

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 July 2006 (13.07.2006)

PCT

(10) International Publication Number
WO 2006/073791 A2

(51) International Patent Classification:
G10L 21/00 (2006.01)

(21) International Application Number:
PCT/US2005/046043

(22) International Filing Date:
20 December 2005 (20.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/905,360 30 December 2004 (30.12.2004) US

(71) Applicant (for all designated States except US): **ALL MEDIA GUIDE, LLC** [US/US]; 1168 Oak Valley Drive, Ann Arbor, MI 48108 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **BOGDAVOV, Vladimir, Askold** [US/US]; 401 River Street, Minneapolis, MN 55401 (US).

(74) Agents: **BERSCHADSKY, Jonathan** et al.; Fitzpatrick, Cella, Harper & Scinto, 30 Rockefeller Plaza, New York, NY 10112-3801 (US).

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

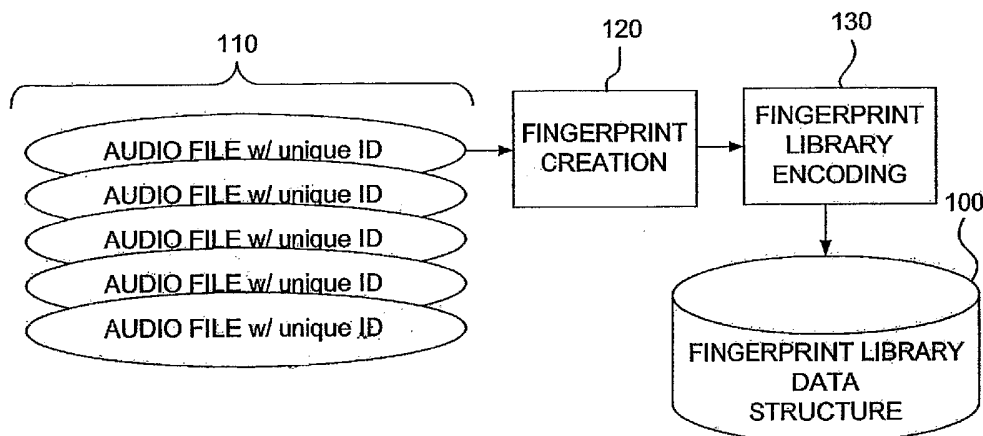
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

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.

(54) Title: METHOD AND APPARATUS FOR IDENTIFYING MEDIA OBJECTS



(57) Abstract: A method, apparatus and computer memory are provided for generating an audio fingerprint of an audio recording. A memory stores stable frequency family data corresponding to a plurality of stable frequency families. A processor curve fits audio recording data to at least one of the stable frequency families, extracts at least one variation from the curve fitted audio recording data, and creates the audio fingerprint of the audio recording from the at least one variation.

WO 2006/073791 A2

Method and Apparatus for Identifying Media Objects

DESCRIPTION

Field of the Invention

[Para 1] The present invention relates generally to delivering supplemental content stored on a database to a user (*e.g.*, supplemental entertainment content relating to an audio recording), and more particularly to determining a fingerprint from a digital file and using the fingerprint to retrieve the supplemental content stored on the database.

Background of the Invention

[Para 2] Recordings can be identified by physically encoding the recording or the media storing one or more recordings, or by analyzing the recording itself. Physical encoding techniques include encoding a recording with a "watermark" or encoding the media storing one or more audio recordings with a TOC (Table of Contents). The watermark or TOC may be extracted during playback and transmitted to a remote database which then matches it to supplemental content to be retrieved. Supplemental content may be, for example, metadata, which is generally understood to mean data that describes other data. In the context of the present invention, metadata may be data that describes the contents of a digital audio compact disc recording. Such metadata may include, for example, artist information (name, birth date, discography, etc.), album information (title, review, track listing, sound samples, etc.), and relational

information (*e.g.*, similar artists and albums), and other types of supplemental information such as advertisements and related images.

[Para 3] With respect to recording analysis, various methods have been proposed. Generally, conventional techniques analyze a recording (or portions of recordings) to extract its “fingerprint,” that is a number derived from a digital audio signal that serves as a unique identifier of that signal. U.S. Patent No. 6,453,252 purports to provide a system that generates an audio fingerprint based on the energy content in frequency subbands. U.S. Application Publication 20040028281 purports to provide a system that utilizes invariant features to generate fingerprints.

[Para 4] Storage space for storing libraries of fingerprints is required for any system utilizing fingerprint technology to provide metadata. Naturally, larger fingerprints require more storage capacity. Larger fingerprints also require more time to create, more time to recognize, and use up more processing power to generate and analyze than do smaller fingerprints.

[Para 5] What is needed is a fingerprinting technology which creates smaller fingerprints, uses less storage space and processing power, is easily scalable and requires relatively little hardware to operate. There also is a need for technology that will enable the management of hundreds or thousands of audio files contained on consumer electronics devices at home, in the car, in portable devices, and the like, which is compact and able to recognize a vast library of music.

Summary of the Invention

- [Para 6] It is an object of the present invention to provide a fingerprinting technology which creates smaller fingerprints, uses less storage space and processing power, is easily scalable and requires relatively little hardware to operate.

[Para 7] It is also an object of the present invention to provide a fingerprint library that will enable the management of hundreds or thousands of audio files contained on consumer electronics devices at home, in the car, in portable devices, and the like, which is compact and able to recognize a vast library of music.

[Para 8] In accordance with one embodiment of the present invention an apparatus for generating an audio fingerprint of an audio recording is provided. The apparatus includes a memory adapted to store stable frequency family data corresponding to a stable frequency families. Also included is a processor operable to curve fit audio recording data to the stable frequency families, extract at least one variation from the curve fitted audio recording data, and create the audio fingerprint of the audio recording from the at least one variation.

[Para 9] In accordance with another embodiment of the present invention a method for generating an audio fingerprint of an audio recording is provided. The method includes curve fitting audio recording data to at least one stable frequency family. The method also includes extracting at least one variation from the curve fitted audio recording data, and creating the audio fingerprint of the audio recording from the at least one variation.

[Para 10] In accordance with yet another embodiment of the present invention computer-readable medium containing code for generating an audio

fingerprint of an audio recording is provided. The code includes code for curve fitting audio recording data to at least one stable frequency family, extracting at least one variation from the curve fitted audio recording data, and creating the audio fingerprint of the audio recording from the at least one variation.

Brief Description of the Drawings

[Para 11] FIG. 1 illustrates a system for creating a fingerprint library data structure on a server.

[Para 12] FIG. 2 illustrates a system for creating a fingerprint from an unknown audio file and for correlating the audio file to a unique audio ID used to retrieve metadata.

[Para 13] FIG. 3 is a flow diagram illustrating how a fingerprint is generated from a multi-frame audio stream.

[Para 14] FIG. 4 illustrates the process performed on an audio frame object.

[Para 15] FIG. 5 is a flowchart illustrating the final steps for creating a fingerprint.

[Para 16] FIG. 6 is an audio file recognition engine for matching the unknown audio fingerprint to known fingerprint data stored in a fingerprint library data structure.

[Para 17] FIG. 7 illustrates a client-server based system for creating a fingerprint from an unknown audio file and for retrieving metadata in accordance with the present invention.

[Para 18] FIG. 8 is device-embedded system for delivering supplemental entertainment content in accordance with the present invention.

Detailed Description of the Invention

[Para 19] As used herein, the term “computer” (also referred to as “processor”) may refer to a single computer or to a system of interacting computers. Generally speaking, a computer is a combination of a hardware system, a software operating system and perhaps one or more software application programs. Examples of computers include, without limitation, IBM-type personal computers (PCs) having an operating system such as DOS, Microsoft Windows, OS/2 or Linux; Apple computers having an operating system such as MAC-OS; hardware having a JAVA-OS operating system; graphical work stations, such as Sun Microsystems and Silicon Graphics Workstations having a UNIX operating system; and other devices such as for example media players (*e.g.*, iPods, PalmPilots Pocket PCs, and mobile telephones).

[Para 20] For the present invention, a software application could be written in substantially any suitable programming language, which could easily be selected by one of ordinary skill in the art. The programming language chosen should be compatible with the computer by which the software application is executed, and in particular with the operating system of that computer. Examples of suitable programming languages include, but are not limited to, Object Pascal, C, C++, CGI, Java and Java Scripts. Furthermore, the functions of the present invention, when described as a series of steps for a method, could be implemented as a series of software instructions for being operated by a data processor, such that the present invention could be implemented as software, firmware or hardware, or a combination thereof.

[Para 21] The present invention uses audio fingerprints to identify audio files encoded in a variety of formats (*e.g.*, WMA, MP3, WAV, and RM) and which have been recorded on different types of physical media (*e.g.*, DVDs, CDs, LPs,

cassette tapes, memory, and hard drives). Once fingerprinted, a retrieval engine may be utilized to match supplemental content to the fingerprints. A computer accessing the recording displays the supplemental content.

[Para 22] The present invention can be implemented in both server-based and client or device-embedded environments. Before the fingerprint algorithm is implemented, the frequency families that exhibit the highest degree of resistance to the compression and/or decompression algorithms ("CODECs") and transformations (such frequency families are also referred to as "stable frequencies") are determined. This determination is made by analyzing a representative set of audio recording files (*e.g.*, several hundred audio files from different genres and styles of music) encoded in common CODECs (*e.g.*, WMA, MP3, WAV, and RM) and different bit rates or processed with other common audio editing software.

[Para 23] The most stable frequency families are determined by analyzing each frequency and its harmonics across the representative set of audio files. First, the range between different renderings for each frequency is measured. The smaller the range, the more stable the frequency. For example, a source file (*e.g.*, one song), is encoded in various formats (*e.g.*, MP3 at 32kbs, 64kbs, 128kbs, etc., WMA at 32kbs, 64kbs, 128kbs, etc.). Ideally, the difference between each rendering would be identical. However, this is not typically the case since compression distorts audio recordings.

