AN ADAPTIVE SINGLE FREQUENCY PHASE VOCODER FOR LOW-POWER HEART RATE DETECTION

Development of a Fast and Low-power Heart Rate Estimation Algorithm for Mobile Phone Applications

Walter Karlen, Guy Dumont
Electrical and Computer Engineering in Medicine (ECEM), The University of British Columbia, Vancouver, Canada

Chris Petersen, Jennifer Gow, J. Mark Ansermino
Pediatric Anesthesia Research Team (PART), BC Children's Hospital, Vancouver, Canada

Keywords: Pulse detection, Heart rate estimation, Mobile phones, Embedded systems, Photoplethysmography, Vocoder.

Abstract: Mobile phones can be used as a platform for clinical decision making in resource-poor and remote areas. Their limited battery and computational resources, however, demand efficient and low-power algorithms. We present a new algorithm for the fast and economical estimation of heart rate (HR) from the photoplethysmogram (PPG) recorded with a pulse oximeter connected to a mobile phone. The new method estimates the HR frequency by adaptively modeling the PPG wave with a sine function using a modified phase vocoder. The obtained wave is also used as an envelope for the detection of peaks in the PPG signal. HR is either computed using the vocoder center frequency or the peak intervals in a histogram. Experiments on a mobile device show comparable speed performance with other time domain algorithms. Preliminary tests show that the HR computed from the vocoder center frequency is robust to Noise in the PPG. The instantaneous HR calculated with the vocoder peak detection method was more sensitive to short-term HR variations than the vocoder center frequency method. These results point to further developments using a combination of both HR estimation methods that will enable the robust implementation of adaptive phase vocoders into mobile phone applications.

1 INTRODUCTION

Mobile health technology is a rapidly advancing field that holds great promise for improving medical services and changing the way that health care is delivered. A common theme in this area is the use of general purpose consumer devices, in particular smart phones. An increasing number of health care applications use these mobile platforms to interface directly to physiological sensors, such as heart rate (HR) detectors. While this reduces or eliminates the cost of custom embedded hardware, features such as an increased noise level, limited battery and computational resources, and the requirement for on-line processing challenge the accurate, real-time detection of physiological parameters, such as HR peaks.

In this paper, we consider the extraction of heart rate (HR) from a photoplethysmography (PPG) signal that originates from an oximeter sensor interfaced to a smart phone. We propose a novel approach to compute HR from the PPG signal based on a dynamically adapted, single frequency phase vocoder algorithm. This algorithm is intended to form a core engine for more complex mobile signal analyzers within a fully functional low-cost, smart phone-based pulse oximeter.

1.1 Related Work

The traditional methods used for peak detection in the ECG signal for HR estimation, which have a long history in biomedical signal processing, can be applied to PPG signals. HR estimation algorithms operate either in the time or the frequency domain. Common time domain algorithms include linear and non-linear filters, artificial neural networks, genetic algorithms, filter banks, and heuristic methods based on nonlinear transforms. Frequency domain algorithms include various wavelet and Fourier transforms. However, not all of these methods are suitable for on-line computa-

Karlen W., Dumont G., Petersen C., Gow J. and Mark Ansermino J..
AN ADAPTIVE SINGLE FREQUENCY PHASE VOCODER FOR LOW-POWER HEART RATE DETECTION - Development of a Fast and Low-power Heart Rate Estimation Algorithm for Mobile Phone Applications.
DOI: 10.5220/0003134200300035
In *Proceedings of the International Conference on Biomedical Electronics and Devices* (BIODEVICES-2011), pages 30-35
ISBN: 978-989-8425-37-9

tion (Kohler et al., 2002). In particular, algorithms must be computationally efficient for mobile applications, where power and computational resources are limited. Algorithms are commonly trimmed in order to achieve computational efficiency (Kohler et al., 2002). However, this comes at the expense of accuracy and performance. Our aim was to design an accurate HR estimation algorithm for PPG that requires low processing power so as to be suitable for battery-powered mobile applications.

2 ALGORITHM DEVELOPMENT

Since cardiac signals are quasi periodic, a time-frequency transformation algorithm could be appropriate to extract HR. However, Fourier or wavelet based transformations on time series require segmentation of the signal, which is not always practical in an on-line, low-cost system. Their relatively high processing power and memory usage requirements add to this unsuitability. Instead, we propose to operate in the time domain, using a methodology inspired by the phase vocoder originally developed for the compression of voice signals in telecommunications (Flanagan et al., 1965). In this method, a vocoder models the input signal with one or multiple sinusoidal waves that vary in time. The parameters that have to be estimated are the time varying amplitude and frequency of each sine wave that comprises the original signal. The phase vocoder can be seen as a filter bank consisting of a series of band-pass filters with successive center frequencies (Dolson, 1986).

