




(22) **Date de dépôt/Filing Date:** 2013/02/15
 (41) **Mise à la disp. pub./Open to Public Insp.:** 2013/08/16
 (45) **Date de délivrance/Issue Date:** 2016/07/19
 (30) **Priorité/Priority:** 2012/02/16 (US61/599,762)

(51) **Cl.Int./Int.Cl. G10L 19/03** (2013.01)
 (72) **Inventeurs/Inventors:**
 HETHERINGTON, PHILLIP ALAN, CA;
 LI, XUEMAN, CA
 (73) **Propriétaires/Owners:**
 QNX SOFTWARE SYSTEMS LIMITED, CA;
 2236008 ONTARIO INC., CA
 (74) **Agent:** RIDOUT & MAYBEE LLP

(54) **Titre : SYSTEME ET PROCEDURE DE MISE EN FORME DE BRUIT RESIDUEL DYNAMIQUE**
 (54) **Title: SYSTEM AND METHOD FOR DYNAMIC RESIDUAL NOISE SHAPING**



(57) **Abrégé/Abstract:**

A system and method for dynamic residual noise shaping configured to reduce hiss noise in an audio signal. The system and method may detect an amount and type of hiss noise. The system and method may limit calculated noise suppression gains responsive to the detected amount and type of hiss noise. The limited noise suppression gains may be applied to the audio signal and may reduce the hiss noise.

ABSTRACT

A system and method for dynamic residual noise shaping configured to reduce hiss noise in an audio signal. The system and method may detect an amount and type of hiss noise. The system and method may limit calculated noise suppression gains responsive to the detected amount and type of hiss noise. The limited noise suppression gains may be applied to the audio signal and may reduce the hiss noise.

SYSTEM AND METHOD FOR DYNAMIC RESIDUAL NOISE SHAPING

BACKGROUND OF THE INVENTION

1. Technical Field

[0001] The present disclosure relates to the field of signal processing. In particular, to a system and method for dynamic residual noise shaping.

2. Related Art

[0002] A high frequency hissing sound is often heard in wideband microphone recordings. While the high frequency hissing sound, or hiss noise, may not be audible when the environment is loud, it becomes noticeable and even annoying when in a quiet environment, or when the recording is amplified. The hiss noise can be caused by a variety of sources, from poor electronic recording devices to background noise in the recording environment from air conditioning, computer fan, or even the lighting in the recording environment.

BRIEF DESCRIPTION OF DRAWINGS

[0003] The system may be better understood with reference to the following drawings and description. The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. Moreover, in the figures, like referenced numerals designate corresponding parts throughout the different views.

[0004] Fig. 1 is a representation of spectrograms of background noise of an audio signal of a raw recording and a conventional noise reduced audio signal.

[0005] Fig. 2 is a schematic representation of an exemplary dynamic residual noise shaping system.

[0006] Fig. 3 is a representation of several exemplary target noise shape functions.

[0007] Fig. 4A is a set of exemplary calculated noise suppression gains.

[0008] Fig. 4B is the set of exemplary limited noise suppression gains.

[0009] Fig. 4C is the set of exemplary hiss noise floored noise suppression gains responsive to the dynamic residual noise shaping process.

[0010] Fig. 5 is a representation of spectrograms of background noise of an audio signal in the same raw recording as represented in Figure 1 processed by a conventionally noise reduced audio signal and a noised reduced audio signal with dynamic residual noise shaping.

[0011] Fig. 6 is flow diagram representing steps in a method for dynamic residual noise shaping in an audio signal.

[0012] Fig. 7 depicts a system for dynamic residual noise shaping in an audio signal.

DETAILED DESCRIPTION

[0013] Disclosed herein are a system and method for dynamic residual noise shaping. Dynamic shaping of residual noise may include, for example, the reduction of hiss noise.

[0014] U.S. Patent Application Serial No. 11/923,358 filed October 24, 2007 and having common inventorship, describes a system and method for dynamic noise reduction. This document discloses principles and techniques to automatically adjust the shape of high frequency residual noise. .

[0015] In a classical additive noise model, a noisy audio signal is given by

$$[0016] \quad y(t) = x(t) + n(t) \quad (1)$$

[0017] where $x(t)$ and $n(t)$ denote a clean audio signal, and a noise signal, respectively.

[0018] Let $|Y_{i,k}|$, $|X_{i,k}|$, and $|N_{i,k}|$ designate, respectively, the short-time spectral magnitudes of the noisy audio signal, the clean audio signal, and noise signal at the i^{th} frame and the k^{th} frequency bin. A noise reduction process involves the application of a suppression gain $G_{i,k}$ to each short-time spectrum value. For the purpose of noise reduction the clean audio signal and the noise signal are both estimates because their exact relationship is unknown. As such, the spectral magnitude of an estimated clean audio signal is given by:

$$[0019] \quad |\hat{X}_{i,k}| = G_{i,k} \cdot |Y_{i,k}| \quad (2)$$

[0020] Where $G_{i,k}$ are the noise suppression gains. Various methods are known in the literature to calculate these gains. One example further described below is a recursive Wiener filter.

[0021] A typical problem with noise reduction methods is that they create audible artifacts such as musical tones in the resulting signal, the estimated clean audio signal $|\hat{X}_{i,k}|$. These audible artifacts are due to errors in signal estimates that cause further errors in the noise suppression gains. For example the noise signal $|N_{i,k}|$ can only be estimated. To mitigate or mask the audible artifacts, the noise suppression gains may be floored (e.g. limited or constrained):

$$[0022] \quad \hat{G}_{i,k} = \max(\sigma, G_{i,k}) \quad (3)$$

[0023] The parameter σ in (3) is a constant noise floor, which defines a maximum amount of noise attenuation in each frequency bin. For example, when σ is set to 0.3, the system will attenuate the noise by a maximum of 10 dB at frequency bin k . The noise reduction process may produce limited noise suppression gains that will range from 0 dB to 10 dB at each frequency bin k .

[0024] The conventional noise reduction method based on the above noise suppression gain limiting applies the same maximum amount of noise attenuation to all frequencies. The constant noise floor in the noise suppression gain limiting may result in good performance for conventional noise reduction in narrowband communication. However,

it is not ideal for reducing hiss noise in high fidelity audio recordings or wideband communications. In order to remove the hiss noise, a lower constant noise floor in the suppression gain limiting may be required but this approach may also impair low frequency voice or music quality. Hiss noise may be caused by, for example, background noise or audio hardware and software limitations within one or more signal processing devices. Any of the noise sources may contribute to residual noise and/or hiss noise.

[0025] Figure 1 is a representation of spectrograms of background noise of an audio signal 102 of a raw recording and a conventional noise reduced audio signal 104. The audio signal 102 is an example raw recording of background noise and the conventional noise reduced audio signal 104 is the same audio signal 102 that has been processed with the noise reduction method where the noise suppression gains have been limited by a constant noise floor as described above. The audio signal 102 shows that a hiss noise 106 component of the background noise occurs mainly above 5 kHz in this example, and the hiss noise 106 in the conventional noise reduced audio signal 104 is a lower magnitude but still remains noticeable. The conventional noise reduction process illustrated in Figure 1 has reduced the level of the entire spectrum by substantially the same amount because the constant noise floor in the noise suppression gain limiting has prevented further attenuation.

[0026] Unlike conventional noise reduction methods that do not change the overall shape of background noise after processing, a dynamic residual noise shaping method may automatically detects hiss noise 106 and once hiss noise 106 is detected, may apply a dynamic attenuation floor to adjust the high frequency noise shape so that the residual noise may sound more natural after processing. For lower frequencies or when no hiss noise is detected in an input signal (e.g. a recording), the method may apply noise reduction similar to conventional noise reduction methods described above. Hiss noise as described herein comprises relatively higher frequency noise components of residual or background noise. Relatively higher frequency noise components may occur, for example, at frequencies above 500Hz in narrowband applications, above 3kHz in wideband applications, or above 5kHz in fullband applications.

