31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

Real-Time, Low-Complexity, Low-Memory Solution to ECG-Based Heart Rate Detection

Sourabh Ravindran, Steven Dunbar and Bhargavi Nisarga

***Abstract*— This paper addresses the issue of heart rate detection from noisy ECG data, and presents a method with low complexity and low memory requirements that can detect QRS complex in the presence of noise and muscle artifacts. On the MIT-BIH arrhythmia database we were able to detect 99.3% of QRS complexes with 0.47% false detection. This method can also be applied to heart rate detection using phonocardio signals.**

I. INTRODUCTION

Heart rate detection from ECG signals has been well studied. One of the popular methods proposed by Pan et al. [1] is easy to implement and is shown to have very good QRS complex detection rates. However, the amount of smoothing could prove to be an issue while detecting elevated heart rates as in the case of tachycardia or heart rate monitoring in fitness equipments. A low computational complexity method proposed previously [2], detects the QRS complex peaks by high pass filtering and subtraction of a fixed threshold. Though straightforward, this method is not robust to variations in relative amplitude. Khaver et al. [3] proposed the use of discrete wavelet transform to remove mean variations and other artifacts. The peaks corresponding to the R-waves were detected using slope detection and thresholding. This method provides good detection rates but is computationally intensive. Some of the other signal processing choices include adaptive filtering to remove artifacts, matched filtering to detect the QRS complex, and frequency tracking.

The autocorrelation method has been successfully employed in speech processing to find the pitch of a speaker [4]. For heart rate detection, Lee et al. [5] have proposed the use of cross-correlation between a segment and another segment that is one cardiac period away. The cardiac period used is an estimate from the previous iteration. Peters et al. [6] proposed using the cross-correlation between ECG signals from different leads to obtain an estimate of the fetal heart rate. This paper proposes a modification of the correlation-based method that can be used for heart rate detection from ECGs, digital stethoscopes and wearable acoustic heart monitoring systems in high noise conditions.

Sourabh Ravindran is with Speech and Audio Lab, Signal Processing Systems R&D Center, Texas Instruments, Dallas, USA. sourabh@ti.com
Steven Dunbar is with Analog/RF Field Applications, Rocky Mountain Region, Texas Instruments, USA. s-dunbar@ti.com
Bhargavi Nisarga is with MSP430 Applications, Texas Instruments, Dallas, USA. b-nisarga@ti.com

II. METHOD

In the proposed method, the input noisy signal is segmented into frames of length N samples with an overlap of $N-1$, while ensuring that N is an odd number. Each frame is folded around the center sample and a dot product is calculated. This operation is similar to a zero lag cross-correlation between the first half of the frame and a time-reversed second half of the frame. The center sample is replaced by the calculated dot product. This procedure is repeated for all the frames, essentially replacing each sample with the cross correlation between the $(N-1)/2$ samples before it and the time-reversed $(N-1)/2$ samples after it.

The basic premise of this method is that the QRS complex has more symmetry around its peaks than do noise and artifacts. Therefore samples around the R-wave peaks will produce a large value of the dot product compared to noise and muscle artifacts. The same concept can be extended to phonocardio signals where the peaks due to $S1$ and $S2$ can be separated from those due to noise. We refer to this as the folded correlation method.

A. Algorithm

The folded correlation method is incorporated into a heart-rate detection algorithm for ECGs as follows: incoming data is first buffered into frames and each of these frames is subjected to a first order difference function which removes the mean as well as low frequency components from the signal. The frame is then processed using the folded correlation method to enhance the R-wave peaks. The envelope of the resulting signal is then input to a peak picking function. Peaks over 5 seconds are collected and subjected to post processing. During post processing, peaks with relatively small amplitudes and small distance (time) separation are discarded (for our application we target a maximum heart rate of 180 beats per minute (BPM)). To save memory, the difference between the peak locations is stored rather than the locations of the peaks themselves.

