Reducing Background Noise through a Stethoscope Cup using Adaptive Filters

Bryce E. Hill MS EE

Abstract—**An adaptive filtering algorithm was used to test the validity of adaptively filtering respiratory signals recorded at the trachea with an external reference microphone. Two different setups were tested. The first used a microphone in open air, the second used a microphone that was housed inside a second stethoscope cup. The primary microphone was affixed to a phantom material. External sounds and music were played via aloud speaker to record additive noise data from within the stethoscope cup. Data showed that adaptive filtration using a secondary stethoscope cup was the most effective method to remove ambient noises.**

Index Terms—**adaptive filters, stethoscope acoustics, noise cancellation, respiratory sound detection.**

I. Introduction

A. Objective

The goal of this project was to find the validity of using noise cancellation to reduce ambient sounds recorded within the precordial stethoscope. Noises generated outside of the stethoscope that are not wide sensse stationary (WSS) can cause signals to be recorded within the stethoscope that are loud enough to be counted as a detected breath. If these sounds are recorded during a period of apnea the acoustic signal will not be able to detect it as a period of apnea and it is classified as a period missed by the acoustic apnea detection algorithm.

For this reason a method to reduce disturbances caused by ambient sounds was explored. The method used was an adaptive filter requiring a secondary microphone to record ambient sounds.

B. Stethoscope acoustics

The stethoscope cup used was a heavy precordial cup shown in figure 1. The stethoscope cup was designed to amplify signals detected within the cup. Physiologically the skin within the stethoscope cup creates a diaphragm that acts like a loud speaker for the vibrations on the skin. The metal stethoscope attenuates external signals from entering the cup. Acoustics of a stethoscope cup attenuation and amplification can be characterized [1] through experimentation, but this can change depending on the placement of the stethoscope cup on the skin, the tightness of the skin within the cup, and the placement of the stethoscope cup on the trachea. Observations from previous research showed that sounds such as talking or machine alarms can be loud enough to be detected as breath sounds.

C. Adaptive noise cancellation overview

Adaptive noise cancellation is a method of signal processing that uses multiple sources of signal to produce a desired signal. A simple adaptive filter with block diagram shown in figure 2 has two inputs. Signal d is the primary input and signal x is the reference input. The reference input is then filtered by the weights $\boldsymbol{w}$ of the adaptive filter and an estimate of d is created called y. The difference of the filtered signal y, and desired signal d is then calculated to create the error signal $e=d-y$. The error signal is then used as an input to the adaptive filter to update the filtering weights $\boldsymbol{w}$ as will be described the next section [2].

The adaptive filter algorithm chosen to adapt the weights $\boldsymbol{w}$ can was the least mean squared (LMS) algorithm. The LMS algorithm is the most widely used adaptive filter algorithm due

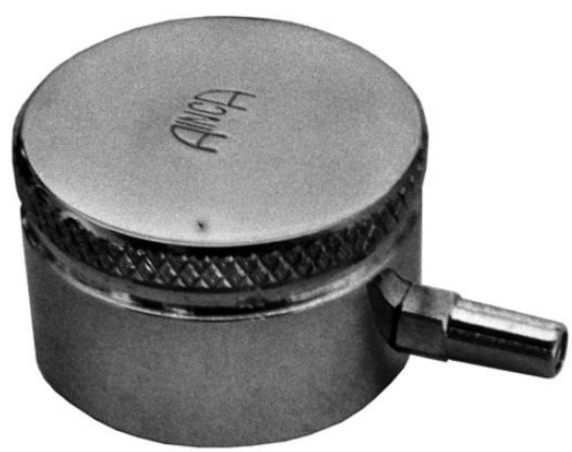

Fig. 1. Metal stethoscope used to house the microphone used for recording the tracheal sounds.

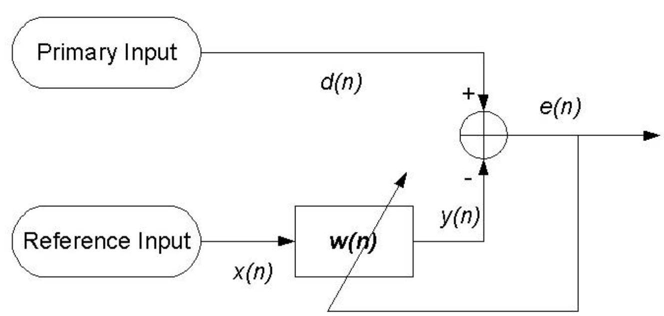

Fig. 2. Block diagram of a typical adaptive filter

Submitted to the NASA Space Grant Consortium April 19, 2010 to its stability, robustness and simplicity [2]. This filter has many variants including the simplified LMS algorithm, and the normalized LMS algorithm.