[Para 24] Only certain frequencies will be less sensitive to the different renderings. For example, it may be the case that 7kHz is 20dB different between a version of MP3 and a version of WMA, and another frequency, *e.g.*, 8kHz, is just 10dB different. In this example, 8kHz is the more stable frequency. The

measurement used to determine the difference can be any common measure of variation such as standard or maximum deviations. Variation in the context of the present invention is a measure of the change in data, a variable, or a function.

[Para 25] As CODECs are changed and updated, this step might need to be performed again. Typically stable frequencies are determined on a server.

[Para 26] The stable frequencies are extracted from the representative set of audio recording files and collected into a table. The table is then stored onto a client device which compares the stable frequencies to the audio recording being fingerprinted. Frequency families are harmonically related frequencies that are inclusive of all the harmonics of any of its member frequencies and as such can be derived from any member frequency taken as a base frequency. Thus, it is not required to store in the table all of the harmonically related stable frequencies or the core frequency of a family of frequencies.

[Para 27] The client maps the elements of the table to the unknown recording in real time. Thus, as a recording is accessed, it is compared to the table for a match. It is not required to read the entire media (*e.g.*, an entire CD) or the entire audio recording to generate a fingerprint. A fingerprint can be generated on the client based only on a portion of the unknown audio recording.

[Para 28] The present invention will now be described in more detail with reference to FIGS. 1–8.

[Para 29] The evaluation of frequency families described below is performed completely in integer math without using frequency domain transformation methods (*e.g.*, Fast Fourier Transform or FFT).

[Para 30] FIG. 1 illustrates a system for creating a fingerprint library data structure 100 on a server. The data structure 100 is used as a reference for the recognition of unknown audio content and is created prior to receiving a fingerprint of an unknown audio file from a client. All of the available audio recordings 110 on the server are assigned unique identifiers (or IDs) and processed by a fingerprint creation module 120 to create corresponding fingerprints. The fingerprint creation module 120 is the same for both creating the reference library and recognizing the unknown audio.

[Para 31] Once the fingerprint creation has been completed, all of the fingerprints are analyzed and encoded into the data structure by a fingerprint encoder 130. The data structure includes a set of fingerprints organized into groups related by some criteria (also referred to as "feature groups," "summary factors," or simply "features") which are designed to optimize fingerprint access.

[Para 32] FIG. 2 illustrates a system for creating a fingerprint from an unknown audio file 220 and for correlating it to a unique audio ID used to retrieve metadata. The fingerprint is generated using a fingerprint creation module 120 which analyzes the unknown audio recording 220 in the same manner as the fingerprint creation module 120 described above with respect to FIG. 1. In the embodiment shown, the query on the fingerprint takes place on a server 200 using a recognition engine 210 that calculates one or more derivatives of the fingerprint and then attempts to match each derivative to one or more fingerprints stored in the fingerprint library data structure 100. The initial search is an "optimistic" approach because the system is optimistic that the one of the derivatives will be identical to or very similar to one of the feature

groups, thereby reducing the number of (server) fingerprints queried in search of a match.

[Para 33] If the optimistic approach fails, then a “pessimistic” approach attempts to match the received fingerprint to those stored in the server database one at a time using heuristic and conventional search techniques.

[Para 34] Once the fingerprint is matched the audio recording’s corresponding unique ID is used to correlate metadata stored on a database. A preferred embodiment of this matching approach is described below with reference to FIG. 6.

[Para 35] FIG. 3 is a flow diagram illustrating how a fingerprint is generated from a multi-frame audio stream 300. A frame in the context of the present invention is a predetermined size of audio data.

[Para 36] Only a portion of the audio stream is used to generate the fingerprint. In the embodiment described herein only 100 frames are analyzed, where each frame has 8192 bytes of data. This embodiment performs the fingerprinting algorithm of the present invention on encoded or compressed audio data which has been converted into a stereo PCM audio stream.

[Para 37] PCM is typically the format into which most consumer electronics products internally uncompress audio data. The present invention can be performed on any type of audio data file or stream, and therefore is not limited to operations on PCM formatted audio streams. Accordingly, any reference to specific memory sizes, number of frames, sampling rates, time, and the like are merely for illustration.

[Para 38] Silence is very common at the beginning of audio tracks and can potentially lower the quality of the audio recognition. Therefore the present

invention skips silence at the beginning of the audio stream 300, as illustrated in step 300a. Silence need not be absolute silence. For example, low amplitude audio can be skipped until the average amplitude level is greater than a percentage (*e.g.*, 1–2%) of the maximum possible and/or present volume for a predetermined time (*e.g.*, 2–3 second period). Another way to skip silence at the beginning of the audio stream is simply to do just that, skip the beginning of the audio stream for a predetermined amount of time (*e.g.*, 10–12 seconds).

[Para 39] Next, each frame of the audio data is read into a memory and processed, as shown in step 400. In the embodiment described herein, each frame size represents roughly 0.18 seconds of standard stereo PCM audio. If other standards are used, the frame size can be adjusted accordingly. Step 400, which is described in more detail with reference to FIG. 4, processes each frame of the audio stream.

[Para 40] FIG. 4 illustrates the process performed on each audio frame object 300b. At step 415, the frame is read. As each sampling point is read, in step 420, left and right channels are combined by summing and averaging the left and right channel data corresponding to each sampling point. For example, in the case of standard PCM audio, each sampling point will occupy four bytes (*i.e.*, two bytes for each channel). Other well-known forms of combining audio channels can be used and still be within the scope of this invention. Alternatively, only one of the channels can be used for the following analysis. This process is repeated until the entire frame has been read, as show in step 425.

[Para 41] At step 426, data points are stored sequentially into integer arrays corresponding to the predefined number of frequency families. More

particularly, each array has a length of a full cycle of one of the predefined frequencies (*i.e.*, stable frequencies) which, as explained above, also corresponds to a family of frequencies. Since a full wavelength can be equated to a given number of points, each array will have a different size. In other words, an array of x points corresponds to a full wave having x points, and an array of y points corresponds to a full wave having y points. The incoming stream of points are accumulated into the arrays by placing the first incoming data point into the first location of each array, the second incoming data point is placed into the second location in each array, and so on. When the end of an array is reached, the next point is added to the first location in that array. Thus, the contents of the arrays are synchronized from the first point, but will eventually differ since each array has a different length (*i.e.*, represents a different wavelength).

[Para 42] After a full frame is processed, at step 430 each one of the accumulated arrays is curve fitted (*i.e.*, compared) to the "model" array of the perfect sine curve for the same stable frequency. To compensate for any potential phase differential, the array being compared is cyclically shifted N times, where N represents the number of points in the array, and then summed with the model array to find the best fit which represents the level of "resonance" between the audio and the model frequency. This allows the strength of the family of frequencies harmonically related to a given frequency to be estimated.

[Para 43] Referring again to FIG. 3, the last step in the frame processing is combining pairs of frequency families, as shown in step 310. This step reduces the number of frequency families by adding the first array with the second, the third with the fourth, and so on. For example, if the predetermined number of

rows in the matrix is 16, then the 16 rows are reduced to 8. In other words, if 155 frames are processed, then each new array includes two of the original sixteen families of frequencies yielding a 155 x 8 matrix of integer numbers from 155 processed frames, where now there are 8 compound frequency families.

[Para 44] Sometimes there are spikes in the audio data (*e.g.*, pops and clicks), which are artifacts. Trimming a percentage (*e.g.*, 5%–10%) of the highest values to the maximum level can improve the overall performance of algorithm by allowing the most variation (*i.e.*, the most significant range) of the audio content. This is accomplished in Step 320 by normalizing the 155 x 8 matrix to fit into a predetermined range of values (*e.g.*, 0...255).

[Para 45] The audio data may be slightly shifted in time due to the way it is read and/or digitized. That is, the recording may start playback a little earlier or later due to the shift of the audio recording. For example, each time a vinyl LP is played the needle transducer may be placed by the user in a different location from one playback to the next. Thus, the audio recording may not start at the same location, which in effect shifts the LP's start time. Similarly, CD players may also shift the audio content differently due to difference in track-gap playback algorithms. Before the fingerprint is created, another summary matrix is created including a subset of the original 155 x 8 matrix, shown at step 325. This step smoothes the frequency patterns and allows fingerprints to be slightly time-shifted, which improves recognition of time altered audio. The frequency patterns are smoothed by summing the initial 155 x 8 matrix. To account for potential time shifts in the audio, a subset of

the resulting summation is used, leaving room for time shifts. The subset is referred to as a summary matrix.

[Para 46] In the embodiment described herein, the resulting summary matrix has 34 points, each representing the sum of 3 points from the initial matrix. Thus, the summary matrix includes $34 \times 3 = 102$ points allowing for 53 points of movement to account for time shifts caused by different playback devices and/or physical media on which audio content is stored (*e.g.*, ± 2.5 seconds). In practice, the shifting operations need not be point by point and may be multiples thereof. Thus, only a small number of data points from the initial 155×8 matrix are used to create each time-shifted fingerprint, which can improve the speed it takes to analyze time-shifted audio data.

[Para 47] FIG. 5 is a flowchart illustrating the final steps for creating a fingerprint. Various analyses are performed on the 34×8 matrix object 325 created in FIG. 3. In step 500, the 34×8 summary matrix is analyzed to determine the extent of any differences between successive values within each one of the compound frequency families. First, the delta of each pair of successive points within one compound frequency family is determined. Next, the value of each element of the 34×8 matrix is increased by double the delta with right and left neighboring elements within the 34 points, thus rewarding the element with high "contrast" to its neighbors (*e.g.*, an abrupt change in amplitude level).

