LOW POWER SIGNAL PROCESSING FOR HEALTH MONITORING
A STUDY OF ALGORITHM DESIGN

T.S. Sun

Sanders
A Lockheed Martin Company
Nashua, NH 03061-0868
tsun@sanders.com

Rajeevan Amirtharajah
Anantha Chandrakasan
Alan Oppenheim

Massachusetts Institute of Technology
Cambridge, MA 02139

Michael V. Scanlon

Army Research Laboratory
Adelphi, MD 20783-1145

ABSTRACT

Low power signal processing is a very important technology for many tactical applications. The success of this technology depends on a well-orchastrated collaborative research effort in the area of sensor, processor and algorithm designs. This paper reports the results of a robust algorithm design for detection and rate estimation of heartbeats from acoustic sensor data, with the objective to demonstrate the feasibility of wearable medical sensor technology. The algorithm has the desirable property that it can be executed efficiently in the ultra-low power processor architecture reported (elsewhere?) with a consistent performance sustaining over a wide range of energy/power-saving operational conditions. These results were obtained using simulated processor functions and experimental acoustic data collected at ARL

INTRODUCTION

Low power signal processing is a very important technology for many battlefield applications. When integrated with a properly designed sensing device, this technology can be used to support a light-weight versatile information/intelligent system. A good example concerns the exploitation of compact acoustic sensors for two diverse applications. First, compact acoustic sensors can be deployed at hard-to-access remote sites by means of ballistic projectile to monitor battlefield activities, such as movements of ground vehicles and airborne helicopters, and locations of artilleries. Second, compact acoustic sensors can be carried or worn by combat personnel to monitor individual's health condition. This latter aspect is particularly invaluable in support of a vital combat casualty care, as has been elucidated in a report earlier (scanlon).

The success of a compact acoustic sensor technology dwells on four critical components: acoustic sensor, low power signal processor, power supply, and processing algorithm. A collaborative research effort has been orchestrated by the participating members to address key technical issues, with (a) ARL focusing on sensor design, and implementation and application-related issues, (b) MIT focusing on processor architecture design and innovative power supply concept, and (c) Sanders focusing on algorithmic design. This paper concerns primarily the results of an enabling algorithm design as it applies to medical monitoring.

One power saving strategy is to make the processing load power/energy scalable so that when the available power is low the system can still perform a simple but crucial function whereas when the available power is high it can perform multiple tasks or complex functions to acquire more detailed information. This strategy is particularly applicable when designing a wearable acoustic sensor. In an early study we have shown that the acoustic sensor can be used to acquire three types of signals regarding the person's condition. These are heartbeat, breath, and voice. Both heartbeat and breath signals provide not only the vital sign but also an indication of physical stress of the individual. The voice signal can provide an indication of the stress condition or can be used directly as a communication medium to provide accurate and detailed information. However, these three types of signals require different amounts of processing power, with heartbeat the least and voice the most. In this regard, it is desirable to design a power/energy scalable processing system such that at minimum it provides a reliable heart rate measurement, with the potential to extend to speech processing when more power/energy sources become available. In this paper we will discuss an algorithm design for heartbeat rate estimation with consideration to this energy/power scalability aspect.

ALGORITHM DESIGN

Primary Design Considerations

The task of algorithm design faces three major challenges. The first one is to achieve a reliable estimate of heartbeat rate in the presence of strong background noise and clutter. The acoustic energy of heartbeat signal is spectrally concentrated below 100 Hz. Although the signal-to-noise ratio (SNR) can be improved in principle by means of a low-pass filter, many common acoustic sources with energy in this region do exist. They include those originating from the wearer, such as breaths, vocal utterances, footsteps, and contacts of equipments, and those originating from the environment, such as wind and most motorized mechanical objects in the field. It is not unusual that the SNR of heartbeat signal as acquired in the field is below 0 dB. Therefore, the algorithm is required to achieve a high SNR gain.

