



- (51) **International Patent Classification:**  
*H04B 15/00* (2006.01)    *H03G 5/00* (2006.01)
- (21) **International Application Number:**  
PCT/US2012/030274
- (22) **International Filing Date:**  
23 March 2012 (23.03.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
61/467,493    25 March 2011 (25.03.2011)    US  
13/205,492    8 August 2011 (08.08.2011)    US
- (71) **Applicant (for all designated States except US):** **THE IN-TELLISIS CORPORATION** [US/US]; 10350 Science Center Drive, Suite 140, San Diego, CA 92121 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **BRADLEY, David, C.** [US/US]; 7775 Exchange Place, La Jolla, CA 92037 (US). **GOLDIN, Daniel, S.** [US/US]; P. O. Box 2765, Malibu, CA 90265 (US). **HILTON, Robert, N.** [US/US]; 8594 Capricorn Way, San Diego, CA 92126 (US). **FISHER, Nicholas, K.** [US/US]; 6230 Caminito Araya, San

Diego, CA 92122 (US). **GATEAU, Rodney** [US/US]; 600 Front Street, Apt. 119, San Diego, CA 92101 (US).

(74) **Agents:** **ESPLIN, D., Benjamin** et al.; Pillsbury Winthrop Shaw Pittman LLP, P. O. Box 10500, McLean, VA 22102 (US).

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

(84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) **Title:** SYSTEMS AND METHODS FOR RECONSTRUCTING AN AUDIO SIGNAL FROM TRANSFORMED AUDIO INFORMATION

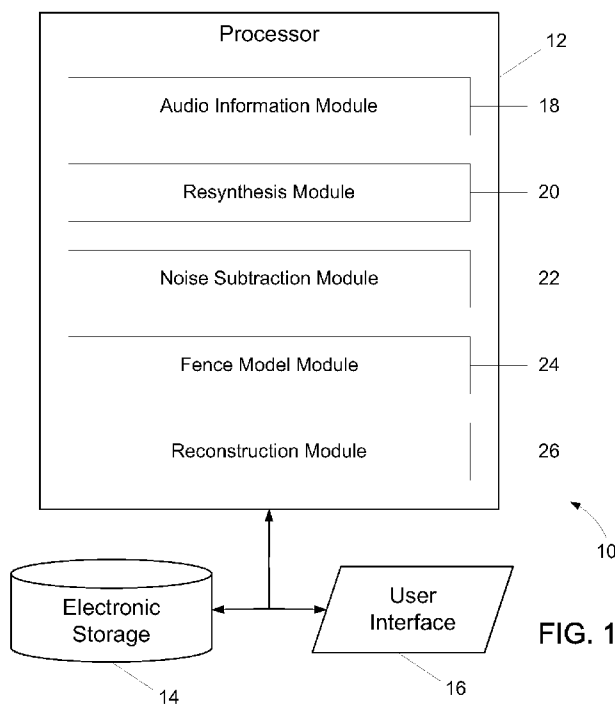


FIG. 1

(57) **Abstract:** A system and method may be configured to reconstruct an audio signal from transformed audio information. The audio signal may be resynthesized based on individual harmonics and corresponding pitches determined from the transformed audio information. Noise may be subtracted from the transformed audio information by interpolating across peak points and across trough points of harmonic pitch paths through the transformed audio information, and subtracting values associated with the trough point interpolations from values associated with the peak point interpolations. Noise between harmonics of the sound may be suppressed in the transformed audio information by centering functions at individual harmonics in the transformed audio information, the functions serving to suppress noise between the harmonics.



SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG). **Published:**

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

**Declarations under Rule 4.17:**

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

## **SYSTEMS AND METHODS FOR RECONSTRUCTING AN AUDIO SIGNAL FROM TRANSFORMED AUDIO INFORMATION**

### ***RELATED APPLICATIONS***

(01) This application claims the benefit of U.S. patent application No. 13/205,492, filed on August 8, 2011, and entitled "Systems and Methods for Reconstructing an Audio Signal from Transformed Audio Information," which claims the priority benefit of U.S. provisional patent application No. 61/467,493, filed on March 25, 2011, and entitled "Spectral Motion Transformation," both of which are hereby incorporated into this disclosure by reference in their entirety.

### ***FIELD***

(02) The disclosure relates to reconstructing an audio signal obtained from a noisy environment based on pitch and harmonic amplitudes determined from transformed audio information and/or suppressing noise between harmonics of the sound in the transformed audio information by centering functions at individual harmonics in the transformed audio information.

### ***BACKGROUND***

(03) Systems and method for reconstructing an audio signal are known. Existing techniques operate with relative accuracy and precision in the best of conditions. However, in "noisy" conditions (e.g., either sound noise or processing noise) the accuracy, precision, and/or speed of conventional techniques may drop off significantly. Since many of the settings and/or audio signals in and on which these techniques are applied may be considered noisy, conventional processing to reconstruct an audio signal may be only marginally useful.

### ***SUMMARY***

(04) One aspect of the disclosure relates to a system and method of reconstructing an audio signal from transformed audio information. The audio signal may be resynthesized based on individual harmonics and corresponding pitches determined from the transformed audio information. Noise may be subtracted from the transformed audio information by interpolating across peak points and across trough points of harmonic pitch paths through the transformed audio information, and subtracting the trough point interpolations from the peak point interpolations. Noise

between harmonics of the sound may be suppressed in the transformed audio information by centering functions at individual harmonics in the transformed audio information, the functions serving to suppress noise between the harmonics.

(05) In some implementations, a system may be configured for reconstructing an audio signal from transformed audio information. The system may comprise one or more processors configured to execute computer program modules. The computer program modules may comprise one or more of an audio information module, a resynthesis module, a noise subtraction module, a fence model module, a reconstruction module, and/or other modules.

