

(51) International Patent Classification:  
*G10L 15/20* (2006.01)(21) International Application Number:  
PCT/US2012/020848(22) International Filing Date:  
10 January 2012 (10.01.2012)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

|            |                              |    |
|------------|------------------------------|----|
| 61/431,395 | 10 January 2011 (10.01.2011) | US |
| 61/431,397 | 10 January 2011 (10.01.2011) | US |
| 13/346,719 | 9 January 2012 (09.01.2012)  | US |
| 13/374,746 | 9 January 2012 (09.01.2012)  | US |

(71) Applicant (for all designated States except US): **ALIPH-COM** [US/US]; 99 Rhode Island Street, Third Floor, San Francisco, California 94103 (US).

(72) Inventor; and

(71) Applicant : **JING, Eric Zhinian** [US/US]; 99 Rhode Island Street, Third Floor, San Francisco, California 94103 (US).(74) Agents: **KOKKA & BACKUS, PC** et al.; 703 High Street, Palo Alto, California 94301 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: DYNAMIC ENHANCEMENT OF AUDIO (DAE) IN HEADSET SYSTEMS




FIG. 1A

(57) Abstract: Various embodiments relate to signal processing and, more particularly, to processing of received speech signals to preserve and enhance speech intelligibility. In one embodiment, a communications apparatus includes a receiving path over which received speech signals traverse in an audio stream, and an dynamic audio enhancement device disposed in the receiving path. The dynamic audio enhancement ("DAE") device is configured to modify an amount of volume and an amount of equalization of the audio stream. The DAE device can include a noise level estimator ("NLE") configured to generate a signal representing a noise level estimate. The noise level estimator can include a non-stationary noise detector and a stationary noise detector. The noise level estimator can be configured to generate the signal representing a first noise level estimate based on detection of the non-stationary noise or a second noise level estimate based on detection of the stationary noise.

## DYNAMIC ENHANCEMENT OF AUDIO (DAE) IN HEADSET SYSTEMS FIELD

The disclosure herein relates generally to signal processing and, more particularly, to processing of received speech signals to preserve and enhance speech intelligibility.

### BACKGROUND

Many handset and headset manufacturers have implemented some forms of Dynamic Audio Enhancement (DAE) technology in their headsets. While functional, conventional headsets are not well-suited to properly determine, for example, when user speech is occurring and thus incorrectly estimating the noise level.

### BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1A is a block diagram of a communications device including a receiver implementing a DAE structure and/or functionality, according to some embodiments;

Figure 1B is a block diagram of a DAE device including an NLE, according to some embodiments;

Figure 1C shows higher order and linear interpolation relationships for the gain of an AVC of a DAE device, according to some embodiments;

Figure 1D shows the magnitude responses of the EQ filter of the DAE, according to some embodiments;

Figure 1E shows the resulting speech responses of the DAE with the EQ filter at different levels of noises, according to some embodiments;

Figure 1F depicts the effects of a limiter, according to some embodiments;

Figure 1G includes plots showing the time-domain waveform change due to Limiter of a DAE device, according to some embodiments;

Figure 1H is a plot of the "loudness" measure of the speech in Sones using a psychoacoustic loudness model, according to some embodiments;

Figure 1I is a plot of the intelligibility measure of the speech using the Limiter and the Speech Intelligibility Index (ANSI S3.5-1997) standard, according to some embodiments;

Figure 2A is a block diagram for an example of a voice activity detector (VAD), according to some embodiments;

Figure 2B is a block diagram for another example of a voice activity detector (VAD), under an alternative embodiment;

Figure 2C is a flow diagram for an example of voice activity detection, according to some embodiments;

Figure 3 is an exemplary SSM signal in time (top) and frequency (0-4 kHz, bottom);

Figure 4 is an exemplary normalized autocorrelation function for the SSM signal with speech present;

Figure 5 is an exemplary normalized autocorrelation function for SSM signal with scratch present;

5        Figure 6 is an exemplary flow chart for an example of an autocorrelation algorithm, according to some embodiments;

Figure 7 is a flow chart for an exemplary cross-correlation algorithm, according to some embodiments;

10       Figure 8 is an example of enhanced denoising performance due to the improvement in SSM VAD, according to some embodiments;

Figure 9 shows an exemplary implementation of the scratch resistant VVAD, according to some embodiments;

Figure 10 is a flow chart of an example of a VAD combination algorithm, according to some embodiments;

15       Figure 11 is an example of a two-microphone adaptive noise suppression system, according to some embodiments;

Figure 12 is an exemplary configuration including an array and speech source (S), according to some embodiments;

20       Figure 13 is a block diagram for an exemplary first order gradient microphone using two omnidirectional elements  $O_1$  and  $O_2$ , according to some embodiments;

Figure 14 is a block diagram for a DOMA including an example of two physical microphones configured to form two virtual microphones  $V_1$  and  $V_2$ , according to some embodiments;

25       Figure 15 is a block diagram for an example of a DOMA including two physical microphones configured to form N virtual microphones  $V_1$  through  $V_N$ , where N is any number greater than one, according to some embodiments;

Figure 16 is an example of a headset or head-worn device that includes the DOMA, as described herein, according to some embodiments;

30       Figure 17 is a flow diagram for examples of denoising acoustic signals using the DOMA, according to some embodiments;

Figure 18 is a flow diagram for forming an exemplary DOMA, according to some embodiments;

Figure 19 is a plot of an exemplary linear response of virtual microphone  $V_2$  to a speech source at a first distance, according to some embodiments;

Figure 20 is a plot of an exemplary linear response of virtual microphone  $V_2$  to a noise source at a second distance, according to some embodiments;

Figure 21 is a plot of an exemplary linear response of virtual microphone  $V_1$  to a speech source at the first distance, according to some embodiments;

5        Figure 22 is a plot of an exemplary linear response of virtual microphone  $V_1$  to a noise source at the second distance, according to some embodiments;

Figure 23 is a plot of linear response of virtual microphone  $V_1$  to a speech source at a distance for exemplary frequencies, according to some embodiments;

10       Figure 24 is a plot showing comparison of frequency responses for speech for the array of an embodiment and for an exemplary cardioid microphone;

Figure 25 is a plot showing exemplary speech responses for  $V_1$  and  $V_2$  versus  $B$  with  $d_s$ , according to some embodiments;

Figure 26 is a plot showing an exemplary ratio of  $V_1/V_2$  speech responses versus  $B$ , according to some embodiments;

15       Figure 27 is a plot of exemplary values of  $B$  versus exemplary values of distance of a speech source,  $d_s$ , according to some embodiments;

Figure 28 is a plot of exemplary values of  $B$  versus exemplary values of  $\theta$ , according to some embodiments;

20       Figure 29 is a diagram of plots of exemplary amplitudes and exemplary phase responses of  $N(s)$ , according to some embodiments;

Figure 30 is a diagram of plots of exemplary amplitudes and exemplary phase responses of  $N(s)$  with non-unity  $B$ , according to some embodiments;

Figure 31 is a diagram of plots of exemplary amplitudes and exemplary phase responses of the effect on the speech cancellation in  $V_2$ , according to some embodiments;

25       Figure 32 is a diagram of plots of other exemplary amplitudes and exemplary phase responses of the effect on the speech cancellation in  $V_2$ , according to some embodiments;

Figure 33 shows experimental results for an exemplary array using a simulator in a loud noise environment, according to some embodiments;

30       Figure 34 is an exemplary configuration of a two-microphone array with speech source  $S$ , according to some embodiments;

Figure 35 is an exemplary block diagram of  $V_2$  construction using a fixed  $\beta(z)$ , according to some embodiments;

Figure 36 is an exemplary block diagram of  $V_2$  construction using an adaptive  $\beta(z)$ , according to some embodiments;

Figure 37 is an exemplary block diagram of  $V_1$  construction, according to some embodiments;

Figure 38 is an exemplary flow diagram of acoustic voice activity detection, according to some embodiments;

5        Figure 39 shows experimental results of an exemplary algorithm using a fixed beta when noise is present, according to some embodiments;

Figure 40 shows experimental results of an exemplary algorithm using a fixed beta when speech is present, according to some embodiments;

10       Figure 41 shows experimental results of an exemplary algorithm using a fixed beta when speech and noise is present, according to some embodiments;

Figure 42 shows experimental results of an exemplary algorithm using an adaptive beta when noise is present, according to some embodiments;

Figure 43 shows experimental results of an exemplary algorithm using an adaptive beta when speech is present, according to some embodiments;

15       Figure 44 shows experimental results of an exemplary algorithm using an adaptive beta when speech and noise is present, according to some embodiments;

Figure 45 is an exemplary block diagram of a NAVSAD system, according to some embodiments;

20       Figure 46 is an exemplary block diagram of a PSAD system, according to some embodiments;

Figure 47 is an exemplary block diagram of a denoising system, referred to herein as the Pathfinder system, according to some embodiments;

Figure 48 is an exemplary flow diagram of a detection algorithm for use in detecting voiced and unvoiced speech, according to some embodiments;

25       Figure 49A is a diagram of a plot depicting a received GEMS signal for an utterance along with the mean correlation, according to some embodiments;

Figure 49B is a diagram of a plot depicting a received GEMS signal for an utterance along with the standard deviation of the GEMS signal and the threshold for voiced speech detection, according to some embodiments;

30       Figure 50 is a diagram of a plot depicting voiced speech detected from an utterance along with the GEMS signal and the acoustic noise, according to some embodiments;

Figure 51 is an exemplary microphone array for use according to some embodiments of the PSAD system;

Figure 52 is a plot of  $\Delta M$  versus  $d_1$  for several  $\Delta d$  exemplary values, according to some embodiments;

Figure 53 shows a plot of an exemplary gain parameter as the sum of the absolute values of  $H_1(z)$  and the acoustic data or audio from microphone 1, according to some embodiments;

5 Figure 54 is an example of an alternative plot of acoustic data presented in Figure 53, according to some embodiments;

Figure 55 is a cross section view of an example of an acoustic vibration sensor, according to some embodiments;

10 Figure 56A is an exploded view of an acoustic vibration sensor, according to the example shown in Figure 55;

Figure 56B is perspective view of an acoustic vibration sensor, according to the example shown in Figure 55;

Figure 57 is a schematic diagram of a coupler of an acoustic vibration sensor, under the embodiment of Figure 55;

15 Figure 58 is an exploded view of an example of an acoustic vibration sensor, under an alternative embodiment;

Figure 59 shows representative areas of sensitivity on the human head appropriate for placement of an example of an acoustic vibration sensor, according to some embodiments;

20 Figure 60 is a headset device that includes an exemplary acoustic vibration sensor placed at any of a number of locations, according to some embodiments; and

Figure 61 is a diagram of a manufacturing method for an exemplary acoustic vibration sensor, according to some embodiments.

#### DETAILED DESCRIPTION

Structures and/or functions implementing Dynamic Audio Enhancement (DAE) technology described in detail herein, may be configured to monitor a noise level of received signals (e.g., continuously or predominantly) and may configured to process received far-end speech signals to preserve and/or enhance speech intelligibility. Thus, the “DAE technology” (the term hereinafter includes structures and/or functions that facilitate the DAE technology) is configured to adjust, modify and/or filter signals representing received audio based on  
25 surrounding acoustic noise level and properties. Therefore, a listener (e.g., a callee or recipient of communication) of the received audio can perceive audio (e.g., voiced audio) to discern said the speech of a far-end speaker (e.g., a caller or originator of communication). Environmental acoustic noise can interfere or mask signals representing the received audio, thereby making the incoming speech more difficult to hear and understand. Devices implementing DAE can reduce  
30

or eliminate the necessity for manual volume adjustment or pushing the listening device (e.g., headset) closer to the ear canal to overcome the level of noise. Thus, DAE devices can enhance intelligibility through the use of various techniques, as described herein, to make the speech more understandable. A DAE device can be configured to perform these functions without the user consciously noticing any change (e.g., without user input). A DAE device can implement proprietary Skin Surface Microphone (SSM) and Acoustic Voice Activity Detection (AVAD) technology, as described in detail herein and in the Related Applications, can perform well in a variety of acoustic environments.

In the following description, numerous specific details are introduced to provide a thorough understanding of, and enabling description for, various embodiments of the DAE devices, system and methods. One skilled in the relevant art, however, will recognize that examples described herein can be practiced without one or more of the specific details, or with other components, systems, etc. In other instances, well-known structures or operations are not shown, or are not described in detail, to avoid obscuring aspects of the disclosed embodiments.

Unless otherwise specified, the following terms have the corresponding meanings in addition to any meaning or understanding they may convey to one skilled in the art.

The term “Dynamic Audio Enhancement” or “DAE” refers to a technique that tailors both the volume and the equalization of an audio stream depending on the noise environment, and can refer to a structure or a function, or both, implemented in association with a communication device.

The term “NLE” refers to Noise Level Estimator.

The term “AVC” refers to Automatic Volume Control.

The term “Frame” refers to a speech segment which, unless otherwise specified, is assumed to be approximately 2.5 msec in duration; other frame sizes can be used.

The term “Far-end” refers to the remote environment of the person or device to which the user of the headset or handset is connected.

The term “Handset” refers to a device held in the hand used to capture near-end speech as well as transmit far-end audio information.

The term “Near-end” refers to the environment of the user.

The term “Noise” refers to environmental acoustic noise.

The term “Microphone” refers to a physical omnidirectional microphone, and can refer, at least in some embodiments, to other types of microphones.

The term “Skin Surface Microphone” or “SSM” refers to a device that is used to detect acoustic vibrations inside the user’s skin. An example of a SSM is described in detail herein and

in the Related Applications (see, for example, US Patent Application Number 12/243,718, filed October 1, 2008, which is a continuation of US Patent Number 7,433,484, issued October 7, 2008, and US Patent Application Number 12/772,947, filed May 3, 2010), and is available from Aliph, Inc. ("Aliph") of San Francisco, California.

5           The term "Voice Activity Detection" or "VAD" refers to voice activity detection signals, and the signals can represent or be voiced or unvoiced speech.

          The term "Acoustic VAD" or "AVAD" refers to voice activity detection that is determined using acoustic input and need not rely on vibration sensors such as the SSM or accelerometers. An example of an AVAD is described in detail herein and in the Related  
10       Applications (see, for example, US Patent Application Number 11/805,987, filed May 25, 2007, which is a continuation of US Patent Number 7,246,058, issued June 27, 2007, and US Patent Application Number 12/606,140, filed October 26, 2009, and US Patent Application Number 12/772,947, filed May 3, 2010), and is available from Aliph, Inc. of San Francisco, California.

          The term "RX" refers to receive, received, and/or receiver.

15           The term "TX" refers to transmit, transmitted, and/or transmitter.

          The term "O<sub>1</sub>" refers to an omnidirectional microphone positioned closest to the user's mouth.

          The term "O<sub>2</sub>" refers to an omnidirectional microphone positioned farthest from the user's mouth.

20           There are many characteristics used to evaluate the performance of a DAE system, device and method. For example, the following characteristics may be used to evaluate, but various implementations are not intended to be limited to these characteristics: the behavior of the system at on-set and off-set of noises, including the response rate; the steady state behavior in stationary noise environments; the behavior in non-stationary noise environments;  
25       the behavior in wind noise environments; the robustness to near-end speech and near-end echo effects.

          A response rate to the on-set and off-set of noises may be similar to that of a person attentively adjusting the volume knob. For example, the time to reach the 3-dB point of a change may be within a range of 1 to 3 seconds. In some cases, a response rate includes a  
30       duration in which a DAE device determines whether a noise is impulsive or steady-state. The system need not respond to impulsive noise, unless there is predetermined threshold and requirement to do so (e.g., there can be a significant number of noise impulses in a short time period that triggers a response to impulsive noises).

In “stationary noise” environments, a DAE system may enhance intelligibility using an amount of gain, dynamical range compression (to emphasis on weak speech segments), and/or emphasis of frequency regions that are beneficial to intelligibility. The amount of enhancement may be a function of the noise level. In some examples, a DAE system or device may operate  
5 for noise levels in the range of 60 – 90 dBA SPL.

A “non-stationary noise” can includes acoustics related to a single impulse (e.g., a clap), to repetitive impulses (e.g., jack-hammer), to nearly stationary noise (e.g., multi-talker babble), but are not so limited. A DAE system and device of the various embodiments are configured to behave and respond similarly to a human. For example, a single impulse may not trigger any  
10 change to gain in a DAE system or device. However, repetitive impulses with a period of less than a certain duration (e.g., one second) may negatively affect speech intelligibility, because speech phonemes typically have durations few hundred milliseconds. For these repetitive noises, then, the system may enhance the incoming speech.

“Wind noise” environments very different from acoustic noise environments. Wind  
15 noise can be very impulsive, with rapid repetition, and can also saturate microphones. Therefore, a DAE system or device is configured to compensate for wind noise, as the noise due to wind can reduce intelligibility or render speech substantially unintelligible. A DAE system or device compensate for wind noise by using a wind detector and wind level estimator.

In some embodiments, a DAE system or device can be configured to forego operations  
20 on received speech (e.g., incoming speech from a remote audio source located at a far-end environment). Therefore, a DAE system or device can include a noise meter to use echo-cancelled microphone signals. Further, a DAE system or device can include a VAD so that a noise level estimate will not update during near-end user speech, thereby foregoing a possible mis-categorization of near-end (e.g., user) speech.

25 Once a desired behavior for a DAE system or device is defined, the following metrics can be used to evaluate the steady state performance of the DAE system or device, but the various embodiments are not intended to be limited: speech intelligibility (e.g., how easy voiced audio from the far-end user is to understand); speech quality (e.g., how pleasant the speech is perceived to be); speech loudness (e.g., how loud speech is perceived to be); and the like.  
30 Examples of measuring speech intelligibility include the Speech Intelligibility Index (SII), the Diagnostic Rhyme Test (DRT), the Modified Diagnostic Rhyme Test (MDRT), and the like; examples of measuring speech quality includes the use of PESQ MOS, the MOS, and the like; and examples of measuring speech loudness include psychoacoustics-based Loudness (Sones)).

In some embodiments, a DAE system and device uses all three metrics to combine good intelligibility with pleasant, easy-to-hear speech.

In one technique, intelligibility of an audio stream can be enhanced by boosting the volume gain. A DAE system and device can include a gain coupled with a compressor to keep a speaker output, for example, linear so that echo suppression is optimal. Compression operates to limit an amplitude of the loudspeaker signal to a maximum amplitude, thereby reducing the nonlinear effects of the loudspeaker. This, in turn, boosts weakly-voiced speech and unvoiced speech that provides a significant amount of information constituting speech that otherwise might be masked by noise.

In another technique, a DAE system and device can be configured to vary the frequency content of a speaker (e.g., loudspeaker) signal to increase intelligibility. A DAE system and device can include an equalizer (EQ) configured to “tilt” its equalization functions to, for example, correct for changes in gain with frequency. Thus, higher frequencies of incoming speech can be boosted with respect to the low frequencies.

Other techniques included a calculation of the effect of noise masking in the frequency domain. In this case, an exemplary DAE system or device subdivides, or breaks down, the speech into Bark-scale (or similar psychoacoustic scale) based on subbands of frequencies (e.g., ranges of frequencies). The DAE system or device then can compute the masking curves of the noise in order to determine how much gain or attenuation each subband may receive. For example, the DAE system or device can operate to emphasize the incoming speech subbands that are competing with the noise environment so that they are more easily perceived by the user. In some cases, the DAE system or device can attenuate speech to preserve loudness. The gains on the subbands may be applied so the speech energies in those subbands rise above the noise making level, resulting in audible speech.

An exemplary DAE system or device may provide for an enhanced user experience or a user experience that is unchanged or better than otherwise might be the case. The various examples of DAE systems, devices, and methods described in detail herein provide for an enhanced user experience. For example, the effects and operational changes of the DAE systems, devices, and methods are generally not noticed by the user (e.g., the volume is perceived to sound about the same, regardless of noise). As another example, the DAE systems, devices, and methods do not necessarily operate to address impulsive and/or short duration noise, and can react quickly enough to longer-term changes in noise level. Thus, a user need not manually adjust his or her volume as the maximum volume level is configured to be low enough not to damage hearing and high enough to be effective in various noise environments.

According to various embodiments, DAE systems, devices, and methods can include one or more of the following: a near-end VAD, a far-end VAD, a noise level estimator, and/or a robust algorithmic flow, all of which are described in detail below.

Figure 1A is a block diagram of a communications device including a receiver  
5 implementing a DAE structure and/or functionality, according to some embodiments. The thin lines represent scalar values (i.e., one value per frame), whereas the thicker lines represent a frame of audio data. TX processing components are not shown in the diagram. Also, an echo canceller, a VAD for near-end speech, and a wind processing module may form part of the TX processing components and are not shown. However, at least some of the TX processing  
10 components are shown as inputs to the RX processing path.