The second challenge is to minimize the processing cost while meeting the processing performance objective. Specifically, the algorithm is required to find the most efficient way of processing under the specific processor hardware environment. The ultra-low power processor design currently developed at MIT (raj) imposes several constraints if an optimal power efficiency is to be maintained.

- Arithmetic and logic operations are limited to multiplications, additions, comparisons, and bit shifting.
- The operands and the results of all operations are limited to integers of fixed bitwidth, whose maximum size including the sign bit is 12 bits.
- The arithmetic unit has an internal maximum dynamical range of 24 bits.
- The bitwidth window at the arithmetic output is limited to four positions on the 24-bit range.

Because of these constraints, many recently advanced high-performance signal processing algorithms, specially those requiring floating-point operations or involving nonlinear functions, are not good candidates for the design model. Algorithms based on convolution appear to be most suitable.

The third challenge is to achieve a high degree of robustness as the processing power is scaled down from maximum to minimum. As discussed in a separate paper (raj), the energy/power consumption of the processor can be reduced in two way: (1) by reducing the bitwidth of operands and (2) by reducing the data rate through decimation. Accordingly the algorithm is required to provide a reliable heartbeat measurement over a sufficient range of changes regarding both bitwidth and data rate.

Baseline Processing Algorithm

The algorithm is divided into four processing stages (Fig. 1). The first stage performs a correlation of the sampled signal sequence with a template to enhance SNR gain. The template is constructed from the average waveform of individual's heartbeat signal over a duration normally less than 100 ms. Optimization of heartbeat detection can be obtained through a careful tuning of the template waveform.

The second processing stage performs a weighted averaging of the output magnitude of the first stage with the purpose to mark the location of the power centroid of each heartbeat. This is done by means of convolution with a bell-shaped window function. We have found that the result is not particularly sensitive to the exact shape of the window function. A Hanning window with a length of 64 ms has been used in this work..

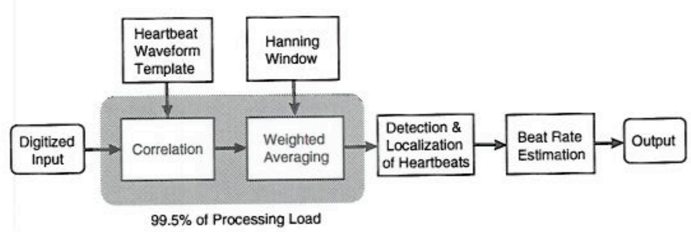

Figure 1. Flow Diagram of Heartbeat Detection and Rate Estimation

The third stage of processing performs detection and localization of heartbeats. This is done by finding the maximum amplitude in the output of the second stage within a preset time interval (60-300 ms). To minimize the problem of false heartbeat detection, an adaptive threshold is used to screen the peaks. The threshold is set at a value proportional to the mean value of the data points within an interval averaged about 2 s. To void the use of division, the sample mean is obtained by first summing the magnitudes of 2^n data points and then performing a right shift of n bits on the sum. The larger the bitwidth is, the less sensitive the performance is to the proportional constant. The constant is set to $1^1/_4$ for bitwidth greater than 5 bits and $^1/_2$ elsewhere. Once a heartbeat peak is detected, its location is saved in memory.

The final stage of processing determines the intervals between consecutive peaks and then estimate the average of heartbeat period from a sample of consecutive intervals. In the actual implementation, we used the median value instead of the mean value as an estimate of the heartbeat period. This has the advantage of making the measurement less sensitive to the influence of misses and false detections and at the same time avoiding the operation of division. The heartbeat period can be converted to the conventional heartbeat rate by multiplying its inverse with a constant.

It is important to point out that more than 99% of the processing load is consumed in the first two stages, which involves two convolution-type processes. The effect of decimation on the processing load can be approximated by the expression.

$$\text{No. of Operations} = 2(M_1+M_2)ND^{-2} - 4ND^{-1} + C$$

where M_1 and M_2 represent the sizes of the correlation template and the Hanning window, N the length of input data, all prior to decimation, D the decimation factor, and C is a constant representing overhead cost. The number of operations decreases by two orders of magnitude if the decimation factor increases eight fold.