In our case, we are solely interested in the dominant frequency in the signal (the heart beat). It is, therefore, sufficient to represent the input signal by only one sinusoidal wave that varies over time. The filter-bank with a distinguished set of frequencies is replaced with a single band-pass filter whose center frequency is adapted over time. This can be achieved by scanning the incoming signal and computing the difference in phase. The obtained frequency parameter, where the system eventually locks-in, can be seen as a filtered frequency of the incoming signal that would correspond to an averaged HR. The method can also find the location in time of each heart beat by using the output sine wave as an envelope to locate maximal peaks in the input waveform.

2.1 Algorithm Description

The raw PPG signal is high-pass filtered to remove the baseline using a second-order Butterworth filter with a cut-off frequency at 0.5 Hz (Figure 1-1). The filtered signal is then routed into two parallel streams. The signal is multiplied by a sine wave in one stream and by a cosine wave in the other (Figure 1-2). Cosine and sine waves have the same unitary amplitude and frequency w_v. The frequency w_v is set to the estimated vocoder frequency of the previous iteration (Figure 1-8). The two parallel streams are, therefore, identical with the exception of the $\pi/2$ shifted phase of the multiplying waveform. This step creates a new signal that is composed of two periodic signals with shifted frequencies of $\pm w_v$ as follows:

$$\begin{aligned}&\cos(w_v * t) * \cos(w_s * t) = \\ &\cos((w_s - w_v) * t) + \cos((w_s + w_v) * t),\end{aligned} \tag{1}$$

where w_s is the frequency of the incoming wave at timestep t. Next, each of the two streams is fed to a moving average low-pass filter (Figure 1-3). The application of this heterodyning step has two effects. First, input frequencies in proximity of the vocoder frequency w_v are shifted down close to DC and are allowed to pass the filter. All other frequency components will also be shifted but they will not go through the low-pass filter. Secondly, the heterodyning provides a way to compute the time-varying amplitude and frequency of the resulting signal in the next step. The two filtered waveforms are subsequently transformed from Cartesian to Polar coordinates to obtain a single amplitude r_v and phase θ_v (Figure 1-4). The amplitude is calculated as the square root of the sum of the squares of the two heterodyned signals as follows:

$$r_v = 4 * \sqrt{y_{sin}^2 * y_{cos}^2}, \tag{2}$$

where y_{sin} and y_{cos} are the heterodyned signals for the sine and the cosine streams, respectively. Similarly, θ_v at each point in time t is the angle whose tangent is the ratio of the vertical to the horizontal position as follows:

$$\theta_v = \arctan(\frac{y_{sin}}{y_{cos}}). \tag{3}$$

The phase is subsequently unwrapped. The real time-varying frequency of the original wave is then estimated by computing the difference between the actual and previous phase (Figure 1-5), and subtracting it from the current center frequency (Figure 1-6) as follows:

$$\hat{w}_s = w_v = w_v^{(t-1)} - \Delta\theta, \tag{4}$$

where $\Delta\theta = \theta_v^{(t-1)} - \theta_v$. The amplitude and the newly computed center frequency are used to compute the vocoder output y_{out} (Figure 1-7) as follows:

$$y_{out} = r_v * \sin(2 * \pi * w_v * t). \tag{5}$$

The output signal frequency w_v is also used to estimate the instantaneous HR. Algorithm 1 shows a possible implementation of the adaptive single-frequency phase vocoder in pseudo C code.

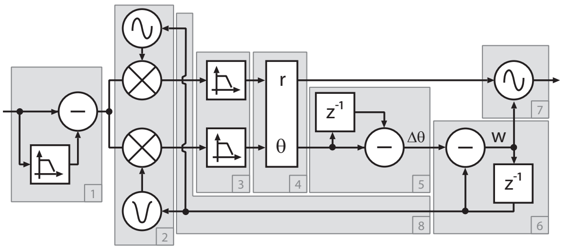

Figure 1: Operation of the adaptive single-frequency phase vocoder: (1) high-pass filter to remove DC value, (2) heterodyning the input with both a sine and a cosine wave in parallel, and a (3) low-pass filter, (4) converting the two signals from rectangular to polar coordinates and unwrapping the angular-position values, (5) subtracting successive unwrapped angular-position values, (6) subtracting the phase-difference from the previous center frequency to obtain the new center frequency, (7) generating the vocoder output wave for peak detection, and (8) feeding the center frequency back into the heterodyne function.