[0027] Figure 2 is a schematic representation of an exemplary dynamic residual noise shaping system. The dynamic residual noise shaping system 200 may begin its signal processing in Figure 2 with subband analysis 202. The system 200 may receive an audio signal 102 that includes speech content, audio content, noise content, or any combination thereof. The subband analysis 202 performs a frequency transformation of the audio signal 102 that can be generated by different methods including a Fast Fourier Transform (FFT), wavelets, time-based filtering, and other known transformation methods. The frequency based transform may also use a windowed add/overlap analysis. The audio signal 102, or audio input signal, after the frequency transformation may be represented by $Y_{i,k}$ at the i^{th} frame and the k^{th} frequency bin or each k^{th} frequency band where a band contains one or more frequency bins. The frequency bands may group frequency bins in different ways including critical bands, bark bands, mel bands, or other similar banding techniques. A signal resynthesis 216 performs an inverse frequency transformation of the frequency transformation performed by the subband analysis 202.

[0028] The frequency transformation of the audio signal 102 may be processed by a subband signal power module 204 to produce the spectral magnitude of the audio signal $|Y_{i,k}|$. The subband signal power module 204 may also perform averaging of frequency bins over time and frequency. The averaging calculation may include simple averages, weighted averages or recursive filtering.

[0029] A subband background noise power module 206 may calculate the spectral magnitude of the estimated background noise $|\tilde{N}_{i,k}|$ in the audio signal 102. The background noise estimate may include signal information from previously processed frames. In one implementation, the spectral magnitude of the background noise is calculated using the background noise estimation techniques disclosed in U.S. Patent No. 7,844,453, except that in the event of any inconsistent disclosure or definition from the present specification, the disclosure or definition herein shall be deemed to prevail. In other implementations, alternative background noise estimation techniques may be used, such as a noise power estimation technique based on minimum statistics.

[0030] A noise reduction module 208 calculates suppression gains $G_{i,k}$ using various methods that are known in the literature to calculate suppression gains. An exemplary noise reduction method is a recursive Wiener filter. The Wiener suppression gain, or noise suppression gains, is defined as:

$$[0031] \quad G_{i,k} = \frac{S\hat{N}R_{\text{priori}_{i,k}}}{S\hat{N}R_{\text{priori}_{i,k}} + 1}. \quad (4)$$

[0032] Where $S\hat{N}R_{\text{priori}_{i,k}}$ is the a priori SNR estimate and is calculated recursively by:

$$[0033] \quad S\hat{N}R_{\text{priori}_{i,k}} = G_{i-1,k} S\hat{N}R_{\text{post}_{i,k}} - 1. \quad (5)$$

[0034] $S\hat{N}R_{\text{post}_{i,k}}$ is the a posteriori SNR estimate given by:

$$[0035] \quad S\hat{N}R_{\text{post}_{i,k}} = \frac{|Y_{i,k}|^2}{|\hat{N}_{i,k}|^2}. \quad (6)$$

[0036] Where $|\hat{N}_{i,k}|$ is the background noise estimate.

[0037] A hiss detector module 210 estimates the amount of hiss noise in the audio signal. The hiss detector module 210 may indicate the presence of hiss noise 106 by analyzing any combination of the audio signal, the spectral magnitude of the audio signal $|Y_{i,k}|$, and the background noise estimate $|\hat{N}_{i,k}|$. An exemplary hiss detector method utilized by the hiss detector module 210 first may convert the short-time power spectrum of a background noise estimation, or background noise level, into the dB domain by:

$$[0038] \quad B(f) = 20 \log_{10} |N(f)|. \quad (7)$$

[0039] The background noise level may be estimated using a background noise level estimator. The dB power spectrum $B(f)$ may be further smoothed in frequency to remove small dips or peaks in the spectrum. A pre-defined hiss cutoff frequency f_0 may be chosen to divide the whole spectrum into a low frequency portion and a high frequency portion. The dynamic hiss noise reduction may be applied to the high frequency portion of the spectrum.

[0040] Hiss noise 106 is usually audible in high frequencies. In order to eliminate or mitigate hiss noise after noise reduction, the residual noise may be constrained to have a target noise shape, or have certain colors. Constraining the residual noise to have certain colors may be achieved by making the residual noise power density to be proportional to $1/f^\beta$. For instance, white noise has a flat spectral density, so $\beta = 0$, while pink noise has $\beta = 1$, and brown noise has $\beta = 2$. The greater the β value, the quieter the noise in high frequencies. In an alternative embodiment, the residual noise power density may be a function that has flatter spectral density at lower frequencies and a more sloped spectral density at higher frequencies.

[0041] The target residual noise dB power spectrum is defined by:

$$[0042] \quad T(f) = B(f_0) - 10\beta \log_{10}(f/f_0). \quad (8)$$

[0043] The difference between the background noise level and the target noise level at a frequency may be calculated with a difference calculator. Whenever the difference between the noise estimation and the target noise defined by:

$$[0044] \quad D(f) = B(f) - T(f) \quad (9)$$

[0045] is greater than a hiss threshold δ , hiss noise is detected and a dynamic floor may be used to do substantial noise suppression to eliminate hiss. A detector may detect when the residual background noise level exceeds the hiss threshold. The dynamic suppression factor for a given frequency above the hiss cutoff frequency f_0 may be given by:

$$[0046] \quad \lambda(f) = \begin{cases} 10^{0.05D(f)}, & \text{if } D(f) > \delta \\ 1, & \text{otherwise} \end{cases}. \quad (10)$$

[0047] Alternatively, for each bin above the hiss cutoff frequency bin k_0 the dynamic suppression factor may be given by:

$$[0048] \quad \lambda(k) = \begin{cases} 10^{0.05D(k_0)}, & \text{if } D(k_0) > \delta \\ 1, & \text{otherwise} \end{cases}. \quad (11)$$

[0049] The dynamic noise floor may be defined as:

$$[0050] \quad \eta(k) = \begin{cases} \sigma * \lambda(k), & \text{when } k \geq k_0 \\ \sigma, & \text{when } k < k_0 \end{cases} \quad (12)$$

[0051] By combining the dynamic floor described above with the conventional noise reduction method, the color of residual noise may be constrained by a pre-defined target noise shape, and the quality of the noise-reduced speech signal may be significantly improved. Below the hiss cutoff frequency f_0 , a constant noise floor may be applied. The hiss cutoff frequency f_0 may be a fixed frequency, or may be adaptive depending on the noise spectral shape.

[0052] A suppression gain limiting module 212 may limit the noise suppression gains according to the result of the hiss detector module 210. In an alternative to flooring the noise suppression gains by a constant floor as in equation (3), the dynamic hiss noise reduction approach may use the dynamic noise floor defined in equation (9) to estimate the noise suppression gains:

$$[0053] \quad \hat{G}_{i,k} = \max(\eta(k), G_{i,k}). \quad (13)$$

[0054] A noise suppression gain applier 214 applies the noise suppression gains to the frequency transformation of the audio signal 102.

[0055] Figure 3 is a representation of several exemplary target noise shape 308 functions. Frequencies above the hiss cutoff frequency 306 may be constrained by the target noise shape 308. The target noise shape 308 may be constrained to have certain colors of residual noise including white, pink and brown. The target noise shape 308 may be adjusted by offsetting the target noise shape 308 by the hiss noise floor 304. Frequencies below the hiss cutoff frequency 306, or conventional noise reduced frequencies 302, may be constrained by the hiss noise floor 304. Values shown in Figure 3 are illustrative in nature and are not intended to be limiting in any way.

