

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
5 December 2002 (05.12.2002)

PCT

(10) International Publication Number
WO 02/097791 A1

(51) International Patent Classification⁷: **G10L 11/00**,
H04N 5/04

[US/US]; 100 Potrero Avenue, San Francisco, CA 94103 (US).

(21) International Application Number: PCT/US02/05806

(72) Inventors; and

(22) International Filing Date: 25 February 2002 (25.02.2002)

(75) Inventors/Applicants (for US only): **CROCKETT, Brett, G.** [US/US]; 100 Potrero Avenue, San Francisco, CA 94103 (US). **SMITHERS, Michael, J.** [AU/US]; 100 Potrero Avenue, San Francisco, CA 94103 (US).

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/293,825	25 May 2001 (25.05.2001)	US
10/045,644	11 January 2002 (11.01.2002)	US
60/351,498	23 January 2002 (23.01.2002)	US
PCT/US02/04317		
	12 February 2002 (12.02.2002)	US

(74) Agents: **GALLAGHER, Thomas** et al.; Gallagher & Lathrop, 601 California Street, Suite 1111, San Francisco, CA 94108-2805 (US).

(81) Designated States (national): AE, AG, AL, AM, AT (utility model), AT, AU (petty patent), AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK (utility model), SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:


US	60/293,825 (CIP)
Filed on	25 May 2001 (25.05.2001)
US	60/351,498 (CIP)
Filed on	23 January 2002 (23.01.2002)
US	10/045,644 (CIP)
Filed on	11 January 2002 (11.01.2002)
US	Not furnished (CIP)
Filed on	12 February 2002 (12.02.2002)

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent

(71) Applicant (for all designated States except US): **DOLBY LABORATORIES LICENSING CORPORATION**

[Continued on next page]

(54) Title: METHOD FOR TIME ALIGNING AUDIO SIGNALS USING CHARACTERIZATIONS BASED ON AUDITORY EVENTS



(57) Abstract: A method for time aligning audio signal, wherein one signal has been derived from the other or both have been derived from another signal, comprises deriving reduced-information characterizations of the audio signals, auditory scene analysis. The time offset of one characterization with respect to the other characterizations is calculated and the temporal relationship of the audio signals with respect to each other is modified in response to the time offset such that the audio signals are coincident with each other. These principles may also be applied to a method for time aligning a video signal and an audio signal that will be subjected to differential time offsets.



WO 02/097791 A1



(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

- 1 -

DESCRIPTION**Method for Time Aligning Audio Signals Using Characterizations
Based on Auditory Events**

5

TECHNICAL FIELD

The invention relates to audio signals. More particularly, the invention relates to characterizing audio signals and using characterizations to time align or synchronize audio signals wherein one signal has been derived from the other or in which both have been derived from the same other signal. Such synchronization is useful, for example, in restoring television audio to video synchronization (lip-sync) and in detecting a watermark embedded in an audio signal (the watermarked signal is compared to an unwatermarked version of the signal). The invention may be implemented so that a low processing power process brings two such audio signals into substantial temporal alignment.

15

BACKGROUND ART

The division of sounds into units perceived as separate is sometimes referred to as "auditory event analysis" or "auditory scene analysis" ("ASA"). An extensive discussion of auditory scene analysis is set forth by Albert S. Bregman in his book *Auditory Scene Analysis - The Perceptual Organization of Sound*, Massachusetts Institute of Technology, 1991, Fourth printing, 2001, Second MIT Press paperback edition. In addition, United States Patent 6,002,776 to Bhadkamkar, et al, December 14, 1999 cites publications dating back to 1976 as "prior art work related to sound separation by auditory scene analysis." However, the Bhadkamkar, et al patent discourages the practical use of auditory scene analysis, concluding that "[t]echniques involving auditory scene analysis, although interesting from a scientific point of view as models of human auditory processing, are currently far too computationally demanding and specialized to be considered practical techniques for sound separation until fundamental progress is made."

30

Bregman notes in one passage that "[w]e hear discrete units when the sound changes abruptly in timbre, pitch, loudness, or (to a lesser extent) location in space."

(*Auditory Scene Analysis - The Perceptual Organization of Sound*, supra at page 469). Bregman also discusses the perception of multiple simultaneous sound streams when, for example, they are separated in frequency.

There are many different methods for extracting characteristics or features from audio. Provided the features or characteristics are suitably defined, their extraction can be performed using automated processes. For example "ISO/IEC JTC 1/SC 29/WG 11" (MPEG) is currently standardizing a variety of audio descriptors as part of the MPEG-7 standard. A common shortcoming of such methods is that they ignore ASA. Such methods seek to measure, periodically, certain "classical" signal processing parameters such as pitch, amplitude, power, harmonic structure and spectral flatness. Such parameters, while providing useful information, do not analyze and characterize audio signals into elements perceived as separate according to human cognition.

Auditory scene analysis attempts to characterize audio signals in a manner similar to human perception by identifying elements that are separate according to human cognition. By developing such methods, one can implement automated processes that accurately perform tasks that heretofore would have required human assistance.