The classic LMS algorithm updates the weights $\boldsymbol{w}$ are updated with the equation:

$$\boldsymbol{w}(n+1)=\boldsymbol{w}(n)+2\mu e(n)\boldsymbol{x}(n) \quad (1)$$

where $\boldsymbol{w}(n+1)$ is the new vector of filtering weights with length N, $\boldsymbol{w}(n)$ is the current vector of filtering weights with length N, μ is the step size parameter, $e(n)$ is the error signal, and $\boldsymbol{x}(n)$ is a vector of recorded reference signals x with length of N.

The simplified LMS algorithm has three forms and is similar to the classic LMS algorithm. The three forms shown here are the signed regressor, sign, and sign-sign algorithms respectively:

$$\boldsymbol{w}(n+1)=\boldsymbol{w}(n)+2\mu\, sign(e(n))\boldsymbol{x}(n) \quad (2a)$$

$$\boldsymbol{w}(n+1)=\boldsymbol{w}(n)+2\mu\, e(n))sign(\boldsymbol{x}(n)) \quad (2b)$$

$$\boldsymbol{w}(n+1)=\boldsymbol{w}(n)+2\mu e\, sign(n))sign(\boldsymbol{x}(n)) \quad (2c)$$

The signed regressor algorithm is favored for its ability to adapt similarly to the classic LMS but requiring less computations. The sign algorithm and the sign-sign algorithm do not converge as quickly [2] but are not much less complicated than the sign algorithm.

The normalized LMS (NLMS) algorithm adds some complexity in order to improve stability. The general form of the NLMS algorithm is

$$\boldsymbol{w}(n+1)=\boldsymbol{w}(n)+\frac{\tilde{\mu}}{\boldsymbol{x}^T(n)\boldsymbol{x}(n)+\psi}e(n)\boldsymbol{x}(n) \quad (3)$$

where $\tilde{\mu}$ is the step size parameter and ψ is a small value to ensure that the denominator of the equation is never zero. The NLMS algorithm improves stability of the adaptation at the expense of computation. The step size parameter is normalized to the values of the input ensuring that if $\tilde{\mu}$ is properly chosen the output will never become unstable.

The step size parameter μ in the LMS algorithm is the value that controls how quickly the filter adapts. If this value is too high the filter can become unstable and the output of the filter becomes useless. If its value is too low the filter will not adapt quickly to changes in the filtering characteristics. The maximum value that μ can be while maintaining the stability of the filter can be calculated with this equation [2]:

$$\mu_{MAX}=\frac{1}{3\mathrm{tr}[\boldsymbol{R}]} \quad (4)$$

where $\boldsymbol{R}$ is the autocorrelation of the input calculated by:

$$\boldsymbol{R}=E[\boldsymbol{x}(n)\boldsymbol{x}(n)^T] \quad (5)$$

and $tr[]$ is the trace of a matrix defined as the sum of the diagonal of the matrix [2]. Although this can be calculated for every sample of the signal being filtered it is computationally expensive. The NLMS algorithm uses this limit to maximize the step size while maintaining stability.

D. Adaptive noise cancellation in the literature

Adaptive noise cancellation in stethoscopes has been performed for very noisy environments. Patel et al. [3] used an adaptive filtering algorithm to filter helicopter noise from cardiac and breathing sounds through a diaphragm stethoscope cup with a second microphone to record ambient sounds. Data were recorded on a subject in a sound proof room using the stethoscope as described and also using a pneumotachometer to measure respiratory flow volume. Sounds simulating being inside a helicopter were played inside the sound proof room. A real time adaptive filter was used to monitor the progress of the filter. Post processing was performed using both an LMS algorithm with $N=40$ and $\mu=0.02$ and an NLMS algorithm with $N=40$ and $\tilde{\mu}=1.2$. Patel found that the NLMS algorithm provided a significant improvement over both non-filtered data and the LMS filtered data.

