The discrete period domain transform – a new method for temporal signal filtering and reconstruction

Dennis E. Bahr[1] and James L. Reuss[2]

[1]HelionX, LLC, Middleton, WI, USA
[2]Prism Clinical Imaging, Elm Grove, WI, USA

Abstract

The Discrete Period Transform (DPT) is an algorithm that maps the time domain to the period domain much as the Discrete Fourier Transform (DFT) maps the time domain to the frequency domain. The DPT can be used to process physiologically generated signals that already have a limited frequency range such as ECG, blood pressure, heart rate, and blood oxygen saturation using pulse oximeters. The DPT algorithm is designed to handle the low pulsatile amplitudes, noise, motion artifact, and respiratory modulation found in these signals. These physiological signals all have an underlying pulsatile characteristic that is quasi-stationary and exhibit semi-ergodic behavior with non-periodic changes in period, amplitude, and wave shape. The DPT can be used to find the underlying pulse period and ensemble averaging can then be used to reconstruct a noise and artifact free signal. The DPT can separate noise and artifacts that share the same band of frequencies as the signal. The efficiency of a DPT incremental algorithm reduces the processing requirements facilitating use in low-power and portable devices.

Keywords: Pulse oximetry, fetal, PPG, ECG, blood pressure, biomedical signal processing, spectral analysis, frequency domain, period domain, ensemble averaging

1. Introduction

The object of this research was to develop an algorithm that could extract signals that were highly contaminated by motion artifact and noise. The research specifically targeted physiological signals produced by the heart including, but not limited to, ECG, blood pressure, heart rate, and blood oxygen saturation using pulse oximeters. These signals all have an underlying pulsatile characteristic that is quasi-stationary and exhibit semi-ergodic behavior with non-periodic changes in period, amplitude, and wave shape. Reconstruction of the underlying temporal waveforms, for further signal processing, is of great interest. The removal of the contaminating noise and artifacts, while not distorting the signal waveform, can be a challenge since often times the corruption is caused by signals that are in the same frequency band as the signal itself. In such cases, simple filtering of the data will not work. One popular way of extracting these data is to use another signal, which has a temporal relationship to the data, as the time frame to ensemble average the data. Ensemble averaging has been applied to photoplethysmographic (PPG) signals, usually employing an external cardiac "trigger" obtained from an ECG source (Palreddy 1997). This has also been attempted in fetal pulse oximetry, although obtaining a reliable fetal ECG signal generally requires use of an invasive fetal scalp electrode.

The specific aim of this research was to develop an algorithm that could be used to determine the underlying periodic frequencies in data corrupted with noise and motion artifact. Once found, this information could be used to ensemble average the data to recover the underlying temporal signal or signals. Therefore, the algorithm would need to

a) generate the period "trigger" from the raw data without the use of any ancillary signal;
b) produce an accurate replica of the data including the correct wave shape and period;
c) have a rapid enough response to track cardiac heart rate period changes in real time;
d) recover rapidly from signal interruptions or from excessive noise or motion artifact;
e) have sufficient computational speed so as not be the limiting factor in determining sampling rate;

f) require moderate storage requirements and be usable in low-power and portable devices.

2. Methods

2.1 Algorithm development and derivation

The Discrete Fourier Transform (DFT) operates over a large frequency range, but requires extremely large internal data arrays in order to obtain high resolution. However, doing so usually smears the resulting frequency data if the data is extracted from living biological systems because such signals may not be stationary over the time required to acquire the data. In order to avoid the data from being smeared an algorithm with high resolution and small internal data arrays was required. Since the specific aim of the research was to determine the period of the data, it followed that frequency should be replaced with period and instead of incrementing frequency as with the DFT, the period should be incremented. Thus was born the Discrete Period Transform (DPT) (Reuss *et al* 2002). Our method has subsequently been used by other researchers (Yan *et al*, 2005).

While the DFT has the frequency increase in a linear fashion, the DPT has the period increase in a linear fashion. This, in a sense, makes the two transforms complements of each other. As such, they both have properties that are useful in different circumstances. The DPT has bins that are separated by the data sample period. If a signal is sampled at a frequency of f_s and processed by the DPT, the spacing between the bins will be $1/f_s$. The abscissa of the DPT has units of increasing period although it can also be plotted in reverse order or converted and plotted as frequency. The DPT is best implemented as an incremental or sliding transform. In doing so, the orthogonal complex elements forming the basis functions can be easily generated and applied as continuous functions in time. The sliding implementation is considered a fast form of the DPT algorithm.