Algorithm 1: Adaptive Phase Vocoder.

initialization
while TRUE **do**
  $val = input$
  $val \leftarrow highpass(val)$
  $y_{sin} \leftarrow lowpass(val * sin(freq_{-1} * t))$;
  $y_{cos} \leftarrow lowpass(val * cos(freq_{-1} * t))$;
  $phase \leftarrow unwrap(atan2(y_{sin}, y_{cos}), phase_{-1})$ ▷ unwrap the phase
  $ampl \leftarrow 4 * \sqrt{y_{sin} * y_{sin} + y_{cos} * y_{cos}}$
  $freq \leftarrow freq_{-1} - (phase_{-1} - phase)$ ▷ adjust frequency to lock onto the waveform
  if $freq < 0.00001$ **then** ▷ clamp to avoid negative frequency values
    $freq \leftarrow 0.00001$
  end if
end while

To detect the location of the HR peaks, the monophone sinusoidal vocoder output y_{out} is used as an indicator of the approximate location of the peaks in the actual PPG waveform. The simple form of the vocoder waveform reduces the problem of detecting peaks in the noisy PPG signal to a straight-forward search for the maximum amplitude within the temporal interval of a positive half cycle of the vocoder output signal. The time elapsed from the previously detected peaks is binned into a moving histogram of heart rate intervals. This serves to eliminate artifacts caused by motion and other spurious effects, with the largest histogram bin giving a direct representation of the instantaneous HR. The sensitivity of this peak detection method can be adjusted by offsetting the sinusoidal vocoder signal; a positive or negative offset adjustment will result in a wider or narrower positive half cycle interval, respectively.

2.2 Algorithm Validation

In order to test the algorithm, following institutional review board approval, data was gathered from nine children (1-7 years old, 14 kg $\pm$ 6.2) who underwent general anesthesia. The recordings obtained included ECG, capnometry, and PPG signals. All signals were recorded with Collect S/5 software (Datex-Ohmeda, Finland) using a sampling frequency of 300Hz. An 8-minute segment of reliable recording of spontaneous breathing was selected from each case. These segments are available from the on-line database CapnoBase.org (Karlen et al., 2010).

The ECG waveform was used as the reference recording for computing HR. A technician independently validated the reference measurement using the CapnoBase Signal Evaluation Tool (Karlen et al., 2010). The mean HR of a sliding window was calculated. The window size was set to 20 s; a size that corresponds to that commonly used in commercial pulse oximeters.

The performance of the adaptive algorithm was assessed using the normalized root mean square (NRMS) error (%) of the calculated HR from the vocoder output frequency. The NRMS error corresponds to the square root of the sum of the squares of the differences between the test and reference HR measurements, divided by the sum of the reference measurements:

$$NRMS\ error = \frac{\sqrt{\sum_{i=1}^{n}(x_i^{ref} - x_i^{alg})^2 * n}}{\sum_{i=1}^{n} x_i^{ref}}. \quad (6)$$

A measurement error was calculated for each test measurement by comparing it with the reference measurement that was nearest to it in time. The first 50s were used to initialize the high-pass filters and the vocoder sliding window, and so were not analyzed.

To evaluate the computational load of the algorithm executed on a mobile device, we built a prototype device consisting of a PureLight medium soft sensor connected to an Xpod OEM module (Nonin, Plymouth, USA) (Figure 2). The module was connected to a 2^{nd} generation iPod Touch (Apple, Cupertino, USA). The iPod Touch displayed the PPG waveform and recorded the continuous data stream. The PPG was recorded with a 16bit resolution at a sampling rate of 75 Hz. The algorithm was implemented in C and embedded into the iOS application software, called *Phone Oximeter*[1], to process the raw PPG signal in real-time. The computational load was mea-

[1]http://www.phoneoximeter.org

Figure 2: Phone Oximeter (Mobile phone pulse oximeter) application. 1) Soft finger probe; 2) Nonin Xpod OEM pulse oximeter module; 3) raw PPG waveform display on iPod Touch; 4) SpO2 value computed by the Xpod module; 5) HR computed by the adaptive vocoder algorithm.

sured in microseconds using interrupts during program execution.

