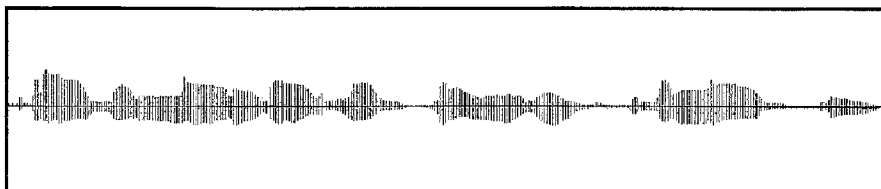




(86) **Date de dépôt PCT/PCT Filing Date:** 2013/03/04
 (87) **Date publication PCT/PCT Publication Date:** 2013/10/03
 (45) **Date de délivrance/Issue Date:** 2020/06/09
 (85) **Entrée phase nationale/National Entry:** 2014/09/04
 (86) **N° demande PCT/PCT Application No.:** US 2013/028788
 (87) **N° publication PCT/PCT Publication No.:** 2013/148069
 (30) **Priorité/Priority:** 2012/03/28 (US13/432,838)

(51) **Cl.Int./Int.Cl. G10L 15/00** (2013.01),
G10L 19/06 (2013.01), **H04M 3/527** (2006.01)
 (72) **Inventeurs/Inventors:**
 VLACK, KEVIN, US;
 WYSS, FELIX IMMANUEL, US
 (73) **Propriétaire/Owner:**
 INTERACTIVE INTELLIGENCE, INC., US
 (74) **Agent:** BROUILLETTE LEGAL INC.

(54) **Titre : SYSTEME ET PROCEDURE DE PRISE D'EMPREINTE D'ENSEMBLES DE DONNEES**
 (54) **Title: SYSTEM AND METHOD FOR FINGERPRINTING DATASETS**



(57) **Abrégé/Abstract:**

Systems and methods for the matching of datasets, such as input audio segments, with known datasets in a database are disclosed. In an illustrative embodiment, the use of the presently disclosed systems and methods is described in conjunction with recognizing known network message recordings encountered during an outbound telephone call. The methodologies include creation of a ternary fingerprint bitmap to make the comparison process more efficient. Also disclosed are automated methodologies for creating the database of known datasets from a larger collection of datasets.

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

(19) World Intellectual Property
Organization
International Bureau(10) International Publication Number
WO 2013/148069 A1(43) International Publication Date
3 October 2013 (03.10.2013)(51) International Patent Classification:
H04M 3/00 (2006.01)

(21) International Application Number:

PCT/US2013/028788

(22) International Filing Date:

4 March 2013 (04.03.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/432,838 28 March 2012 (28.03.2012) US


(71) Applicant: **INTERACTIVE INTELLIGENCE, INC.**
[US/US]; 7601 Interactive Way, Indianapolis, IN 46278
(US).(72) Inventors: **VLACK, Kevin**; 2077 Boulder Drive, Apt. H-
5, Indianapolis, IN 46260 (US). **WYSS, Felix, Immanuel**;
722 Morningside Drive, Zionsville, IN 46077 (US).(74) Agents: **COLE, Troy, J.** et al.; Ice Miller, LLP, One
American Square, Suite 2900, Indianapolis, IN 46282-
0200 (US).(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, 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, PA, 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, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
ML, MR, NE, SN, TD, TG).**Declarations under Rule 4.17:**— as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))**Published:**

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

— before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments (Rule 48.2(h))

(54) Title: SYSTEM AND METHOD FOR FINGERPRINTING DATASETS

Fig. 1A



(57) Abstract: Systems and methods for the matching of datasets, such as input audio segments, with known datasets in a database are disclosed. In an illustrative embodiment, the use of the presently disclosed systems and methods is described in conjunction with recognizing known network message recordings encountered during an outbound telephone call. The methodologies include creation of a ternary fingerprint bitmap to make the comparison process more efficient. Also disclosed are automated methodologies for creating the database of known datasets from a larger collection of datasets.



WO 2013/148069 A1

WO 2013/148069

PCT/US2013/028788

SYSTEM AND METHOD FOR FINGERPRINTING DATASETS

TECHNICAL FIELD OF THE DISCLOSED EMBODIMENTS

[1] The present invention generally relates to identifying known datasets, such as audio signals and, more particularly, to systems and methods for fingerprinting datasets.

BACKGROUND OF THE DISCLOSED EMBODIMENTS

[2] The background of the present disclosure and the illustrative embodiments disclosed herein are described in the context of identifying known audio recordings encountered during an outbound telephone call, for example during a call placed from a contact center. However, the present invention has applicability to the identification of any segment of audio or an image (as used herein, the term “image” is intended to encompass both still and moving images), regardless of the type or source of the audio or image, and regardless of in what circumstances the audio or image is encountered. Furthermore, the present invention also has applicability to the identification of any segment of data such as, for example, data obtained from any type of sensor. Therefore, as used herein, the term “dataset” shall encompass a collection of any type of data, whether comprising audio, image, or other type of data.

[3] In a classic contact center scenario, outbound calls are made either automatically (by a class of devices known as “automated dialers” or “autodialers”) or manually. A number of human “agents” are available to join into calls that are determined to reach a live person at the called end. In this way, efficiencies are obtained by not having agents involved in a call until it is determined that there is a live person at the called end with whom the agent may speak. The use of automated equipment to monitor the telephone line during the outbound call is referred to as call progress analysis (CPA). CPA is a class of algorithms that operate on audio and network

WO 2013/148069

PCT/US2013/028788

signaling during call setup. The goal of CPA is to determine the nature of the callee, or the outcome of call setup to an external network (traditional public switched telephone network or Voice over Internet Protocol (VoIP)). Specifically, when a call or session is being established, the caller or initiator must determine whether it was answered by a live speaker, if the line is busy, etc. When the caller is an automated application, such as an automated dialer or message broadcasting system, CPA algorithms are used to perform the classification automatically. CPA is used to interpret so-called call-progress tones, such as ring back and busy, that are delivered by the telephone network to the calling entity. Traditional CPA is performed using low- and high-pass frequency discriminators together with energy measurements over time to qualify in-band signaling tones.

[4] Another method for classifying audio on an outbound call is known as Voice Activity Detection (VAD), which is a class of audio processing algorithms that identify where speech is present in an audio stream. The detected speech may originate from any source, including a live speaker or a prerecorded message. Modern VAD algorithms use spectral analysis to distinguish the utterance of a primary speaker from background noise.

[5] A subclass of CPA algorithms that extract speaking patterns using VAD, and determine whether the patterns originate from a live speaker or a prerecorded message, is known as Answering Machine Detection (AMD). By identifying calls that do not connect to a live speaker, an accurate AMD algorithm can significantly increase throughput of an automated dialer. However, false positives from AMD lead to silent or abandoned calls, causing revenue loss for the contact center, and negative impressions amongst the public. The quality of an AMD algorithm is a function of the accuracy and response time, and some regions of the world (notably the U.S. and U.K.) impose strict legal requirements on both.

WO 2013/148069

PCT/US2013/028788

[6] AMD is not an exact science, and the optimal approach is an open problem. To achieve acceptable accuracy, speed, and flexibility, AMD algorithms use a combination of heuristics and statistical models such as neural networks to classify an utterance as live or pre-recorded.

Although many commercial AMD systems available on the market report high accuracy rates in the marketing literature (e.g., 95% or more), there is no independent auditor for these figures, and the actual accuracy rate is typically much lower in practice (e.g., 80% or less), as reflected by continued widespread complaints. A general ban has been proposed by some consumer advocacy groups, and some contact centers simply cannot use AMD because of its limitations.

[7] A relatively new science of audio identification is known as Acoustic Fingerprinting, in which a system generates a “fingerprint” of a candidate audio stream, and compares it against a database of known fingerprints, analogous to human fingerprinting used in forensics. In this context, a “fingerprint” is a condensed digest of an audio stream that can quickly establish perceptual equality with other audio streams. A database of known fingerprints may associate known fingerprints with meta-data such as “title”, “artist”, etc. The past ten years have seen a rapidly growing scientific and industrial interest in fingerprinting technology for audio and images. Applications include identifying songs and advertisements, media library management, and copyright compliance.

[8] Various acoustic fingerprinting algorithm classes have been proposed, and the most prevalent today are those based on either “landmarks” or “bitmaps”. Landmark-based algorithms extract discrete features from an audio stream called “landmarks”, such as spectral peaks, sudden changes in tone, pitch, loudness, etc. The optimal choice of landmark is an open question guided mostly by heuristics. The acoustic fingerprint is stored as a sequence of data

structures that describe each landmark. At runtime, landmarks extracted from a candidate audio stream are compared to a database of fingerprints based on a distance metric.

[9] Bitmap-based algorithms analyze an audio stream as a sequence of frames, and use a filter bank to quantize each frame into a bit vector of size N , where N is typically
5 chosen for convenience as the number of bits in a C-style integer, e.g., $N \in \{8, 16, 32, \text{ or } 64\}$. A popular and well-studied example is known as the "Haitsma-Kalker algorithm", which computes a binary bitmap using a filter that compares short-term differences in both time and frequency. The Haitsma-Kalker Algorithm has been well-studied in the literature. Its inventors, Jaap Haitsma and Ton Kalker, have published a report of use of
10 the Haitsma-Kalker Algorithm and the comparison of binary acoustic fingerprint bitmaps to identify three (3) second recordings of songs from a database of millions of songs (Haitsma and Kalker, "A Highly Robust Audio Fingerprinting System," *Journal of New Music Research*, Vol. 32, No. 2 (2003), pp. 211-221). The complete acoustic fingerprint is stored as a sequence of bit vectors, or a bitmap. As illustrated in FIG. 1A-C (Prior
15 Art), there are shown three images of an audio stream containing a message from a telephone network saying "This number has been disconnected". FIG 1A (Prior Art) shows the original audio wave signal, with 1.5 seconds of audio sampled at 8000 KHz. FIG 1B (Prior Art) shows a spectrogram of the original audio input signal, with dark regions indicating high energy at a particular frequency. FIG. 1C (Prior Art) shows a
20 binary acoustic fingerprint bitmap created using the Haitsma-Kalker algorithm, with height $N=16$. The height is determined by the number of bits computed at each frame, and the width is determined by the number of frames in the audio stream. At runtime, the bitmap computed from a candidate audio stream is compared to a database of bitmaps based on the number of non-matching bits, also known as the Hamming distance.

