Wavelet Signal Processing of Physiologic Waveforms

by R. G. Hohlfeld[1,2], C. Rajagopalan[3], and G. W. Neff[1]

ABSTRACT

Many physiological signals may be described either as isolated pulses or as quasi-periodic sequences of isolated pulses. Wavelets are a powerful tool for the representation and analysis of such physiologic waveforms because a wavelet has finite duration (compact support) as contrasted with Fourier methods based on sinusoids of infinite duration. We show two examples of physiological signal processing using wavelet bases. The first example is compression of electrocardiogram (ECG) signals using an Associated Hermite wavelet basis and the second example shows removal of artifact from non-invasive blood pressure (NIBP) measurements.

I. Introduction

Many physiologic signals can be reasonably characterized as isolated pulses or as sequences of pulses. For this reason, wavelet-based signal processing techniques appear as attractive alternatives to Fourier-based techniques for these signals. At the very least, the fact that wavelet bases have "compact support" or "essentially compact support" (that is, basis elements of finite duration or of effectively finite duration), guarantees that at least some candidate wavelet representations of these pulse-like signals will be more rapidly convergent than their corresponding Fourier series representations. This elementary property has been expressed in other more sophisticated ways; for example the process of taking a wavelet transform "decorrelates" a signal by concentrating signal energy in a relatively small number of coefficients (Daubechies and Sweldens 1996).

In any event, it is this property of reducing a signal to a comparatively small number of components that makes wavelet-based techniques potentially powerful for signal processing algorithms. In particular, this property has motivated much of the effort for development of wavelet-based signal compression algorithms, particularly for electrocardiogram (ECG) signals

[1] Wavelet Technologies, Inc., 664 Pike Ave., Attleboro, MA 02703, USA
[2] Center for Computational Science, Boston University, 3 Cummington St., Boston, MA 02215, USA
[3] Datascope Corporation, 800 MacArthur Blvd., Mahwah, NJ 07430, USA

(Ramakrishnan and Saha 1997; Hilton 1997; Djohan, Nguyen, and Tompkins 1995; Dong, Kim, and Pearlman 2000; Cárdenas-Barrera and Lorenzo-Ginori 1999).

From this elementary standpoint, it is clear that the greatest benefit from application of wavelet-based signal processing techniques is obtained when proper choice of a wavelet basis (in effect the shape of the wavelet) is made using detailed physiologic knowledge of the signal under consideration. To illustrate this idea, we shall show two physiologic signal processing algorithms using wavelets, although in comparatively unconventional ways. The applications are in ECG data compression and in artifact removal in Non-Invasive Blood Pressure (NIBP) measurements. In both applications the emphasis will be on obtaining pragmatically useful results rather than in achieving a high degree of mathematical rigor or elegance.

II. ECG Data Compression

II.1 Overview and Objectives

There is considerable commercial motivation for the development of effective ECG data compression. For example, in a typical clinical setting, a single lead (or channel) of ECG data will yield 250 samples per second of 12 bit data, resulting in 32.4 Mbytes of packed binary data per day per channel. Usually a patient will have two or three leads of ECG data being taken, with many more leads in some clinical situations. Furthermore, a cardiac ward will have many beds leading to a single central station, as many as 64 in a large state-of-the-art cardiac ward. This implies an amount of data accumulated in a central station of tens of Gbytes per day. Portions of this data may need to be archived for years. The other example is that of ambulatory patients on telepacks. One or multiple leads of ECG data need to be transmitted, and in order to minimize bandwidth and power consumption, data compression becomes a necessity.

Two arguments balancing this pressure for high compression of ECG data are: (1) the continuing drop in costs of mass-storage of data in computers, which is perceived as reducing the necessity for ECG data compression, and (2) concerns that a compression algorithm may distort clinically significant information and limit the usefulness of previously compressed data to physicians. In practice, the first argument is not compelling because requirements for data storage have historically grown at least as fast as the technologies for data storage. The second objection is more serious, and in the past some ECG compression algorithms have been used only for strictly limited diagnostic objectives, as in Holter monitors.

Our objective was to develop a high-fidelity compression algorithm that would not impair later physician diagnoses. To achieve this goal, we determined that we would be willing to trade off the very highest levels of performance in terms of compression ratio. We initially considered the conventional measure of fidelity of the compression and reconstruction process, the percentage root-mean-square difference (PRD) (Tompkins 1993)

$$\mathrm{PRD} = \left(\sqrt{\frac{\sum_{1}^{n} (x_i - \bar{x}_i)^2}{\sum_{1}^{n} (x_i)^2}} \right) \times 100\% \qquad (1)$$

where x_i is the initial signal value, $\bar{x}_i$ is the reconstructed signal value, and n is the number of samples in the ECG signal. However, we found that in practice users were unhappy with reconstructions with comparatively low PRD scores. Furthermore, other workers have noted that classical wavelet compression algorithms yielding small PRD values often are clinically unacceptable due to blocking and ringing in the reconstructed waveforms (see e.g. Strang and Nguyen 1996). These considerations motivated our search for a fidelity criterion closer to clinical practice.

The performance of ECG analysis algorithms for beat detection and beat typing is measured using databases like the MIT-BIH Arrhythmia Database and / or the AHA Database. These databases contain multiple files of ECG signals and each beat has been classified (normal, or one of several abnormal types) by cardiologists. The ability of the algorithm to classify these beats as accurately as the clinicians is a measure of its performance. Parameters like sensitivity and positive predictivity are used to gauge performance. Obviously any distortions introduced in the signals will tend to increase the possibility of mis-classification. In order to measure the effectiveness of the wavelet based compression scheme, the performance statistics (sensitivity, positive predictivity) of the ECG algorithm on the raw signal was compared to the numbers obtained using the compressed/decompressed signals for the standard databases. The effectiveness of compression is judged by how close the statistics for the two cases get. The rationale for this procedure was that by using a beat-typing software package we were addressing more of the issues related to the physiologically relevant features of the ECG signal than are addressed by the simple PRD criterion.

II.2 Selection of Wavelet

Considerable attention has been given to the compression of ECG data using classical methods based on discrete wavelet transforms (Strang and Nguyen 1996; Cetin, Tewfik, and Yardimci 1994; Sastry and Rajgopal 1996; Chen, Itoh, and Hashimoto 1994). Usually these approaches have been based on application of the Daubechies wavelet and usually with fixed length buffers of data. We chose instead to work with a basis of wavelet-like objects (in the sense of possessing an essentially compact support) and which possess the usual completeness and orthonormality properties desirable in a functional basis. The choice was further motivated by the objective that this wavelet-derived basis set should be rapidly convergent over a broad range of ECG beat types. Existence of representations in terms of this basis set for arbitrary physiologically meaningful ECG waveforms (e.g. ventricular ectopic beats) is guaranteed by the completeness of the wavelet-derived basis set.