The heart rate obtained from each 5 second segment is subjected to a weighted average over the past 4 heart rate values. The resulting value is the reported heart rate for last added frame. The information from the oldest frame is discarded and a new frame is added to obtain the next estimate. The algorithm is designed to output heart rate values once per second. Figures 1-3 show the working of the algorithm for a noisy ECG input. Figures 4-6 show the algorithm's performance for an input with strong T-wave components.

The specific application targeted was heart rate detection using ECG for fitness equipments. In this application the acquired signal is often of poor quality, and it is important to design a solution with low memory and low computational requirements to minimize cost. With a sampling rate of 60 Hz and correlation window length, N=5, the folded correlation consists of 2 multiplies and 2 adds. Since the method does not rely on smoothing to remove artifacts, it results in a relatively high temporal resolution of the peaks which makes it ideal for dealing with elevated heart rates. The algorithm has been shown to work with as little as 500 msec of buffered data, this leads to very low memory requirements (as low as 128 bytes in our implementation with input data resolution of 8 bits).

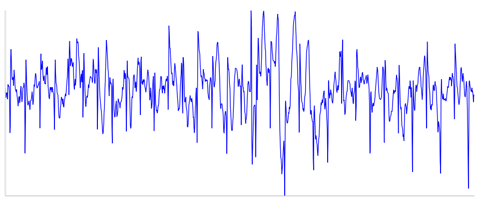

Fig. 1. Noisy input ECG.

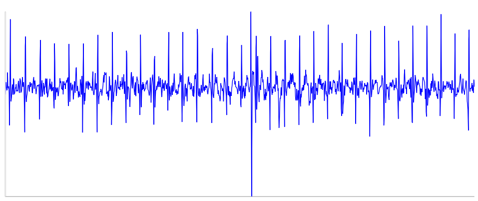

Fig. 2. Result of first order difference operation that removes the DC offset and slowly varying components (from waveform in Fig. 1).

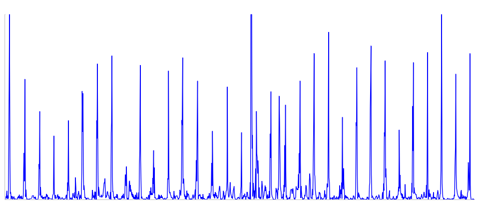

Fig. 3. Output of the folded correlation (absolute value). The R-wave peaks are enhanced while most of the artifacts are suppressed.

B. Hardware

The system consists of a hardware circuit that detects the presence of valid data and enables data logging. The detection circuit consists of a dual comparator with a tunable threshold voltage. The output of the comparator is used to generate a pulse that begins or terminates data logging. The signal path consists of an amplifier that provides gain while rejecting line noise due to its low common mode error, a high-pass filter for DC removal and a low-pass filter to reject high frequency noise. The signal is then fed to an MSP430 mixed signal microcontroller, which computes the heart rate. This solution can run in as little as 128 bytes of RAM and 2 KB of code space. The hardware system is shown in Figure 7.

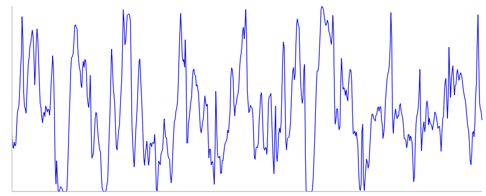

Fig. 4. Input ECG where the T-wave component dominates.

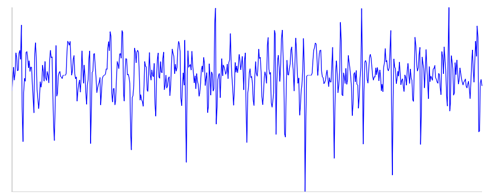

Fig. 5. Result of first order difference operation that removes the DC offset and slowly varying components (from waveform in Fig. 4).

