



US 20020116195A1

(19) **United States**

(12) **Patent Application Publication**
Pitman et al.

(10) **Pub. No.: US 2002/0116195 A1**

(43) **Pub. Date: Aug. 22, 2002**

(54) **SYSTEM FOR SELLING A PRODUCT
UTILIZING AUDIO CONTENT
IDENTIFICATION**

(60) Provisional application No. 60/245,799, filed on Nov. 3, 2000.

(75) Inventors: **Michael C. Pitman**, Newburgh, NY
(US); **Blake G. Fitch**, White Plains,
NY (US); **Steven Abrams**, New City,
NY (US); **Robert S. Germain**,
Larchmont, NY (US)

Publication Classification

(51) **Int. Cl.⁷** **G10L 13/08**

(52) **U.S. Cl.** **704/260**

Correspondence Address:

**FLEIT, KAIN, GIBBONS,
GUTMAN & BONGINI, P.L.
ONE BOCA COMMERCE CENTER
551 NORTHWEST 77TH STREET, SUITE 111
BOCA RATON, FL 33487 (US)**

(73) Assignee: **INTERNATIONAL BUSINESS
MACHINES CORPORATION,**
ARMONK, NY

(21) Appl. No.: **10/080,097**

(22) Filed: **Feb. 21, 2002**

Related U.S. Application Data

(63) Continuation-in-part of application No. 09/803,298,
filed on Mar. 9, 2001.

(57) **ABSTRACT**

A method is provided for selling products containing or relating to audio content. According to the method, it is determined whether audio identifying information generated for an audio content image matches audio identifying information in an audio content database. If the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, at least one product containing or relating to audio content that corresponds to the matching audio identifying information is identified. In one embodiment, the audio content image is received, and the audio identifying information is generated for the audio content image. In another embodiment, the audio identifying information for the audio content image is received. Also provided is a system for selling products.




FIG. 1B




FIG. 1A




FIG. 2A




FIG. 2B





FIG. 3

FIG. 4A




FIG. 4B




FIG. 5




FIG. 6




FIG. 7





FIG. 8 800



SYSTEM FOR SELLING A PRODUCT UTILIZING AUDIO CONTENT IDENTIFICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation-in-part of prior U.S. application Ser. No. 09/803,298, filed Mar. 9, 2001, now _____, which is based upon and claims priority from prior U.S. Provisional Application No. 60/245,799, filed Nov. 3, 2000. The entire disclosures of application Ser. No. 09/803,298 and Provisional Application No. 60/245,799 are herein incorporated by reference.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention relates to a system for selling a product, and more specifically to systems and methods for selling a product containing audio content in which the audio content is automatically identified.

[0004] 2. Description of Related Art

[0005] Audio content such as music is broadcast to listeners over various mediums that include radio, television, cable television, satellite television, and Internet web sites. When hearing a song from such a broadcast, a listener frequently does not know the artist and/or title of the song. If the listener enjoys the song and desires to purchase a recording of the song or a product containing the song (e. g., an album or video), the inability to identify the song currently playing prevents the listener from making the purchase. This leads to a missed sales opportunity to the detriment of the artist, all copyright holders, the product distributor, and the product retailer.

[0006] To identify a song that is heard, a listener currently must rely on some type of manual review that can only be successful if the chosen reviewer knows the identity of the song. For example, the listener could attempt to sing a portion of the song from memory for another person that the listener believes may know the identity of the song. Alternatively, the listener could record the song (or a portion of it) and play it back for another person to try to identify. Such identification techniques relying on manual review require the listener to find another person that is knowledgeable about such songs. Further, the listener's initial desire to purchase may wane by the time the song can be identified. Thus, there is a need for a system and method for allowing a listener to automatically identify audio content so that a product containing the audio content can be purchased.

[0007] One difficulty in developing a practical system for selling a product that provides automatic audio content identification is providing a mechanism for automatically identifying audio content. The identification of music from any source is not a trivial problem. Different encoding schemes will yield a different bit stream for the same song. Even if the same encoding scheme is used to encode the same song (i.e., sound recording) and create two digital audio files, the files will not necessarily match at the bit level.

[0008] Further, various effects can lead to differentiation of the bit stream even though the resulting sound differences as judged by human perception are negligible. These effects

include: subtle differences in the overall frequency response of the recording system, digital to analog conversion effects, acoustic environmental effects such as reverb, and slight differences in the recording start time. Further, the bit stream that results from a recording of a song will vary depending on the type of audio source. For example, the bitstream for a song created by encoding the output of one stereo receiver will generally not match the bitstream for the same song created by encoding the output of another stereo receiver.

[0009] In addition, there are forms of noise and distortion that are quite audible to humans, but that do not impede our ability to recognize music. FM broadcasts and audio cassettes both have a lower bandwidth than CD recordings, and many of the MP3 files on the Internet are of relatively low quality. Furthermore, some recording systems may alter or distort the music, such as through slight time-stretching or time-compressing. In such cases, not only may the start and stop times be different, but the song duration may be different as well. All such differences may be barely noticeable to humans (if at all), but can foil many identification schemes.