We term the wavelet-derived basis set we used, the Associated-Hermite wavelets, defined by

$$U_n^{\lambda}(x) = \frac{1}{\sqrt{2^n\, n!\,\lambda\sqrt{\pi}}} H_n(x/\lambda)\exp\left(-x^2/2\lambda^2\right), \quad n = 0,1,2,\ldots \qquad (2)$$

where λ is a positive scaling parameter, and H_n is the nth order Hermite polynomial. The reader will note that $U_n^\lambda(x)$ are a Gaussian function ($n = 0$) and its derivatives (n 1), suitably normalized. A plot of the first few members of this basis set is given in Figure 1. An arbitrary signal $f(x)$ may be written in terms of this basis set by

$$f(x) = \sum_{n=0}^{\infty} \alpha_n \, U_n^\lambda(x) \tag{3}$$

with the weights α_n determined by the usual orthogonality condition

$$\alpha_n = \int_{-\infty}^{\infty} f(x) \, U_n^\lambda(x) \, dx \, . \tag{4}$$

We construct the compressed representation of the signal in terms of the α_n values rather than in terms of the original sampled data values. If the basis set is correctly chosen, the range of n will be fairly small (*i.e.* the series representation in equation (3) is rapidly convergent), the number of bits in the α_n values, together with the overhead required for organization of the data (see the following section) will be significantly smaller than the bits in the original data stream.

The constant λ may be chosen as convenient, though it must remain fixed in any given series expansion. The choice of this constant is generally made so as to yield a more rapidly convergent representation of the signal in terms of the Associated-Hermite wavelet basis. Information about the choice of λ must be carried along with the wavelet expansion coefficients (the α_n values) and the overhead information in the compression algorithm.

II.3. Implementation of the Compression Algorithm

We now discuss briefly the implementation of the ECG compression algorithm using the Associated-Hermite wavelets. Several specific issues in the implementation arise directly from the mathematical properties of the Associated-Hermite wavelets. These technical points, required to achieve a very high fidelity algorithm with an acceptable compression ratio are discussed in the following two subsections. The third subsection addresses the issue of overhead in the implementation of the ECG compression algorithm.

II.3.a. Segmentation of Data into Pulses

Although the Associated-Hermite wavelet basis is complete and thus can represent any signal in a convergent series, the lower order wavelets have an effective support limited to within a few times $\pm\lambda$. So expansion of waveforms with significant structure outside this range is possible in principle, but very slowly convergent in practice and thus not attractive in compression applications. For this reason, our compression algorithm operates by separating the ECG signal beat-by-beat into segments that are compressed separately. This increases the amount of overhead information because the start and stop times of the segments must be included in the overhead of the compression algorithm, but the overall gain in the efficiency of the expansion is substantial.

Identification of individual ECG beats is accomplished using a conventional QRS detection algorithm, such as the Pan and Tompkins algorithm (Tompkins 1993; Pan and Tompkins 1985). The origin for carrying out evaluation of the wavelet expansion coefficients by evaluation of Eq. (4) is usually taken at the location of the R-wave.

If a beat is not detected for a significant stretch of time, the algorithm automatically converts to processing data segments of fixed length. This allows the faithful representation of periods of asystole, ventricular fibrillation, and other similar events. The completeness of the Associated-Hermite wavelet basis guarantees that these waveforms can also be accommodated without change to the core algorithm.

The fundamental reason that ECG compression is regarded as a difficult problem is that the ECG waveform contains clinically significant information on a wide variety of time scales. In Fourier language, we would say that the ECG spectrum is very broad and for this reason Fourier methods are unattractive for ECG compression. This problem is mitigated somewhat, but not entirely eliminated by the choice of the Associated-Hermite wavelet basis set. In particular, the R-wave typically has a much shorter characteristic time scale than the rest of the ECG waveform.

We chose to address this problem in a pragmatic fashion rather than with maximum mathematical elegance. We chose to subtract off the R-wave in each QRS complex and linearly interpolate the missing R-wave component in the remaining ECG pulse. Then the R-wave and the remaining ECG pulse are compressed separately, each with its own appropriate value of λ. This adds to the overall overhead in bits of the compression algorithm, but leads to a substantial gain in compression ratio and reconstruction fidelity for two reasons:

1. The R-wave expansion in Associated-Hermite wavelets is very rapidly convergent and almost always requires only a few terms.
2. The value of λ chosen for compression of the rest of the ECG pulse may then be chosen so as to improve the convergence (*i.e.* minimize the number of terms required).

II.3.b. Removal of Signal Baseline

Due to the completeness property of the Associated Hermite wavelet basis set, it is possible to represent a constant offset or nonzero values at the ends of a data segment, however it is very inefficient to use the Associated Hermite wavelet expansion in this way. This fact is immediately clear from an examination of the wavelet basis functions shown in Figure 1. The Associated Hermite wavelet bases (with the exception of the zeroth basis function) all oscillate about zero and the nth wavelet basis goes to zero like $x^n \exp(-x^2)$ as $x \to \pm\infty$. The effective width of the wavelet basis functions increases as the order increases, which is why it is possible to represent nonzero values at the ends of a data interval, but the number of terms required is prohibitive (with corresponding poor compression performance).

A practical solution is to subtract a linear baseline from the data so that the baseline-subtracted pulse begins and ends with a zero value. This leads to a more rapidly convergent wavelet basis representation and improved compression ratio. The overhead in the data header for each pulse is increased because it is necessary to include two pieces of information to allow reconstruction of the linear baseline (*i.e.* the endpoint values or the initial value and a value of the slope of the linear baseline). For practical ECG waveforms, the increase in compression ratio performance more than compensates for the increase in size of the data header.

II.3.c. Preparation of Data Headers

Each pulse must have associated with it the full set of information required for its reconstruction and insertion at the proper point in the ECG datastream. Waveform reconstruction requires more information than simply the wavelet coefficient values (the α_n's). At a minimum, we also require the beginning and ending times of the data segment, the origin of time within the pulse used for the Associated-Hermite wavelet expansion (usually the time of occurrence of the R-wave), the λ value used in the Associate-Hermite expansion, and the information required to reconstruct the baseline of the data segment. If the R-wave is to be compressed separately, the additional information for waveform construction must also include the beginning and endpoints of the R-wave and the corresponding value of λ used for the R-wave. Most of this information can be adequately represented as 16-bit integers, but the information to locate the pulse absolutely in the datastream requires more bits, typically a 32-bit integer.[4]

The highest average compression ratio will be achieved by using no more Associated-Hermite coefficients than necessary for a desired level of precision. This requires a format that can accommodate a variable number of coefficients, for example a specification of the number of coefficients in the expansion (up to some maximum order) followed by the coefficient values themselves.