25

WO 2013/148069

PCT/US2013/028788

[10] The use of bitmap matching and the process of acoustic fingerprinting is a powerful emerging tool in the science of audio recognition; however, it is computationally intense and requires several seconds of sampled audio to make a match in many cases. This delay makes it not well suited for use in call progress analysis. Accordingly, there remains a need for faster and more accurate systems and methods for identifying audio, both in the general case and during an outbound call attempt.

SUMMARY OF THE DISCLOSED EMBODIMENTS

[11] Systems and methods for the matching of datasets, such as input audio segments, with known datasets in a database are disclosed. In an illustrative embodiment, the use of the presently disclosed systems and methods is described in conjunction with recognizing known network message recordings encountered during an outbound telephone call. The methodologies include creation of a ternary fingerprint bitmap to make the comparison process more efficient. Also disclosed are automated methodologies for creating the database of known datasets from a larger collection of datasets.

[12] In one embodiment, a method for identifying a candidate audio segment from an outbound telephone call is disclosed, the method comprising the steps of: a) creating a spectrogram of the candidate audio segment; b) creating a candidate binary acoustic fingerprint bitmap of the spectrogram; c) comparing the candidate binary acoustic fingerprint bitmap to at least one known binary acoustic fingerprint bitmap of a known network message; d) if the candidate binary acoustic fingerprint bitmap matches one of said at least one known binary acoustic fingerprint bitmaps within a predetermined threshold, declaring the match; and e) if the candidate binary acoustic fingerprint bitmap does not match one of said at least one known

binary acoustic fingerprint bitmaps within the predetermined threshold, using an answering machine detection algorithm to analyze the candidate audio segment.

[13] In another embodiment, a method for identifying a candidate audio segment from an outbound telephone call is disclosed, the method comprising the steps of: a) creating a spectrogram of the candidate audio segment; b) creating a candidate binary fingerprint bitmap of the spectrogram; c) comparing the candidate binary fingerprint bitmap to at least one known binary fingerprint bitmap of a known recording; d) if the candidate binary fingerprint bitmap matches one of said at least one known binary fingerprint bitmaps within a predetermined threshold, declaring the match; and e) if the candidate binary fingerprint bitmap does not match one of said at least one known binary fingerprint bitmaps within the predetermined threshold, using an alternate process to analyze the candidate audio segment.

[14] In a further embodiment, a method for creating a ternary bitmap of a dataset is disclosed, the method comprising the steps of: a) computing a binary fingerprint bitmap of the dataset; b) deleting a first number of samples from the dataset; c) after step (b), computing another binary fingerprint bitmap of the dataset; d) repeating steps (b) and (c) a plurality of times to create a plurality of binary fingerprint bitmaps; and e) combining the plurality of binary fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows: e.1) If a bit is 0 (zero) in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0 (zero); e.2) If a bit is 1 (one) in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1 (one); and e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care").

[15] In yet another embodiment, a method for identifying a candidate dataset is disclosed, the method comprising the steps of: a) computing a binary fingerprint bitmap of a known dataset in a

WO 2013/148069

PCT/US2013/028788

known dataset database; b) deleting a first number of samples from the known dataset; c) after step (b), computing another binary fingerprint bitmap of the known dataset; d) repeating steps (b) and (c) a plurality of times to create a plurality of binary fingerprint bitmaps; and e) combining the plurality of binary fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows: e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0; e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1; and e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care"); f) saving the ternary bitmap into a ternary bitmap database; g) repeating steps (a) – (f) for all known datasets in the known dataset database; h) creating a candidate dataset binary fingerprint bitmap from the candidate dataset; and i) comparing the candidate dataset binary fingerprint bitmap to each ternary bitmap in the ternary bitmap database, wherein said comparison ignores the Don't Care bits.

[16] In a further embodiment, a method for creating a ternary bitmap of an audio segment is disclosed, the method comprising the steps of: a) computing a binary acoustic fingerprint bitmap of the audio segment; b) deleting a first number of samples from the audio segment; c) after step (b), computing another binary acoustic fingerprint bitmap of the audio segment; d) repeating steps (b) and (c) a plurality of times to create a plurality of binary acoustic fingerprint bitmaps; and e) combining the plurality of binary acoustic fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows: e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0; e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit

WO 2013/148069

PCT/US2013/028788

in the ternary bitmap to 1; and e.3) Otherwise, set the bit of the ternary bitmap to * (“Don’t Care”).

[17] In still another embodiment, a method for identifying a candidate audio segment is disclosed, the method comprising the steps of: a) computing a binary acoustic fingerprint bitmap of a known audio segment in a known audio segment database; b) deleting a first number of samples from the known audio segment; c) after step (b), computing another binary acoustic fingerprint bitmap of the known audio segment; d) repeating steps (b) and (c) a plurality of times to create a plurality of binary acoustic fingerprint bitmaps; and e) combining the plurality of binary acoustic fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows: e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0; e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1; and e.3) Otherwise, set the bit of the ternary bitmap to * (“Don’t Care”); f) saving the ternary bitmap into a ternary bitmap database; g) repeating steps (a) – (f) for all known audio segments in the known audio segment database; h) creating a candidate audio segment binary acoustic fingerprinting bitmap from the candidate audio segment; and i) comparing the candidate audio segment binary acoustic fingerprint bitmap to each ternary bitmap in the ternary bitmap database, wherein said comparison ignores the Don’t Care bits.

[18] In yet another embodiment, a method for creating weighted compressed representation of a dataset is disclosed, the method comprising the steps of a) computing a compressed representation of the dataset; b) applying a transformation to the dataset; c) after step (b), computing another compressed representation of the dataset; d) repeating steps (b) and (c) a plurality of times to create a plurality of compressed representations; and e) combining the

WO 2013/148069

PCT/US2013/028788

plurality of compressed representations into a weighted compressed representation, where each weight in the weighted compressed representation is computed as a function of the elements in the plurality of compressed representations.

[19] Other embodiments are also disclosed.

BRIEF DESCRIPTION OF THE DRAWINGS

[20] **FIGs. 1A-C** are an audio wave, spectrogram and binary acoustic fingerprint bitmap, respectively.

[21] **FIG. 2** is a schematic flow diagram of one embodiment of a method for building an acoustic fingerprint.

[22] **FIG. 3A** is a graph of amplitude vs. time for a set of exemplary sliding windows used in a transformation according to one embodiment.

[23] **FIG. 3B** is a spectrogram of an audio sample.

[24] **FIG. 4** is a schematic representation of the Mel scale;

[25] **FIG. 5** is a graphical representation of the Haitsma-Kalkor Algorithm;

[26] **FIG. 6** is a schematic flow diagram illustrating a method for creating ternary bitmaps according to one embodiment;

[27] **FIG. 7** is a schematic graphical representation of the creation of a ternary bitmap from a binary bitmap according to one embodiment;

[28] **FIG. 8** is a schematic flow diagram of the creation of a hash key from a ternary bitmap according to one embodiment;

[29] **FIG. 9** is a graphical representation of the process of FIG. 7;

[30] **FIG. 10** is a schematic flow diagram of a matching procedure using hash keys according to one embodiment; and

[31] FIG. 11 is a schematic flow diagram of a method for extracting fingerprints for common recordings from large audio datasets according to one embodiment.

DETAILED DESCRIPTION OF THE DISCLOSED EMBODIMENTS

[32] For the purposes of promoting an understanding of the principles of the invention, reference will now be made to the embodiment illustrated in the drawings, and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope of the invention is intended. Alterations and modifications in the illustrated systems and methods, and further applications of the principles of the invention as illustrated therein, as would normally occur to one skilled in the art to which the invention relates are contemplated, are desired to be protected. Such alternative embodiments require certain adaptations to the embodiments discussed herein that would be obvious to those skilled in the art.

[33] Although the various embodiments disclosed herein will be described in the context of identifying known recordings encountered during an outbound telephone call, for example during a call placed from a contact center, the present invention has applicability to the identification of any segment of audio, image, or other type of data, regardless of the type or source of the audio, image, or other type of data, and regardless of in what circumstances the audio, image, or other type of data is encountered. Additionally, the present invention has applicability to the recognition of any type of dataset having two or more dimensions. The predominant reliance on the recognition of audio recordings herein is simply for convenience of description.

[34] The use of acoustic fingerprinting in a Call Progress Analysis system.

WO 2013/148069

PCT/US2013/028788

[35] Call Progress Analysis may be improved by using a system that augments conventional AMD with an acoustic fingerprinting system to identify specific call progress events of interest including, but not limited to:

[36] a. Telephone network messages (e.g., “We’re sorry”; “The number or code you have dialed is incorrect”)

[37] b. Voicemail greetings shared by multiple subscribers (e.g., “Hello, the party you have dialed”)

[38] c. Colored tones, jingles, chimes

[39] d. Ringback tones containing music or speech (e.g., “Please hold while your call is completed”)

[40] e. Call screeners

[41] f. Privacy managers (e.g., “You have reached a number that does not accept solicitations”)

[42] g. Interactive Voice Response (IVR) prompts

[43] These are collectively referred to herein as “network messages.” All such network messages share the common trait that the same audio is used by the network in certain situations for calls placed to multiple called parties. Therefore, a dialing program should expect to encounter these network messages in the future, and identifying them as such will help the call progress analysis software determine that a live speaker has not been reached. It is desired that the call progress analysis software exhibit the following characteristics:

[44] 1. Efficient (low use of CPU and memory)

[45] 2. Fast (low latency and delay)

[46] 3. Robust (low rate of false negatives)

WO 2013/148069

PCT/US2013/028788

[47] 4. Accurate (low rate of false positives)

[48] 5. Scalable (low lookup cost)

[49] 6. Repeatable (low learning curve)

[50] In one embodiment, during the real-time progress of an outbound call, the CPA system processes the audio stream using both a conventional AMD algorithm and an acoustic fingerprinting system. As the audio proceeds, the acoustic fingerprinting system identifies whether there is a likely match in the database of acoustic fingerprints of known network messages. If so, any output from the AMD algorithm is preempted: instead, the CPA system reports specific call progress events based on the matched acoustic fingerprint. For the general case where the acoustic fingerprint system finds no matches in the database, the conventional AMD algorithm is used to detect a greeting, and report the event of either a live speaker or an answering machine. All events are interpreted by the user of the CPA system to choose an appropriate response, such as whether to proceed or terminate the call, dispatch to an agent, adjust the dialer plan, etc. In other embodiments, the acoustic fingerprinting system is used without an additional CPA mechanism.