A DAE device is depicted as Dynamic Audio Enhancer 102, which resides in a RX processing path. DAE device 102 is configured to receive an echo-free microphone signal (from  $O_2$ ) via path 101 and a near-end VAD signal via path 103 as inputs. As shown, DAE device includes, but is not limited to, NLE 112, AVC 114, Equalizer 116, and Limiter 118. In some  
15 embodiments, NLE 112 can be configured to receive the echo-free microphone signal (from  $O_2$ ) and the near-end VAD signal to generate a signal representing a noise level estimate on path 115. AVC 114 receives the signal representing a noise level estimate and is configured to adjust the gain ("gain 3") as applied to the RX speech path 107 based on the noise level. In some examples, the applied gain onto the signal from AVC 114 can be within 0 to 18 dB. Optionally,  
20 a signal representing a volume gain is applied via path 117 to attenuate or modify the output of AVC 114 to form a signal on path 121a. In turn, the signal on path 121a is processed with a signal on RX speech path 107 to form a final gain signal on path 121b. For example, the final gain for the signal on path 121b can range from -14 to 46 dB.

Signals representing speech on RX speech path 107 is generated by an RX AGC 110,  
25 which has a gain ("gain 2") applied to the speech-related signals. For example, the gain for speech-related signals on path 107 can range from -9 to +9 dB. Speech-related signals on path 107 are derived from the received speech ("RX Speech") 105 passing through a far-end VAD 104 and a noise gate 106 and an equalizer 108. Noise gate 106 can be configured, at least in some embodiments, to gate out or remove a noise, such as static noise (e.g., when no or  
30 insufficient RX speech 105 is present). Equalizer 108 is configured to receive the signal representing the noise level estimate on path 109.

Equalizer 116 can operate as an equalizer filter configured to adapt based on the noise level to boost the high frequencies in noise, and can be configured to receive the signal representing the noise level estimate from path 115 and a signal on path 113a. The signal on

path 113a can be based on the noised-gated speech 105 and a signal on path 111. The signal on path 111 can represent signals representing tones or other audio announcements, such as voice announcements (e.g., power on tones, connect tones, battery life announcement, and other like audio-based functions).

5           Limiter 118 is configured to receive the output of AVC 114 in combination with speech signals on path 107 and a compression gain via path 119. An example of a compression gain is 22 dB. Thus, limiter 118 is configured to dynamically adjust the RX speech level on a per sample basis to reduce or prevent hard clipping due to the gain to be applied. Further, limiter 118 can be configured to reduce the speech dynamic range and enhance the weakly-voiced  
10 sections of the speech. Limiter 118 is configured to “look ahead” at the gain to be applied to the speech, and if clipping is to be expected, then the limiter attenuates 118 the speech enough to prevent the hard clipping. While in some cases the AVC gain change is a slow process, and can be thought of as emulating a person adjusting the volume with a knob, limiter 118 operates much faster, such as operating in the range of milliseconds rather than seconds.

15           Wind generator 120 introduces wind comfort noise that are signals representing wind or wind-like noise that can included in the received signal to alert the user to the presence of wind without unduly affecting the communication intelligibility.

          In addition to the NLE 112, an algorithm implements a near-end VAD (both SSM, examples of which are described in detail herein and in the Related Applications (see, for  
20 example, US Patent Application Number 12/243,718, filed October 1, 2008, which is a continuation of US Patent Number 7,433,484, issued October 7, 2008, and US Patent Application Number 12/772,947, filed May 3, 2010)), an AVAD (see the detailed description herein and in the Related Applications, for example, US Patent Application Number 11/805,987, filed May 25, 2007, which is a continuation of US Patent Number 7,246,058, issued June 27,  
25 2007, and US Patent Application Number 12/772,947, filed May 3, 2010), and a wind algorithm (see the Related Applications, for example, US Patent Application Number 12/772,963, filed May 3, 2010). Furthermore, an implementation may includes a VAD and a wind detector algorithm in the DAE technique.

          Figure 1B is a block diagram of a DAE device, according to some embodiments. As  
30 shown, a DAE device can include NLE 124. A first input to the NLE 124, represented as “WindPresent” signal 125 is the output of a wind detector and/or wind detector algorithm, which indicates whether there is wind currently present or not. If wind is present, then NLE 125 freezes noise estimation updates, and outputs a signal 127 representing the current noise level using the noise level from a frame that precedes the current frame (e.g., a frame that precedes by

two (2) frames (approximately 5 msec)) to ensure that the noise level estimate is not polluted with wind noise.

Another input to NLE 124 of an embodiment, represented as “VadState” signal 129 is the output of the near-end VAD algorithm, which indicates whether there is near-end speech at present. If near-end speech is detected, then NLE 124 freezes noise estimation updates, and outputs signal 127 representing the noise level from before a “guard” period. The “guard” period of an embodiment, can be, for example, approximately 50 msec, as a time period just before the VAD transitions from 0 to 1 (i.e., just before the start of detecting speech). The guard period can be used in an implementation to ensure that near-end speech does not affect or otherwise pollute the noise level estimate.

When there is no wind or near-end speech present, a signal 126 is activated to operate “switch” 128, thereby gating the echo-cancelled  $O_2$  signal as the primary input to NLE 124. In an example of a headset configuration (see the Related Applications, for example, US Patent Application Number 12/139,333, filed June 13, 2008), the  $O_2$  signal is used as it is the microphone closest to the ear and as such provides a slightly more accurate measure of the noise heard by the user. The  $O_2$  signal is to be sufficiently echo cancelled, so that any echo is at least 20 – 30 dB (or more) below the incoming far-end speech.

The  $O_2$  signal of an embodiment is first passed thru an A-weighting-like filter 130 so the noise meter accounts for the human hearing response, and low frequency noise that is not sensitive to human hearing is appropriately de-emphasized. The standard A-weighting, which may be configured for low level noise, need not be used. Instead, a filter that has a response similar to B-weighting can be used.

The filtered  $O_2$  signal is then input into a VAD 131, which can be built using any commonly known techniques, such as an energy based technique. In some embodiments, VAD 131 is built using an energy-verses-threshold technique. VAD 131 outputs a binary scalar value once per frame, which is then filtered using a first-order infinite impulse response (IIR) low-pass filter (“LPF”) 132 (e.g., with time constant of approximately 2.5 seconds). The function of the LPF 132 is to slow down a reaction time of NLE 124 to be on the order of, for example, 2 seconds. A hysteresis module 133 can be configured to apply some “stickiness” of the current state (e.g., the hysteresis can equal 1 when input > 0.35 and remain at 1 until input < 0.1). A binary signal from hysteresis module 133 indicates which of two modes of noise is present—either stationary noise or speech-like non-stationary noise. Speech-like non-stationary noise detector 134 detects the speech-like non-stationary noise, whereas stationary noise detector 136 detects stationary noise.

This bi-modal classification and the resulting processing provides for enhanced DAE techniques. As such, NLE 124 can be described in some embodiments as a bi-modal NLE as it generates noise level estimates differently based on the detect type of noise. When “stationary noise” is detected, the noise level is computed from the energy in a first mode. In some  
 5   embodiments, an energy computation module computes the energy from which the estimated noise level is derived.

When speech-like noise is detected, a peak tracker is used to estimate the noise level in a second mode. The peak tracker works by first finding a largest absolute value in the current frame. Then the 1-per-frame value goes through a fast-attack and slow-delay asymmetric  
 10   exponential averaging filter, governed by, for example, the following equation:

$$y(n) = y(n-1) * \alpha + x * (1.0 - \alpha);$$

*$\alpha = 0.989981$  when  $x > y(n-1)$ , corresponds to a time constant of 0.262 sec,*

*$\alpha = 0.998998$  when  $x < y(n-1)$ , corresponds to a time constant of 2.62 sec.*

After the exponential averaging filter (or any other nonlinear peak tracking filter), NLE 124 uses  
 15   a final low-pass filter (“LPF”) in speech-like non-stationary noise detector 134 to smooth the signal. The final filter is built with an exponential averaging filter:

$$y(n) = y(n-1) * \alpha + x * (1.0 - \alpha);$$

$$\alpha = 0.8990384.$$

In some embodiments, use of the nonlinear peak tracking filter for speech-like noise can lead to  
 20   enhanced and more stable noise level estimates in the presence of nonstationary noise.

The parameters used in the moving average filter, peak tracker and the low pass filters of an embodiment are such that there is a smooth transition between the two modes. That is, for noise that is semi-stationary that may be classified as either speech-like noise or stationary noise, NLE 124 may provide output noise values from either mode. Examples of a moving average  
 25   filter include those filters that average a number of input samples to produce an output sample, as well as FIR filters, or any other like type of filter.

For optimal performance, the near-end user speech and far-end speech echo can be excluded from the noise level estimate. NLE 124 freezes its calculation if near-end speech is detected, and need not determine if the far-end signal is desired speech or undesired echo. Thus, a highly  
 30   accurate VAD and echo cancellation are used for optimal performance, and the SSM VAD and AVAD of an embodiment provide a robust NLE 124.

NLE 124 of various examples can be insensitive to impulsive noises, unless the noise repeats continuously (such as jackhammer noise). Any single burst of noise of duration less than approximately 2-3 seconds is considered an impulsive noise. Tuning of the low pass filter

132 (positioned between the output of the VAD and the input of the hysteresis) and hysteresis are largely responsible for the noise insensitivity of an embodiment.

For speech and other non-stationary noise types which can change energy rapidly, NLE 124 can track the noise level over a longer period than when more stationary noise is detected.

5 This is accomplished by the fast-track and slow decay nature of the peak tracker, which is more suitable than an energy-based approach for non-stationary noise types.

NLE 124 is configured to ramp up to 71% (-3dB) of the true noise level in approximately 2-3 seconds at the onset of noise (both stationary and non-stationary noise). NLE 124 then can ramp down to 71% (-3dB) of the new noise level in less than approximately 1 second at the end of stationary noise, and approximately 2-3 seconds at the end of non-stationary noise. This results in excellent performance; most users need not be aware that the volume and/or equalization of the system may be changing to suit the environment, and the users may perceive the far-end speech in a variety of noise environments without manually adjusting the volume control.

15 Automatic Volume Control (AVC) 114 of FIG. 1A can use the noise level of NLE 124 as input, and can compute the gain to be applied to the RX speech signal (see Figure 1A) to preserve or enhance the perceived loudness and intelligibility.

AVC 114 adjusts from 0 – 18 dB, for NLE noise levels of 60 to 90 dB SPL, using a function of

20 
$$\text{gain\_avc} = (18 - 0) * (\text{noise\_level} - 60 \text{ dB}).$$

The operating range can be based on a noise level that above approximately 60 dB SPL starts to negatively affect a perceived loudness and intelligibility. The maximum gain of can be capped at 18 dB, when noise level is at 90 dB or higher, to protect against possible damage to hearing due too much loudspeaker volume. In some implementations, DAE device 102 may include a higher order interpolation function.

Figure 1C shows higher order and linear interpolation relationships for the gain of an AVC of a DAE device, according to some embodiments;

Also shown is an example of a higher order interpolation curve 141 (e.g., 2<sup>nd</sup> order curve) of the type that may be used in the AVC of an alternative embodiment.

30 Adaptive equalization of the audio to emphasize high frequencies may be applied in an embodiment. As the noise level of is applied to AVC 114 of figure 1A, EQ 116 of figure 1A filter adapts for NLE 112 noise levels of 60 to 90 dB SPL. The EQ filter can be implemented as a finite impulse response (FIR), and can adapts according to the following exemplary function:

For noise\_level ≤ 60 dB SPL):

$$h[k] = h\_base[k]$$

For  $60 \leq \text{noise\_level} \leq 90$  dB SPL :

$$h[k] = h\_base[k] + h\_delta[k] * [(\text{noise\_level}[k] - 60 \text{ dB}) / 30]$$

For  $\text{noise\_level} \geq 90$  dB SPL

5  $h[k] = h\_base[k] + h\_delta[k]$

where  $h\_base[k]$  represents the baseline filter coefficient vector at time  $k$ .

Figure 1D shows magnitude responses of an EQ filter of a DAE device, according to some embodiments. EQ filter coefficients may be computed continuously (or for predominate amounts of time) by interpolating between the 60 dB noise curve, and the 90 dB noise curve, based on the noise level. At noise level of 90 dB SPL, or higher (e.g., as determined by the NLE), the EQ filter has an added tilt of about 12 dB over the frequency range of 0-4000 Hz, with a predominant amount of the tilt occurring in the 2000-3500 Hz range, where information useful for speech intelligibility is located.

Figure 1E shows exemplary resulting speech responses of a DAE device with a EQ filter at different levels of noises, according to at least some embodiments. In this example, a B&K Head and Torso Simulator (HATS) has been used with a Jawbone headset (available from Aliph, Inc.) to generate exemplary measurements under the guidance of specification of ITU-810. The mask of ITU-810 is also shown. The response of the headset at 90 dB noise level falls outside the ITU mask, an acceptable condition because the ITU-810 specification does not require the response to be within the mask for noisy situations. At high noise levels, the response exceeded the upper mask in a range of approximately 2 – 3.4 kHz. Response 152 represents responses less than or equal to 60 dB and response 152 represents responses greater than or equal to 90 dB.

Regarding dynamic range compression, the concept of range compression to enhance unvoiced and weakly voiced speech as a method to enhance intelligibility can be performed. An adaptive compressor may operate to vary a compressor input gain via path 119 of figure 1A as a function of the noise level. When the noise level is high, the input gain is high and the compressor is operating in the nonlinear region. Compression of the high peaks and boosting of low speech levels enhances speech intelligibility.

Limiter 118 of figures 1A can be configured to prevent clipping on the loudspeaker of the Jawbone headset. However, with the proper input adjustment and limiter parameters, the limiter is a powerful tool in improving speech intelligibility. The Limiter looks ahead by introducing a frame of delay. Limiter 118 can be fine-tuned for speech, with fast logarithmic attack and slow geometric decay.

Figure 1F depicts the effects of a limiter, according to some embodiments. The thicker line shows the limiter gain value 156 in relation to the actual time-domain speech signals, according to some embodiments. Limiter 118 of figure 1A can be used to enhance speech by pre-gaining the speech, which has already been normalized to a fixed level by AGC 110, to be just below or near the operating point of Limiter 118. Then any AVC gain would cause Limiter 118 to reduce the strongly-voiced sections of speech, thus amplifying the unvoiced and weakly voiced sections, relative to the voiced sections. In effect, the Limiter operates as a fast adapting dynamic range compressor.

AGC 110 of figure 1A, while not part of the DAE system (Figure 1A), contributes to the dynamic compression effect of the limiter. AGC 110 can remove speech level variability due to the source (e.g., the talker) and the communication channel, thus allowing the limiter to operate near the optimal point.

After AGC 110 stage, the RX speech via path 107 can be normalized to a nominal level. This affords for the determination of the headroom left in the speech signal, before Limiter operations (i.e., clipping starts). This headroom is about 12dB on a communication device, such as a Jawbone. In other words, applying a 12 dB gain to the RX speech signal on path 107 can influence the RX speech level relative to the threshold of clipping. Any additional gain required by AVC 114 would result in the limiter compressing the signal.

Limiter 118 reduces the dynamic range of the speech signal by approximately  $r \cdot \text{avcGain}$ . By maintaining the peaks at the same level as before the Limiter, the weakly and unvoiced speech sections are effectively brought up by  $r \cdot \text{avcGain}$ , where  $r$  is ratio determined in Limiter (for example, 0.9). So, the weak sections of speech can be increased by as much as  $0.9 \cdot 18 \text{ dB} \sim 16 \text{ dB}$ . Since weak sections of speech are most adversely masked by the noise, and they tend to contain relevant speech information, the limiter greatly improves speech intelligibility.

Figures 1G to 1I are plots that illustrate the effect of Limiter 118 on speech. The speech used to produce these plots is the spoken phrase “Star Trek, rated PG 13, press 8”. The utterance of “star trek” is between 49 and 50 seconds. “Rated PG 13” is between 50 and 52 seconds. “Press 8” is between 52 and 53.5 seconds.

Figure 1G includes plots showing the time-domain waveform change due the Limiter of the DAE, according to some embodiments. Plot 160 shows speech waveforms before implementation of limiter 118, and plot 162 show speech waveforms after implementation of limiter 118. Operation of Limiter 118 can be configured to reduce the dynamic range by approximately 10 dB. The unvoiced and weakly-voiced phonemes can be boosted by the limiter. For example, fricative ‘S’ of the word “Star” and voiced stop ‘K’ of the work “Trek”

are significantly boosted. Also, portion 161 of “Rated PG 13” speech is enhanced as portion 163 in plot 162. Without limiter 118, the phrase “Star Trek” may be at least partially masked by noise, potentially making it unintelligible.

Figure 1H is a plot 170 of the “loudness” measure of the speech in Sones, using a psychoacoustic loudness model, according to some embodiments. The psychoacoustic loudness model of an embodiment is the MBSD model, but the embodiment is not so limited (see Wonho, Y., Enhanced modified bark spectral distortion (EMBSD): An objective speech quality measure based on audible distortion and cognition model, Ph. D. thesis, Temple University, Ft. Washington, USA, 1999; see also Moore, B., Glasberg, B., Baer, T., ‘A model for the prediction of thresholds, loudness and partial loudness’, J. Audio Eng. Soc.45, 1997). Alternatively, the Moore and Glas model can provide similar results (see J. Timoney, T. Lysaght, Marc Schoenwiesner, L. McManus, IMPLEMENTING LOUDNESS MODELS IN MATLAB, Proc. Of the 7th Int. Conference on Digital Audio Effects (DAFX-04), Naples, Italy, October 5-8, 2004).

A loudness model can simulate the perceived loudness rating by a human. Three curves are shown in the plot, representing the original speech 173, speech modified by the limiter 171, and original speech with a simple 10 dB gain applied 172. For weak-speech sections, the limiter-processed speech may achieve loudness levels similar to that of the 10 dB gain. While for strong speech sections, the limiter processed speech can be similar to that of the original speech. This behavior is desirable in that the gain in intelligibility is achieved without sounding too loud and/or driving the loudspeaker to distortion.

Figure 1I is a plot 180 of the intelligibility measure of the speech using the Limiter and the Speech Intelligibility Index (ANSI S3.5-1997) standard, according to some embodiments. In this plot, the intelligibility of the Limiter-processed speech can be similar to that of the 10 dB gain version, for weakly-voiced sections. Three curves are shown in the plot, representing the original speech 183, speech modified by the limiter 181, and original speech with a simple 10 dB gain applied 182.

Referring back to Figure 1A, limiter 118 of various embodiments can include an input as a single gain, which is the sum of the AGC 110, AVC 114 and volume setting gains. In addition to its range compression effect for intelligibility enhancement, limiter 118 also reduces or prevents hard clipping. Thus, in a speech path (or a portion thereof) of the DAE, gain may be applied at the limiter, which results in an algorithm with relatively more stability. In a specific embodiment, the limiter provides for the only gain to influence stability of the algorithm.

Masking is a well-understood frequency-based psychoacoustic property that can be exploited to enhance speech intelligibility. MP3 and other modern audio codecs take advantage of the masking effect to reduce bit rate by “hiding” quantization noise in frequency bands that are masked and inaudible. In an embodiment, far-end speech is filtered so that frequencies that would be masked by environmental noise are made clear. When done properly, higher intelligibility may be achieved for the same given loudness, when compared to, for example, increases in gain across the frequency range.

A voice activity detector (VAD) or detection system is described for use in electronic systems. The VAD of an embodiment combines the use of an acoustic VAD and a vibration sensor VAD as appropriate to the environment or conditions in which a user is operating a host device, as described below. An accurate VAD is critical to the noise suppression performance of any noise suppression system, as speech that is not properly detected could be removed, resulting in devoicing. In addition, if speech is improperly thought to be present, noise suppression performance can be reduced. Also, other algorithms such as speech recognition, speaker verification, and others require accurate VAD signals for best performance. Traditional single microphone-based VADs can have high error rates in non-stationary, windy, or loud noise environments, resulting in poor performance of algorithms that depend on an accurate VAD. Any italicized text herein generally refers to the name of a variable in an algorithm described herein.

Figure 2A is a block diagram 200 for an example of a voice activity detector (VAD), according to some embodiments. VAD 202 of an embodiment includes a first detector (“DET 1”) 204 that receives a first signal and a second detector (“DET 2”) 206 that receives a second signal that is different from the first signal. VAD 202 includes a first voice activity detector (VVAD) component 208 coupled to the first detector and the second detector. First VAD component 208 determines that the first signal corresponds to voiced speech when energy resulting from at least one operation on the first signal exceeds a first threshold. VAD 202 includes a second VAD component (“AVAD”) 210 coupled to the second detector. Second VAD component 210 determines that the second signal corresponds to voiced speech when a ratio of a second parameter corresponding to the second signal and a first parameter corresponding to the first signal exceeds a second threshold.

VAD 202 of an embodiment includes a contact detector 212 coupled to the first VAD component and the second VAD component. Contact detector 212 determines a state of contact of the first detector with skin of a user, as described in detail herein.

VAD 202 of an embodiment includes a selector 214 coupled to the first VAD component and the second VAD component. The selector generates a VAD signal to indicate a presence of voiced speech when the first signal corresponds to voiced speech and the state of contact is a first state. Alternatively, selector 214 generates the VAD signal when either of the first signal and the second signal corresponds to voiced speech and the state of contact is a second state.

Figure 2B is a block diagram 250 for another example of a voice activity detector (VAD), according to various embodiments. VAD 252 includes a first detector 254 that receives a first signal and a second detector 256 that receives a second signal that is different from the first signal. Second detector 256 of this alternative embodiment is an acoustic sensor that comprises two omnidirectional microphones (e.g., omnidirectional microphones ("O<sub>1</sub>") 255 and omnidirectional microphone ("O<sub>2</sub>") 257), but the embodiment is not so limited.