II.4. Results

We now present results indicating the performance of the Associated-Hermite ECG compression algorithm. Because the algorithm operates on beats as units, we first show the operation of the algorithm on a single beat. An example from the MIT-BIH Arrhythmia Database Record 105 is shown in Figure 2. Eight terms were used in the Associated-Hermite expansion of the R-wave and sixteen terms were used in the expansion of the remainder of the ECG waveform. Overall correspondence between the original and reconstructed waveform is very high, with an *rms* error of only 0.0298 millivolt. The reconstructed waveform is adequate for all clinical purposes since the primary difference between the original and reconstructed waveform is the removal of electrical noise in the reconstructed waveform. A single ventricular ectopic beat (VEB) taken from MIT-BIH Arrhythmia Database Record 124 is shown together with its reconstruction after compression in Figure 3. Although the shape of this beat is very different from the previous figure, the fidelity of reconstruction is very high and relevant clinical features of the beat are well preserved.

A longer stretch of data from Record 116 of the MIT-BIH Arrhythmia Database is shown in Figure 4. Both leads of ECG data are shown in the figure, with the traces reconstructed from the compressed data below the original data traces. Fidelity of the compression algorithm is seen to be very high for normal sinus ECG waveforms. A voltage step artifact is seen at the end of the lead 1 data in this trace, which is also represented in the compressed data (though some ringing is apparent because of the extreme rate of change in voltage which cannot be accommodated by a truncated Associated-Hermite expansion).

Figure 5 shows ECG data of varying beat morphology, normal sinus beats and ventricular ectopic beats from Record 124 of the MIT-BIH Arrhythmia Database. The fidelity of reconstruction of the Associated-Hermite algorithm is seen to be very good for both beat types.

An average compression ratio slightly better than 12:1 was obtained for the records used in this study. Reconstructed waveforms were processed with an automated beat-typing program and the results compared to those obtained using the original data to determine whether the

---

[4] Consider that at a modest sampling rate of 100 Hz, there are 8.64 million samples per day, and thus representing a unique time to that precision requires 24 bits.

process of compression had distorted the ECG waveforms in a clinically significant way. The results of this test are shown in Table 1. The results indicate that the difference in performance for the two cases is within 5% and strongly suggests that no significant distortion of the ECG waveform has resulted from the compression with the Associated-Hermite algorithm.

III. Artifact Removal in Non-Invasive Blood Pressure (NIBP) Waveforms

We turn now to a more conventional application of wavelet methods to processing of a medical waveform. This application is more conventional because it uses a wavelet transform based on the application of a single wavelet, rather than a basis set constructed from a family of mathematically related wavelets. Again, the choice of a wavelet with appropriate morphological characteristics relative to the physiological signal under consideration is crucial to the success of the application.

III.1 General Description of NIBP Artifact Removal Problem

The NIBP measurement system used in most patient monitors measures the small fluctuations in pressure in a blood pressure cuff (applied to one of the patient's limbs) to obtain a determination of the patient's systolic and diastolic pressure. Usually the mean arterial pressure and pulse rate are obtained as well. These pressure fluctuations are usually termed "oscillometric pulses" and, as the cuff pressure is slowly decreased (either in steps or by a continuous bleed), the oscillometric pulses increase in amplitude as the systolic pressure is passed, rising to a maximum at the mean arterial pressure, and decrease in amplitude to a low level as the cuff pressure decreases to the patient's diastolic pressure. (L.A Geddes, 1991) A typical oscillometric pulse profile is shown in Figure 5, which is obtained by applying a bandpass filter (with corner frequencies at ½ Hz and 15 Hz) to the cuff pressure.

The oscillometric pulse profile has a very small amplitude, typically a maximum amplitude only on the order of 1 mm Hg. Small fluctuations in cuff pressure can also arise due to patient motion during measurement that are of comparable amplitude, as well as "transport artifacts" arising when a patient is being transported by ambulance or medevac helicopter.

Particularly problematic is patient tremor, frequently occurring in operating rooms, which can have large amplitude and a frequency close to the patient heart rate. This close proximity in frequency makes removal of tremor artifact from NIBP oscillometric signals especially difficult by Fourier techniques.

III.2 The Dyadic Discrete Wavelet Transform (DWT)

The NIBP artifact removal algorithm we describe here makes use of a conventional discrete wavelet transform. A dyadic discrete wavelet transform is performed on the data segment (see *e.g.* [Strang and Nguyen 1996, Mallat 1999]). Each level of the dyadic wavelet transform is implemented using a highpass and a lowpass filter constructed from the coefficient wavelet. Usually these highpass and lowpass filters are referred to as a pair of quadrature mirror filters. Also, in carrying out the forward wavelet transform described here, this filter pair is referred to as "analysis filters", because they break down the signal into wavelet components. The highpass and lowpass filters generate a "detail" and an "average" signal, respectively, which contain small-scale and longer-scale information of the original signal. The high-pass and low-pass signals are, at this stage, each equal to the length of the original signal (neglecting end effects, which may add slightly to the length depending on the type of boundary conditions imposed, but which we will not consider in detail here). So the information in the high-pass and low-pass signals has a redundancy of 2×. This redundancy is removed by decimating both the detail and average signals by 2×.

The next level of the discrete wavelet transform operates on the average signal generated from the previous level to generate a new detail and a new average signal. Because of the decimation step performed previously, effectively the wavelet has been dilated by 2×. In this way, the dyadic discrete wavelet transform implements the classic description of a wavelet transform as a decomposition of a signal using a basis composed of dilated and translated copies of the "mother" wavelet basis function. Note that the new average and detail signals generated at this level of the transform are 4× smaller than the original signal.

This dyadic wavelet transform process may be continued until the average signal reaches the length of a single sample (or a length too small for further application of the analysis filter pair, in some types of boundary conditions). The wavelet transform may also be truncated at any level of the process before an average signal of length of one sample is reached. In any event, the dyadic discrete wavelet transform consists of the set of detail signals generated at each level of the transform, together with the average signal generated at the highest level (shortest length signals) of the transform.