[51] FIG. 2 schematically illustrates one embodiment of a method for building an acoustic fingerprint. In order to build an acoustic fingerprint, either of a known audio segment or of a candidate audio segment from a current telephone call, a time varying spectral representation of the signal (referred to herein as a "spectrogram") is created (as indicated at step 100) using any desired transform, such as Discrete Fourier Transform (DFT), Discrete Cosine Transform (DCT), wavelet transform, or even just a set of filter banks, to name just four non-limiting examples. At each frame, the power spectrum in frequency space is computed. For example, FIG. 3A illustrates amplitude vs. time for the sliding windows used in a Fast Fourier Transform (FFT), a

WO 2013/148069

PCT/US2013/028788

particularly efficient algorithm to implement Discrete Fourier Transform) of an audio sample using the following parameters:

[52] Sample rate: 8 kHz

[53] Frame duration: 30 milliseconds

[54] Step size: 20 milliseconds

[55] Overlap: 1/3

[56] FFT size: 256

[57] A sample spectrogram of an audio sample using this technique is illustrated in FIG. 3B.

[58] The next step is to create a binary acoustic fingerprint bitmap, as indicated at step 102.

The power spectrogram may be used as is, or it may be optionally transformed in order to condense the data. Any transformation technique that will condense the data may be used. To give just one non-limiting example, the Haitsma-Kalker Algorithm may be used, where a sequence of frames are created and are combined to build the acoustic fingerprint bitmap. A filter bank can be used to quantize each frame into a bit vector of size N, where N may be chosen for convenience as the number of bits in a C-style integer (8, 16, 32, or 64). In one embodiment, a Mel-scale filter bank is used to transform the power spectrum data into Mel space with (N+1) bands, as indicated at step 104. The Mel scale is a perceptual scale of pitches judged by listeners to be equal in distance from one another, as illustrated in FIG. 4.

[59] From the sequence of (N+1)-band spectrums in Mel space, a sequence of N-bit binary fingerprint frame values based on band energy differences over successive frames is computed at step 106. In one embodiment, this is done using the Haitsma-Kalker Algorithm as follows:

$$[60] \quad F(n, m) = \begin{cases} 1 & \text{if } E(n, m) - E(n, m + 1) - (E(n - 1, m) - E(n - 1, m + 1)) > 0 \\ 0 & \text{if } E(n, m) - E(n, m + 1) - (E(n - 1, m) - E(n - 1, m + 1)) \leq 0 \end{cases}$$

[61] where: $E(n, m)$ is the energy of frequency band m of frame n, and

[62] $F(n, m)$ is the m th bit of fingerprint frame n

[63] This is illustrated schematically in FIG. 5. Thus, the acoustic fingerprint is built up as a vector of N -bit integers, one integer per frame. FIG. 1C illustrates an exemplary 16-bit binary acoustic fingerprint bitmap.

[64] For use during call progress analysis, a database is maintained that contains binary acoustic fingerprint bitmaps for known network messages. In one embodiment, during the real-time progress of an outbound call, the CPA system processes the audio stream using both a conventional AMD algorithm and an acoustic fingerprinting system. As the audio proceeds, the acoustic fingerprinting system creates binary acoustic fingerprint bitmaps of the incoming audio and compares those bitmaps to known network message bitmaps stored in a database (step 108). The CPA system identifies whether there is a match in the database of binary bitmaps (step 110), indicating that the outbound call has resulted in receipt of a known network message instead of a live speaker (step 112). If so, any output from the AMD algorithm is preempted: instead, the CPA system reports specific call progress events based on the matched acoustic fingerprint. For the general case where the acoustic fingerprint system finds no matches in the database, the conventional AMD algorithm is used to detect a greeting, and report the event of either a live speaker or an answering machine (step 114). All events are interpreted by the user of the CPA system to choose an appropriate response, such as whether to proceed or terminate the call, dispatch to an agent, adjust the dialer plan, etc. In other embodiments, the binary acoustic fingerprinting bitmap matching system is used without an additional CPA mechanism.

[65] The Haitsma-Kalker Algorithm has been well-studied in the literature. In Haitsma and Kalker's published report of use of the Haitsma-Kalker Algorithm and the comparison of binary acoustic fingerprint bitmaps to identify three (3) second recordings of songs from a database of

WO 2013/148069

PCT/US2013/028788

millions of songs (Haitsma and Kalker, "A Highly Robust Audio Fingerprinting System," *Journal of New Music Research*, Vol. 32, No. 2 (2003), pp. 211-221). Their system required a large frame size, large overlap/small step size, and large fingerprints in order to achieve good results. The parameters they used were:

- [66] Sample rate: 5 kHz
- [67] Frame duration: 370 milliseconds
- [68] Step size: 10 milliseconds
- [69] Overlap: 31/32
- [70] FFT size: 2048

[71] The Haitsma-Kalker algorithm computes a binary acoustic fingerprint bitmap with a relatively low Signal-to-Noise ratio, with bits highly sensitive to noise and windowing artifacts. To achieve an acoustic fingerprinting system with acceptably high accuracy and low false positive rate typically requires a relatively long segment of audio (~3 seconds) with large frame size (370 ms) and large overlap between frames (31/32, or about 97%).

[72] Such a system is computationally intensive and requires a relatively large audio sample to make reliable comparisons, both of which are undesirable in many audio matching scenarios. The present inventors have observed that many bits in a binary acoustic fingerprint bitmap are sensitive to noise and artifacts resulting from the transform into the spectral representation (windowing), especially where the energy of the signal is relatively small. A proposed solution is to mask out bits of little value due to their sensitivity to noise and windowing artifacts, and compute the error rate of the bitmap comparisons using only the bits that are in the mask (i.e., the bits that are not masked out).

WO 2013/148069

PCT/US2013/028788

[73] Ternary bitmap acoustic fingerprint using mask bits to obtain high Signal-to-Noise Ration (SNR)

[74] To improve the bitmap matching speed without sacrificing accuracy, one embodiment of the present invention makes the following modification to any acoustic fingerprinting algorithm that generates a binary bitmap. In the training stage where the database of known network message bitmaps is created, the methodology of FIG. 6 may be implemented. At step 200, the binary acoustic fingerprint bitmap is created for the audio segment using the same process described hereinabove with respect to FIG. 2. At step 202, in the illustrated embodiment the first sample is deleted from the audio segment (although other segment modification schemes may be employed, as discussed below), and at step 204 another binary acoustic fingerprint bitmap is created for the modified audio segment using the same process described hereinabove with respect to FIG. 2. Steps 202-204 are repeated X times (Step 206), where X is the number of samples in each frame used by the acoustic fingerprinting process of FIG. 2. This process will generate X binary acoustic fingerprint bitmaps of the audio segment.

[75] The X binary acoustic fingerprint bitmaps are combined into a ternary bitmap at step 208 as follows:

[76] • If a bit is 0 in all X binary bitmaps, set the bit in the ternary bitmap to 0

[77] • If a bit is 1 in all X binary bitmaps, set the bit in the ternary bitmap to 1

[78] • Otherwise, set the bit of the ternary bitmap to * ("Don't Care")

[79] The "Don't Care" bits change in the successive binary acoustic fingerprint bitmaps because they are sensitive to framing, noise, compression, and other effects that introduce signal distortion. It will be appreciated from the above description that the "Don't Care" bits may be defined in other ways, such as a predetermined percentage of bit values changing across the

WO 2013/148069

PCT/US2013/028788

bitmaps, etc. If the ternary bitmap is used for bitmap comparison (step 210) and the “Don’t Care” bits are ignored during the comparison process, the “Don’t Care” bits mask out these regions in frequency and time of the original binary acoustic fingerprint bitmap that introduce signal distortion. FIG. 7 illustrates an example comparison between a binary bitmap and its associated ternary bitmap. It will also be appreciated from the above description that the input signal can be artificially degraded before computing the acoustic fingerprint by adding noise or other artifacts commonly introduced by the communication channel in order to make the acoustic fingerprint and mask more sensitive.

[80] At runtime, these “Don’t Care” bits are excluded from the evaluated “masked Hamming distance” between the candidate audio stream and the database of known fingerprints. Use of the ternary bitmap mask eliminates false negatives due to framing misalignment between the input candidate audio stream and the recordings in the database. By excluding the bits most sensitive to windowing and noise, the ternary fingerprint system is more robust than its binary counterpart, and achieves comparable accuracy and false positive rate with much fewer bits. It will be appreciated that the above steps create many fingerprints of the same signal, with the analysis window shifted by a fraction of the frame size, and those fingerprints are then used to identify which parts of the fingerprint change. Therefore, in step 202 the analysis window may be shifted by any fraction of the frame size, rather than the illustrated single-sample shift. It will additionally be appreciated that the acoustic fingerprint and ternary bitmap (mask) may be represented in ways other than as vectors of vectors where each element is a bit with a value of zero or one, but instead more generally as vectors of vectors where each element comprises numeric scores and weights (and quantized into multi-bit representations).

WO 2013/148069

PCT/US2013/028788

[81] For example, a method for creating a weighted compressed representation of a dataset may comprise the steps of a) computing a compressed representation of the dataset; b) deleting a first number of elements from the dataset to create a modified dataset; c) computing another compressed representation of the modified dataset; d) repeating steps (b) and (c) a plurality of times to create a plurality of compressed representations; and e) combining the plurality of compressed representations into a weighted compressed representation, where each weight in the weighted compressed representation is computed as a function of the elements in the plurality of compressed representations.