The identification of separately perceived elements would allow the unique identification of an audio signal using substantially less information than the full signal itself. Compact and unique identifications based on auditory events may be employed, for example, to identify a signal that is copied from another signal (or is copied from the same original signal as another signal).

25 ***DISCLOSURE OF THE INVENTION***

A method is described that generates a unique reduced-information characterization of an audio signal that may be used to identify the audio signal. The characterization may be considered a "signature" or "fingerprint" of the audio signal. According to the present invention, an auditory scene analysis (ASA) is performed to identify auditory events as the basis for characterizing an audio signal. Ideally, the

auditory scene analysis identifies auditory events that are most likely to be perceived by a human listener even after the audio has undergone processing, such as low bit rate coding or acoustic transmission through a loudspeaker. The audio signal may be characterized by the boundary locations of auditory events and, optionally, by the dominant frequency subband of each auditory event. The resulting information pattern, constitutes a compact audio fingerprint or signature that may be compared to the fingerprint or signature of a related audio signal to determine quickly and/or with low processing power the time offset between the original audio signals. The reduced-information characteristics have substantially the same relative timing as the audio signals they represent.

The auditory scene analysis method according to the present invention provides a fast and accurate method of time aligning two audio signals, particularly music, by comparing signatures containing auditory event information. ASA extracts information underlying the perception of similarity, in contrast to traditional methods that extract features less fundamental to perceiving similarities between audio signals (such as pitch amplitude, power, and harmonic structure). The use of ASA improves the chance of finding similarity in, and hence time aligning, material that has undergone significant processing, such as low bit coding or acoustic transmission through a loudspeaker.

In the embodiments discussed below, it is assumed that the two audio signals under discussion are derived from a common source. The method of the present invention determines the time offset of one such audio signal with respect to the other so that they may be brought into approximate synchronism with respect to each other.

Although in principle the invention may be practiced either in the analog or digital domain (or some combination of the two), in practical embodiments of the invention, audio signals are represented by samples in blocks of data and processing is done in the digital domain.

Referring to FIG. 1A, auditory scene analysis 2 is applied to an audio signal in order to produce a "signature" or "fingerprint," related to that signal. In this case,

- 4 -

there are two audio signals of interest. They are similar in that one is derived from the other or both have been previously derived from the same original signal. Thus, auditory scene analysis is applied to both signals. For simplicity, FIG. 1A shows only the application of ASA to one signal. As shown in FIG. 1B, the signatures for the two audio signals, Signature 1 and Signature 2, are applied to a time offset calculation function 4 that calculates an "offset" output that is a measure of the relative time offset between the two signatures.

Because the signatures are representative of the audio signals but are substantially shorter (*i.e.*, they are more compact or have fewer bits) than the audio signals from which they were derived, the time offset between the signatures can be determined much faster than it would take to determine the time offset between the audio signals. Moreover, because the signatures retain substantially the same relative timing relationship as the audio signals from which they are derived, a calculation of the offset between the signatures is usable to time align the original audio signals.

Thus, the offset output of function 4 is applied to a time alignment function 6. The time alignment function also receives the two audio signals, Audio signal 1 and Audio signal 2 (from which Signature 1 and 2 were derived), and provides two audio signal outputs, Audio signal 3 and Audio signal 4. It is desired to adjust the relative timing of Audio signal 1 with respect to Audio signal 2 so that they are in time alignment (synchronism) or are nearly in time alignment. To accomplish this, one may be time shifted with respect to the other or, in principle, both may be time shifted. In practice, one of the audio signals is a "pass through" of Audio signal 1 or Audio signal 2 (*i.e.*, it is substantially the same signal) and the other is a time shifted version of the other audio signal that has been temporally modified so that Audio Signal 3 and Audio Signal 4 are in time synchronism or nearly in time synchronism with each other, depending on the resolution accuracy of the offset calculation and time alignment functions. If greater alignment accuracy is desired, further processing may be applied to Audio Signal 3 and/or Audio Signal 4 by one or more other processes that form no part of the present invention.

- 5 -

The time alignment of the signals may be useful, for example, in restoring television audio to video synchronization (lip-sync) and in detecting a watermark embedded in an audio signal. In the former case, a signature of the audio is embedded in the video signal prior to transmission or storage that may result in the audio and video getting out of synchronism. At a reproduction point, a signature may be derived from the audio signal and compared to the signature embedded in the video signal in order to restore their synchronism. Systems of that type not employing characterizations based on auditory scene analysis are described in U.S. Patents Re 33,535, 5,202,761, 6,211,919, and 6,246,439, all of which are incorporated herein by reference in their entireties. In the second case, an original version of an audio signal is compared to a watermarked version of the audio signal in order to recover the watermark. Such recovery requires close temporal alignment of the two audio signals. This may be achieved, at least to a first degree of alignment by deriving a signature of each audio signal to aid in time alignment of the original audio signals, as explained herein. Further details of FIGS. 1A and 1B are set forth below.