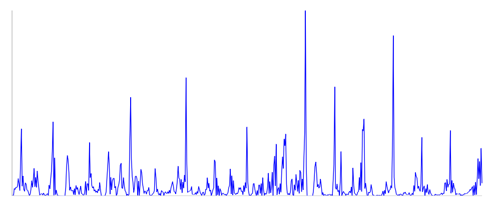

Fig. 6. Output of the folded correlation (absolute value). The R-wave peaks are enhanced while most of the artifacts are suppressed.

III. Results

Figures 8-9 show the results from a working system. The input is an ECG that varies from 170 BPM to 90 BPM within roughly 2 minutes. As the plot shows, the system is able to quickly ramp up to the 170 BPM level and track the heart rate accurately as it decreases.

The results from running the algorithm on the MIT-BIH arrhythmia database is shown in Table I. We were able to detect 99.3% of the QRS complexes with 0.47 % of false detections. Figures 10-12 show a few examples from the MIT-BIH database. As can be seen from these plots there is fair amount of agreement between the expert QRS location markings and the QRS locations detected by the proposed algorithm.

IV. Phonocardio Data

This section presents results from using the folded correlation method to detecting primary heart sounds (S1 and S2) in the presence of noise. The background noise may come from environments such as airports, exhibition halls, subways and trains. One of the targeted applications is to reliably extract the heart rate at the scene of an accident using a digital stethoscope. Another application is audio-based heart rate monitoring of elderly people in a home environment, in the presence of different types of noise. The method does not need any kind of training, which makes it ideal for multiple noise environments.

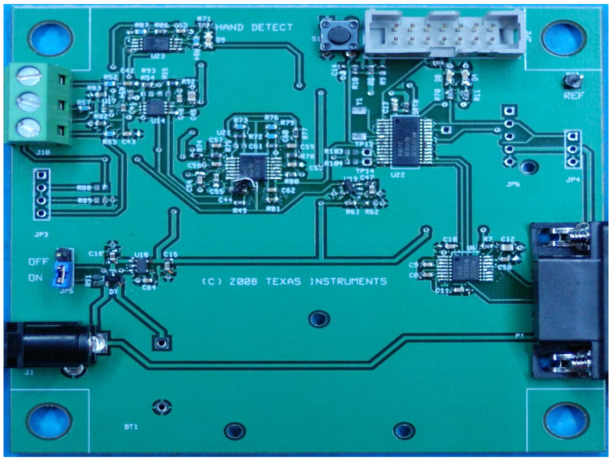

Fig. 7. Figure showing the system hardware.

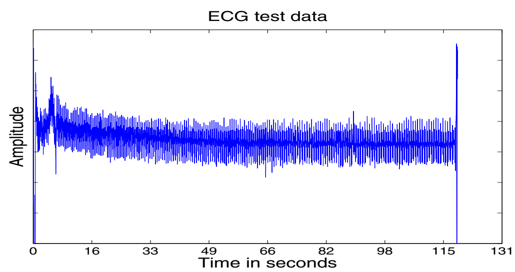

Fig. 8. Input ECG with heart rate varying from 180 BPM to 90 BPM over roughly 2 minutes.

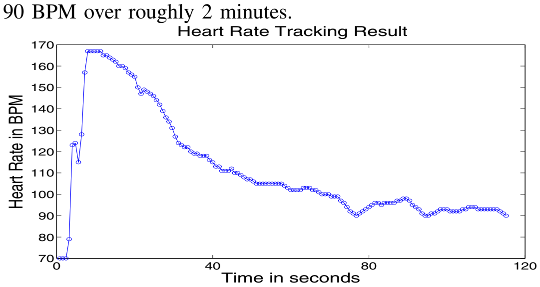

Fig. 9. Plot showing the heart rate detection system tracking the heart rate.