[Para 48] Step 510 determines, for each point in the 34×8 matrix, which frequencies are predominant (*e.g.*, frequency with highest amplitude) or with very little presence. First, two 8 member arrays are created, where each member of an array is a 4 byte integer. For the first 32 points of each row of the 34×8 summary matrix, a bit in one of the newly created arrays (SGN) is set

to "on" (i.e., a bit is set to one) if a value in the row of the summary matrix exceeds the average of the entire matrix plus a fraction of its standard deviation. For each of the first 32 points in the 34 x 8 summary matrix that is below the average of the entire matrix minus a fraction of its standard deviation a corresponding bit in the second newly created array (SGN₂) is set to "on." The result of this procedure is the two 8 member arrays indicating the distributional values of the original integer matrix, thereby reducing the amount of information necessary to indicate which frequencies are predominant or not present, which in turn helps make processing more efficient.

[Para 49] In step 520, the 8 frequency families are summed together resulting in one 32 point array. From this array, the average and deviation can be calculated and a determination made as to which points exceed the average plus its deviation. For each point in the 32 point array that exceeds the average plus a fraction of the standard deviation, a corresponding bit in another 4-byte integer (SGN₁) is set "on."

[Para 50] Some types of music have very little, if any, variation within a particular span within the audio stream (e.g., within 34 points of audio data). In step 530, a measurement of the quality or "quality measurement factor" (QL) for the fingerprint is defined as the sum of the total variation of the 3 highest variation frequency families. Stated differently, the sum of all differences for each one of the eight combined frequency families results in 8 values representing a total change within a given frequency family. The 3 highest values of the 8 values are those with the most overall change. When added together, the 3 highest values become the QL factor. The QL factor is thus a measurement of the overall variation of the audio as it relates to the model frequency families. If

~~PCT/US2005/046043~~
there is not enough variation, the fingerprint may not be distinctive enough to generate a unique fingerprint, and therefore, may not be sufficient for the audio recognition. The QL factor is thus used to determine if another set of 155 frames from the audio stream should be read and another fingerprint created.

[Para 51] In step 540, a 1 byte integer (SGN2) is created. This value is a bitmap where 5 of its bits correspond to the 5 frequency families with the highest level of variation. The bits corresponding to the frequency families with the highest variation are set on. The variation determination for step 540 and step 530 are the same. For example, the variation can be defined as the sum of differences between values across all of the (time) points. The total of the differences is the variation.

[Para 52] Finally, in step 550, a 1 byte integer value (SGN3) is created to store the translation of the total running time of the audio file (if known) to the 0...255 integer. This translation can take into account the actual running time distribution of the audio content. For example, popular songs typically average in time from 2.5 to 4 minutes. Therefore the majority of the 0...255 range should be allocated to these times. The distribution could be quite different for classical music or for spoken word.

[Para 53] One audio file can potentially have multiple fingerprints associated with it. This might be necessary if the initial QL value is low. The fingerprint creation program continues to read the audio stream and create additional fingerprints until the QL value reaches an acceptable level.

[Para 54] Once the fingerprints have been created for all the available audio files they can be put into the fingerprint library which includes a data structure optimized for the recognition process. As a first step the fingerprints are

clustered into 255 clusters based on the SGN and SGN_ values (*i.e.*, the two integer arrays discussed above with respect to step 510 in FIG. 5). The center point of each cluster is written to the library. Then the whole set of fingerprints is ordered by SGN2 which corresponds to the five frequency families with the highest level of variation.

[Para 55] All fingerprints are written into the library as binary data in an order based on SGN2. As discussed above, SGN and SGN_ represent the most predominant and least present frequencies, respectively. Out of 8 frequency families there are five frequency bands that exhibit the highest level of variation, which are denoted by the bits set in SGN2. Instead of storing 8 integers from each of the SGN and SGN_ arrays, only 5 each are written based of the bits set in SGN2 (*i.e.*, those corresponding to the highest variation frequency families). Advantageously, this saves storage space since the 3 frequency families with the lowest variation are much less likely to contribute to the recognition.

[Para 56] The variation data that remain have the most information. The record in the database is as follows: 1 byte for SGN2, 1 Byte for cluster number, 4 bytes for SGN1, 20 bytes for 5 SGN numbers, 20 bytes for 5 SGN_ numbers, 3 bytes for the audio ID, and 1 byte for SGN3. The size of each fingerprint is thus 50 bytes.

[Para 57] FIG. 6 is an audio file recognition engine for matching the unknown audio fingerprint to known fingerprint data stored in the fingerprint library data structure. As discussed above, the fingerprint for the unknown audio file is created the same way as for the fingerprint library and passed on to the recognition engine. First, the recognition engine determines any potential

clusters the fingerprint could fall into by matching its SGN and SGN_ values against 255 cluster center points, as shown is 610.

[Para 58] In step 620, the recognition engine attempts to recognize the audio in a series of data scans starting with the most direct and therefore the most immediate match cases. The “instant” method assumes that SGN1 matches precisely and SGN2 matches with only a minor difference (*e.g.*, a one bit variation). If the “instant” method does not yield a match, then a “quick” method is invoked in step 630 which allows a difference (*e.g.*, up to a 2 bit variation) on SGN2 and no direct matches on SGN1.

[Para 59] If still no match is found, in step 640 a “standard” scan is used, which may or may not match SGN2, but uses SGN2, SGN1 and potential fingerprint cluster numbers as a quick heuristic to reject a large number of records as a potential match. If still no match is found in step 650 a “full” scan of the database is evoked as the last resort.

[Para 60] Each method keeps a running list of the best matches and the corresponding match levels. If the purpose of recognition is to return a single ID, the process can be interrupted at any point once an acceptable level of match is reached, thus allowing for very fast and efficient recognition. If on the other hand, all possible matches need to be returned, the “standard” and “full” scan should be used.

[Para 61] FIG. 7 illustrates a client-server based system for creating a fingerprint from an unknown audio file and for retrieving metadata in accordance with the present invention. The client PC 700 may be any computer connected to a network 760.

[Para 62] The exchange of information between a client and a recognition server 750 include returning a web page with metadata based on a fingerprint. The exchange can be automatic, triggered for example when an audio recording is uploaded onto a computer (or a CD placed into a CD player), a fingerprint is automatically generated using a fingerprint creation module (not shown), which analyzes the unknown audio recording in the same manner as described above. After the fingerprint creation engine generates a fingerprint 710, the client PC 700 transmits the fingerprint onto the network 760 to a recognition server 750, which for example may be a Web server. Alternatively, the fingerprint creation and recognition process can be triggered manually, for instance by a user selecting a menu option on a computer which instructs the creation and recognition process to begin.

[Para 63] The network can be any type of connection between any two or more computers, which permits the transmission of data. An example of a network, although it is by no means the only example, is the Internet.

[Para 64] A query on the fingerprint takes place on a recognition server 750 by calculating one or more derivatives of the fingerprint and matching each derivative to one or more fingerprints stored in a fingerprint library data structure. Upon recognition of the fingerprint, the recognition server 750 transmits audio identification and metadata via the network 760 to the client PC 700. Internet protocols may be used to return data to the application which runs the client, which for example may be implemented in a web browser, such as Internet Explorer, Mozilla or Netscape Navigator, or on a proprietary media viewer.

[Para 65] Alternatively, the invention may be implemented without client-server architecture and/or without a network. Instead, all software and data necessary for the practice of the present invention may be stored on a storage device associated with the computer (also referred to as a device-embedded system). In a most preferred embodiment the computer is an embedded media player. For example, the device may use a CD/DVD drive, hard drive, or memory to playback audio recordings. Since the present invention uses simple arithmetic operations to perform audio analysis and fingerprint creation, the device's computing capabilities can be quite modest and the bulk of the device's storage space can be utilized more effectively for storing more audio recordings and corresponding metadata.

[Para 66] As illustrated in FIG. 8, a recognition engine 830 may be installed onto the device 800, which includes embedded data stored on a CD drive, hard drive, or in memory. The embedded data may contain a complete set or a subset of the information available in the databases on a recognition server 750 such as the one described above with respect to FIG. 7. Updated databases may be loaded onto the device using well known techniques for data transfer (*e.g.*, FTP protocol). Thus, instead of connecting to a remote database server each time fingerprint recognition is sought, databases may be downloaded and updated occasionally from a remote host via a network. The databases may be downloaded from a Web site via the Internet through a WI-FI, WAP or BlueTooth connection, or by docking the device to a PC and synchronizing it with a remote server.

[Para 67] More particularly, after the fingerprint creation engine 810 generates a fingerprint 840, the device 800 internally communicates the fingerprint 840 to

an internal recognition engine 830 which includes a library for storing metadata and audio recording identifiers (IDs). The recognition engine 830 recognizes a match, and communicates an audio ID and metadata corresponding to the audio recording. Other variations exist as well.

[Para 68] While the present invention has been described with respect to what is presently considered to be the preferred embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. To the contrary, the invention is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims. The scope of the following claims is to be accorded the broadest interpretation so as to encompass all such modifications and equivalent structures and functions.

What is claimed is:

[Claim 1] An apparatus for generating an audio fingerprint of an audio recording, comprising:

memory adapted to store stable frequency family data corresponding to a plurality of stable frequency families;

processor operable to curve fit audio recording data to at least one of the stable frequency families, extract at least one variation from the curve fitted audio recording data, and create an audio fingerprint of the audio recording from the at least one variation.

[Claim 2] An apparatus according to Claim 1, the processor operable to combine the frequency families of the curve fitted audio recording, create a summary matrix from a subset of the combined frequency families, and detect the at least one variation from the summary matrix.