VAD 252 of this alternative embodiment includes a first voice activity detector (VVAD) component 258 coupled to the first detector and the second detector. First VAD component 258 determines that the first signal corresponds to voiced speech when energy resulting from at least one operation on the first signal exceeds a first threshold. VAD 252 includes a second VAD component 260 coupled to the second detector. Second VAD component 260 determines that the second signal corresponds to voiced speech when a ratio of a second parameter corresponding to the second signal and a first parameter corresponding to the first signal exceeds a second threshold.

VAD 252 of this alternative embodiment includes a contact detector 262 coupled to the first VAD component and the second VAD component. Contact detector 262 determines a state of contact of the first detector with skin of a user, as described in detail herein.

VAD 252 of this alternative embodiment includes a selector 264 coupled to the first VAD component and the second VAD component and the contact detector. Selector 264 generates a VAD signal to indicate a presence of voiced speech when the first signal corresponds to voiced speech and the state of contact is a first state. Alternatively, the selector generates the VAD signal when either of the first signal and the second signal corresponds to voiced speech and the state of contact is a second state.

Figure 2C is a flow diagram for an example of voice activity detection 290, according to some embodiments. The voice activity detection receives a first signal at a first detector and a second signal at a second detector at 292. The first signal is different from the second signal. The voice activity detection determines the first signal corresponds to voiced speech when energy resulting from at least one operation on the first signal exceeds a first threshold at 294. The voice activity detection determines a state of contact of the first detector with skin of a user

at 296. The voice activity detection determines the second signal corresponds to voiced speech when a ratio of a second parameter corresponding to the second signal and a first parameter corresponding to the first signal exceeds a second threshold at 298. The voice activity detection algorithm generates a voice activity detection (VAD) signal to indicate a presence of voiced  
5 speech when the first signal corresponds to voiced speech and the state of contact is a first state, and/or generates the VAD signal when either of the first signal and the second signal correspond to voiced speech and the state of contact is a second state at 299.

The acoustic VAD (AVAD) algorithm described below (see section "Acoustic Voice Activity Detection (AVAD) Algorithm for use with Electronic Systems" below) uses two  
10 omnidirectional microphones combined in way that significantly increases VAD accuracy over convention one- and two-microphone systems, but it is limited by its acoustic-based architecture and may begin to exhibit degraded performance in loud, impulsive, and/or reflective noise environments. The vibration sensor VAD (VVAD) described below (see section "Detecting Voiced and Unvoiced Speech Using Both Acoustic and Nonacoustic Sensors" and section  
15 "Acoustic Vibration Sensor" below) works very well in almost any noise environment but can exhibit degraded performance if contact with the skin is not maintained or if the speech is very low in energy. In some cases, a VAD may be influenced by movement errors where the vibration sensor moves with respect to the user's skin due to user movement.

A combination of AVAD and VVAD, though, is able to mitigate many of the problems  
20 associated with the individual algorithms. Also, extra processing to remove gross movement errors has significantly increased the accuracy of the combined VAD.

The communications headset example used in this disclosure is the Jawbone Prime Bluetooth headset, produced by AliphCom in San Francisco, Calif. This headset uses two omnidirectional microphones to form two virtual microphones using the system described below  
25 (see section "Dual Omnidirectional Microphone Array (DOMA)" below) as well as a third vibration sensor to detect human speech inside the cheek on the face of the user. Although the cheek location is preferred, any sensor that is capable of detecting vibrations reliably (such is an accelerometer or radiovibration detector (see section "Detecting Voiced and Unvoiced Speech Using Both Acoustic and Nonacoustic Sensors" below) can be used as well.

30 Unless specifically stated, the following acronyms and terms are defined as follows.

Denoising is the removal of unwanted noise from an electronic signal.

Devoicing is the removal of desired speech from an electronic signal.

False Negative is a VAD error when the VAD indicates that speech is not present when speech is present.

False Positive is a VAD error when the VAD indicates that speech is present when speech is not present.

Microphone is a physical acoustic sensing element.

5 Normalized Least Mean Square (NLMS) adaptive filter is a common adaptive filter used to determine correlation between the microphone signals. Any similar adaptive filter may be used.

The term  $O_1$  represents the first physical omnidirectional microphone

The term  $O_2$  represents the second physical omnidirectional microphone

10 Skin Surface Microphone (SSM) is a microphone adapted to detect human speech on the surface of the skin (see section "Acoustic Vibration Sensor" below). Any similar sensor that is capable of detecting speech vibrations in the skin of the user can be substituted.

Voice Activity Detection (VAD) signal is a signal that contains information regarding the location in time of voiced and/or unvoiced speech.

15 Virtual microphone is a microphone signal comprised of combinations of physical microphone signals.

The VVAD of an embodiment uses the Skin Surface Microphone (SSM) produced by AliphCom, based in San Francisco, Calif. The SSM is an acoustic microphone modified to enable it to respond to vibrations in the cheek of a user (see section "Acoustic Vibration Sensor" below) rather than airborne acoustic sources. Any similar sensor that responds to vibrations  
20 (such as an accelerometer or radiovibrometer (see section "Detecting Voiced and Unvoiced Speech Using Both Acoustic and Nonacoustic Sensors" below)) can also be used. These sensors allow accurate detection of user speech even in the presence of loud environmental acoustic noise, but are susceptible to false positives due to gross movement of the sensor with respect to the user. These non-speech movements (generally referred to a "scratches" below) can be  
25 generated when the user walks, chews, or is physically located in a vibrating space such a car or train. The algorithms below limit the occurrences of false positives due to these movements.

Figure 3 is an exemplary SSM signal 302 in time (top) and frequency (0-4 kHz, bottom). Figure 4 is an exemplary normalized autocorrelation function for the SSM signal with speech present. Figure 5 is an exemplary normalized autocorrelation function for SSM signal with  
30 scratch present.

An energy based algorithm has been used for the SSM VAD (see section "Detecting Voiced and Unvoiced Speech Using Both Acoustic and Nonacoustic Sensors" below). It worked quite well in most noise environments, but could have performance issues with non-speech scratches resulting in false positives. These false positives reduced the effectiveness of

the noise suppression and a way was sought to minimize them. The result is that the SSM VAD of an embodiment uses a non-energy based method since scratches often generate more SSM signal energy than speech does.

The SSM VAD decision of an embodiment is computed in two steps. The first is the existing energy-based decision technique. When the energy-based technique determines there is speech present is the second step applied in an attempt to reduce false positives.

Before examining the algorithms used to reduce false positives, the following description presents a review of the properties of the SSM and similar vibration sensor signals that operate on the cheek of the user. One property of the SSM and similar vibration sensor signals is that sensor signals for voiced speech are detectable but can be very weak; unvoiced speech is typically too weak to be detected. Another property of an exemplary SSM and similar vibration sensor signals is that they are effectively low-pass filtered, and may have various amounts of energy (e.g., some to significant amounts of energy) below 600-700 Hz. A further property of the SSM and similar vibration sensor signals is that they vary significantly from person to person as well as phoneme to phoneme. Yet another property of the SSM and similar vibration sensor signals is that the relationship between the strength of the sensor signal and the acoustically recorded speech signal is normally inverse--high energy vibration sensor signals correspond to a significant amount of energy inside the mouth of the user (such as an "ee") and a low amount of radiated acoustic energy. In the same manner, low energy vibration sensor signals correlate with high energy acoustic output.

Two main classes of algorithms are used in an embodiment to differentiate between speech signals and "scratch" signals: Pitch detection of the SSM signal and cross-correlation of SSM signal with microphone signal(s). Pitch detection is used because the voiced speech detected by the SSM always has a fundamental and harmonics present, and cross-correlation is used to ensure that speech is being produced by the user. Cross-correlation alone is insufficient as there can be other speech sources in the environment with similar spectral properties.

Pitch detection can simply and effectively implemented by computing the normalized autocorrelation function, finding the peak of it, and comparing it to a threshold.

The autocorrelation sequence used in an embodiment for a window of size  $N$  is:

$$R_k = \sum_{i=0}^{N-1-k} S_i S_{i+k} e^{-i/\tau}$$

where  $i$  is the sample in the window,  $S$  is the SSM signal, and  $e^{-i/\tau}$  (the exponential decay factor) is applied to provide faster onset of the detection of a speech frame and a smoothing effect. Also,  $k$  is the lag, and is computed for the range of 20 to 120 samples, corresponding to pitch

frequency range of 400 Hz to 67 Hz. The window size used in computing the autocorrelation function is a fixed size of  $2 \times 120 = 240$  samples. This is to ensure that there are at least two complete periods of the wave in the computation.

In actual implementation, to reduce MIPS, the SSM signal is first downsampled by a factor of 4 from 8 kHz to 2 kHz. This is acceptable because the SSM signal has little useful speech energy above 1 kHz. This means that the range of  $k$  can be reduced to 5 to 30 samples, and the window size is  $2 \times 30 = 60$  samples. This still covers the range from 67 to 400 Hz.

Figure 6 shows a flow chart 600 for an example of an autocorrelation algorithm, according to some embodiments. The data in a history buffer 606 may be applied with an exponential gain 608 and delayed 611, and then a new frame 602 of down-sampled 604 (e.g., by four) SSM signal gets stored 610 in it.  $R(0)$  614 is calculated once during the current frame.  $R(k)$  616 may be calculated for the range of lags. The maximum  $R(k)$  620 is then compared to  $T \times R(0)$ , and if it is greater than  $T \times R(0)$ , then the current frame is denoted as containing speech.

Cross-correlation of the sensor signal with the microphone signal(s) is also very useful, since the microphone signal will not contain a scratch signal. However, detailed examination shows that there are multiple challenges with this method.

The microphone signal and the SSM signal need not necessarily be synchronized, and thus time alignment of the signals may be performed, according to some examples.  $O_1$  or  $O_2$  are susceptible to acoustic noise which is not present in the SSM signal, thus in low SNR environments, the signals may have a low correlation value even when speech is present. Also, environmental noise may contain speech elements that correlate with the SSM signal. However, the autocorrelation has been shown to be useful in reducing false positives.

Figure 7 shows the flow chart 700 of an exemplary cross-correlation algorithm, according to some embodiments. The  $O_1$  signal 702 and  $O_2$  signal 704 first pass through a noise-suppressor 706 (NS, it may be single channel or dual-channel noise suppression) and are then low-pass filtered (LPF) 708 to make the speech signal to look similar to a SSM signal 710. LPF 708 can model the static response of the SSM signal 710, both in magnitude and phase response. Then the speech signal may be filtered by an adaptive filter (H) 712 that models the dynamic response of the SSM signal when speech is present. The error residual on path 715 drives the adaptation of the filter, and the adaptation may take place when the AVAD detects speech. When speech dominates the SSM signal, the residual energy 716 may be small. When scratch dominates the SSM signal, the residual energy 716 may be large.

Figure 8 shows an exemplary effect of scratch resistant VVAD on noise suppression performance, according to some embodiments. That is, Figure 8 depicts an example of the improved denoising performance due to the improvement in SSM VAD. Top figure 802 shows that the noise suppression system having trouble denoising well due to the false positives of the original VVAD, because it is triggering on scratch due to chewing gum. Bottom figure 852 shows the same noise suppression system, with the improved scratch resistant VVAD implemented. The denoising performance may be enhanced, at least in some examples, because a VVAD need not trigger on scratch and thus allowing the denoising system to adapt and remove noise.

Figure 9 shows an exemplary implementation of the scratch resistant VVAD, according to some embodiments. Figure 9 shows the VVAD (solid black line), the adaptive threshold (dashed black line), and the SSM energy (dashed gray line) during periods of speech (as detected), scratch noise due to moving the SSM across a face (as ignored except for a single frame), and scratch noise due to walking (as ignored), according to some embodiments. Solid black line 902 in the figure is an indicator of the output of the VVAD, dashed black line 904 is the adaptive energy threshold, and dashed gray line 906 is the energy of the SSM signal. In this embodiment, to be classified as speech using energy, the energy of the SSM may be more than the adaptive energy threshold. Note how the system correctly identifies the speech segment, but rejects all but a single window of the scratch noise segments, even though most of the scratch energy is well above the adaptive energy threshold. Without the improvements in the VAD algorithm as described herein, many of the high-energy scratch SSM signals may have generated false positive indications, reducing the ability of the system to remove environmental acoustic noise. Thus, this algorithm may significantly reduce or eliminate the number of false positives associated with non-speech vibration sensor signals without significantly affecting the ability of system to identify speech.

In some examples, the combined VAD algorithm can influence and/or contribute to the VAD selection process. For example, neither the AVAD nor the VVAD can be relied upon (e.g., either AVAD or VVAD can be emphasized or deemphasized in structure and/or functionality), so care may be taken to select the combination that is most likely to be correct.

The combination of the AVAD and VVAD of an embodiment is an "OR" combination - if either VVAD or AVAD indicates that the user is producing speech, then the VAD state is set to TRUE. While effective in reducing false negatives, this increases false positives. This is especially true for the AVAD, which is more susceptible to false positive errors, especially in high noise and reflective environments.

To reduce false positive errors, it is useful to attempt to determine how well the SSM is in contact with the skin. In some examples, if there is good contact and the SSM is reliable, then the VVAD may be used. If there is not good contact, then the "OR" combination above may be beneficial, in other examples.

5       Without a dedicated (hardware) contact sensor, there is no simple way to know in real-time that whether the SSM contact is good or not. The method below uses a conservative version of the AVAD, and whenever the conservative AVAD (CAVAD) detects speech it compares its VAD to the SSM VAD output. If the SSM VAD also detects speech consistently when CAVAD triggers, then SSM contact is determined to be good. Conservative means the  
10   AVAD is unlikely to falsely trigger (false-positive) due to noise, but may be very prone to false negatives to speech. The AVAD works by comparing the V1/V2 ratio against a threshold, and AVAD is set to TRUE whenever V1/V2 is greater than the threshold (e.g., approximately 3-6 dB). The CAVAD has a relatively higher (for example, 9+ dB) threshold. At this level, it is extremely unlikely to return false positives but sensitive enough to trigger on speech a  
15   significant percentage of the time. It is possible to set this up practically because of the very large dynamic range of the V1/V2 ratio given by the DOMA technique.

      However, if the AVAD is not functioning properly for some reason, this technique can fail and render the algorithm (and the headset) useless. So, the conservative AVAD is also compared to the VVAD to see if the AVAD is working. Figure 10 is a flow chart of an example  
20   of a VAD combination algorithm, according to some embodiments. The details of this algorithm are shown in Figure 10, where the *SSM\_contact\_state* is the final output. It takes one of the three values: GOOD, POOR or INDETERMINATE. If GOOD, the AVAD output is ignored. If POOR or INDETERMINATE, it is used in the "OR" combination with the VVAD as described above.

25       Several improvements to the VAD system of a headset that uses dual omnidirectional microphones and a vibration sensor have been described herein. False positives caused by large-energy spurious sensor signals due to relative non-speech movement between the headset and face have been reduced by using both the autocorrelation of the sensor signal and the cross-correlation between the sensor signal and one or both of the microphone signals. False positives  
30   caused by the "OR" combination of the acoustic microphone-based VAD and the sensor VAD have been reduced by testing the performance of each against the other and adjusting the combination depending on which is the more reliable sensor.

      A dual omnidirectional microphone array (DOMA) that provides improved noise suppression is described herein. Compared to conventional arrays and algorithms, which seek to

reduce noise by nulling out noise sources, the array of an embodiment is used to form two distinct virtual directional microphones that are configured to have very similar noise responses and very dissimilar speech responses. The null formed by the DOMA is one used to remove the speech of the user from  $V_2$ . The two virtual microphones of an embodiment can be paired with an adaptive filter algorithm and/or VAD algorithm to significantly reduce the noise without distorting the speech, significantly improving the SNR of the desired speech over conventional noise suppression systems. The embodiments described herein are stable in operation, flexible with respect to virtual microphone pattern choice, and have proven to be robust with respect to speech source-to-array distance and orientation as well as temperature and calibration techniques.

In the following description, numerous specific details are introduced to provide a thorough understanding of, and enabling description for, embodiments of the DOMA. One skilled in the relevant art, however, will recognize that these embodiments can be practiced without one or more of the specific details, or with other components, systems, etc. In other instances, well-known structures or operations are not shown, or are not described in detail, to avoid obscuring aspects of the disclosed embodiments.

Unless otherwise specified, the following terms have the corresponding meanings in addition to any meaning or understanding they may convey to one skilled in the art.

The term “bleedthrough” means the undesired presence of noise during speech.

The term “denoising” means removing unwanted noise from Mic1, and refers to the amount of reduction of noise energy in a signal in decibels (dB).

The term “devoicing” means removing/distorting the desired speech from Mic1.

The term “directional microphone (DM)” means a physical directional microphone that is vented on both sides of the sensing diaphragm.

The term “Mic1 (M1)” means a general designation for an adaptive noise suppression system microphone that usually contains more speech than noise.

The term “Mic2 (M2)” means a general designation for an adaptive noise suppression system microphone that usually contains more noise than speech.

The term “noise” means unwanted environmental acoustic noise.

The term “null” means a zero or minima in the spatial response of a physical or virtual directional microphone.

The term “ $O_1$ ” means a first physical omnidirectional microphone used to form a microphone array.

The term “O<sub>2</sub>” means a second physical omnidirectional microphone used to form a microphone array.

The term “speech” means desired speech of the user.

The term “Skin Surface Microphone (SSM)” is a microphone used in an earpiece (e.g.,  
5 the Jawbone earpiece available from Aliph of San Francisco, Calif.) to detect speech vibrations on the user’s skin.

The term “V<sub>1</sub>” means the virtual directional “speech” microphone, which has no nulls.

The term “V<sub>2</sub>” means the virtual directional “noise” microphone, which has a null for the user’s speech.

10 The term “Voice Activity Detection (VAD) signal” means a signal indicating when user speech is detected.

The term “virtual microphones (VM)” or “virtual directional microphones” means a microphone constructed using two or more omnidirectional microphones and associated signal processing.

15 Figure 11 is an example of a two-microphone adaptive noise suppression system 1100, according to some embodiments. The two-microphone system 1100 including the combination of physical microphones (“MIC 1”) 1102 and (“MIC 2”) 1103 along with the processing or circuitry components to which the microphones couple (described in detail below, but not shown in this figure) is referred to herein as the dual omnidirectional microphone array (DOMA) 1110,  
20 but the embodiment is not so limited. Referring to FIG. 11, in analyzing the single noise source 1101b and the direct path to the microphones, the total acoustic information coming into MIC 1 (1102, which can be a physical or virtual microphone) is denoted by  $m_{\text{sub.1}}(n)$ . The total acoustic information coming into MIC 2 (1103, which can also be a physical or virtual microphone) is similarly labeled  $m_2(n)$ . In the  $z$  (digital frequency) domain, these are  
25 represented as  $M_1(z)$  and  $M_2(z)$ . Then,

$$M_1(z) = S(z) + N_2(z)$$

$$M_2(z) = N(z) + S_2(z)$$

with

$$N_2(z) = N(z)H_1(z)$$

$$30 \quad S_2(z) = S(z)H_2(z),$$

so that

$$M_1(z) = S(z) + N(z)H_1(z)$$

$$M_2(z) = N(z) + S(z)H_2(z).$$

Eq. 1

This is the general case for all two microphone systems. Equation 1 has four unknowns and two known relationships and therefore cannot be solved explicitly.

However, there is another way to solve for some of the unknowns in Equation 1. The analysis starts with an examination of the case where the speech is not being generated, that is, where a signal from the VAD subsystem 1104 (optional) equals zero. In this case,  $s(n)=S(z)=0$ , and Equation 1 reduces to

$$M_{1N}(z) = N(z)H_1(z)$$

$$M_{2N}(z) = N(z) ,$$

where the N subscript on the M variables indicate that noise is being received. This leads to

$$\begin{aligned} M_{1N}(z) &= M_{2N}(z)H_1(z) \\ H_1(z) &= \frac{M_{1N}(z)}{M_{2N}(z)} \end{aligned} \quad \text{Eq. 2}$$

The function  $H_1(z)$  1122 can be calculated using any of the available system identification algorithms and the microphone outputs when the system is certain that noise is being received. The calculation can be done adaptively, so that the system can react to changes in the noise.

A solution is now available for  $H_1(z)$ , one of the unknowns in Equation 1. The final unknown,  $H_2(z)$  1120, can be determined by using the instances where speech is being produced and the VAD equals one. When this is occurring, but the recent (perhaps less than 1 second) history of the microphones indicate low levels of noise, it can be assumed that  $n(s) = N(z) \sim 0$ . Then Equation 1 reduces to

$$\begin{aligned} M_{1S}(z) &= S(z) \\ M_{2S}(z) &= S(z)H_2(z) , \end{aligned}$$

which, in turn, leads to:

$$\begin{aligned} M_{2S}(z) &= M_{1S}(z)H_2(z) \\ H_2(z) &= \frac{M_{2S}(z)}{M_{1S}(z)} , \end{aligned}$$

which is the inverse of the  $H_1(z)$  1122 calculation. However, it is noted that different inputs are being used (now the speech is occurring whereas before the noise was occurring). While calculating  $H_2(z)$  1120, the values calculated for  $H_1(z)$  are held constant (and vice versa) and it is assumed that the noise level is not high enough to cause errors in the  $H_2(z)$  calculation.