A remarkable feature of many useful wavelet transforms, is that they obey a perfect reconstruction theorem. That is the dyadic discrete wavelet transform may be inverted to recover the original signal **exactly**. (In practice, the exact reconstruction will be limited by the details of representing the data in finite-length computer memory, but this limitation can be shown not to lead to intractable difficulties such as numerical instabilities in the inversion process.) The inversion process is carried out first by upsampling (or expanding) the highest level detail and average signals. Upsampling is carried out by inserting zeros between samples of the signal to be upsampled. Then the upsampled average and detail signals are run through synthesis filters (which are generally constructed mathematically from the analysis filters) and added together. The sum signal is the average signal for the next lowest level of the wavelet transform. This process is carried out at each lower level until the original signal is recovered at the lowest level as the zero level average signal. For the details of the algorithm, the reader is referred to any of the excellent standard texts. [See *e.g.* (Kaiser 1994; Strang and Nguyen 1996; Mallat 1998).]

III.3 Separation of Artifact Energy from the Physiological Waveform by the DWT

The wavelet-based artifact elimination algorithm is based on the observation that the dyadic discrete wavelet transform puts the physiologic oscillometric waveform in a very different region of the transform plane than the signal components attributable to artifact. This is best shown by a plot of the discrete wavelet transform in which the absolute values of the wavelet transform coefficients, as a function of time (sample number) and level of the transform (scale), constitute a time-scale plot. Usually plots of this type are referred to as "scalograms".

When artifact is present, usually most prominent in the short scale portions of the scalogram, this portion of the discrete wavelet transform may be zeroed out. The modified discrete wavelet transform may then be inverted to yield a reconstruction of the oscillometric signal with artifact substantially reduced. The reconstructed oscillometric signal may then be used as an input to an NIBP pressure determination algorithm in the usual way for the measurement of desired patient pressure values.

III.4 Results on NIBP Artifact Removal

We now briefly discuss some results on NIBP artifact removal illustrating the wavelet-based artifact removal procedure. Figure 6 shows the bandpass filtered pressure signal (oscillometric signal) for a human subject in the upper pane. The independent variable axis in this and following plots is data samples (taken at 100 samples/second) after cuff deflation begins. The character of the oscillometric signal is clear with one pulse per heartbeat. A pressure determination algorithm would assign the systolic pressure to the cuff pressure corresponding to sample 700 and diastolic pressure to around sample 1700 (approximately). The scalogram for the oscillometric trace is shown in the lower pane with the color coding the absolute value of the discrete wavelet transform, maximum values in magenta. Note that the time information is preserved in the wavelet transform, since oscillometric pulses are evident across a range of scales.

NIBP artifact data is difficult to obtain in clinical situations because of considerations of patient safety under difficult conditions, so we use data generated by a simulator (Bio-Tek BP Pump). Figure 7 shows data heavily contaminated by noise from transporting a patient over a gravel road. In the upper pane the underlying oscillometric signal is almost completely obscured by the artifact signal. In the scalogram, the artifact component is most prominent in the shorter scale (lower portion) of the scalogram.

In Figure 8 the scalogram of Figure 7 is modified by zeroing out the lower portion of the scalogram (levels 4 through 7). Usually a smaller number of wavelet scales may be removed, but we have chosen an extreme case for illustrative purposes. The upper pane shows the reconstructed oscillometric signal generated by performing the inverse wavelet transform on the remaining wavelet transform data. Although some small distortion of the oscillometric signal has been introduced, the reconstructed oscillometric signal may be used for determination of NIBP values in the usual way.

IV. Summary

We have given two illustrative examples applying wavelet techniques to the representation and analysis of physiological waveforms. The first example applied the Associated Hermite wavelet basis to the compression of ECG data. The objective of the algorithm was to provide a very high fidelity compression and reconstruction of the ECG waveform so as not to introduce artifacts limiting clinical use of the reconstructed waveform. Towards this end, we used a commercially available beat classification program to evaluate the performance of the compression program, rather than the *rms* error of the reconstruction. When applied to the MIT-BIH Arrhythmia Database, the classification error rate of reconstructed waveforms was only 0.33% greater than for the original data.

We also applied a wavelet-based algorithm for the removal of artifact from NIBP oscillometric data. This algorithm exploited the fact that the physiological and artifact components of the NIBP signal are mapped to different portions of the discrete wavelet transform plane. Removal of those components associated with artifact allows reconstruction of the physiological waveform by performing the inverse wavelet transform.

REFERENCES

AAMI, **Testing and Reporting Performance Results of Cardiac Rhythm and ST Segment Measurement Algorithms**, AAMI EC-57, Recommended practice, (1998).

The American Heart Association Database for Evluation of Ventricular Arrhythmia Detectors, ECRI, 5200 Butler Pike, Plymouth PA 19462 (USA).

Cárdenas-Barrera, J. and Lorenzo-Ginori, J., “Mean-Shape Vector Quantizer for ECG Signal Compression”, *IEEE Trans. Biomed. Eng.*, **46**, 62-70 (1999).

Cetin, A.E.; Tewfik, A.H.; Yardimci, Y., “Coding of ECG signals by wavelet transform extrema”, Time-Frequency and Time-Scale Analysis, 1994., *Proceedings of the IEEE-SP International Symposium on* , Page(s): 544 -547

Chen, J.; Itoh, S.; Hashimoto, T., “Wavelet transform based ECG data compression with desired reconstruction signal quality”, *Information Theory and Statistics*, 1994. Proceedings., 1994 IEEE-IMS Workshop on , Page(s): 84.

Daubechies, I. and Sweldens, W., “Factoring Wavelet Transforms into Lifting Steps”, Preprint, Bell Laboratories, Lucent Technologies (1996).

Djohan, A., Nguyen, G., and Tompkins, W., “ECG Compression Using Discrete Symmetric Wavelet Transform”, presented at the 17th Int. Conf. IEEE Medicine and Biology (1995).

Dong, L., Kim, D., and Pearlman, W., “Wavelet Compression of ECG Signals by the Set Partitioning in Hierarchical Trees Algorithm”, *I*

Geddes, L.A, **Handbook of Blod Pressure Measurement**, The Humana Press, Clifton, NJ (1991)

Hilton, M., “Wavelet and Wavelet Packet Compression of Electrocardiograms”, *IEEE Trans. Biomed. Eng.*, **44**, 394-402 (1997).

Lux, Robert L., “Karhunen-Loeve Representation of ECG Data”, *Journal of Electrocardiology*, vol. 125 Supplement, pp. 195-198.

The Massachusetts Institute of Technology-Beth Israel Hospital Arrhythmia Datbase, MIT Room E25-505, Cambridge MA 02139 (USA).

Mallat, S., **A Wavelet Tour of Signal Processing**, Academic Press, San Diego, 1998.

Pan, J. and Tompkins, W. J., “A Real-Time QRS Detection Algorithm”, *IEEE Trans. Biomed. Eng.*, **BME-32**:230-236 (1985).