[82] In some embodiments, the compressed representation is a two-dimensional vector of first numbers. In some embodiments, the first numbers are each represented with a first number of bits, where the first number of bits may be one bit in some embodiments, and greater than one bit in other embodiments. In other embodiments, the weighted compressed representation is a two-dimensional vector of second numbers. In some embodiments the second numbers are each represented with a second number of bits. In some embodiments, the function computes each said weight as correlation coefficient of the elements in the plurality of compressed representations. In some embodiments, the correlation coefficient is computed as follows:

[83] 1) if an element is below a first predetermined threshold in a first predetermined number of the plurality of compressed representations, set the correlation coefficient to a first value;

[84] 2) if an element is above a second predetermined threshold in a second predetermined number of the plurality of compressed representations, set the correlation coefficient to a second value;

[85] 3) otherwise set the correlation coefficient to a third value.

[86] For example, in one embodiment, the first value is +1, the second value is -1, and the third value is 0 (zero). Those skilled in the art will recognize from the present disclosure that the correlation coefficient can be any value between -1 and +1 (or any other desired number range). In the present embodiment, the values correspond to the ternary bitmap discussed herein, where a set bit is +1, a cleared bit is -1 and a Don't Care bit is 0. The thresholds can be set to any desired level. For example, the first predetermined threshold may be 1 (one) and the second predetermined threshold may be 0 (zero) in one embodiment. When computing the correlation coefficient in one embodiment, the first predetermined number comprises all of the plurality of compressed representations and the second predetermined number comprises all of the plurality of compressed representations.

[87] The methodology of FIG. 6 can be compared to the Haitsma-Kalker algorithm binary acoustic fingerprint bitmap comparison approach of FIG. 2, which computes a binary acoustic fingerprint bitmap with a relatively low Signal-to-Noise ratio, with bits highly sensitive to noise and windowing artifacts. To achieve a fingerprinting system with acceptably high accuracy and low false positive rate using this approach typically requires a relatively long segment of audio (~ 3 seconds) with large frame size (370 ms) and large overlap between frames (31/32, or about 97%). By comparison, the ternary acoustic fingerprint bitmap methodology disclosed herein, such as that illustrated in FIG. 6, can obtain comparable matching accuracy and precision much more efficiently, with a short segment of low-bandwidth audio (~0.25 seconds at 8 kHz), using a small frame size of 30 ms and a small frame overlap of 33%. Not only does this achieve matching in approximately 1/12th the amount of time, it is much less computationally intensive and works well with low quality audio samples.

[88] Real-time matching of acoustic fingerprints in a database with minimal delay using exhaustive search

[89] Most existing acoustic fingerprinting systems are not well adapted for real-time applications because they require a large segment of a source audio stream to achieve a valid match. The requirement imposes a time delay of several seconds from the beginning of the matched segment before a matching fingerprint can be confirmed. However, to use fingerprinting in real-time in tandem with other signal processing algorithms such as conventional AMD, the fingerprinting system must identify whether a match is likely within a fraction of a second. To achieve minimal delay, the present inventors propose an approach using an optimized exhaustive search to match an audio stream in real-time against a database on the order of thousands of fingerprints.

[90] As shown in FIGs. 8 and 9, during preprocessing a collection of ternary acoustic fingerprint bitmaps, one for each audio object to detect, with N ternary values per frame, is generated (step 300). For each fingerprint, the ternary bitmap is subdivided into a plurality of equal segments (step 302). In one embodiment, the ternary bitmap is subdivided into four equal segments as shown in FIG. 9. From each segment, a hash key is obtained (step 304) by extracting the ternary values from $128 / N$ frames of the subdivided fingerprint, for a total of 128 ternary values in each hash key (i.e., a 128-bit ternary hash key). It will be appreciated that the hash key may have fewer or a greater number of values. All of the hash keys thus obtained from the ternary acoustic fingerprint bitmaps are then aggregated into a single lookup table (step 306). The selection of hash keys from various segments of the acoustic fingerprint allows a match to be made even if the candidate audio stream is an interrupted recording and the beginning of the recording is not received. For example, if using this methodology for identifying samples of

WO 2013/148069

PCT/US2013/028788

music, the person wishing to identify a song may not have (and most likely did not) start recording a sample of the song to be identified until sometime after the song started playing. Use of hash keys from different segments of the acoustic fingerprint allows the system to still make a match between the candidate audio stream and the acoustic fingerprints in the database.

[91] FIG. 10 illustrates one embodiment of a matching procedure that may be used at runtime after the lookup table of hash keys has been assembled. For each new frame of the input audio stream, a binary bitmap acoustic fingerprinting algorithm is used to generate a new bit vector of length N (step 400). The bit vectors from the previous $128/N$ frames of the input audio stream are combined into a 128-bit binary hash key for lookup (step 402). The current 128-bit binary hash key of the input audio stream is compared against all of the ternary hash keys in the lookup table by computing the masked Hamming distance (step 404). On modern microprocessors, this calculation can easily be parallelized and optimized with Single-Instruction-Multiple-Data instructions (such as the SSE or AVX SIMD instruction set extensions to the Intel x86 microprocessor architecture) and/or the “population count” instruction. If any hash keys in the lookup table match (step 406) with a sufficiently low masked Hamming distance, the audio fingerprint corresponding to the hash key is identified as a “candidate” (step 408). The threshold bit error rate (BER) for establishing a candidate may be based on a relatively low value of $\sigma_{\text{candidate}}$ (e.g., 3 standard deviations, although other thresholds may be used). If the overall BER for a candidate exceeds a relatively large value of σ_{notify} , (e.g., 9 standard deviations, although other thresholds may be used) (step 410) this candidate is determined to be a match (step 412). If no match has been determined, the process returns to step 400 where the next frame is used to generate a new bit vector of length N . By continuing to analyze subsequent frames, the process

is able to resolve all candidates as a match or non-match, based upon where the masked Hamming distance falls with respect to the thresholds.

[92] Method for extracting fingerprints for common recordings from large audio collections

[93] As will be appreciated from the above description, systems and methods are provided for matching an audio sample to a database of known audio recordings. The systems and methods described above are computationally efficient and are able to identify a match with an acoustic fingerprint in a database, but a problem remains in how to efficiently identify the audio recordings that will be represented by acoustic fingerprints in the database. The methodology described hereinbelow for doing so is discussed in the context of a telephonic contact center; however, those skilled in the art will recognize from the present disclosure that the methodology may be applied to compiling fingerprint databases relating to any type of source audio, images, or other types of data.

[94] A sizeable campaign from a contact center may generate thousands of digital recordings of outbound calls each day. From this recording collection, the objective is to extract all audio segments containing call progress events of interest, in order to assist CPA in future call campaigns. For example, many calls in a campaign are dialed to numbers on the same network or sub-network, which plays an identical recorded greeting for each subscriber. In the absence of an automated data driven technique, the large volumes of data must be listened to by a human subject, to identify candidate segments from which acoustic fingerprints may be extracted. The present inventors propose an offline process to automatically identify common acoustic recordings in a large collection of recordings, so they may be detected in future call campaigns to improve the accuracy of CPA. The process used for automated identification of fingerprints in our invention is illustrated schematically in FIG. 11.

WO 2013/148069

PCT/US2013/028788

[95] At step 500, the CPA system, combined with AMD and acoustic fingerprinting (if the database already contains acoustic fingerprints for some known recordings), is used to classify all audio recordings in the collection based on events detected therein. For all audio recordings classified as a live speaker or answering machine using conventional AMD, the endpoints of the audio segments that do not comprise silence, background noise, or simple tones are identified at step 502. For example, audio segments containing speaking patterns, music, non-speech signals, etc., are identified at step 502. At step 504, a ternary acoustic fingerprint bitmap is generated for all identified segments using the methodology described hereinabove (it will be appreciated that a binary acoustic fingerprint bitmap can be used instead of the ternary bitmap). These new acoustic fingerprints are added to the acoustic fingerprint database at step 506. All of the audio recordings in the collection are then re-processed through the CPA/AMD system using the newly-augmented ternary acoustic fingerprint database at step 508. At step 510, the system identifies the augmented fingerprints that are not unique and were detected multiple times (i.e., a recording in the collection being processed matched multiple ones of the augmented acoustic fingerprints, indicating that these are acoustic fingerprints of recordings that were encountered multiple times in the collection). Any acoustic fingerprints discovered multiple times are likely candidates for a call progress event of interest, and are therefore left in the acoustic fingerprint database at step 512 for use in future CPA tasks. All of the other augmented fingerprints (i.e., those new acoustic fingerprints that were not detected multiple times) are removed from the database at step 514, since they represent unique audio segments of vocal utterances, and therefore cannot be determined to be recorded network messages. In this way, an automated system may be used to look through a large collection of recordings (or images or other data) and extract therefrom fingerprints of recordings that were encountered multiple times. This

WO 2013/148069

PCT/US2013/028788

methodology allows the acoustic fingerprint database to be built automatically and continuously updated to account for new network messages that may be brought on line by the various telephone companies.

[96] As can be seen from the above description, the various embodiments allow for the matching of input audio segments with known audio segments in a database. The audio segments may originate from any source and contain any type of audio, such as speech, tones, music, or any other audio that is desired to be recognized. In an illustrative embodiment, the use of the presently disclosed systems and methods was described in conjunction with recognizing known network message recordings encountered during an outbound telephone call. However, those skilled in the art will recognize that the disclosed systems and methods will find application in recognition of any type of two- or more-dimensional dataset, such as any form of audio, image, or other type of data.

[97] While the organization of steps, software blocks, data and data structures have been illustrated as clearly delineated, a person skilled in the art will appreciate that the delineation between steps, blocks and data is somewhat arbitrary. Numerous other arrangements of steps, software blocks and data are possible.

[98] Finally, it will be understood that the invention is not limited to the embodiments described herein which are merely illustrative of several embodiments for carrying out the invention, and which are susceptible to modification of form, arrangement of parts, steps, details and order of operation. The invention, rather, is intended to encompass all such modifications within its spirit and scope, as defined by the claims.

Claims

- 1) A method for creating a ternary bitmap of a dataset, the method comprising the steps of:
 - a) computing a binary fingerprint bitmap of the dataset;
 - 5 b) deleting a first number of samples from the dataset;
 - c) after step (b), computing another binary fingerprint bitmap of the dataset;
 - d) repeating steps (b) and (c) a plurality of times to create a plurality of binary fingerprint bitmaps; and
 - 10 e) combining the plurality of binary fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows:
 - e.1) If a bit is 0 (zero) in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0 (zero);
 - e.2) If a bit is 1 (one) in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1 (one); and
 - 15 e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care").
- 2) The method of claim 1, wherein said first number of samples comprises a first sample from the dataset.
- 3) The method of claim 1, wherein said first predetermined number comprises all of the plurality of binary bitmaps and said second predetermined number comprises all
20 of the plurality of binary bitmaps.
- 4) The method of claim 1, wherein said first predetermined number and said second predetermined number are equal.
- 5) The method of claim 1, wherein the dataset is artificially degraded prior to step (a).