[Claim 3] An apparatus according to Claim 2, wherein the processor is further operable to determine the difference between successive values within each one of the combined frequency families and increase the value of each element of the summary matrix.

[Claim 4] An apparatus according to Claim 2, wherein the at least one variation is based on at least one of the predominance and presence of the combined frequency families.

[Claim 5] An apparatus according to Claim 2, the processor further operable to sum the frequency families of the summary matrix, average the summed frequency and add a deviation to the average of the summed frequency families, wherein the at least one variation is based on the average of the summed frequency families plus the deviation.

[Claim 6] An apparatus according to Claim 2, wherein the at least one variation is the sum of the total variation of a predetermined number of the highest variation frequency families.

[Claim 7] An apparatus according to Claim 2, wherein the at least one variation is the sum of a predetermined number of the frequency families having the highest level of variation.

[Claim 8] An apparatus according to Claim 2, wherein the memory is further adapted to store an integer, wherein a predetermined number of bits of the integer are set to indicate the frequency families with the highest level of variation.

[Claim 9] An apparatus according to Claim 2, wherein the at least one variation is based on the translation of the total running time of an audio file.

[Claim 10] An apparatus according to Claim 1, wherein the processor is further operable to measure the range of differences of a plurality of audio recordings between different renderings of said audio recordings and select a predetermined number of frequencies having the highest degree of resistance to the different renderings, thereby determining the stable frequency families.

[Claim 11] An apparatus according to Claim 10, wherein the processor is further operable to store in the memory the data corresponding to the stable frequency families.

[Claim 12] An apparatus according to Claim 1, wherein the processor is further operable to sequentially store the audio recording data into a plurality of integer arrays corresponding to the stable frequency families.

[Claim 13] An apparatus according to Claim 12, wherein each one of the integer arrays has a length of a full cycle corresponding to one of the stable frequency families.

[Claim 14] An apparatus according to Claim 1, wherein the processor is further operable to skip a predetermined amount of the audio recording data.

[Claim 15] An apparatus according to Claim 2, wherein the processor is further operable to normalize the combined frequency families into a predetermined range of values.

[Claim 16] An apparatus according to Claim 2, wherein the processor is further operable to shift a plurality of points in the summary matrix.

[Claim 17] An apparatus according to Claim 2, wherein the processor is further operable to compensate for a time shift in the audio recording.

[Claim 18] An apparatus according to Claim 1, wherein the processor is further operable to match the audio fingerprint to a known fingerprint stored in a database.

[Claim 19] An apparatus according to Claim 1, wherein the processor is further operable to recognize the audio fingerprint based on at least one of the variations.

[Claim 20] An apparatus according to Claim 18, wherein the processor is further adapted to retrieve metadata from the database corresponding to the audio fingerprint.

[Claim 21] A network computer system, comprising the apparatus for generating an audio fingerprint of Claim 1.

[Claim 22] A device-embedded system, comprising the apparatus for generating an audio fingerprint of Claim 1.

[Claim 23] A method for generating an audio fingerprint of an audio recording, comprising:

curve fitting audio recording data to at least one stable frequency family;
extracting at least one variation from the curve fitted audio recording data; and
creating the audio fingerprint of the audio recording from the at least one variation.

[Claim 24] A method according to Claim 23, further comprising:

combining the frequency families of the curve fitted audio recording;
creating a summary matrix from a subset of the combined frequency families; and
detecting the at least one variation from the summary matrix.

[Claim 25] A method according to Claim 24, further comprising:

determining the difference between successive values within each one of the combined frequency families; and
increasing the value of each element of the summary matrix.

[Claim 26] A method according to Claim 24, further comprising:

determining the predominance and presence of the combined frequency families, wherein
the at least one variation is based on the determination.

[Claim 27] A method according to Claim 24, further comprising:

summing the frequency families of the summary matrix; and
averaging the result of the summing step; and
adding a deviation to the averaging step,
wherein the at least one variation is based on the average of the summed frequency families plus a deviation.

[Claim 28] A method according to Claim 24, further comprising:

summing the total variation of a predetermined number of the highest variation frequency families, wherein the at least one variation is the result of the summing step.

[Claim 29] A method according to Claim 24, further comprising:
summing a predetermined number of the frequency families having the highest level of variation, wherein the at least one variation is the result of the summing step.

[Claim 30] A method according to Claim 24, further comprising:
setting a predetermined number of bits of an integer array to indicate the frequency families with the highest level of variation.

[Claim 31] A method according to Claim 24, further comprising:
translating the total running time of an audio file, wherein the at least one variation is based on the result of the translating step.

[Claim 32] A method according to Claim 23, further comprising:
measuring the range of differences of a plurality of audio recordings between different renderings of said audio recordings; and
selecting a predetermined number of frequencies having the highest degree of resistance to the different renderings, thereby determining the stable frequency families.

[Claim 33] A method according to Claim 32, further comprising:
recording the data corresponding to the stable frequency families.

[Claim 34] A method according to Claim 23, further comprising:
sequentially storing the audio recording data into a plurality of integer arrays corresponding to the stable frequency families.

[Claim 35] A method according to Claim 34, wherein each one of the integer arrays
has a length of a full cycle corresponding to one of the stable frequency families.

[Claim 36] A method according to Claim 23, further comprising:
skipping a predetermined amount of the audio recording data.

[Claim 37] A method according to Claim 24, further comprising:

normalizing the combined frequency families into a predetermined range of values.

[Claim 38] A method according to Claim 24, further comprising:
shifting a plurality of points in the summary matrix.

[Claim 39] A method according to Claim 24, further comprising:
compensating for a time shift in the audio recording.

[Claim 40] A method according to Claim 23, further comprising:
matching the audio fingerprint to a known fingerprint stored in a database.

[Claim 41] A method according to Claim 23, further comprising:
recognizing the audio fingerprint based on at least one of the variations.

[Claim 42] A method according to Claim 41, further comprising:
retrieving metadata from the database corresponding to the audio fingerprint.

[Claim 43] An apparatus for generating an audio fingerprint of an audio recording,
comprising:
means for curve fitting audio recording data to at least one stable frequency family;
means for extracting at least one variation from the curve fitted audio recording data; and
means for creating the audio fingerprint of the audio recording from the at least one
variation.

[Claim 44] An apparatus according to Claim 43, further comprising:
means for combining the frequency families of the curve fitted audio recording;
means for creating a summary matrix from a subset of the combined frequency families;
and
means for detecting the at least one variation from the summary matrix.

[Claim 45] An apparatus according to Claim 44, further comprising:
means for determining the difference between successive values within each one of the
combined frequency families; and

means for increasing the value of each element of the summary matrix.

[Claim 46] An apparatus according to Claim 44, further comprising:
means for determining the predominance and presence of the combined frequency families,
wherein the at least one variation is based the determination.

[Claim 47] An apparatus according to Claim 44, further comprising:
means for summing the frequency families of the summary matrix; and
means for averaging the result of the summing step; and
means for adding a deviation to the averaging step,
wherein the at least one variation is based on the average of the summed frequency families
plus a deviation.

[Claim 48] An apparatus according to Claim 44, further comprising:
means for summing the total variation of a predetermined number of the highest variation
frequency families to obtain the at least one variation.

[Claim 49] An apparatus according to Claim 44, further comprising:
means for summing a predetermined number of the frequency families having the highest
level of variation to obtain the at least one variation.

[Claim 50] An apparatus according to Claim 44, further comprising:
means for setting a predetermined number of bits of an integer array to indicate the
frequency families with the highest level of variation.

[Claim 51] An apparatus according to Claim 44, further comprising:
means for translating the total running time of an audio file to obtain the at least one
variation.

[Claim 52] An apparatus according to Claim 43, further comprising:
means for measuring the range of differences of a plurality of audio recordings between
different renderings of said audio recordings; and

means for selecting a predetermined number of frequencies having the highest degree of resistance to the different renderings, thereby determining the stable frequency families.

[Claim 53] An apparatus according to Claim 52, further comprising:
means for recording the data corresponding to the stable frequency families.

[Claim 54] An apparatus according to Claim 43, further comprising:
means for sequentially storing the audio recording data into a plurality of integer means for arrays corresponding to the stable frequency families.

[Claim 55] An apparatus according to Claim 54, wherein each one of the integer arrays has a length of a full cycle corresponding to one of the stable frequency families.

[Claim 56] An apparatus according to Claim 43, further comprising:
means for skipping a predetermined amount of the audio recording data.

[Claim 57] An apparatus according to Claim 44, further comprising:
means for normalizing the combined frequency families into a predetermined range of values.

[Claim 58] An apparatus according to Claim 44, further comprising:
means for shifting a plurality of points in the summary matrix.

[Claim 59] An apparatus according to Claim 44, further comprising:
means for compensating for a time shift in the audio recording.

[Claim 60] An apparatus according to Claim 43, further comprising:
means for matching the audio fingerprint to a known fingerprint stored in a database.

[Claim 61] An apparatus according to Claim 43, further comprising:
means for recognizing the audio fingerprint based on at least one of the variations.

[Claim 62] An apparatus according to Claim 61, further comprising:
means for retrieving metadata from the database corresponding to the audio fingerprint.

[Claim 63] Computer-readable medium containing code for generating an audio fingerprint of an audio recording, said code for:

curve fitting audio recording data to at least one stable frequency family;
extracting at least one variation from the curve fitted audio recording data; and
creating the audio fingerprint of the audio recording from the at least one variation.

[Claim 64] Computer-readable medium containing code according to Claim 63,
further including code for:

combining the frequency families of the curve fitted audio recording;
creating a summary matrix from a subset of the combined frequency families; and
detecting the at least one variation from the summary matrix.