(06) The audio information module may be configured to obtain transformed audio information representing one or more sounds. The audio signal may have a duration. That is, the audio signal may span a discrete period of time. The transformed audio information may have been transformed in discrete time sample windows over the audio signal. The time sample windows may be overlapping or non-overlapping in time. The transformed audio information may include pitch and/or pitch information associated with the audio signal. In some implementations, pitch and/or pitch information may be determined as described in one or both of U.S. Patent Application Serial No. 13/205,483, filed August 8, 2011, and entitled "System And Method For Tracking Sound Pitch Across An Audio Signal", and/or U.S. Patent Application Serial No. 13/205,521, filed August 8, 2011, and entitled "System And Method For Tracking Sound Pitch Across An Audio Signal Using Harmonic Envelope," which are hereby incorporated by reference into the present application in their entireties. The transformed audio information may specify magnitude of a coefficient related to signal intensity as a function of frequency for an audio signal and time. In some implementations, the transformed audio information for the time sample window may include a plurality of sets of transformed audio information. The individual sets of transformed audio information may correspond to different fractional chirp rates. Obtaining the transformed audio information may include transforming the audio signal, receiving the transformed audio information in a communications transmission, accessing stored transformed audio information, and/or other techniques for obtaining information.

(07) The resynthesis module may be configured to resynthesize the audio signal based on individual harmonics and corresponding pitches determined from the

transformed audio information. According to some implementations, resynthesizing the audio signal may include tracking one or more pitches of the sound to determine individual pitches and corresponding amplitudes as a function of time for individual harmonics of the sound. Individual harmonics may be synthesized using oscillators corresponding to individual harmonics. Synthesizing individual harmonics may include, for a given harmonic, integrating a corresponding pitch over time to determine a phase of the given harmonic. Individual ones of the oscillators may be based on a cosine function. The synthesized harmonics may be summed to obtain the resynthesized audio signal.

(08) In some implementations, resynthesis module may be configured to solve any phase problems because the audio signal may be built through integration, where phase is a consequence of the audio signal and not something that needs to be factored in. Also, the degree of compression of the resynthesized audio signal may go below a kB per second for voice.

(09) The resynthesized audio signal may be built from oscillators and parameters that specify pitch and harmonic amplitudes as a function of time. One or more of these parameters may be adjusted independently of the others without altering the phase and without harmonics suddenly dropping out. In some implementations, one or more parameters may be adjusted for pitch shifting, where the same timbre is preserved by preserving the pattern of harmonic amplitudes. Whisper synthesis may be performed, according to some implementations, where a set of evenly spaced harmonics is replaced by a white-noise source. This may preserve vowel shapes to give an indication as to what a given person would sound like if he were whispering.

(10) The noise subtraction module may be configured to subtract noise from the transformed audio information. Subtracting noise may include interpolating across peak points and trough points of harmonic pitch paths through the transformed audio information. The peak points may lie along harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and time for a given harmonic. The trough points may be positioned midway between peak points of adjacent harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and/or time. Values associated with individual ones of the trough point interpolations may be subtracted from values associated with

individual ones of the peak point interpolations to yield noise-reduced transformed audio information.

(11) The fence model module may be configured to suppress noise between harmonics of the sound in the transformed audio information by centering functions at individual harmonics in the transformed audio information. The functions may serve to suppress noise between the harmonics in order to yield noise-reduced transformed audio information. The width of a given function may be based on a bandwidth of a corresponding harmonic. In some implementations, individual ones of the functions utilized by the fence model module may include a Gaussian function. According to some implementations, individual ones of the functions may include a rectangular function.

(12) The reconstruction module may be configured to reconstruct an audio signal. In some implementations, one or more reverse transformations may be performed on transformed audio information and/or other non-time-domain information to obtain a reconstructed audio signal. The reconstruction module may be configured to reconstruct noise-reduced transformed audio information obtained from the noise subtraction module, the fence model module, and/or another source of noise-reduced transformed audio information. A reverse transformation used by the reconstruction module may correspond to a reverse and/or inverse of a transform performed on the original audio signal to produce the transformed audio information.

(13) These and other objects, features, and characteristics of the system and/or method disclosed herein, as well as the methods of operation and functions of the related elements of structure and the combination of parts and economies of manufacture, will become more apparent upon consideration of the following description and the appended claims with reference to the accompanying drawings, all of which form a part of this specification, wherein like reference numerals designate corresponding parts in the various figures. It is to be expressly understood, however, that the drawings are for the purpose of illustration and description only and are not intended as a definition of the limits of the invention. As used in the specification and in the claims, the singular form of "a", "an", and "the" include plural referents unless the context clearly dictates otherwise.

**BRIEF DESCRIPTION OF THE DRAWINGS**

(14) FIG. 1 illustrates a system configured for reconstructing an audio signal from transformed audio information, in accordance with one or more implementations.

(15) FIG. 2 illustrates an exemplary plot of transformed audio information.

(16) FIG. 3 illustrates a method for reconstructing an audio signal from transformed audio information by resynthesizing the audio signal based on individual harmonics and corresponding pitches determined from the transformed audio information, in accordance with one or more implementations.

(17) FIG. 4 illustrates a method for reconstructing an audio signal from transformed audio information where noise is subtracted from the transformed audio information, in accordance with one or more implementations.

(18) FIG. 5 illustrates a method for reconstructing an audio signal from transformed audio information where noise is suppressed between harmonics of the sound in the transformed audio information, in accordance with one or more implementations.

