



- (51) **International Patent Classification:**
G10L 21/0208 (2013.01) *H04M 9/08* (2006.01)
- (21) **International Application Number:**
PCT/US2017/016274
- (22) **International Filing Date:**
2 February 2017 (02.02.2017)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
PCT/CN2016/073188
2 February 2016 (02.02.2016) CN
62/298,371 22 February 2016 (22.02.2016) US
16161565.3 22 March 2016 (22.03.2016) EP
- (71) **Applicant: DOLBY LABORATORIES LICENSING CORPORATION** [US/US]; 1275 Market Street, San Francisco, California 94103 (US).
- (72) **Inventors: DICKINS, Glenn N.**; c/o Dolby Australia Pty Ltd, 35 - 51 Mitchell Street, McMahon's Point, Sydney, New South Wales NSW 2060 (AU). **SHI, Dong**; Room 908, No. 2 Lone 567, Chezhan Bei Road, Shanghai, Shanghai 200434 (CN). **GUNAWAN, David**; c/o Dolby Australia Pty Ltd, 35 - 51 Mitchell Street, McMahon's Point, Sydney, New South Wales NSW 2060 (AU).
- (74) **Agents: DOLBY LABORATORIES, INC.** et al.; Intellectual Property Group, 1275 Market Street, San Francisco, California 94103 (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, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))

(54) **Title:** ADAPTIVE SUPPRESSION FOR REMOVING NUISANCE AUDIO

(57) **Abstract:** Described herein is a method and audio capture system for suppressing noise due to mechanical disturbances in an audio capture system. In one embodiment a method includes the steps of: a) receiving an input audio signal from a microphone; b) receiving an input reference signal separate from the input audio signal; c) processing the input reference signal to generate a control signal for identifying a noise event; d) in response to the control signal, estimating a contribution of the noise event to the input audio signal by applying an adaptive filter to a sliding temporal window of the control signal in the frequency domain; and e) based on the estimated contribution, selectively modifying the input audio signal to generate an output audio signal in which the noise arising from the noise event is at least partially suppressed.



ADAPTIVE SUPPRESSION FOR REMOVING NUISANCE AUDIO

TECHNOLOGY

[0001] The present application relates to audio signal processing. More specifically, embodiments of the present invention relate to suppression of acoustic noise due to mechanical disturbance of microphones.

[0002] While some embodiments will be described herein with particular reference to that application, it will be appreciated that the invention is not limited to such a field of use, and is applicable in broader contexts.

BACKGROUND

[0003] Any discussion of the background art throughout the specification should in no way be considered as an admission that such art is widely known or forms part of common general knowledge in the field.

[0004] In audio conferencing, the source of an audio signal is generally a microphone which is sensitive to acoustic stimulus and other mechanically coupled vibration. Generally for capturing audio from participants in a room, the desired acoustic sources are relatively distant, for example more than 500mm from the microphones. However, the nature of a compact device with user interface and the possibility of other noise sources located closer to the device creates the potential for very loud and unwanted signals as detected by the microphone. Specifically, mechanical disturbances, such as the physical manipulation of one of the microphones or the operation of any user interface on an audio conference module, can give rise to associated acoustic noise or nuisance audio signals. Although the acoustic noise radiated into the room from such direct mechanical interference may be low and of no concern to those in the room, the interfering noise signals as detected and amplified by the microphone close to that disturbance can significantly affect far end conference participants. Undesired nearby acoustic interference or unwanted mechanical vibration will be picked up by the microphone as noise and then may be preferentially selected over the desired conversation audio.

[0005] Techniques are known for suppressing nuisance audio signals in systems of microphones positioned at known locations. However, to the inventors' knowledge, there is no known effective technique for suppressing the presence of nuisance audio on a single microphone module where the audio is not related to the output or echo of a known or well characterized stimulus. In particular, there is no work to the knowledge of the inventors that deals with attempting to work adaptively to optimize the suppression and avoid the loss of

desired signal when the interfering nuisance sound is caused by an independent acoustical or mechanical disturbance. This is a particularly important use case where a system is designed to work in proximity to people and activity which may be closer to the capturing device than the desired acoustical object. This is often the case as described in a conference phone placed on a meeting table with people speaking from a range of distance, and some people operating the device or working on the table near the device.

US 8,867,757 B1 discloses system including a first housing that houses a plurality of mechanical keys. A first microphone that is configured to detect a dynamic noise is located within the first housing and under the mechanical keys. A second microphone is configured to detect acoustic waves that include speech and to convert the acoustic waves into an electrical audio signal. The dynamic noise is not associated with the detected speech. The system further includes a dynamic audio signal filter that is configured to suppress, in the electrical audio signal, dynamic noise, and the dynamic audio signal filter is activated in response to the first microphone detecting the dynamic noise.

US 2010/145689 A1 discloses that an audio signal is received that might include keyboard noise and speech. The audio signal is digitized and transformed from a time domain to a frequency domain. The transformed audio is analyzed to determine whether there is likelihood that keystroke noise is present. If it is determined there is high likelihood that the audio signal contains keystroke noise, a determination is made as to whether a keyboard event occurred around the time of the likely keystroke noise. If it is determined that a keyboard event occurred around the time of the likely keystroke noise, a determination is made as to whether speech is present in the audio signal around the time of the likely keystroke noise. If no speech is present, the keystroke noise is suppressed in the audio signal. If speech is detected in the audio signal or if the keystroke noise abates, the suppression gain is removed from the audio signal.

US 2004/213419 A1 discloses various embodiments to reduce noise within a particular environment, while isolating and capturing speech in a manner that allows operation within an otherwise noisy environment. In one embodiment, an array of one or more microphones is used to selectively eliminate noise emanating from known, generally fixed locations, and pass signals from a pre-specified region or regions with reduced distortion.

SUMMARY OF THE INVENTION

[0006] In accordance with a first aspect of the present invention there is provided a method of suppressing noise due to mechanical disturbances in an audio capture system, the method including the steps of:

- a. receiving an input audio signal from a microphone;
- b. receiving an input reference signal separate from the input audio signal;

- c. processing the input reference signal to generate a control signal for identifying a noise event;
- d. in response to the control signal, estimating a contribution of the noise event to the input audio signal by applying an adaptive filter to a sliding temporal window of the control signal in the frequency domain; and
- e. based on the estimated contribution of the noise event, selectively modifying the input audio signal to generate an output audio signal in which the noise arising from the noise event is at least partially suppressed.

[0007] Preferably the method also includes the step of:

b2. analyzing the input reference signal to classify potential noise events likely to be associated with an external disturbance; and

the processing in step c. is based on the classifying in step b2).

[0008] In one embodiment the step of classifying potential noise events includes specifying a noise event when the reference signal power is greater than a predetermined threshold.