- 6) The method of claim 1, wherein said binary fingerprint bitmap comprises an acoustic fingerprint bitmap.
- 7) The method of claim 1, wherein said dataset comprises an audio segment.
- 8) The method of claim 7, wherein said audio segment comprises an audio segment from an outbound telephone call.
- 9) A method for identifying a candidate dataset, the method comprising the steps of:
 - a) computing a binary fingerprint bitmap of a known dataset in a known dataset database;
 - b) deleting a first number of samples from the known dataset;
 - 10 c) after step (b), computing another binary fingerprint bitmap of the known dataset;
 - d) repeating steps (b) and (c) a plurality of times to create a plurality of binary fingerprint bitmaps;
 - e) combining the plurality of binary fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows,
 - 15 e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0;
 - e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1; and
 - 20 e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care");
 - f) saving the ternary bitmap into a ternary bitmap database;
 - g) repeating steps (a) - (f) for all known datasets in the known dataset database;
 - h) creating a candidate dataset binary fingerprint bitmap from the candidate dataset; and

- i) comparing the candidate dataset binary fingerprint bitmap to each ternary bitmap in the ternary bitmap database, wherein said comparison ignores the Don't Care bits.
- 10) The method of claim 9, wherein said first number of samples comprises a first
5 sample from the dataset.
- 11) The method of claim 9, wherein said first predetermined number comprises all of the plurality of binary bitmaps and said second predetermined number comprises all of the plurality of binary bitmaps.
- 12) The method of claim 9, wherein said first predetermined number and said second
10 predetermined number are equal.
- 13) The method of claim 9, wherein the dataset is artificially degraded prior to step (a).
- 14) The method of claim 9, wherein said binary fingerprint bitmap and said candidate binary fingerprint bitmap each comprise an acoustic fingerprint bitmap.
- 15) The method of claim 9, wherein said dataset comprises an audio segment.
- 15 16) The method of claim 15, wherein said audio segment comprises an audio segment from an outbound telephone call.
- 17) A method for creating a ternary bitmap of an audio segment, the method comprising the steps of:
- a) computing a binary acoustic fingerprint bitmap of the audio segment;
 - 20 b) deleting a first number of samples from the audio segment;
 - c) after step (b), computing another binary acoustic fingerprint bitmap of the audio segment;
 - d) repeating steps (b) and (c) a plurality of times to create a plurality of binary acoustic fingerprint bitmaps; and

- e) combining the plurality of binary acoustic fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows:
 - e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0;
 - 5 e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1; and
 - e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care").
- 18) The method of claim 17, wherein said first number of samples comprises a first sample from the audio segment.
- 10 19) The method of claim 17, wherein said first predetermined number comprises all of the plurality of binary bitmaps and said second predetermined number comprises all of the plurality of binary bitmaps.
- 20) The method of claim 17, wherein said first predetermined number and said second predetermined number are equal.
- 15 21) The method of claim 17, wherein the dataset is artificially degraded prior to step (a).
- 22) The method of claim 17, wherein said audio segment comprises an audio segment from an outbound telephone call.
- 23) A method for identifying a candidate audio segment, the method comprising the steps of:
 - 20 a) computing a binary acoustic fingerprint bitmap of a known audio segment in a known audio segment database;
 - b) deleting a first number of samples from the known audio segment;
 - c) after step (b), computing another binary acoustic fingerprint bitmap of the known audio segment;

- d) repeating steps (b) and (c) a plurality of times to create a plurality of binary acoustic fingerprint bitmaps;
 - e) combining the plurality of binary acoustic fingerprint bitmaps into a ternary bitmap, where each bit in the ternary bitmap is determined as follows:
 - 5 e.1) If a bit is 0 in a first predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 0;
 - e.2) If a bit is 1 in a second predetermined number of the plurality of binary bitmaps, set the bit in the ternary bitmap to 1; and
 - e.3) Otherwise, set the bit of the ternary bitmap to * ("Don't Care");
 - 10 f) saving the ternary bitmap into a ternary bitmap database;
 - g) repeating steps (a) - (f) for all known audio segments in the known audio segment database;
 - h) creating a candidate audio segment binary acoustic fingerprint bitmap from the candidate audio segment; and
 - 15 i) comparing the candidate audio segment binary acoustic fingerprint bitmap to each ternary bitmap in the ternary bitmap database, wherein said comparison ignores the Don't Care bits.
- 24) The method of claim 23, wherein said first number of samples comprises a first sample from the dataset.
- 20 25) The method of claim 23, wherein said first predetermined number comprises all of the plurality of binary bitmaps and said second predetermined number comprises all of the plurality of binary bitmaps.
- 26) The method of claim 23, wherein said first predetermined number and said second predetermined number are equal.
- 25 27) The method of claim 23, wherein the dataset is artificially degraded prior to step (a).

28) The method of claim 23, wherein said audio segment comprises an audio segment from an outbound telephone call.