**DETAILED DESCRIPTION**

(19) FIG. 1 illustrates a system 10 configured for reconstructing an audio signal from transformed audio information, in accordance with one or more implementations. The transformed audio signal may be obtained from the audio signal through Fourier Transform, Fast Fourier Transform, Short Time Fourier Transform, Spectral Motion Transform, and/or other transforms. The audio signal may be resynthesized based on individual harmonics and corresponding pitches determined from the transformed audio information. Noise may be subtracted from the transformed audio information by interpolating across peak points and across trough points of harmonic pitch paths through the transformed audio information, and subtracting the trough point interpolations from the peak point interpolations. Noise between harmonics of the sound may be suppressed in the transformed audio information by centering functions at individual harmonics in the transformed audio information, the functions serving to suppress noise between the harmonics.

(20) The system 10 may be implemented in an overarching system (not shown) configured to process the audio signal. For example, the overarching system may be configured to segment sounds represented in the audio signal (e.g., divide sounds into groups corresponding to different sources, such as human speakers, within the audio signal), classify sounds represented in the audio signal (e.g., attribute sounds to

specific sources, such as specific human speakers), reconstruct sounds represented in the audio signal, and/or process the audio signal in other ways. In some implementations, system 10 may include one or more of one or more processors 12, electronic storage 14, a user interface 16, and/or other components.

(21) The processor 12 may be configured to execute one or more computer program modules. The computer program modules may be configured to execute the computer program module(s) by software; hardware; firmware; some combination of software, hardware, and/or firmware; and/or other mechanisms for configuring processing capabilities on processor 12. In some implementations, the one or more computer program modules may include one or more of an audio information module 18, a resynthesis module 20, a noise subtraction module 22, a fence model module 24, a reconstruction module 26, and/or other modules.

(22) The audio information module 18 may be configured to obtain transformed audio information representing one or more sounds. The transformed audio information may include a transformation of an audio signal into the frequency domain, a pseudo-frequency domain, a dynamical-frequency domain, and/or other non-temporal domains. By way of non-limiting example, the transformed audio information may be transformed from the audio signal by way of a Fourier transform, a fast-Fourier transform, a short-time-Fourier transform, and/or other transforms. The transformed audio information may include pitch and/or pitch information associated with the audio signal. In some implementations, pitch and/or pitch information may be determined as described one or both of U.S. Patent Application Serial No. 13/205,483, filed August 8, 2011, and entitled "System And Method For Tracking Sound Pitch Across An Audio Signal", and/or U.S. Patent Application Serial No. 13/205,521, filed August 8, 2011, and entitled "System And Method For Tracking Sound Pitch Across An Audio Signal Using Harmonic Envelope," which are hereby incorporated by reference into the present application in their entireties. The transformed audio information may include a transformation of an audio signal into a frequency-chirp domain, such as that described in U.S. Patent Application No. 13/205,424, filed August 8, 2011, and entitled "System And Method For Processing Sound Signals Implementing A Spectral Motion Transform" ("the 'YYY Application") which is hereby incorporated into this disclosure by reference in its entirety. The transformed audio information may have been transformed in discrete time sample windows over the audio signal. The time sample

windows may be overlapping or non-overlapping in time. Generally, the transformed audio information may specify magnitude of a coefficient related to signal intensity as a function of frequency, time, chirp, and/or other parameters for an audio signal within a time sample window.

(23) By way of illustration, FIG. 2 depicts an exemplary plot 28 of transformed audio information. The plot 28 may depict a magnitude of a coefficient related to energy as a function of frequency. The transformed audio information represented by plot 28 may include a harmonic sound, represented by a series of spikes 30 in the magnitude of the coefficient at the frequencies of the harmonics of the harmonic sound. Assuming that the sound is harmonic, spikes 30 may be spaced apart at intervals that correspond to the pitch ( $\phi$ ) of the harmonic sound. As such, individual spikes 30 may correspond to individual ones of the overtones of the harmonic sound.

(24) Other spikes (e.g., spikes 32 and/or 34) may be present in the transformed audio information. These spikes may not be associated with harmonic sound corresponding to spikes 30. The difference between spikes 30 and spike(s) 32 and/or 34 may not be amplitude, but instead frequency, as spike(s) 32 and/or 34 may not be at a harmonic frequency of the harmonic sound. As such, these spikes 32 and/or 34, and the rest of the amplitude between spikes 30 may be a manifestation of noise in the audio signal. As used in this instance, "noise" may not refer to a single auditory noise, but instead to sound (whether or not such sound is harmonic, diffuse, white, or of some other type) other than the harmonic sound associated with spikes 30.

(25) The transformation that yields the transformed audio information from the audio signal may result in the coefficient related to signal intensity being a complex number. The transformation may include an operation to make the complex number a real number. This may include, for example, taking the square of the modulus of the complex number, and/or other operations for making the complex number a real number. In some implementations, the complex number for the coefficient generated by the transform may be preserved. In such implementations, for example, the real and imaginary portions of the coefficient may be analyzed separately, at least at first. By way of illustration, plot 28 may represent the real portion of the coefficient, and a separate plot (not shown) may represent the imaginary portion of the coefficient as a function of frequency. The plot representing the imaginary portion of the coefficient as

a function of frequency may have spikes at the harmonics of the harmonic sound that corresponds to spikes 30.

(26) In some implementations, the transformed audio information may represent all of the energy present in the audio signal, or a portion of the energy present in the audio signal. For example, if the transformation of the audio signal places the audio signal into the frequency-chirp domain, the coefficient related to signal intensity may be specified as a function of frequency and fractional chirp rate. In such examples, the transformed audio information may include a representation of the energy present in the audio signal having a common fractional chirp rate (e.g., a one dimensional slice through a two-dimensional frequency-chirp domain to produce a frequency domain representation with a fixed chirp rate and/or other fixed parameter).