For some applications, the processes of FIGS. 1A and 1B should be real-time. For other applications, they need not be real-time. In a real-time application, the process stores a history (a few seconds, for example) of the auditory scene analysis for each input signal. Periodically, that event history is employed to update the offset calculation in order to continually correct the time offset. The auditory scene analysis information for each of the input signals may be generated in real time, or the information for either of the signals may already be present (assuming that some offline auditory scene analysis processing has already been performed). One use for a real-time system is, for example, an audio/video aligner as mentioned above. One series of event boundaries is derived from the audio; the other series of event boundaries is recovered from the video (assuming some previous embedding of the audio event boundaries into the video). The two event boundary sequences can be periodically compared to determine the time offset between the audio and video in order to improve the lip sync, for example.

Thus, both signatures may be generated from the audio signals at nearly the same time that the time offset of the signatures is calculated and used to modify the alignment of the audio signals to achieve their substantial coincidence. Alternatively, one of the signatures to be compared may be carried along with the audio signal from which it was derived, for example, by embedding the signature in another signal, such as a video signal as in the case of audio and video alignment as just described. As a further alternative, both signatures may be generated in advance and only the comparison and timing modification performed in real time. For example, in the case of two sources of the same television program (with both video and audio), both with embedded audio signatures; the respective television signals (with accompanying audio) could be synchronized (both video and audio) by comparing the recovered signatures. The relative timing relationship of the video and audio in each television signal would remain unaltered. The television signal synchronization would occur in real time, but neither signature would be generated at that time nor simultaneously with each other.

In accordance with aspects of the present invention, a computationally efficient process for dividing audio into temporal segments or "auditory events" that tend to be perceived as separate is provided.

A powerful indicator of the beginning or end of a perceived auditory event is believed to be a change in spectral content. In order to detect changes in timbre and pitch (spectral content) and, as an ancillary result, certain changes in amplitude, the audio event detection process according to an aspect of the present invention detects changes in spectral composition with respect to time. Optionally, according to a further aspect of the present invention, the process may also detect changes in amplitude with respect to time that would not be detected by detecting changes in spectral composition with respect to time.

In its least computationally demanding implementation, the process divides audio into time segments by analyzing the entire frequency band of the audio signal (full bandwidth audio) or substantially the entire frequency band (in practical implementations, band limiting filtering at the ends of the spectrum are often

- 7 -

employed) and giving the greatest weight to the loudest audio signal components.

This approach takes advantage of a psychoacoustic phenomenon in which at smaller time scales (20 milliseconds (msec) and less) the ear may tend to focus on a single auditory event at a given time. This implies that while multiple events may be

5 occurring at the same time, one component tends to be perceptually most prominent and may be processed individually as though it were the only event taking place.

Taking advantage of this effect also allows the auditory event detection to scale with the complexity of the audio being processed. For example, if the input audio signal being processed is a solo instrument, the audio events that are identified will likely be

10 the individual notes being played. Similarly for an input voice signal, the individual components of speech, the vowels and consonants for example, will likely be identified as individual audio elements. As the complexity of the audio increases, such as music with a drumbeat or multiple instruments and voice, the auditory event detection identifies the most prominent (*i.e.*, the loudest) audio element at any given
15 moment. Alternatively, the "most prominent" audio element may be determined by taking hearing threshold and frequency response into consideration.

Optionally, according to further aspects of the present invention, at the expense of greater computational complexity, the process may also take into consideration changes in spectral composition with respect to time in discrete
20 frequency bands (fixed or dynamically determined or both fixed and dynamically determined bands) rather than the full bandwidth. This alternative approach would take into account more than one audio stream in different frequency bands rather than assuming that only a single stream is perceptible at a particular time.

Even a simple and computationally efficient process according to an aspect of
25 the present invention for segmenting audio has been found useful to identify auditory events.

An auditory event detecting process of the present invention may be implemented by dividing a time domain audio waveform into time intervals or blocks and then converting the data in each block to the frequency domain, using either a
30 filter bank or a time-frequency transformation, such as a Discrete Fourier Transform

- 8 -

(DFT) (implemented as a Fast Fourier Transform (FFT) for speed). The amplitude of the spectral content of each block may be normalized in order to eliminate or reduce the effect of amplitude changes. Each resulting frequency domain representation provides an indication of the spectral content (amplitude as a function of frequency) of the audio in the particular block. The spectral content of successive blocks is compared and each change greater than a threshold may be taken to indicate the temporal start or temporal end of an auditory event.

In order to minimize the computational complexity, only a single band of frequencies of the time domain audio waveform may be processed, preferably either the entire frequency band of the spectrum (which may be about 50 Hz to 15 kHz in the case of an average quality music system) or substantially the entire frequency band (for example, a band defining filter may exclude the high and low frequency extremes).

Preferably, the frequency domain data is normalized, as is described below. The degree to which the frequency domain data needs to be normalized gives an indication of amplitude. Hence, if a change in this degree exceeds a predetermined threshold, that too may be taken to indicate an event boundary. Event start and end points resulting from spectral changes and from amplitude changes may be ORed together so that event boundaries resulting from either type of change are identified.

