Photoplethysmogram Processing Using an Adaptive Single Frequency Phase Vocoder Algorithm

Walter Karlen[1], Chris Petersen[2], Jennifer Gow[2],
J. Mark Ansermino[2], and Guy Dumont[1]

[1] Electrical and Computer Engineering in Medicine (ECEM)
The University of British Columbia, Vancouver, Canada
{walterk,guyd}@ece.ubc.ca
[2] Pediatric Anesthesia Research Team (PART)
British Columbia Children's Hospital, Vancouver, Canada
{CPetersen,JGow}@cw.bc.ca,
anserminos@yahoo.ca
http://ecem.ece.ubc.ca,
http://part.cfri.ca

Abstract. We have previously designed a pulse oximeter connected to a mobile phone, called the *Phone Oximeter*, for clinical decision making based on photoplethysmography. The limited battery and computational resources demand efficient and low-power algorithms for the *Phone Oximeter* to be effective in resource-poor and remote areas. We present two new algorithms for the fast and economical estimation of heart rate (HR) from the photoplethysmogram (PPG). One method estimates the HR frequency by adaptively modeling the PPG wave with a sine function using a modified phase vocoder. The other method uses the obtained wave as an envelope for the detection of peaks in the PPG signal. HR is computed using the vocoder center frequency or the peak intervals in a histogram, respectively. PPG data obtained from 42 subjects were processed with the vocoder algorithms and, for comparison, with two traditional methods that use filtering algorithms (Pan-Tompkins) and frequency domain transformations (Fast-Fourier Transform). We compared HR estimates obtained from these four methods to the reference HR obtained from a electrocardiogram. The two vocoder methods performed at least as well as the two traditional methods in terms of normalized root mean sqare error and robustness towards artifacts. Experiments on a mobile device prototype showed comparable speed performance of the vocoder algorithms with the Pan-Tompkins algorithm while the frequency domain approach was nearly two orders of magnitude slower. These results point to further developments using a combination of both vocoder HR estimation methods that will enable the robust implementation of adaptive phase vocoders into mobile device health applications.

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

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

A. Fred, J. Filipe, and H. Gamboa (Eds.): BIOSTEC 2011, CCIS 273, pp. 31–42, 2012.

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. This reduces or eliminates the cost of custom embedded hardware. However, features such as 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 mobile device. 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, mobile phone-based pulse oximeter.

1.1 Related Work

The traditional methods used for peak detection in the electrocardiogram (ECG) signal for HR estimation, which have a long history in biomedical signal processing, can be applied to PPG signals. These 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 computation [6]. In particular, algorithms must be computationally efficient for mobile applications, where battery power and computational resources are limited. Algorithms are commonly trimmed in order to achieve computational efficiency [6]. 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 and so would 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 [2]. 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 comprise 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 [1].

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:

$$\hat{w_s} = w_v = w_v^{(t_{-1})} - \Delta\theta, \tag{4}$$

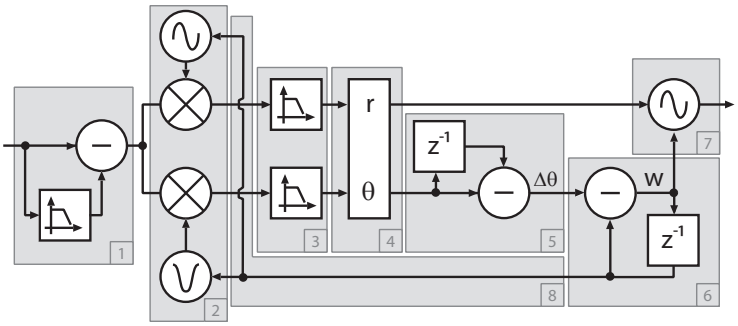

Fig. 1. Operation of the adaptive single-frequency phase vocoder: (1) high-pass filter to remove DC value, (2) heterodyning the input with both a sine and a cosine wave in parallel, (3) a low-pass filter, (4) converting the two signals from rectangular to polar coordinates and unwrapping the angular-position 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.

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.

