

AN ULTRA-LOW POWER SUBBAND-BASED ELECTRONIC STETHOSCOPE

Julie Johnson¹, David Hermann¹, Melody Witter¹, Etienne Cornu¹, Robert Brennan¹ and Alain Dufaux²

¹AMIS Canada, 611 Kumpf Drive, Unit 200, Waterloo, Ontario, Canada N2V 1K8

²AMIS Switzerland, Champs-Montants 12a, 2074 Marin, Switzerland

ABSTRACT

Electronic stethoscopes are able to offer signal amplification and other benefits over traditional stethoscopes. However, many electronic stethoscopes rely on a personal computer for their signal processing, which reduces portability and requires a relatively large amount of power. This paper presents a low power, portable electronic stethoscope system that is based on a signal processing approach using an over-sampled filterbank. This system is implemented on an ultra-low resource DSP system. The stethoscope incorporates multiple filtering modes as well as audio record and playback (full and half speed) functionality.

1. INTRODUCTION

Auscultation continues to be a routine part of cardiovascular and pulmonary examinations. Recently, there have been a number of advances in auscultation technology, including the introduction of a number of electronic stethoscopes [1,2,3]. Many of these stethoscopes rely on the use of a personal computer to filter, record and replay the signal. As such, these systems are large and have limited portability. Little emphasis has been placed on generating high quality audio in a small, portable unit. Clearly, an electronic stethoscope system with the features of these larger systems that is portable and low power would be advantageous.

An electronic stethoscope system has been implemented on a highly flexible ultra-low power, miniature DSP system. At the heart of this system is an over-sampled, weighted overlap-add (WOLA) filterbank specifically designed for subband signal processing. Multiple filtering modes, specialized for cardiovascular and pulmonary auscultation applications, have been designed. For recording, regular speed playback and half speed playback, a low resource audio codec was developed specifically for low frequency signals. Sufficient resources remain for additional algorithms including heartbeat monitoring and wireless functionality.

One of the most challenging aspects of the design is the integration of algorithms requiring different filterbanks. A method of fulfilling the requirements of two filterbanks while minimizing computations and thus power requirements, is presented. Filterbank selection is further complicated by the low frequency nature of the input signals. A method of combining the two required filterbanks while effectively addressing this issue will be presented. This method also accommodates the application of large gain adjustments to decoded signals and is an essential component of the half speed playback algorithm.

In the following sections, a description of the DSP system architecture is presented, followed by an overview of the stethoscope system. In Section 4, the implementation of the stethoscope functionality is discussed in greater detail showing innovative solutions to the design challenges discussed above. Section 5 provides a characterization of the stethoscope's performance and conclusions are presented in Section 6.

2. DSP SYSTEM

The DSP system consists of three major components: an 18-bit block floating point weighted overlap-add (WOLA) filterbank coprocessor, a 16-bit fixed-point DSP core, and an input-output processor (IOP). These components run in parallel and communicate through interrupts and shared memory. The parallel operation of these components enables the implementation of complex signal processing algorithms with low system clock rates and low resource usage and is particularly adept at subband signal processing.

The WOLA coprocessor implements a flexible over-sampled filterbank. Although initially designed for analysis and synthesis involving over-sampled, complex subband signals, it may be adapted to generate critically-sampled, real-valued filterbanks as required for the codec in this application [4].

The algorithms are implemented on the DSP system using a 16-band, 4-times over-sampled WOLA filterbank configuration with odd-stacking. The selected configuration generates a group delay of 17 ms, consumes 4.1 mW, has a system clock frequency of 5.12 MHz and a sampling frequency (F_s) of 8 kHz.

3. STETHOSCOPE SYSTEM

The stethoscope system is deployed on the DSP system described in Section 2. The high-level block diagram for the overall stethoscope system is shown in Fig. 1.

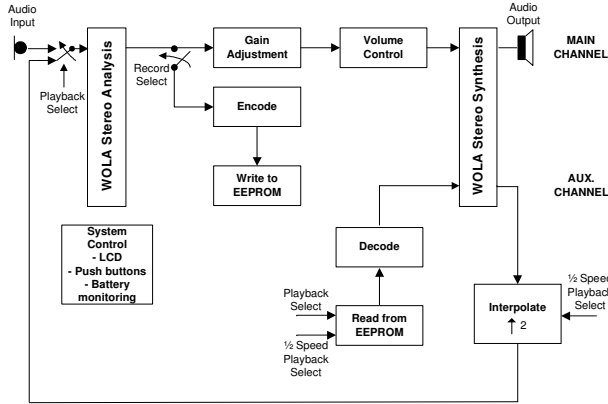


Figure 1. High-level block diagram of the stethoscope system.