[0056] Figure 4A is a set of exemplary calculated noise suppression gains 402. The exemplary calculated noise suppression gains 402 may be the output of the recursive Wiener filter described in equation 4. Figure 4B is a set of limited noise suppression gains 404. The limited noise suppression gains 404 are the calculated noise suppression

gains 402 that have been floored as described in equation 3. Limiting the calculated noise suppression gains 402 may mitigate audible artifacts caused by the noise reduction process. Figure 4C is a set of exemplary modified noise suppression gains 406 responsive to the dynamic residual noise shaping process. The modified noise suppression gains 406 are the calculated noise suppression gains 402 that have been floored as described in equation 12.

[0057] Figure 5 is a representation of spectrograms of background noise of an audio signal 102 in the same raw recording as represented in Figure 1 processed by a conventionally noise reduced audio signal 104 and a noise reduced audio signal processed by dynamic residual noise shaping 502. The example hiss cutoff frequency 306 is set to approximately 5 kHz. It can be observed that at frequencies above the hiss cutoff frequency 306 that the noise reduced audio signal with dynamic residual noise shaping 502 may produce a lower noise floor than the noise floor produced by the conventionally noise reduced audio signal 104.

[0058] Figure 6 is flow diagram representing steps in a method for dynamic residual noise shaping in an audio signal 102. In step 602, the amount and type of hiss noise is detected in the audio signal 102. In step 604, a noise reduction process is used to calculate noise suppression gains 402. In step 606, the noise suppression gains 402 are modified responsive to the detected amount and type of hiss noise 106. Different modifications may be applied to noise suppression gains 402 associated with frequencies below and above a hiss cutoff frequency 306. In step 608, the modified noise suppression gains 406 are applied to the audio signal 102.

[0059] The method according to the present description may be implemented by computer executable program instructions stored on a computer-readable storage medium. A system for dynamic hiss reduction may comprise electronic components, analog and/or digital, for implementing the processes described above. In some embodiments the system may comprise a processor and memory for storing instructions that, when executed by the processor, enact the processes described above.

[0060] Figure 7 depicts a system for dynamic residual noise shaping in an audio signal 102. The system 702 comprises a processor 704 (aka CPU), input and output interfaces 706 (aka I/O) and memory 708. The processor 704 may comprise a single processor or multiple processors that may be disposed on a single chip, on multiple devices or distribute over more than one system. The processor 704 may be hardware that executes computer executable instructions or computer code embodied in the memory 708 or in other memory to perform one or more features of the system. The processor 704 may include a general processor, a central processing unit, a graphics processing unit, an application specific integrated circuit (ASIC), a digital signal processor, a field programmable gate array (FPGA), a digital circuit, an analog circuit, a microcontroller, any other type of processor, or any combination thereof.

[0061] The memory 708 may comprise a device for storing and retrieving data or any combination thereof. The memory 708 may include non-volatile and/or volatile memory, such as a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM), or a flash memory. The memory 708 may comprise a single device or multiple devices that may be disposed on one or more dedicated memory devices or on a processor or other similar device. Alternatively or in addition, the memory 708 may include an optical, magnetic (hard-drive) or any other form of data storage device.

[0062] The memory 708 may store computer code, such as the hiss detector 210, the noise reduction filter 208 and/or any component. The computer code may include instructions executable with the processor 704. The computer code may be written in any computer language, such as C, C++, assembly language, channel program code, and/or any combination of computer languages. The memory 708 may store information in data structures such as the calculated noise suppression gains 402 and the modified noise suppression gains 406.

[0063] The memory 708 may store instructions 710 that when executed by the processor, configure the system to enact the system and method for reducing hiss noise described herein with reference to any of the preceding Figures 1-6. The instructions 710

may include the following. Detecting an amount and type of hiss noise 106 in an audio signal of step 602. Calculating noise suppression gains 402 by applying a noise reduction process to the audio signal 102 of step 604. Modifying the noise suppression gains 402 responsive to the detected amount and type of hiss noise 102 of step 606. Applying the modified noise suppression gains 406 to the audio signal 102 of step 608.

[0064] All of the disclosure, regardless of the particular implementation described, is exemplary in nature, rather than limiting. The system 200 may include more, fewer, or different components than illustrated in Figure 2. Furthermore, each one of the components of system 200 may include more, fewer, or different elements than is illustrated in Figure 2. Flags, data, databases, tables, entities, and other data structures may be separately stored and managed, may be incorporated into a single memory or database, may be distributed, or may be logically and physically organized in many different ways. The components may operate independently or be part of a same program or hardware. The components may be resident on separate hardware, such as separate removable circuit boards, or share common hardware, such as a same memory and processor for implementing instructions from the memory. Programs may be parts of a single program, separate programs, or distributed across several memories and processors.

[0065] The functions, acts or tasks illustrated in the figures or described may be executed in response to one or more sets of logic or instructions stored in or on computer readable media. The functions, acts or tasks are independent of the particular type of instructions set, storage media, processor or processing strategy and may be performed by software, hardware, integrated circuits, firmware, micro code and the like, operating alone or in combination. Likewise, processing strategies may include multiprocessing, multitasking, parallel processing, distributed processing, and/or any other type of processing. In one embodiment, the instructions are stored on a removable media device for reading by local or remote systems. In other embodiments, the logic or instructions are stored in a remote location for transfer through a computer network or over telephone lines. In yet other embodiments, the logic or instructions may be stored within a given computer such as, for example, a central processing unit ("CPU").

[0066] While various embodiments of the invention have been described, it will be apparent to those of ordinary skill in the art that many more embodiments and implementations are possible within the scope of the present invention. Accordingly, the invention is not to be restricted except in light of the attached claims and their equivalents.

CLAIMS

What is claimed is:

1. A dynamic residual noise shaping method, comprising:
 - detecting an amount of high frequency hiss noise in an audio signal by a computer processor;
 - calculating noise suppression gains by the computer processor by applying a noise reduction filter to the audio signal;
 - modifying the calculated noise suppression gains by the computer processor responsive to the detected amount of high frequency hiss noise; and
 - applying the modified noise suppression gains by the computer processor to the audio signal.
2. The method of claim 1, where modifying the calculated noise suppression gains responsive to the detected amount of high frequency hiss noise comprises modifying the calculated noise suppression gains above a hiss cutoff frequency.
3. The method of claim 1, where detecting the amount and type of high frequency hiss noise in an audio signal comprises:
 - estimating a background noise level for each of a plurality of frequency bins of the audio signal;
 - calculating a difference between the background noise level and a target noise shape for each of the plurality of frequency bins of the audio signal; and
 - detecting when the difference exceeds a hiss threshold for each of the plurality of frequency bins of the audio signal.
4. The method of claim 3, where the target noise shape is adjusted by a hiss noise floor offset.
5. The method of claim 3, where detecting when the difference exceeds the hiss threshold for each of the plurality of frequency bins further comprises calculating the

hiss threshold responsive to any one or more of an audio signal level, the background noise level and an associated frequency bin.