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 wave
  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

Table 1. Algorithm configuration parameters

Parameter	**Value**
Heterodyne low-pass cut-off frequency	0.039 Hz
Initial vocoder frequency $w_v^{(0)}$	50 Hz
Initial vocoder amplitude $r_v^{(0)}$	1
Number of histogram bins	32

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 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 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, data was gathered from 29 children (4.8 years $\pm$ 5.4, 18.5 kg $\pm$ 23.4) and 13 adults (46.3 years $\pm$ 9.0, 73.5 kg $\pm$ 24.2) who underwent general anesthesia, following institutional review board approval. 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 300 Hz. An 8-min segment of reliable recording of spontaneous or controlled breathing was selected from each case. The mean HR for these segments ranged from 53 to 142 bpm. The segments are available for download from the on-line database CapnoBase.org [5].

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 [5]. The mean ECG 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. In addition, the technician labeled the beginning and end of all artifacts in the PPG waveform.

The vocoder algorithms were written in C programming language and compiled to be embedded in a Matlab (Mathworks, Natick, USA) development framework. The adaptive vocoder parameters were tuned using the in-vivo dataset from Capnobase that was different from our test data set. The configuration parameters are shown in Table 1.

To compare the performance of the two vocoder algorithms, we implemented two other HR estimation methods; one that uses traditional filtering methods and another that uses a frequency domain-based approach. The first used the Pan-Tompkins algorithm [7]. 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. Again, mean HR using a sliding window of 20 s duration was calculated. The second method calculated the power spectral density with a Fast Fourier Transform

Table 2. Normalized root mean square (NRMS) error between the reference HR, obtained from ECG, and the HR estimated by the vocoder frequency, vocoder peak detection, Pan-Tompkins, or FFT method, respectively

Case	**HR NRMS error** [%]				**mean HR** [bpm]
	vocoder peak	vocoder freq	Pan-Tompkins	FFT	ECG
0009l	1.97	3.22	19.99	1.39	102.46
0015l	1.28	0.55	0.07	0.45	120.07
0016l	1.87	2.19	0.94	1.59	126.33
0018l	1.62	0.79	2.14	0.28	141.15
0023l	2.58	2.28	0.99	0.84	102.17
0028l	2.49	2.33	0.98	1.68	73.27
0029l	3.90	2.48	0.62	2.49	68.53
0030l	2.44	6.39	3.57	5.18	117.00
0031l	3.21	4.00	4.89	4.02	67.76
0032l	2.58	2.25	5.39	2.72	85.63
0035l	1.59	2.87	0.88	1.55	112.62
0038l	2.37	8.31	1.88	1.96	120.39
0103l	1.32	0.86	0.06	0.55	103.36
0104l	2.90	1.94	0.21	0.45	113.61
0105l	1.32	1.14	3.10	1.19	66.27
0115l	2.65	5.67	6.69	3.34	96.01
0121l	2.23	2.07	0.26	2.29	72.37
0122l	2.13	3.29	4.17	3.55	73.03
0123l	2.31	3.92	1.19	1.30	90.16
0125l	1.82	1.14	5.12	1.04	78.27
0127l	2.34	2.39	4.08	1.79	77.11
0128l	2.26	3.13	0.34	3.34	67.99
0133l	1.63	1.87	2.61	0.87	71.32
0134l	1.92	1.98	2.29	1.49	72.38
0142l	1.69	1.89	9.61	1.44	92.03
0147l	1.33	7.62	3.07	13.67	66.00
0148l	1.52	1.32	0.16	1.55	77.92
0149l	1.42	1.63	4.97	1.60	57.13
0150l	1.64	3.34	1.72	1.16	63.55
0309l	1.23	2.00	2.17	1.30	63.92
0311l	1.36	4.34	0.29	3.06	68.97
0312l	2.91	3.97	5.00	3.38	53.91
0313l	0.92	1.40	0.08	0.88	73.53
0322l	0.83	1.99	0.15	1.11	73.72
0325l	0.86	2.69	0.76	1.11	73.20
0328l	1.14	5.69	74.38	1.72	78.61
0329l	2.99	12.76	45.11	5.81	91.08
0330l	0.68	0.77	0.09	0.60	75.06
0331l	5.11	12.53	1.67	9.25	66.68
0332l	1.22	4.90	0.14	1.54	73.86
0333l	8.73	24.09	8.51	9.61	73.84
0370l	1.85	3.78	2.53	2.46	72.77
Average	**2.15**	**3.90**	**5.54**	**2.54**	**83.69**

(FFT) algorithm [8] that was available through the C library [3]. A 50 s window size was selected, border effects were reduced using a Hanning windowing function, and the maximum power band larger than 0.4 Hz was chosen to estimate the HR. The FFT was recalculated every second by shifting the window by 75 samples.