Although the work done in [3] is closely related to the noise cancellation performed in this project, there are significant differences in the methods performed. The comparison of the respiratory rate of the acoustic data to that of the respiratory flow data was not discussed by Patel. In addition the adaptive filter was not applied when the reference signal was in a quiet situation. In a quiet enough setting the signal detected by the reference input can be uncorrelated with the noise on the primary input. When the two input signals of the adaptive filter are uncorrelated the reference input can increase the noise on the output when compared with the primary input. This paper defines this phenomenon as contamination. Another difference between the research performed in [3] and the current research was that Patel only used one kind of additive noise at one amplitude. In this paper several different types of additive noise were used at several different additive gain amplitudes. An increase in the number of different kinds of sources and the amplitudes of these sources allows for the production of a more robust adaptive filter. Finally, this research uses an automatic breath detection algorithm to determine the validity of the adaptive filtering algorithm rather than a subjective argument.

II. Methods

A. Data set

Data recorded from the data set as described in [4] were used to test adaptive filter algorithms. From the periods of apnea identified by the respiratory flow meter, forty periods of apnea detected were selected on a manual basis. The periods were selected if the signal fifteen seconds prior to the apnea showed normal breathing on the flow volume signal. The data used included the respiratory flow volume from the CosmoII+ (Respironics) and the raw acoustic signal recorded at the trachea. Conventionally the data recorded at the trachea would be collected on a subject as he/she lies in an an-echoic chamber as different sounds are played and recorded in the stethoscope cup, while the subject performs breath-hold periods. This method was not pursued due to necessity of obtaining IRB approval, the lack of time to receive that approval, and budget constraints. It is noted that the method used involving subjective selection of data was not optimal.

Each data segment was marked as the breathing segment and the apnea segment. This distinction was made so that during breath detection the number of breaths detected could be counted toward a period of breathing or a period of apnea.

The noise floor was calculated for each breath sound segment using the method described in [4]. The noise floor served as a reference for the amplitude of the additive the noise sound that will be described in the next section. The average noise floor measurement for each segment was calculated as σ_{NF}.

The audio data in the forty sets described above were assumed to be free of any major ambient artifact. These data were also recorded without the use of a second reference microphone to record the disturbances coming from outside of the stethoscope cup. For this reason additive noise was recorded at different time using two microphones and a phantom material representing human tissue. The additive noise was added to the acoustic breathing signal after the recordings were finished. This method allowed for the original signal to be known before noise was added and also allowed for different amplitude levels of additive noise to be tested without the need of multiple recordings. This method is also similar to a common image processing technique where an image assumed to be noise free is corrupted with noise and filtering algorithms are tested by comparison of the filtered image to the original image.

Eight sound segments were used consisting of simulated Gaussian noise, talking, and several kinds of music. These sounds were chosen because they are common in an operating room environment.

A microphone (WM-56A103 Panasonic) was placed inside the stethoscope cup (Wenger #00-390-C, AINCA, San Marcos, CA) as was done during the recording of the breathing sounds. The cup was affixed to a gelatin phantom made of edible gelatin formed inside a latex balloon, by a double stick disk (#2181 3M). Gelatin was chosen as cheap phantom that has similar properties to humsan soft-tissue. The balloon was suspended above a table with the second microphone resting outside the stethoscope cup. Speakers were placed on a second table approximately twenty five centimeters from the microphone assembly. The balloon suspension and use of two separate tables was an attempt to minimize mechanical coupling of the loud speakers to the microphones. The speakers played the eight noise segments described above. The amplitude of the sound played by the speakers was adjusted to ensure that the external microphone was not saturated and that the signal was detected by the microphone inside the stethoscope cup. The data recorded at the microphones were digitized via an audio soundcard (SoundBlaster Audigy, Creative, Singapore) at a sample rate of 22 kHz directly to a computer hard-drive. A diagram of this setup is shown in figure 3.

After performing using this filter on the data described it was noticed that the type of additive noise used affected how well the filter performed. The sound sources that performed better were Gaussian noise, talking, or symphonic music. The method for determining how well the filters worked will be described later. Rock music caused some problems because the contamination of the signal even when the gain, G_i, was zero. After looking directly at the resulting waveforms and listening to them, the sounds that were not able to be filtered were strong impulses related to a drum beat or similar sound. It was concluded that the impulse disturbance had a high enough frequency that the adaptive filter could not adapt quickly enough to remove the sound.