Traditional time domain techniques employed to process PPG signals include peak detection and fiducially point determination for cardiac period calculation, and peak-valley measurement for the pulsatile amplitude measurement used in blood oxygen saturation calculation. Recently, frequency domain analysis has been used for determination of the fundamental cardiac frequency and, to some extent, selective removal of noise components based upon frequency content (Rusch *et al* 1996). These algorithms work in conjunction with time domain techniques, rather than replacing them. It is proposed that the PPG signals be processed in the period domain, i.e., determining the relative contributions of different periods in the signal (Bahr *et al* 2002). The advantages of this method are improved resolution for low frequency biomedical signals, and compatibility with time domain algorithms. By deriving a reliable cardiac period estimate from period domain analysis, ensemble averaging may proceed without an external trigger source.

The Discrete Fourier Transform (DFT) and incremental or "sliding" DFT are fundamental algorithms (Rabiner *et al* 1975). For sampling frequency f_s the frequency "bin" k of the N-point DFT corresponds to frequency $f_k = k \cdot f_s / N$ Hz, and

$$X^i(k) = \sum_{n=0}^{N-1} x(i+n)\, e^{-j2\pi kn/N} \quad \text{k = 0, 1, …, N-1} \tag{1}$$

is the expression for the DFT of the k^{th} frequency "bin" for the sample sequence $x_i \ldots x_{i+N-1}$. At i+1, the "sliding" or incremental DFT is calculated as

$$X^{i+1}(k) = e^{j2\pi k/N} \left[X^i(k) + x(i) - x(i-N) \right] \tag{2}$$

To derive the Discrete Period Transform (DPT), let s = 1, 2, …, N samples be the range of periods possible in the sequence $x_i \ldots x_{i+N}$. Frequency f_k corresponds to period $s_k = 1/f_k = N / (k \cdot f_s)$ seconds = N / k samples, so $k = N / s_k$. Substituting into (1) and (2), for the period s,

$$T^i(s) = \sum_{n=0}^{N-1} x(i+n)\, e^{-j2\pi n/s} \qquad \text{s = 1, 2, …, N} \tag{3}$$

therefore the sliding period transform can be written as:

$$T^{i+1}(s) = e^{j2\pi/s}\left[T^{i}(s) + x(i) - x(i-N)\right] \tag{4}$$

The DPT calculates the period spectrum with a resolution of $1/f_s$. Over the relatively small frequency range of PPG signals (approx. 0.1-10Hz), this resolution is achieved with modest processing power and memory. No conversion from frequency is necessary for compatibility with time domain algorithms, an advantage where period measurements are interchanged between power spectrum and time domains. The value of N does not need to be a power of two even when using the fast form of the transform.

2.2 Algorithm implementation

The mathematics is straight forward, but the implementation is difficult because the basis functions are comprised of sets of complex functions that are incommensurate and that differ one from another by the sample period. The most basic form is a set of complex sinusoidal functions, an example of ten sinusoidal functions where the maximum period is 300 sample points is shown in Figure 1.