[Claim 65] Computer-readable medium containing code according to Claim 64,
further including code for:

determining the difference between successive values within each one of the combined frequency families; and
increasing the value of each element of the summary matrix.

[Claim 66] Computer-readable medium containing code according to Claim 64,
further including code for:

determining the predominance and presence of the combined frequency families, wherein
the at least one variation is based on the determination step.

[Claim 67] Computer-readable medium containing code according to Claim 64,
further including code for:

summing the frequency families of the summary matrix; and
averaging the result of the summing step; and
adding a deviation to the averaging step,

wherein the at least one variation is based on the average of the summed frequency families plus a deviation.

[Claim 68] Computer-readable medium containing code according to Claim 64,
further including code for:

summing the total variation of a predetermined number of the highest variation frequency families, wherein the at least one variation is the result of the summing step.

[Claim 69] Computer-readable medium containing code according to Claim 64,
further including code for:

summing a predetermined number of the frequency families having the highest level of variation, wherein the at least one variation is the result of the summing step.

[Claim 70] Computer-readable medium containing code according to Claim 64,
further including code for:

setting a predetermined number of bits of an integer array to indicate the frequency families with the highest level of variation.

[Claim 71] Computer-readable medium containing code according to Claim 64,
further including code for:

translating the total running time of an audio file, wherein the at least one variation is based on the result of the translating step.

[Claim 72] Computer-readable medium containing code according to Claim 63,
further including code for:

measuring the range of differences of a plurality of audio recordings between different renderings of said audio recordings; and

selecting a predetermined number of frequencies having the highest degree of resistance to the different renderings, thereby determining the stable frequency families.

[Claim 73] Computer-readable medium containing code according to Claim 72,
further including code for:
recording the data corresponding to the stable frequency families.

[Claim 74] Computer-readable medium containing code according to Claim 63,
further including code for:
sequentially storing the audio recording data into a plurality of integer arrays corresponding
to the stable frequency families.

[Claim 75] Computer-readable medium containing code according to Claim 74,
wherein each one of the integer arrays has a length of a full cycle
corresponding to one of the stable frequency families.

[Claim 76] Computer-readable medium containing code according to Claim 63,
further including code for:
skipping a predetermined amount of the audio recording data.

[Claim 77] Computer-readable medium containing code according to Claim 64,
further including code for:
normalizing the combined frequency families into a predetermined range of values.

[Claim 78] Computer-readable medium containing code according to Claim 64,
further including code for:
shifting a plurality of points in the summary matrix.

[Claim 79] Computer-readable medium containing code according to Claim 64,
further including code for:
compensating for a time shift in the audio recording.

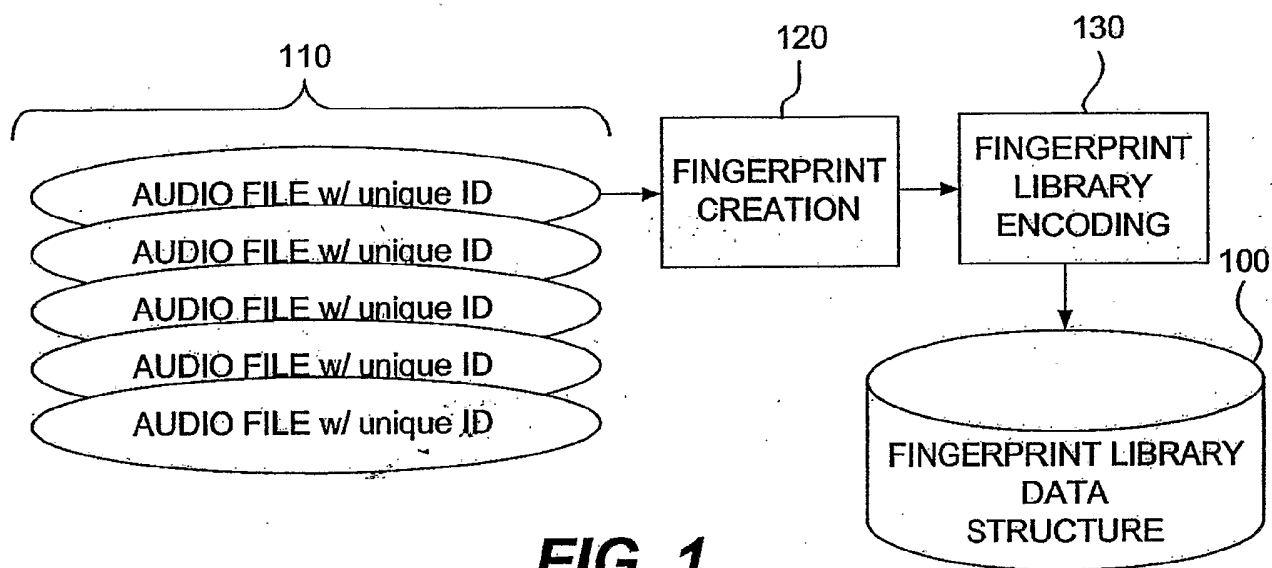
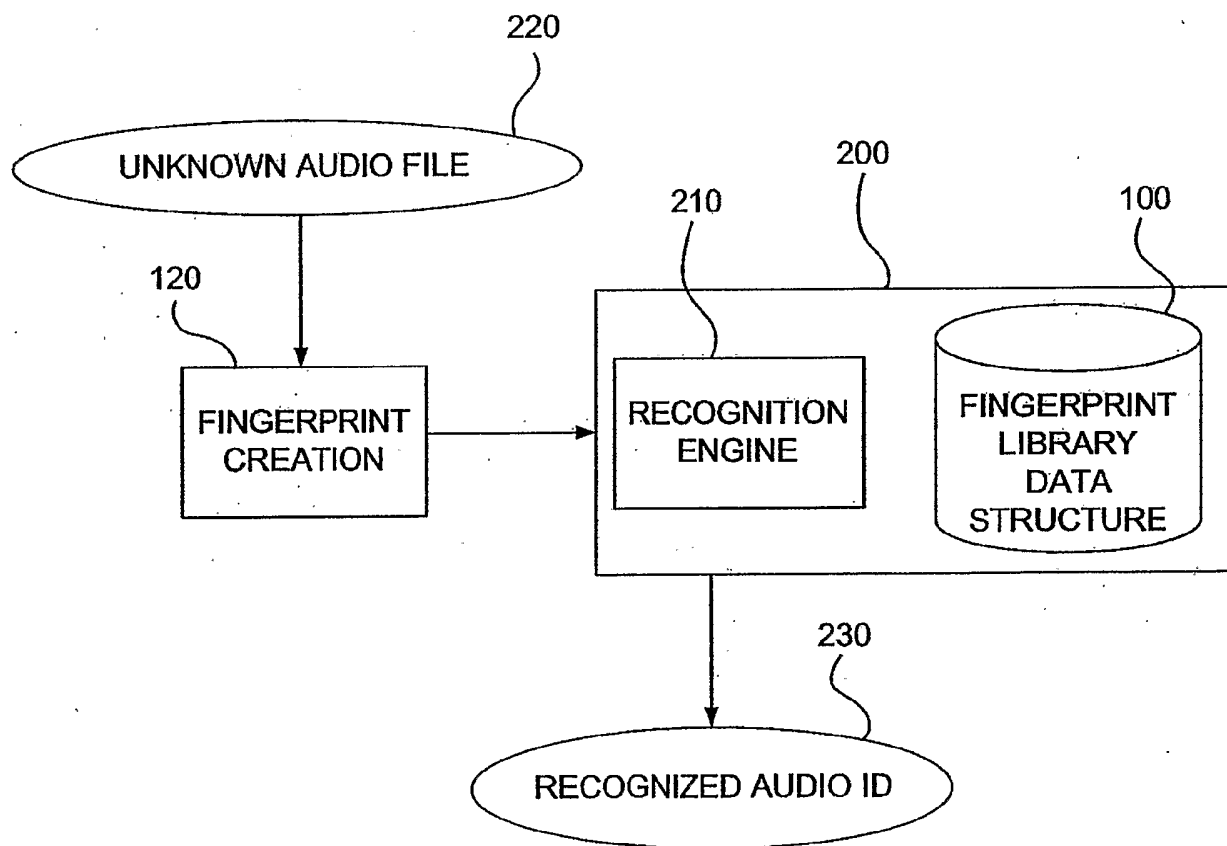
[Claim 80] Computer-readable medium containing code according to Claim 63,
further including code for:
matching the audio fingerprint to a known fingerprint stored in a database.