6. The method of claim 1, where modifying the noise suppression gains responsive to the detected amount of high frequency hiss noise comprises modifying the noise suppression gains to substantially correlate to a target noise shape for each of a plurality of frequency bins of the audio signal.
7. The method of claim 6, where the target noise shape comprises one of a white, a pink or a brown noise.
8. The method of claim 6, where the target noise shape comprises an increasing gain with an increasing frequency.
9. The method of claim 1, where calculating noise suppression gains by applying the noise reduction filter to the audio signal comprises averaging the audio signal in time and frequency.
10. The method of claim 1, further comprising generating a set of subbands of the audio signal through a subband filter or a Fast Fourier Transform.
11. The method of claim 10, further comprising generating the set of subbands of the audio signal according to a critical, an octave, a mel, or a bark band spacing technique.
12. A system for dynamic residual noise shaping, comprising:
 - a hiss noise detector to detect an amount of high frequency hiss noise in an audio signal;
 - a noise reduction filter applied to the audio signal to calculate noise suppression gains;
 - a noise suppression gain modifier to modify the noise suppression gains responsive to the amount of high frequency hiss noise detected; and
 - a suppression gain applier to apply the modified noise suppression gains to the audio signal.

13. The system of claim 12, where the noise suppression gain modifier, responsive to the amount of high frequency hiss noise, modifies the calculated noise suppression gains above a hiss cutoff frequency.
14. The system of claim 12, where the hiss noise detector to detect an amount of high frequency hiss noise in an audio signal comprises:
 - a background noise level estimator to estimate a background noise level for each of a plurality of frequency bins of the audio signal;
 - a difference calculator to calculate a difference between the background noise level and a target noise shape for each of the plurality of frequency bins of the audio signal; and
 - a detector to detect when the difference exceeds a hiss threshold for each of the plurality of frequency bins of the audio signal.
15. The system of claim 14, where the target noise shape is adjusted by a hiss noise floor offset.
16. The system of claim 14, where the detector to detect when the difference exceeds the hiss threshold for each of the plurality of frequency bins further comprises a calculator to calculate the hiss threshold responsive to any one or more of an audio signal level, the background noise level and an associated frequency bin.
17. The system of claim 12, where the noise suppression gain modifier to modify the noise suppression gains responsive to the amount of high frequency hiss noise comprises a modifier to modify the noise suppression gains to substantially correlate to a target noise shape for each of a plurality of frequency bins of the audio signal.
18. The system of claim 17, where the target noise shape comprises one of a white, a pink or a brown noise.
19. The system of claim 17, where the target noise shape comprises increasing gain with increasing frequency.

20. The system claim 12, where the noise reduction filter applied to the audio signal to calculate noise suppression gains comprises averaging the audio signal in time and frequency.
21. The system of claim 12, further comprising a subband analyzer to generate a set of subbands of the audio signal through a subband filter or a Fast Fourier Transform.
22. The system of claim 21, further comprising a subband analyzer to generate the set of subbands of the audio signal according to a critical, octave, mel, or bark band spacing technique.