Ramakrishnan, A. and Saha, S., “ECG Coding by Wavelet-Based Linear Prediction, *IEEE Trans. Biomed. Eng.*, **44**, 1253-1261 (1997).

Sastry, R.V.S.; Rajgopal, K., “ECG compression using wavelet transform”, *Engineering in Medicine and Biology Society*, 1996 and 14th Conference of the Biomedical Engineering Society of India. An International Meeting, Proceedings of the First Regional Conference., IEEE , Page(s): 2/52 -2/53

Strang, Gilbert, and Nguyen, Truong, **Wavelets and Filter Banks**, Wellesley-Cambridge Press, Wellesley, MA (1996).

Tompkins, Willis J. (ed.), **Biomedical Digital Signal Processing: C-Language Examples and Laboratory Experiments for the IBM PC**, Prentice Hall, Englewood Cliffs, NJ (1993).

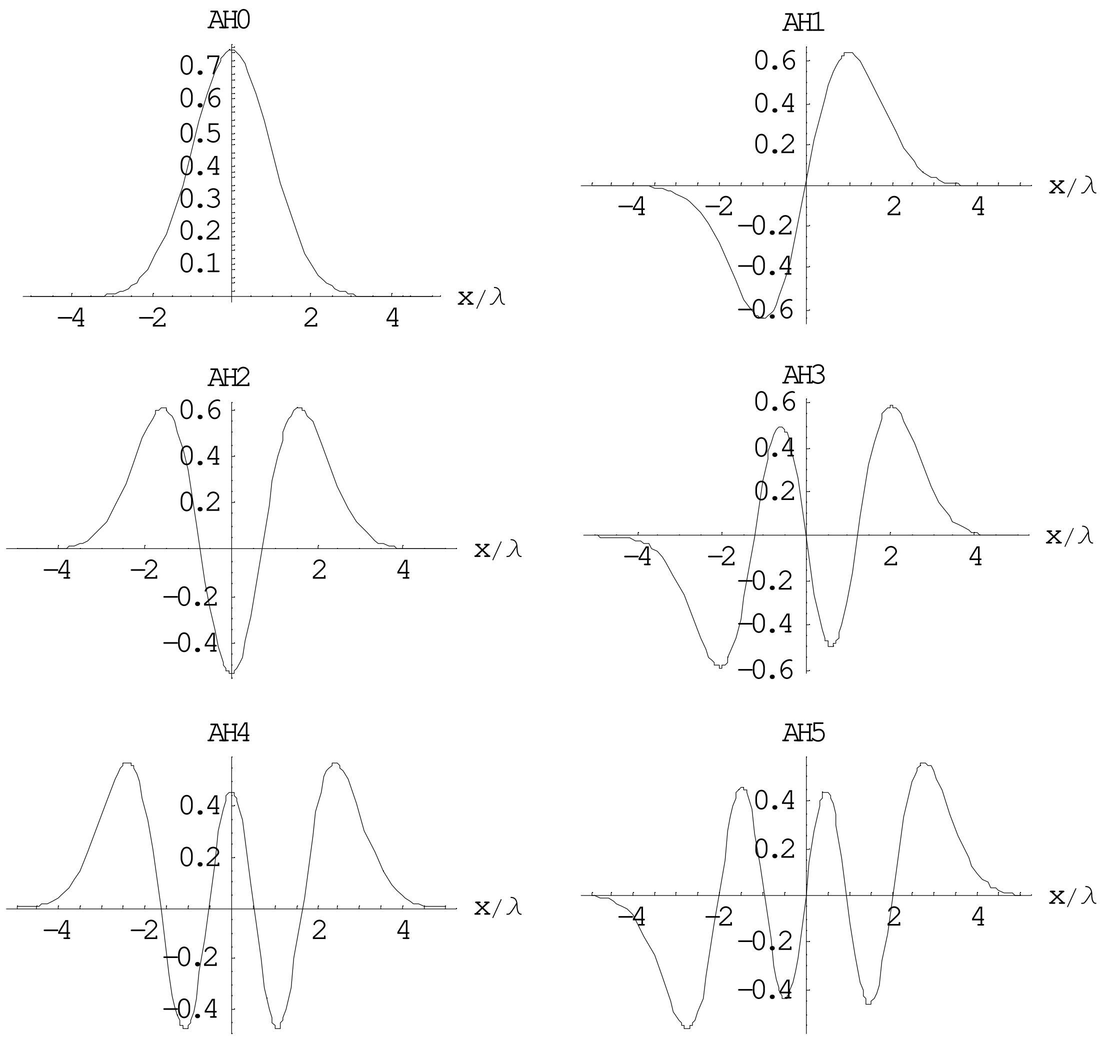

Figure 1: The first six Associated Hermite wavelets $U_n^\lambda(x)$, $n = 0,\ldots,5$ *plotted as a function of* x/λ. *The individual wavelets are labeled as AHn.*