* * *




Fig. 1A

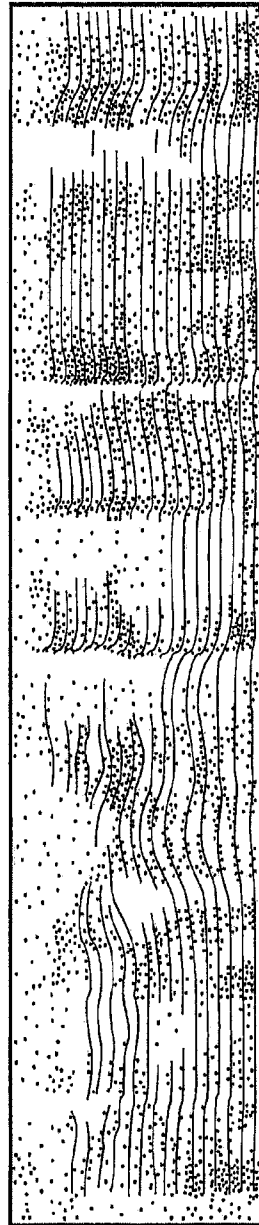


Fig. 1B




Fig. 1C



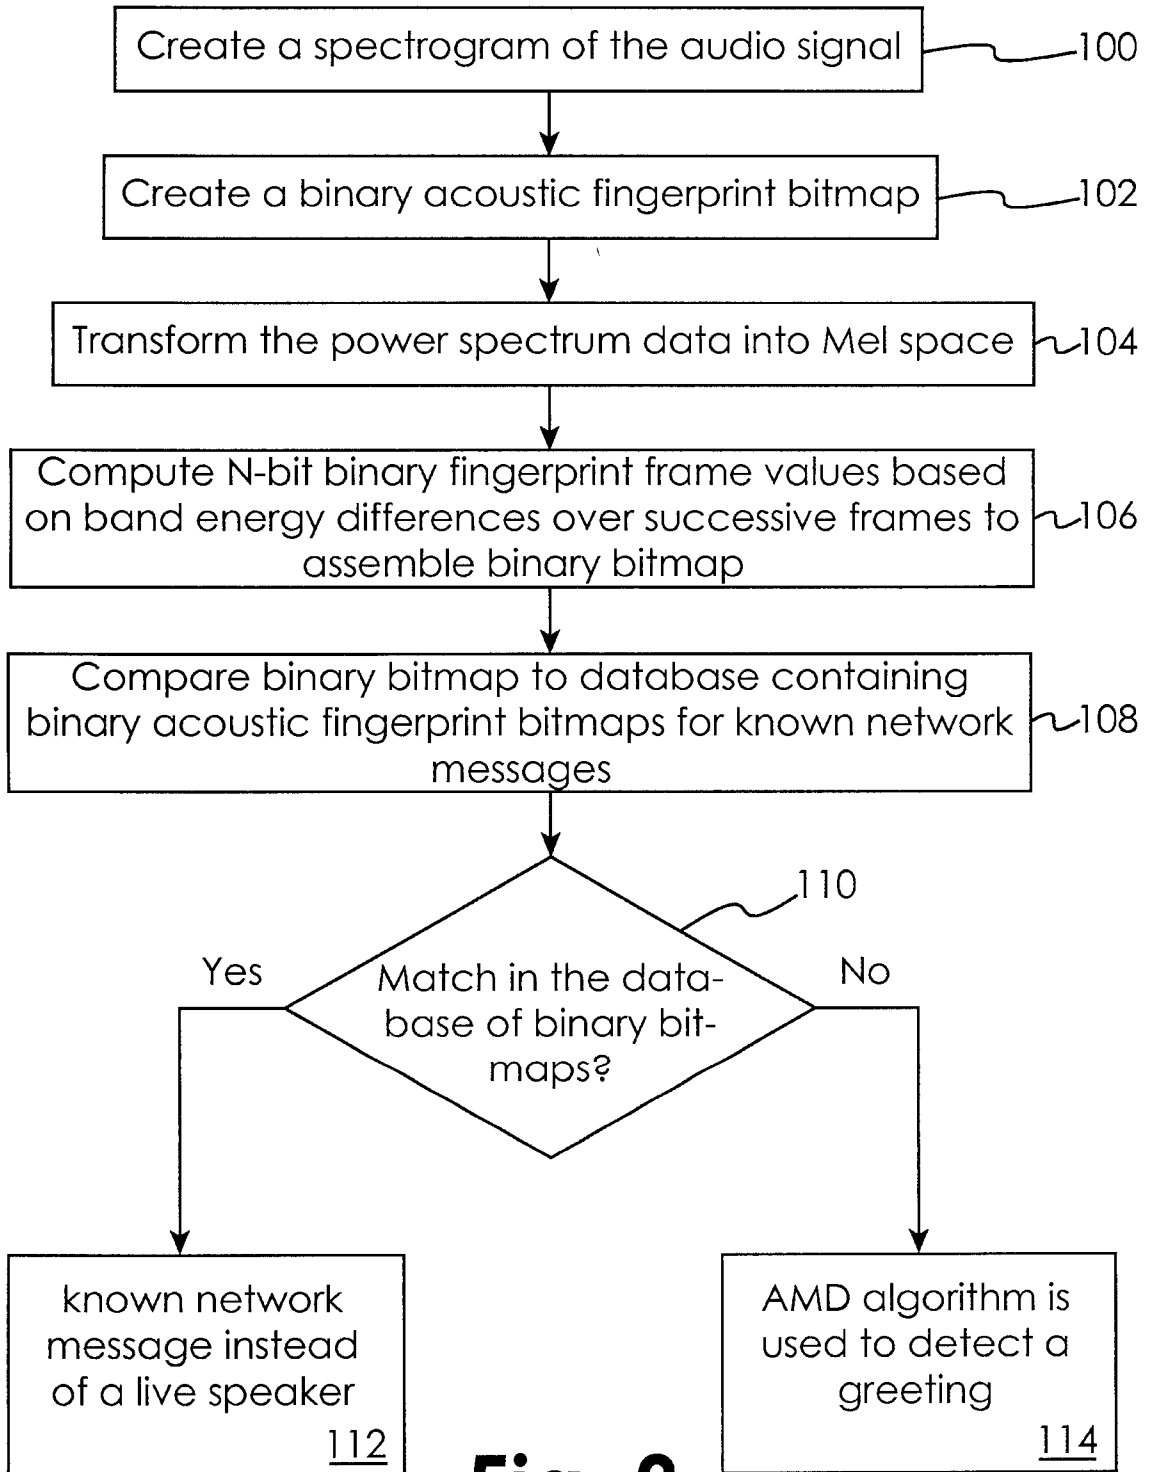


Fig. 2

Overlapped Frames with Hanning Window

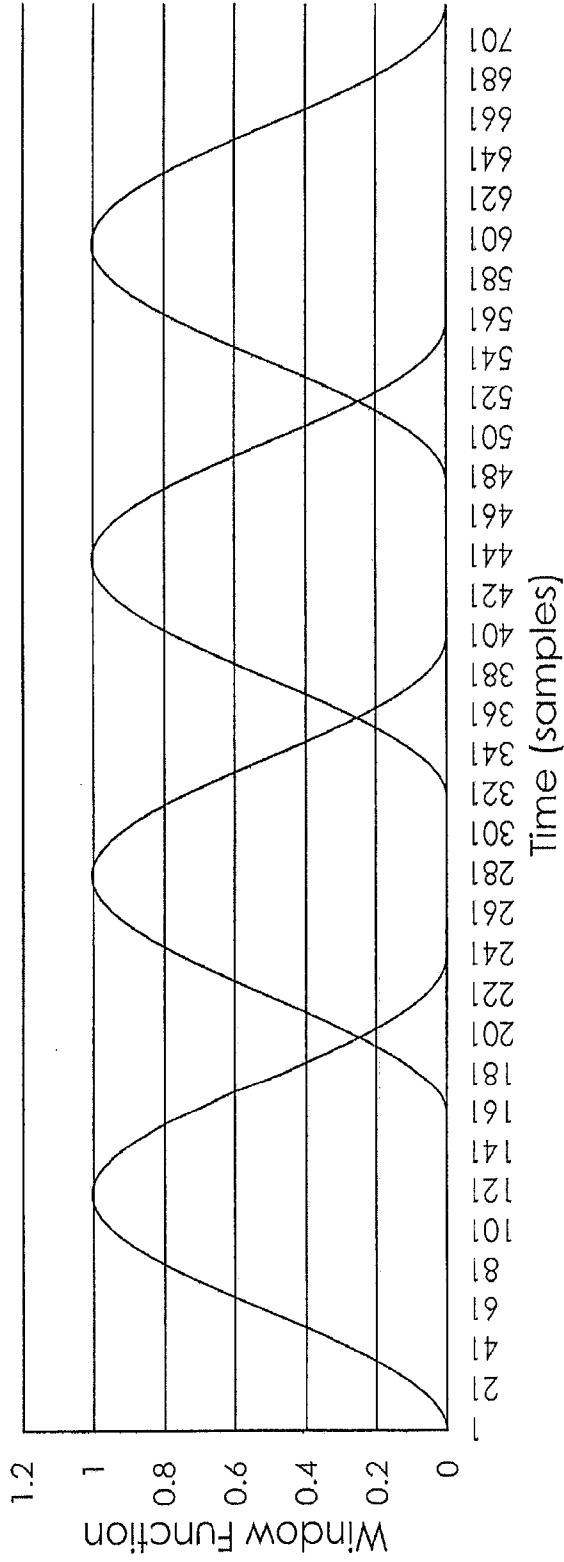