In practical embodiments in which the audio is represented by samples divided into blocks, each auditory event temporal start and stop point boundary necessarily coincides with a boundary of the block into which the time domain audio waveform is divided. There is a trade off between real-time processing requirements (as larger blocks require less processing overhead) and resolution of event location (smaller blocks provide more detailed information on the location of auditory events).

As a further option, as suggested above, but at the expense of greater computational complexity, instead of processing the spectral content of the time domain waveform in a single band of frequencies, the spectrum of the time domain waveform prior to frequency domain conversion may be divided into two or more frequency bands. Each of the frequency bands may then be converted to the

- 9 -

frequency domain and processed as though it were an independent channel. The resulting event boundaries may then be ORed together to define the event boundaries for that channel. The multiple frequency bands may be fixed, adaptive, or a combination of fixed and adaptive. Tracking filter techniques employed in audio noise reduction and other arts, for example, may be employed to define adaptive frequency bands (*e.g.*, dominant simultaneous sine waves at 800 Hz and 2 kHz could result in two adaptively-determined bands centered on those two frequencies).

Other techniques for providing auditory scene analysis may be employed to identify auditory events in the present invention.

DESCRIPTION OF THE DRAWINGS

FIG. 1A is a flow chart showing the process of extraction of a signature from an audio signal in accordance with the present invention. The audio signal may, for example, represent music (*e.g.*, a musical composition or "song").

FIG. 1B is a flow chart illustrating a process for the time alignment of two audio signal signals in accordance with the present invention.

FIG. 2 is a flow chart showing the process of extraction of audio event locations and the optional extraction of dominant subbands from an audio signal in accordance with the present invention.

FIG. 3 is a conceptual schematic representation depicting the step of spectral analysis in accordance with the present invention.

FIGS. 4A and 4B are idealized audio waveforms showing a plurality of auditory event locations and auditory event boundaries in accordance with the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

In a practical embodiment of the invention, the audio signal is represented by samples that are processed in blocks of 512 samples, which corresponds to about 11.6 msec of input audio at a sampling rate of 44.1 kHz. A block length having a time less than the duration of the shortest perceivable auditory event (about 20 msec)

is desirable. It will be understood that the aspects of the invention are not limited to such a practical embodiment. The principles of the invention do not require arranging the audio into sample blocks prior to determining auditory events, nor, if they are, of providing blocks of constant length. However, to minimize complexity, a fixed block length of 512 samples (or some other power of two number of samples) is useful for three primary reasons. First, it provides low enough latency to be acceptable for real-time processing applications. Second, it is a power-of-two number of samples, which is useful for fast Fourier transform (FFT) analysis. Third, it provides a suitably large window size to perform useful auditory scene analysis.

In the following discussions, the input signals are assumed to be data with amplitude values in the range $[-1, +1]$.

Auditory Scene Analysis 2 (FIG. 1A)

Following audio input data blocking (not shown), the input audio signal is divided into auditory events, each of which tends to be perceived as separate, in process 2 ("Auditory Scene Analysis") of FIG. 1A. Auditory scene analysis may be accomplished by an auditory scene analysis (ASA) process discussed above. Although one suitable process for performing auditory scene analysis is described in further detail below, the invention contemplates that other useful techniques for performing ASA may be employed.

FIG. 2 outlines a process in accordance with techniques of the present invention that may be used as the auditory scene analysis process of FIG. 1A. The ASA step or process 2 is composed of three general processing substeps. The first substep 2-1 ("Perform Spectral Analysis") takes the audio signal, divides it into blocks and calculates a spectral profile or spectral content for each of the blocks. Spectral analysis transforms the audio signal into the short-term frequency domain. This can be performed using any filterbank; either based on transforms or banks of band-pass filters, and in either linear or warped frequency space (such as the Bark scale or critical band, which better approximate the characteristics of the human ear). With any filterbank there exists a tradeoff between time and frequency. Greater time resolution, and hence shorter time intervals, leads to lower frequency resolution.

- 11 -

Greater frequency resolution, and hence narrower subbands, leads to longer time intervals.

The first substep calculates the spectral content of successive time segments of the audio signal. In a practical embodiment, described below, the ASA block size is 512 samples of the input audio signal (FIG.3). In the second substep 2-2, the differences in spectral content from block to block are determined ("Perform spectral profile difference measurements"). Thus, the second substep calculates the difference in spectral content between successive time segments of the audio signal. In the third substep 2-3 ("Identify location of auditory event boundaries"), when the spectral difference between one spectral-profile block and the next is greater than a threshold, the block boundary is taken to be an auditory event boundary. Thus, the third substep sets an auditory event boundary between successive time segments when the difference in the spectral profile content between such successive time segments exceeds a threshold. As discussed above, a powerful indicator of the beginning or end of a perceived auditory event is believed to be a change in spectral content. The locations of event boundaries are stored as a signature. An optional process step 2-4 ("Identify dominant subband") uses the spectral analysis to identify a dominant frequency subband that may also be stored as part of the signature.

In this embodiment, auditory event boundaries define auditory events having a length that is an integral multiple of spectral profile blocks with a minimum length of one spectral profile block (512 samples in this example). In principle, event boundaries need not be so limited.