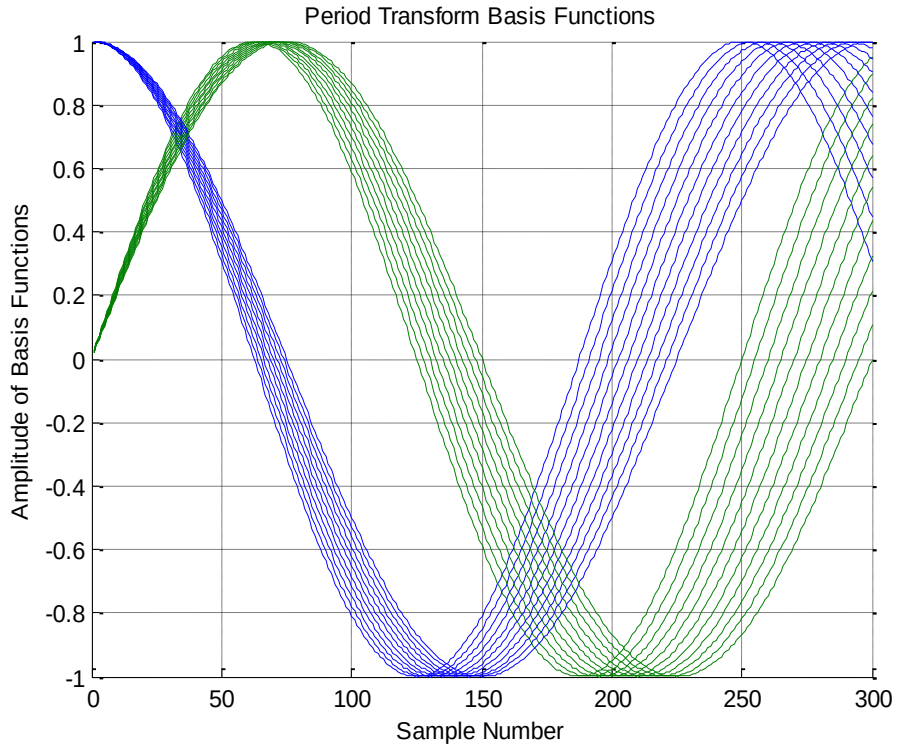

Figure 1. (color) Period Transform Sinusoidal Basis Functions with a maximum period of 300 samples.

Another implementation utilizes a set of basis functions that are the incremental phase angles of the complex sinusoids shown in Figure 1. They are easily derived using equation (5) where s represents the period.

$$e(s) = e^{j2\pi/s} \qquad \text{s} = 20, 21, \ldots, 300 \tag{5}$$

Figure 2 shows the amplitudes of the basis functions for periods from 20 sample points to 300 sample points. Because phase is incremented, this implementation also uses a different technique to do the correlation than the functions shown in Figure 1.

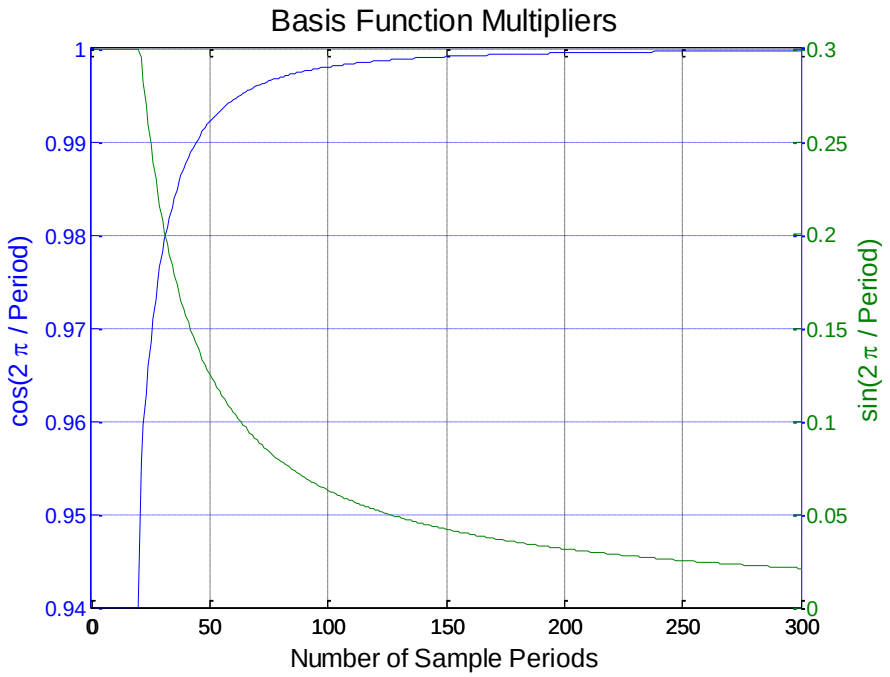

Figure 2. (color) Period Transform Basis Functions showing the values of the multipliers for increasing period.

The simplest implementation is to combine equation (3) above and either one of the two sets of basis functions. Using this method requires that all of the data points to be transformed are sampled and then processed in batch mode. If the data is lengthy, the data could be broken up into blocks and the blocks transformed one at a time. Calculating the basis functions in the main correlation loop forces the implementation to be somewhat slow but easy to develop and understand. The basis functions could be calculated before the correlation loop and therefore only need to be done once, thus speeding up the data processing. Using either of these two examples produces exactly the same result, as it should. See Figure 3.