Fig. 3A

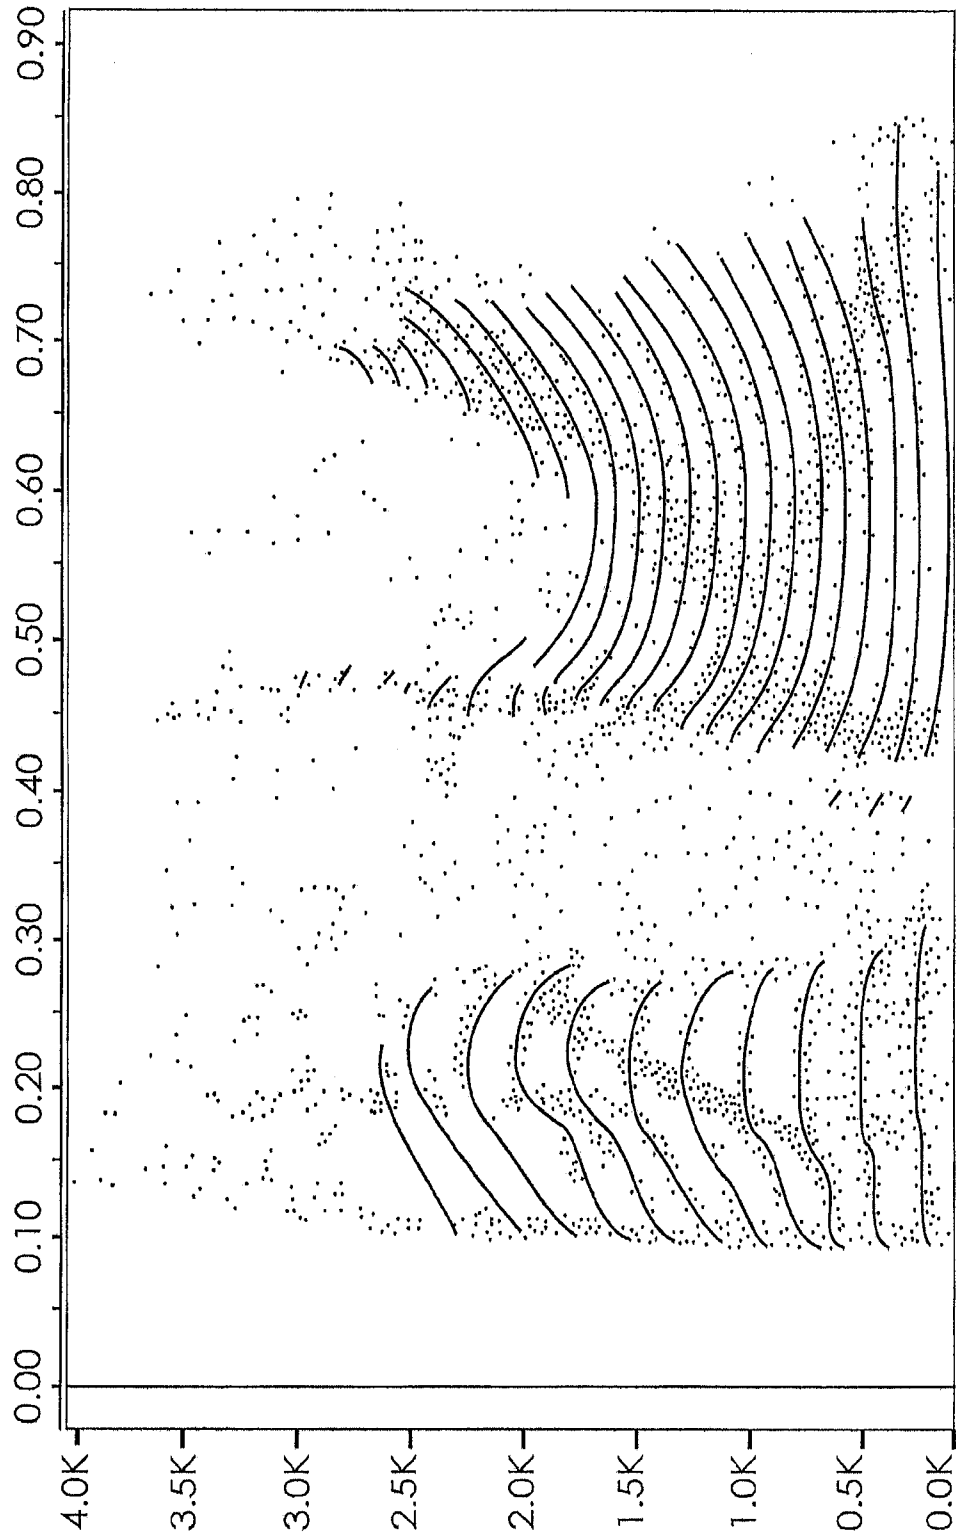


Fig. 3B

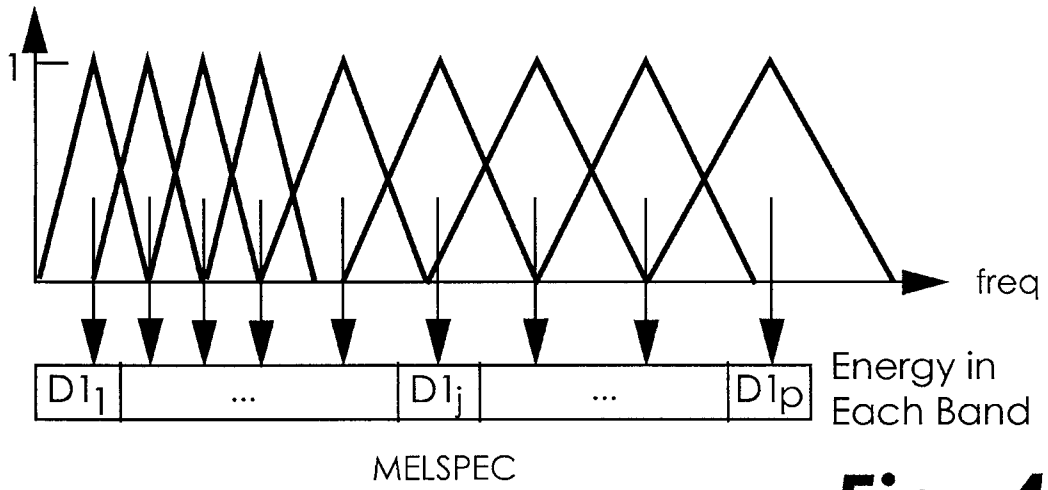


Fig. 4

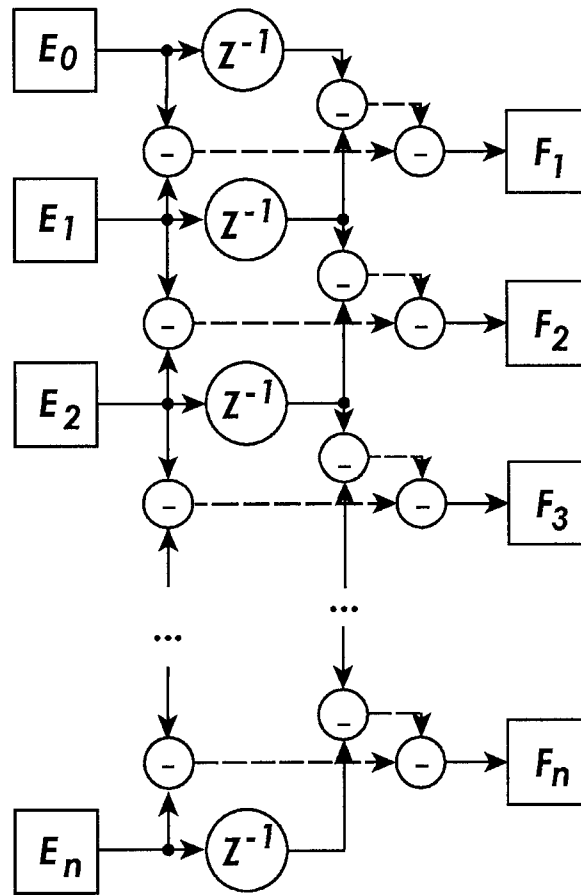
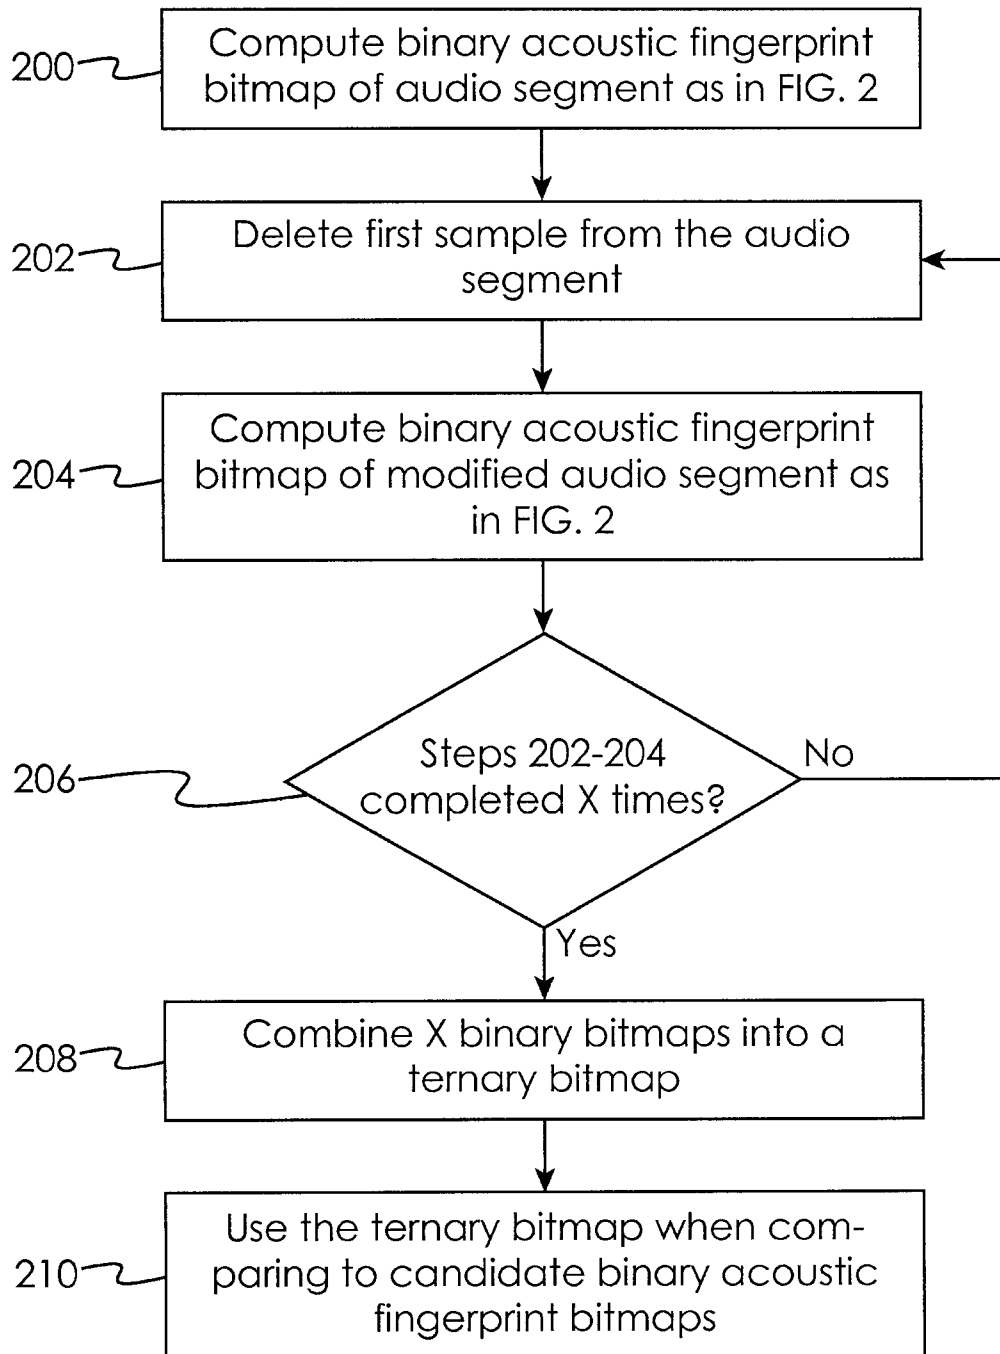