The performance of the four algorithms was then assessed using the normalized root mean square (NRMS) error (%). 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:

Table 3. Normalized root mean square (NRMS) error for the labeled artifact sections between the reference HR, obtained from ECG, and the HR estimated by the vocoder frequency vocoder peak detection, Pan-Tompkins, or FFT method, respectively

Case	**HR NRMS error** [%]				**mean HR** [bpm]
	vocoder peak	vocoder freq	Pan-Tompkins	FFT	ECG
00161	2.35	1.97	1.85	1.08	128.71
00301	0.34	5.76	0.64	1.58	107.15
00321	3.15	1.00	4.71	3.29	85.69
00351	3.20	5.61	2.39	3.21	118.43
01051	1.29	1.00	13.34	0.48	65.69
01151	2.22	3.32	7.37	2.49	97.56
01231	3.52	3.84	4.92	0.63	89.64
01271	4.40	0.66	5.71	0.63	75.19
01471	2.62	2.23	13.48	0.78	64.82
01501	1.78	5.68	6.06	0.90	61.41
03091	2.88	2.55	2.63	1.52	62.78
03121	5.10	3.54	8.78	4.01	54.35
03701	2.62	0.90	7.68	1.15	76.16
Average	**2.73**	**2.93**	**6.12**	**1.67**	**83.66**

Table 4. Algorithm execution time per sample on the iPod Touch calculated over 16 measurements of one minute

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

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

The NRMS error was calculated for each test measurement by comparing it with the reference measurement that was nearest to it in time. The first 50 s, which were used to initialize the high-pass filters and the vocoder sliding window, were not analyzed. The artifact labels from the technician were ignored by the algorithm in this assessment.

To evaluate the robustness of each algorithm with respect to artifacts in the PPG, the NRMS error was then calculated on the sections that were labeled as artifacts. Processing delays were taken into consideration by extending the sections by 5 s. The statistacl tests (Lilliefors,T-test, and Wilcoxon) were performed at the 5% significance level.

We built a prototype device, called the *Phone Oximeter* [4] to evaluate the computational load of the algorithm when executed on a mobile device (Figure 2). This consists

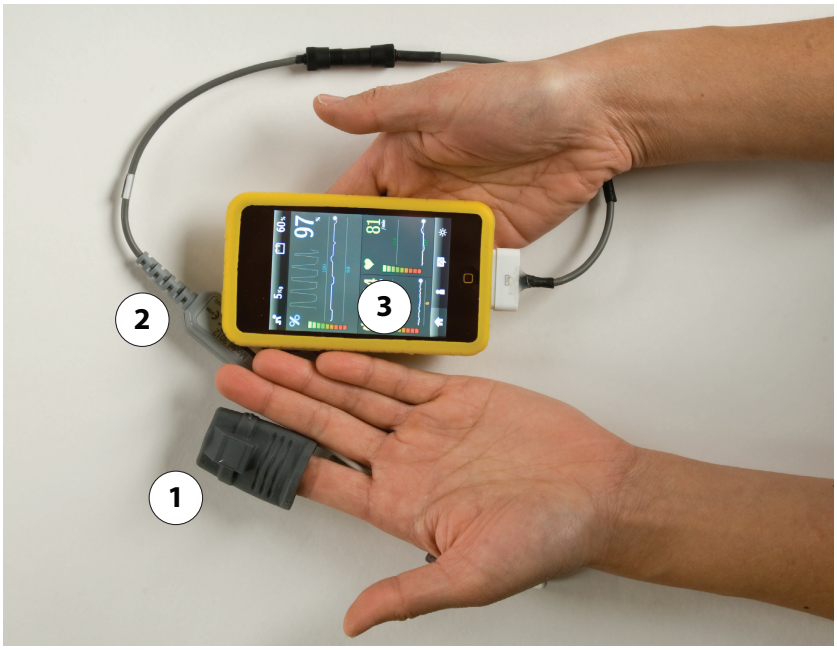

Fig. 2. The *Phone Oximeter* (mobile phone pulse oximeter) application. 1) Soft finger probe; 2) Nonin Xpod OEM pulse oximeter module; 3) iPod Touch with raw PPG waveform, SpO2 and HR trend display.

of a PureLight medium soft sensor connected to an Xpod OEM module (Nonin, Plymouth, USA), which was in turn connected to a 2^{nd} generation iPod Touch (Apple, Cupertino, USA) that features an ARM11 620 MHz microprocessor running underclocked at 533 MHz. 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 to process the raw PPG signal in real-time [4]. The computational load was measured in μs dynamic program analysis using an integrated low-overhead flat profiler.