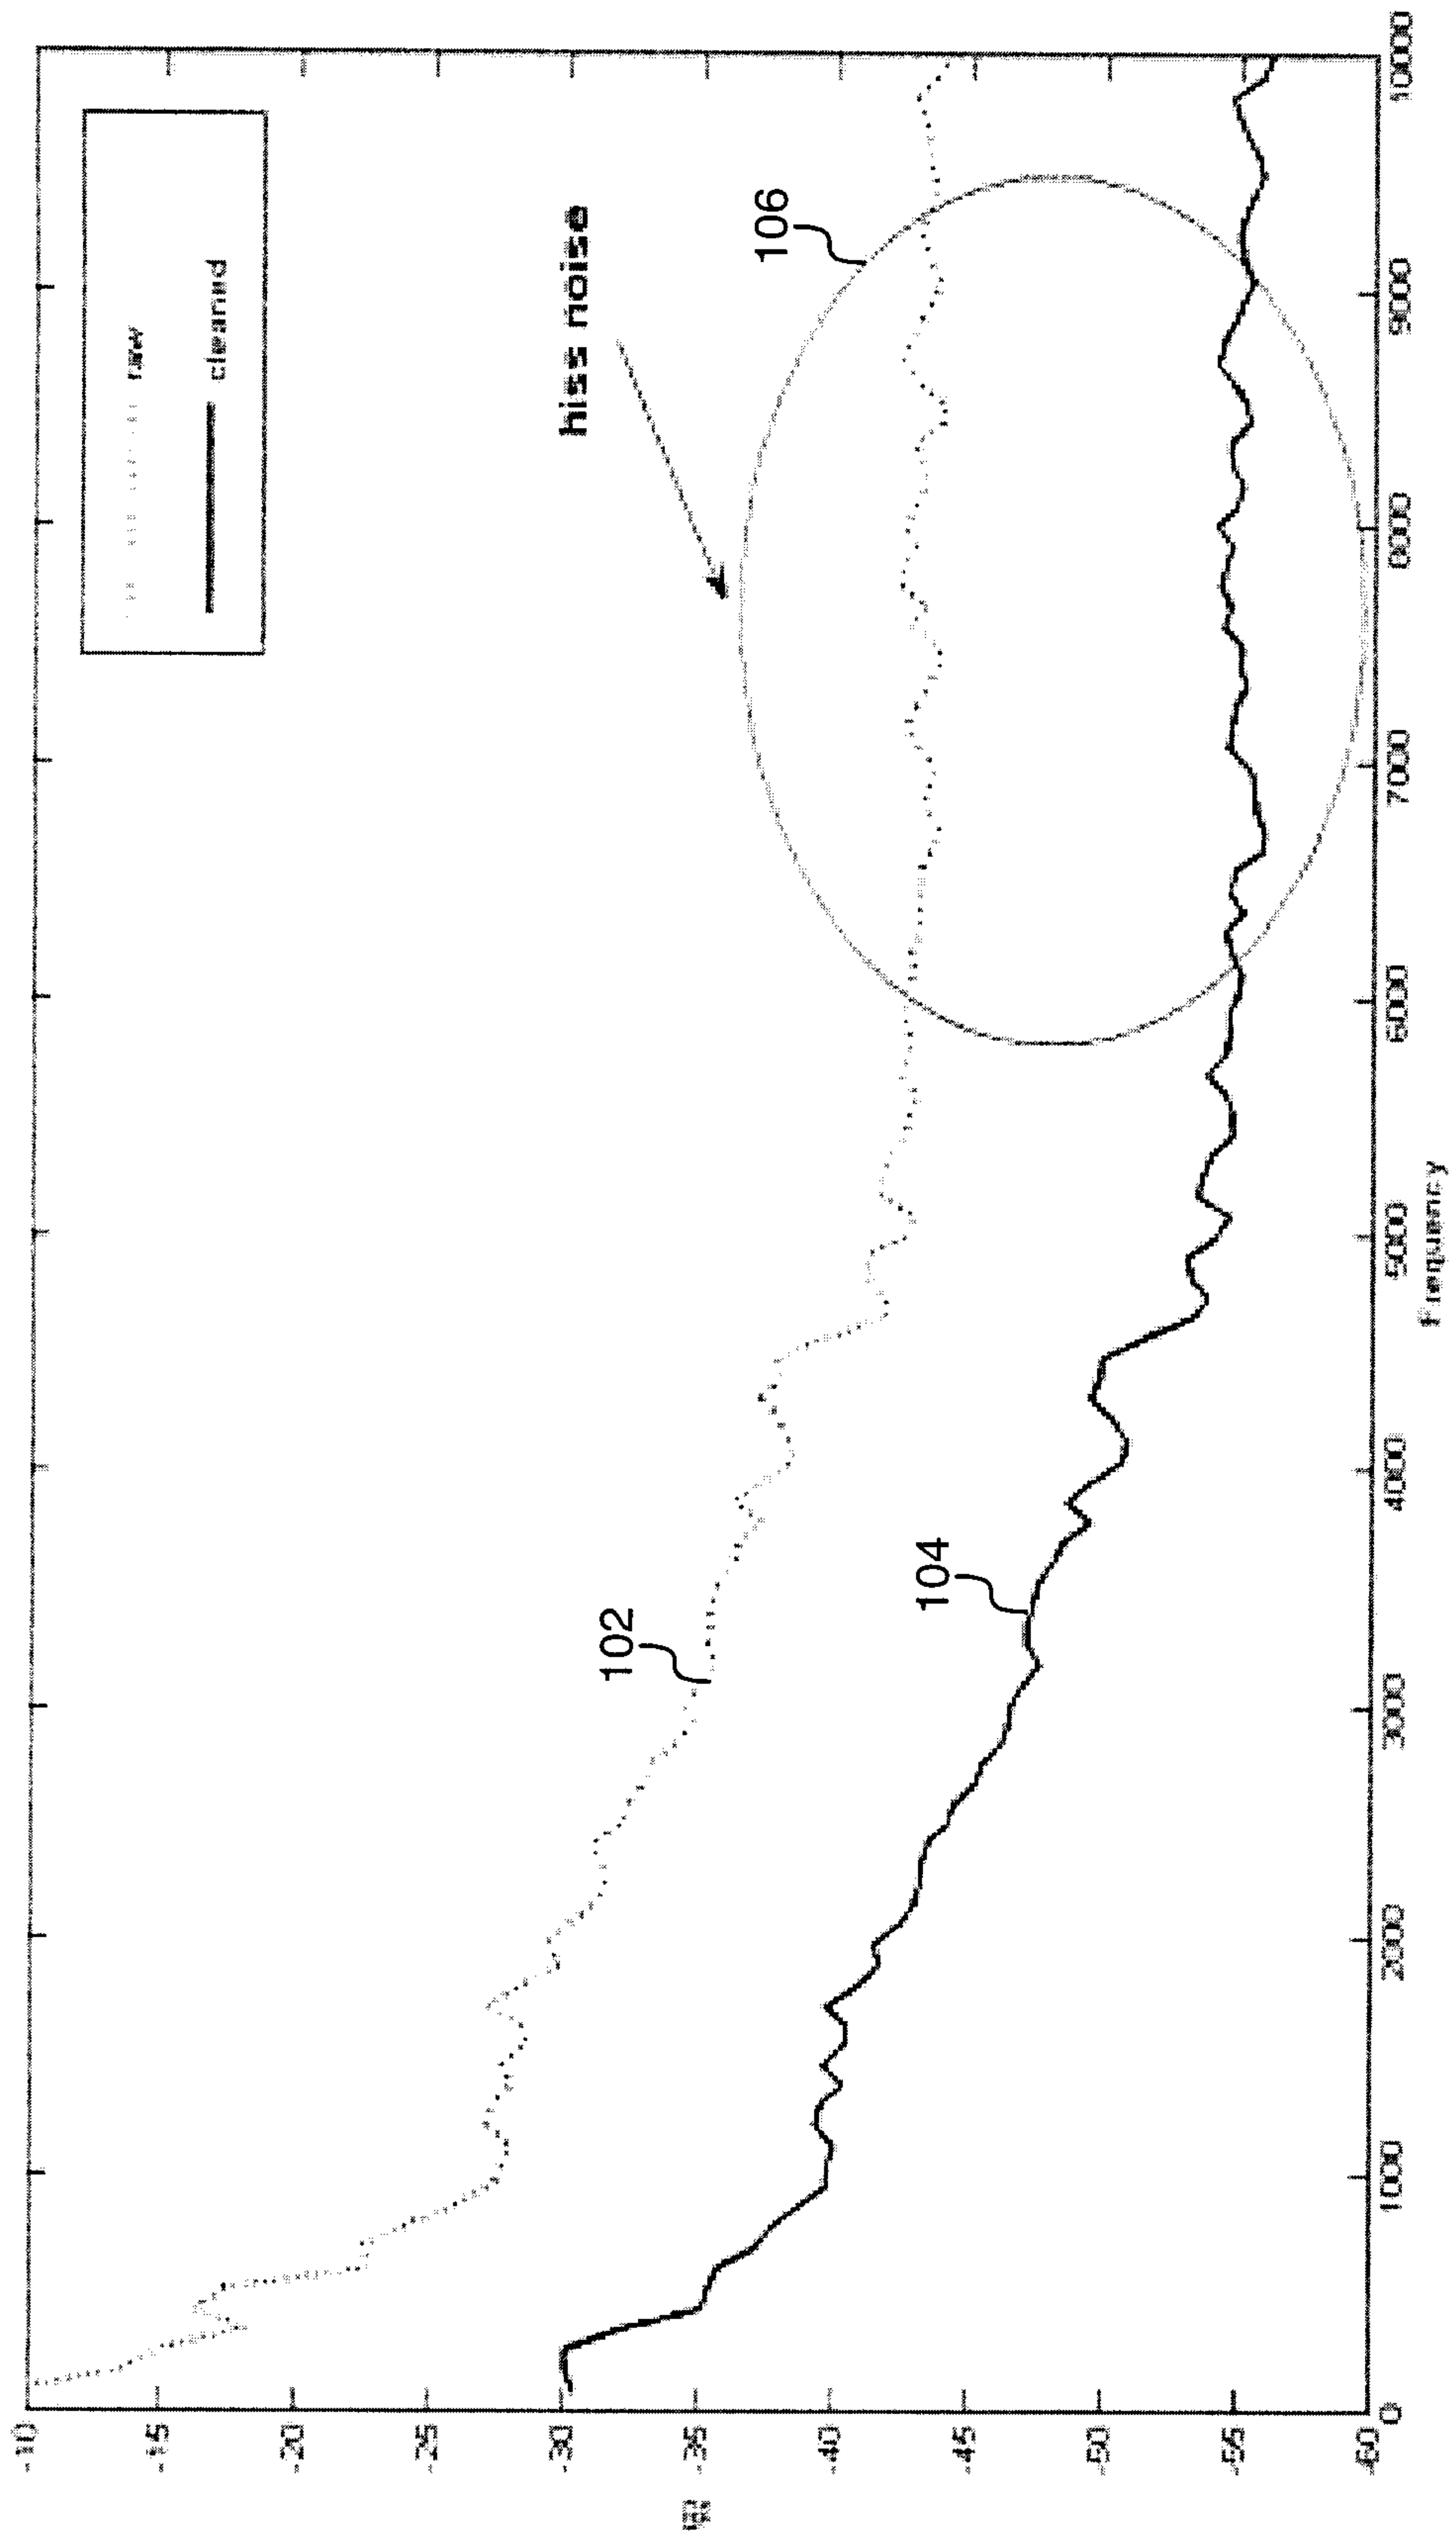



Figure 1



200