After calculating  $H_1(z)$  1122 and  $H_2(z)$  1120, they are used to remove the noise from the signal. If Equation 1 is rewritten as

$$\begin{aligned} S(z) &= M_1(z) - N(z)H_1(z) \\ N(z) &= M_2(z) - S(z)H_2(z) \\ S(z) &= M_1(z) - [M_2(z) - S(z)H_2(z)]H_1(z) \end{aligned}$$

$$S(z)[1 - H_2(z)H_1(z)] = M_1(z) - M_2(z)H_1(z),$$

then  $N(z)$  may be substituted as shown to solve for  $S(z)$  as

$$S(z) = \frac{M_1(z) - M_2(z)H_1(z)}{1 - H_1(z)H_2(z)}. \quad \text{Eq. 3}$$

If the transfer functions  $H_1(z)$  1122 and  $H_2(z)$  1120 can be described with sufficient accuracy, then the noise can be completely removed and the original signal recovered. This remains true without respect to the amplitude or spectral characteristics of the noise. If there is very little or no leakage from the speech source into  $M_2$ , then  $H_2(z) \approx 0$  and Equation 3 reduces to

$$S(z) \approx M_1(z) - M_2(z)H_1(z). \quad \text{Eq. 4}$$

Equation 4 is much simpler to implement and is very stable, assuming  $H_1(z)$  is stable. However, if significant speech energy is in  $M_2(z)$ , devoicing can occur. In order to construct a well-performing system and use Equation 4, consideration is given to the following conditions:

- R1. Availability of a perfect (or at least very good) VAD in noisy conditions
- R2. Sufficiently accurate  $H_1(z)$
- R3. Very small (ideally zero)  $H_2(z)$ .
- R4. During speech production,  $H_1(z)$  cannot change substantially.
- R5. During noise,  $H_2(z)$  cannot change substantially.

Condition R1 is easy to satisfy if the SNR of the desired speech to the unwanted noise is high enough. "Enough" means different things depending on the method of VAD generation. If a VAD vibration sensor is used, as in Burnett 7,256,048, accurate VAD in very low SNRs (-10 dB or less) is possible. Acoustic-related methods using information from  $O_1$  and  $O_2$  can also return accurate VADs, but are limited to SNRs of -3 dB or greater for adequate performance.

Condition R5 is normally simple to satisfy because for most applications the microphones will not change position with respect to the user's mouth very often or rapidly. In those applications where it may happen (such as hands-free conferencing systems) it can be satisfied by configuring Mic2 1103 so that  $H_2(z) \approx 0$ .

Satisfying conditions R2, R3, and R4 are more difficult but are possible given the right combination of  $V_1$  and  $V_2$ . Methods are examined below that have proven to be effective in satisfying the above, resulting in excellent noise suppression performance and minimal speech removal and distortion in an embodiment.

The DOMA, in various embodiments, can be used with the Pathfinder system as the adaptive filter system or noise removal. The Pathfinder system, available from AliphCom, San Francisco, Calif., is described in detail in other patents and patent applications referenced herein.

Alternatively, any adaptive filter or noise removal algorithm can be used with the DOMA in one or more various alternative embodiments or configurations

When the DOMA is used with the Pathfinder system, the Pathfinder system generally provides adaptive noise cancellation by combining the two microphone signals (e.g., Mic1, Mic2) by filtering and summing in the time domain. The adaptive filter generally uses the signal received from a first microphone of the DOMA to remove noise from the speech received from at least one other microphone of the DOMA, which relies on a slowly varying linear transfer function between the two microphones for sources of noise. Following processing of the two channels of the DOMA, an output signal is generated in which the noise content is attenuated with respect to the speech content, as described in detail below.

Figure 12 is an exemplary two-microphone array (DOMA) including an array of omnidirectional microphone ("O<sub>2</sub>") 1202 and omnidirectional microphone ("O<sub>1</sub>") 1201, and speech source S configuration 1200, according to some embodiments. The microphones can be separated by a distance approximately equal to  $2d_0$ , and the speech source is located a distance,  $d_s$ , away from a midpoint of the array at an angle  $\theta$ . The system may be axially symmetric so  $d_s$  and  $\theta$  may be specified.

Figure 13 is a system 1300 for generating, producing or emulating an exemplary first order gradient microphone V using two omnidirectional elements ("O<sub>1</sub>") 1201 and O<sub>2</sub> 1202, according to some embodiments. The array of an embodiment includes two physical microphones 1201 and 1202 (e.g., omnidirectional microphones) placed a distance  $2d_0$  apart and a speech source 1200 is located a distance  $d_s$  away from a reference point 1210 at an angle  $\theta$ . This array is axially symmetric (at least in free space), so no other angle is needed. The output from each microphone 1201 and 1202 can be delayed by delay ("z<sub>1</sub>") 1304 and delay ("z<sub>2</sub>") 1306, multiplied by a gain ("A<sub>1</sub>") 1308 and again ("A<sub>2</sub>") 1310, and then summed with the other as demonstrated in Figure 13. The output of the array is or forms at least one virtual microphone, as described in detail below. This operation can be over any frequency range desired. By varying the magnitude and sign of the delays and gains, a wide variety of virtual microphones (VMs), also referred to herein as virtual directional microphones, can be realized, at least in some embodiments. There are other methods known to those skilled in the art for constructing VMs but this is a common one and will be used in the enablement below.

As an example, Figure 14 is a block diagram for a system 1400 (e.g., a DOMA system) including an example of two physical microphones configured to form two virtual microphones V<sub>1</sub> 1410 and V<sub>2</sub> 1412, according to some embodiments. DOMA 1400 includes two first order gradient microphones V<sub>1</sub> and V<sub>2</sub> formed using the outputs of two microphones or elements O<sub>1</sub>

and  $O_2$  (1201 and 1202), according to some embodiments. DOMA 1400 of an embodiment includes two physical microphones (" $O_1$ ") 1201 and (" $O_2$ ") 1202 that can be omnidirectional microphones, as described above with reference to Figures 12 and 13. The output from each microphone is coupled to a processing component 1402, or circuitry, and the processing component outputs signals representing or corresponding to the virtual microphones  $V_1$  and  $V_2$ .

In this example system 1400, the output of physical microphone 1201 is coupled to processing component 1402 that includes a first processing path that includes application of a first delay (" $z_{11}$ ") 1304a and a first gain (" $A_{11}$ ") 1308a and a second processing path that includes application of a second delay (" $z_{12}$ ") 1304b and a second gain (" $A_{12}$ ") 1306b. The output of physical microphone 1202 is coupled to a third processing path of the processing component 1402 that includes application of a third delay (" $z_{21}$ ") 1306a and a third gain (" $A_{21}$ ") 1310a and a fourth processing path that includes application of a fourth delay (" $z_{22}$ ") 1306b and a fourth gain (" $A_{22}$ ") 1316b. The output of the first and third processing paths is summed to form virtual microphone (" $V_1$ ") 1410, and the output of the second and fourth processing paths is summed to form virtual microphone (" $V_2$ ") 1412.

As described in detail below, varying the magnitude and sign of the delays and gains of the processing paths leads to a wide variety of virtual microphones (VMs), also referred to herein as virtual directional microphones, can be realized. While the processing component 1402 described in this example includes four processing paths generating two virtual microphones or microphone signals, the embodiment is not so limited.

For example, Figure 15 is a block diagram for an example of a DOMA 1500 including two physical microphones 1201 and 1202 configured to form  $N$  virtual microphones  $V_1$  through  $V_N$ , where  $N$  is any number greater than one, according to some embodiments. Thus, the DOMA can include a processing component 1502 having any number of processing paths as appropriate to form a number  $N$  of virtual microphones.

The DOMA of an embodiment can be coupled or connected to one or more remote devices. In a system configuration, the DOMA outputs signals to the remote devices. The remote devices include, but are not limited to, at least one of cellular telephones, satellite telephones, portable telephones, wireline telephones, Internet telephones, wireless transceivers, wireless communication radios, personal digital assistants (PDAs), personal computers (PCs), headset devices, head-worn devices, and earpieces.

Furthermore, the DOMA of an embodiment can be a component or subsystem integrated with a host device. In this system configuration, the DOMA outputs signals to components or subsystems of the host device. The host device includes, but is not limited to, at least one of

cellular telephones, satellite telephones, portable telephones, wireline telephones, Internet telephones, wireless transceivers, wireless communication radios, personal digital assistants (PDAs), personal computers (PCs), headset devices, head-worn devices, and earpieces.

As an example, Figure 16 is an example of a headset or head-worn device 1600 that includes the DOMA, as described herein, according to some embodiments. The headset 1600 of an embodiment includes a housing having two areas or receptacles (not shown) that receive and hold two microphones (e.g.,  $O_1$  and  $O_2$ ). The headset 1600 is generally a device that can be worn by a speaker 1602, for example, a headset or earpiece that positions or holds the microphones in the vicinity of the speaker's mouth. The headset 1600 of an embodiment places a first physical microphone (e.g., physical microphone  $O_1$ ) in a vicinity of a speaker's lips. A second physical microphone (e.g., physical microphone  $O_2$ ) is placed a distance behind the first physical microphone. In some examples, the distance of may be in a range of a few centimeters behind the first physical microphone or as described herein (e.g., described with reference to FIGS. 11-15). The DOMA is symmetric and is used in the same configuration or manner as a single close-talk microphone, but is not so limited.

Figure 17 is a flow diagram for examples of denoising 1700 acoustic signals using the DOMA, according to some embodiments. The denoising 1700 begins by receiving 1702 acoustic signals at a first physical microphone and a second physical microphone. In response to the acoustic signals, a first microphone signal is output from the first physical microphone and a second microphone signal is output from the second physical microphone 1704. A first virtual microphone is formed 1706 by generating a first combination of the first microphone signal and the second microphone signal. A second virtual microphone is formed 1708 by generating a second combination of the first microphone signal and the second microphone signal, and the second combination is different from the first combination. The first virtual microphone and the second virtual microphone are distinct virtual directional microphones with substantially similar responses to noise and substantially dissimilar responses to speech. The denoising 1700 generates 1710 output signals by combining signals from the first virtual microphone and the second virtual microphone, and the output signals include less acoustic noise than the acoustic signals.

Figure 18 is a flow diagram for forming 1800 an exemplary DOMA, according to some embodiments. Formation 1800 of the DOMA includes forming 1802 a physical microphone array including a first physical microphone and a second physical microphone. The first physical microphone outputs a first microphone signal and the second physical microphone outputs a second microphone signal. A virtual microphone array is formed 1804 comprising a

first virtual microphone and a second virtual microphone. The first virtual microphone comprises a first combination of the first microphone signal and the second microphone signal. The second virtual microphone comprises a second combination of the first microphone signal and the second microphone signal, and the second combination is different from the first combination. The virtual microphone array including a single null oriented in a direction toward a source of speech of a human speaker.

The construction of VMs for the adaptive noise suppression system of an embodiment includes substantially similar noise response in  $V_1$  and  $V_2$ . Substantially similar noise response as used herein means that  $H_1(z)$  is simple to model and will not change much during speech, satisfying conditions R2 and R4 described above and allowing strong denoising and minimized bleedthrough.

The construction of VMs for the adaptive noise suppression system of an embodiment includes relatively small speech response for  $V_2$ . The relatively small speech response for  $V_2$  means that  $H_2(z) \approx 0$ , which will satisfy conditions R3 and R5 described above.

The construction of VMs for the adaptive noise suppression system of an embodiment further includes sufficient speech response for  $V_1$  so that the cleaned speech will have significantly higher SNR than the original speech captured by  $O_1$ .

The description that follows assumes that the responses of the omnidirectional microphones  $O_1$  and  $O_2$  to an identical acoustic source have been normalized so that they have exactly the same response (amplitude and phase) to that source. This can be accomplished using standard microphone array methods (such as frequency-based calibration) well known to those versed in the art.

Referring to the condition that construction of VMs for the adaptive noise suppression system of an embodiment includes relatively small speech response for  $V_2$ , it is seen that for discrete systems  $V_2(z)$  can be represented as:

$$V_2(z) = O_2(z) - z^\gamma \beta O_1(z)$$

where

$$B = \frac{d_1}{d_2}$$

$$\gamma = \frac{d_2 - d_1}{c} \cdot f_s \quad (\text{samples})$$

$$d_1 = \sqrt{d_s^2 - 2d_s d_0 \cos(\theta) + d_0^2}$$

$$d_2 = \sqrt{d_s^2 + 2d_s d_0 \cos(\theta) + d_0^2}$$

The distances  $d_1$  and  $d_2$  are the distance from  $O_1$  and  $O_2$  to the speech source (see Figure 12), respectively, and  $\gamma$  is their difference divided by  $c$ , the speed of sound, and multiplied by the

sampling frequency  $f_s$ . Thus  $\gamma$  is in samples, but need not be an integer. For non-integer  $\gamma$ , fractional-delay filters (well known to those versed in the art) may be used.

Note that the  $\beta$  above is not the conventional  $\beta$  used to denote the mixing of VMs in adaptive beamforming; it is a physical variable of the system that depends on the intra-  
 5 microphone distance  $d_0$  (which is fixed) and the distance  $d_s$  and angle  $\theta$ , which can vary. As shown below, for properly calibrated microphones, it is not necessary for the system to be programmed with the exact  $\beta$  of the array. Errors of approximately 10-15% in the actual  $\beta$ .  
 (i.e. the  $\beta$  used by the algorithm is not the  $\beta$  of the physical array) have been used with very little degradation in quality. The algorithmic value of  $\beta$  may be calculated and set for a  
 10 particular user or may be calculated adaptively during speech production when little or no noise is present. However, adaptation during use is not required for nominal performance.

Figure 19 is a plot of an exemplary linear response of virtual microphone  $V_2$  to a speech source at a first distance, according to some embodiments. In the example shown, Figure 19 is a  
 15 plot of a linear response 1900 of virtual microphone  $V_2$  with  $\beta = 0.8$  to a 1 kHz speech source at a distance of 0.1 m, according to some embodiments. The null in the linear response of virtual microphone  $V_2$  to speech is located at 0 degrees, where the speech is typically expected to be located. Thus, the null is at 0 degrees, where the speech is normally located.

Figure 20 is a plot of an exemplary linear response of virtual microphone  $V_2$  to a noise source at a second distance, according to some embodiments. In the example shown, Figure 20  
 20 depicts a plot of a linear response 2000 of virtual microphone  $V_2$  with  $\beta = 0.8$  to a 1 kHz noise source at a distance of 1.0 m, according to some embodiments. The linear response of  $V_2$  to noise may be devoid of or may include no null. For example, significant or all noise sources may be detected. Thus, there may be no null and most or all noise sources may be detected.

The above formulation for  $V_2(z)$  has a null at the speech location and will therefore  
 25 exhibit minimal response to the speech. This is shown in Figure 19 for an array with  $d_0=10.7$  mm and a speech source on the axis of the array ( $\theta = 0$ ) at 10 cm ( $\beta = 0.8$ ). Note that the speech null at zero degrees is not present for noise in the far field for the same microphone, as shown in Figure 20 with a noise source distance of approximately 1 meter. This insures that noise in front of the user will be detected so that it can be removed. This differs from conventional systems  
 30 that can have difficulty removing noise in the direction of the mouth of the user.

The  $V_1(z)$  can be formulated using the general form for  $V_1(z)$ :

$$V_1(z) = \alpha_A O_1(z) \cdot z^{-d_A} - \alpha_B O_2(z) \cdot z^{-d_B}$$

Since

$$V_2(z) = O_2(z) - z^{\gamma} \beta O_1(z)$$

and, since for noise in the forward direction

$$O_{2N}(z) = O_{1N}(z) \cdot z^{-\gamma},$$

then

$$V_{2N}(z) = O_{1N}(z) \cdot z^{-\gamma} - z^{-\gamma} \beta O_{1N}(z)$$

$$V_{2N}(z) = (1 - \beta)(O_{1N}(z) \cdot z^{-\gamma})$$

If this is then set equal to  $V_1(z)$  above, the result is

$$V_{1N}(z) = \alpha_A O_{1N}(z) \cdot z^{-d_A} - \alpha_B O_{1N}(z) \cdot z^{-\gamma} \cdot z^{-d_B} = (1 - \beta)(O_{1N}(z) \cdot z^{-\gamma})$$

thus the following may set

$$d_A = \gamma$$

$$d_B = 0$$

$$\alpha_A = 1$$

$$\alpha_B = \beta$$

to get

$$V_1(z) = O_1(z) \cdot z^{-\gamma} - \beta O_2(z)$$

The definitions for  $V_1$  and  $V_2$  above mean that for noise  $H_1(z)$  is:

$$H_1(z) = \frac{V_1(z)}{V_2(z)} = \frac{-\beta O_2(z) + O_1(z) \cdot z^{-\gamma}}{O_2(z) - z^{-\gamma} \beta O_1(z)}$$

which, if the amplitude noise responses are about the same, has the form of an allpass filter.

Accordingly, the noise can be accurately modeled, especially in magnitude response, satisfying R2. This formulation assures that the noise response will be as similar as possible and that the speech response will be proportional to  $(1 - \beta^2)$ . Since  $\beta$  is the ratio of the distances from  $O_1$  and  $O_2$  to the speech source, it is affected by the size of the array and the distance from the array to the speech source.

Figure 21 is a plot of an exemplary linear response of virtual microphone  $V_1$  to a speech source at the first distance, according to some embodiments. In the example shown, Figure 21 depicts a plot of a linear response 2100 of virtual microphone  $V_1$  with  $\beta = 0.8$  to a 1 kHz speech source at a distance of 0.1 m, according to some embodiments. The linear response of virtual microphone  $V_1$  to speech is devoid of or includes no null and the response for speech is greater than that shown in Figure 19. In some examples, there is a negligible or no null and the response for speech may be greater than that shown in Figure 19.

Figure 22 is a plot of an exemplary linear response of virtual microphone  $V_1$  to a noise source at the second distance, according to some embodiments. In the example shown, Figure 22 is a plot of a linear response 2200 of virtual microphone  $V_1$  with  $\beta = 0.8$  to a 1 kHz noise source at a distance of 1.0 m, according to some embodiments. The linear response of virtual

microphone  $V_1$  to noise is devoid of or includes no null and the response is very similar to  $V_2$  shown in Figure 20. Figure 22 is a plot of an exemplary linear response of virtual microphone  $V_1$  to a 1 kHz noise source at a distance of 1.0 m, according to some embodiments. There is no null and the response is very similar to  $V_2$  shown in Figure 20.

5 Figure 23 is a plot of linear response of virtual microphone  $V_1$  to a speech source at a distance for exemplary frequencies, according to some embodiments. In the example shown, Figure 23 is a plot of linear responses 2300 of virtual microphone  $V_1$  with  $\beta = 0.8$  to a speech source at a distance of 0.1 m for frequencies of 100, 500, 1000, 2000, 3000, and 4000 Hz, according to some embodiments.

10 Figure 24 is a plot showing comparison of a frequency response 2402 for speech for the array of an embodiment and of a frequency response 2404 for an exemplary cardioid microphone.

The response of  $V_1$  to speech is shown in Figure 21, and the response to noise in Figure 22. Note the difference in speech response compared to  $V_2$  shown in Figure 19 and the  
15 similarity of noise response shown in Figure 20. Also note that the orientation of the speech response for  $V_1$  shown in Figure 21 is completely opposite the orientation of conventional systems, where the main lobe of response is normally oriented toward the speech source. The orientation of an embodiment, in which the main lobe of the speech response of  $V_1$  is oriented away from the speech source, means that the speech sensitivity of  $V_1$  is lower than a normal  
20 directional microphone but is flat for all frequencies within approximately  $\pm 30$  degrees of the axis of the array, as shown in Figure 23. A flatness of response for speech means that, for example, no shaping postfilter may be needed to restore omnidirectional frequency response. This does come at a price - as shown in Figure 24, which shows the speech response of  $V_1$  with  
25  $\beta = 0.8$  and the speech response of a cardioid microphone. The speech response of  $V_1$  is approximately 0 to  $\sim 13$  dB less than a normal directional microphone between approximately 500 and 7500 Hz and approximately 0 to 10+ dB greater than a directional microphone below approximately 500 Hz and above 7500 Hz for a sampling frequency of approximately 16000 Hz. However, the superior noise suppression made possible using this system more than compensates for the initially poorer SNR.

30 It should be noted that Figures 19-22 assume the speech is located at approximately 0 degrees and approximately 10 cm,  $\beta = 0.8$ , and the noise at all angles is located approximately 1.0 meter away from the midpoint of the array. Generally, the noise distance is not required to be 1 m or more, but the denoising is the best for those distances. For distances less than approximately 1 m, denoising will not be as effective due to the greater dissimilarity in the noise

responses of  $V_1$  and  $V_2$ . This has not proven to be an impediment in practical use - in fact, it can be seen as a feature. Any "noise" source that is ~10 cm away from the earpiece may be likely to be desired to be captured and transmitted, at least in some embodiments.