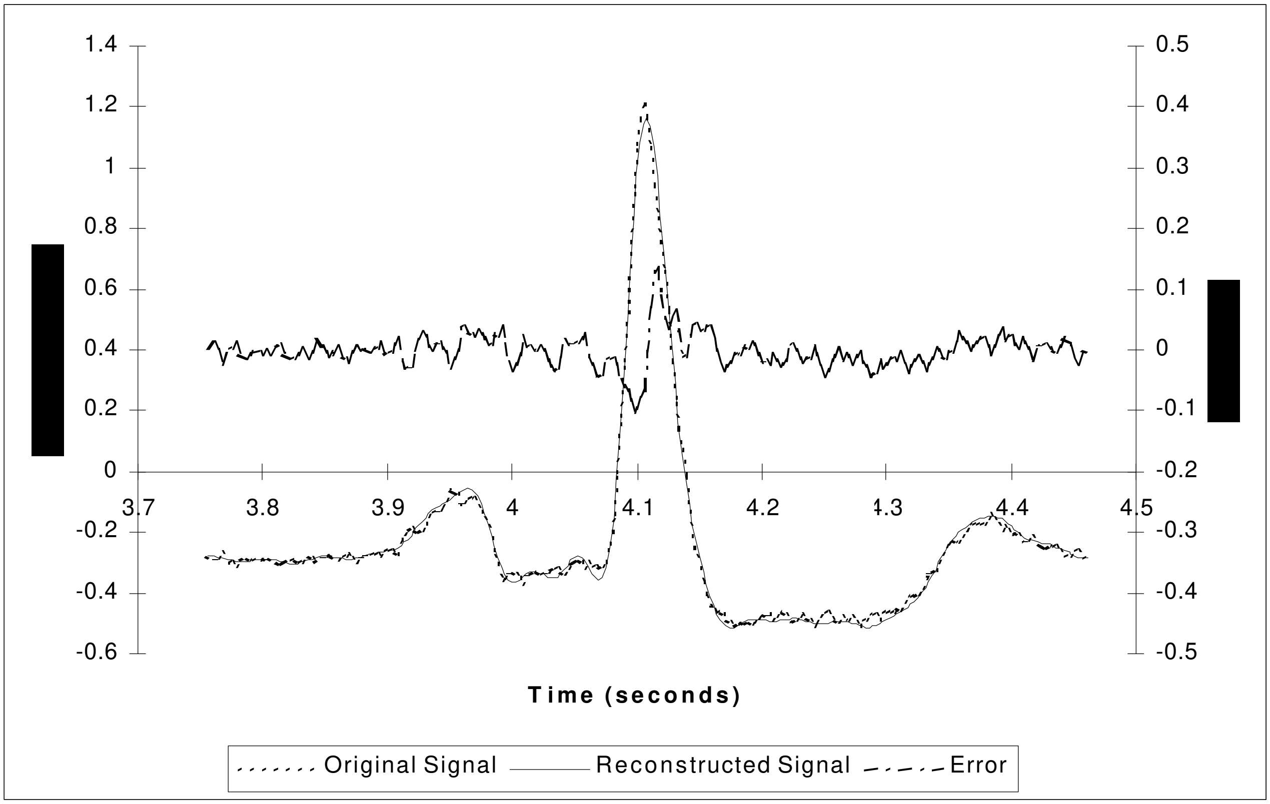

Figure 2: Application of the Associated Hermite wavelet compression algorithm to an ECG wave from the MIT-BIH Arrhythmia Database, Record 105. The original data is shown with the waveform obtained compression and reconstruction. The fidelity of the waveform after reconstruction is excellent and adequate for diagnostic purposes such as ST segment analysis. The rms error in the reconstruction is 0.0298 millivolts. The error in the reconstruction is shown with the scale on the right hand side of the figure.

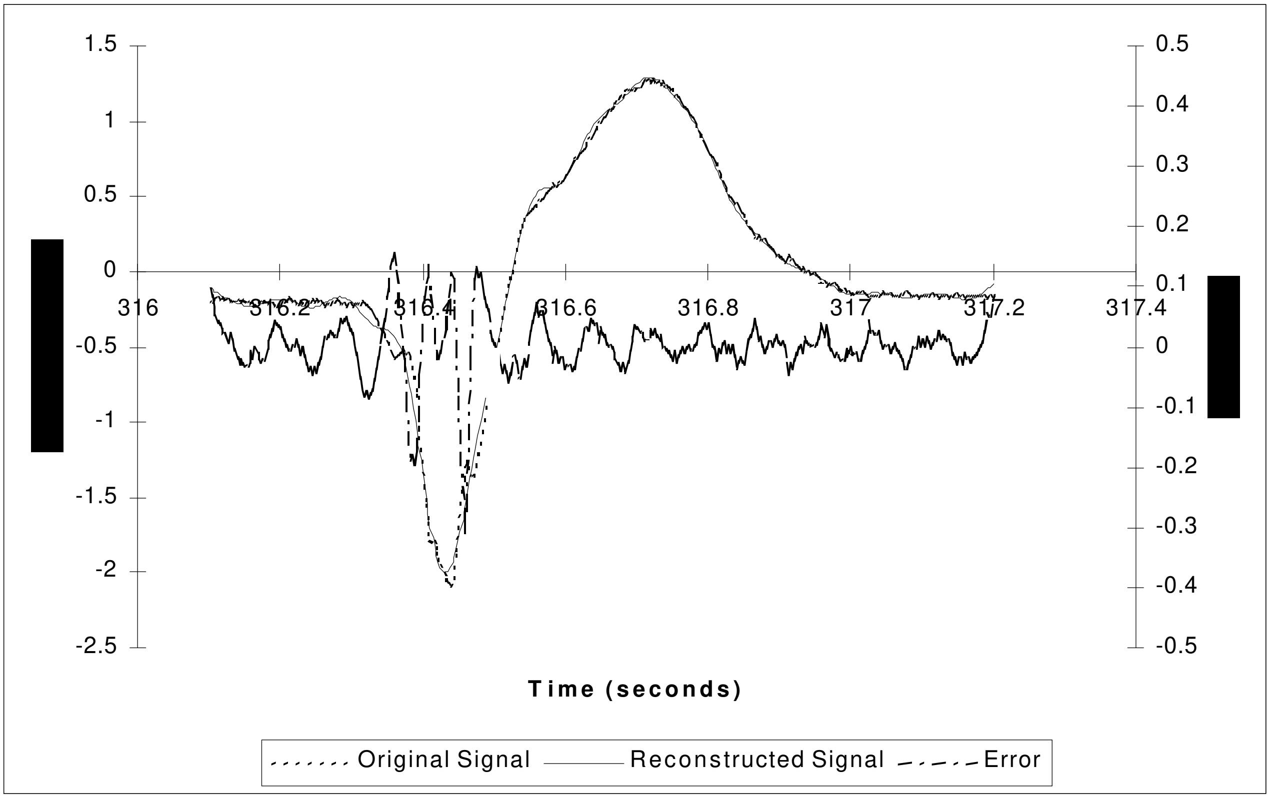

Figure 3: Results from the compression and reconstruction of a single ventricular ectopic beat (VEB) taken from Record 124 of the MIT-BIH Arrhythmia Database. Although the waveform shape is very different from a normal sinus beat, the fidelity of the compression and reconstruction algorithm remains very high. Again the error signal axis is on the right hand side of the figure.