An additional problem that was noticed was that the signal of the reference microphone contained much higher frequency signals. It is also assumed that the stethoscope cup attenuates signals in a non-linear manner over both frequency and amplitude. This is a problem because the filter used was a linear filter.

Matching a non-linear filter such as the stethoscope cup has considerable challenges. Although this is possible a single order linear filter such as the NLMS adaptive filter algorithm described would not be sufficient. A non-linear adaptive filter such as an adaptive polynomial filter could be a solution to this problem [5], but it was hypothesized that physically filtering the reference signal with a similar stethoscope cup would be a simpler solution.

The additive noise was recorded again as described above, with the one difference. The difference was that the reference microphone was placed inside an identical stethoscope cup as the primary microphone and affixed to the back of the primary microphone cup via a double stick disk. A diagram of this setup is shown in figure 4. This setup did not test the amount

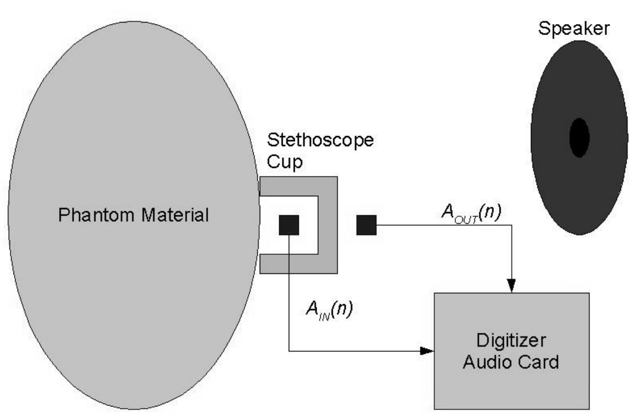

Fig. 3. Diagram of experimental setup for recording ambient room noises using an open air microphone.

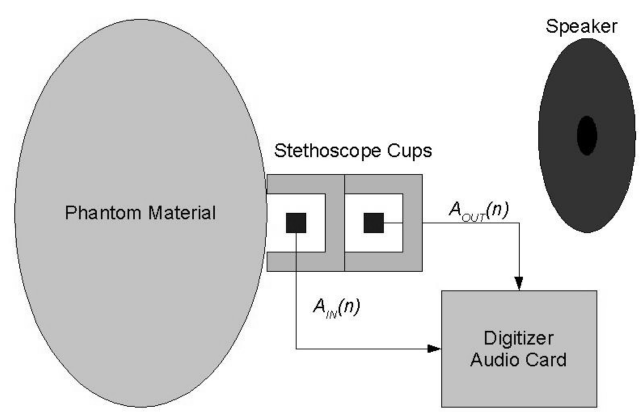

Fig. 4. Diagram of experimental setup for recording ambient room noises using a microphone inside a stethoscope cup.

of desired sound signal such as breathing, that would pass through the primary stethoscope cup to the reference microphone. It is assumed that the reference microphone and cup can be sufficiently isolated to eliminate any of the sounds recorded on the trachea.. The most important features of the reference microphone in this situation are that it is in close proximity of the primary microphone, physically filtered similarly to the primary microphone, and the same kind of microphone as the primary microphone. This is because the additive noise is easiest to filter from the primary source when the reference source is as close as possible to the additive noise signal.

B. Mixing the signals

The standard deviation, σ_{IN}, of the additive noise segment recorded inside the stethoscope cup, $A_{IN}(n)$, was calculated for each of the eight segments. The corresponding reference input will be called $A_{OUT}(n)$. The segment $A_{IN}(n)$ and the segment $A_{OUT}(n)$ were divided by σ_{IN} to normalize the sounds recorded inside the stethoscope. The signal $A_{IN}(n)$ was additionally divided by the standard deviation of the noise floor of the breathing signal σ_{NF} described above to normalize the amplitude of the additive noise signal to the amplitude of the noise floor of the breathing signal.

The normalized sound $A_{IN}(n)$ was added to the breathing sound $B(n)$ resulting in $X(n)$ with equation:

$$X(n)=B(n)+G_i\times A_{IN}(n) \quad (6)$$

where G_i is the gain applied to the noise. The adaptive filter was processed with G_i having values of 0, 2, 4, 6, and 8. The value of 0 was chosen to test the effect of having an uncorrelated signal as the reference input. The other values of G_i were chosen to add a range of sounds that would be detected by the breath detection algorithm.

