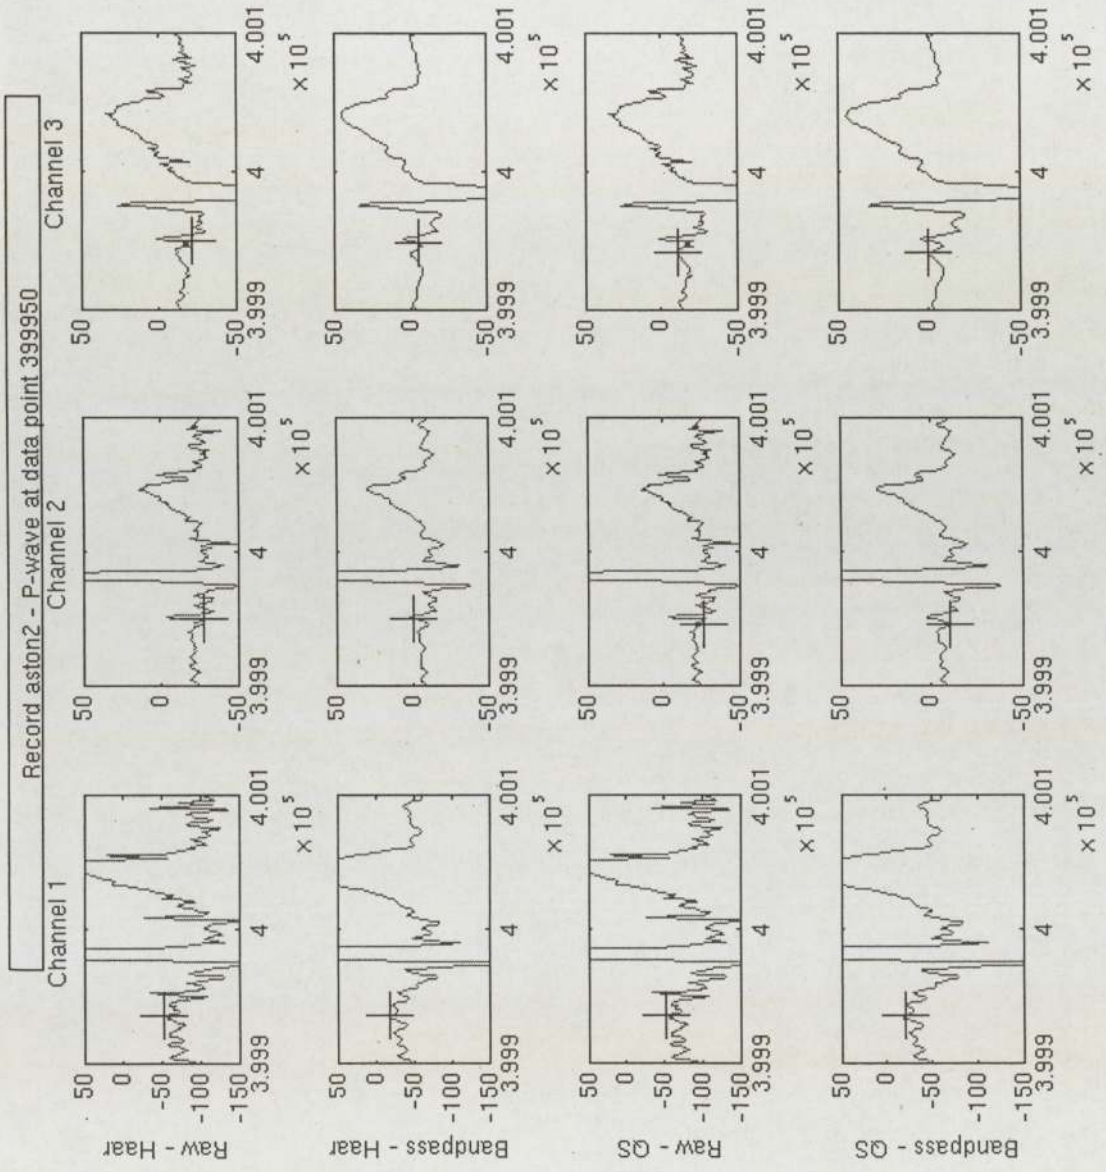

Figure 6.11: Beat Number 1682

Chapter 7

Summary

7.1 Introduction

The aim of this research was to develop a method for the reliable detection of P–waves in ambulatory ECG recordings. A review of the recent literature was carried out to discover which techniques had previously been used for overcoming this problem. This revealed that a variety of methods had been tried, support vector machines [Kikawa and Oguri 2004], time–frequency analysis [Stridh *et al.* 2004; Tu *et al.* 2004], genetic algorithms [Tu *et al.* 2004], fractional order differentiation [Goutas *et al.* 2005], entropy [Tu, Zeng, and Yang 2004] and wavelet transforms [Carrault *et al.* 2003; Clavier *et al.* 2002; Martínez *et al.* 2004; Senhadji *et al.* 2002; Sovilj, Jeras, and Magjarevic 2004; Sternickel 2002]. As can be seen from the above cited papers, wavelet transforms have been a popular choice for the analysis of P–waves. This fact can be attributed to their suitability for the analysis of non–stationary signals, of which the ECG is one such example.

After examination of the papers discovered during the literature search, it was decided that the wavelet transform offered the most likely means of successfully detecting P–waves in ambulatory ECG recordings. Although several researchers had previously, successfully, utilised wavelet transforms, there was no consensus as to which wavelet was most suited to the analysis of P–waves, with the Haar [Clavier *et al.* 2002], Coiflet [Sternickel 2002], quadratic spline [Martínez *et al.* 2004], Morlet [Clavier *et al.* 2002], and Daubechies [Senhadji *et al.* 1995] wavelets all having been used.

7.2 Analysis

The analysis comprised three distinct parts; denoising; the implementation of the algorithme à trous; and P–wave detection.

7.2.1 Denoising