For comparison, we implemented into the iOS application two other HR estimation methods. The first used the Pan-Tompkins algorithm (Pan and Tompkins, 1985). This time domain based method makes use of a cascade of band-pass filter, integrator, squaring, and differentiators. An adaptive threshold was applied to detect the heart beat pulses. The second method calculated the power spectral density with a Fast Fourier Transform (FFT) algorithm from (Press et al., 1992) that was implemented in the C library from (Goldberger et al., 2000). A 50s window size was selected, border effects were reduced using a Hamming windowing function, and the maximum power band was chosen to estimate the HR. The FFT was recalculated every second by shifting the window by 75 samples.

3 RESULTS

The parameters used in the adaptive vocoder algorithm are shown in Table 1.

With an average NRMS error of 3.33% (Table 2), the HR calculated directly from the vocoder output frequency was accurate. However, two cases exhibited higher NRMS error than the others (00301

Table 1: Algorithm configuration parameters.

Parameter	**Value**
heterodyne low-pass cut-off frequency	0.04 Hz
initial vocoder frequency $w_v^{(0)}$	50 Hz
initial vocoder amplitude $r_v^{(0)}$	1
number of histogram bins	32

Table 2: Normalized root mean square (NRMS) error between the reference HR obtained from ECG and the HR estimated by the vocoder frequency or vocoder peak detection method respectively. The last column is the mean ECG HR used for calculating the NRMS error.

Case	**HR NRMS error** [%]		**mean HR** [bpm]
	vocoder freq	vocoder peak	ECG
00091	3.4979	1.7496	92.5807
00151	0.5915	0.6440	108.454
00181	0.5170	1.2345	127.149
00231	2.3822	1.3859	92.3054
00281	2.4448	2.8312	66.4234
00301	6.8719	4.3432	107.406
00321	2.3679	2.8300	77.2644
00351	2.9794	1.4727	102.547
00381	8.3220	1.2452	108.462
Average	**3.3305**	**1.9707**	**98.0658**

and 00381). Closer inspection of these cases illustrates that reduced performance of the vocoder frequency HR is mainly due to its low responsiveness to a rapidly changing HR of more than 10bpm (Figure 3 and 4).

With an average NRMS error of 1.97% (Table 2), the HR calculated from the vocoder peak detection using histograms was more accurate than the vocoder frequency estimations. The disagreement between ECG reference HR and HR obtained from vocoder peak detection resulted from a delay in the HR estimation of the latter; a 5s to 10s delay observed throughout these cases can be attributed to the size of the histogram for HR calculation. Cases 00281, 00301, and 00321 showed greater decreases in performance than the mean of the cases. In these cases, the poor quality of the PPG signal prevented the vocoder peak detection algorithm from accurately estimating HR (Figure 3). The vocoder peak detection was able to track fast variations and compensate for the low responsiveness of the direct vocoder frequency calculation.

The adaptive vocoder algorithm including peak detection required an average of 47.33 us to process a new value, which was slower than the Pan-Tompkins algorithm (Table 3). Both were significantly faster than the inter-sampling distance of 13.3

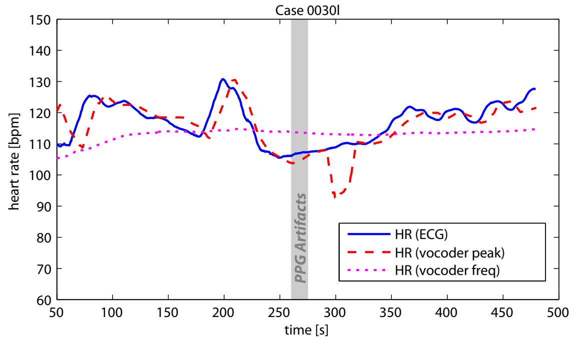

Figure 3: HR agreement comparison for Case 0030l. The HR calculated by the vocoder peak method is degraded after a time offset (at 300s) when the raw PPG signal was corrupted with noise over a certain period of time (gray bar). The HR estimated by the vocoder frequency is smoothed out and does not respond quickly to large variations in HR.