The stethoscope is comprised of a number of modules including subband gain adjustment, record functionality, regular and half speed playback functionality, as well as system-level features including battery monitoring, volume control and a LCD display. This paper focuses on two algorithms operating on the subband signal: gain adjustment and subband coding/decoding.

The subband gain adjustment algorithm provides frequency shaping as required by the listening modes. Generally, the components of heartbeat and lung sounds useful for diagnostic purposes are in the range of 20-1000 Hz [5]. The first through fourth heart sounds fall in the range of 20-115 Hz [5]. Disorders such heart murmurs occur in the range of 140-600 Hz [5]. Thus, a suitable listening range for heart sounds is approximately 20-600 Hz. For breathing sounds, the strongest part of the signal is typically under 100 Hz, although the signal can have useful components up to 1.2 kHz [5,6].

The subband codec is used as part of the record/playback functionality. During recording, the signal is captured, encoded, packed and written to non-volatile memory (EEPROM). During playback, the packed signal is read, unpacked, decoded and re-synthesized in real-time.

The filterbank requirements of the subband gain adjustment algorithm and the subband coding algorithm are significantly different. Subband gain adjustment requires low delay and optimal filter responses to reduce the level of uncanceled aliasing that is generated when gains are varied in different subbands. The WOLA filterbank uses over-sampling to achieve high levels of aliasing reduction without increasing the filter length and consequently the group delay [7]. To keep the group delay as low as possible, a relatively high sampling frequency of 8 kHz was selected.

Since the bulk of the usable signal falls under approximately 1 kHz, the gain adjustments required by the different listening modes are quite large. Consequently, an over-sampling factor of at least 4 is desirable to minimize group delay and aliasing artifacts.

In contrast, the subband coding algorithm requires a critically-sampled, real-valued filterbank to achieve minimal data rates [8]. As described in [4], critically-sampled, real-valued subband signals can be obtained by postprocessing the over-sampled complex WOLA subband signals.

A WOLA filterbank configuration was required to simultaneously accommodate the large adjustment requirements of the signal processing algorithms at the lowest group delay while also providing a critically-sampled, real-valued filterbank for coding purposes. Further complicating this task is the stringent demands on aliasing performance placed on the WOLA filterbank when the most common configurations are used with low-frequency signals. Furthermore, the selected filterbank configuration must leave sufficient free cycles so that additional algorithms may be integrated into the system in the future.

Another challenging problem is the half speed playback algorithm. This feature presents an interesting real-time sample rate conversion problem and also involves issues related to the fixed sampling rate of the DSP architecture.

4. SYSTEM DESCRIPTION

This section describes how the stethoscope was designed to overcome the challenges addressed in Section 3.

4.1. Over-sampled filterbank parameter setup

As mentioned above, the WOLA filterbank is required to have low group delay and an over-sampling factor of 4. A filterbank configuration which satisfies these criteria uses an analysis window length (L_a) and a synthesis window length (L_s) of 128 samples, an input block size of $R = 8$ samples and an FFT size of $N = 32$. The inputs to this filterbank are lung and heartbeat sounds, both of which are very low frequency sounds. Using typical prototype filters, these low frequency signals produce aliasing distortion in the WOLA filterbank. The aliasing can be seen at 1 kHz, 2 kHz and 3 kHz in Fig. 2. This aliasing is particularly strong for two reasons. First, the signal contains a strong low frequency component that is very close to one side of a subband where the filter behaviour is least ideal, i.e., the filter's transition band. Secondly, this low frequency component is very strong compared to the remainder of the spectrum, which makes the aliasing more obvious.

The solution is to develop a customized prototype filter optimized for low frequency sounds that provides a sharper cut-off while maintaining passband flatness. The resulting filterbank configuration yields a group delay of 17 ms and




Figure 2. Output when using lung sounds as input and a WOLA filterbank configuration of $R = 8$, $N = 32$ and $L_a = L_s = 128$.

achieves over 90 dB of aliasing rejection.