(27) Referring back to FIG. 1, resynthesis module 20 may be configured to resynthesize the audio signal based on individual harmonics and corresponding pitches determined from the transformed audio information. According to some implementations, resynthesizing the audio signal may include tracking one or more pitches of the sound to determine individual pitches and corresponding amplitudes as a function of time for individual harmonics of the sound. Individual harmonics may be synthesized using oscillators corresponding to individual harmonics. Synthesizing individual harmonics may include, for a given harmonic, integrating a corresponding pitch over time to determine the unwrapped phase of the given harmonic. Individual ones of the oscillators may be based on a cosine function. The synthesized harmonics may be summed to obtain the resynthesized audio signal.

(28) According to some implementations, the output  $y$  as a function of time  $t$  of the  $i^{\text{th}}$  oscillator may be expressed as, or similar to,

$$y_i(t) = \cos \int_0^t i \phi(\tau) d\tau ,$$

where  $\phi$  is pitch (first harmonic) as a function of time. This equation may be fixed, so the entire representation of a sound is stored in the pitch and harmonic amplitude parameters. Time courses may be represented sparsely because pitch and envelope (the set of harmonic amplitudes) change slowly per time relative to the sampling rate. For example, a cubic spline with 20 knots may provide an accurate fit to the pitch time course over one second for a human voice. Similarly, the harmonic amplitudes may be

represented with about 10 knots along the frequency dimension and 20 per second in time to form an "amplitude surface" (e.g., amplitude as a function of frequency and time, and/or transformed audio information) expressing the changing envelope. Some or all harmonic amplitudes and envelopes for synthesizing consonants with a white noise source may be shaped by such an amplitude surface.

(29) In some implementations, resynthesis module 20 may be configured to solve any phase problems because the audio signal may be built through integration, where phase is a consequence of the audio signal and not something that needs to be factored in. Also, the degree of compression of the resynthesized audio signal may go below a kB per second for voice, which is far better than the current mp3 standard.

(30) The resynthesized audio signal may be built from oscillators and parameters that specify pitch and harmonic amplitudes as a function of time. One or more of these parameters may be adjusted independently of the others without altering the phase and without harmonics suddenly dropping out.

(31) In some implementations, individual ones of the oscillators may include a white noise source to simulate a whispered version of the voice that retains word shaping and speech rhythms. Parameters may be altered to adjust for known channel distortions. For example, cell phones vary subtly in their pass-band, but generally have the same approximate high and low roll-offs. A correction may be made by dividing the transformed audio information by the roll-off transfer function.

(32) The noise subtraction module 22 may be configured to subtract noise from the transformed audio information. Subtracting noise may include interpolating across peak points of harmonic pitch paths through the transformed audio information. The peak points may lie along harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and time for a given harmonic. In some implementations, interpolation across the peak points may include polynomial interpolation, use of splines, and/or other interpolation techniques.

(33) Subtracting noise may further include interpolating across trough points of harmonic pitch paths through the transformed audio information. The trough points may be positioned midway between peak points of adjacent harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and time. In some implementations, interpolation across the trough points may include polynomial interpolation, use of splines, and/or other interpolation techniques. Such

splines may include linear, quadratic, cubic, and/or other splines. Values associated with individual ones of the trough point interpolations may be subtracted from values associated with individual ones of the peak point interpolations to yield noise-reduced transformed audio information.

(34) The fence model module 24 may be configured to suppress noise between harmonics of the sound in the transformed audio information by centering functions at individual harmonics in the transformed audio information. The functions may serve to suppress noise between the harmonics in order to yield noise-reduced transformed audio information. The width of a given function may be based on a bandwidth of a corresponding harmonic.

(35) In some implementations, individual ones of the functions utilized by fence model module 24 may include a Gaussian function. Such a Gaussian function may be configured to suppress information between the harmonics. The Gaussian function may be configured to replace information associated with individual harmonics with Gaussian (or other) curves to provide noise-reduced transformed audio information. A given Gaussian curve may be fitted to a corresponding harmonic.

(36) An audio signal may be reconstructed from the noise-reduced transformed audio information, as discussed in connection with the reconstruction module 26. Such a reconstructed audio signal may closely resemble the undistorted original audio signal, even down to 3dB noise. Additionally, the reconstructed audio signal may be more compact relative to the original audio signal because only the harmonic frequencies and corresponding amplitudes need to be transmitted to resynthesize the reconstructed audio signal.

(37) According to some implementations, individual ones of the functions may include a rectangular fence. Such a fence may be configured to zero information between the harmonics while preserving information associated with the harmonics. In some implementations, one or more functions utilized by fence model module 24 may be separately applied to real and imaginary components of the transformed audio information.

(38) The reconstruction module 26 may be configured to reconstruct an audio signal and/or portions of an audio signal (e.g., vowel and/or consonant sounds). In some implementations, one or more reverse transformations may be performed on transformed audio information and/or other non-time-domain information to obtain a

reconstructed audio signal. That is, reconstruction may include converting a frequency domain representation and/or frequency-chirp domain representation to a time-domain representation, according to some implementations. The reconstruction module 26 may be configured to reconstruct noise-reduced transformed audio information obtained from noise subtraction module 22, fence model module 24, and/or another source of noise-reduced transformed audio information. A reverse transformation used by reconstruction module 26 may correspond to a reverse and/or inverse of a transform performed on the original audio signal to produce the transformed audio information.

(39) Processor 12 may be configured to provide information processing capabilities in system 10. As such, processor 12 may include one or more of a digital processor, an analog processor, a digital circuit designed to process information, an analog circuit designed to process information, a state machine, and/or other mechanisms for electronically processing information. Although processor 12 is shown in FIG. 1 as a single entity, this is for illustrative purposes only. In some implementations, processor 12 may include a plurality of processing units. These processing units may be physically located within the same device, or processor 12 may represent processing functionality of a plurality of devices operating in coordination (e.g., "in the cloud", and/or other virtualized processing solutions).