C. Adaptive filter implementation

An NLMS adaptive filter was used to filter additive noise from the signal $X(n)$. The NLMS algorithm was chosen over the others to ensure stability of the filter. Using the adaptive filter shown in figure 2, the signal $X(n)$ was used as the primary input d and the signal $A_{OUT}(n)$ was used as the reference input x. The signal used as output of the adaptive filter was the error signal e.

This configuration of the adaptive filter works for the following reasons. If the reference signal $A_{OUT}(n)$ was perfectly filtered to match the external noise signal within the stethoscope cup $A_{IN}(n)$, the error signal between y and the input signal $X(n)$ would result in just the breathing signal, $B(n)$.

The value of $\tilde{\mu}$ was chosen experimentally by performing adaptation on the signals of $A_{IN}(n)$ and $A_{OUT}(n)$ before adding the additive noise to the respiratory sound. The signal $A_{IN}(n)$ was used as the primary signal d and the signal $A_{OUT}(n)$ was used as the reference signal x. The output signal used was the error signal e.

The filtration process was calculated using each of the eight noise signals and varying the value of $\tilde{\mu}$ from 0 to 2 in steps of 0.01. The standard deviation of the filtered signal e was calculated. The minimum standard deviation was chosen as the best value of $\tilde{\mu}$. The optimized value of $\tilde{\mu}$ ranged from 0.3 to 0.8 over all of the signals processed. The value of 0.65 was chosen to be the optimal value for the step parameter. The length of the filter N was chosen to be 51 from experimental observations. Figure 5 Shows an example of the original signal $A_{IN}(n)$ and the signal after it was filtered.

After the samples were filtered they were compared to the original input sample. The output sample $e(n)$ was compared to the input sample $X(n)$. If $|X(n)|<|e(n)|$, then the value of the input sample $X(n)$ was chosen as the output sample rather than the filtered signal. This was done to minimize contamination of the output signal with the reference signal. Contamination is clearly evident when the gain G_i is 0 and the output signal $e(n)$ is not the same as the input signal $X(n)$. For the purposes of breath detection this is only a problem when the amplitude of the contamination increases the absolute amplitude of the signal.

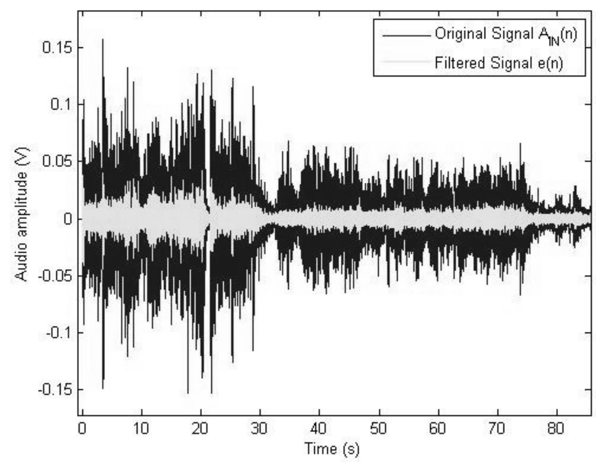

Fig. 5. Example of the adaptive filtering of the additive noise..

D. Breath detection

Breath detection was performed exactly as described in [4]. The standard deviation of the noise floor σ_{NF} and the standard deviation of the detected signal σ_S were calculated using the expectation maximization (EM) algorithm. The breath detection threshold τ was calculated from the two signals σ_{NF} and σ_S. An audio envelope was calculated using the audio signal $X(n)$, and also the filtered signal e for each audio segment, gain G_i, and additive noise signal. Sounds were detected as breaths that had an envelope that rose above the noise floor threshold for a period of 0.3 seconds or more.

The breath detection was calculated using the original respiratory signal $X(n)$. This calculation was used as a reference for the number of sounds detected in the breathing segment and apnea segment before noise was added and the adaptive filter algorithm was performed. The number of sounds detected were counted during the respiratory section of each segment, and the number of sounds detected were counted during the apnea section of each segment. The number of sounds detected in each section was compared to the number of sounds detected when it was not filtered, when it was filtered using an open-air reference signal and when using a reference signal inside a stethoscope cup.