[0009] In one embodiment the step of classifying potential noise events includes specifying a noise event when the reference signal contains predetermined spectral characteristics.

[0010] In one embodiment the step of processing the input reference signal includes filtering the input reference signal to generate the control signal. In one embodiment the filtering includes performing signal gating on the reference signal to generate the control signal. The signal gating is preferably triggered by detection of a threshold amplitude of the reference signal. In another embodiment the processing the input reference signal includes extending the spectral content of the reference signal in the spectral domain to generate the control signal.

[0011] In one embodiment the step of processing the input reference signal includes adjusting the dynamic range of the input reference signal to generate the control signal.

[0012] In another embodiment the step of processing the input reference signal includes modulating the input reference signal with an external signal to generate the control signal.

[0013] In a further embodiment the step of processing the input reference signal includes generating a synthetic control signal having characteristics based on the input reference signal.

[0014] In one embodiment, in the absence of detecting a potential noise event, the control signal is the input reference signal. In another embodiment, in the absence of detecting a potential noise event, the control signal is zero.

[0015] In one embodiment the step of processing the input reference signal includes transforming the input reference signal and/or the input audio signal into the frequency domain. In one embodiment transforming the reference signal includes iteratively transforming the reference signal over predetermined time windows to calculate respective temporally evolving reference power spectra indicative of the reference signal, each reference power spectrum including reference power values separated into a plurality of frequency bins. The method preferably further includes the step of iteratively transforming the input audio signal into the frequency domain over the predetermined time windows to calculate respective temporally evolving signal power spectra indicative of the input audio signal, each signal power spectrum including signal power values separated into a like plurality of frequency bins to the reference power spectra. The reference power spectra and the signal power spectra may be banded power spectra where frequency bins are grouped into perceptual bands.

[0016] In one embodiment the predetermined time windows have a period in the range of 5 ms to 500 ms.

[0017] In one embodiment the step of estimating a contribution of the noise event to the input audio signal includes applying the adaptive filter in the frequency domain to the reference power spectra. The adaptive filter may be a linear filter calculating, for each frequency bin, a weighted sum of present and past reference power values in the sliding window for that frequency bin. The step of estimating a contribution of the noise event may further comprise determining parameters of the adaptive filter by comparing the signal power spectra and reference power spectra over a predetermined period of time. The predetermined period of time preferably includes time windows during a detected noise event or shortly before and/or after a detected noise event. In one embodiment the predetermined period of time includes a period covering 5 to 15 predetermined time windows.

[0018] In one embodiment comparing the power spectra includes iteratively performing a spectral subtraction of the reference power spectra from the signal power spectra across different predetermined time windows to determine a relationship between fluctuations of the respective power spectra.

[0019] In one embodiment comparing the power spectra includes performing a least mean squares method on the respective reference power spectra and signal power spectra over the

predetermined period of time. In particular, the least mean squares method may be used to determine filter coefficients for the adaptive filter. For example, the adaptive filter may be updated using a normalized least mean squares method that iteratively updates filter coefficients based on a difference between the signal power spectra and the spectral contribution of the noise event to the input audio signal as estimated by the adaptive filter in the frequency domain.

[0020] In one embodiment the method includes the step of calculating estimated noise event power contribution values for each frequency bin in each spectrum within the predetermined period of time e.g. by applying the adaptive filter on each of the reference power values.

[0021] In one embodiment the step of selectively modifying the audio signal includes calculating suppression gain coefficients for each frequency bin within the predetermined period of time based on a comparison of the estimated noise event power contribution values with the signal power values within the predetermined period of time. Optionally, the noise event power contribution values and the signal power values may be smoothed before the comparison.

[0022] In one embodiment the method includes the step of multiplying the signal power values by the suppression gain coefficients to generate modified signal power values.

[0023] In one embodiment the method includes the step of inverse transforming the modified signal power values into the time domain to produce a modified audio signal.

[0024] In one embodiment the input reference signal is a vibration signal from a vibration sensor. In another embodiment the input reference signal is a detection signal from a button press on a user interface. In a further embodiment the input reference signal originates from a second microphone. In another embodiment input reference signal is a filtered version of an audio signal from the second microphone. In some embodiments the input audio signal and input reference signal are incoherent.

[0025] In accordance with a second aspect of the present invention there is provided an audio capture system for suppressing noise due to mechanical disturbances, the system including:

- a first input port for receiving an input audio signal from a microphone;
- a second input port for receiving an input reference signal separate from the input audio signal; and
- a processor configured for:

processing the input reference signal to generate a control signal for identifying a noise event;

in response to the control signal, estimating a contribution of the noise event to the input audio signal by applying an adaptive filter to a sliding temporal window of the control signal in the frequency domain;

based on the estimated contribution of the noise event, selectively modifying the input audio signal to generate an output audio signal in which the noise arising from the noise event is at least partially suppressed.

[0026] Preferably the processor is further configured for:

analyzing the input reference signal to classify potential noise events likely to be associated with an external disturbance; and

the processing of the input reference signal is based on the classifying.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] Example embodiments of the disclosure will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic functional representation of an audio capture system for suppressing noise due to mechanical disturbances according to an embodiment of the present invention;

Figure 2 is a process flow diagram illustrating the primary steps in a method of suppressing noise due to mechanical disturbances in an audio capture system;

Figure 3 is a process flow diagram of sub-steps of a processing step in the method of Figure 2;

Figure 4 is a schematic illustration of two scenarios of signal gating;

Figure 5 is a process flow diagram of sub-steps of modifying the audio signal in the method of Figure 2;

Figure 6 is a schematic summary functional view of a subset of the steps in the method of Figure 2;

Figure 7 illustrates three embodiments of creating or modifying a reference signal for the adaptive suppressor based on an impulsive activity signal; and

Figure 8 illustrates schematically, the use of a reference signal to modify the envelope of a synthesized noise burst.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Overview of the invention

[0028] The present invention is adapted for adaptively suppressing audio signals based on the simultaneous measurement and processing of a reference signal. The invention includes
5 some links to particular adaptive echo suppression techniques but can be applied to a broader range of applications. Adaptive detection and suppression of echo signals is performed in audio conferencing systems to improve the overall audio quality. In these echo suppression systems, the echo represents a known nuisance audio signal and so the characteristics of the signal to be suppressed are known. By definition, the echo is primarily linearly related to the
10 output or reference signal. In this case an adaptive filter is used to match the relationship that exists between the physical acoustic output of the speaker and a microphone input, and subtract this estimate of the actual signal to reduce the impact of echo.

[0029] In the present invention, access to such a direct reference is not available, and the disturbance is largely incident on the system and created by some external agent or activity.