(40) It should be appreciated that although modules 18, 20, 22, 24, and 26 are illustrated in FIG. 1 as being co-located within a single processing unit, in implementations in which processor 12 includes multiple processing units, one or more of modules 18, 20, 22, 24, and/or 26 may be located remotely from the other modules. The description of the functionality provided by the different modules 18, 20, 22, 24, and/or 26 described below is for illustrative purposes, and is not intended to be limiting, as any of modules 18, 20, 22, 24, and/or 26 may provide more or less functionality than is described. For example, one or more of modules 18, 20, 22, 24, and/or 26 may be eliminated, and some or all of its functionality may be provided by other ones of modules 18, 20, 22, 24, and/or 26. As another example, processor 12 may be configured to execute one or more additional modules that may perform some or all of the functionality attributed below to one of modules 18, 20, 22, 24, and/or 26.

(41) Electronic storage 14 may comprise electronic storage media that stores information. The electronic storage media of electronic storage 14 may include one or

both of system storage that is provided integrally (i.e., substantially non-removable) with system 10 and/or removable storage that is removably connectable to system 10 via, for example, a port (e.g., a USB port, a firewire port, etc.) or a drive (e.g., a disk drive, etc.). Electronic storage 14 may include one or more of optically readable storage media (e.g., optical disks, etc.), magnetically readable storage media (e.g., magnetic tape, magnetic hard drive, floppy drive, etc.), electrical charge-based storage media (e.g., EEPROM, RAM, etc.), solid-state storage media (e.g., flash drive, etc.), and/or other electronically readable storage media. Electronic storage 14 may include virtual storage resources, such as storage resources provided via a cloud and/or a virtual private network. Electronic storage 14 may store software algorithms, information determined by processor 12, information received via user interface 16, and/or other information that enables system 10 to function properly. Electronic storage 14 may be a separate component within system 10, or electronic storage 14 may be provided integrally with one or more other components of system 10 (e.g., processor 12).

(42) User interface 16 may be configured to provide an interface between system 10 and users. This may enable data, results, and/or instructions and any other communicable items, collectively referred to as "information," to be communicated between the users and system 10. Examples of interface devices suitable for inclusion in user interface 16 include a keypad, buttons, switches, a keyboard, knobs, levers, a display screen, a touch screen, speakers, a microphone, an indicator light, an audible alarm, and a printer. It is to be understood that other communication techniques, either hard-wired or wireless, are also contemplated by the present invention as user interface 16. For example, the present invention contemplates that user interface 16 may be integrated with a removable storage interface provided by electronic storage 14. In this example, information may be loaded into system 10 from removable storage (e.g., a smart card, a flash drive, a removable disk, etc.) that enables the user(s) to customize the implementation of system 10. Other exemplary input devices and techniques adapted for use with system 10 as user interface 14 include, but are not limited to, an RS-232 port, RF link, an IR link, modem (telephone, cable or other). In short, any technique for communicating information with system 10 is contemplated by the present invention as user interface 14.

(43) FIG. 3 illustrates a method 36 for reconstructing an audio signal from transformed audio information by resynthesizing the audio signal based on individual harmonics and corresponding pitches determined from the transformed audio information, in accordance with one or more implementations. The operations of method 36 presented below are intended to be illustrative. In some implementations, method 36 may be accomplished with one or more additional operations not described, and/or without one or more of the operations discussed. Additionally, the order in which the operations of method 36 are illustrated in FIG. 3 and described below is not intended to be limiting.

(44) At operation 38, transformed audio information representing a sound may be obtained. The transformed audio information may specify magnitude of a coefficient related to signal intensity as a function of frequency for the audio signal and time. In some implementations, operation 38 may be performed by an audio information module that is the same as or similar to audio information module 18 (shown in FIG. 1 and described above).

(45) At operation 40, one or more pitches of the sound may be tracked to determine individual pitches and corresponding amplitudes as a function of time for individual harmonics of the sound. In some implementations, operation 40 may be performed by a resynthesis module that is the same as or similar to resynthesis module 20 (shown in FIG. 1 and described above).

(46) At operation 42, individual harmonics may be synthesized using oscillators corresponding to individual harmonics. According to some implementations, only harmonics associated with a desired sound (e.g., speech from a particular speaker) within the original audio signal may be synthesized, thus excluding unwanted sounds. In some implementations, operation 42 may be performed by a resynthesis module that is the same as or similar to resynthesis module 20 (shown in FIG. 1 and described above).

(47) At operation 44, the synthesized harmonics may be summed to obtain the resynthesized audio signal. In some implementations, operation 44 may be performed by a resynthesis module that is the same as or similar to resynthesis module 20 (shown in FIG. 1 and described above).

(48) FIG. 4 illustrates a method 46 for reconstructing an audio signal from transformed audio information where noise is subtracted from the transformed audio

information, in accordance with one or more implementations. The operations of method 46 presented below are intended to be illustrative. In some implementations, method 46 may be accomplished with one or more additional operations not described, and/or without one or more of the operations discussed. Additionally, the order in which the operations of method 46 are illustrated in FIG. 4 and described below is not intended to be limiting.

(49) At operation 48, transformed audio information representing a sound may be obtained. The transformed audio information may specify magnitude of a coefficient related to signal intensity as a function of frequency for the audio signal and time. In some implementations, operation 48 may be performed by an audio information module that is the same as or similar to audio information module 18 (shown in FIG. 1 and described above).

(50) At operation 50, peak points of harmonic pitch paths may be interpolated through the transformed audio information. The peak points may lie along harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and time for a given harmonic. In some implementations, operation 50 may be performed by a noise subtraction module that is the same as or similar to noise subtraction module 22 (shown in FIG. 1 and described above).