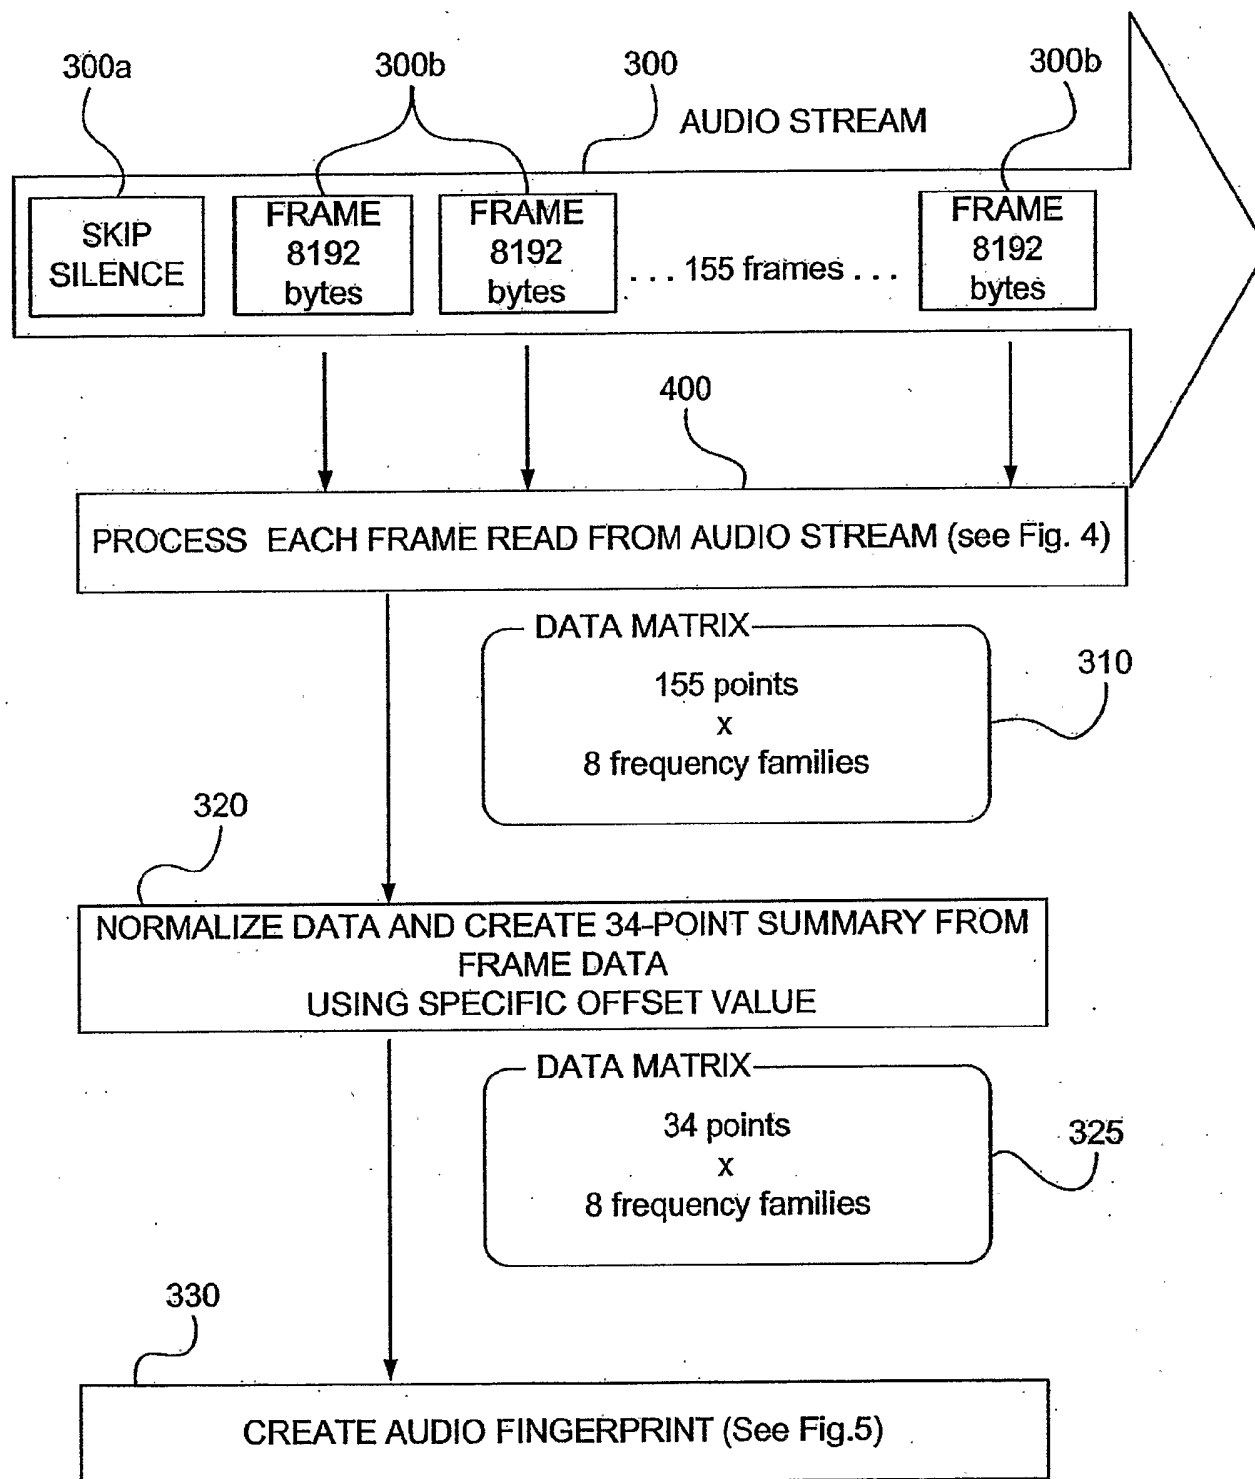
[Claim 81] Computer-readable medium containing code according to Claim 63,
further including code for:
recognizing the audio fingerprint based on at least one of the variations.

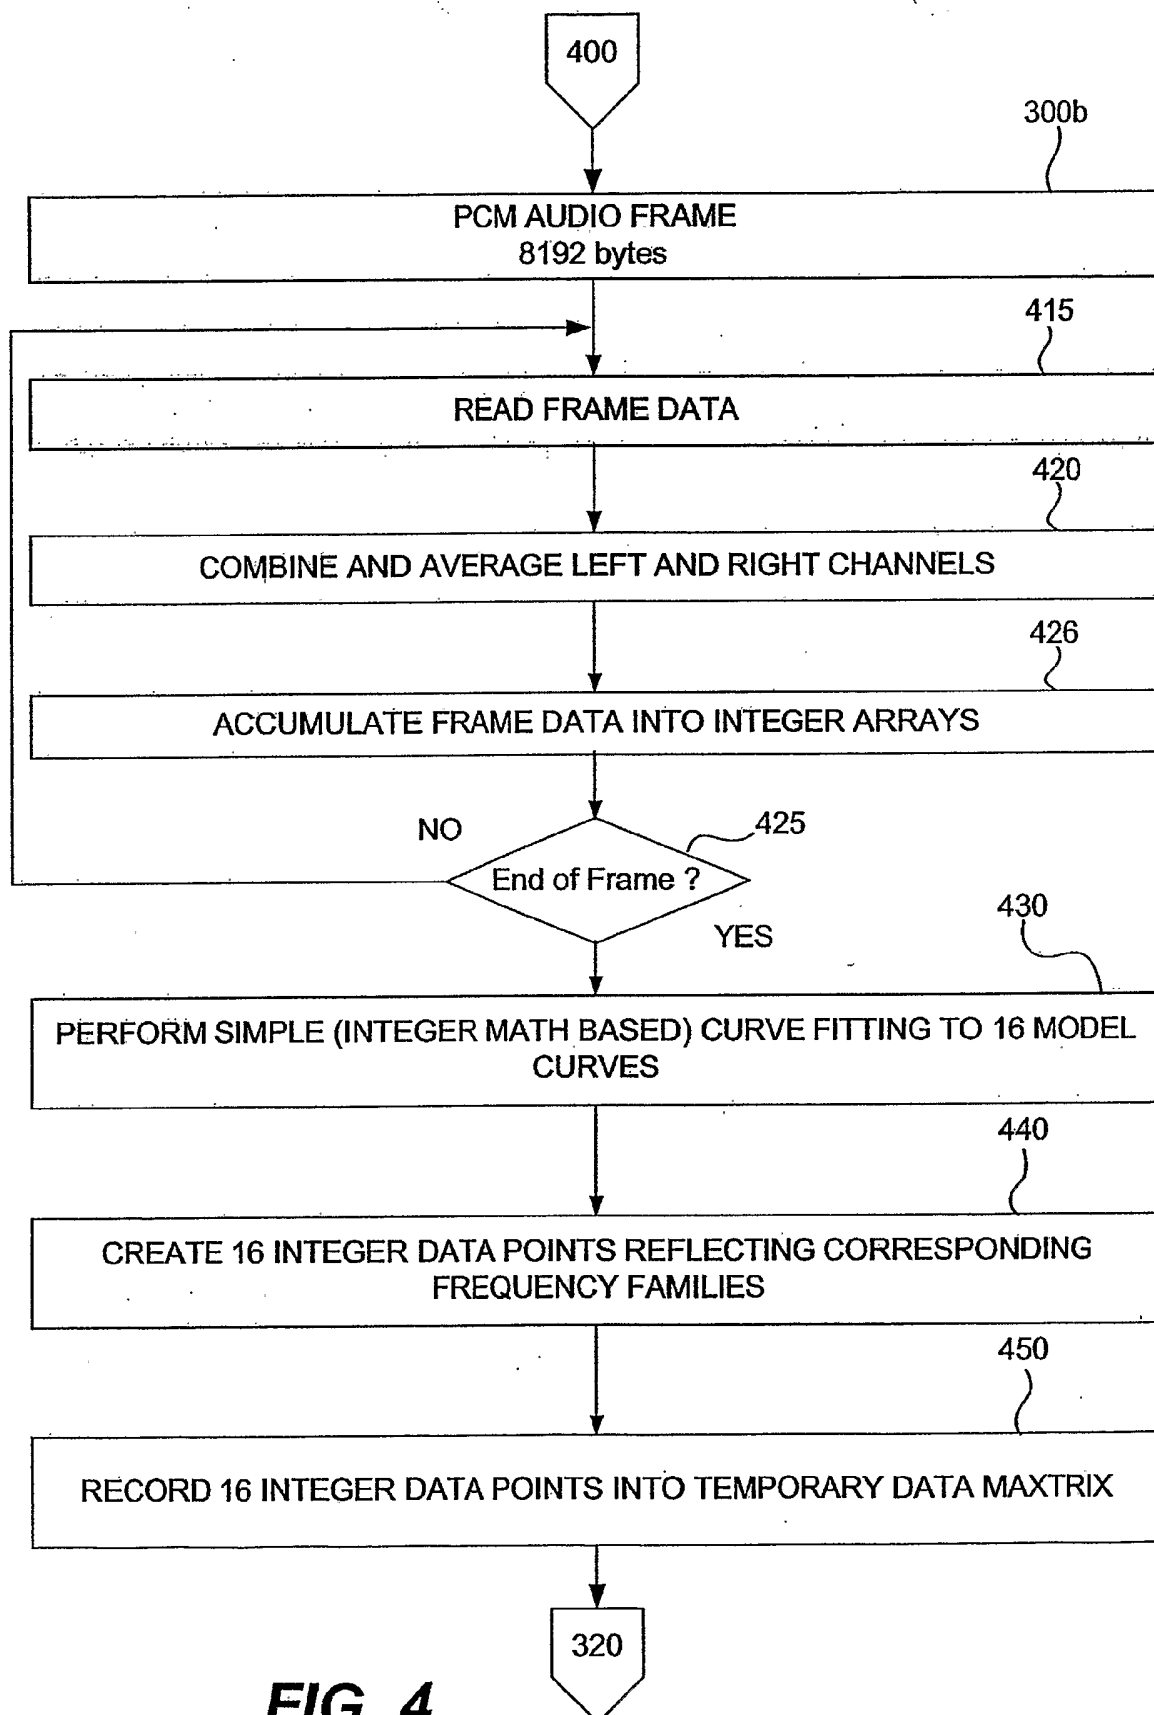
[Claim 82] Computer-readable medium containing code according to Claim 81,
further including code for:
retrieving metadata from the database corresponding to the audio fingerprint.