15 This is a very different application and departs substantially from anything related to acoustic or vibrational echo where the disturbance is created by the device itself or some other related and connected system and audio output. In this case of adaptive signal processing for nuisance signals, both the acoustic signal at the microphone and a form of reference for the disturbance are detected by the device. Due to the similarity of some of the primitive
20 components and signal processing in this approach, it can be seen that aspects of this embodiment may be integrated with the general function of echo and noise removal. It is generally known in the art that signal processing algorithms can be combined within the transform domain rather than only formed as concatenation of processing blocks. The embodiments of the invention described here focus on the generalized adaption for
25 suppression of nuisance audio signals arising from mechanical disturbances.

[0030] A fundamental difference between echo suppression and cancellation is that a suppressor does not subtract an estimate of the echo component from the microphone signal. Rather the suppressor applies an attenuation profile (G_N) in the frequency domain to remove the spectral content that is due to echo. The suppression approach offers the advantage of
30 being more robust to non-linearity and variation in the relationship between the echo reference and microphone pickup. On the other hand, suppressive algorithms will have impact on the desired signal at times where the echo is of a comparable level at any instant in time and frequency.

[0031] Referring to Figure 1 there is illustrated a schematic functional diagram of an audio capture system 1 for suppressing noise due to mechanical disturbances. An example application of system 1 is a central audio mixing module for an audio conference. In one embodiment, system 1 is implemented on the hardware of a personal computer which
5 leverages the associated computer processor.

[0032] System 1 includes a first input port 5 for receiving an input audio signal $y(n)$ from a microphone 7. A second input port 9 is provided for receiving an input reference signal $x(n)$ that is separate from input audio signal $y(n)$. Here the term 'separate' indicates that the reference signal is not derived directly or indirectly from the input audio signal such
10 as a duplicate of the audio signal used in echo suppression. Although the reference signal may include information relating to associated audio objects or origins, the actual signal is captured by a separate sensor, device or input to the microphone used to capture the input audio signal.

[0033] Ports 5 and 9 can be any type of audio connection known in the art such as a
15 standard 3-5 mm audio jack. An example reference signal is a vibration signal received from a vibration sensor 11 co-located within a module 13 together with microphone 7. However, various other types of reference signal are possible, such as a control or detection signal from a button press on a user interface (keyboard presses, mouse clicks etc), an activity signal, audio signal or filtered audio signal from a second microphone. In some embodiments, input
20 audio signal $y(n)$ may be indicative of audio from more than one microphone collocated within module 13. It is envisaged and a general extension that this approach would work with a wider range of spatial microphones, microphone arrays, and reasonable diversity where the one input reference is seen to relate to an expected disturbance at several microphones.

[0034] System 1 is adapted for performing a method 200 as illustrated in Figure 2 and the
25 operation of system 1 will be described with reference to method 200. Steps 201 and 202 relate respectively to the receiving of the input audio signal $y(n)$ and input reference signal $x(n)$ as mentioned above.

[0035] At step 203, reference signal $x(n)$ is analyzed to classify potential noise events that
30 are likely to be associated with the external disturbance. Noise events include mechanical handling of microphone, movement of objects proximate to the microphone, actuation of buttons on an interface associated with system 1 or reverb associated with a second microphone. Classification may be a binary ('noise event' or 'no noise event') classification

by specifying a noise event when the reference signal power is greater than a predetermined threshold, or a higher order classification to classify the type or degree of potential noise events such as when the reference signal contains predetermined spectral characteristics. Classifying may include creating a continuous probability measure (say something that varies from 0 to 1).

[0036] At step 204, based on the classifying in step 203, the input reference signal is processed to generate a reference control signal $\hat{x}(t)$. This processing includes several sub-steps as illustrated in Figure 3. At sub-step 204a, signal preprocessing is performed including amplification by amplifier 15 and analog to digital conversion by ADC 17.

[0037] At sub-step 204b, a set of signal processing is carried out to generate the reference control signal. The signal processing includes one or more of:

- a) Filtering the reference signal in the form of linear or non-linear operations on the sampled audio;
- b) Dynamic range expansion of the reference signal including compression, expansion or gating in order to selectively pass more or less of the signal based on some local measure of the signal activity;
- c) Classification of reference signal activity and events in order to synthetically create the reference control signal that is more representative of a set of events rather than specific mechanical or acoustical vibration; or
- d) Substantially replacing the reference signal or spectral estimation of the reference source with an appropriate reference control signal based on classification, non-linear amplitude mapping.

[0038] In the absence of detecting a potential noise event, the control signal may be set to be the input reference signal or set to zero. The general rationale and goal of the processing in step 204b is towards creating a broadband stimulation of the adaptive suppression system when significant nuisance is likely to be present at the acoustic microphone pickup.

[0039] For example, in the case of signal gating used in some embodiments reference is made to the process of selectively muting, deactivating or ignoring an input signal if the amplitude of that signal falls below a predetermined threshold. This concept is illustrated schematically in the context of the present invention in Figure 4. The top panel of Figure 4 illustrates a scenario where mechanical handling of module 13 is in progress. This mechanical handling gives rise to significant mechanical vibrations which are detected by vibration sensor 11, which generates reference signal $x(n)$ as a corresponding vibration signal. As the detected vibrations are significant, the reference signal exceeds the gating

threshold and the reference signal is further processed. The lower panel of Figure 4 illustrates the normal scenario of a conference participant speaking into microphone 7. This speaking produces only small mechanical vibrations which fall below the predetermined gating threshold. The gating threshold is preferably calibrated such that the threshold is significantly higher than the normal acoustic vibrations detected during speech.

[0040] For example, in one embodiment, it was found that for a contact microphone signal levels as referred in a digital domain of up to -30dB full scale were apparent in the presence of acoustical activity, whilst vibration activity related to nearby contact or nuisance achieved levels substantially above this. By inspection, a threshold was set such that any pickup of the contact (mechanically coupled) microphone was ignored to the point of a user shouting loudly centimeters from the device. This threshold of -30dB then allowed for significant mechanical activity to pass through and be utilized for adaptively reducing the impact of this nuisance on the microphone capture.

[0041] By gating in this way, the reference signal becomes largely an impulse and very much non-linearly related to the audio content in the input audio signal. Furthermore, much of the actual signal coming from the vibration sensor may be replaced with a gated or amplitude modulated broadband noise. The reason and benefit for this will be explained below.

[0042] In alternative embodiments, step 203 can be skipped such that the processing of step 204 is performed directly on the input reference signal.