4.2. Subband gain adjustment algorithm

Three different filter modes have been designed based upon the characteristics of heart and lung sounds: a bell mode, which amplifies low frequency heart sounds in the range 0-500 Hz, a diaphragm mode, which amplifies lung sounds in the range 0-1000 Hz and an extended range mode which amplifies sounds between 0-1500 Hz.

The use of an over-sampled filterbank permits the application of extremely efficient gain adjustments. The gain application is a vector process in which each subband is multiplied by a real-valued gain. The gain application process occurs on dedicated, efficient hardware, namely, the WOLA coprocessor, which was described in Section 2.

The number of subbands used in this design is 16. Since F_s is 8 kHz, the bandwidth of each band is 250 Hz. The system utilizes odd-stacking which means that the first band encompasses the frequencies from 0-250 Hz. To implement the bell mode, for example, gains greater than zero are provided for the first two subbands while gains of zero are provided for the remaining subbands. A larger number of subbands would provide improved frequency resolution but would increase the computational complexity and require longer filters (and more group delay) for equal levels of aliasing reduction.

4.3. Low resource subband codec

The codec used in this application is a subband fixed-point implementation of the one word memory APCM quantizer developed by Jayant [8]. This codec was selected due to its low resource usage and good performance with slowly varying input signals. The implementation used here has a data rate of 20 kbps. The bottom 7 subbands of the signal, which cover the range between 0 -1750 Hz were quantized

This subband codec requires critically-sampled, real-valued subband signals as input. Since the filterbank

required by the gain adjustment algorithm has an over-sampling factor of 4, the analysis results must be down-sampled by a factor of 4.

The analysis filterbank suitable for the subband codec has the form described by (1), where $h_m(n)$ is the subband analysis filter, m is the subband index, $M=16$ is the number of subbands and $h_p(n)$ is the prototype low-pass filter. The filter length, L was set to $L_a = L_s$. Note that this filterbank, a cosine-modulated filterbank, uses odd-stacking and that the WOLA filterbank is also configured for odd-stacking.

$$h_m(n) = h_p(n) \cos \left[\frac{\pi}{M} \left(m + \frac{1}{2} \right) \left(n - \frac{M}{2} \right) \right] \quad n = 0 \dots L-1 \quad (1)$$

To achieve this filterbank, two steps are required. First, to reduce the data by a factor of 2, every other input block is skipped. This effectively doubles the block size (R). The resulting subband signals do not contain additional aliased images because the original over-sampled subband signal is band-limited to $\pi/4$. Secondly, to obtain critically-sampled, real-valued data, the method described in [4] is used to generate the cosine-modulated filterbank in (1).

In the design shown in Fig. 3, analysis results in the main channel are decimated, encoded, packed and stored during the record operation. During playback, these signals are unpacked, decoded and interpolated into an auxiliary channel. The reconstructed signals are synthesized in this auxiliary channel and then copied from the output of this channel to the input of the main channel.

This method was necessary in order to overcome the distortion that is created when a decoded signal undergoes gain adjustment. The decimation of the signal before coding generates aliasing that becomes audible once amplified by the gain adjustment algorithm. To overcome this, the reconstructed signal must be filtered prior to gain adjustment in order to remove the unwanted aliasing. Rather than implementing a filter on the DSP core, a more efficient solution within this architecture is to synthesize and then re-analyze the data using another over-sampled filterbank prior to gain adjustment.

This two-channel approach was selected because the two separate analysis and synthesis chains are more efficiently implemented on two channels than two completely separate filterbanks can be implemented on a single channel. The chosen method takes advantage of available capabilities on the WOLA coprocessor while minimally increasing resource usage.