EXPERIMENTAL DESCRIPTION

The acoustic data of heartbeats under various stress conditions were collected in a typical indoor laboratory environment without anechoic surrounding. The stress condition was created by asking a test subject to walk on a stationary treadmill at various speeds. Two male volunteers (simply referred as subjects A and B) of different ages and physical builds participated in the experiment. Each participant went through a speeding up period and a cooling off period. During each session, the data were recorded simultaneously from two acoustic sensors worn on the neck of the participant. Other sensors were also used to establish the ground truth and to monitor the background noise and interference. They include a blood pulse oximeter for timing the heartbeats, an ambient mike for recording the sound of footfalls, and an accelerometer for measuring the vibration level of the walker. We also recorded the voice of each participant during walk with the objective of exploiting speech signal for acquiring additional health information. However, this part of study has not been completed and will not be discussed here.

The acoustic data collected in this experiment, although free from major external noise sources, do not contain clean heartbeat signals because of two problems. First, the emplacement of acoustic sensors was not optimized to enable a good transmission of heartbeat signal to the sensors. The SNR in most part of recording was well below 10 dB. Second, the sound level of footfalls during the speeding up period frequently arose above the heartbeat signal level.

Interference by footfalls has been a very difficult obstacle to overcome in light of two considerations.

- Both the heartbeat and the footfall signals have similar spectral and temporal characteristics. Therefore, it is difficult to design a linear filter to separate the two.
- The heartbeat and footfall signals may synchronize at certain walking speed.

We have carefully analyzed the correlation between heartbeats and footfalls in order to minimize the likelihood of false detection. It was discovered that the heartbeat signal of Subject B was completely masked by his footfall signal during the speeding up period. For this reason, this part of data was excluded from the following analysis.

SUMMARY OF RESULTS

Baseline Performance

As a performance reference, we first present the result of heartbeat measurements under the baseline condition.

The algorithm is executed in MATLAB using 64-bit floating-point operations. The acoustic data were sampled at 2 kHz with a 1kHz-cutoff low-pass filter. This bandwidth is a reasonable compromise between meeting low power requirement and preserving the vocal information.