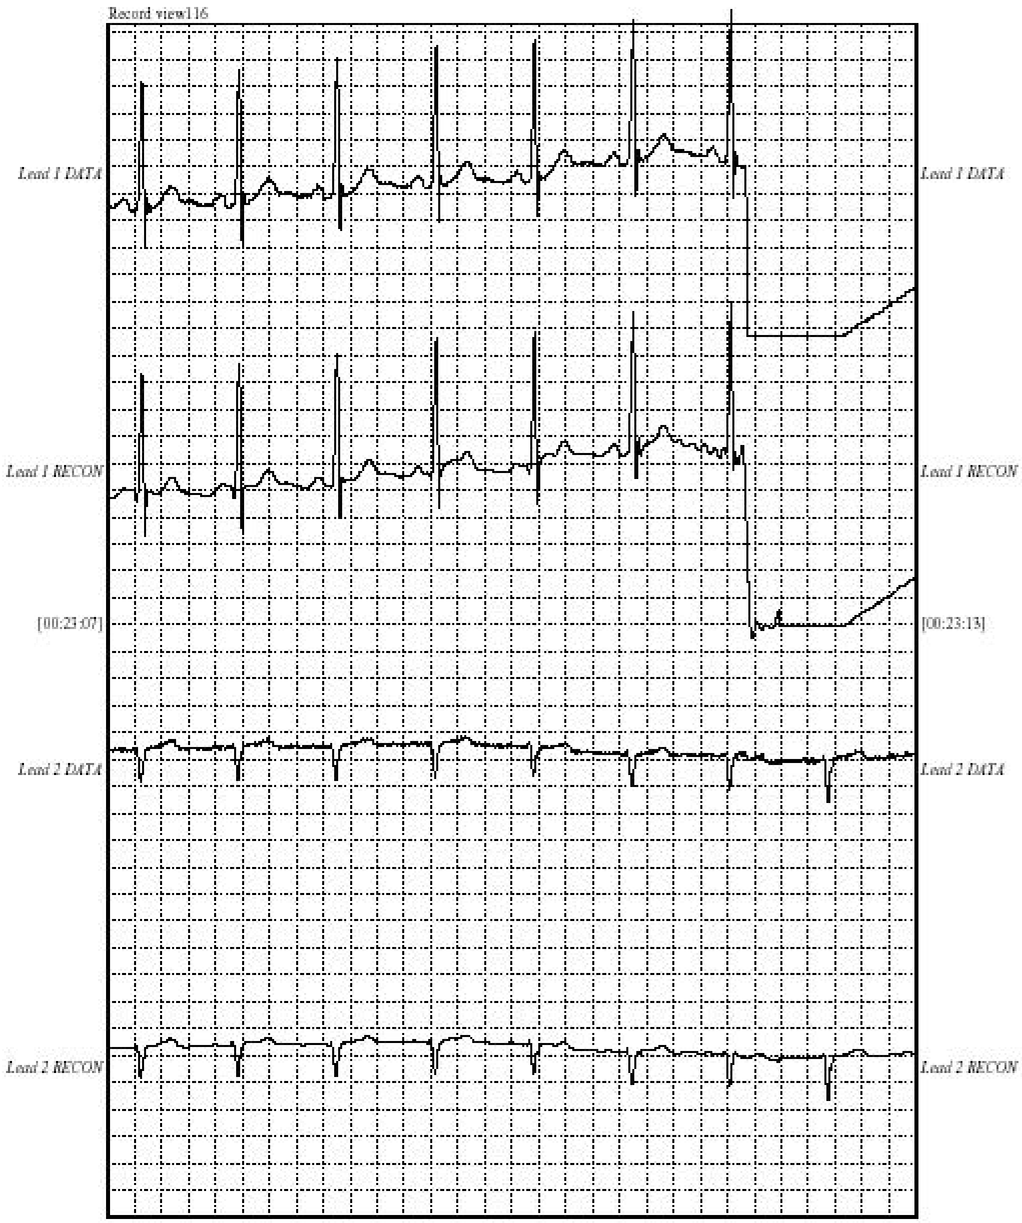

Figure 4: Original and reconstructed data from MIT-BIH Arrhythmia Database Record 116. Both leads from the database are shown, with original data traces above the corresponding reconstructed traces from data compressed by the Associated-Hermite algorithm. Performance of the algorithm over normal sinus data is seen to be excellent. At the end of a plot, lead 1 shows an artifact which is also captured in the reconstruction.

Figure 5: A segment of Record 124 of the MIT-BIH Arrhythmia Database showing normal and ventricular ectopic beats. Compression fidelity is seen to be very high for both beat types.

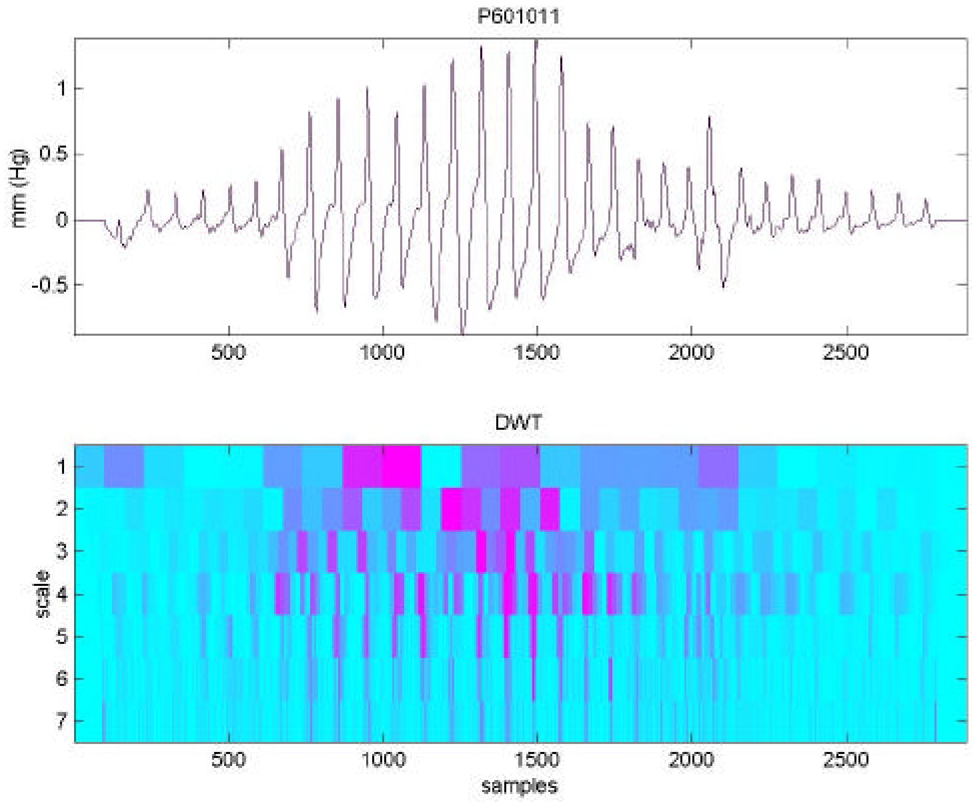

Figure 6: Oscillometric data and the corresponding discrete wavelet transform (lower pane) The independent variable in both plots is the sample number after the maximum in cuff pressure. Each pulse in the upper plot corresponds to one heartbeat. The maximum coefficient amplitude in the discrete wavelet transform is indicated by magenta. The wavelet transform is carried out through 7 scales related by factors of 2 in decimation between scales.