3 Results

The two vocoder methods performed at least as well as the two traditional methods in terms of overall NRMS error (Table 2). The HR calculated from the vocoder peak detection using histograms (average NRMS error of 2.15%) was more accurate than the vocoder frequency estimations (average NRMS error of 3.9% , T-test, $p<0.01$). Both vocoder methods did not show a clear difference in their distribution from the Pan-Tompkins (average NRMS error of 2.15%) and the FFT (average NRMS error of 2.15%) method (T-test, $p>0.1$).

The vocoder methods were robust towards artefacts (Table 3). Thirteen of the 42 obtained recordings contained PPG artifacts. The total duration of artifacts in the dataset

was 113.83 s. The NRMS errors for the artifact experiment were not normally distributed (Lilliefors, $p>0.07$). The Pan-Tompkins algorithm was the algorithm most vulnerable to artifacts, with an average NRMS error of 6.12% (Wilcoxon, $p<0.031$). Both vocoder algorithms and the FFT method showed similar NRMS error distributions (Wilcoxon, $p>0.06$).

All the algorithms showed a processing speed that allows real-time HR computation. The adaptive vocoder algorithms including peak detection required an average of 47.33 μs to process a new value (Table 4). While this was slower than the Pan-Tompkins algorithm (28.84 μs), both were significantly faster than the inter-sampling distance of 13.3 ms. In contrast, the FFT HR estimation method was about 40 times slower than the vocoder algorithms. The algorithms can, therefore, be computed in real-time on the *Phone Oximeter* with sufficient processing power remaining for other computing tasks.

4 Discussion

The major design criteria of accuracy and on-line capabilities for a mobile phone HR estimation algorithm were met by the two newly developed vocoder algorithms.

4.1 Overall Accuracy

With an average accuracy of 2.15% NRMS, the HR prediction from vocoder peak detection is accurate for a mobile device. Any disagreement between it and the ECG reference HR arose from delays in the vocoder peak detection HR estimation; a 5 s to 10 s delay observed in some cases can be attributed to the size of the histogram for HR calculation. Only the cases 0331l and 0333l showed a decrease in performance greater than 4%. In these cases, the peak detection algorithm and the histogram method were not able to fully compensate for the poor performance of the vocoder frequency method which was required for the peak detection algorithm to accurately estimate HR. The vocoder peak detection was able to track fast variations and compensate for the low responsiveness of the direct vocoder frequency calculation.

The HR calculated directly from the vocoder output frequency was also accurate. However, six cases (0030l, 0038l, 0147l, 0329l, 0331l, and 0333l) exhibited higher NRMS errors than the others. Closer inspection illustrates that reduced performance of the vocoder frequency HR in these cases was mainly due to the low responsiveness of this algorithm to a rapidly changing HR of more than 10 bpm (Figure 3). The vocoder responsiveness could be improved by tuning the cut-off frequency of the low-pass filter shown in Figure 1-3. In the special case of 0333l, the vocoder locked into the double frequency of the original signal. This may have been due to a strong presence of a dicrotic notch that was amplified during the Butterworth filtering. A possible solution to circumvent this problem is to validate the vocoder frequency using the output of the peak detection algorithm.

A desired improvement to the vocoder peak HR estimation algorithm is to reduce the delay that was introduced by the histogram detection algorithm. It is, therefore, evident that the advantages and drawbacks of the two vocoder 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.

The Pan-Tompkins algorithm showed three cases with very large NRMS errors (0009l, 0328l, and 0329l). These were the results of repeated double detection of peaks. To avoid these errors in the future, rejection of too closely located peaks will be necessary. For the FFT method, NRMS errors larger than 5% were caused by the detection of larger frequency content in the frequency band spanning over the double HR.