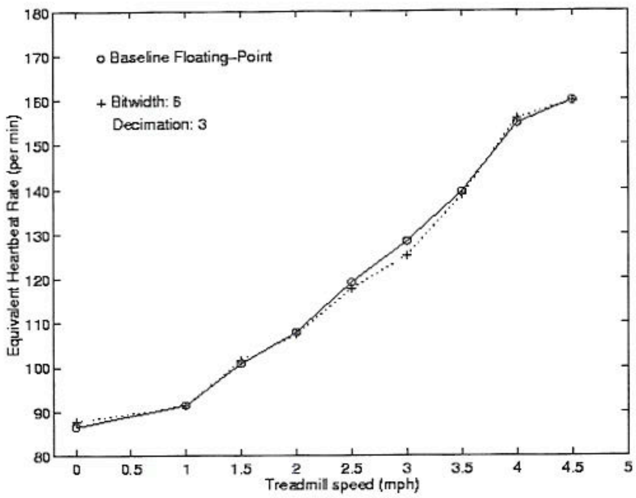

Figure 2. Heartbeat Rate of Subject A as Function of Treadmill speed

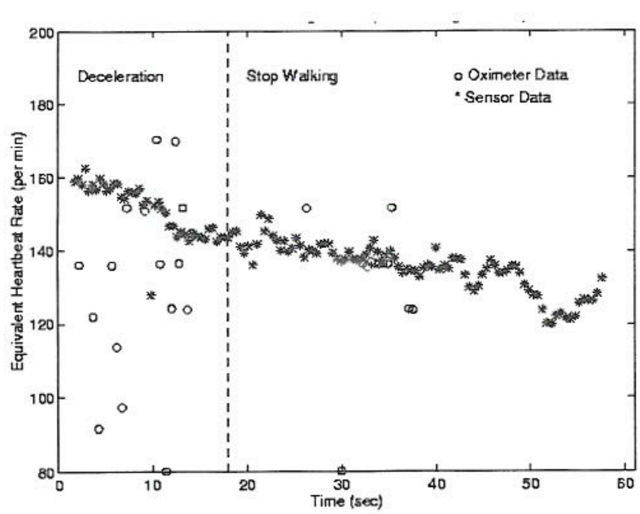

Figure 3. Change of Heartbeat Rate for Subject A during Cooling-Off Period

Figure 2 shows that the heartbeat rate as function of treadmill speed for Subject A. Figures 3 and 4 show the changes of heartbeat rate of Subjects A and B during their cooling-off period respectively. These two participants appeared to relax in very different patterns as a reflection of their physiological difference. Figure 4 also shows the data of oximeter, which is in excellent agreement with the acoustic sensor data. These results confirm the suitability of the algorithm for heartbeat rate monitoring.

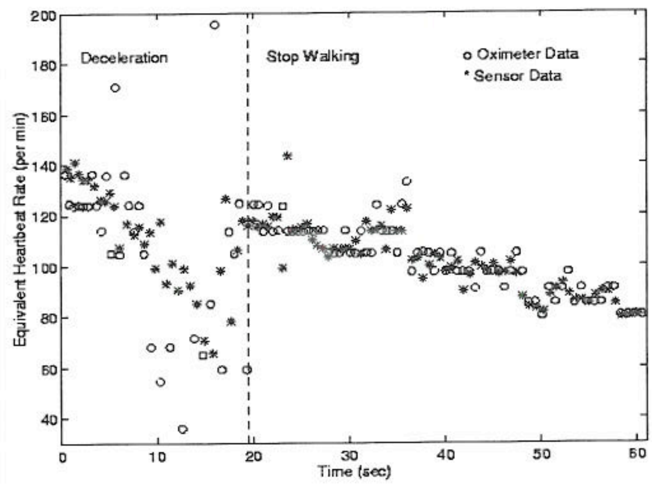

Figure 4. Change of Heartbeat Rate for Subject B during Cooling-Off Period

Effect of Power/Energy Saving on Performance

We next show how the processing algorithm performs subject to the hardware constraints and under a wide range of power saving conditions. More specifically, we examine how replacing the floating-point operations with fixed-point operations and how reductions in data rate and bitwidth affect the performance. In this analysis, we simulated the processor functions in a C-coded program in which the bitwidths of the operands and the output of every arithmetic operation are tailored in accordance with the processor constraints mentioned earlier. The results of this exercise are best presented by plotting the following three performance indicators as functions of decimation factor and bitwidth (Figures 5 through 7).

- The number of heartbeats missed from detection.
- Timing Error of individual beats
- Error in the heartbeat period estimation

The abscissa in the plots is the bitwidth pf operands and each curve represents a result using decimation factor, as marked directly on the curve.

A number of observations can be made from these figures.

- The change from floating-point operations to fixed-point operations has resulted in a noticeable degradation in the performance factors. It is noted that 12% of heartbeats have escaped detection.

- The performance appears to be relatively stable over a wide range of bitwidth and decimation factors. In particular, the error in the heartbeat period estimation remains very low even the data rate is reduced by eight fold and the bitwidth is reduced to 6 bits.
- The performance begins to degrade rapidly as the bitwidth is reduced below 6. Reductions in both bitwidth and data rate have an additive effect in the low bitwidth (< 6) region.