4.4. Half speed playback

An algorithm that plays a decoded recording at half speed was developed to allow medical professionals to hear the finer details of heart and lung sounds. The challenge of this feature is that it requires data to be decoded at one rate and

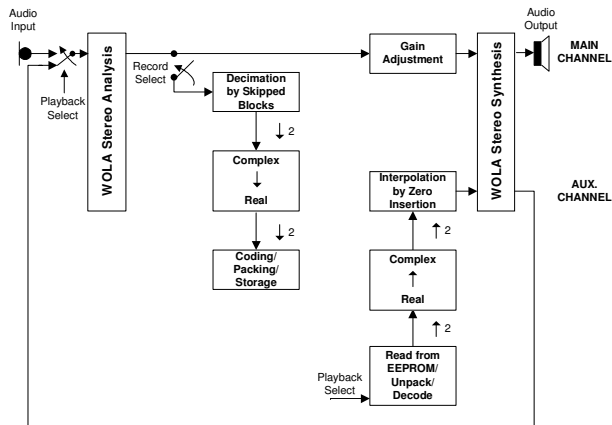


Figure 3. Method of combining filterbanks.

played back at another, in an architecture that has a fixed sampling rate. Halving the speed of a signal doubles the amount of the data to be processed. Due to memory restrictions, storing this extra data is not an option. Thus, the problem must be solved by changing the effective sampling rate of the recorded data while playing it back at the normal fixed sampling rate.

The playback speed is halved by interpolating the decoded signal by a factor of 2 in the time domain while keeping the system's sampling rate constant. Although this method does not preserve the pitch of the signal (it is halved), the details in the sounds are more clearly heard by the medical professional in this mode. One issue is that the interpolation of the time domain signal creates an image of the entire spectrum. Conveniently, the gain adjustment algorithm always removes the top half of the spectrum. Thus, mode filtering that is already in place can be used to eliminate this imaging.

5. SYSTEM CHARACTERIZATION

To characterize the audio performance of system, a prototype stethoscope was developed. The prototype consists of a microphone positioned behind the diaphragm of the chestpiece and a hearing aid receiver positioned at the base of the ear tubing. This prototype is displayed in Fig. 4.

The dynamic range of the stethoscope was found to be approximately 75 dB electrically. The amplification provided by the DSP in all three filter modes was approximately 21 dB electrically. This amplification provides a significant improvement over non-electronic stethoscopes. For example, many diaphragm chestpieces on mechanical stethoscopes actually attenuate the signal [5]. The DSP operates on 1.8 V and during typical operation, consumes 4.1 mW. The prototype consumes 47 mW, of which approximately 43 mW are consumed by an LCD



Figure 4. Stethoscope prototype.

circuit and voltage regulators. Two AAA batteries will provide 72 hours of continuous operation.

6. CONCLUSIONS

A portable, high fidelity electronic stethoscope has been successfully implemented on an ultra-low power DSP system. The system incorporates multiple listening modes as well as record/playback functionality. The system successfully integrates an over-sampled filterbank with a critically-sampled, real-valued filterbank and is able to handle very low frequency input signals. The maximum signal amplification is approximately 21 dB and the electrical dynamic range of the system is approximately 75 dB. Future work will involve the integration of other algorithms such as noise reduction, a heartbeat detector and wireless functionality.

7. REFERENCES

- [1] S. Lukkarinen et. al., "A New Phonocardiographic Recording System", *Computers in Cardiology*, pp. 117-120, Vol 24, 1997.
- [2] A. S. Brown et. al., "PhiSAS: A Low-Cost Medical System for the Observation of Respiratory Dysfunction", *IEE Seminar on Appropriate Medical Technology for Developing Countries*, pp. 10/1-10/5, 2002.
- [3] S. Lukkarinen et. al., "Multimedia Personal Computer Based Phonocardiography", *Proc. 18th Annual International Conference of the IEEE*, pp. 2303-230, Vol. 5, 1996.
- [4] D. Hermann et. al., "Low-Power Implementation of the Bluetooth Subband Audio Codec", *Proc. ICASSP 2004*.
- [5] M. Abella and J. Formolo, "Comparison of the Acoustic Properties of Six Popular Stethoscopes", *J. Acoust. Soc. Am.* 91 (4), Pt. 1, pp. 2224-2228, April 1992.
- [6] I. Mazic, S. Sovilj and R. Magjarevic, "Analysis of Respiratory Sounds in Asthmatic Infants", *Measurement Science Review*, Vol. 3, Section 2, 2003.
- [7] R. Brennan and T. Schneider, "Filterbank Structure and Method for Filtering and Separating an Information Signal into Different Bands, Particularly for Audio Signal in Hearing Aids", *United States Patent 6,236,731*. WO 98/47313. April 16, 1997.
- [8] N. S. Jayant and P. Noll, "Digital Coding of Waveforms", Prentice Hall, 1984.