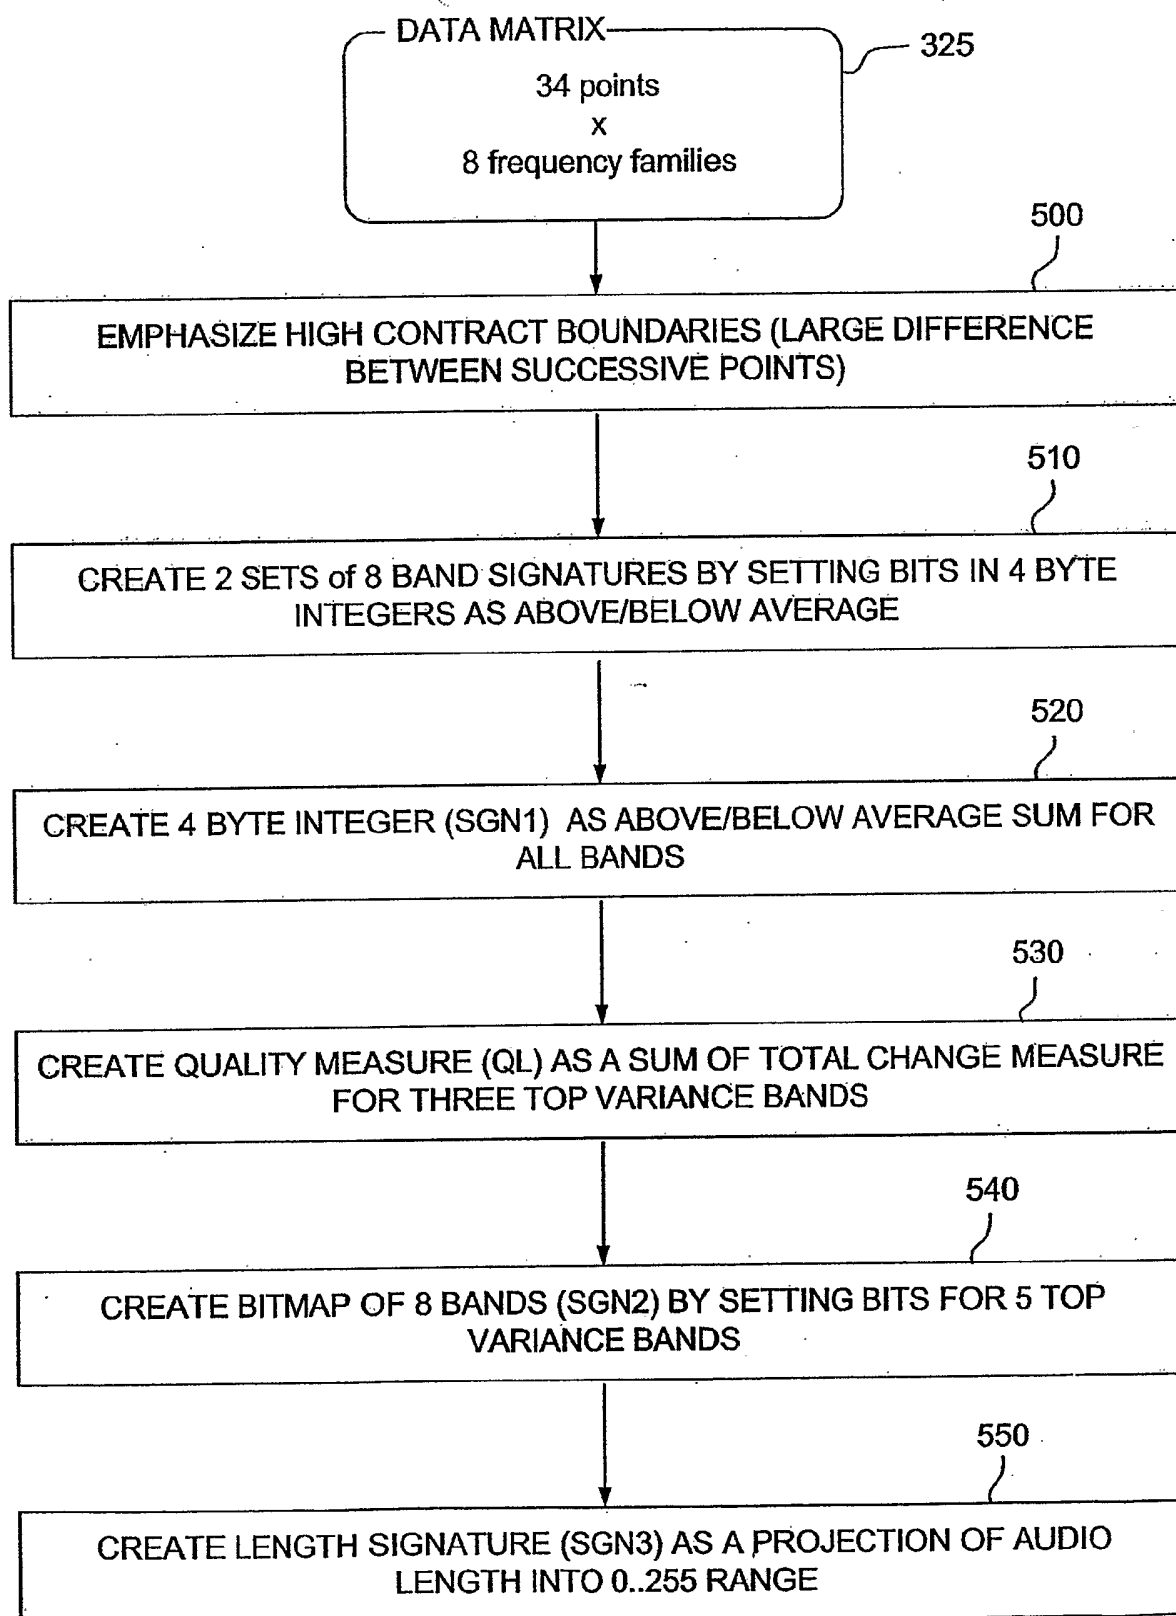
[Claim 83] A network computer system executing the computer-readable medium of
Claim 63.