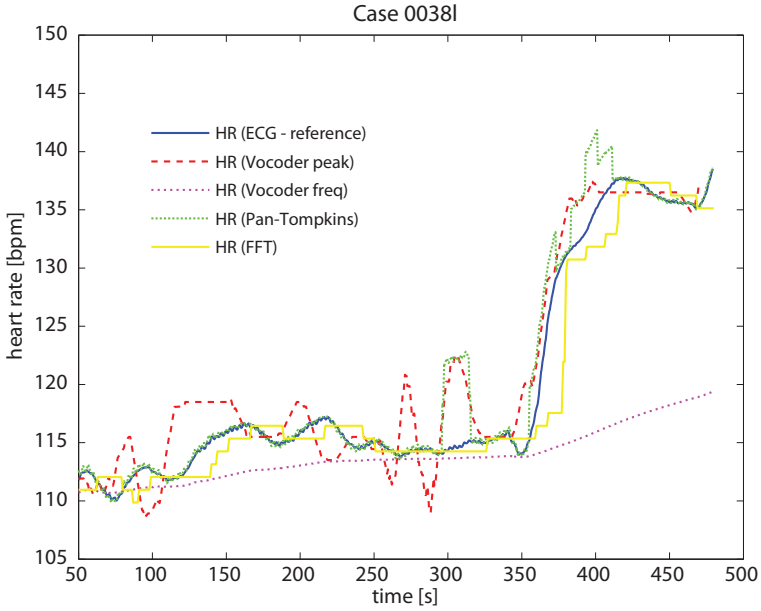

Fig. 3. HR agreement comparison for case 0038l. The HR estimated by the vocoder frequency does not respond rapidly to a large permanent change in HR after 350 s. The HR calculated by the other methods follows the ECG HR trend more closely.

4.2 Robustness to Artifacts

The adaptive vocoder frequency output filtered short variations in HR and PPG artifacts effectively. 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. On the other hand, the vocoder peak detection method was able to rapidly track changes in HR. Nevertheless, this attribute made it more vulnerable to errors when the PPG was corrupted with noise. Using a histogram to determine HR limited theses errors, but increased the detection delay. In the future, we plan to increase its robustness to poor quality PPG signals by adding an additional computation of a signal quality index to the algorithm. Active detection of poor signal will prevent the erroneous HR output and increase overall estimation performance.

The Pan-Tompkins algorithm was less robust to artifacts than the other methods. Its simple method of using a moving-average to filter the HR was not able to compensate for these artifacts, and propagated the error for the averaging window duration (Figures 4 and 5). The FFT algorithm was robust to shorter artifacts, but had a relative low time resolution, since it analyzed all frequencies inside a 50 s window. This was also

expressed by it having the largest delay of all methods, which can easily be seen in Figures 3 and 4.

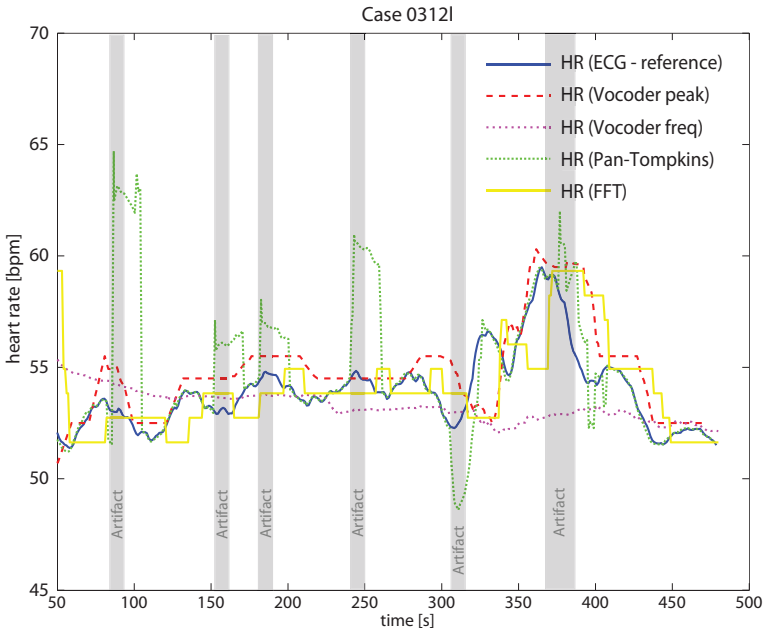

Fig. 4. HR agreement comparison for case 0312l. This case illustrates how the Pan-Tompkins algorithm is impacted by artifacts in the plethysmogram (grey bars) and the propagation of the error for the duration of the smoothing window.