[0043] Returning to Figure 3, at sub-step 204c, the reference control signal $\hat{x}(t)$ is divided into time windows of a predefined time duration (window size). Exemplary window sizes are in the range of 5 ms to 500 ms and preferably about 20 ms. At sub-step 203e, each time window is transformed into the frequency domain by a form of transform or subband filtering (for example the Short-Time Fourier Transform (STFT)). The result is a reference power spectrum X_N indicative of the power-frequency distribution of the reference signal during each time window. Each reference power spectrum includes reference power values separated into a plurality of frequency bins. At sub-step 204e, the frequency bins may be grouped into B perceptual bands, spaced on a perceptually motivated scale with an inter-band distance of approximately 1 Bark. In other embodiments, the frequency bins may be grouped into bands based on other acoustic scales. A typical embodiment has 30-40 bands over the operational audio frequency range. Given the frequency domain coefficients X_k of the reference power spectrum, the banded power values are computed as

$$X_i = \sum_{k=0}^{N-1} w_{i,k} |X_k|^2$$

[0044] where $w_{i,k}$ represents the k^{th} bin weight for the i^{th} spectral band. This has many advantages, including reduced computational complexity, matching the critical bandwidth of the human auditory system, and a robustness to slight spectral variations or sampling rate shifts between the input and output. Step 204 is performed iteratively on the reference signal to calculate respective reference power spectra indicative of the temporal spectral evolution of the reference signal. The resulting banded reference power spectrum is designated as X_B .

[0045] Returning to Figure 2, at step 205, the input audio signal $y(n)$ is similarly processed by an amplifier 19 and ADC 21, and transformed to the frequency domain in a similar manner to the reference signal $x(n)$ described above. That is, the input audio signal is divided into time windows having the same start times, end times and duration as the windows for the reference signal. Each time window of the input audio signal is transformed into the frequency domain to form signal power spectra Y_N indicative of the power-frequency distribution of the input audio signal during each time window. Each signal power spectrum includes signal power values separated into a plurality of frequency bins similar to the reference power spectra. The resulting banded signal spectrum is designated as Y_B . The subsequent smoothed signal power estimates of Y_B are designated as S_B . By way of an example embodiment, the previous 3 paragraphs set out one approach for creating a set of reasonably spaced frequency selective estimators for the power in each of the input and where required also the processed (output of 203) reference signals. Such approaches are well known in the field and may be referred to, for example and not exhaustively, filter-banks, frequency analysis, lapped transforms, sub-band filtering, and many forms of banded power short-time frequency transform.

[0046] At step 206, in response to the control signal the contribution of the noise event to the input audio signal is estimated. This includes comparing the signal power spectra and reference power spectra over a predetermined period of time using an adaptive filter procedure. In this procedure, the current and previous time samples of the power spectrum of the reference signal are passed through a filter which is adaptively estimated and this is used to drive a spectral suppressor via the application of an attenuation profile across frequency to the captured input audio signal.

[0047] In one embodiment, step 206 is performed by the reference signal being passed through a perceptual band domain block adaptive filter with L filter taps. The output (E_B)

represents the estimated contribution of the reference power representing the noise event to the input audio signal for the current time window (frame) of audio.

$$E_B = \sum_{l=0}^{L-1} F_{B,l} X_{B,l}$$

[0048] Where $F_{B,l}$ is the vector of B filter taps for the l^{th} time lag, and $X_{B,l}$ is the vector of

5 B banded reference powers for the l^{th} time lag. The linear filter estimates, separately for each band, the contribution of the reference power to the noise event in the input audio signal by a weighted sum of the present reference power of the band and previous reference powers of the band within a sliding temporal window defined by L . The sliding temporal window is applied, separately for each band, to banded reference power values $X_{B,l}$. A typical value of

10 L for a 20ms frame size is 10 providing an overall temporal window of 200ms for estimating the contribution of the noise event to the input audio signal. In the case of nuisance and disturbance that is close to the capturing microphone, it is reasonable that this length be kept short in order to improve the filter stability. Provided that the bulk delay is aligned between the reference signal and the first expected disturbance at the microphone, then the filter can

15 be shortened. Note that the effects of acoustic reflections in the room are largely irrelevant in this application as we are dealing with disturbance that is generally near one or more microphones and whilst significant at that microphone is not a problem to the wider room. Nuisance signals that are loud enough in the room for the reverberation time of the room to be a factor would generally be noticed and avoided by the participants. A further note is that

20 the filtering need not be causal, in that where the reference signal is expected to arrive with some delay after the microphone disturbance, a suitable delay can be applied to the normal microphone signal in order to align the reference and microphone activity. Whilst undesirable, such delays are necessary and may be further controlled based on the ongoing presence and extent of nuisance as detected at a given microphone and reference. In this

25 exemplary implementation, the adaptive filter is updated using the normalised least mean squares (NLMS) method. The update of the filter parameters (e.g. filter tabs) may be for each frame or multiple frames such as the temporal window of the filter. However, it will be appreciated that other iterative data fitting routines can be implemented. Using the NLMS method, an updated filter coefficient $F'_{B,l}$ is calculated from the previous coefficient $F_{B,l}$ with

30 the following update rule where μ is a parameter to control the update speed:

$$F'_{B,l} = F_{B,l} + \mu \frac{(Y_B - E_B) X_{B,l}}{\sum_{l'=0}^{L-1} X_{B,l'}^2 + X_{\text{sens}}}$$

[0049] The parameter X_{sens} is used to ensure the adaption is stable for low signal levels (typical values of X_{sens} are in the order of -85dB). The estimated contribution of the reference power E_B to the noise event in the input audio signal is applied with an asymmetric smoothing rule with instantaneous attack and release controlled by a constant α_N (typical value of 0.5 for a 20ms buffer size) to calculate smoothed estimates of the noise event power contribution N_B indicative of the spectral power contribution of the noise event to the input audio signal. The subsequent estimate of the smoothed noise event power contribution N'_B for the next window is calculated from the previous estimate N_B :

$$N'_B = \begin{cases} \text{if } E_B \geq N_B : & N'_B = E_B \\ \text{if } E_B < N_B : & N'_B = \alpha_N E_B + (1 - \alpha_N) N_B \end{cases}$$

[0050] At step 207, the input audio signal is selectively modified to generate an output audio signal in which the noise arising from the noise event is at least partially suppressed. This step includes a number of sub-steps as illustrated in Figure 5. At sub-step 207a, the smoothed noise event power contribution values N_B are compared with the corresponding smoothed signal power values S_B at each time window. At sub-step 207b, this comparison is used to calculate suppression gain coefficients G_B for an attenuation profile G_N :

$$G_B = \min \left(G_{\min}, \frac{\max(0, S_B - \gamma_N N_B)}{S_B} \right)^2$$

[0051] The parameter γ_N is an over subtraction factor to control the aggressiveness of the echo suppression (typical value of 3dB). Parameter $G_{\min}$ defines a minimum attenuation that will be applied to each band (typical valid of -60dB). The banded gain coefficients are interpolated to produce a set of N frequency bin gains defining the attenuation profile G_N by inverting the mix matrix used to form the bands.