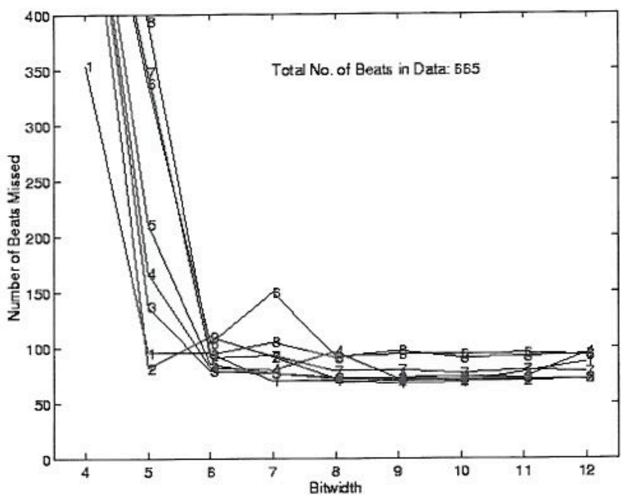

Figure 5. Number of Misses as Functions of Bitwidth for Various Decimation Factors

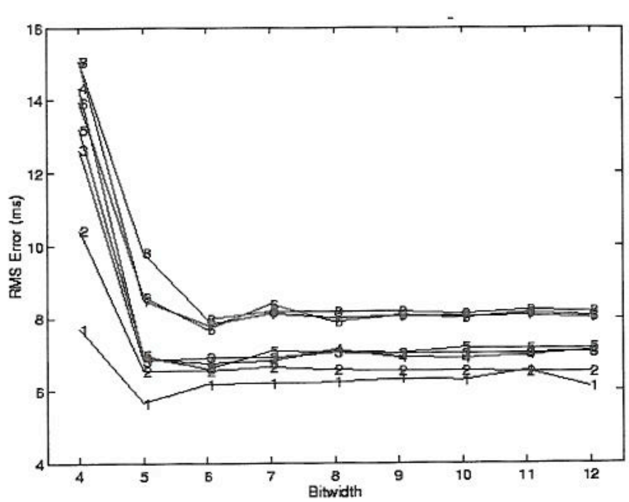

Figure 6. Error of Individual Beat Timing as Functions of Bitwidth for Various Decimation Factors

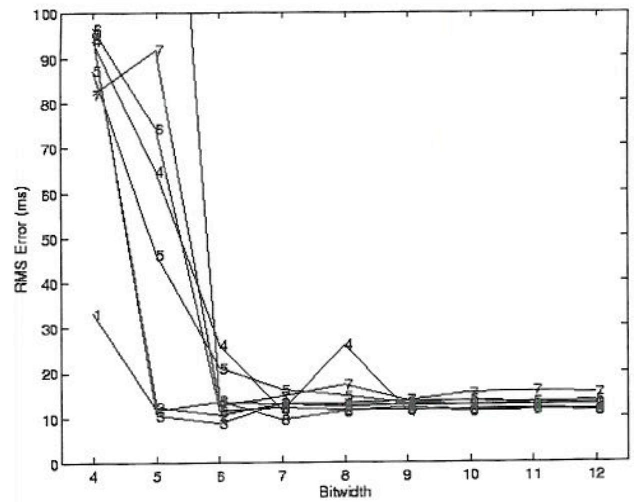

Figure 7. Error of Heartbeat Period Estimation as Functions of Bitwidth for Various Decimation Factors

- The degradation pattern does not follow a monotonic relationship with either bitwidth or decimation factor. In particular, the best result is not the one based on the largest bitwidth and data rate as expected from a simple theoretical consideration. This is attributed primarily to an interplay between the restricted dynamic range in the processor and three data-dependant factors. The first factor concerns the relatively large size of the kernel in the correlation process. The second factor concerns the small size of the data set. The third factor is due to a large fluctuation in the acoustic power of individual heartbeats.

CONCLUSION

We have demonstrated a robust algorithm design for detection and rate estimation of heartbeats in acoustic signals with two desirable properties

- It can be implemented efficiently subject to low power processor hardware constraints.
- Its performance remains reliable for a wide range of data rates and bitwidths.

These properties make this algorithm design a viable component in the development of wearable sensor technology for heath monitoring applications. However, this study is only the first step toward a demonstration of its applicability. A more rigorous test of the algorithm will be performed with the actual processor developed at MIT when it becomes available.