(51) At operation 52, trough points of harmonic pitch paths may be interpolated through the transformed audio information. The trough points may be positioned midway between peak points of adjacent harmonic frequencies in the transformed audio information, and may be determined as a function of frequency and time. In some implementations, operation 52 may be performed by a noise subtraction module that is the same as or similar to noise subtraction module 22 (shown in FIG. 1 and described above).

(52) At operation 54, values associated with individual ones of the trough point interpolations may be subtracted from values associated with individual ones of the peak point interpolations to yield noise-reduced transformed audio information. In some implementations, operation 54 may be performed by a noise subtraction module that is the same as or similar to noise subtraction module 22 (shown in FIG. 1 and described above).

(53) At operation 56, the audio signal may be reconstructed based on a reverse transformation of the noise-reduced transformed audio information. In some

implementations, operation 56 may be performed by a reconstruction module that is the same as or similar to reconstruction module 26 (shown in FIG. 1 and described above).

(54) FIG. 5 illustrates a method 58 for reconstructing an audio signal from transformed audio information where noise is suppressed between harmonics of the sound in the transformed audio information, in accordance with one or more implementations. The operations of method 58 presented below are intended to be illustrative. In some implementations, method 58 may be accomplished with one or more additional operations not described, and/or without one or more of the operations discussed. Additionally, the order in which the operations of method 58 are illustrated in FIG. 5 and described below is not intended to be limiting.

(55) At operation 60, transformed audio information representing a sound may be obtained. The transformed audio information may specify magnitude of a coefficient related to signal energy as a function of frequency for the audio signal and time. In some implementations, operation 60 may be performed by an audio information module that is the same as or similar to audio information module 18 (shown in FIG. 1 and described above).

(56) At operation 62, noise between harmonics of the sound may be suppressed in the transformed audio information by centering functions at individual harmonics in the transformed audio information. The functions may serve to suppress noise between the harmonics to yield noise-reduced transformed audio information. The width of a given function may be based on a bandwidth of a corresponding harmonic. In some implementations, operation 62 may be performed by a fence model module that is the same as or similar to fence model module 24 (shown in FIG. 1 and described above).

(57) At operation 64, the audio signal may be reconstructed based on a reverse transformation of the noise-reduced transformed audio information. In some implementations, operation 64 may be performed by a reconstruction module that is the same as or similar to reconstruction module 26 (shown in FIG. 1 and described above).

(58) In some implementations, methods 36, 46, and/or 58 may be implemented in one or more processing devices (e.g., a digital processor, an analog processor, a digital circuit designed to process information, an analog circuit designed to process information, a state machine, and/or other mechanisms for electronically processing

information). The one or more processing devices may include one or more devices executing some or all of the operations of methods 36, 46, and/or 58 in response to instructions stored electronically on an electronic storage medium. The one or more processing devices may include one or more devices configured through hardware, firmware, and/or software to be specifically designed for execution of one or more of the operations of methods 36, 46, and/or 58.

(59) Although the system(s) and/or method(s) of this disclosure have been described in detail for the purpose of illustration based on what is currently considered to be the most practical and preferred implementations, it is to be understood that such detail is solely for that purpose and that the disclosure is not limited to the disclosed implementations, but, on the contrary, is intended to cover modifications and equivalent arrangements that are within the spirit and scope of the appended claims. For example, it is to be understood that the present disclosure contemplates that, to the extent possible, one or more features of any implementation can be combined with one or more features of any other implementation.

**What is claimed is:**

1. A method for reconstructing an audio signal from transformed audio information, the method comprising:
  - obtaining transformed audio information representing a sound, wherein the transformed audio information specifies magnitude of a coefficient related to signal intensity as a function of frequency for the audio signal and time; and
  - resynthesizing the audio signal based on individual harmonics and corresponding pitches determined from the transformed audio information, wherein resynthesizing the audio signal includes:
    - tracking one or more pitches of the sound to determine individual pitches and corresponding amplitudes as a function of time for individual harmonics of the sound,
    - synthesizing individual harmonics using oscillators corresponding to individual harmonics, and
    - summing the synthesized harmonics to obtain the resynthesized audio signal.
2. The method of claim 1, wherein synthesizing individual harmonics includes, for a given harmonic, integrating a corresponding pitch over time to determine a phase of the given harmonic.
3. The method of claim 1, wherein individual ones of the oscillators are based on a cosine function.
4. The method of claim 1, wherein individual ones of the oscillators include a white noise source.
5. A method for reconstructing an audio signal from transformed audio information, the method comprising:
  - obtaining transformed audio information representing a sound, wherein the transformed audio information specifies magnitude of a coefficient related to signal intensity as a function of frequency for the audio signal and time;

subtracting noise from the transformed audio information, wherein subtracting the noise includes:

interpolating across peak points of harmonic pitch paths through the transformed audio information, the peak points lying along harmonic frequencies in the transformed audio information,

interpolating across trough points of harmonic pitch paths through the transformed audio information, the trough points being positioned midway between peak points of adjacent harmonic frequencies in the transformed audio information, and

subtracting values associated with individual ones of the trough point interpolations from values associated with individual ones of the peak point interpolations to yield noise-reduced transformed audio information; and

reconstructing the audio signal based on a reverse transformation of the noise-reduced transformed audio information.

6. The method of claim 5, wherein splines are used to interpolate across the peak points.
7. The method of claim 5, wherein splines are used to interpolate across the trough points.
8. The method of claim 5, wherein polynomial interpolation is used to interpolate across the peak points.
9. The method of claim 5, wherein polynomial interpolation is used to interpolate across the trough points.
10. A method for reconstructing an audio signal from transformed audio information, the method comprising:
  - obtaining transformed audio information representing a sound, wherein the transformed audio information specifies magnitude of a coefficient related to signal intensity as a function of frequency for the audio signal and time;

suppressing noise between harmonics of the sound in the transformed audio information by centering functions at individual harmonics in the transformed audio information, the functions serving to suppress noise between the harmonics to yield noise-reduced transformed audio information, wherein the width of a given function is based on a bandwidth of a corresponding harmonic; and

reconstructing the audio signal based on a reverse transformation of the noise-reduced transformed audio information.