Figures 13-15 show the operation of the algorithm in different kinds of noise and various signal-to-noise ratios (SNRs) ranging from 0 to -15 dB. In each case, the noise was added synthetically, using noise files from the Aurora database. A peak picking algorithm was used to identify the peaks once the signal is processed using the folded correlation method. S1/S2 in the figures are identified using timing information (i.e. S1-to-S2 distance is smaller than S2-to-S1 distance).

V. FUTURE WORK

Future work will apply the folded correlation method to identifying pathological conditions found in phonocardio signals. Murmurs presumably will have less correlation than S1 and S2 sounds which can be exploited to distinguish between peaks due to murmurs and those due to primary heart sounds.

Record No.	Number of QRS	False Negatives	False Positives
100	35	0	0
101	32	0	0
102	34	0	0
103	32	0	0
104	35	0	0
105	38	0	0
106	31	0	0
107	33	0	0
108	28	0	2
109	44	0	0
111	32	0	0
112	40	0	0
113	27	0	0
114	25	0	0
115	29	0	0
116	36	0	0
117	23	0	0
118	34	0	0
119	30	0	0
121	28	0	0
122	41	0	0
123	22	0	0
124	23	0	0
200	41	0	0
201	40	0	0
202	24	0	0
203	49	3	2
205	41	0	0
207	27	0	0
208	47	3	0
209	43	0	0
210	43	4	0
212	42	0	0
213	51	0	0
214	35	0	0
215	52	1	0
217	33	0	0
219	33	0	0
220	34	0	0
221	35	0	0
222	35	0	0
223	37	0	0
228	37	1	0
230	38	0	0
231	29	0	0
232	25	0	4
233	48	0	0
234	42	0	0
	1693	12	8

TABLE I
TABLE SHOWING THE FALSE NEGATIVES AND FALSE POSITIVES FOR EACH OF THE RECORDS FROM THE MIT-BIH ARRHYTHMIA DATABASE.

VI. ACKNOWLEDGMENTS

The authors would like to thank Randy Cole for fruitful discussions and helpful hints in designing the algorithm.

REFERENCES

[1] Jiapu Pan and Willis Tompkins, “A Real-Time QRS Detection Algorithm, *IEEE Transactions on Biomedical Engineering*”, vol. BME-32, no. 3, Mar 1985.

[2] Murugavel Raju, “Heart-Rate and EKG Monitor using the MSP430FG439”, *Texas Instruments Application Report*, Sept 2007.

[3] M.A. Khaver amd M.A. Haque, “ECG Peak Detection using Wavelet Transform”, *ICECE*, 2004.

[4] Lawrence R. Rabiner, “On the use of autocorrelation for pitch detection”, *IEEE Transactions on ASSP*, vol. 25, no. 1, Feb 1977.

[5] Jung H. Lee, Sun I. Kim, and Doo S. Lee, “Fast cross-correlation method for real time detection of fetal heart rate”, *IEEE Engineering In Medicine and Biology Society*, 1998.

[6] Chris Peters, Rik Vullings, Jan Bergmans, Guid Oei, and Pieter Wijn, “Heart Rate Detection in Low Amplitude Non-Invasive Fetal ECG Recordings”, *IEEE Engineering In Medicine and Biology Society*, 2006.

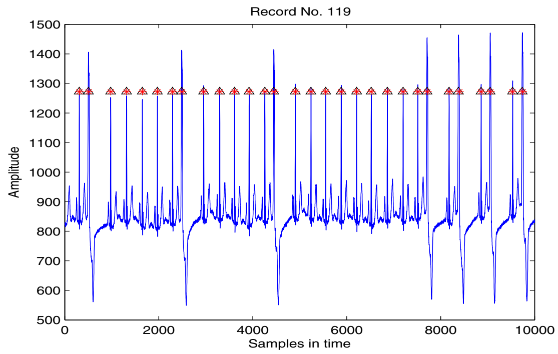

Fig. 10. Figure showing the QRS detection for record no. 119 from MIT-BIH database. The black triangles show the expert's marking and the red asterix show the detection output of the proposed algorithm.