Initially analysis was carried out on the denoising of ECG signals using wavelet transforms. Individual beat sections were isolated (100 beats) and padded with the end values to give a signal length of 256 data points, a signal length of 2^n ($n \in \mathbb{Z}^+$) being required in order to carry out the wavelet analysis. The Haar, Coiflet, and Daubechies wavelets were compared for their ability to successfully denoise the signals, the Haar because it is the simplest of all the wavelets, and the Coiflet and Daubechies because they had previously been used for denoising by other researchers [Sternickel 2002; Cuesta-Frau *et al.* 2000; Sánchez *et al.* 2002; Senhadji *et al.* 1995; Tasoulis *et al.* 2004; Li *et al.* 1995]. The software used for the denoising was WaveLab802, compatible with MATLAB, and developed by, and freely available from, the Department of Statistics at Stanford University[1].

The Haar wavelet proved to be the least successful at denoising the ECG signals, with the reconstructed beats having a very blocky appearance. Both the Coiflet and Daubechies wavelets produced very similar results for the removal of noise, giving far smoother reconstructions of the signals than the Haar wavelet. This visual impression was confirmed by an examination of the RMSE for each wavelet, with that of the Haar wavelet being higher than those of the Coiflet and Daubechies wavelets, which were both very similar, suggesting that both of these wavelets were equally suitable for denoising ECG recordings.

7.2.2 Algorithme À Trous

The reason the algorithme à trous was used for the detection of P–waves, instead of the usual method of performing wavelet analysis, which involves downsampling, was because this algorithm is time invariant. This time invariance stems from the fact that, unlike the algorithms normally used for wavelet analysis, there is no downsampling of the filter outputs; it being the downsampling which creates uncertainty in the temporal domain. The lack of downsampling, therefore,allows temporal information to be retained, hence allowing the timing of the P–wave to be accurately determined. This advantage, however, comes with the disadvantage of an increased computational load.