Fig. 5

**Fig. 6**




Binary Bitmap
Acoustic Fingerprint

X



"Don't care" bits

=



Ternary Bitmap
Acoustic Fingerprint

Fig. 7

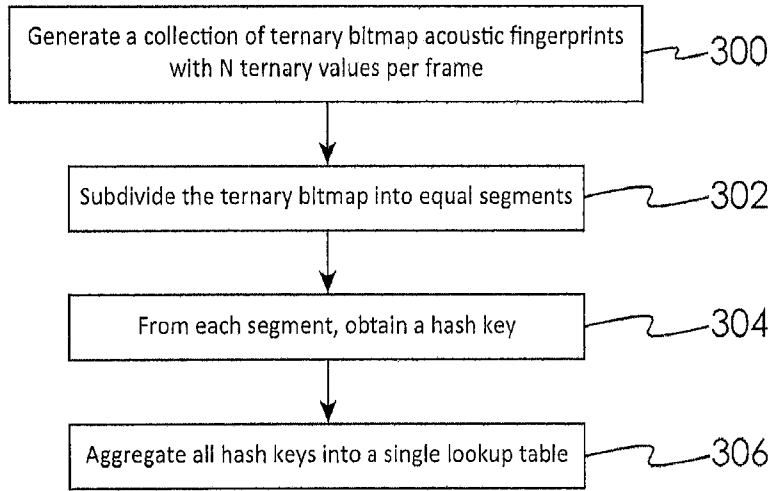


Fig. 8

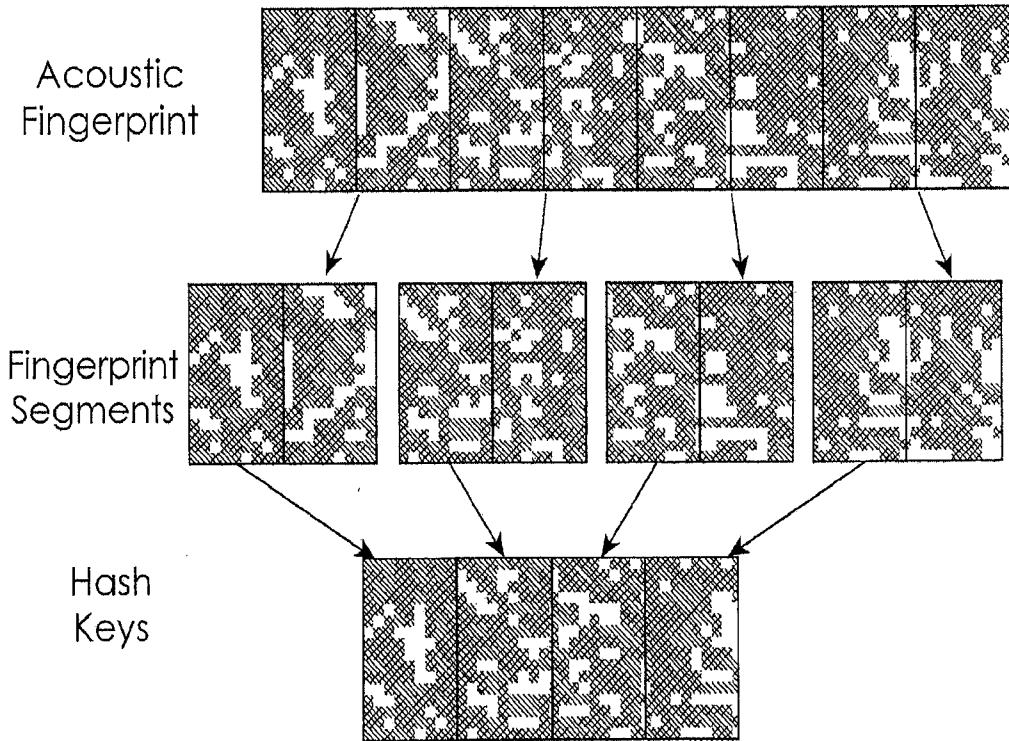


Fig. 9

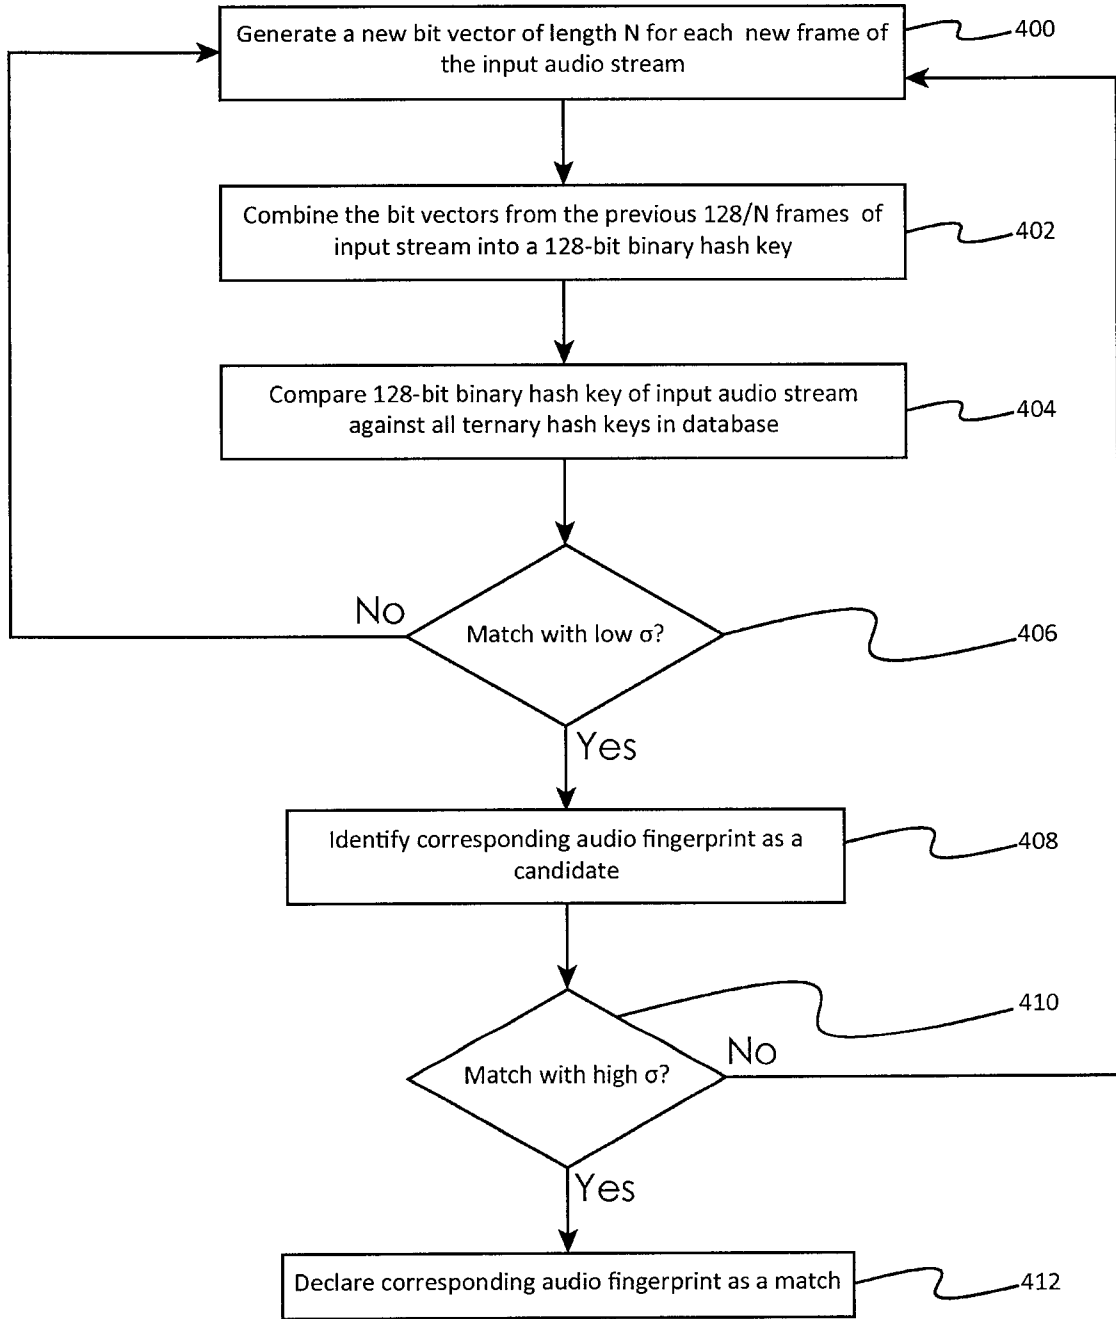
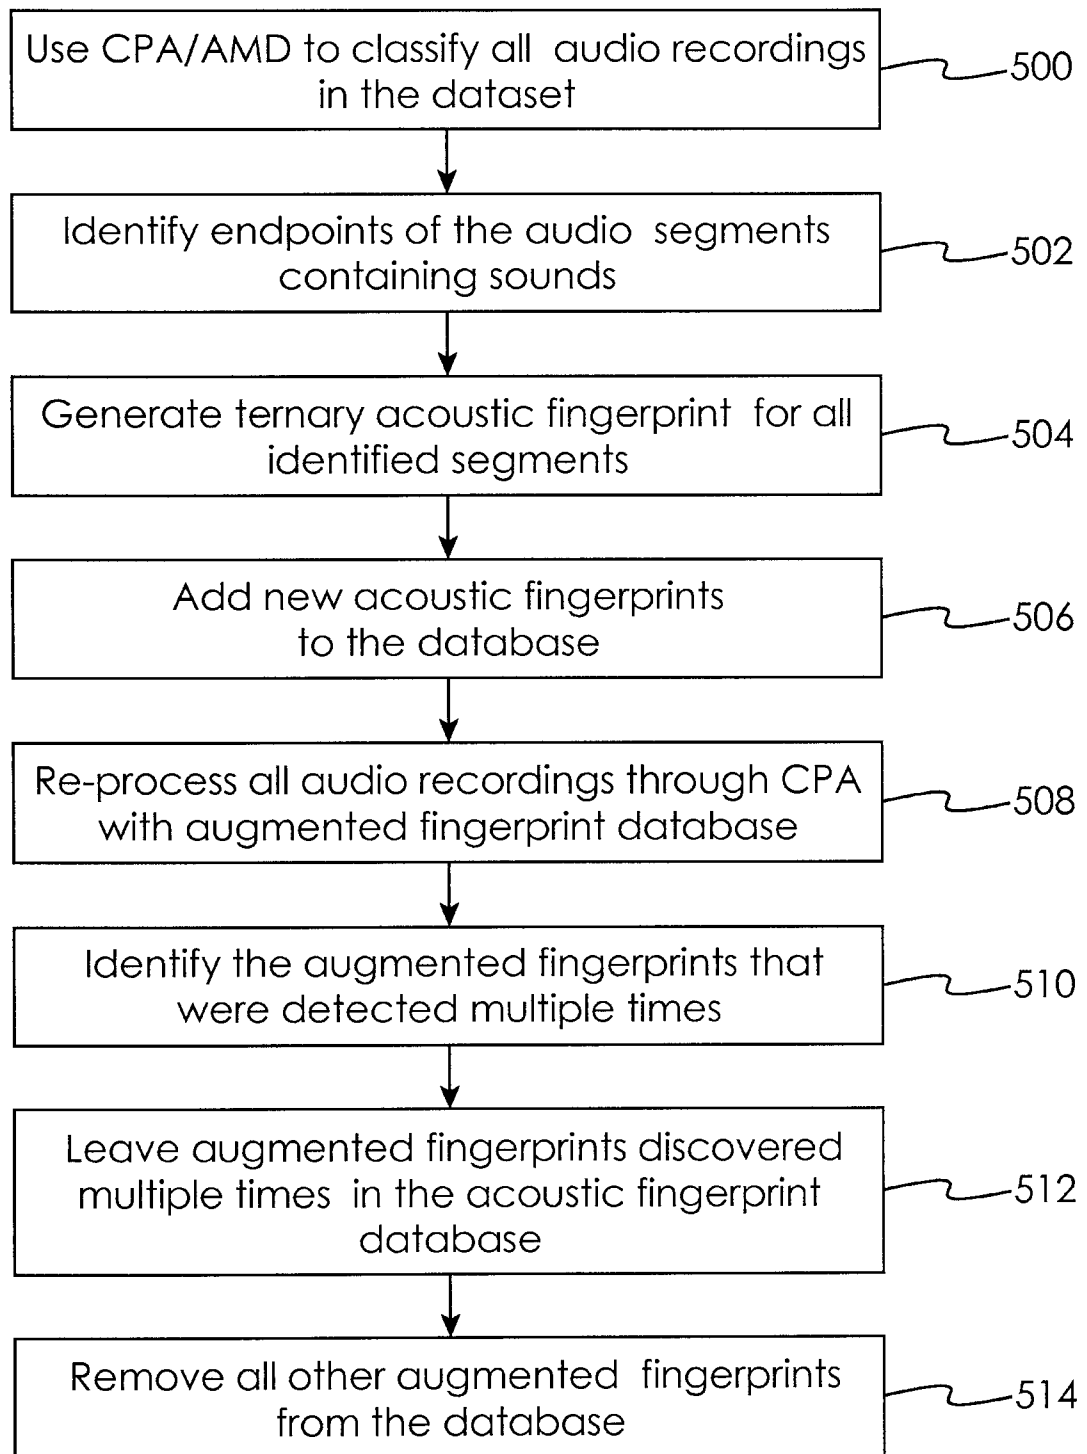


Fig. 10

**Fig. 11**