Either overlapping or non-overlapping segments of the audio may be windowed and used to compute spectral profiles of the input audio. Overlap results in finer resolution as to the location of auditory events and, also, makes it less likely to miss an event, such as a transient. However, as time resolution increases, frequency resolution decreases. Overlap also increases computational complexity. Thus, overlap may be omitted. FIG. 3 shows a conceptual representation of non-overlapping 512 sample blocks being windowed and transformed into the frequency domain by the Discrete Fourier Transform (DFT). Each block may be windowed and

- 12 -

transformed into the frequency domain, such as by using the DFT, preferably implemented as a Fast Fourier Transform (FFT) for speed.

The following variables may be used to compute the spectral profile of the input block:

- 5 N = number of samples in the input signal
- M = number of windowed samples used to compute spectral profile
- P = number of samples of spectral computation overlap
- Q = number of spectral windows/regions computed

In general, any integer numbers may be used for the variables above.

- 10 However, the implementation will be more efficient if M is set equal to a power of 2 so that standard FFTs may be used for the spectral profile calculations. In a practical embodiment of the auditory scene analysis process, the parameters listed may be set to:

- M = 512 samples (or 11.6 msec at 44.1 kHz)
- 15 P = 0 samples (no overlap)

- The above-listed values were determined experimentally and were found generally to identify with sufficient accuracy the location and duration of auditory events. However, setting the value of P to 256 samples (50% overlap) has been found to be useful in identifying some hard-to-find events. While many different
- 20 types of windows may be used to minimize spectral artifacts due to windowing, the window used in the spectral profile calculations is an M-point Hanning, Kaiser-Bessel or other suitable, preferably non-rectangular, window. The above-indicated values and a Hanning window type were selected after extensive experimental analysis as they have shown to provide excellent results across a wide range of audio
 - 25 material. Non-rectangular windowing is preferred for the processing of audio signals with predominantly low frequency content. Rectangular windowing produces spectral artifacts that may cause incorrect detection of events. Unlike certain codec applications where an overall overlap/add process must provide a constant level, such a constraint does not apply here and the window may be chosen for characteristics
 - 30 such as its time/frequency resolution and stop-band rejection.

- 13 -

In substep 2-1 (FIG. 2), the spectrum of each M-sample block may be computed by windowing the data by an M-point Hanning, Kaiser-Bessel or other suitable window, converting to the frequency domain using an M-point Fast Fourier Transform, and calculating the magnitude of the FFT coefficients. The resultant data is normalized so that the largest magnitude is set to unity, and the normalized array of M numbers is converted to the log domain. The array need not be converted to the log domain, but the conversion simplifies the calculation of the difference measure in substep 2-2. Furthermore, the log domain more closely matches the log domain amplitude nature of the human auditory system. The resulting log domain values have a range of minus infinity to zero. In a practical embodiment, a lower limit can be imposed on the range of values; the limit may be fixed, for example -60 dB, or be frequency-dependent to reflect the lower audibility of quiet sounds at low and very high frequencies. (Note that it would be possible to reduce the size of the array to M/2 in that the FFT represents negative as well as positive frequencies).

Substep 2-2 calculates a measure of the difference between the spectra of adjacent blocks. For each block, each of the M (log) spectral coefficients from substep 2-1 is subtracted from the corresponding coefficient for the preceding block, and the magnitude of the difference calculated (the sign is ignored). These M differences are then summed to one number. Hence, for the whole audio signal, the result is an array of Q positive numbers; the greater the number the more a block differs in spectrum from the preceding block. This difference measure could also be expressed as an average difference per spectral coefficient by dividing the difference measure by the number of spectral coefficients used in the sum (in this case M coefficients).

Substep 2-3 identifies the locations of auditory event boundaries by applying a threshold to the array of difference measures from substep 2-2 with a threshold value. When a difference measure exceeds a threshold, the change in spectrum is deemed sufficient to signal a new event and the block number of the change is recorded as an event boundary. For the values of M and P given above and for log domain values (in substep 2-1) expressed in units of dB, the threshold may be set equal to 2500 if

- 14 -

the whole magnitude FFT (including the mirrored part) is compared or 1250 if half the FFT is compared (as noted above, the FFT represents negative as well as positive frequencies — for the magnitude of the FFT, one is the mirror image of the other). This value was chosen experimentally and it provides good auditory event boundary detection. This parameter value may be changed to reduce (increase the threshold) or increase (decrease the threshold) the detection of events.

The details of this practical embodiment are not critical. Other ways to calculate the spectral content of successive time segments of the audio signal, calculate the differences between successive time segments, and set auditory event boundaries at the respective boundaries between successive time segments when the difference in the spectral profile content between such successive time segments exceeds a threshold may be employed.