The speech null of  $V_2$  means that the VAD signal is no longer a critical component. The VAD's purpose was to ensure that the system would not train on speech and then subsequently remove it, resulting in speech distortion. If, however,  $V_2$  contains no speech, the adaptive system cannot train on the speech and cannot remove it. As a result, the system can denoise all the time without fear of devoicing, and the resulting clean audio can then be used to generate a VAD signal for use in subsequent single-channel noise suppression algorithms such as spectral subtraction. In addition, constraints on the absolute value of  $H_1(z)$  (i.e., restricting it to absolute values less than two) can keep the system from fully training on speech even if it is detected. In reality, though, speech can be present due to a mis-located  $V_2$  null and/or echoes or other phenomena, and a VAD sensor or other acoustic-related VAD is recommended to minimize speech distortion.

Depending on the application,  $\beta$  and  $\gamma$  may be fixed in the noise suppression algorithm or they can be estimated when the algorithm indicates that speech production is taking place in the presence of little or no noise. In either case, there may be an error in the estimate of the actual  $\beta$  and  $\gamma$  of the system. The following description examines these errors and their effect on the performance of the system. As above, "good performance" of the system indicates that there is sufficient denoising and minimal devoicing.

The effect of an incorrect  $\beta$  and  $\gamma$  on the response of  $V_1$  and  $V_2$  can be seen by examining the definitions above:

$$V_1(z) = O_1(z) \cdot z^{-\gamma_T} - \beta_T O_2(z)$$

$$V_2(z) = O_2(z) \cdot z^{-\gamma_T} - \beta_T O_1(z)$$

where  $\beta_T$  and  $\gamma_T$  denote the theoretical estimates of  $\beta$  and  $\gamma$  used in the noise suppression algorithm. In reality, the speech response of  $O_2$  is

$$O_{2s}(z) = \beta_R O_{1s}(z) \cdot z^{-\gamma_R}$$

where  $\beta_R$  and  $\gamma_R$  denote the real  $\beta$  and  $\gamma$  of the physical system. The differences between the theoretical and actual values of  $\beta$  and  $\gamma$  can be due to mis-location of a speech source (e.g., the speech source may not be where it is assumed to be) and/or a change in air temperature (which changes the speed of sound). Inserting the actual response of  $O_2$  for speech into the above equations for  $V_1$  and  $V_2$  yields

$$V_{1s}(z) = O_{1s}(z)[z^{-\gamma_T} - \beta_T \beta_R z^{-\gamma_R}]$$

$$V_{2s}(z) = O_{1s}(z)[\beta_R z^{-\gamma_R} - \beta_T z^{-\gamma_T}]$$

If the difference in phase is represented by

$$\gamma_R = \gamma_T + \gamma_D$$

And the difference in amplitude as

$$\beta_R = B\beta_T$$

5 then

$$\begin{aligned} V_{1S}(z) &= O_{1S}(z)z^{-\gamma_T}[1 - B\beta_T^2z^{-\gamma_D}] \\ V_{2S}(z) &= \beta_T O_{1S}(z)z^{-\gamma_T}[Bz^{-\gamma_D} - 1]. \end{aligned} \quad \text{Eq. 5}$$

The speech cancellation in  $V_2$  (which directly affects the degree of devoicing) and the speech response of  $V_1$  will be dependent on both  $B$  and  $D$ . An examination of the case where  $D = 0$  follows.

Figure 25 is a plot showing exemplary speech responses for  $V_1$  and  $V_2$  versus  $B$  with  $d_s$ , according to some embodiments. In this example, figure 25 is a plot showing speech response 2502 for  $V_1$  (top, dashed) and speech response 2504  $V_2$  (bottom, solid) versus  $B$  with  $d_s$ , assumed to be 0.1 m, according to some embodiments. In this example, this plot shows the spatial null in  $V_2$  to be relatively broad.

Figure 26 is a plot showing an exemplary ratio of  $V_1/V_2$  speech responses versus  $B$ , according to some embodiments. In the example shown, figure 26 is a plot showing a ratio 2602 of  $V_1/V_2$  speech responses shown in Figure 20 versus  $B$ , according to some embodiments. The ratio of  $V_1/V_2$  may be above 10 dB for all  $0.8 < B < 1.1$ , and this means that the physical  $\beta$  of the system need not be exactly modeled for good performance.

Figure 27 is a plot of exemplary values of  $B$  versus exemplary values of distance of a speech source,  $d_s$ , according to some embodiments. In the example shown, the plot depicts a relationship 2700 between  $B$  and actual  $d_s$  assuming that  $d_s = 10$  cm and  $\theta = 0$ , according to some embodiments.

Figure 28 is a plot of exemplary values of  $B$  versus exemplary values of  $\theta$ , according to some embodiments. In the example shown, the plot depicts a relationship 2800 between  $B$  and  $\theta$  with  $d_s = 10$  cm, according to some embodiments.

In Figure 25, the speech response for  $V_1$  (upper, dashed) and  $V_2$  (lower, solid) compared to  $O_1$  is shown versus  $B$  when  $d_s$  is thought to be approximately 10 cm and  $\theta = 0$ . When  $B = 1$ , the speech is absent from  $V_2$ . In Figure 26, the ratio of the speech responses in Figure 20 is shown. When  $0.8 < B < 1.1$ , the  $V_1/V_2$  ratio is above approximately 10 dB - enough for good performance. Clearly, if  $D = 0$ ,  $B$  can vary significantly without adversely affecting the performance of the system. Again, this assumes that calibration of the microphones so that, for example, the amplitude(s) and phase response(s) are the same for an identical source.

The B factor can be non-unity for a variety of reasons. Either the distance to the speech source or the relative orientation of the array axis and the speech source or both can be different than expected. If both distance and angle mismatches are included for B, then

$$B = \frac{\beta_R}{\beta_T} \frac{\sqrt{d_{SR}^2 - 2d_{SR}d_0 \cos(\theta_R) + d_0^2}}{\sqrt{d_{SR}^2 + 2d_{SR}d_0 \cos(\theta_R) + d_0^2}} \cdot \frac{\sqrt{d_{ST}^2 + 2d_{ST}d_0 \cos(\theta_T) + d_0^2}}{\sqrt{d_{ST}^2 - 2d_{ST}d_0 \cos(\theta_T) + d_0^2}}$$

where again the T subscripts indicate the theorized values and R the actual values. In Figure 27,

5 the factor B is plotted with respect to the actual  $d_s$  with the assumption that  $d_s = 10$  cm and  $\theta = 0$ . So, if the speech source is on-axis of the array, the actual distance can vary from approximately 5 cm to 18 cm without significantly affecting performance--a significant amount. Similarly, Figure 28 shows what happens if the speech source is located at a distance of approximately 10 cm but not on the axis of the array. In this case, the angle can vary up to approximately  $\pm 55$   
10 degrees and still result in a B less than 1.1, assuring good performance. This is a significant amount of allowable angular deviation. If there is both angular and distance errors, the equation above may be used to determine if the deviations will result in adequate performance. Of course, if the value for  $\beta_T$  is allowed to update during speech, essentially tracking the speech source, then B can be kept near unity for almost all configurations.

15 An examination follows of the case where B is unity but D is nonzero. This can happen if the speech source is not where it is thought to be or if the speed of sound is different from what it is believed to be. From Equation 5 above, a factor that reduces the speech null in  $V_2$  for speech is expressed as:

$$N(z) = Bz^{-\gamma_D} - 1$$

or in the continuous s domain as:

$$20 \quad N(s) = Be^{-Ds} - 1.$$

Since  $\gamma$  is the time difference between arrival of speech at  $V_1$  compared to  $V_2$ , it can be errors in estimation of the angular location of the speech source with respect to the axis of the array and/or by temperature changes. Examining the temperature sensitivity, the speed of sound varies with temperature as

$$25 \quad c = 331.3 + (0.606T) \text{ m/s}$$

where T is degrees Celsius. As the temperature decreases, the speed of sound also decreases. Setting 20 °C as a design temperature and a maximum expected temperature range to -40 °C to +60 °C (-40 °F to 140 °F). The design speed of sound at 20 °C is 343 m/s and the slowest speed of sound will be 307 m/s at -40 °C with the fastest speed of sound 362 m/s at 60 °C. Set the  
30 array length ( $2d_0$ ) to be 21 mm. For speech sources on the axis of the array, the difference in travel time for the largest change in the speed of sound is

$$\nabla t_{\text{MAX}} = \frac{d}{c_1} - \frac{d}{c_2} = 0.021 \text{ m} \left( \frac{1}{343 \text{ m/s}} - \frac{1}{307 \text{ m/s}} \right) = -7.2 \times 10^{-6} \text{ sec}$$

or approximately 7 microseconds. The response for  $N(s)$  given  $B = 1$  and  $D = 7.2 \mu\text{sec}$  is shown in Figure 29.

Figure 29 is a diagram of plots of exemplary amplitudes and exemplary phase responses of  $N(s)$ , according to some embodiments. In the example shown, plots of amplitude 2902 (top) and phase 2904 (bottom) response of  $N(s)$  with  $B = 1$  and  $D = -7.2 \mu\text{sec}$ , according to some embodiments. The resulting phase difference may affect some high frequencies more than some low frequencies. The amplitude response is less than approximately -10 dB for all frequencies less than 7 kHz and may be about -9 dB at 8 kHz. Therefore, assuming  $B = 1$ , this system would likely perform well at frequencies up to approximately 8 kHz. This means that a properly compensated system would work well even up to 8 kHz in an exceptionally wide (e.g., -40 C to 80 C) temperature range. Note that the phase mismatch due to the delay estimation error causes  $N(s)$  to be much larger at high frequencies compared to low. If  $B$  is not unity, the robustness of the system may be reduced since the effect from non-unity  $B$  is cumulative with that of non-zero  $D$ . Figure 30 shows the amplitude 3002 and phase response 3004 for  $B = 1.2$  and  $D = 7.2 \mu\text{sec}$ .

Figure 30 is a diagram of plots of exemplary amplitudes and exemplary phase responses of  $N(s)$  with non-unity  $B$ , according to some embodiments. In the example shown, the plots amplitude (top) and phase (bottom) response of  $N(s)$  with  $B = 1.2$  and  $D = -7.2 \mu\text{sec}$ , according to some embodiments. Non-unity  $B$  may affect some or the entire frequency range. Now  $N(s)$  is below approximately -10 dB for frequencies less than approximately 5 kHz and the response at low frequencies is larger. Such a system would still perform well below 5 kHz and may suffer from slightly elevated devoicing for frequencies above 5 kHz. For ultimate performance, a temperature sensor may be integrated into the system to allow the algorithm to adjust  $\gamma_T$  as the temperature varies.

Another way in which  $D$  can be non-zero is when the speech source is not where it is believed to be - specifically, the angle from the axis of the array to the speech source is incorrect. The distance to the source may be incorrect as well, but that introduces an error in  $B$ , not  $D$ .

Referring to Figure 12, it can be seen that for two speech sources (each with their own  $d_s$  and  $\theta$ ) that the time difference between the arrival of the speech at  $O_1$  and the arrival at  $O_2$  is

$$\Delta t = \frac{1}{c} (d_{12} - d_{11} - d_{22} + d_{21})$$

where

$$d_{11} = \sqrt{d_{s1}^2 - 2d_{s1}d_0 \cos(\theta_1) + d_0^2}$$

$$d_{12} = \sqrt{d_{s1}^2 + 2d_{s1}d_0 \cos(\theta_1) + d_0^2}$$

$$d_{21} = \sqrt{d_{s2}^2 - 2d_{s2}d_0 \cos(\theta_2) + d_0^2}$$

$$d_{22} = \sqrt{d_{s2}^2 + 2d_{s2}d_0 \cos(\theta_2) + d_0^2}$$

The  $V_2$  speech cancellation response for  $\theta_1 = 0$  degrees and  $\theta_2 = 30$  degrees and assuming that  $B = 1$  is shown in Figure 31.

Figure 31 is a diagram of plots of exemplary amplitudes and exemplary phase responses of the effect on the speech cancellation in  $V_2$ , according to some embodiments. In the example shown, the plots of amplitude 3102 (top) and phase (bottom) response 3104 of the effect on the speech cancellation in  $V_2$  may be due to a mistake in the location of the speech source with  $q_1 = 0$  degrees and  $q_2 = 30$  degrees, according to some embodiments. Note that the cancellation may be below -10 dB for frequencies below 6 kHz. The cancellation is still below approximately -10 dB for frequencies below approximately 6 kHz, so an error of this type will not significantly affect the performance of the system. However, if  $\theta_2$  is increased to approximately 45 degrees, as shown in Figure 32, the cancellation is below approximately -10 dB for frequencies below approximately 2.8 kHz.

Figure 32 is a diagram of plots of other exemplary amplitudes and exemplary phase responses of the effect on the speech cancellation in  $V_2$ , according to some embodiments. In the example shown, the plots of amplitude 3202 (top) and phase (bottom) response 3204 represent exemplary effects on the speech cancellation in  $V_2$  due to a mistake in the location of the speech source with  $q_1 = 0$  degrees and  $q_2 = 45$  degrees, according to some embodiments. The cancellation may be below -10 dB for frequencies below about 2.8 kHz and a reduction in performance is expected. Less desirable speech cancellation above approximately 4 kHz may result in significant devoicing for those frequencies.

The description above has assumed that the microphones  $O_1$  and  $O_2$  were calibrated so that their response to a source located the same distance away was identical for both amplitude and phase. This is not always feasible, so a more practical calibration procedure is presented below. It is not as accurate, but is much simpler to implement. Begin by defining a filter  $\alpha(z)$  such that:

$$O_{1c}(z) = \alpha(z)O_{2c}(z)$$

where the “C” subscript indicates the use of a known calibration source. The simplest one to use is the speech of the user. Then

$$O_{1S}(z) = \alpha(z)O_{2C}(z)$$

The microphone definitions may be expressed as follows:

$$V_1(z) = O_1(z) \cdot z^{-\gamma} - \beta(z)\alpha(z)O_2(z)$$

$$V_2(z) = \alpha(z)O_2(z) - z^{-\gamma}\beta(z)O_1(z)$$

The  $\beta$  of the system may be fixed and as close to the real value as possible. In practice, the system is not sensitive to changes in  $\beta$  and errors of approximately  $\pm 5\%$  are easily tolerated. During times when the user is producing speech but there is little or no noise, the system can train  $\alpha(z)$  to remove as much speech as possible. This is accomplished by:

1. Construct an adaptive system as shown in Figure 11 with  $\beta O_{1S}(z)z^{-\gamma}$  in the “MIC1” position,  $O_{2S}(z)$  in the “MIC2” position, and  $\alpha(z)$  in the  $H_1(z)$  position.

2. During speech, adapt  $\alpha(z)$  to minimize the residual of the system.

3. Construct  $V_1(z)$  and  $V_2(z)$  as above.

A simple adaptive filter can be used for  $\alpha(z)$  so that the relationship between the microphones is well modeled. The system of an embodiment trains when speech may be produced by the user. A sensor, such as the SSM, can be used to determine when speech is being produced in the absence of noise. If the speech source is fixed in position and will not vary significantly during use (such as when the array is on an earpiece), the adaptation may be infrequent and slow to update in order to minimize any errors introduced by noise present during training.

The above formulation works very well because the noise (far-field) responses of  $V_1$  and  $V_2$  are very similar while the speech (near-field) responses are very different. However, the formulations for  $V_1$  and  $V_2$  can be varied and still result in good performance of the system as a whole. If the definitions for  $V_1$  and  $V_2$  are taken from above and new variables  $B_1$  and  $B_2$  are inserted, the result is:

$$V_1(z) = O_1(z) \cdot z^{-\gamma} - B_1\beta_T O_2(z)$$

$$V_2(z) = O_2(z) - z^{-\gamma}B_2\beta_T O_1(z)$$

where  $B_1$  and  $B_2$  are both positive numbers or zero. If  $B_1$  and  $B_2$  are set equal to unity, the optimal system results as described above. If  $B_1$  is allowed to vary from unity, the response of  $V_1$  is affected. An examination of the case where  $B_2$  is left at 1 and  $B_1$  is decreased follows. As  $B_1$  drops to approximately zero,  $V_1$  becomes less and less directional, until it becomes a simple omnidirectional microphone when  $B_1 = 0$ . Since  $B_2 = 1$ , a speech null remains in  $V_2$ , so very different speech responses remain for  $V_1$  and  $V_2$ . However, the noise responses are much less

similar, so denoising will not be as effective. Practically, though, the system still performs well. B1 can also be increased from unity and once again the system will still denoise well, just not as well as with  $B1 = 1$ .

If B2 is allowed to vary, the speech null in  $V_2$  is affected. As long as the speech null is still sufficiently deep, the system will still perform well. Practically values down to approximately  $B2 = 0.6$  have shown sufficient performance, but it is recommended to set B2 close to unity for optimal performance.

Similarly, variables  $\epsilon$  and  $\Delta$  may be introduced so that:

$$V_1(z) = (\epsilon - \beta)O_{2N}(z) + (1 + \Delta)O_{1N}(z)z^{-\gamma}$$

$$V_2(z) = (1 + \Delta)O_{2N}(z) + (\epsilon - \beta)O_{1N}(z)z^{-\gamma}$$

This formulation also allows the virtual microphone responses to be varied but retains the all-pass characteristic of  $H_1(z)$ . In conclusion, the system is flexible enough to operate well at a variety of B1 values, but note that, at least in some examples, B2 values may be selected to be close to unity to limit devoicing for enhanced performance.

Figure 33 shows experimental results for an exemplary array using a simulator in a loud noise environment, according to some embodiments. In the example shown, experimental results are depicted for a  $2d_0 = 19$  mm array using a linear  $\beta$  of 0.83 on a Bruel and Kjaer Head and Torso Simulator (HATS) in very loud ( $\sim 85$  dBA) music/speech noise environment, according to some embodiments. The noise may be reduced by about 25 dB and the speech may be negligibly or unaffected, with minimal or no noticeable distortion. An alternate microphone calibration technique discussed above may be used to calibrate the microphones. The noise in diagram 3302 may be reduced by about 25 dB and the speech hardly affected, with no noticeable distortion as depicted in diagram 3304. In various examples, the technique significantly increases the SNR of the original speech, thereby outperforming at least some conventional noise suppression techniques.

The DOMA can be a component of a single system, multiple systems, and/or geographically separate systems. The DOMA can also be a subcomponent or subsystem of a single system, multiple systems, and/or geographically separate systems. The DOMA can be coupled to one or more other components (not shown) of a host system or a system coupled to the host system.

One or more components of the DOMA and/or a corresponding system or application to which the DOMA is coupled or connected includes and/or runs under and/or in association with a processing system. The processing system includes any collection of processor-based devices or computing devices operating together, or components of processing systems or devices, as is

known in the art. For example, the processing system can include one or more of a portable computer, portable communication device operating in a communication network, and/or a network server. The portable computer can be any of a number and/or combination of devices selected from among personal computers, cellular telephones, personal digital assistants,  
5 portable computing devices, and portable communication devices, but is not so limited. The processing system can include components within a larger computer system.

Acoustic Voice Activity Detection (AVAD) methods and systems are described herein. The AVAD methods and systems, which include algorithms or programs, use microphones to generate virtual directional microphones that have very similar noise responses and very  
10 dissimilar speech responses. The ratio of the energies of the virtual microphones is then calculated over a given window size and the ratio can then be used with a variety of methods to generate a VAD signal. The virtual microphones can be constructed using either a fixed or an adaptive filter. The adaptive filter generally results in a more accurate and noise-robust VAD signal but requires training. In addition, restrictions can be placed on the filter to ensure that it is  
15 training on speech and not on environmental noise.

In the following description, numerous specific details are introduced to provide a thorough understanding of, and enabling description for, embodiments. One skilled in the relevant art, however, will recognize that these embodiments can be practiced without one or more of the specific details, or with other components, systems, etc. In other instances, well-  
20 known structures or operations are not shown, or are not described in detail, to avoid obscuring aspects of the disclosed embodiments.

Figure 34 is an exemplary configuration of a two-microphone array of the AVAD with speech source S 3406, according to some embodiments. The AVAD of an embodiment uses two physical microphones 3402 ("O<sub>1</sub>") and 3404 ("O<sub>2</sub>") to form two virtual microphones (V<sub>1</sub> and  
25 V<sub>2</sub>). The virtual microphones of an embodiment are directional microphones, but the embodiment is not so limited. The physical microphones of an embodiment include omnidirectional microphones, but the embodiments described herein are not limited to omnidirectional microphones. The virtual microphone (VM) V<sub>2</sub> is configured in such a way that it has minimal response to the speech of the user, while V<sub>1</sub> is configured so that it does respond  
30 to the user's speech but has a very similar noise magnitude response to V<sub>2</sub>, as described in detail herein. The PSAD VAD methods can then be used to determine when speech is taking place. A further refinement is the use of an adaptive filter to further minimize the speech response of V<sub>2</sub>, thereby increasing the speech energy ratio used in PSAD and resulting in better overall performance of the AVAD.

The PSAD algorithm as described herein calculates the ratio of the energies of two directional microphones  $M_1$  and  $M_2$ :

$$R = \sum_i \sqrt{\frac{M_1(z_i)^2}{M_2(z_i)^2}}$$

where the “z” indicates the discrete frequency domain and “i” ranges from the beginning of the window of interest to the end, but the same relationship holds in the time domain. The

5 summation can occur over a window of any length; 200 samples at a sampling rate of 8 kHz has been used to good effect. Microphone  $M_1$  is assumed to have a greater speech response than microphone  $M_2$ . The ratio  $R$  depends on the relative strength of the acoustic signal of interest as detected by the microphones.