11. The method of claim 10, wherein individual ones of the functions include a Gaussian function, the Gaussian function being configured to suppress information between the harmonics and replace information associated with the harmonics with a Gaussian curve.
12. The method of claim 11, wherein a given Gaussian curve is fitted to a corresponding harmonic.
13. The method of claim 10, wherein individual ones of the functions include a rectangular function, the rectangular function being configured to zero information between the harmonics while preserving information associated with the harmonics.
14. The method of claim 10, wherein the functions are separately applied to real and imaginary components of the transformed audio information.

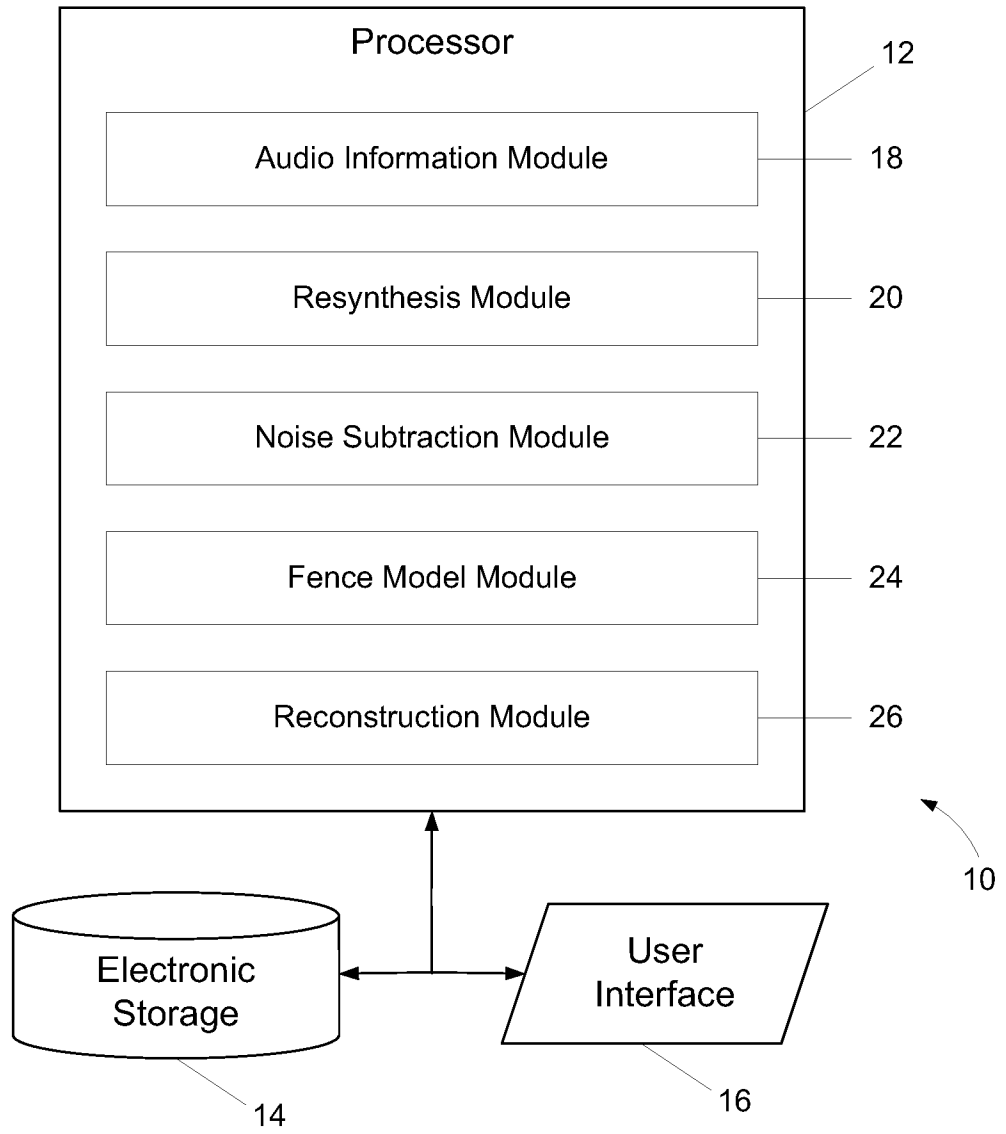


FIG. 1




FIG. 2

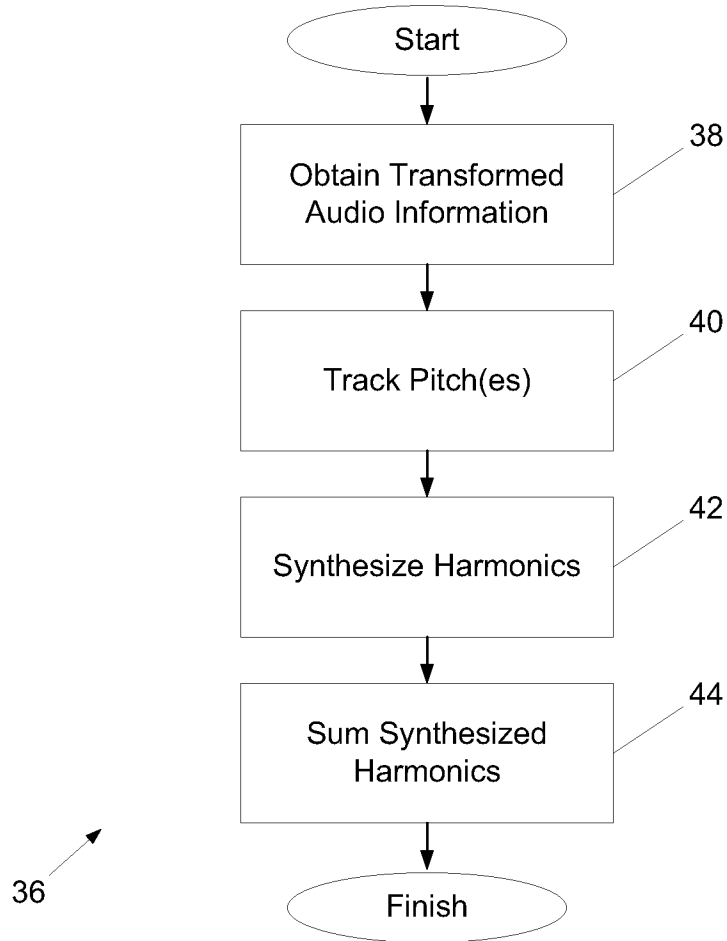


FIG. 3

4/5

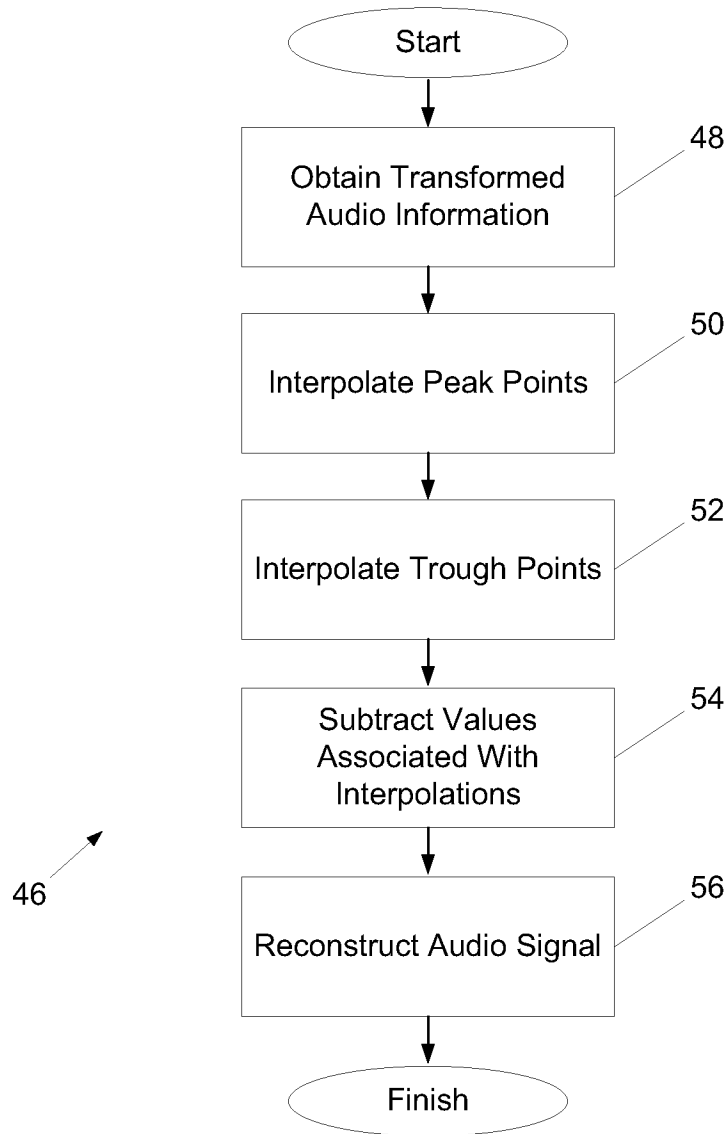


FIG. 4

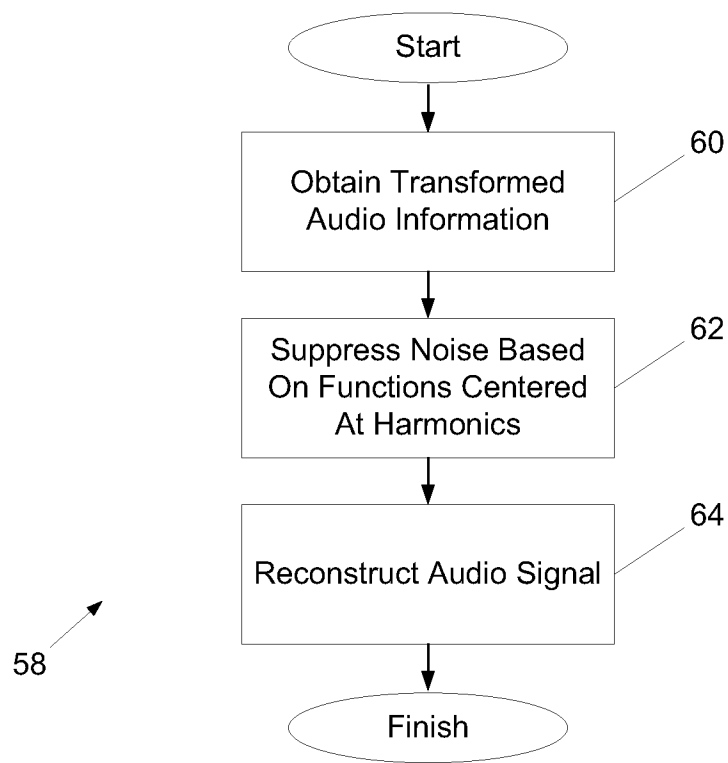


FIG. 5