Figure 2




Figure 3

Figure 4A




Figure 4B





Figure 4C






Figure 5

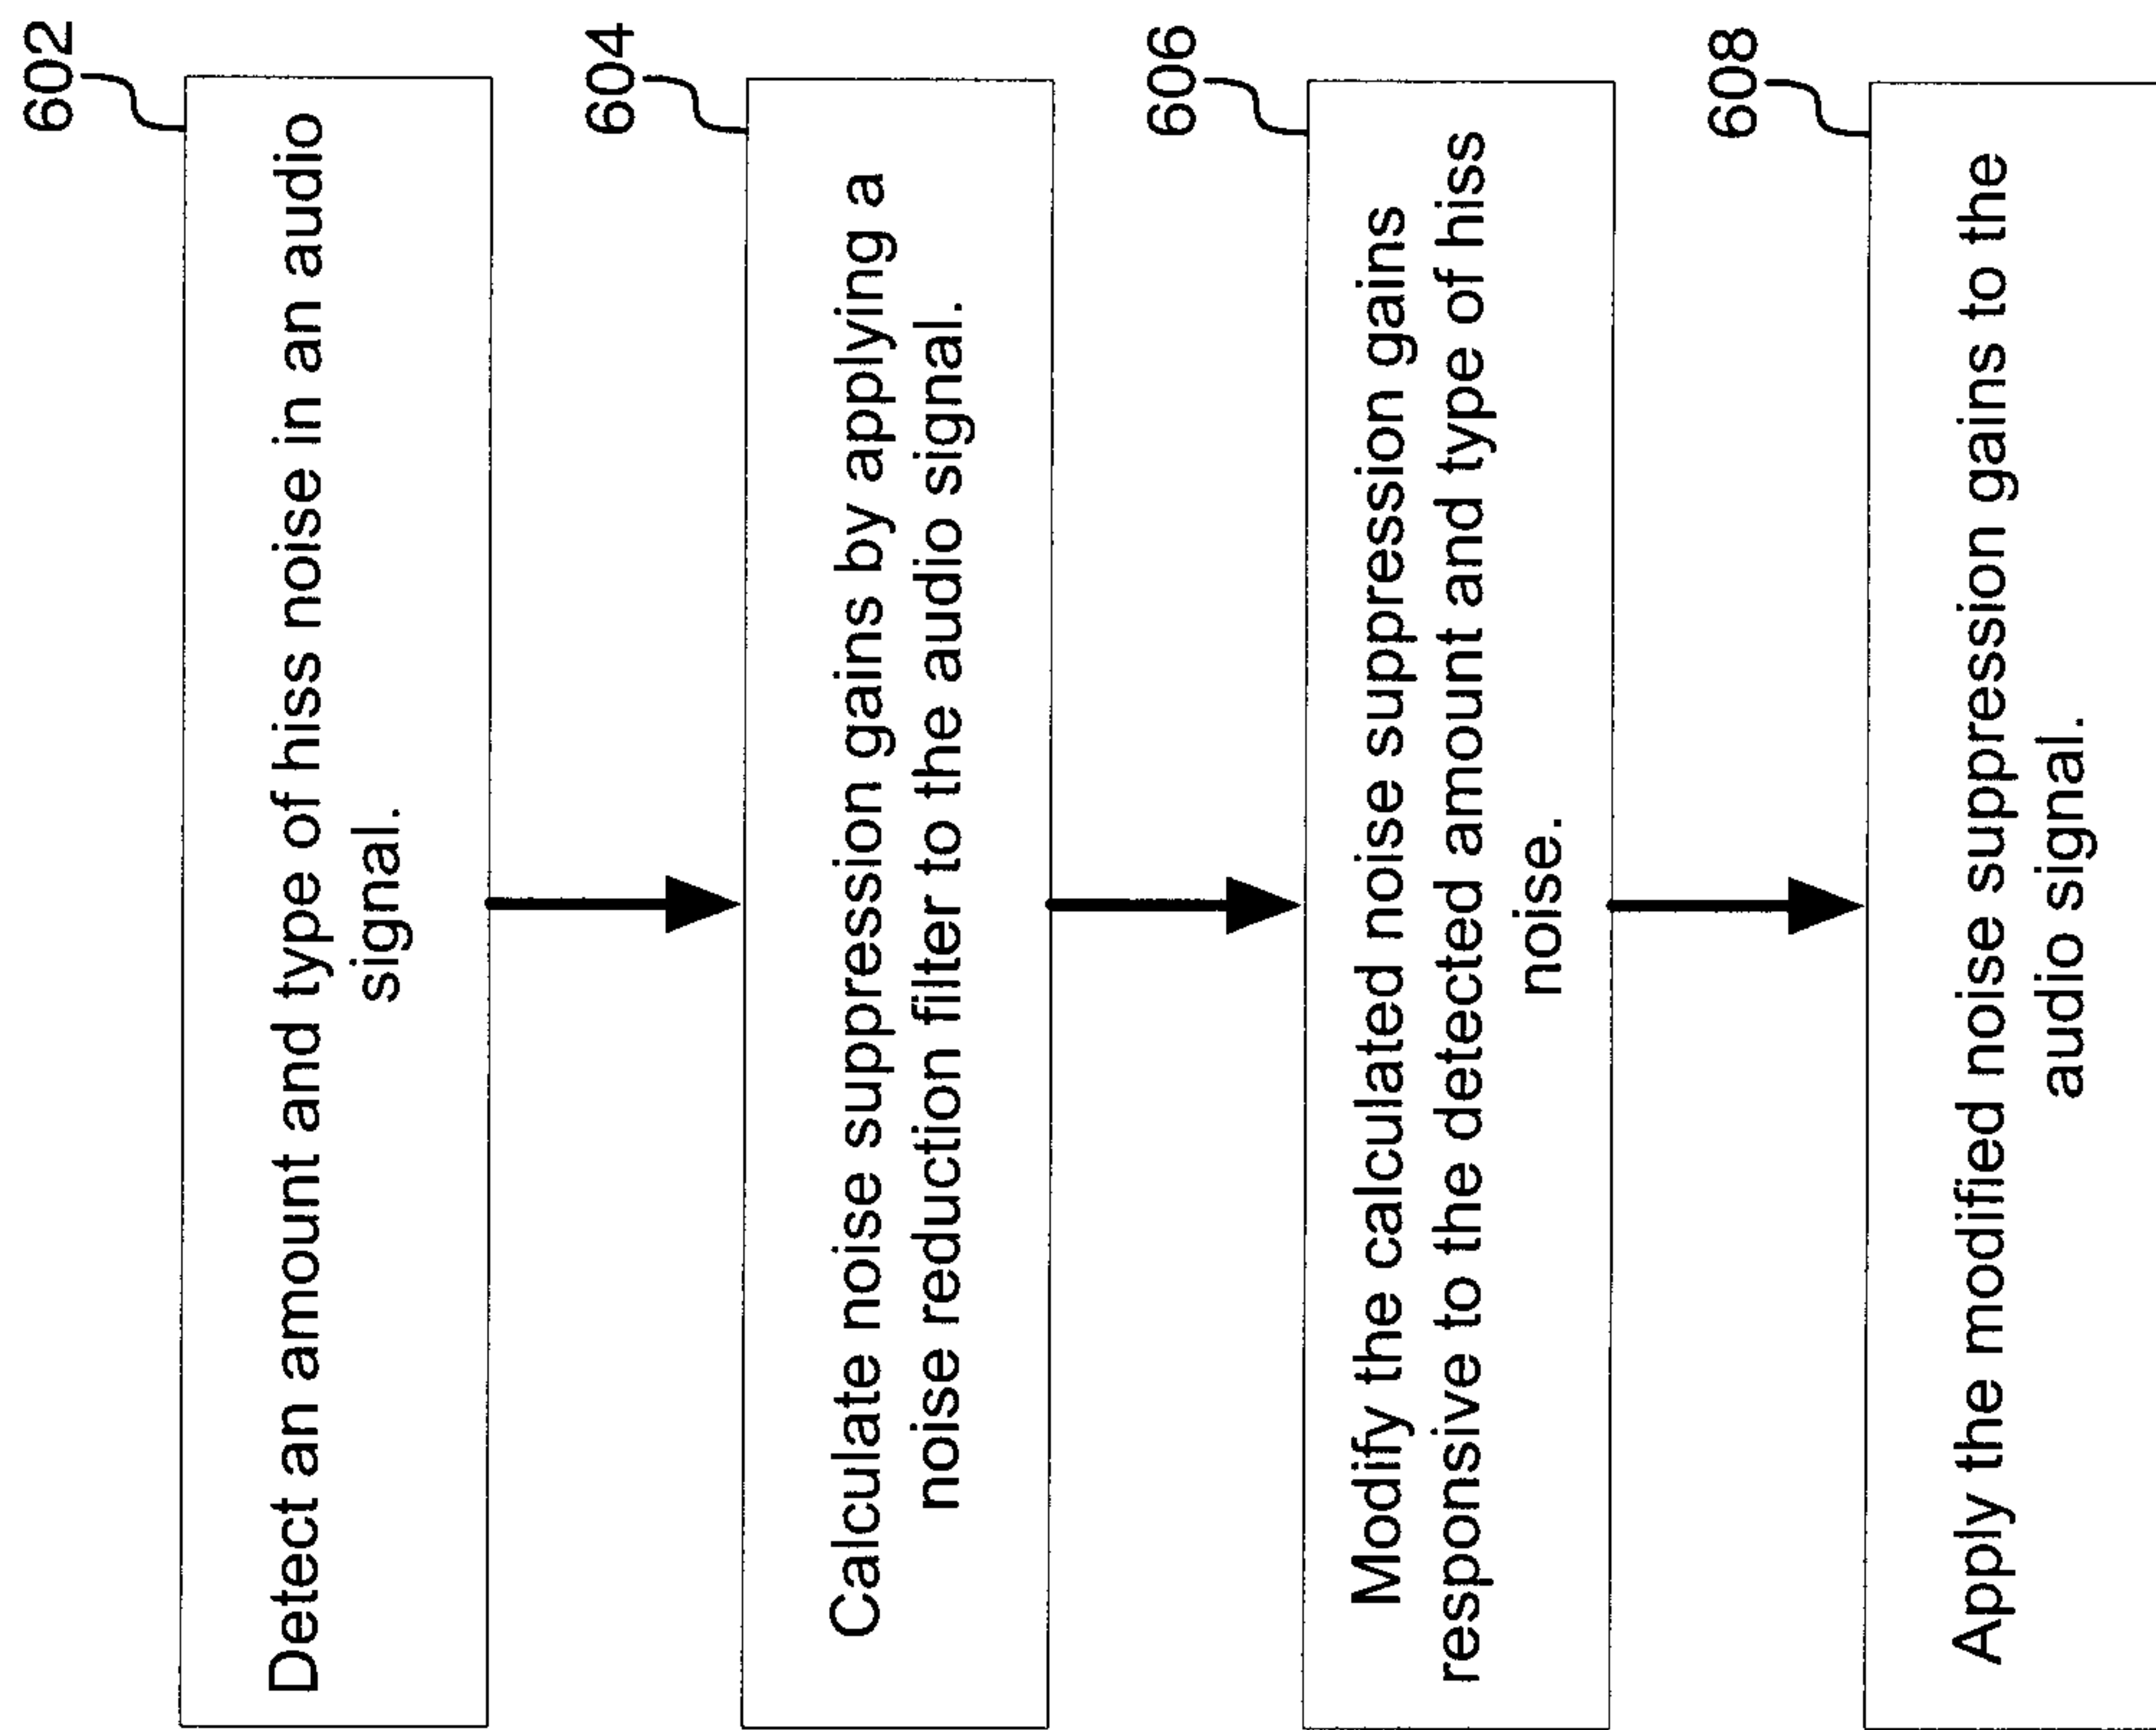