The algorithms for implementing the algorithme à trous with both the Haar and the quadratic

[1] URL:http:/www-stat.stanford.edu/~wavelab/

spline wavelet (as used by Martínez *et al.* [2004]) were tested on artificially produced ECG signals and were found to accurately determine the P–wave maxima. When noise was added to these artificial signals it was discovered that, whilst a moderate amount of noise had little effect on the P–wave detection, excessive noise caused the Haar wavelet to fail to detect the P–wave maxima. Excessive noise also caused the failure of the level 3 detail coefficients of the quadratic spline wavelet to correctly determine the position of the P–wave. The ability of the level 4 and 5 detail coefficients of the quadratic spline to detect the P–wave maxima was, however, relatively unaffected. The lack of success of the quadratic spline level 3 detail coefficients in detecting P–waves, can be attributed to the fact that these coefficients contain information concerning higher frequencies, such as are present in noise.

7.2.3 P–Wave Detection

The Haar and the quadratic spline wavelets were compared for their ability to accurately detect coupled P–waves using the algorithme à trous. The comparison was made on two separate sets of 100 random beats (consisting of 25 beats each from records *aston1, aston2, aston3* & *aston5*) and a continuous recording of the first half an hour of record *aston2*. The level four detail coefficients were used for the comparison as they had previously been found to give the best indication of the P–wave maxima (section 6.3.1).

Problems were encountered with the isolated P–wave sections from record *aston3*. This record differed from the other available records, both in its amplitude and in the first derivative of the lowpass filtered data, which caused difficulties in the correct isolation of the P–wave sections (see section 3.4).

Results for the other records were encouraging, with the vast majority of the P–waves being detected for all three channels, both for the raw and the bandpass filtered data, and with both wavelets. As a rough guide, if the standard deviation of the positions of the P–wave marker for all twelve signal/wavelet combinations of a single beat was greater than five, then at least one of the P–wave markers for that beat was likely to be inaccurate.

7.3 Conclusions

The Coiflet and Daubechies wavelets are both suitable for the removal of noise from ambulatory ECG recordings, whilst the Haar wavelet is fairly poor at denoising these signals. This can be attributed to the Haar wavelet not being as smooth as the other two wavelets.

From the research carried out, wavelet transforms appear to offer a valuable tool for the detection

of P–waves in ambulatory recordings, with the quadratic spline wavelet appearing to offer excellent detection possibilities, even in exceptionally noisy signals. The level 4 detail coefficients contain the most useful information concerning the position of the P–waves, being least susceptible to both noise and baseline drift. The algorithme à trous is a suitable method of implementing these wavelet transforms, being time–invariant, and therefore being able to accurately determine the temporal position of the P–wave maxima.

7.4 Future Work

The algorithm which has resulted from this research has only been developed on a few ECG recordings and it is desirable that further recordings should be obtained in order to reach a better assessment of its performance. The algorithm should also be tested on the standard ECG databases to give a true indication of its usefulness, and to provide a comparison with other P–wave detection algorithms.

Whilst the detection of coupled P–waves with the use of temporal windows derived from the location of the QRS complex has been relatively straightforward, it is often the case that the P–waves unassociated with a QRS complex are of the most interest. The algorithm will therefore have to be adapted to avoid the use of a search window, as there can be no certainty of an unassociated P–wave being contained within such a window.

As it will be necessary to detect P–waves without knowledge of the relative position of the QRS complex, some kind of threshold will have to be applied to the wavelet coefficients, necessitating the analysis to be performed on normalised signals. Work will therefore have to be carried out to determine the appropriate threshold which will maximise the number of correctly detected P–waves, whilst minimising the number of P–waves being erroneously detected.

Further work will have to be carried out to determine whether unassociated P–waves hidden beneath the T–waves of preceding beats can be adequately detected using the quadratic spline wavelet at the current sampling rate of 250Hz.

Once accurate algorithms have been developed for the detection of uncoupled P–waves, they will need to be incorporated into the existing Medilog Darwin software. This will involve converting the programs written in MATLAB into C, in order to be compatible with the existing software.