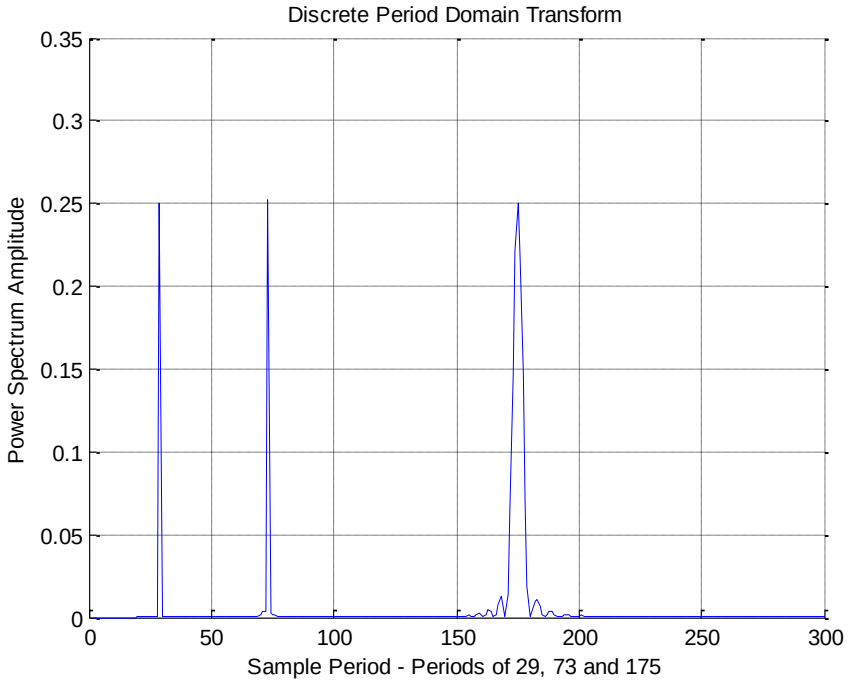

Figure 3. (color) Transform Power Spectrum from three sinusoidal signals that are incommensurate relative to each other.

The plot shows only half of the resulting spectrum. A second identical set exists as it does with the DFT. Since the input data are sinusoidal waveforms with unity amplitudes, the amplitudes of each are 0.5, because of the duplicity, and their respective powers are 0.25. Since the abscissa is plotted as period, one can see the great resolution that this transform offers. The results using this transform will display the periods of the periodic signals buried in the data and suppress the noise and motion artifact.

Implementing the algorithm as an incremental or sliding transform, as shown in equation (4), is considerably more difficult. The sliding transform implementation has a signal flow diagram of a comb filter followed by a resonator, as does the implementation of the sliding form of the DFT (Jacobsen *et al* 2003, Jacobsen *et al* 2004). The comb filter is implemented by a fixed length buffer known as the recurrence buffer and the resonator by the correlation with the complex basis functions.

Since the components of the DPT complex basis functions are not harmonically related, the end points of these functions do not always form continuous functions in the sample space as the DFT does. This produces the small side lobes shown in Figure 3. However, implementing the DPT as a sliding transform wraps the basis functions allowing the component basis functions to become continuous in nature as shown in Figure 4. As the data and basis functions temporally slide by together, and the correlation is computed, basis function and data continuity are maintained. This is the algorithm that is used by the instrumentation described below.