[0052] At sub-step 207c, the attenuation profile G_N is multiplied by complex frequency bin vector Y_N to derive modified signal power values. Finally, at sub-step 207d, and the resulting data is passed through an inverse transform to produce a modified audio signal in the time domain.

[0053] Steps 206 and 207 of method 200 can be performed in software by an adaptive suppressor routine, as illustrated in Figure 1. The functional steps of such an adaptive suppression routine are summarized schematically in Figure 6. The adaptive suppressor will converge to find a gain and temporal relationship between any fluctuations in the reference

power spectra and signal power spectra, and from this determine gains where there is a predicted and seemingly causal relationship. That is; where activity on the reference seems to be followed by energy with a correlated temporal and spectral envelope on the input. This component is able to be tuned to be quite robust to errors in spectral activity, and by nature will adapt to estimate the temporal relationship between the two signals.

[0054] In other works it has been observed that specific types of impulsive noises, for example button clicks, have a fairly consistent spectral shape and often consistent amplitude. Whilst at first it may seem that the gating and conversion of the reference to something broadband and impulsive (or amplitude modulated) is ineffective and problematic for the adaptive suppressor, in fact gating the reference signal is advantageous and the following is noted:

- The adaptive suppressor can converge reasonably quickly, and when it is enabled by energy on the reference, will be aggressive by design in finding and suppressing temporally related content. Thus, gating the reference signal can help to more quickly identify the noise contribution to the input audio signal and suppress that noise.
- Whilst the above technique is aggressive, it is also much more selective in both time and frequency than conventional and prior art methods which may totally attenuate or suppress the captured audio signal for a conservative period of time when a noise event is detected.
- Impulsive and handling noise is often repeated and similar in nature (for example repeated taps on a screen or of a volume button) and hence the learning process of the adaptive suppressor can be very effective in determining a good estimation of the suppression.
- The depth and regularization of the suppression can be tuned to allow a maximal level of suppression rather than complete distortion of the input signal, allowing the system to perform well when there is both a disturbance (as indicated by the gated reference) and some other desired signal. Specifically, the system may choose to limit the suppression of nuisance signals as identified in the adaptive system by an amount. In one embodiment a limit on the suppression of 10 dB was found to have a good compromise of reducing the impact of the nuisance without sacrificing ongoing audio capture in the room.

[0055] It will be understood that the above methodology may be performed simultaneously on multiple input audio signals from different microphones in an audio conferencing system.

Similarly, the above methodology may be performed simultaneously using multiple reference signals adapted for detecting different distinct types of mechanical disturbances.

Applications and extensions of the invention

Impulsive Reference - Discrete Button / User Interface Activity

5 [0056] In many situations, for example pressing of a volume button, there is a direct and unambiguous indication of the mechanical activity in the device from the resulting actuation signal. This type of signal and activity creates a single valued reference signal, though quite often represents a very consistent spectral, temporal and amplitude disturbance with a temporal relationship to the button activity signal.

10 [0057] This case is labeled an 'impulsive' reference, where it has been already thresholded or gated so that there is a very high confidence that the event in this signal is or will be associated with some disturbance or nuisance detection at the microphones.

[0058] For suppressing impulsive references, additional signal processing may be advantageous over and above that described above. As discussed above, the adaptive filter
15 will adjust to match the temporal presence of spectral energy in its output (e.g. E_B), where there is some excitation in the same spectral bands in its input (e.g. X_B). Therefore, it may be advantageous to process impulsive references by modifying them to create a broadband excitation. The following embodiments are possible and illustrated in Figure 7:

a) The impulse itself is broadband, however if coming from a button or key scan code it
20 may have a different duration for certain presses. The mechanical noise of a button is generally associated with the press action (and also the release). In this scenario, the rising and/or falling edge of a particular button press is modified into a fixed length to define a constant energy reference signal. This is illustrated in the top scenario of Figure 7.

25 b) Each noise event designated by a rising and/or falling edge is replaced with a burst of noise. This may be broadband or alternatively shaped to match the expected spectra of the disturbance. Since the suppressive adaptive filter is only sensitive to the power spectra, it is possible to use an exemplary audio capture of the disturbance itself. Such an approach can aid in the adaption and suppression of the disturbance. This is
30 illustrated in the middle scenario in Figure 7. It is a reasonable and expected extension that different levels of burst or synthetic signal may be used for each of the rising and falling edges, and where possible the use of different scan codes or characteristics of key events may be further used through the use of different reference stimulus or by use of the reference input to an alternate adaptive filter.

- c) Since the goal here is to excite the reference power spectrum, it is feasible to skip the input transform step of the reference signal and proceed directly to an impulsive activity on the banded reference power spectrum X_B (labeled as 3 in Figure 6). Again the spectra here may be broadband (unity or other across all bands) or shaped to match the expected or an exemplary sample of the expected disturbance correlated with the event. This is illustrated in the bottom scenario in Figure 7.

Amplitude Varying Impulsive Reference

[0059] In some instances the indication of activity may include some representation of amplitude that would correlate with the power as seen by the microphone(s) related to the disturbance. For example, some devices such as a piano keyboard, include velocity sensors on the keys, and other UIs such as a touch screen may estimate contact velocity and/or force applied to the screen.

[0060] In such cases, the reference is not necessarily an impulse, and the indication of activity may not directly map to the power level of the disturbance. In these cases, often a simple mapping can be empirically obtained through observation, trial tuning or machine learning. For the case of a stateless mapping from the indication of activity to the expected level of the disturbance, the embodiment illustrated in Figure 8 can be utilized, whereby the incoming reference signal that is already gated, is used to modify the envelope of a synthesized noise burst having an appropriate spectral range and potentially shape to match the expected disturbance.

Extensions on Potential Source of the Reference Control

[0061] In the previous two sections, the suggested embodiments have extended the approach to start with inputs into the suppressive system that are no longer specifically related or coherent with the microphone response and activity caused by the disturbance. A wide set of extensions to the invention are envisaged where the indication of likely disturbance at the microphone is substantially extended.

[0062] In each case, the suggested input can be reasonably transformed through gating, feature extraction, classification, stateless mapping and even temporal filtering. In the most general example, some indication is provided that an event has happened, and in principle a component is suggested that makes an approximation of what sound or mechanical disturbance would relate to that event. Noting that the approximation of the signal need not be exact in any sense of spectral envelope or temporal modulation, and the best results are often obtained with a simple and short (20ms) impulsive excitation of the reference signal.