References

Anant, K. S., F. U. Dowla, and G. H. Rodrigue 1994, Apr. Detection of the Electrocardiogram P–Wave Using Wavelet Analysis. In H. H. Szu (Ed.), *Proceedings of the SPIE Conference on Intelligent Information Systems*, Volume 2242, pp. 744–749.

Andreoli, K. G., V. K. Fowkes, D. P. Zipes, and A. G. Wallace 1983. *Comprehensive Cardiac Care A Text for Nurses, Physicians, and Other Health Practitioners* (Fifth ed.). St. Louis: C V Mosby. editors.

Arrhythmia Alliance 2005. *Atrial Fibrillation (AF) Patient Information.* Stratford upon Avon: Arrhythmia Alliance.

Aytemir, K., B. Amasyali, G. Abali, S. Kose, A. Kilic, O. Orhan, L. Tokgozoglu, G. Kabakci, H. Ozkutlu, N. Nazli, E. Isik, and A. Oto 2005, Aug. The Signal-Averaged P-Wave Duration is Longer in Hypertensive Patients with History of Paroxysmal Atrial Fibrillation as Compared to Those Without. *International Journal of Cardiology* **103** (1), 37–40.

Bahoura, M., M. Hassani, and M. Hubin 1997. DSP Implementation of Wavelet Transform for Real Time ECG Wave Forms Detection and Heart Rate Analysis. *Computer Methods and Programs in Biomedicine* **52**, 35–44.

Burke Hubbard, B. 1998. *The World According to Wavelets* (Second ed.). Natick, USA: A K Peters.

Cañizares, M., N. Gómez, R. I. González, M. M. Rivero, J. Folgueras, and G. Meissimilly 2003, Sep. A New Method for Electrocardiogram Study. In *Proceedings of the 25th IEEE Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Volume 3, pp. 2643–2646.

Carrault, G., M.-O. Cordier, R. Quiniou, and F. Wang 2003. Temporal Abstraction and Inductive Logic Programming for Arrhythmia Recognition from Electrocardiograms. *Artificial Intelligence in Medicine* **28** (3), 231–263.

Choi, K.-J., D. C. Shah, P. Jais, M. Hocini, L. Macle, C. Scavee, R. Weerasooriya, F. Raybaud, J. Clementy, and M. Haissaguerre 2002, Dec. QRST Subtraction Combined with a Pacemap Catalogue for the Prediction of Ectopy Source by Surface Electrocardiogram in Patients with Paroxysmal Atrial Fibrillation. *Journal of the American College of Cardiology* **40** (11), 2013–2021.

Clavier, L., J.-M. Boucher, J.-J. Lepage, and J.-C. Cornily 2002. Automatic P-Wave Analysis of Patients Prone to Atrial Fibrillation. *Medical and Biological Engineering and Computing* **40** (1), 63–71.

Cohen, A., I. Daubechies, B. Jawerth, and P. Vial 1993. Multiresolution Analysis, Wavelets and Fast Algorithms on an Interval. *Comptes Rendus de L'Academie des Sciences - Series 1 - Mathematics* **316**, 417–421.

Cohen, A., I. Daubechies, and P. Vial 1993. Wavelets on the Interval and Fast Wavelet Transforms. *Applied and Computational Harmonic Analysis* **1**, 54–81.

Cuesta-Frau, D., D. Novák, V. Eck, J. C. Pérez-Cortés, and G. Andreu-García 2000. Electrocardiogram Baseline Removal Using Wavelet Approximations. *Analysis of Biomedical Signals and Images* **1**, 136–138.

Donoho, D. L. and I. M. Johnstone 1994, Sep. Ideal Spatial Adaptation by Wavelet Shrinkage. *Biometrika* **81** (3), 425–455.

Dunn, M. I. and B. S. Lipman 1989. *Lipman-Massie Clinical Electrocardiography* (Eighth ed.). Chicago: Year Book Medical Publishers.