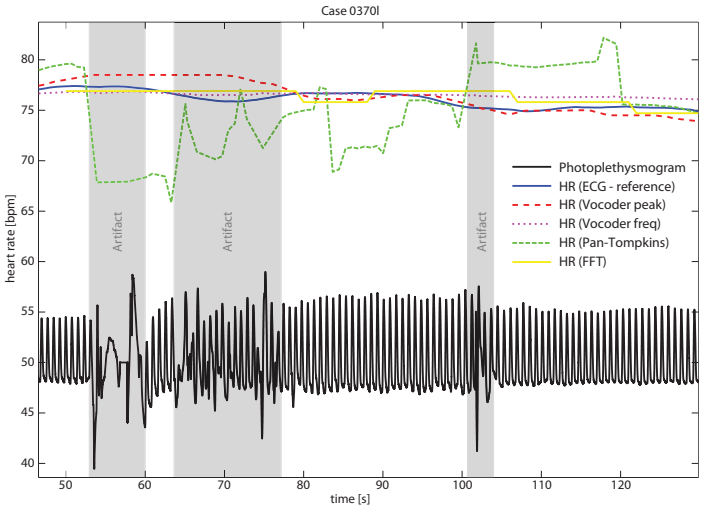

Fig. 5. HR agreement comparison for case 0370l. This case illustrates how the Pan-Tompkins algorithm is impacted by artifacts (grey bars) in the plethysmogram waveform.

4.3 Computational Costs

The speed tests on the *Phone Oximeter* showed that the proposed algorithms achieve the computation within the range of traditional filtering methods like the Pan-Tompkins algorithm and faster than the frequency domain-based approach of the FFT. Processing speed and theoverall accuracy achieved by the four tested algorithms suggested that they are all valid methods for HR estimation on mobile devices. However, the computing intensive FFT method is not optimized for low-power systems.

5 Conclusions

Two novel approaches to compute HR based on a dynamically adapted, single frequency phase vocoder algorithm were proposed. The computed HR was accurate with respect to the reference HR computed from ECG and compared favorably with other algorithms in terms of overall accuracy, robustness to artifacts and computational costs. This makes the algorithms suitable to process on-line PPG signals that are recorded from an oximeter sensor interfaced to a mobile device. We intend to use these algorithms as a core engine for more complex mobile signal analyzers (i.e. for the estimation of heart rate variability or respiratory rate) within a mobile phone based pulse oximeter called the *Phone Oximeter*. The suggested algorithms have the potential to be applied to other periodic signals whose frequency range needs to be determined in real-time.

References

1. Dolson, M.: The Phase Vocoder: A Tutorial. Computer Music Journal 10(4), 14–27 (1986)
2. Flanagan, J.L., Meinhart, D.I.S., Golden, R.M., Sondhi, M.M.: Phase Vocoder. The Journal of the Acoustical Society of America 38(5), 939–940 (1965)
3. 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., Stanley, H.E.: PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. Circulation 101(23), e215–e220 (2000), `http://www.physionet.org/physiotools/wfdb/psd/fft.c`
4. Karlen, W., Dumont, G., Petersen, C., Gow, J., Lim, J., Sleiman, J., Ansermino, J.M.: Human-centered Phone Oximeter Interface Design for the Operating Room. In: Traver, V., Fred, A., Filipe, J., Gamboa, H. (eds.) HEALTHINF 2011 - Proceedings of the International Conference on Health Informatics, pp. 4333–4337. SciTePress, Rome (2011)
5. Karlen, W., Turner, M., Cooke, E., Dumont, G., Ansermino, J.M.: CapnoBase: Signal database and tools to collect, share and annotate respiratory signals. In: Annual Meeting of the Society for Technology in Anesthesia (STA), p. 25. West Palm Beach (2010), `http://www.capnobase.org`
6. Kohler, B., Hennig, C., Orglmeister, R.: The principles of software QRS detection. IEEE Engineering in Medicine and Biology Magazine 21(1), 42–57 (2002)
7. Pan, J., Tompkins, W.J.: A real-time QRS detection algorithm. IEEE Transactions on Bio-Medical Engineering 32(3), 230–236 (1985)
8. Press, W., Teukolsky, S., Vetterling, W., Flannery, B.: Numerical recipes in C: the art of scientific computing, 2nd edn. Cambridge University Press, New York (1992)