Self Noise

[0063] In some cases, the vibration corrupting the microphone may be created by activity within the device itself. Some examples here include:

- Vibration or other haptic feedback to the user;
- 5 ➤ Mechanical manipulation such as lens adjustment or focusing etc.; or
- Haptic and vibrational feedback devices and actuators.

[0064] In these cases, the control signal to the output mechanical transducer can be used as a reference. As in other embodiments, the extension of the signal can be performed to create a suitable excitation of reference that could match all possible signal components
10 caused by the mechanical activity.

[0065] In the diagrams and descriptions above, it is reasonable to infer that all acoustic microphones may in fact be one or more microphones in each module, allowing for the individual capture of spatial audio from each microphone module, and the overall output of a multichannel audio representation of the scene. Generally, only a single vibration sensor
15 would be required for each mechanically separate microphone module.

Pre-Emptive Estimation

[0066] Using the present invention, there is also the potential for other methods of detecting or pre-emptively detecting handling and vibration. Examples include:

- accelerometer for detecting device movement and reorientation (for example a
20 falling device prior to hitting a hard surface);
- capacitive sensor that can detect proximate hand or user gestures;
- visual or optical sensing to detect movement approaching or nearby the device and choose to remove it from potential selection or audio input in the mix; or
- acoustic ultrasonic ranging from the microphone and some output transducer
25 (piezo or speaker) based on Doppler or ranging techniques.

[0067] Given the adaptive process of the invention, there will be a natural reduction in the estimated gain of the disturbance within the adaptive filter N_B if there are repeated occurrences of the reference signal occurring without an associated activity at the microphone. This can be managed with tuning and adaption constants. In particular a slow
30 adaption can deal with false positives, an asymmetric adaption can have greater retention of events that did correlate with high energy at the microphone, the overestimation parameter γ_B in the suppression equation can be set high.

Keyboard Activity Detection

[0068] The idea of using a keyboard scan code to suppress or gate the microphone has been seen in some prior art, however the approaches are generally coarse gating of the microphone signal in order to effectively mute microphone capture when the user is typing.

5 The present invention is able to provide an improved solution to this problem:

➤ The system is able to learn the typical characteristics of the vibrational coupling, resonance and general disturbance arising from a keyclick, this adaption removes the need for any apriori knowledge of key click noises or training.

10 ➤ The specific key scan code can be used, for example, to give an indication of proximity to a microphone in the system (e.g. laptop) and again provide a more selective suppression avoiding the removal of desired input.

➤ There is indication that different users will have characteristic pressure and stroke velocity of keys, and this will be naturally learned by the adaptive system.

15

Mouse Click Detection

[0069] Similar to the detection of keyboard strokes, the present invention can also be used to detect the event of mouse clicking by using the 'clicking' flag from the system as the input to the adaptive system. Compared to a keyboard, mouse clicking usually gives a more consistent sound, regardless of the users. Also, the number of buttons on a mouse is much fewer than that of a keyboard, resulting easier classification and learning of the spectral curve.

20

Feature Extraction Detection

[0070] In the case of a microphone module having more than one microphone, it can be noted that a degree of vibration sensing can be obtained from consideration of the covariance or spatial parameters detected by those microphones. In particular, for microphones that face different directions or are somewhat separated, any acoustical signal should present differently at the microphones whilst mechanical vibration may appear quite correlated and time aligned. In this way, where the microphone module includes more than one acoustic microphone, it is possible in this invention to derive the detection of a vibration signal from some filtering or feature detection on the microphone array. At this point, such a signal represents an indication of vibration and can be employed in a similar way to the impulsive reference signal detailed elsewhere in this invention.

25

30

Conclusions

[0071] It will be appreciated that the above described invention provides significant systems and methods for suppressing acoustic noise associated with events of mechanical disturbance.

5 [0072] The present invention is a novel approach to the suppressive removal of nuisance audio signals where there is some additional reference signal that is temporally correlated with the nuisance but not coherent or predominantly linearly related to the disturbance signal occurring at the microphone. In particular, a key problem is the nuisance audio signal detected at the microphone due to activity around, near or on the microphone module itself.
10 When the reference signal originates from a vibration sensor collocated within a microphone module, mechanical noise arising from the handling of the microphone module or the operation of any user interface and buttons that may be located on the module can be detected and suppressed efficiently.

Interpretation

15 [0073] Unless specifically stated otherwise, as apparent from the following discussions, it is appreciated that throughout the specification discussions utilizing terms such as "processing," "computing," "calculating," "determining", analyzing" or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as physical, such as electronic,
20 quantities into other data similarly represented as physical quantities.

[0074] In a similar manner, the term "processor" may refer to any device or portion of a device that processes electronic data, e.g., from registers and/or memory to transform that electronic data into other electronic data that, e.g., may be stored in registers and/or memory. A "computer" or a "computing machine" or a "computing platform" may include one or more
25 processors.

[0075] The methodologies described herein are, in one embodiment, performable by one or more processors that accept computer-readable (also called machine-readable) code containing a set of instructions that when executed by one or more of the processors carry out at least one of the methods described herein. Any processor capable of executing a set of
30 instructions (sequential or otherwise) that specify actions to be taken are included. Thus, one example is a typical processing system that includes one or more processors. Each processor may include one or more of a CPU, a graphics processing unit, and a programmable DSP unit. The processing system further may include a memory subsystem including main RAM and/or a static RAM, and/or ROM. A bus subsystem may be included for communicating

between the components. The processing system further may be a distributed processing system with processors coupled by a network. If the processing system requires a display, such a display may be included, e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT) display. If manual data entry is required, the processing system also includes an input device such as one or more of an alphanumeric input unit such as a keyboard, a pointing control device such as a mouse, and so forth. The term memory unit as used herein, if clear from the context and unless explicitly stated otherwise, also encompasses a storage system such as a disk drive unit. The processing system in some configurations may include a sound output device, and a network interface device. The memory subsystem thus includes a computer-readable carrier medium that carries computer-readable code (e.g., software) including a set of instructions to cause performing, when executed by one or more processors, one of more of the methods described herein. Note that when the method includes several elements, e.g., several steps, no ordering of such elements is implied, unless specifically stated. The software may reside in the hard disk, or may also reside, completely or at least partially, within the RAM and/or within the processor during execution thereof by the computer system. Thus, the memory and the processor also constitute computer-readable carrier medium carrying computer-readable code.

[0076] Furthermore, a computer-readable carrier medium may form, or be included in a computer program product.

[0077] In alternative embodiments, the one or more processors operate as a standalone device or may be connected, e.g., networked to other processor(s), in a networked deployment, the one or more processors may operate in the capacity of a server or a user machine in server-user network environment, or as a peer machine in a peer-to-peer or distributed network environment. The one or more processors may form a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine.