Farrell, R. M., J. Q. Xue, and B. J. Young 2003, Sep. Enhanced Rhythm Analysis for Resting ECG Using Spectral and Time-Domain Techniques. In *Computers in Cardiology 2003*, Volume 30, pp. 733–736. IEEE.

Ferdi, Y., J.-P. Herbeuval, A. Charef, and B. Boucheham 2003, Sep. R Wave Detection Using Fractional Digital Differentiation. *ITBM-RBM* **24** (2), 273–280.

Fourier, J. 1955. *The Analytical Theory of Heat.* New York: Dover. Translated, with notes, by Freeman, A (Théorie Analytique de la Chaleur, originally published in 1822).

Goldberger, A. L., L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley 2000, Jun. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* **101** (23), e215–e220. Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215.

Goldman, M. J. 1986. *Principles of Clinical Electrocardiography* (Twelfth ed.). Los Altos, USA: Lange Medical Publications.

Goutas, A., Y. Ferdi, J.-P. Herbeuval, M. Boudraa, and B. Boucheham 2005, Apr. Digital Fractional Order Differentiation-Based Algorithm for P and T-Waves Detection and Delineation. *ITBM-RBM* **26** (2), 127–132.

Haar, A. 2006. On the Theory of Orthogonal Function Systems. In C. Heil and D. F. Walnut (Eds.), *Fundamental Papers in Wavelet Theory*, pp. 155–188. Princetown: Princetown University Press. Translated by Zimmermann, G (Zur Theorie der Orthogonalen Funktionensysteme originally published in 1910).

Ho, C.-S. and Y.-F. Yeh 2005. Electrocardiogram Diagnosis Using Hybrid Case-Based Reasoning. *Applied Artificial Intelligence* **19** (1), 43–68.

Hoffman, G. S., M. M. Miller, M. Kabrisky, P. S. Maybeck, and J. F. Raquet 2002, Dec. A Novel Electrocardiogram Segmentation Algorithm Using a Multiple Model Adaptive Estimator. In *Proceedings of the 41st IEEE Conference on Decision and Control*, Volume 3, pp. 2524–2529. IEEE.

Holschneider, M., R. Kronland-Martinet, J. Morlet, and P. Tchamitchian 1988, May. The "Algorithme A Trous". Technical Report CPT–88/P.2115, Centre de Physique Théorique.

Holschneider, M., R. Kronland-Martinet, J. Morlet, and P. Tchamitchian 1989. A Real–Time Algorithm for Signal Analysis with the Help of the Wavelet Transform. In J. M. Combes, A. Grossmann, and P. Tchamitchian (Eds.), *Wavelets: Time–Frequency Methods and Phase Space*, pp. 286–297. Springer–Verlag.

Holter, N. J. 1961, Oct. A New Method for Heart Studies. *Science* **134**, 1214–1220.

Kikawa, Y. and K. Oguri 2004, Nov. A Study for Excluding Incorrect Detections of Holter ECG Data Using SVM. In N. R. Pal, N. Kasabov, and R. K. Mudi (Eds.), *Proceedings of the 11th Internationl Conference on Neural Information Processing ,ICONIP 2004, Calcutta, India*, Volume 3316, pp. 1223–1228. Springer–Verlag.

Li, C., C. Zheng, and C. Tai 1995, Jan. Detection of ECG Characteristic Points Using Wavelet Transforms. *IEEE Transactions on Biomedical Engineering* **42** (1), 21–28.

Madias, J. E. 2004, Aug. Comparison of P Waves Recorded on the Standard Electrocardiogram, the "Lewis Lead," and "Saline-Filled Central Venous Catheter"-Based Intercardiac Electrocardiogram. *American Journal of Cardiology* **94**, 474–478.

Mallat, S. 1989, Jul. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11** (7), 674–693.