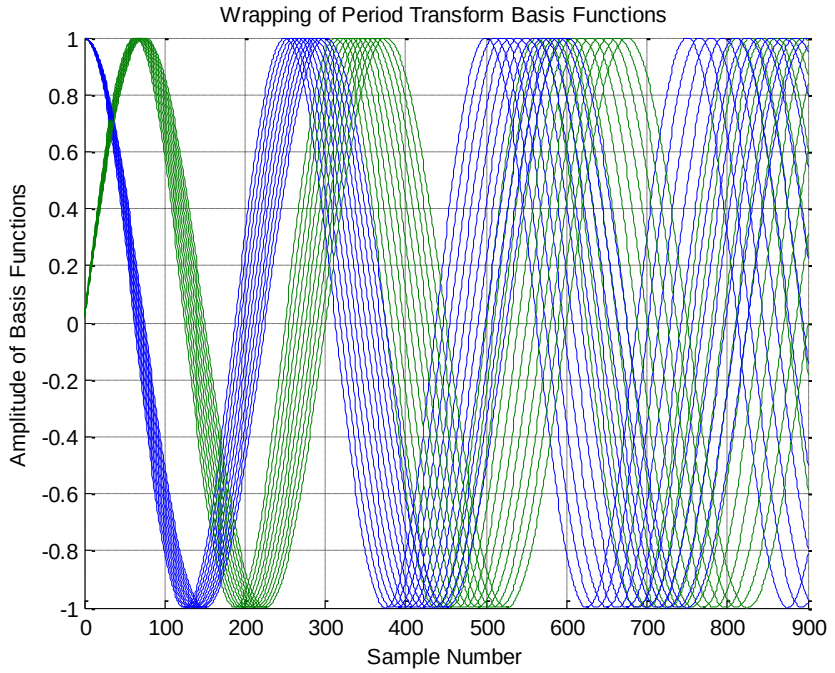

Figure 4. (color) Complex sinusoidal basis functions showing wrapping as the number of sample points is increased.

The most efficient sliding transform uses the basis functions shown in Figure 2. Assume that the section where the correlation is done is fixed in length and that the new data replaces the old data that was generated N samples earlier and saved in the recurrence buffer. To complete the sliding transform picture the updated contents of the recurrence buffer, that are the length of the period bin being processed, are rotated by the basis function for that period. The length of this buffer determines the overall resolution and once enough data has entered the process to fill this buffer the transform results reach a stable limit. The data is scaled by the ratio of the recurrence buffer size and the maximum period times the period bin k.

2.3 Algorithm application

Pulse oximetry (Jubran 1999) is the non-invasive measurement of arterial oxygen saturation (SpO_2) based upon the relative absorbance of multiple light wavelengths by different species of hemoglobin. Many advances

have been made since its introduction in the 1970's, primarily aimed at reducing the impact of low perfusion and motion on availability and reliability. The application of pulse oximetry has also been extended into new areas such as intrapartum fetal monitoring.

Due to placental versus respiratory supply, the fetus functions at a much broader and lower oxygen saturation range, typically SpO_2 = 40%-75%, with clinically significant desaturations occurring below 30% (Dildy *et al* 1996). For intrapartum monitoring of oxygen saturation via fetal pulse oximetry (FPO), the sensor must either be placed on the presenting part of the fetus, or through the birth canal onto the fetus torso. The design and calibration of these devices is substantially more difficult, since tissue conditions at the monitoring site, such as blood volume, are more critical at low saturations (Reuss *et al* 2001).

The few commercially available fetal pulse oximeters utilize the reflectance, or backscattering, mode of operation, as opposed to the transmittance mode of most adult sensors. The fetal sensor must be placed trans-cervically upon the fetal body in utero (Luttkus *et al* 2001). The pulse amplitude of the photoplethysmographic (PPG) signals is small compared to adults or neonates. Although the intrauterine placement generally eliminates ambient light as a source of interference, the intimate proximity of the mother's body introduces a new problem, possibility of maternal modulation of the fetal signals. Maternally vascularized tissue near the fetal sensor can result in a strong modulation of the fetal PPG at the maternal pulse rate, potentially interfering with not only pulse rate detection but SpO_2 calculation as well.

Utilizing the OB Scientific™ OBS-500 Fetal Pulse Oximeter, a fetal signal database was collected to aid in development of pulse oximetry algorithms. The OBS-900 Fetal Oxygen Sensor was placed through the birth canal onto the fetal torso. Red, infrared, and dark signals from the sensor, digitized at a rate of 120 samples/second and 21-bit resolution, were transmitted from the oximeter to a personal-computer-based Fetal Oximetry Platform (FOP). This software incorporates the oximeters algorithms as well as a user interface for graphical and textual display and analysis of the oximeter operation. The FOP permits a record of unprocessed signal data to be re-run with alterations to the algorithms and pulse-to-pulse comparison of the results.

3. Results

Two examples that show the utility of the DPT include extracting the PPG signals from corrupted data where the DPT signals are masked by frequencies in the same band as the signals and the extraction of the PPG signals from a fetus where the maternal PPG signal dominates the data.

3.1. Signal extraction from noisy data

Figure 5 illustrates an example of weak signal data (red and infrared (IR)) in the time domain (a-c), and the corresponding period domain spectrum (c). The peak in (c) is the cardiac period. The fetal heart rate is 137 beats per minute.