For an audio signal consisting of Q blocks (of size M samples), the output of the auditory scene analysis process of function 2 of FIG. 1A is an array $B(q)$ of information representing the location of auditory event boundaries where $q = 0, 1, \dots, Q-1$. For a block size of $M = 512$ samples, overlap of $P = 0$ samples and a signal-sampling rate of 44.1kHz, the auditory scene analysis function 2 outputs approximately 86 values a second. Preferably, the array $B(q)$ is stored as the signature, such that, in its basic form, without the optional dominant subband frequency information, the audio signal's signature is an array $B(q)$ representing a string of auditory event boundaries.

An example of the results of auditory scene analysis for two different signals is shown in FIGS. 4A and 4B. The top plot, FIG. 4A, shows the results of auditory scene processing where auditory event boundaries have been identified at samples 1024 and 1536. The bottom plot, FIG. 4B, shows the identification of event boundaries at samples 1024, 2048 and 3072.

Identify dominant subband (optional)

For each block, an optional additional step in the ASA processing (shown in FIG. 2) is to extract information from the audio signal denoting the dominant frequency "subband" of the block (conversion of the data in each block to the

frequency domain results in information divided into frequency subbands). This block-based information may be converted to auditory-event based information, so that the dominant frequency subband is identified for every auditory event. This information for every auditory event provides the correlation processing (described below) with further information in addition to the auditory event boundary information. The dominant (largest amplitude) subband may be chosen from a plurality of subbands, three or four, for example, that are within the range or band of frequencies where the human ear is most sensitive. Alternatively, other criteria may be used to select the subbands.

The spectrum may be divided, for example, into three subbands. The preferred frequency range of the subbands is:

Subband 1	301Hz to 560Hz
Subband 2	560Hz to 1938Hz
Subband 3	1938Hz to 9948Hz

To determine the dominant subband, the square of the magnitude spectrum (or the power magnitude spectrum) is summed for each subband. This resulting sum for each subband is calculated and the largest is chosen. The subbands may also be weighted prior to selecting the largest. The weighting may take the form of dividing the sum for each subband by the number of spectral values in the subband, or alternatively may take the form of an addition or multiplication to emphasize the importance of a band over another. This can be useful where some subbands have more energy on average than other subbands but are less perceptually important.

Considering an audio signal consisting of Q blocks, the output of the dominant subband processing is an array $DS(q)$ of information representing the dominant subband in each block ($q = 0, 1, \dots, Q-1$). Preferably, the array $DS(q)$ is stored in the signature along with the array $B(q)$. Thus, with the optional dominant subband information, the audio signal's signature is two arrays $B(q)$ and $DS(q)$, representing, respectively, a string of auditory event boundaries and a dominant frequency subband within each block. Thus, in an idealized example, the two arrays could have the following values (for a case in which there are three possible dominant subbands).

- 16 -

1 0 1 0 0 0 1 0 0 1 0 0 0 0 0 1 0 (Event Boundaries)
 1 1 2 2 2 2 1 1 1 3 3 3 3 3 3 1 1 (Dominant Subbands)

In most cases, the dominant subband remains the same within each auditory event, as shown in this example, or has an average value if it is not uniform for all blocks within the event. Thus, a dominant subband may be determined for each auditory event and the array $DS(q)$ may be modified to provide that the same dominant subband is assigned to each block within an event.

Time Offset Calculation

The output of the Signature Extraction (FIG. 1A) is one or more arrays of auditory scene analysis information that are stored as a signature, as described above. The Time Offset Calculation function (FIG. 1B) takes two signatures and calculates a measure of their time offset. This is performed using known cross correlation methods.

Let S_1 (length Q_1) be an array from Signature 1 and S_2 (length Q_2) an array from Signature 2. First, calculate the cross-correlation array $R_{E_1E_2}$ (see, for example, John G. Proakis, Dimitris G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, Macmillan Publishing Company, 1992, ISBN 0-02-396815-X).

$$R_{E_1E_2}(l) = \sum_{q=-\infty}^{\infty} S_1(q) \cdot S_2(q-l) \quad l = 0, \pm 1, \pm 2, \dots \quad (1)$$

In a practical embodiment, the cross-correlation is performed using standard FFT based techniques to reduce execution time.

- 17 -

Since both S_1 and S_2 are finite in length, the non-zero component of $R_{E_1E_2}$ has a length of $Q_1 + Q_2 - 1$. The lag l corresponding to the maximum element in $R_{E_1E_2}$ represents the time offset of S_2 relative to S_1 .

$$l_{peak} = l \text{ for } MAX(R_{E_1E_2}(l)) \quad (2)$$

This offset has the same units as the signature arrays S_1 and S_2 . In a practical implementation, the elements of S_1 and S_2 have an update rate equivalent to the audio block size used to generate the arrays minus the overlap of adjacent blocks: that is, $M - P = 512 - 0 = 512$ samples. Therefore the offset has units of 512 audio samples.

Time Alignment

The Time Alignment function 6 (FIG. 1B) uses the calculated offset to time align the two audio signals. It takes as inputs, Audio Signals 1 and 2 (used to generate the two signatures) and offsets one in relation to the other such that they are both more closely aligned in time. The two aligned signals are output as Audio Signals 3 and 4. The amount of delay or offset applied is the product of the relative signature delay l_{peak} between signature S_2 and S_1 , and the resolution $M-P$, in samples, of the signatures.