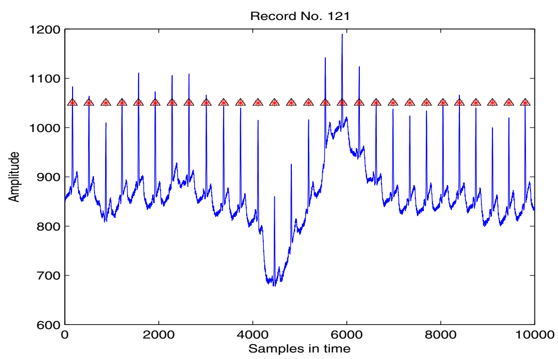

Fig. 11. Figure showing the QRS detection for record no. 121 from MIT-BIH database. The black triangles show the expert's marking and the red asterix show the detection output of the proposed algorithm.

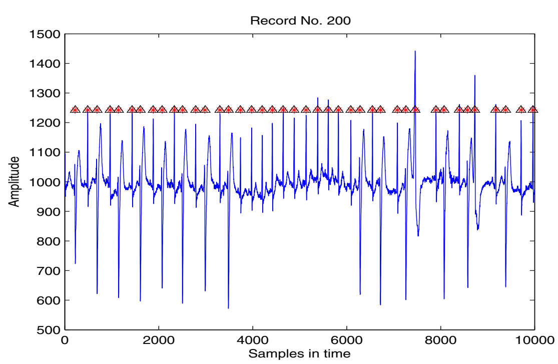

Fig. 12. Figure showing the QRS detection for record no. 200 from MIT-BIH database. The black triangles show the expert's marking and the red asterix show the detection output of the proposed algorithm.

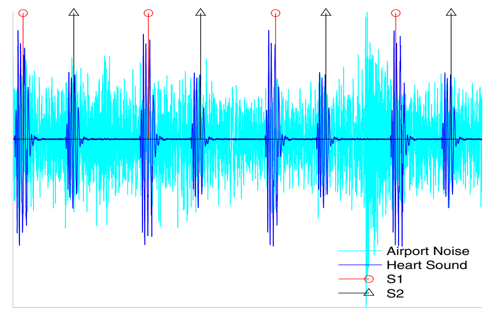

Fig. 13. Figure showing the result of the heart rate detection algorithm for heart sound in the presence of synthetically added airport noise at 0 dB SNR.

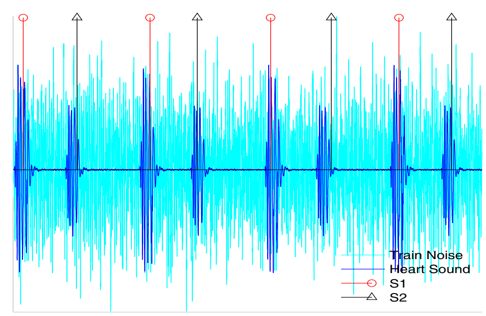

Fig. 14. Figure showing the result of the heart rate detection algorithm for heart sound in the presence of synthetically added train noise at -5 dB SNR.

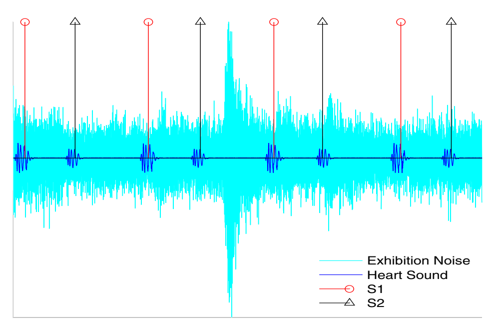

Fig. 15. Figure showing the result of the heart rate detection algorithm for heart sound in the presence of synthetically added exhibition noise at -15 dB SNR.