10 For matched omnidirectional microphones (i.e. they have the same response to acoustic signals for all spatial orientations and frequencies), the size of  $R$  can be calculated for speech and noise by approximating the propagation of speech and noise waves as spherically symmetric sources. For these the energy of the propagating wave decreases as  $1/r^2$ :

$$R = \sum_i \sqrt{\frac{M_1(z_i)^2}{M_2(z_i)^2}} = \frac{d_2}{d_1} = \frac{d_1 + d}{d_1}$$

The distance  $d_1$  is the distance from the acoustic source to  $M_1$ ,  $d_2$  is the distance from the acoustic source to  $M_2$ , and  $d = d_2 - d_1$  (see Figure 34). It is assumed that microphone (“O<sub>1</sub>”) 3402  
15 is closer to the speech source (the user’s mouth) so that  $d$  is always positive. If the microphones and the user’s mouth are all on a line, then  $d = 2d_0$ , the distance between the microphones. For matched omnidirectional microphones, the magnitude of  $R$ , depends on the relative distance between the microphones and the acoustic source. For noise sources, the distances are typically a meter or more, and for speech sources, the distances are on the order of 10 cm, but the  
20 distances are not so limited. Therefore for a 2-cm array typical values of  $R$  are:

$$R_S = \frac{d_2}{d_1} \approx \frac{12cm}{10cm} = 1.2$$

$$R_N = \frac{d_2}{d_1} \approx \frac{102cm}{100cm} = 1.02$$

where the “S” subscript denotes the ratio for speech sources and “N” the ratio for noise sources. There is not a significant amount of separation between noise and speech sources in this case, and therefore it would be difficult to implement a robust solution using simple omnidirectional microphones.

A better implementation is to use directional microphones where the second microphone has minimal speech response. As described herein, such microphones can be constructed using omnidirectional microphones  $O_1$  and  $O_2$ :

$$V_1(z) = -\beta(z)\alpha(z)O_2(z) + O_1(z)z^{-\gamma} \quad [1]$$

$$V_2(z) = \alpha(z)O_2(z) - \beta(z)O_1(z)z^{-\gamma}$$

where  $\alpha(z)$  is a calibration filter used to compensate  $O_2$ 's response so that it is the same as  $O_1$ ,

- 5  $\beta(z)$  is a filter that describes the relationship between  $O_1$  and calibrated  $O_2$  for speech, and  $\gamma$  is a fixed delay that depends on the size of the array. There is no loss of generality in defining  $\alpha(z)$  as above, as either microphone may be compensated to match the other. For this configuration  $V_1$  and  $V_2$  have very similar noise response magnitudes and very dissimilar speech response magnitudes if

$$\gamma = \frac{d}{c}$$

- 10 where again  $d = 2d_0$  and  $c$  is the speed of sound in air, which is temperature dependent and approximately

$$c = 331.3 \sqrt{1 + \frac{T}{273.15}} \frac{m}{sec}$$

where  $T$  is the temperature of the air in Celsius.

The filter  $\beta(z)$  can be calculated using wave theory to be

$$\beta(z) = \frac{d_1}{d_2} = \frac{d_1}{d_1 + d} \quad [2]$$

- 15 where again  $d_k$  is the distance from the user's mouth to  $O_k$ .

Figure 35 is an exemplary block diagram of  $V_2$  construction using a fixed  $\beta(z)$  filter 3508, according to some embodiments. This fixed (or static)  $\beta$  works sufficiently well if the calibration filter  $\alpha(z)$  3506 is accurate and  $d_1$  and  $d_2$  are accurate for the user. The output of filter 3508 is transmitted to delay filter 3510. This fixed- $\beta$  algorithm, however, neglects effects  
20 such as reflection, diffraction, poor array orientation (i.e. the microphones and the mouth of the user are not all on a line), and the possibility of different  $d_1$  and  $d_2$  values for different users.

The filter  $\beta(z)$  can also be determined experimentally using an adaptive filter.

Figure 36 is an exemplary block diagram of  $V_2$  construction using an adaptive  $\beta(z)$  3608, according to some embodiments, where:

$$\tilde{\beta}(z) = \frac{\alpha(z)O_2(z)}{z^{-\gamma}O_1(z)}$$

- 25 The adaptive process varies  $\tilde{\beta}(z)$  to minimize the output of  $V_2$  when speech is being received by  $O_1$  and  $O_2$ . A small amount of noise may be tolerated with little ill effect, but it is preferred that

speech is being received when the coefficients of  $\tilde{\beta}(z)$  are calculated. Any adaptive process may be used; a normalized least-mean squares (NLMS) algorithm was used in the examples below.

5 The  $V_1$  can be constructed using the current value for  $\tilde{\beta}(z)$  or the fixed filter  $\beta(z)$  can be used for simplicity.

Figure 37 is an exemplary block diagram of  $V_1$  construction, according to some embodiments. An example of ratio  $R$  can be expressed as follows:

$$R = \frac{\|V_1(z)\|}{\|V_2(z)\|} = \sqrt{\frac{(-\tilde{\beta}(z)\alpha(z)O_2(z) + O_1(z)z^{-\gamma})^2}{(\alpha(z)O_2(z) - \tilde{\beta}(z)O_1(z)z^{-\gamma})^2}}$$

where double bar indicates norm and again any size window may be used. If  $\tilde{\beta}(z)$  has been accurately calculated, the ratio for speech may be relatively high (e.g., greater than approximately 2) and the ratio for noise may be relatively low (e.g., less than approximately 1.1). The ratio calculated will depend on both the relative energies of the speech and noise as well as the orientation of the noise and the reverberance of the environment. In practice, either the adapted filter  $\tilde{\beta}(z)$  or the static filter  $b(z)$  may be used for  $V_1(z)$  with little effect on  $R$  - but note the use of the adapted filter  $\tilde{\beta}(z)$  in  $V_2(z)$  for best performance. Many techniques known to those skilled in the art (e.g., smoothing, etc.) can be used to make  $R$  more amenable to use in generating a VAD and the embodiments herein are not so limited.

As shown, Figure 37 includes a calibration filter 3706, a filter 3710, and a delay 3708. The ratio  $R$  can be calculated for the entire frequency band of interest, or can be calculated in frequency subbands. One effective subband discovered was 250 Hz to 1250 Hz, another was 200 Hz to 3000 Hz, but many others are possible and useful.

Once generated, the vector of the ratio  $R$  versus time (or the matrix of  $R$  versus time if multiple subbands are used) can be used with any detection system (such as one that uses fixed and/or adaptive thresholds) to determine when speech is occurring. While many detection systems and methods are known to exist by those skilled in the art and may be used, the method described herein for generating an  $R$  so that the speech is easily discernable is novel. Note that the  $R$  does not depend on the type of noise or its orientation or frequency content;  $R$  simply depends on the  $V_1$  and  $V_2$  spatial response similarity for noise and spatial response dissimilarity for speech. In this way it is very robust and can operate smoothly in a variety of noisy acoustic environments.

30 Figure 38 is an exemplary flow diagram of acoustic voice activity detection 3800, according to some embodiments. The detection comprises forming a first virtual microphone by

combining a first signal of a first physical microphone and a second signal of a second physical microphone 3802. The detection comprises forming a filter that describes a relationship for speech between the first physical microphone and the second physical microphone 3804. The detection comprises forming a second virtual microphone by applying the filter to the first signal to generate a first intermediate signal, and summing the first intermediate signal and the second signal 3806. The detection comprises generating an energy ratio of energies of the first virtual microphone and the second virtual microphone 3808. The detection comprises detecting acoustic voice activity of a speaker when the energy ratio is greater than a threshold value 3810.

The accuracy of the adaptation to the  $\beta(z)$  of the system is a factor in determining the effectiveness of the AVAD. A more accurate adaptation to the actual  $\beta(z)$  of the system leads to lower energy of the speech response in  $V_2$ , and a higher ratio  $R$ . The noise (far-field) magnitude response is largely unchanged by the adaptation process, so the ratio  $R$  will be near unity for accurately adapted beta. For purposes of accuracy, the system can be trained on speech alone, or the noise may be low enough in energy so as not to affect or to have a minimal affect the training.

To make the training as accurate as possible, the coefficients of the filter  $\beta(z)$  of an embodiment are generally updated under the following conditions, but the embodiment is not so limited: speech is being produced (requires a relatively high SNR or other method of detection such as an Aliph Skin Surface Microphone (SSM) as described in U.S. patent application Ser. No. 10/769,302, filed Jan. 30, 2004, which is incorporated by reference herein in its entirety); no wind is detected (wind can be detected using many different methods known in the art, such as examining the microphones for uncorrelated low-frequency noise); and the current value of  $R$  is much larger than a smoothed history of  $R$  values (this ensures that training occurs when strong speech is present). These procedures are flexible and others may be used without significantly affecting the performance of the system. These restrictions can make the system relatively more robust.

Even with these precautions, it is possible that the system accidentally trains on noise (e.g., there may be a higher likelihood of this without use of a non-acoustic VAD device such as the SSM used in the Jawbone headset produced by Aliph, San Francisco, Calif.). Thus, an embodiment includes a further failsafe system to preclude accidental training from significantly disrupting the system. The adaptive  $\beta$  is limited to certain values expected for speech. For example, values for  $d_1$  for an ear-mounted headset will normally fall between 9 and 14 centimeters, so using an array length of  $2d_0 = 2.0$  cm and Equation 2 above,

$$|\beta(z)| = \frac{d_1}{d_2} \approx \frac{d_1}{d_1 + 2d_0}$$

which means that

$$0.82 < |\beta(z)| < 0.88.$$

The magnitude of the  $\beta$  filter can therefore be limited to between approximately 0.82 and 0.88 to preclude problems if noise is present during training. Looser limits can be used to compensate for inaccurate calibrations (the response of omnidirectional microphones is usually calibrated to one another so that their frequency response is the same to the same acoustic source--if the calibration is not completely accurate the virtual microphones may not form properly).

Similarly, the phase of the  $\beta$  filter can be limited to be what is expected from a speech source within  $\pm 30$  degrees from the axis of the array. As described herein, and with reference to Figure 34,

$$\gamma = \frac{d_2 - d_1}{c} \quad (\text{seconds})$$

$$d_1 = \sqrt{d_s^2 - 2d_s d_0 \cos(\theta) + d_0^2}$$

$$d_2 = \sqrt{d_s^2 + 2d_s d_0 \cos(\theta) + d_0^2}$$

where  $d_s$  is the distance from the midpoint of the array to the speech source. Varying  $d_s$  from 10 to 15 cm and allowing  $\theta$  to vary between 0 and  $\pm 30$  degrees, the maximum difference in  $\gamma$  results from the difference of  $\gamma$  at 0 degrees (58.8  $\mu\text{sec}$ ) and  $\gamma$  at  $\pm 30$  degrees for  $d_s = 10$  cm (50.8  $\mu\text{sec}$ ). This means that the maximum expected phase difference is  $58.8 - 50.8 = 8.0 \mu\text{sec}$ , or 0.064 samples at an 8 kHz sampling rate. Since

$$\varphi(f) = 2\pi f t = 2\pi f (8.0 \times 10^{-6}) \text{ rad}$$

the maximum phase difference realized at 4 kHz may be 0.2 rad or about 11.4 degrees, a small amount, but not a negligible one. Therefore the  $\beta$  filter may almost linear phase, but some allowance made for differences in position and angle. In practice a slightly larger amount was used (0.071 samples at 8 kHz) in order to compensate for poor calibration and diffraction effects, and this worked well. The limit on the phase in the example below was implemented as the ratio of the central tap energy to the combined energy of the other taps:

$$\text{phase limit ratio} = \frac{(\text{center tap})^2}{\|\beta\|}$$

where  $\beta$  is the current estimate. This limits the phase by restricting the effects of the non-center taps. Other ways of limiting the phase of the  $\beta$  filter are known to those skilled in the art and the algorithm presented here is not so limited.

Embodiments are presented herein that use both a fixed  $\beta(z)$  and an adaptive  $\beta(z)$ , as described in detail above. In both cases, R was calculated using frequencies between 250 and 3000 Hz using a window size of 200 samples at 8 kHz. The results for  $V_1$  (top plot),  $V_2$  (middle plot), R (bottom plot, solid line, windowed using a 200 sample rectangular window at 8 kHz) and the VAD (bottom plot, dashed line) are shown in Figures 39-44. Figures 39-44 demonstrate the use of a fixed beta filter  $\beta(z)$  in conditions of noise (street and bus noise, approximately 70 dB SPL at the ear), speech (normalized to 94 dB SPL at the mouth reference point (MRP)), and mixed noise and speech, respectively. A Brüel & Kjær Head and Torso Simulator (HATS) was used for the tests and the omnidirectional microphones mounted on HATS' ear with the midline of the array approximately 11 cm from the MRP. The fixed beta filter used was  $\beta_F(z)=0.82$ , where the "F" subscript indicates a fixed filter. The VAD was calculated using a fixed threshold of 1.5.

Figure 39 shows experimental results of an exemplary algorithm using a fixed beta when noise is present, according to some embodiments. The top plot is  $V_1$  3902, the middle plot is  $V_2$  3904, and the bottom plot is R 3906 (solid line) and the VAD result (dashed line) versus time. Examining Figure 39, the response of both  $V_1$  and  $V_2$  are very similar, and the ratio R is very near unity for the entire sample. The VAD response has occasional false positives denoted by spikes in the R plot (windows that are identified by the algorithm as containing speech when they do not), but these are easily removed using standard pulse removal algorithms and/or smoothing of the R results.

Figure 40 shows experimental results of an exemplary algorithm using a fixed beta when speech is present, according to some embodiments. The top plot is  $V_1$  4002, the middle plot is  $V_2$  4004, and the bottom plot is R 4006 (solid line) and the VAD result (dashed line) versus time. The R ratio is between approximately 2 and approximately 7 on average, and the speech is easily discernable using the fixed threshold. These results show that the response of the two virtual microphones to speech are very different, and indeed that ratio R varies from 2-7 during speech. In some examples, there may be few false positives and few false negatives (i.e., windows that contain speech but are not identified as speech windows). Accordingly, the speech may be accurately detected.

Figure 41 shows experimental results of an exemplary algorithm using a fixed beta when speech and noise is present, according to some embodiments. The top plot is  $V_1$  4102, the middle plot is  $V_2$  4104, and the bottom plot is R 4106 (solid line) and the VAD result (dashed line) versus time. The R ratio is lower than when no noise is present, but the VAD remains accurate with a few false positives. There are more false negatives than with no noise, but the

speech remains easily detectable using standard thresholding algorithms. Even in a moderately loud noise environment (Figure 41) the R ratio remains significantly above unity, and the VAD once again returns few false positives. More false negatives are observed, but these may be reduced using standard methods such as smoothing of R and allowing the VAD to continue reporting voiced windows for a few windows after R is under the threshold.

Results using the adaptive beta filter are shown in Figures 42-44. The adaptive filter used was a five-tap NLMS FIR filter using the frequency band from 100 Hz to 3500 Hz. A fixed filter of  $z^{-0.43}$  is used to filter  $O_1$  so that  $O_1$  and  $O_2$  are aligned for speech before the adaptive filter is calculated. The adaptive filter was constrained using the methods above using a low  $\beta$  limit of 0.73, a high  $\beta$  limit of 0.98, and a phase limit ratio of 0.98. Again a fixed threshold was used to generate the VAD result from the ratio R, but in this case a threshold value of 2.5 was used since the R values using the adaptive beta filter are normally greater than when the fixed filter is used. This allows for a reduction of false positives without significantly increasing false negatives.

Figure 42 shows experimental results of an exemplary algorithm using an adaptive beta when noise is present, according to some embodiments. The top plot is  $V_1$  4202, the middle plot is  $V_2$  4204, and the bottom plot is R 4206 (solid line) and the VAD result (dashed line) versus time, with the y-axis expanded to 0-50. Again,  $V_1$  and  $V_2$  are very close in energy and the R ratio is near unity. A single false positive may be generated.

Figure 43 shows experimental results of an exemplary algorithm using an adaptive beta when speech is present, according to some embodiments. The top plot is  $V_1$  4302, the middle plot is  $V_2$  4304, and the bottom plot is R 4306 (solid line) and the VAD result (dashed line) versus time, expanded to 0-50. The  $V_2$  response is greatly reduced using the adaptive beta, and the R ratio has increased from the range of approximately 2-7 to the range of approximately 5-30 on average, making the speech even simpler to detect using standard thresholding algorithms. There are almost no false positives or false negatives. Therefore, the response of  $V_2$  to speech is minimal, R is very high, and all of the speech is easily detected with almost no false positives.

Figure 44 shows experimental results of an exemplary algorithm using an adaptive beta when speech and noise is present, according to some embodiments. The top plot is  $V_1$  4402, the middle plot is  $V_2$  4404, and the bottom plot is R 4406 (solid line) and the VAD result (dashed line) versus time, with the y-axis expanded to 0-50. The R ratio is again lower than when no noise is present, but this R with significant noise present results in a VAD signal that is about the same as the case using the fixed beta with no noise present. This shows that use of the adaptive beta allows the system to perform well in higher noise environments than the fixed beta.

Therefore, with mixed noise and speech, there are again very few false positives and fewer false negatives than in the results of Figure 41, demonstrating that the adaptive filter can outperform the fixed filter in the same noise environment. In practice, the adaptive filter has proven to be significantly more sensitive to speech and less sensitive to noise.

5 Systems and methods for discriminating voiced and unvoiced speech from background noise are provided below including a Non-Acoustic Sensor Voiced Speech Activity Detection (NAVSAD) system and a Pathfinder Speech Activity Detection (PSAD) system. The noise removal and reduction methods provided herein, while allowing for the separation and classification of unvoiced and voiced human speech from background noise, address the  
10 shortcomings of typical systems known in the art by cleaning acoustic signals of interest without distortion.

Figure 45 is an exemplary block diagram of a NAVSAD system 4500, according to some embodiments. The NAVSAD system couples microphones 4510 and sensors 4520 to at least one processor 4530. The sensors 4520 of an embodiment include voicing activity detectors or  
15 non-acoustic sensors. The processor 4530 controls subsystems including a detection subsystem 4550, referred to herein as a detection algorithm, and a denoising subsystem 4540. Operation of the denoising subsystem 4540 is described in detail in the Related Applications. The NAVSAD system works extremely well in any background acoustic noise environment.

Figure 46 is an exemplary block diagram of a PSAD system 4600, according to some  
20 embodiments. The PSAD system couples microphones 4610 to at least one processor 4630. The processor 4630 includes a detection subsystem 4650, referred to herein as a detection algorithm, and a denoising subsystem 4640. The PSAD system is highly sensitive in low acoustic noise environments and relatively insensitive in high acoustic noise environments. The PSAD can operate independently or as a backup to the NAVSAD, detecting voiced speech if the  
25 NAVSAD fails.

Note that the detection subsystems 4650 and denoising subsystems 4640 of both the NAVSAD and PSAD systems of an embodiment are algorithms controlled by the processor 4630, but are not so limited. Alternative embodiments of the NAVSAD and PSAD systems can include detection subsystems 4650 and/or denoising subsystems 4640 that comprise additional  
30 hardware, firmware, software, and/or combinations of hardware, firmware, and software. Furthermore, functions of the detection subsystems 4650 and denoising subsystems 4640 may be distributed across numerous components of the NAVSAD and PSAD systems.

Figure 47 is an exemplary block diagram of a denoising subsystem 4700, referred to herein as the Pathfinder system, according to some embodiments. An example of the Pathfinder

system is briefly described below, and is described in detail in the Related Applications. Two microphones Mic 1 and Mic 2 are used in the Pathfinder system, and Mic 1 is considered the “signal” microphone. With reference to Figure 45, the Pathfinder system 4700 is equivalent to the NAVSAD system 4500 when the voicing activity detector (VAD) 4720 is a non-acoustic voicing sensor 20 and the noise removal subsystem 4740 includes the detection subsystem 50 and the denoising subsystem 40. With reference to Figure 46, the Pathfinder system 4700 is equivalent to the PSAD system 4600 in the absence of the VAD 4720, and when the noise removal subsystem 4740 includes the detection subsystem 50 and the denoising subsystem 40.

The NAVSAD and PSAD systems support a two-level commercial approach in which (i) a relatively less expensive PSAD system supports an acoustic approach that functions in most low- to medium-noise environments, and (ii) a NAVSAD system adds a non-acoustic sensor to enable detection of voiced speech in any environment. Unvoiced speech is normally not detected using the sensor, as it normally does not sufficiently vibrate human tissue. However, in high noise situations detecting the unvoiced speech is not as important, as it is normally very low in energy and easily washed out by the noise. Therefore in high noise environments the unvoiced speech is unlikely to affect the voiced speech denoising. Unvoiced speech information may occur in the presence of little to no noise and, therefore, the unvoiced detection may be highly sensitive in low noise situations, and insensitive in high noise situations. This is not easily accomplished, and comparable acoustic unvoiced detectors known in the art are incapable of operating under these environmental constraints.

The NAVSAD and PSAD systems include an array algorithm for speech detection that uses the difference in frequency content between two microphones to calculate a relationship between the signals of the two microphones. This is in contrast to conventional arrays that attempt to use the time/phase difference of each microphone to remove the noise outside of an “area of sensitivity”. The methods described herein provide a significant advantage, as they do not require a specific orientation of the array with respect to the signal.