For applications where only the passage common to the two sources is of interest (as in the case of watermark detection where unmarked and marked signals are to be directly compared), the two sources may be truncated to retain only that common passage.

For applications where no information is to be lost, one signal may be offset by the insertion of leading samples. For example let $x_1(n)$ be the samples of Audio Signal 1 with a length of N_1 samples, and $x_2(n)$ be the samples of Audio Signal 2 with a length of N_2 samples. Also l_{peak} represents the offset of S_2 relative to S_1 in units of $M-P$ audio samples.

The sample offset D_{21} of Audio Signal 2 relative to Audio Signal 1 is the product of the signature offset l_{peak} and $M-P$.

$$D_{21} = l_{peak} \cdot (M - P) \quad (3)$$

5

If D_{21} is zero, the both input signals are output unmodified as signals 3 and 4 (see FIG. 1B). If D_{21} is positive then input signal $x_1(n)$ is modified by inserting leading samples.

$$x'_1(m) = \begin{cases} 0 & 0 \leq m < D_{21} \\ x_1(n) & 0 \leq n < L_1 \quad m = n + D_{21} \end{cases} \quad (4)$$

10

Signals $x'_1(n)$ and $x_2(n)$ are output as Signals 3 and 4 (see FIG. 1B).

If D_{21} is negative then input signal $x_2(n)$ is modified by inserting leading samples.

$$x'_2(m) = \begin{cases} 0 & 0 \leq m < -D_{21} \\ x_2(n) & 0 \leq n < L_2 \quad m = n - D_{21} \end{cases} \quad (5)$$

15

Computation Complexity and Accuracy

The computational power required to calculate the offset is proportional to the lengths of the signature arrays, Q_1 and Q_2 . Because the process described has some offset error, the time alignment process of the present invention may be followed by a conventional process having a finer resolution that works directly with the audio signals, rather than signatures. For example such a process may take sections of the aligned audio signals (slightly longer than the offset error to ensure some overlap) and cross correlate the sections directly to determine the exact sample error or fine offset.

25

Since the signature arrays are used to calculate the sample offset, the accuracy of the time alignment method is limited to the audio block size used to generate the signatures: in this implementation, 512 samples. In other words this method will

- 19 -

have error in the sample offset of approximately plus/minus half the block size: in this implementation ± 256 samples.

This error can be reduced by increasing the resolution of the signatures; however there exists a tradeoff between accuracy and computational complexity.

- 5 Lower offset error requires finer resolution in the signature arrays (more array elements) and this requires higher processing power in computing the cross correlation. Higher offset error requires coarser resolution in the signature arrays (less array elements) and this requires lower processing power in computing the cross correlation.

10 *Applications*

Watermarking involves embedding information in a signal by altering the signal in some predefined way, including the addition of other signals, to create a marked signal. The detection or extraction of embedded information often relies on a comparison of the marked signal with the original source. Also the marked signal
15 often undergoes other processing including audio coding and speaker/microphone acoustic path transmission. The present invention provides a way of time aligning a marked signal with the original source to then facilitate the extraction of embedded information.

Subjective and objective methods for determining audio coder quality compare
20 a coded signal with the original source, used to generate the coded signal, in order to create a measure of the signal degradation (for example an ITU-R 5 point impairment score). The comparison relies on time alignment of the coded audio signal with the original source signal. This method provides a means of time aligning the source and coded signals.

25 Other applications of the invention are possible, for example, improving the lip-syncing of audio and video signals, as mentioned above.

It should be understood that implementation of other variations and modifications of the invention and its various aspects will be apparent to those skilled in the art, and that the invention is not limited by these specific embodiments
30 described. It is therefore contemplated to cover by the present invention any and all

modifications, variations, or equivalents that fall within the true spirit and scope of the basic underlying principles disclosed and claimed herein.

The present invention and its various aspects may be implemented as software functions performed in digital signal processors, programmed general-purpose digital computers, and/or special purpose digital computers. Interfaces between analog and digital signal streams may be performed in appropriate hardware and/or as functions in software and/or firmware.

CLAIMS

1. A method for time aligning audio signals, wherein one signal has been derived from the other or both have been derived from another signal, comprising
5 deriving reduced-information characterizations of said audio signals, wherein said reduced-information characterizations are based on auditory scene analysis, calculating the time offset of one characterization with respect to the other characterization,
 modifying the temporal relationship of said audio signals with respect to each
10 other in response to said time offset such that said audio signals are substantially coincident with each other.
2. The method of claim 1 wherein said reduced-information characterizations are derived from said audio signals and embedded in respective other signals that are
15 carried with the audio signals from which they were derived prior to said calculating and modifying.
3. The method of claim 2 wherein said other signals are the video portion of a television signal and said audio signal are the audio portion of the respective
20 television signal.
4. A method for time aligning an audio signal and another signal, comprising
 deriving a reduced-information characterization of the audio signal and
 embedding said characterization in the other signal when the audio signal and other
25 signal are substantially in synchronism, wherein said characterization is based on auditory scene analysis,
 recovering the embedded characterization of said audio signal from said other signal and deriving a reduced-information characterization of said audio signal from said audio signal in the same way the embedded characterization of the audio signal

was derived based on auditory scene analysis, after said audio signal and said other signal have been subjected to differential time offsets,

calculating the time offset of one characterization with respect to the other characterization,

5 modifying the temporal relationship of the audio signal with respect to the other signal in response to said time offset such that the audio signal and video signal are substantially in synchronism with each other.