III. Results

To measure how well each algorithm performed the number of periods of apnea of the forty data sets was calculated. The percent error of detecting apnea out of the forty data sets were calculated for non-filtered audio data, data filtered with the reference microphone in open air and the data filtered with the reference microphone in a stethoscope cup. Table 1 shows the percent missed apnea detections for each of these groups with respect to the kind of noise used and the gain G_i of the additive noise.

TABLE I

		Type of additive noise			
		Gaussian	Talking	Symphonic	Jazz
Gain G_i	0	0	0	0	0
	2	0	52.5	7.5	37.5
	4	0	82.5	32.5	75
	6	0	92.5	47.5	85
	8	0	92.5	50	90
		Jazz	Vocal	Rock	Rock
Gain G_i	0	0	0	0	0
	2	60	100	95	92.5
	4	92.5	100	97.5	100
	6	100	100	97.5	100
	8	100	100	100	100

Percent error when detecting apnea when no filtration was performed

		Type of additive noise			
		Gaussian	Talking	Symphonic	Jazz
Gain G_i	0	17.5	2.5	20	42.5
	2	17.5	2.5	20	42.5
	4	25	2.5	20	42.5
	6	67.5	2.5	20	42.5
	8	82.5	2.5	20	42.5
		Jazz	Vocal	Rock	Rock
Gain G_i	0	100	90	22.5	92.5
	2	100	90	22.5	92.5
	4	100	90	22.5	92.5
	6	100	90	22.5	92.5
	8	100	90	22.5	92.5

Percent error when detecting apnea when filtration using an open air reference microphone was used

		Type of additive noise			
		Gaussian	Talking	Symphonic	Jazz
Gain G_i	0	0	0	0	0
	2	0	2.5	0	0
	4	0	15	5	0
	6	0	40	27.5	0
	8	0	57.5	32.5	0
		Jazz	Vocal	Rock	Rock
Gain G_i	0	0	0	0	0
	2	0	0	0	0
	4	2.5	10	2.5	10
	6	30	30	2.5	25
	8	45	42.5	10	37.5

Percent error when detecting apnea when filtration using a reference microphone inside of a stethoscope was used

The open air reference microphone did improve the detection of apnea in most cases but was subject to contamination with a gain G_i of zero. This was especially prominent when the additive noise contained a strong beat such as in rock music.

The reference microphone inside of the stethoscope cup improved apnea detection at every gain for every kind of noise when compared to not filtering at all. The worst percentage rate was a percentage miss rate of 57.5%.

IV. Discussion

Additive Gaussian noise did not affect the apnea detection when there was no filtering performed. This can be explained by the way the breath sounds were detected. The noise floor of the signal was modeled as a Gaussian signal and independent Gaussian signals summed also produce a Gaussian signal. Therefore an additive noise that is close do Gaussian will not affect the apnea detection but may affect the breath detection.

The open air microphone does improve the apnea detection when the amplitude of the additive signal is several times the noise floor. This is because the contamination from the filter is enough to add signals to the output even when the reference input is uncorrelated with the primary input.

The cupped reference microphone produces a reference signal as close as possible to the additive noise. This allows the adaptive filter to not have to adapt as quickly for nonlinearities in the frequency and amplitude of the stethoscope cup. The closer the reference input is to the additive noise signal, the easier it is to filter out.

The cupped reference microphone clearly performed the best and further testing are needed. A future addition to this algorithm will be to only apply the adaptive filter when the signals are correlated and when a signal is detected by the reference microphone.

References

[1] P.Y. Ertel, M. Lawrence, R.K. Brown, and A.M. Stern, "Stethoscope acoustics. II. Transmission and filtration patterns.," *Circulation*, vol. 34, 1966, pp. 899-909.

[2] B. Farhang-Boroujeny, *Adaptive filters*, Wiley, 1998.

[3] S.B. Patel, T.F. Callahan, M.G. Callahan, J.T. Jones, G.P. Graber, K.S. Foster, K. Glifort, and G.R. Wodicka, "An adaptive noise reduction stethoscope for auscultation in high noise environments," *The Journal of the Acoustical Society of America*, vol. 103, May. 1998, pp. 2483-2491.

[4] B. Hill, "Acoustic Breath Detection and Classification: Modeling Respiratory Events," *Rocky Mountain Space Grant Consorcium*, Salt Lake Community College, Salt Lake City, UT: 2008.

[5] V.J. Mathews and G.L. Sicuranza, *Polynomial signal processing*, Wiley, 2000.