[0078] Note that while diagrams only show a single processor and a single memory that carries the computer-readable code, those in the art will understand that many of the components described above are included, but not explicitly shown or described in order not to obscure the inventive aspect. For example, while only a single machine is illustrated, the term "machine" shall also be taken to include any collection of machines that individually or

jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0079] Thus, one embodiment of each of the methods described herein is in the form of a computer-readable carrier medium carrying a set of instructions, e.g., a computer program that is for execution on one or more processors, e.g., one or more processors that are part of web server arrangement. Thus, as will be appreciated by those skilled in the art, embodiments of the present invention may be embodied as a method, an apparatus such as a special purpose apparatus, an apparatus such as a data processing system, or a computer-readable carrier medium, e.g., a computer program product. The computer-readable carrier medium carries computer readable code including a set of instructions that when executed on one or more processors cause the processor or processors to implement a method. Accordingly, aspects of the present invention may take the form of a method, an entirely hardware embodiment, an entirely software embodiment or an embodiment combining software and hardware aspects. Furthermore, the present invention may take the form of carrier medium (e.g., a computer program product on a computer-readable storage medium) carrying computer-readable program code embodied in the medium.

[0080] The software may further be transmitted or received over a network via a network interface device. While the carrier medium is shown in an example embodiment to be a single medium, the term "carrier medium" should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term "carrier medium" shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by one or more of the processors and that cause the one or more processors to perform any one or more of the methodologies of the present invention. A carrier medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical, magnetic disks, and magneto-optical disks. Volatile media includes dynamic memory, such as main memory. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise a bus subsystem. Transmission media also may also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications. For example, the term "carrier medium" shall accordingly be taken to include, but not be limited to, solid-state memories, a computer product embodied in optical and magnetic media; a medium bearing a propagated signal detectable by at least one processor or one or more processors and representing a set of instructions that, when executed, implement a

method; and a transmission medium in a network bearing a propagated signal detectable by at least one processor of the one or more processors and representing the set of instructions.

[0081] It will be understood that the steps of methods discussed are performed in one embodiment by an appropriate processor (or processors) of a processing (e.g., computer) system executing instructions (computer-readable code) stored in storage. It will also be understood that the invention is not limited to any particular implementation or programming technique and that the invention may be implemented using any appropriate techniques for implementing the functionality described herein. The invention is not limited to any particular programming language or operating system.

[0082] Reference throughout this specification to “one embodiment”, “some embodiments” or “an embodiment” means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the present disclosure. Thus, appearances of the phrases “in one embodiment”, “in some embodiments” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures or characteristics may be combined in any suitable manner, as would be apparent to one of ordinary skill in the art from this disclosure, in one or more embodiments.

[0083] As used herein, unless otherwise specified the use of the ordinal adjectives “first”, “second”, “third”, etc., to describe a common object, merely indicate that different instances of like objects are being referred to, and are not intended to imply that the objects so described must be in a given sequence, either temporally, spatially, in ranking, or in any other manner.

[0084] In the claims below and the description herein, any one of the terms comprising, comprised of or which comprises is an open term that means including at least the elements/features that follow, but not excluding others. Thus, the term comprising, when used in the claims, should not be interpreted as being limitative to the means or elements or steps listed thereafter. For example, the scope of the expression a device comprising A and B should not be limited to devices consisting only of elements A and B. Any one of the terms including or which includes or that includes as used herein is also an open term that also means including at least the elements/features that follow the term, but not excluding others. Thus, including is synonymous with and means comprising.

[0085] It should be appreciated that in the above description of example embodiments of the disclosure, various features of the disclosure are sometimes grouped together in a single embodiment, Fig., or description thereof for the purpose of streamlining the disclosure and

aiding in the understanding of one or more of the various inventive aspects. This method of disclosure, however, is not to be interpreted as reflecting an intention that the claims require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive aspects lie in less than all features of a single foregoing disclosed embodiment.

5 Thus, the claims following the Detailed Description are hereby expressly incorporated into this Detailed Description, with each claim standing on its own as a separate embodiment of this disclosure.

[0086] Furthermore, while some embodiments described herein include some but not other features included in other embodiments, combinations of features of different embodiments
10 are meant to be within the scope of the disclosure, and form different embodiments, as would be understood by those skilled in the art. For example, in the following claims, any of the claimed embodiments can be used in any combination.

[0087] In the description provided herein, numerous specific details are set forth. However, it is understood that embodiments of the disclosure may be practiced without these
15 specific details. In other instances, well-known methods, structures and techniques have not been shown in detail in order not to obscure an understanding of this description.

[0088] Similarly, it is to be noticed that the term coupled, when used in the claims, should not be interpreted as being limited to direct connections only. The terms "coupled" and "connected," along with their derivatives, may be used. It should be understood that these
20 terms are not intended as synonyms for each other. Thus, the scope of the expression a device A coupled to a device B should not be limited to devices or systems wherein an output of device A is directly connected to an input of device B. It means that there exists a path between an output of A and an input of B which may be a path including other devices or means. "Coupled" may mean that two or more elements are either in direct physical,
25 electrical or optical contact, or that two or more elements are not in direct contact with each other but yet still co-operate or interact with each other.

[0089] Thus, while there has been described what are believed to be the best modes of the disclosure, those skilled in the art will recognize that other and further modifications may be made thereto without departing from the spirit of the disclosure, and it is intended to claim all
30 such changes and modifications as fall within the scope of the disclosure. For example, any formulas given above are merely representative of procedures that may be used. Functionality may be added or deleted from the block diagrams and operations may be interchanged among functional blocks. Steps may be added or deleted to methods described within the scope of the present disclosure.

CLAIMS

We claim:

1. A method of suppressing noise due to mechanical disturbances in an audio capture system,
5 the method including the steps of:
 - a. receiving an input audio signal from a microphone;
 - b. receiving an input reference signal separate from the input audio signal;
 - c. processing the input reference signal to generate a control signal for identifying a
noise event;
 - 10 d. in response to the control signal, estimating a contribution of the noise event to the
input audio signal by applying an adaptive filter to a sliding temporal window of
the control signal in the frequency domain; and
 - e. based on the estimated contribution of the noise event, selectively modifying the
input audio signal to generate an output audio signal in which the noise arising
15 from the noise event is at least partially suppressed.
2. A method according to claim 1 including the step of:
 - b2. analyzing the input reference signal to classify potential noise events likely to be
associated with an external disturbance; and
 - 20 wherein the processing in step c is based on the classifying in step b2).
3. A method according to claim 2 wherein the step of classifying potential noise events
includes specifying a noise event when the reference signal power is greater than a
predetermined threshold.
25
4. A method according to claim 2 or claim 3 wherein step of classifying potential noise
events includes specifying a noise event when the reference signal contains predetermined
spectral characteristics.
- 30 5. A method according to any one of claims 1 to 4 wherein the step of processing the input
reference signal includes filtering the input reference signal to generate the control signal.
6. A method according to claim 5 wherein the filtering includes performing signal gating on
the reference signal to generate the control signal.