5. The method of claim 4 wherein said other signal is a video signal.

10

6. The method of claim 1 or claim 4 wherein calculating the time offset includes performing a cross-correlation of said characterizations.

7. The method of any one of claims 1-6 wherein said reduced-information
15 characterizations based on auditory scene analysis are arrays of information representing at least the location of auditory event boundaries.

8. The method of claim 7 wherein said auditory event boundaries are determined by

20 calculating the spectral content of successive time segments of said audio signal,



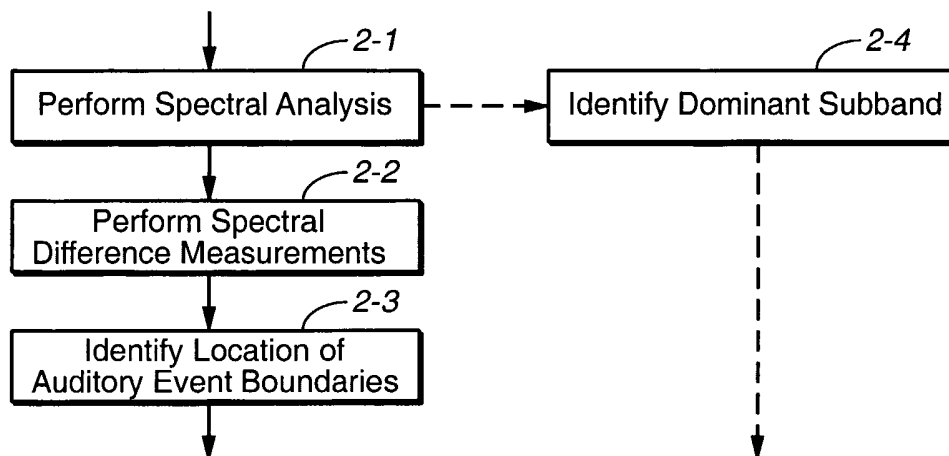
calculating the difference in spectral content between successive time segments of said audio signal, and




25 identifying an auditory event boundary as the boundary between successive time segments when the difference in the spectral content between such successive time segments exceeds a threshold.

9. The method of claim 7 or claim 8 wherein said arrays of information also represent the dominant frequency subband of each of said auditory events.

30

1 / 2

**FIG._ 1A****FIG._ 1B****FIG._ 2**



INTERNATIONAL SEARCH REPORT

Int. l. Application No

PCT/US 02/05806

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 G10L11/00 H04N5/04

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G10L H04N

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 practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 040 081 A (MCCUTCHEN DAVID) 13 August 1991 (1991-08-13) column 3, line 4 - line 21 column 3, line 51 - line 64 ---	1, 4, 6
A	WO 91 19989 A (REYNOLDS SOFTWARE INC) 26 December 1991 (1991-12-26) the whole document ---	1, 4, 7, 8
A	US 6 211 919 B1 (BAKER DANIEL G ET AL) 3 April 2001 (2001-04-03) cited in the application column 1, line 28 - line 44 column 8, line 10 - line 20 --- -/--	1-5

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in 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 document 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.
- *G* document member of the same patent family

Date of the actual completion of the international search

26 September 2002

Date of mailing of the international search report

07/10/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Ogor, M

INTERNATIONAL SEARCH REPORT

Int. Patent Application No.

PCT/US 02/05806

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 372 155 A (GRAVEREAUX DANIEL W ;KARAMON JOHN J (US)) 13 June 1990 (1990-06-13) column 1, line 1 - line 6 column 3, line 45 -column 4, line 7 column 12, line 36 -column 13, line 3 -----	1,6
A	US 4 624 009 A (GLENN JAMES W ET AL) 18 November 1986 (1986-11-18) -----	

Information on patent family members

International Application No
PCT/US 02/05806

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5040081	A	13-08-1991	NONE	
WO 9119989	A	26-12-1991	CA 2085887 A1 DE 4191297 T GB 2262992 A ,B GB 2282456 A ,B JP 5509409 T WO 9119989 A1 US 5400261 A US 5276629 A	22-12-1991 15-07-1993 07-07-1993 05-04-1995 22-12-1993 26-12-1991 21-03-1995 04-01-1994
US 6211919	B1	03-04-2001	EP 0868082 A2 JP 10304322 A US 6246439 B1	30-09-1998 13-11-1998 12-06-2001
EP 0372155	A	13-06-1990	US 5055939 A AT 151539 T CA 1307745 A1 DE 68927950 D1 EP 0372155 A2	08-10-1991 15-04-1997 22-09-1992 15-05-1997 13-06-1990
US 4624009	A	18-11-1986	NONE	