Figure 6

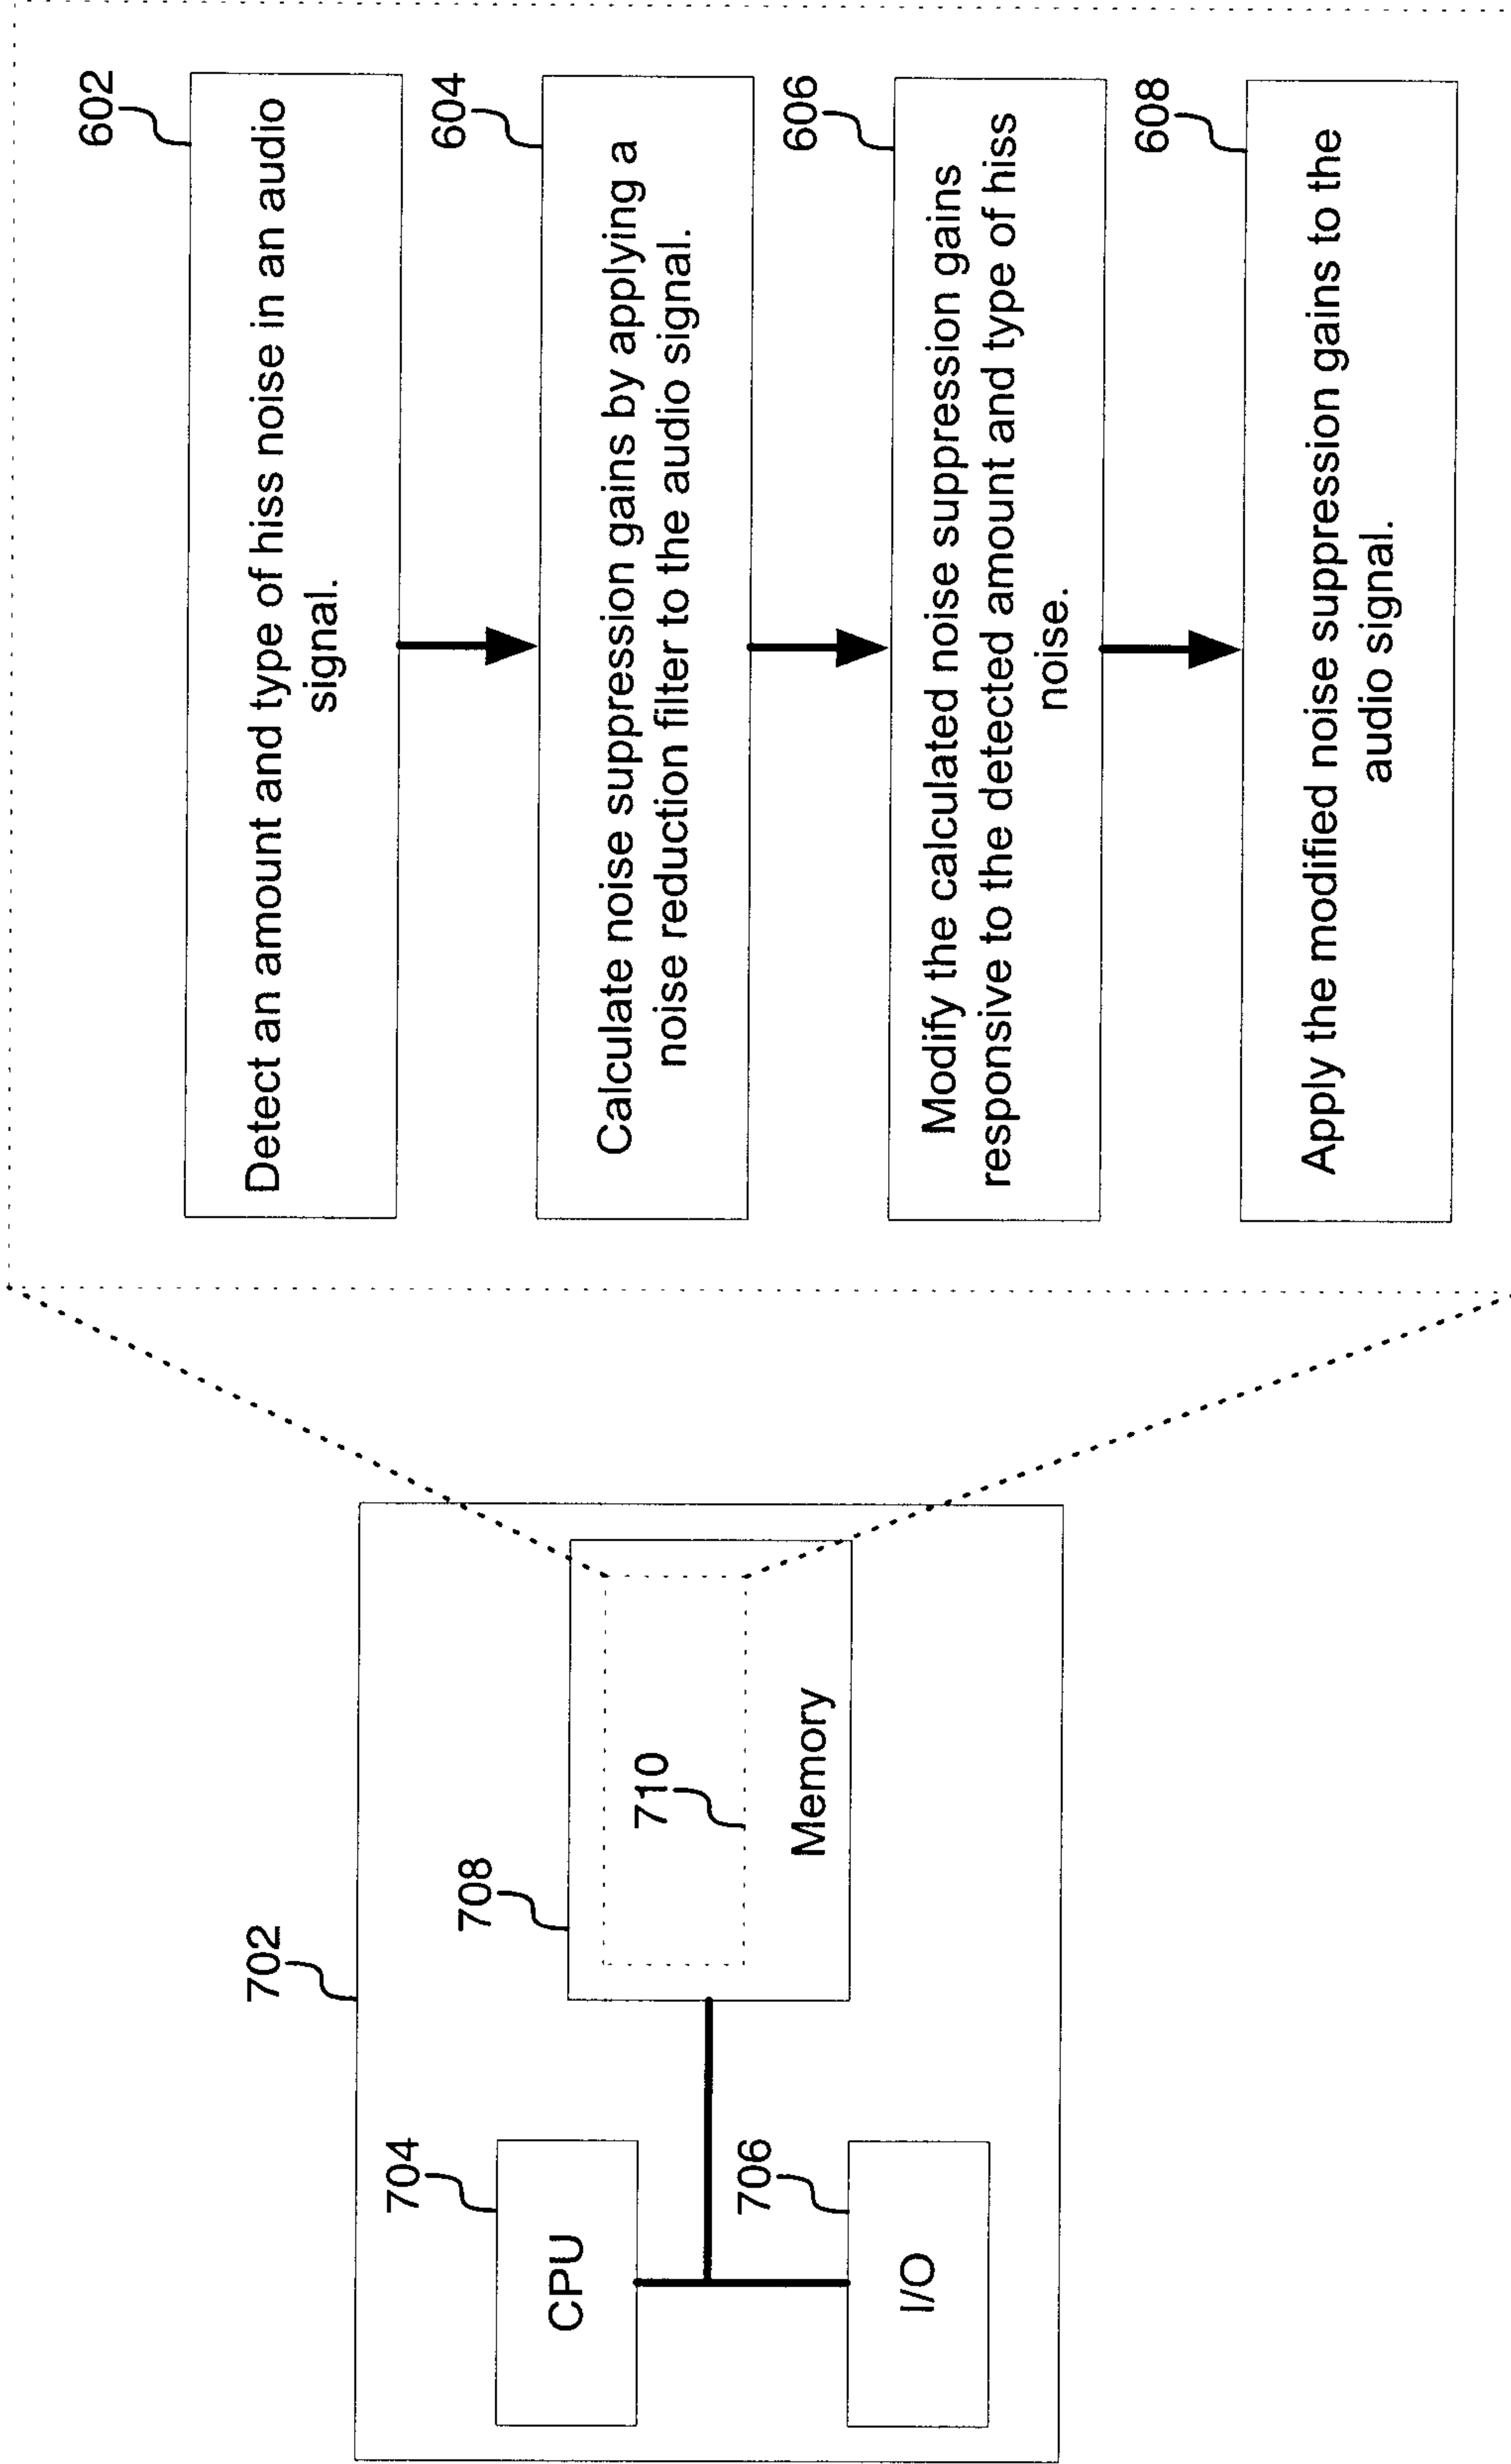


Figure 7