Marenco, J. P., H. Nakagawa, S. Yang, D. MacAdam, L. C. Xu, D. S. He, M. S. Link, M. K. Homoud, N. A. M. Estes, and P. J. Wang 2003, Jan. Testing of a New T-Wave Subtraction Algorithm as an Aid to Localizing Ectopic Atrial Beats. *Annals of Noninvasive Electrocardiology* **8** (1), 55–59.

Marriott, H. J. L. and M. B. Conover 1998. *Advanced Concepts in Arrhythmias* (Third ed.). St. Louis: Mosby.

Martínez, J. P., R. Almeida, S. Olmos, A. P. Rocha, and P. Laguna 2004, Apr. A Wavelet-Based ECG Delineator: Evaluation on Standard Databases. *IEEE Transactions on Biomedical Engineering* **51** (4), 570–581.

Martínez, J. P., S. Olmos, and P. Laguna 2000, Sep. Evaluation of a Wavelet–Based ECG Waveform Detector on the QT Database. In *Computers in Cardiology 2000*, Volume 27, pp. 81–84. IEEE.

Office for National Statistics 2004. *Mortality Statistics Cause* (30 ed.). DH2. London: ONS. Review of the Registrar General on Deaths by Cause, Sex and Age, in England and Wales, 2003.

Office for National Statistics 2005. *Mortality Statistics Cause* (31 ed.). DH2. London: ONS. Review of the Registrar General on Deaths by Cause, Sex and Age, in England and Wales, 2004.

Office for National Statistics 2006. *Mortality Statistics Cause* (32 ed.). DH2. London: ONS. Review of the Registrar General on Deaths by Cause, Sex and Age, in England and Wales, 2005.

Oldham, K. B. and J. Spanier 1974. *The Fractional Calculus*, Volume 111 of *Mathematics in Science and Engineering*. New York: Academic Press.

Papaloukas, C., D. I. Fotiadis, A. Liskas, and L. K. Michalis 2002. An Ischemia Detection Method Based on Artificial Neural Networks. *Artificial Intelligence in Medicine* **24**, 167–178.

Ros, E., S. Mota, F. J. Fernández, F. J. Toro, and J. L. Bernier 2004. ECG Characterisation of Paroxysmal Atrial Fibrillation: Parameter Extraction and Automatic Diagnosis Algorithm. *Computers in Biology and Medicine* **34**, 679–696.

Ryan, J., C. O'Sullivan, and C. Bell 2004. Real-Time Interactive Volumetric Animation of the Heart's Electrical Cycle from Automatically Synchronized ECG. *Computer Animation and Virtual Worlds* **15** (3–4), 353–360.

Sahambi, J. S., S. N. Tandon, and R. K. P. Bhatt 1997, Jan. Using Wavelet Transforms for ECG Characterisation. An On–Line Digital Signal Processing System. *IEEE Engineering in Medicine and Biology* **16** (1), 77–83.

Sánchez, C., J. Millet, J. J. Rieta, F. Castells, J. Ródenas, R. Ruiz-Granell, and V. Ruiz 2002, Sep. Packet Wavelet Decomposition: An Approach for Atrial Activity Extraction. In *Computers in Cardiology 2002*, Volume 29, pp. 33–36.

Schamroth, L. 1990. *An Introduction to Electrocardiography* (Seventh ed.). Oxford: Blackwell. Revised by Scamroth, C.

Senhadji, L., G. Carrault, J. J. Bellanger, and G. Passariello 1995, Mar/Apr. Comparing Wavelet Transforms for Recognizing Cardiac Patterns. *IEEE Engineering in Medicine and Biology Magazine* **14** (2), 167–173.

Senhadji, L., F. Wang, A. I. Hernandez, and G. Carrault 2002, Sep. Wavelets Extrema Representation for QRS-T Cancellation and P Wave Detection. In *Computers in Cardiology 2002*, Volume 29, pp. 37–40. IEEE.