Further, the systems described herein are sensitive to noise of every type and every orientation, unlike conventional arrays that depend on specific noise orientations. Consequently, the frequency-based arrays presented herein are unique as they depend on the relative orientation of the two microphones themselves with no dependence on the orientation of the noise and signal with respect to the microphones. This results in a robust signal processing system with respect to the type of noise, microphones, and orientation between the noise/signal source and the microphones.

The systems described herein use the information derived from the Pathfinder noise suppression system and/or a non-acoustic sensor described in the Related Applications to determine the voicing state of an input signal, as described in detail below. The voicing state includes silent, voiced, and unvoiced states. The NAVSAD system, for example, includes a non-acoustic sensor to detect the vibration of human tissue associated with speech. The non-acoustic sensor of an embodiment is a General Electromagnetic Movement Sensor (GEMS) as described briefly below and in detail in the Related Applications, but is not so limited. Alternative embodiments, however, may use any sensor that is able to detect human tissue motion associated with speech and is unaffected by environmental acoustic noise.

The GEMS is a radio frequency device (2.4 GHz) that allows the detection of moving human tissue dielectric interfaces. The GEMS includes an RF interferometer that uses homodyne mixing to detect small phase shifts associated with target motion. In essence, the sensor sends out weak electromagnetic waves (less than 1 milliwatt) that reflect off of whatever is around the sensor. The reflected waves are mixed with the original transmitted waves and the results analyzed for any change in position of the targets. Anything that moves near the sensor will cause a change in phase of the reflected wave that will be amplified and displayed as a change in voltage output from the sensor. A similar sensor is described by Gregory C. Burnett (1999) in "The physiological basis of glottal electromagnetic micropower sensors (GEMS) and their use in defining an excitation function for the human vocal tract"; Ph.D. Thesis, University of California at Davis.

Figure 48 is an exemplary flow diagram of a detection flow 4800 for use in detecting voiced and unvoiced speech, according to some embodiments. With reference to Figures 45 and 46, both the NAVSAD and PSAD systems of an embodiment include the detection flow 4800 as the detection subsystem 4550 of figure 45. This detection flow 4500 operates in real-time and, in an embodiment, operates on 20 millisecond windows and steps 10 milliseconds at a time, but is not so limited. The voice activity determination is recorded for the first 10 milliseconds, and the second 10 milliseconds functions as a "look-ahead" buffer. While an embodiment uses the 20/10 windows, alternative embodiments may use numerous other combinations of window values.

Consideration was given to a number of multi-dimensional factors in developing the detection flow 4800. The biggest consideration was to maintaining the effectiveness of the Pathfinder denoising technique, described in detail in the Related Applications and reviewed herein. Pathfinder performance can be compromised if the adaptive filter training is conducted

on speech rather than on noise. Note that exclusion of any significant amount of speech from the VAD may be avoided in various examples to keep such disturbances to a minimum

Consideration was also given to the accuracy of the characterization between voiced and unvoiced speech signals, and distinguishing each of these speech signals from noise signals.

- 5 This type of characterization can be useful in such applications as speech recognition and speaker verification.

Furthermore, the systems using the detection algorithm of an embodiment function in environments containing varying amounts of background acoustic noise. If the non-acoustic sensor is available, this external noise is not a problem for voiced speech. However, for  
10 unvoiced speech (and voiced if the non-acoustic sensor is not available or has malfunctioned) reliance is placed on acoustic data alone to separate noise from unvoiced speech. An advantage inheres in the use of two microphones in an embodiment of the Pathfinder noise suppression system, and the spatial relationship between the microphones is exploited to assist in the detection of unvoiced speech. However, there may occasionally be noise levels high enough  
15 that the speech will be nearly undetectable and the acoustic-related method may fail. In these situations, the non-acoustic sensor (or hereafter just the sensor) will be required to ensure good performance.

In the two-microphone system, the speech source may be relatively louder in one designated microphone when compared to the other microphone. Tests have shown that this  
20 requirement is easily met with conventional microphones when the microphones are placed on the head, as any noise may result in an  $H_1$  with a gain near unity.

Regarding the NAVSAD system, and with reference to Figure 45 and Figure 47, the NAVSAD relies on two parameters to detect voiced speech. These two parameters include the energy of the sensor in the window of interest, determined in an embodiment by the standard  
25 deviation (SD), and optionally the cross-correlation (XCORR) between the acoustic signal from microphone 1 and the sensor data. The energy of the sensor can be determined in any one of a number of ways, and the SD is just one convenient way to determine the energy.

For the sensor, the SD is akin to the energy of the signal, which normally corresponds quite accurately to the voicing state, but may be susceptible to movement noise (relative motion  
30 of the sensor with respect to the human user) and/or electromagnetic noise. To further differentiate sensor noise from tissue motion, the XCORR can be used. The XCORR is calculated to 15 delays, which corresponds to just under 2 milliseconds at 8000 Hz.

The XCORR can also be useful when the sensor signal is distorted or modulated in some fashion. For example, there are sensor locations (such as the jaw or back of the neck) where

speech production can be detected but where the signal may have incorrect or distorted time-based information. That is, they may not have well defined features in time that will match with the acoustic waveform. However, XCORR is more susceptible to errors from acoustic noise, and in high (<0 dB SNR) environments is almost useless. Therefore it may not be the sole  
5 source of voicing information.

The sensor detects human tissue motion associated with the closure of the vocal folds, so the acoustic signal produced by the closure of the folds is highly correlated with the closures. Therefore, sensor data that correlates highly with the acoustic signal is declared as speech, and sensor data that does not correlate well is termed noise. The acoustic data is expected to lag  
10 behind the sensor data by about 0.1 to 0.8 milliseconds (or about 1-7 samples) as a result of the delay time due to the relatively slower speed of sound (around 330 m/s). However, an embodiment uses a 15-sample correlation, as the acoustic wave shape varies significantly depending on the sound produced, and a larger correlation width is needed to ensure detection.

The SD and XCORR signals are related, but are sufficiently different so that the voiced  
15 speech detection is more reliable. For simplicity, though, either parameter may be used. The values for the SD and XCORR are compared to empirical thresholds, and if both are above their threshold, voiced speech is declared. Example data is presented and described below.

Figures 49A, 49B, and 50 show data plots for an example in which a subject twice speaks the phrase “pop pan”, according to some embodiments. Figure 49A is a diagram of a  
20 plot depicting a received GEMS signal for an utterance along with the mean correlation, according to some embodiments. In the example shown, GEMS signal 4902 for this utterance is depicted along with the mean correlation 4904 between the GEMS signal and the Mic 1 signal and the threshold T1 used for voiced speech detection.

Figure 49B is a diagram of a plot depicting a received GEMS signal for an utterance  
25 along with the standard deviation of the GEMS signal and the threshold for voiced speech detection, according to some embodiments. In the example shown, the plot depicts a received GEMS signal 4902 for this utterance along with the standard deviation 4906 of the GEMS signal and the threshold T2 used for voiced speech detection.

Figure 50 is a diagram of a plot depicting voiced speech detected from an utterance along  
30 with the GEMS signal and the acoustic noise, according to some embodiments. In the example shown, voiced speech 5002 detected from the acoustic or audio signal 5008 is depicted, along with the GEMS signal 5004 and the acoustic noise 5006; no unvoiced speech is detected in this example because of the heavy background babble noise 5006. The thresholds have been set so that there are virtually no false negatives, and occasional false positives. A voiced speech

activity detection accuracy of greater than 99% has been attained under any acoustic background noise conditions.

The NAVSAD can determine when voiced speech is occurring with high degrees of accuracy due to the non-acoustic sensor data. However, the sensor offers little assistance in separating unvoiced speech from noise, as unvoiced speech normally causes no detectable signal in most non-acoustic sensors. If there is a detectable signal, the NAVSAD can be used. In the absence of a detectable signal use is made of the system and methods of the Pathfinder noise removal algorithm in determining when unvoiced speech is occurring. A brief review of the Pathfinder algorithm is described below, while a detailed description is provided in the Related Applications.

With reference to Figure 47, the acoustic information coming into Microphone 1 is denoted by  $m_{\text{sub.1}}(n)$ , the information coming into Microphone 2 is similarly labeled  $m_{\text{sub.2}}(n)$ , and the GEMS sensor is assumed available to determine voiced speech areas. In the  $z$  (digital frequency) domain, these signals are represented as  $M_{\text{sub.1}}(z)$  and  $M_{\text{sub.2}}(z)$ . Then

$$M_1(z) = S(z) + N_2(z)$$

$$M_2(z) = N(z) + S_2(z)$$

with

$$N_2(z) = N(z)H_1(z)$$

$$S_2(z) = S(z)H_2(z)$$

so that

$$M_1(z) = S(z) + N(z)H_1(z)$$

$$M_2(z) = N(z) + S(z)H_2(z)$$

This is the general case for all two microphone systems. There is always going to be some leakage of noise into Mic 1, and some leakage of signal into Mic 2. Equation 1 has four unknowns and two relationships and cannot be solved explicitly.

However, there is another way to solve for some of the unknowns in Equation 1. Examine the case where the signal is not being generated--that is, where the GEMS signal indicates voicing is not occurring. In this case,  $s(n) = S(z) = 0$ , and Equation 1 reduces to

$$M_{1n}(z) = N(z)H_1(z)$$

$$M_{2n}(z) = N(z)$$

where the  $n$  subscript on the  $M$  variables indicate that noise is being received. This leads to

$$M_{1n}(z) = M_{2n}(z)H_1(z)$$

$$H_1(z) = \frac{M_{1n}(z)}{M_{2n}(z)}$$

$H_1(z)$  can be calculated using any of the available system identification algorithms and the microphone outputs when noise is being received. The calculation can be done adaptively, so that if the noise changes significantly  $H_1(z)$  can be recalculated quickly.

- 5 With a solution for one of the unknowns in Equation 1, solutions can be found for another,  $H_2(z)$ , by using the amplitude of the GEMS or similar device along with the amplitude of the two microphones. When the GEMS indicates voicing, but the recent (less than 1 second) history of the microphones indicate low levels of noise, assume that  $n(s) = N(z) \sim 0$ . Then Equation 1 reduces to

$$\begin{aligned} M_{1s}(z) &= S(z) \\ M_{2s}(z) &= S(z)H_2(z) \end{aligned}$$

which in turn leads to

$$\begin{aligned} M_{2s}(z) &= M_{1s}(z)H_2(z) \\ H_2(z) &= \frac{M_{2s}(z)}{M_{1s}(z)} \end{aligned}$$

- 10 which is the inverse of the  $H_1(z)$  calculation, but note that different inputs are being used.

After calculating  $H_1(z)$  and  $H_2(z)$  above, they are used to remove the noise from the signal. Rewrite Equation 1 as

$$\begin{aligned} S(z) &= M_1(z) - N(z)H_1(z) \\ N(z) &= M_2(z) - S(z)H_2(z) \\ S(z) &= M_1(z) - [M_2(z) - S(z)H_2(z)]H_1(z) , \\ S(z)[1 - H_2(z)H_1(z)] &= M_1(z) - M_2(z)H_1(z) \end{aligned}$$

and solve for  $S(z)$  as:

$$S(z) = \frac{M_1(z) - M_2(z)H_1(z)}{1 - H_2(z)H_1(z)}$$

In practice  $H_2(z)$  is usually quite small, so that  $H_2(z)H_1(z) \ll 1$ , and

$$S(z) \approx M_1(z) - M_2(z)H_1(z) ,$$

- 15 obviating the need for the  $H_2(z)$  calculation.

With reference to Figure 46 and Figure 47, the PSAD system is described. As sound waves propagate, they normally lose energy as they travel due to diffraction and dispersion. Assuming the sound waves originate from a point source and radiate isotropically, their amplitude will decrease as a function of  $1/r$ , where  $r$  is the distance from the originating point.

- 20 This function of  $1/r$  proportional to amplitude is the worst case, if confined to a smaller area the reduction will be less. However it is an adequate model for the configurations of interest, specifically the propagation of noise and speech to microphones located somewhere on the user's head.

Figure 51 is an exemplary microphone array for use according to some embodiments of the PSAD system. Placing the microphones Mic 1 and Mic 2 in a linear array with the mouth on the array midline, the difference in signal strength in Mic 1 and Mic 2 (assuming the microphones have identical frequency responses) will be proportional to both  $d_1$  and  $\Delta d$ .

5 Assuming a  $1/r$  (or in this case  $1/d$ ) relationship, it is seen that

$$\Delta M = \frac{|\text{Mic1}|}{|\text{Mic 2}|} = \Delta H_1(z) \propto \frac{d_1 + \Delta d}{d_1},$$

where  $\Delta M$  is the difference in gain between Mic 1 and Mic 2 and therefore  $H_1(z)$ , as above in Equation 2. The variable  $d_{\text{sub.1}}$  is the distance from Mic 1 to the speech or noise source.

Figure 52 is a plot 5200 of  $\Delta M$  versus  $d_1$  for several  $\Delta d$  exemplary values, according to some embodiments. It is clear that as  $\Delta d$  becomes larger and the noise source is closer,  $\Delta M$  becomes larger. The variable  $\Delta d$  will change depending on the orientation to the speech/noise source, from the maximum value on the array midline to zero perpendicular to the array midline. From the plot 5200 it is clear that for small  $\Delta d$  and for distances over approximately 30 centimeters (cm),  $\Delta M$  is close to unity. Since most noise sources are farther away than 30 cm and are unlikely to be on the midline on the array, it is probable that when calculating  $H_1(z)$  as above in Equation 2,  $\Delta M$  (or equivalently the gain of  $H_1(z)$ ) will be close to unity. Conversely, for noise sources that are close (within a few centimeters), there could be a substantial difference in gain depending on which microphone is closer to the noise.

If the “noise” is the user speaking, and Mic 1 is closer to the mouth than Mic 2, the gain increases. Since environmental noise normally originates much farther away from the user’s head than speech, noise will be found during the time when the gain of  $H_1(z)$  is near unity or some fixed value, and speech can be found after a sharp rise in gain. The speech can be unvoiced or voiced, as long as it is of sufficient volume compared to the surrounding noise. The gain will stay somewhat high during the speech portions, then descend quickly after speech ceases. The rapid increase and decrease in the gain of  $H_1(z)$  may be sufficient to allow the detection of speech under almost any circumstances. The gain in this example is calculated by the sum of the absolute value of the filter coefficients. This sum is not equivalent to the gain, but the two are related in that a rise in the sum of the absolute value reflects a rise in the gain.

As an example of this behavior, Figure 53 shows a plot 5300 of an exemplary gain parameter 5302 as the sum of the absolute values of  $H_1(z)$  and the acoustic data 5304 or audio from microphone 1, according to some embodiments. The speech signal was an utterance of the phrase “pop pan”, repeated twice. The evaluated bandwidth included the frequency range from 2500 Hz to 3500 Hz, although 1500 Hz to 2500 Hz was additionally used in practice. Note the

rapid increase in the gain when the unvoiced speech is first encountered, then the rapid return to normal when the speech ends. The large changes in gain that result from transitions between noise and speech can be detected by any standard signal processing techniques. The standard deviation of the last few gain calculations is used, with thresholds being defined by a running average of the standard deviations and the standard deviation noise floor. The later changes in gain for the voiced speech are suppressed in this plot 5300 for clarity.

Figure 54 is an example of an alternative plot 5400 of acoustic data presented in Figure 53, according to some embodiments. The data used to form plot 5300 is presented again in this plot 5400, along with audio data 5404 and GEMS data 5406 without noise to make the unvoiced speech apparent. The voiced signal 5402 has three possible values: 0 for noise, 1 for unvoiced, and 2 for voiced. Denoising may be accomplished when  $V = 0$ . It is clear that the unvoiced speech is captured very well, aside from two single dropouts in the unvoiced detection near the end of each "pop". However, these single-window dropouts are not common and do not significantly affect the denoising algorithm. They can easily be removed using standard smoothing techniques.

What is not clear from this plot 5400 is that the PSAD system functions as an automatic backup to the NAVSAD. This is because the voiced speech (since it has the same spatial relationship to the microphones as the unvoiced speech) will be detected as unvoiced if the sensor or NAVSAD system fail for any reason. The voiced speech will be misclassified as unvoiced, but the denoising will still not take place, preserving the quality of the speech signal.

However, this automatic backup of the NAVSAD system functions best in an environment with low noise (approximately 10+ dB SNR), as high amounts (10 dB of SNR or less) of acoustic noise can quickly overwhelm any acoustic-related unvoiced detector, including the PSAD. This is evident in the difference in the voiced signal data 5002 and 5402 shown in plots 5000 and 5400 of Figures 50 and 54, respectively, where the same utterance is spoken, but the data of plot 5000 shows no unvoiced speech because the unvoiced speech is undetectable. This is the desired behavior when performing denoising, since if the unvoiced speech is not detectable then it will not significantly affect the denoising process. Using the Pathfinder system to detect unvoiced speech ensures detection of any unvoiced speech loud enough to distort the denoising.

Regarding hardware considerations, and with reference to Figure 51, the configuration of the microphones can have an effect on the change in gain associated with speech and the thresholds needed to detect speech. In general, each configuration will require testing to determine the proper thresholds, but tests with two very different microphone configurations

showed the same thresholds and other parameters to work well. The first microphone set had the signal microphone near the mouth and the noise microphone several centimeters away at the ear, while the second configuration placed the noise and signal microphones back-to-back within a few centimeters of the mouth. The results presented herein were derived using the first  
5 microphone configuration, but the results using the other set are virtually identical, so the detection algorithm is relatively robust with respect to microphone placement.

A number of configurations are possible using the NAVSAD and PSAD systems to detect voiced and unvoiced speech. One configuration uses the NAVSAD system (non-acoustic) to detect voiced speech along with the PSAD system to detect unvoiced speech; the  
10 PSAD also functions as a backup to the NAVSAD system for detecting voiced speech. An alternative configuration uses the NAVSAD system (non-acoustic correlated with acoustic) to detect voiced speech along with the PSAD system to detect unvoiced speech; the PSAD also functions as a backup to the NAVSAD system for detecting voiced speech. Another alternative configuration uses the PSAD system to detect both voiced and unvoiced speech.

15 While the systems described above have been described with reference to separating voiced and unvoiced speech from background acoustic noise, there are no reasons more complex classifications cannot be made. For more in-depth characterization of speech, the system can band pass the information from Mic 1 and Mic 2 so that it is possible to see which bands in the Mic 1 data are more heavily composed of noise and which are more weighted with speech.

20 Using this knowledge, it is possible to group the utterances by their spectral characteristics similar to conventional acoustic methods; this method would work better in noisy environments.

As an example, the “k” in “kick” has significant frequency content from 500 Hz to 4000 Hz, but a “sh” in “she” contains significant energy from 1700-4000 Hz. Voiced speech could be classified in a similar manner. For instance, an /i/ (“ee”) has significant energy around 300 Hz  
25 and 2500 Hz, and an /a/ (“ah”) has energy at around 900 Hz and 1200 Hz. This ability to discriminate unvoiced and voiced speech in the presence of noise is, thus, very useful.

An acoustic vibration sensor, also referred to as a speech sensing device, is described below. The acoustic vibration sensor is similar to a microphone in that it captures speech information from the head area of a human talker or talker in noisy environments. Previous  
30 solutions to this problem have either been vulnerable to noise, physically too large for certain applications, or cost prohibitive. In contrast, the acoustic vibration sensor described herein accurately detects and captures speech vibrations in the presence of substantial airborne acoustic noise, yet within a smaller and cheaper physical package. The noise-immune speech information provided by the acoustic vibration sensor can subsequently be used in downstream

speech processing applications (speech enhancement and noise suppression, speech encoding, speech recognition, talker verification, etc.) to improve the performance of those applications.

Figure 55 is a cross section view of an example of an acoustic vibration sensor 5500, also referred to herein as the sensor 5500, according to some embodiments. Figure 56A is an exploded view of an acoustic vibration sensor 5500, according to the example shown in Figure 55. Figure 56B is perspective view of an acoustic vibration sensor 5500, according to the example shown in Figure 55. The sensor 5500 includes an enclosure 5502 having a first port 5504 on a first side and at least one second port 5506 on a second side of the enclosure 5502. A diaphragm 5508, also referred to as a sensing diaphragm 5508, is positioned between the first and second ports. A coupler 5510, also referred to as the shroud 5510 or cap 5510, forms an acoustic seal around the enclosure 5502 so that the first port 5504 and the side of the diaphragm facing the first port 5504 are isolated from the airborne acoustic environment of the human talker. The coupler 5510 of an embodiment is contiguous, but is not so limited. The second port 5506 couples a second side of the diaphragm to the external environment.

The sensor also includes electret material 5520 and the associated components and electronics coupled to receive acoustic signals from the talker via the coupler 5510 and the diaphragm 5508 and convert the acoustic signals to electrical signals representative of human speech. Electrical contacts 5530 provide the electrical signals as an output. Alternative embodiments can use any type/combination of materials and/or electronics to convert the acoustic signals to electrical signals representative of human speech and output the electrical signals.