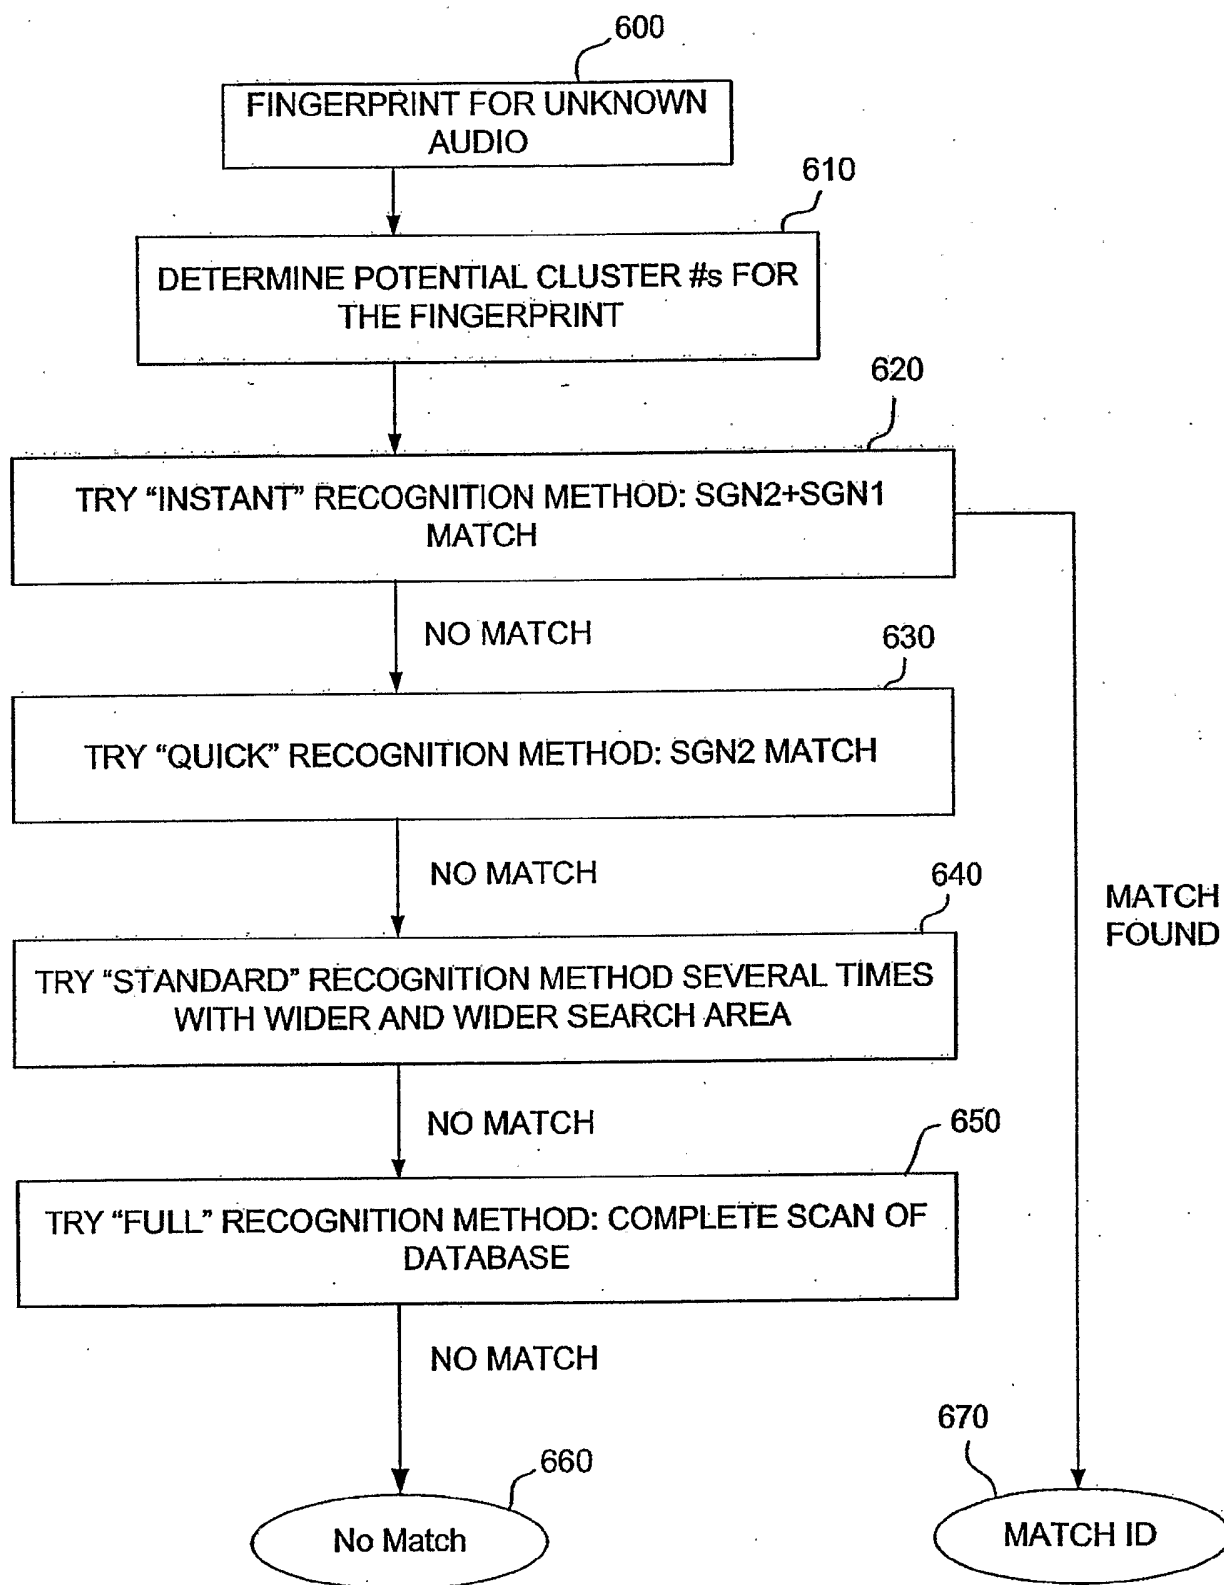
[Claim 84] A device-embedded system executing the computer-readable medium of
Claim 63.

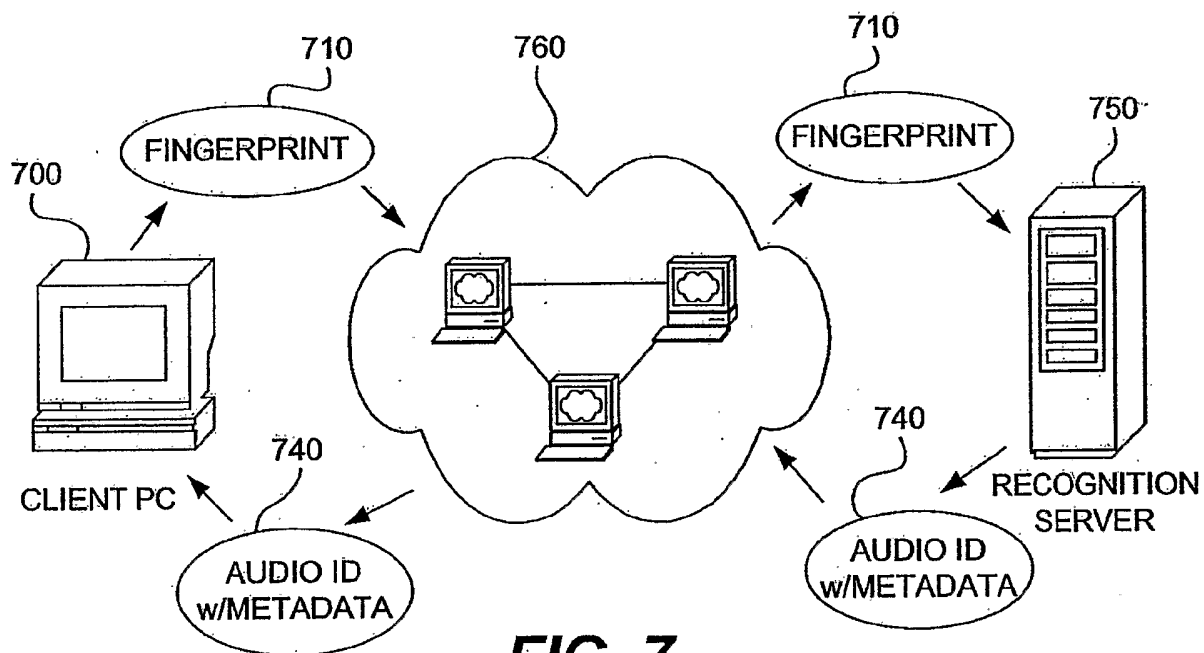
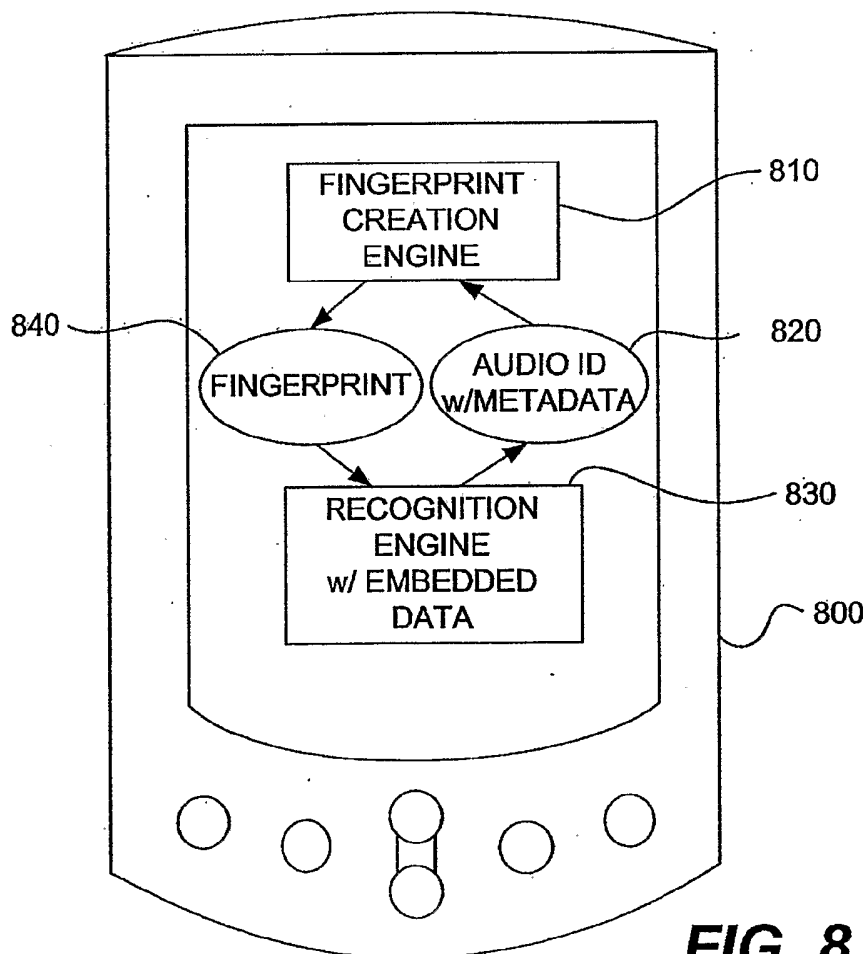
**FIG. 1****FIG. 2**

**FIG. 3**

**FIG. 4**

**FIG. 5**

**FIG. 6**

**FIG. 7****FIG. 8**