7. A method according to claim 6 wherein the signal gating is triggered by detection of a threshold amplitude of the reference signal.

5 8. A method according to any one of claims 1 to 7 wherein the step of processing the input reference signal includes extending the spectral content of the reference signal in the spectral domain to generate the control signal.

9. A method according to any one of claims 1 to 8 wherein the step of processing the input
10 reference signal includes adjusting the dynamic range of the input reference signal to generate the control signal.

10. A method according to any one of claims 1 to 9 wherein the step of processing the input reference signal includes modulating the input reference signal with an external signal to
15 generate the control signal.

11. A method according to any one of claims 1 to 10 wherein the step of processing the input reference signal includes generating a synthetic control signal having characteristics based on the input reference signal.
20

12. A method according to any one of claims 2 to 11 wherein, in the absence of detecting a potential noise event, the control signal is the input reference signal.

13. A method according to any one of claims 2 to 11 wherein, in the absence of detecting a
25 potential noise event, the control signal is zero.

14. A method according to any one of the preceding claims wherein the step of processing the input reference signal includes transforming the input reference signal into the frequency domain.
30

15. A method according to claim 14 wherein transforming the reference signal includes iteratively transforming the reference signal over predetermined time windows to calculate respective temporally evolving reference power spectra indicative of the reference signal,

each reference power spectrum including reference power values separated into a plurality of frequency bins.

16. A method according to any one of the preceding claims including the step of transforming
5 the input audio signal into the frequency domain.

17. A method according to claim 16 wherein transforming the input audio signal includes
iteratively transforming the input audio signal into the frequency domain over the
predetermined time windows to calculate respective temporally evolving signal power spectra
10 indicative of the input audio signal, each signal power spectrum including signal power
values separated into a like plurality of frequency bins to the reference power spectra.

18. A method according to claim 16 or 17 wherein the step of estimating a contribution of the
noise event to the input audio signal includes determining parameters of the adaptive filter by
15 comparing the signal power spectra and reference power spectra over a predetermined period
of time.

19. A method according to claim 18 wherein the predetermined period of time includes time
windows during a detected noise event or shortly before and/or after a detected noise event.
20

20. A method according to claim 18 or claim 19 wherein the predetermined period of time
includes a period covering 5 to 15 predetermined time windows.

21. A method according to any one of claims 18 to 20 wherein comparing the power spectra
25 includes iteratively performing a spectral subtraction of the reference power spectra from the
signal power spectra across different predetermined time windows to determine a relationship
between fluctuations of the respective power spectra.

22. A method according to any one of claims 18 to 21 wherein comparing the power spectra
30 includes performing a least mean squares method on the respective reference power spectra
and signal power spectra over the predetermined period of time.

23. A method according to any one of claims 14 to 22 including the step of calculating estimated noise event power contribution values for each frequency bin in each spectrum within the predetermined period of time.

5 24. A method according to claim 23 wherein the step of selectively modifying the audio signal includes calculating suppression gain coefficients for each frequency bin within the predetermined period of time based on a comparison of the estimated noise event power contribution values with the signal power values within the predetermined period of time.

10 25. A method according to claim 24 including the step of multiplying the signal power values by the suppression gain coefficients to generate modified signal power values.

26. A method according to claim 25 including the step of inverse transforming the modified signal power values into the time domain to produce a modified audio signal.

15

27. A method according to any one of the preceding claims wherein the input reference signal is a vibration signal from a vibration sensor.

28. A method according to any one of claims 1 to 26 wherein the input reference signal is a
20 detection signal from a button press on a user interface.

29. A method according to any one of claims 1 to 26 wherein the input reference signal originates from a second microphone.

25 30. A method according to claim 29 wherein the input reference signal is a filtered version of an audio signal from the second microphone.

31. A method according to any one of the preceding claims wherein the input audio signal and input reference signal are incoherent.

30

32. An audio capture system for suppressing noise due to mechanical disturbances, the system including:

a first input port for receiving an input audio signal from a microphone;

a second input port for receiving an input reference signal separate from the input audio signal; and

a processor configured for:

processing the input reference signal to generate a control signal for identifying a noise event;

5

in response to the control signal, estimating a contribution of the noise event to the input audio signal by applying an adaptive filter to a sliding temporal window of the control signal in the frequency domain;

based on the estimated contribution of the noise event, selectively modifying the input audio signal to generate an output audio signal in which the noise arising from the noise event is at least partially suppressed.

10

33. An audio capture system according to claim 32 wherein the processor is further configured for:

analyzing the input reference signal to classify potential noise events likely to be associated with an external disturbance; and

15

wherein the processing of the input reference signal is based on the classifying.






FIG. 1

2 / 8



3 / 8

**FIG. 3**

4 / 8





FIG. 4

5 / 8

**FIG. 5**






FIG. 6





INTERNATIONAL SEARCH REPORT

International application No

PCT/US2017/016274

A. CLASSIFICATION OF SUBJECT MATTER
INV. G10L21/0208 H04M9/08
ADD.

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)

G10L H04M

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 8 867 757 B1 (OOI LENG [US]) 21 October 2014 (2014-10-21) column 1, line 20 - line 26 column 1, line 34 - column 2, line 61 column 4, line 20 - line 24 column 5, line 2 - column 6, line 28 column 6, line 53 - column 8, line 26 -----	1-33
X	US 2013/132076 A1 (YANG JUN [US] ET AL) 23 May 2013 (2013-05-23) paragraphs [0030], [0053] - [0058], [0068] - [0074] -----	1-33
A	US 2010/145689 A1 (LI QIN [US] ET AL) 10 June 2010 (2010-06-10) paragraphs [0015] - [0039]; figures 1,2 ----- -/--	1-33



Further documents are listed in the continuation of Box C.



See patent family annex.

* 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

"&" document member of the same patent family

Date of the actual completion of the international search

30 March 2017

Date of mailing of the international search report

06/04/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Sorrentino, Andrea

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2017/016274

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004/213419 A1 (VARMA ANKUR [US] ET AL) 28 October 2004 (2004-10-28) paragraphs [0051] - [0066]; figures 3-6 -----	1-33

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/016274

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8867757	B1	21-10-2014	NONE
US 2013132076	A1	23-05-2013	NONE
US 2010145689	A1	10-06-2010	NONE
US 2004213419	A1	28-10-2004	US 2004213419 A1 28-10-2004
			US 2009175462 A1 09-07-2009