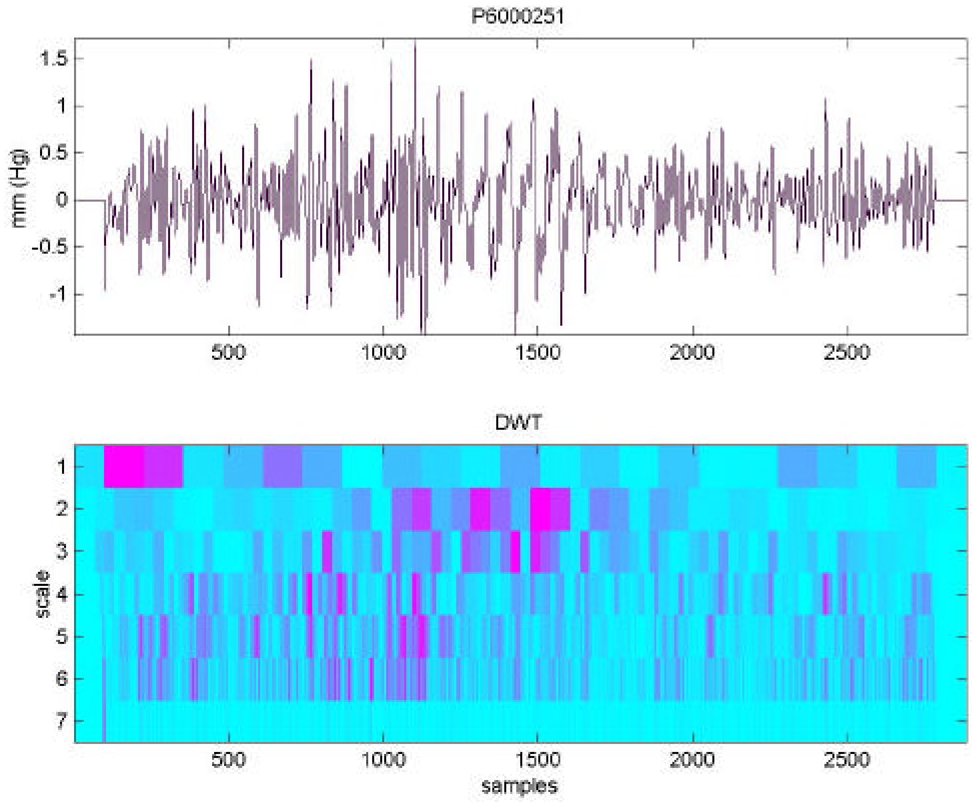

Figure 7: Oscillometric data generated using a BP Pump simulator to show a high level of motion artifact due to transport over a gravel road. The artifact nearly completely obscures the oscillometric signal in the upper trace. Note that the contributions from the gravel road artifact are most important in the shorter scales (larger scale value) at the bottom of the discrete wavelet transform. By comparison with the discrete wavelet transform in the previous signal, it can be seen that the physiological signal and the artifact are concentrated in different regions of the discrete wavelet transform.

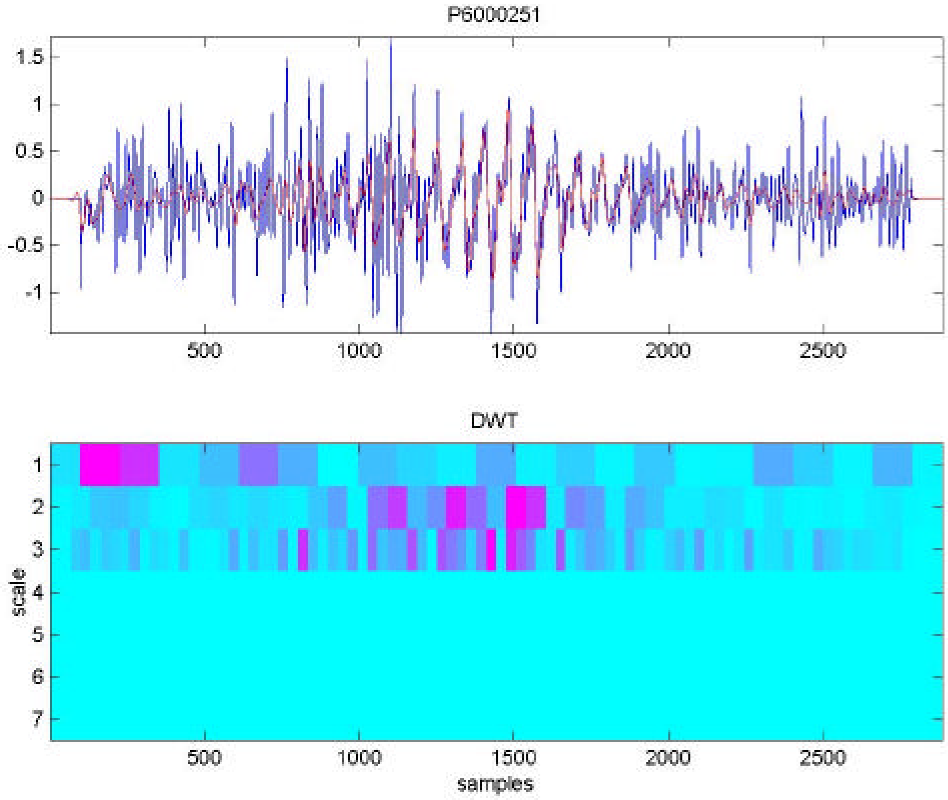

Figure 8: The upper plot shows the reconstructed oscillometric data from Figure 7 obtained by zeroing out the scales 4 through 7 in the discrete wavelet transform, as shown in the lower plot, and applying the inverse discrete wavelet transform. The original data is shown in the upper plot in blue and the reconstructed signal from the truncated wavelet transform is shown in red. The reconstructed signal has the artifact component much reduced and a usable oscillometric signal has been obtained despite the presence of severe artifact.

MIT Database test results for	QRS Se	QRS Pp	VEB Se	VEB Se	VEB FPR
Raw Data	99.66%	99.55%	92.39%	90.49%	0.74%
Data after compression/decompression	98.36%	99.58%	89.00%	86.52%	1.07%

Table 1: Results for selectivity (Se) and positive predictivity (Pp) in beat classification for ECG data before and after compression and reconstruction by the Associated-Hermite algorithm. Normal beats are denoted by "QRS" and ventricular ectopic beats by "VEB". The Associated-Hermite compression algorithm increases the fraction of mis-classified beats by only 0.33%, demonstrating the high level of fidelity of the compression and reconstruction for clinically significant features of the ECG waveform.