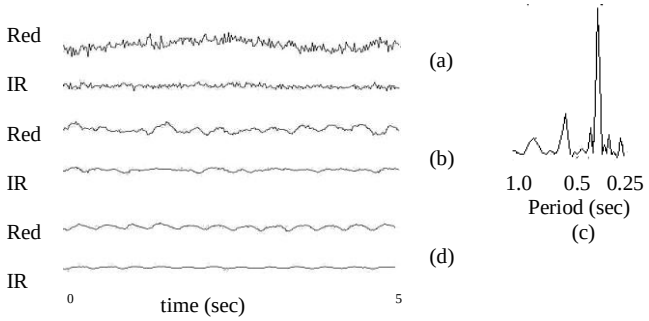

Figure 5. Weak photoplethysmographic signals: (a) unprocessed, (b) band-passed, (c) period domain transform, (d) ensemble averaged

Figure 5(a) is the unfiltered signal data from the sensor. Simple band-pass filtering results in the pulsatile signals shown in 5(b). Reliable determination of the pulse rate from (b) would be difficult. The cardiac period is apparent as the spectral peak in (c). The cardiac period estimates of (c) can be used to direct ensemble averaging to extract the fetal signal components, as shown in (d). The period domain spectrum yields a useful cardiac period for rate calculation that can be tracked through all but extreme, prolonged noise.

3.2 Signal extraction from data with maternal modulation

A case of significant maternal modulation is illustrated in Figure 6, showing the time domain (a-b) and the period domain (c). The right peak in (c) is the cardiac period of the fetal heart rate (126 beats/minute), whereas the left peak is the maternal cardiac period (97 beats/minute heart rate).

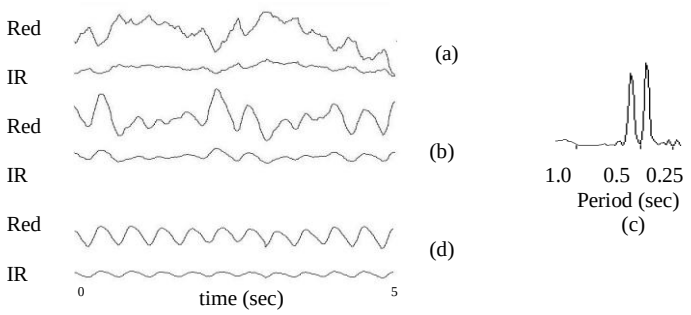

Figure 6. Maternal modulation in PPG signals: (a) unprocessed, (b) band-passed, (c) period domain transform, (d) ensemble averaged

Again, simple band-pass filtering to obtain the pulsatile portion of each signal is shown in Figure 6(b). Errors in pulse rate and oxygen saturation calculations could result from processing the signals in (b). Ensemble averaging utilizing the cardiac period derived from period domain analysis extracts the fetal signal components more effectively, as illustrated in (d). These signals will produce accurate oxygen saturation calculations.

4. Discussion

The PPG period domain spectrum obtained with the incremental DPT algorithm tracks relatively fast heart rate changes (such as heart rate accelerations and decelerations), but may not produce useful results in the presence of arrhythmias. Because it is used in conjunction with time domain pulse detection, the period domain algorithm may be disabled automatically in such circumstances.

In the case of fetal pulse oximetry, highly irregular rhythms are relatively uncommon. When bigeminy is present with regularity of pulse spacing, period tracking may "lock on" and track the normal (hemodynamically strong) pulses at half the actual pulse rate. Ensemble averaging at the halved rate will still be effective, if the temporal relationship of normal and PVC beats is consistent. It is important that the system be able to recover from signal interruptions or excessive noise and motion artifact. There are many ways to accomplish this including restarting the system or freezing the recurrence buffer until the interference has subsided.

5. Conclusion

Period domain analysis utilizing an incremental DPT algorithm is an effective and efficient way to process periodic biomedical signals for spectral content. It provides the capabilities of frequency domain analysis, with certain advantages in implementation. Processing of photoplethysmographic fetal pulse oximeter signals with period domain analysis improves pulse rate availability and accuracy, and permits removal of interference by techniques such as ensemble averaging without an external noise reference. Work to be done in the future will include the ability to modify the length of the recurrence buffer using some measure of the signal to noise ratio.