The coupler 5510 of an embodiment is formed using materials having acoustic impedances matched to the impedance of human skin (characteristic acoustic impedance of skin is approximately  $1.5 \times 10^6 \text{ Pa} \times \text{s/m}$ ). The coupler 5510 therefore, is formed using a material that includes at least one of silicone gel, dielectric gel, thermoplastic elastomers (TPE), and rubber compounds, but is not so limited. As an example, the coupler 5510 of an embodiment is formed using Kraiburg TPE products. As another example, the coupler 5510 of an embodiment is formed using Sylgard® Silicone products.

The coupler 5510 of an embodiment includes a contact device 5512 that includes, for example, a nipple or protrusion that protrudes from either or both sides of the coupler 5510. In operation, a contact device 5512 that protrudes from both sides of the coupler 5510 includes one side of the contact device 5512 that is in contact with the skin surface of the talker and another side of the contact device 5512 that is in contact with the diaphragm, but the embodiment is not

so limited. The coupler 5510 and the contact device 5512 can be formed from the same or different materials.

The coupler 5510 transfers acoustic energy efficiently from skin/flesh of a talker to the diaphragm, and seals the diaphragm from ambient airborne acoustic signals. Consequently, the  
5 coupler 5510 with the contact device 5512 efficiently transfers acoustic signals directly from the talker's body (speech vibrations) to the diaphragm while isolating the diaphragm from acoustic signals in the airborne environment of the talker (characteristic acoustic impedance of air is approximately  $415 \text{ Pa} \cdot \text{s/m}$ ). The diaphragm is isolated from acoustic signals in the airborne environment of the talker by the coupler 5510 because the coupler 5510 prevents the signals  
10 from reaching the diaphragm, thereby reflecting and/or dissipating much of the energy of the acoustic signals in the airborne environment. Consequently, the sensor 5500 responds primarily to acoustic energy transferred from the skin of the talker, not air. When placed against the head of the talker, the sensor 5500 picks up speech-induced acoustic signals on the surface of the skin while airborne acoustic noise signals are largely rejected, thereby increasing the signal-to-noise  
15 ratio and providing a very reliable source of speech information.

Performance of the sensor 5500 is enhanced through the use of the seal provided between the diaphragm and the airborne environment of the talker. The seal is provided by the coupler 5510. A modified gradient microphone is used in an embodiment because it has pressure ports on both ends. Thus, when the first port 5504 is sealed by the coupler 5510, the  
20 second port 5506 provides a vent for air movement through the sensor 5500.

Figure 57 is a schematic diagram of a coupler 5510 of an acoustic vibration sensor, according to the example shown in Figure 55. The dimensions shown are in millimeters and are intended to serve as an example for one embodiment. Alternative embodiments of the coupler can have different configurations and/or dimensions. The dimensions of the coupler 5510 show  
25 that the acoustic vibration sensor 5500 is small in that the sensor 5500 of an embodiment is approximately the same size as typical microphone capsules found in mobile communication devices. This small form factor allows for use of the sensor 5510 in highly mobile miniaturized applications, where some example applications include at least one of cellular telephones, satellite telephones, portable telephones, wireline telephones, Internet telephones, wireless  
30 transceivers, wireless communication radios, personal digital assistants (PDAs), personal computers (PCs), headset devices, head-worn devices, and earpieces.

The acoustic vibration sensor provides very accurate Voice Activity Detection (VAD) in high noise environments, where high noise environments include airborne acoustic environments in which the noise amplitude is as large if not larger than the speech amplitude as

would be measured by conventional omnidirectional microphones. Accurate VAD information provides significant performance and efficiency benefits in a number of speech processing applications including but not limited to: noise suppression algorithms such as the Pathfinder algorithm available from Aliph, Brisbane, Calif. and described in the Related Applications; speech compression algorithms such as the Enhanced Variable Rate Coder (EVRC) deployed in many commercial systems; and speech recognition systems.

In addition to providing signals having an improved signal-to-noise ratio, the acoustic vibration sensor uses minimal power to operate (on the order of 200 micro Amps, for example). In contrast to alternative solutions that require power, filtering, and/or significant amplification, the acoustic vibration sensor uses a standard microphone interface to connect with signal processing devices. The use of the standard microphone interface avoids the additional expense and size of interface circuitry in a host device and supports for of the sensor in highly mobile applications where power usage is an issue.

Figure 58 is an exploded view of an example of an acoustic vibration sensor 5800, according to at least some embodiments. The sensor 5800 includes an enclosure 5802 having a first port 5804 on a first side and at least one second port (not shown) on a second side of the enclosure 5802. A diaphragm 5808 is positioned between the first and second ports. A layer of silicone gel 5809 or other similar substance is formed in contact with at least a portion of the diaphragm 5808. A coupler 5810 or shroud 5810 is formed around the enclosure 5802 and the silicon gel 5809 where a portion of the coupler 5810 is in contact with the silicon gel 5809. The coupler 5810 and silicon gel 5809 in combination form an acoustic seal around the enclosure 5802 so that the first port 5804 and the side of the diaphragm facing the first port 5804 are isolated from the acoustic environment of the human talker. The second port couples a second side of the diaphragm to the acoustic environment.

As described above, the sensor includes additional electronic materials as appropriate that couple to receive acoustic signals from the talker via the coupler 5810, the silicon gel 5809, and the diaphragm 5808 and convert the acoustic signals to electrical signals representative of human speech. Alternative embodiments can use any type/combination of materials and/or electronics to convert the acoustic signals to electrical signals representative of human speech.

The coupler 5810 and/or gel 5809 of an embodiment are formed using materials having impedances matched to the impedance of human skin. As such, the coupler 5810 is formed using a material that includes at least one of silicone gel, dielectric gel, thermoplastic elastomers (TPE), and rubber compounds, but is not so limited. The coupler 5810 transfers acoustic energy efficiently from skin/flesh of a talker to the diaphragm, and seals the diaphragm from ambient

airborne acoustic signals. Consequently, the coupler 5810 efficiently transfers acoustic signals directly from the talker's body (speech vibrations) to the diaphragm while isolating the diaphragm from acoustic signals in the airborne environment of the talker. The diaphragm is isolated from acoustic signals in the airborne environment of the talker by the silicon gel 5809/coupler 5810 because the silicon gel 5809/coupler 5810 prevents the signals from reaching the diaphragm, thereby reflecting and/or dissipating much of the energy of the acoustic signals in the airborne environment. Consequently, the sensor 5800 responds primarily to acoustic energy transferred from the skin of the talker, not air. When placed against the head of the talker, the sensor 5800 picks up speech-induced acoustic signals on the surface of the skin while airborne acoustic noise signals are largely rejected, thereby increasing the signal-to-noise ratio and providing a very reliable source of speech information.

There are many locations outside the ear from which the acoustic vibration sensor can detect skin vibrations associated with the production of speech. The sensor can be mounted in a device, handset, or earpiece in any manner, the restriction being that reliable skin contact is used to detect the skin-borne vibrations associated with the production of speech.

Figure 59 shows representative areas of sensitivity 5900-5920 on the human head appropriate for placement of an example of an acoustic vibration sensor 5500/5800, according to some embodiments. The areas of sensitivity 5900-5920 include numerous locations 5902-5908 in an area behind the ear 5900, at least one location 5912 in an area in front of the ear 5910, and in numerous locations 5922-5928 in the ear canal area 5920. The areas of sensitivity 5900-5920 are the same for both sides of the human head. These representative areas of sensitivity 5900-5920 are provided as examples and do not limit the embodiments described herein to use in these areas.

Figure 60 is a headset device 6000 that includes an exemplary acoustic vibration sensor 5500/5800 placed at any of a number of locations 6002-6010, according to some embodiments. Generally, placement of the acoustic vibration sensor 5500/5800 can be on any part of the device 6000 that corresponds to the areas of sensitivity 5900-5920 (Figure 59) on the human head. While a headset device is shown as an example, any number of communication devices known in the art can carry and/or couple to an acoustic vibration sensor 5500/5800.

Figure 61 is a diagram of a manufacturing method 6100 for an exemplary acoustic vibration sensor, according to some embodiments. Operation begins with, for example, a uni-directional microphone 6120, at block 6102. Silicon gel 6122 is formed over/on the diaphragm (not shown) and the associated port, at block 6104. A material 6124, for example polyurethane film, is formed or placed over the microphone 6120/silicone gel 6122 combination, at block

6106, to form a coupler or shroud. A snug fit collar or other device is placed on the microphone to secure the material of the coupler during curing, at block 6108.

Note that the silicon gel (block 6102) is an optional component that depends on the embodiment of the sensor being manufactured, as described above. Consequently, the manufacture of an acoustic vibration sensor 5500 that includes a contact device 5512 (referring to Figure 55) will not include the formation of silicon gel 6122 over/on the diaphragm. Further, the coupler formed over the microphone for this sensor 5500 will include the contact device 5512 or formation of the contact device 5512.

The systems and methods described herein include and/or run under and/or in association with a processing system. The processing system includes any collection of processor-based devices or computing devices operating together, or components of processing systems or devices, as is known in the art. For example, the processing system can include one or more of a portable computer, portable communication device operating in a communication network, and/or a network server. The portable computer can be any of a number and/or combination of devices selected from among personal computers, cellular telephones, personal digital assistants, portable computing devices, and portable communication devices, but is not so limited. The processing system can include components within a larger computer system.

The processing system of an embodiment includes at least one processor and at least one memory device or subsystem. The processing system can also include or be coupled to at least one database. The term "processor" as generally used herein refers to any logic processing unit, such as one or more central processing units (CPUs), digital signal processors (DSPs), application-specific integrated circuits (ASIC), etc. The processor and memory can be monolithically integrated onto a single chip, distributed among a number of chips or components of a host system, and/or provided by some combination of algorithms. The methods described herein can be implemented in one or more of software algorithm(s), programs, firmware, hardware, components, circuitry, in any combination.

System components embodying the systems and methods described herein can be located together or in separate locations. Consequently, system components embodying the systems and methods described herein can be components of a single system, multiple systems, and/or geographically separate systems. These components can also be subcomponents or subsystems of a single system, multiple systems, and/or geographically separate systems. These components can be coupled to one or more other components of a host system or a system coupled to the host system.

Communication paths couple the system components and include any medium for communicating or transferring files among the components. The communication paths include wireless connections, wired connections, and hybrid wireless/wired connections. The communication paths also include couplings or connections to networks including local area  
5 networks (LANs), metropolitan area networks (MANs), wide area networks (WANs), proprietary networks, interoffice or backend networks, and the Internet. Furthermore, the communication paths include removable fixed mediums like floppy disks, hard disk drives, and CD-ROM disks, as well as flash RAM, Universal Serial Bus (USB) connections, RS-232 connections, telephone lines, buses, and electronic mail messages.

10 Unless the context clearly requires otherwise, throughout the description, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of “including, but not limited to.” Additionally, the words “herein,” “hereunder,” “above,” “below,” and words of similar import refer to this application as a whole and not to any particular portions of this application. When  
15 the word “or” is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

The above description of embodiments is not intended to be exhaustive or to limit the systems and methods described to the precise form disclosed. While specific embodiments and  
20 examples are described herein for illustrative purposes, various equivalent modifications are possible within the scope of other systems and methods, as those skilled in the relevant art will recognize. The teachings provided herein can be applied to other processing systems and methods, not for the systems and methods described above.

The elements and acts of the various embodiments described above can be combined to  
25 provide further embodiments. These and other changes can be made to the embodiments in light of the above detailed description.

In at least some of the embodiments of the invention, the structures and/or functions of any of the above-described interfaces and panels can be implemented in software, hardware, firmware, circuitry, or a combination thereof. Note that the structures and constituent elements  
30 shown throughout, as well as their functionality, can be aggregated with one or more other structures or elements.

Alternatively, the elements and their functionality can be subdivided into constituent sub-elements, if any. As software, the above-described described techniques can be implemented using various types of programming or formatting languages, frameworks, syntax,

applications, protocols, objects, or techniques, including C, Objective C, C++, C#, Flex™, Fireworks®, Java™, Javascript™, AJAX, COBOL, Fortran, ADA, XML, HTML, DHTML, XHTML, HTTP, XMPP, and others. These can be varied and are not limited to the examples or descriptions provided.

5           In at least some of the embodiments of the invention, one or more of the structures and/or functions of any of the above-described features can be implemented in software, hardware, firmware, circuitry, or a combination thereof. Note that the structures and constituent elements above, as well as their functionality, can be aggregated with one or more other structures or elements. Alternatively, the elements and their functionality can be subdivided into  
10       constituent sub-elements, if any.

          The foregoing description, for purposes of explanation, used specific nomenclature to provide a thorough understanding of the invention. However, it will be apparent to one skilled in the art that specific details are not required in order to practice the invention. In fact, this description may not be read to limit any feature or aspect of the present invention to any  
15       embodiment; rather features and aspects of one embodiment can readily be interchanged with other embodiments.

          Thus, the foregoing descriptions of specific embodiments of the invention are presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed; many alternatives, modifications, equivalents, and  
20       variations are possible in view of the above teachings. For the purpose of clarity, technical material that is known in the technical fields related to the embodiments has not been described in detail to avoid unnecessarily obscuring the description. Thus, the various embodiments can be modified within the scope and equivalents of the appended claims. Further, the embodiments were chosen and described in order to best explain the principles of the invention and its  
25       practical applications; they thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. Notably, not every benefit described herein need be realized by each embodiment of the present invention; rather any specific embodiment can provide one or more of the advantages discussed above. In the claims, elements and/or operations do not imply any  
30       particular order of operation, unless explicitly stated in the claims. It is intended that the following claims and their equivalents define the scope of the invention.

In the claims:

- I.       A communications apparatus comprising:  
          a receiving path over which received speech signals traverse in an audio stream;

an dynamic audio enhancement device disposed in the receiving path, the dynamic audio enhancement device configured to modify an amount of volume and an amount of equalization of the audio stream and comprising:

a noise level estimator ("NLE") configured to generate a signal representing a noise level estimate, the noise level estimator including:

a non-stationary noise detector to detect non-stationary noise, and  
a stationary noise detector detects stationary noise,

wherein the noise level estimator is configured to generate the signal representing a first noise level estimate based on detection of the non-stationary noise or a second noise level estimate based on detection of the stationary noise.

2. The apparatus of claim 1 wherein the noise level estimator further comprises:

a bi-modal noise level estimator configured to generate noise level estimates differently based on the detect type of noise.

3. The apparatus of claim 1 further comprising:

an automatic volume control ("AVC") configured to implement either the first noise level estimate or the second noise level estimate.

4. The apparatus of claim 1 wherein the received speech signals comprise:

a signal ("O<sub>2</sub>") including data generated by an omnidirectional microphone.

5. The apparatus of claim 1 wherein the non-stationary noise detector comprises:

a nonlinear peak tracking filter configured to receive a signal representing the non-stationary noise as speech-like non-stationary noise and to filter a signal representing a largest absolute value in a frame; and

a low pass filter configured to smooth the result of the nonlinear peak tracking filter.

6. The apparatus of claim 1 wherein the stationary noise detector comprises:

an energy computation module configured to compute the energy of a one of the received speech signals; and

a moving average filter.

7. The apparatus of claim 1 wherein the dynamic audio enhancement device further comprises:

a limiter configured to introduce a frame of delay to detect a gain to be applied to the received speech signals, and to attenuate a magnitude of the received speech signals.

8. The apparatus of claim 7 wherein the limiter is further configured to prevent clipping at a speaker by attenuating the magnitude of the received speech signals.

9. The apparatus of claim 7 wherein the limiter is further configured to reduce strongly-voiced portions of the received speech signals to amplify unvoiced and weakly-voiced portions of the received speech signals.
10. The apparatus of claim 1 further comprises:
- 5 either a headset or a handset.
11. A method comprising:
- receiving received speech signals in an audio stream via a receiving processing path;
- detecting a type of noise;
- generating a signal representing either a first noise level estimate based on detection of
- 10 non-stationary noise as the type of noise or a second noise level estimate based on detection of stationary noise as the type of noise;
- modifying an amount of volume of the audio stream based on either the first noise level estimate or the second noise level estimate to reduce or eliminate the effects of environmental acoustic noise associated with the received speech signals.
- 15 12. The method of claim 11 further comprising:
- introducing a delay to form a delayed portion of the received speech signals;
- detecting a gain to be applied to the delayed portion of the received speech signals at a first point in time; and
- attenuating a magnitude of one of the received speech signals at a second point in time.
- 20 13. The method of claim 12 wherein attenuating the magnitude of the one of the received speech signals further comprises:
- preventing clipping at a speaker.
14. The method of claim 12 wherein attenuating the magnitude of the one of the received speech signals further comprises:
- 25 reducing magnitudes of strongly-voiced portions of the received speech signals; and
- amplifying unvoiced and weakly-voiced portions of the received speech signals.
15. The method of claim 11 further comprising:
- applying the first noise level estimate or the second noise level estimate to an automatic volume control ("AVC") and an equalizer.
- 30 16. The method of claim 11 wherein receiving received speech signals comprises:
- receiving a signal ("O<sub>2</sub>") including data generated by an omnidirectional microphone.
17. The method of claim 11 further comprising:
- identifying a first mode of operation;

identifying a largest absolute value in a frame for a received speech signal of the received speech signals;

filtering the largest absolute value of the received speech signal nonlinearly form a result; and

5        applying a low pass filter to smooth the result of nonlinearly filtering.

18.     The method of claim 11 wherein the stationary noise detector comprises:

identifying a second mode of operation;

computing the energy of a one of the received speech signals; and

filtering the one of the received speech signals using a moving average filter.

10    19.     The method of claim 11 further comprising:

receiving a compressor gain; and

compressing high speech levels and boosting low speech levels.

20.     The method of claim 11 further comprising:

adding wind comfort noise.

15     DYNAMIC ENHANCEMENT OF AUDIO (DAE) IN HEADSET SYSTEMS




FIG. 1A

124




FIG. 1B

140




FIG. 1C

145




FIG. 1D

150




FIG. 1E

155




FIG. 1F




FIG. 1G

170




FIG. 1H

180 →




FIG. 1I




FIG. 2A




FIG. 2B

290

11/60




FIG. 2C




FIG.3




FIG. 4




FIG. 5





FIG.6






FIG.8




FIG.9

At initialization:

SSM\_contact\_state = FALSE; SSMVAD\_cnt = 0; AVAD\_cnt = 0;

Steady State Operation:




FIG.10




FIG.11




FIG.12




FIG.13




FIG.14




FIG.15




FIG.16




FIG.17




FIG.18

Linear response of V2 to a speech source at 0.10 meters




FIG.19

Linear response of V2 to a noise source at 1 meters




FIG.20

Linear response of V1 to a speech source at 0.10 meters




FIG.21

Linear response of V1 to a noise source at 1 meters




FIG.22

Linear response of V1 to a speech source at 0.1 meters




FIG.23




FIG.24




FIG.25




FIG.26


B factor vs. actual  $d_s$  assuming  $d_s = 0.1\text{m}$  and  $\theta = 0$ 

FIG.27


B versus  $\theta$  assuming  $d_s = 0.1\text{m}$ 

FIG.28




FIG.29




FIG.30




FIG.31




FIG.32

35/60




FIG.33




FIG.34




FIG.35




FIG.36




FIG.37




FIG.38

V1 (top) and V2 (bottom) for fixed beta in noise




FIG.39

40/60




FIG.40

V1 (top) and V2 (bottom) for fixed beta speech in noise




FIG.41




FIG.42

43/60




FIG.43

44/60

V1 (top) and V2 (bottom) for adaptive beta speech in noise




FIG.44




FIG.45




FIG.46




FIG.47




FIG. 48




FIG. 49A




FIG. 49B




FIG.50




FIG.51




FIG.52




FIG.53







FIG.54

54/60


FIG.55





56/60

FIG.57



57/60

FIG.58







FIG.59





60/60



## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2012/020848

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G10L 15/20 (2012.01)

USPC - 704/233

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G06K 9/40; G10L 15/20, 19/00, 14, 21/02; H03G 3/00, 9/00; H04N 5/21; H04R 3/02, 7/10 (2012.01)

USPC - 381/56, 71.11, 94.1, 94.8, 103, 107, 122; 704/233, E15.039, E19.001

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, ProQuest

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|------------------------------------------------------------------------------------|-----------------------|
| Y         | US 2009/0281805 A1 (LEBLANC et al) 12 November 2009 (12.11.2009) entire document   | 1-20                  |
| Y         | US 5,579,435 A (JANSSON) 26 November 1996 (26.11.1996) entire document             | 1-20                  |
| Y         | US 7,139,393 B1 (MURASE) 21 November 2006 (21.11.2006) entire document             | 1-20                  |
| Y         | US 7,245,965 B1 (PEI et al) 17 July 2007 (17.07.2007) entire document              | 2                     |
| Y         | US 6,819,804 B2 (TENZE et al) 16 November 2004 (16.11.2004) entire document        | 3, 15                 |
| Y         | US 7,174,022 B1 (ZHANG et al) 06 February 2007 (06.02.2007) entire document        | 4, 16                 |
| Y         | US 7,349,547 B1 (ISVAN) 25 March 2008 (25.03.2008) entire document                 | 20                    |

☐ Further documents are listed in the continuation of Box C.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;" document member of the same patent family

Date of the actual completion of the international search

04 May 2012

Date of mailing of the international search report

23 MAY 2012

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents  
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300  
PCT OSP: 571-272-7774