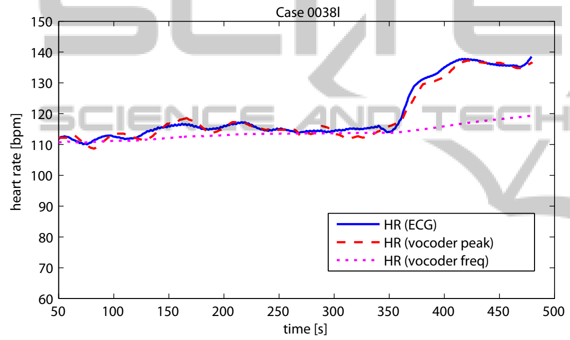

Figure 4: HR agreement comparison for Case 0038l. The HR estimated by the vocoder frequency does not respond rapidly to a large permanent change in HR after 350s. The HR calculated by the vocoder peak method is accurately following the ECG HR trend.

ms. The algorithms can be computed in real-time on the phone while sufficient processing power is available for other computing tasks. The FFT HR estimation method was about 40 times slower than the vocoder algorithm (Table 3).

4 DISCUSSION

We can conclude from these preliminary tests that the major design criteria of accuracy and on-line capabilities for a mobile phone HR estimation algorithm were met. With an average accuracy of 1.97% NRMS, the HR prediction is acceptable for a mobile device. The speed tests show that the proposed algorithm achieves this within the range of traditional filtering methods and faster than frequency domain-based approaches.

The adaptive vocoder frequency output filtered short variations in HR and PPG artifacts effectively.

Table 3: Algorithm execution time per sample on the iPod Touch calculated over 16 measurements of one minute.

	Vocoder freq	Vocoder peak	Pan-Tompkins	FFT
	[us]	[us]	[us]	[us]
Mean	41.25	47.33	28.84	1937
SD	0.9	0.8	0.45	4.06

However, the algorithm was not able to detect large, rapid changes in HR and lock into the new frequency in a reasonable amount of time. The vocoder responsiveness could be improved by tuning the cut-off frequency of the low-pass filter shown in Figure 1-3. On the other hand, the vocoder peak detection method was able to rapidly track changes in HR. However, this attribute made it more vulnerable to errors when the PPG was corrupted with noise.

It is, therefore, evident that the advantages and drawbacks of the two approaches are complementary. A logical step would be the combined use of these vocoder based HR extraction methods. A system that detects short-term variations with the peak detection and long-term trends with the vocoder frequency could provide increased diagnostic information.

A desired improvement to the vocoder peak HR estimation algorithm is to reduce the delay that was introduced by the histogram detection algorithm. We also plan to increase its robustness to poor quality PPG signals by adding a signal quality index to the algorithm. Active detection of poor signal will prevent the erroneous HR output and increase overall estimation performance.

Although we compared processing speed of different HR estimation algorithms when embedded into a mobile device, we did not compare the estimation performance between these methods in this study. We plan to proceed with a performance comparison as soon we have implemented the suggested improvements in a prototype device. This will allow us to conduct clinical studies for validation on an increased number of cases.

5 CONCLUSIONS

A novel approach to compute HR based on a dynamically adapted, single frequency phase vocoder algorithm was proposed. Initial experiments showed good accuracy and low computational costs for the suggested approach. The robustness of the algorithm towards noise in the plethysmogram waveform was another strength. This makes the algorithm suitable to process on-line PPG signals that are recorded from an oximeter sensor interfaced to a smart phone. We intend to use the algorithm as a core engine for more complex mobile signal analyzers (i.e. for the estimation of heart rate variability or respiratory rate) within a smart phone based pulse oximeter. The suggested algorithm has the potential to be applied to other periodic signals whose frequency range has to be determined in real-time.

REFERENCES

Dolson, M. (1986). The Phase Vocoder: A Tutorial. *Computer Music Journal*, 10(4):14–27.

Flanagan, J. L., Meinhart, D. I. S., Golden, R. M., and Sondhi, M. M. (1965). Phase Vocoder. *The Journal of the Acoustical Society of America*, 38(5):939–940.

Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 101(23):e215—-e220.

Karlen, W., Turner, M., Cooke, E., Dumont, G., and Ansermino, J. M. (2010). CapnoBase: Signal database and tools to collect, share and annotate respiratory signals. In *Annual Meeting of the Society for Technology in Anesthesia (STA)*, page 25, West Palm Beach.

Kohler, B., Hennig, C., and Orglmeister, R. (2002). The principles of software QRS detection. *IEEE Engineering in Medicine and Biology Magazine*, 21(1):42–57.

Pan, J. and Tompkins, W. J. (1985). A real-time QRS detection algorithm. *IEEE transactions on bio-medical engineering*, 32(3):230–6.

Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. (1992). *Numerical recipes in C: the art of scientific computing.* Cambridge University Press, New York, 2nd edition.