Shyu, L.-Y., Y.-H. Wu, and W. Hu 2004, Jul. Using Wavelet Transform and Fuzzy Neural Network for VPC Detection from the Holter ECG. *IEEE Transactions on Biomedical Engineering* **51** (7), 1269–1273.

Sovilj, S., M. Jeras, and R. Magjarevic 2004, May. Real Time P-Wave Detector Based on Wavelet Analysis. In *Proceedings of the 12th IEEE Mediterranean Electrotechnical Conference, 2004, MELECON 2004*, Volume 1, pp. 403–406.

Sternickel, K. 2002, May. Automatic Pattern Recognition in ECG Time Series. *Computer Methods and Programs in Biomedicine* **68** (2), 109–115.

Stridh, M., L. Sörnmo, C. J. Meurling, and S. B. Olsson 2004, Jan. Sequential Characterization of Atrial Tachyarrhythmias Based on ECG Time-Frequency Analysis. *IEEE Transactions on Biomedical Engineering* **51** (1), 100–114.

Tasoulis, D. K., L. Vladutu, V. P. Plagianakos, A. Bezerianos, and M. N. Vrahatis 2004, Sep. Online Neural Network training for Automatic Ischemia Episode Detection. In L. Rutkowski, J. Siekmann, and R. Tadeusiewicz (Eds.), *Proceedings of the 7th International Conference on Artificial Intelligence and Soft Computing*, Volume 3070/2004, pp. 1062–1068. Springer–Verlag.

Tu, C., Y. Zeng, and X. Yang 2004. Nonlinear Processing and Analysis of ECG Data. *Technology and Health Care* **12** (1), 1–9.

Tu, C. Y., Y. J. Zeng, X. Y. Ren, S. C. Wu, and X. C. Yang 2004, Sep. Hybrid Processing and Time-Frequency Analysis of ECG Signal. In *Proceedings of the 26th IEEE Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Volume 1, pp. 361–364.

Viterbi, A. J. 1967, Apr. Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. *IEEE Transactions on Information Theory* **13** (2), 260–269.

Wagner, G. S. 2001. *Marriott's Practical Electrocardiography* (Tenth ed.). Philadelphia: Lippencott, Williams & Wilkins.

World Health Organization 1999. *The World Health Report 1999 Making A Difference.* Geneva: WHO.

Appendix A

Software

A.1 Missing Software

The WaveLab software attempted to call three MATLAB functions which were missing from the version of MATLAB being used (version 7.0). These functions were **choices.m** and **choicex.m**, both of which had been present in MATLAB6.5, and **fmin.m**, included in version 5.0. The Internet was searched in an attempt to find copies of the missing programs. Copies of **choices.m** and **choicex.m** were located on the Wave Analysis for Fatigue and Oceanography website[1], whilst a copy of **fmin.m** was found on the koders website[2]. Copies of all three programs were installed in the MATLAB folder in the home directory.

A.2 Program Editing

Problems were encountered on attempting to run `WLBrowser`. On investigation it was found that four of the WaveLab programs contained errors, **do_action.m**, **do_shrinkage.m**, **MakeSignal.m**, and **WLBrowser.m** itself.

do_action.m

This program used `switch` as a variable name. However, switch is a reserved keyword in MATLAB. The variable name was altered to `switc`.

do_shrinkage.m

This program contained an extra `end` on line 38, which was commented out.

[1] URL:http:/www.maths.lth.se/matstat/wafo/docomentation/wafodoc/index.html
[2] URL:http:/www.koders.com/

MakeSignal.m

This program used `i`, both for the square root of minus one, and also as a user defined variable. Its use as a user defined variable was replaced with `EYE` whenever it occurred.

WLBrowser.m

There was a problem with the backgrounds of the Signal/Reconstruction, Transform, and Auxiliary windows being black, meaning that no signals were visible on them. This was solved by commenting out the following three lines of code.

line 394 (`set(fig_hand1, 'Color', back_color);`),

line 400 (`set(fig_hand2, 'Color', back_color);`),

line 406 (`set(fig_hand3, 'Color', back_color);`).