[0010] There is a need for systems and methods that enable a listener to purchase unidentifiable audio content, such as by automatically and effectively identifying the audio content that was heard and then presenting one or more products containing the audio content.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIGS. 1A and 1B are schematics of an exemplary system on which a preferred embodiment of the present invention can be implemented.

[0012] FIGS. 2A and 2B are block diagrams of systems for selling a product according to two embodiments of the present invention.

[0013] FIG. 3 is a flow diagram of a process for selling a product containing audio content according to a preferred embodiment of the present invention.

[0014] FIGS. 4A and 4B are a flow diagram of a process for identifying events from an audio segment in an exemplary embodiment of the present invention.

[0015] FIG. 5 is a flow diagram of a process for generating keys from the events produced by the process shown in FIGS. 4A and 4B.

[0016] FIG. 6 is a flow diagram of a process for generating keys from the content of a key generator buffer in an exemplary embodiment of the present invention.

[0017] FIG. 7 is a flow diagram of a process for filtering percussion events in an exemplary embodiment of the present invention.

[0018] FIG. 8 is a flow diagram of a process for using keys to compare two audio segments in an exemplary embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0019] Objects, features, and advantages of the present invention will become apparent from the following detailed description. It should be understood, however, that the detailed description and specific examples, while indicating

preferred embodiments of the present invention, are given by way of illustration only and various modifications may naturally be performed without deviating from the present invention.

[0020] FIGS. 1A and 1B are schematics of an exemplary system on which the present invention can be implemented. The system includes a content source 102 and a content recorder 104 that are coupled together through a first communications link 106. Additionally, the system includes a content player 112 and an audio identification computer (AIC) 120 that are coupled together through a second communications link 122. A computer readable memory medium 124 such as a CD-ROM is provided for loading software onto the AIC 120 for carrying out methods such as those described in detail below. For example, the software can operate to automatically identify audio content from an audio content image supplied by the content player 112, and then to display one or more products containing the identified audio content. In some embodiments, there is additionally or instead displayed one or more products relating to the identified audio content (e.g., a biography or poster of the artist who recorded the identified song).

[0021] The content source 102 can be any source that supplies audio content, either alone or as part of a multimedia presentation. The content recorder 104 can be any device that can receive and record an image of the audio content from the content source 102, and the content player 112 can be any device that is capable of outputting the audio content image recorded by the content recorder 104 in a suitable form. The communications links 106 and 122 can be any mechanisms that allow the transfer of audio content in a format suitable for the devices (e.g., sound waves traveling through air, an analog signal traveling through a radio frequency link or local cable connection, or a digital datastream traveling through a local cable connection or long-distance network connection).

[0022] The AIC 120 can be any conventional computer system such as an IBM PC-compatible computer. As is known, an IBM PC-compatible computer can include a microprocessor, basic input/output system read-only memory (BIOS ROM), random access memory (RAM), hard disk drive storage, removable computer readable medium storage (e.g., a CD-ROM drive), a video display adapter card, a video monitor, a network interface (e.g., modem), a keyboard, a pointing device (e.g., mouse), a sound card, and speakers. The AIC 120 is loaded with an operating system (such as Windows or UNIX) and an AIC application, such as one that operates as described below.

[0023] For example, in one embodiment, the content source 102 is an FM radio station, the content recorder 104 is a stereo having an FM receiver and a cassette deck, the content player 112 is a portable cassette player, and the AIC 120 is a computer-powered kiosk at a record store. Audio content received from the radio station over a radio frequency link is recorded onto a cassette by the stereo. The cassette is then played by the portable cassette player into a microphone of the kiosk. Alternatively, the kiosk could include a cassette player for playing the cassette.

[0024] In another embodiment, a recorded video playback device (e.g., VCR or DVD player) is the content source 102, a personal computer is both the content recorder 104 and content player 112, and a remote web server is the AIC 120.

Audio content included in the recorded video (e.g., a song from a movie or background music from a television show) is supplied to the sound card of the personal computer over a local cable. An image of the audio content is stored on the computer as an audio file (e.g., in MP3 format), and then the computer system uploads the audio file to the web server of a retailer's web site over an Internet connection. Further embodiments include many different combinations of content sources, recorders, and players that enable some type of image of the audio content of interest to be stored and then transferred to the AIC in the proper format.

[0025] FIG. 2A is a block diagram of software for a system for selling a product according to one embodiment of the present invention. An audio content image supplied by the content player is coupled through communication link 208 to an input interface 210. The audio content image can be supplied in any format including an analog signal that is modulated by the audio content, a digital datastream, or a computer file. The input interface 210 receives the audio content image and outputs an audio datastream in the same or a different format (e.g., digital data in MIDI, WAV, or MP3 format) for use by an audio identification and product presentation application 212.

[0026] The audio identification and product presentation application 212 includes a content identifier module 218. The content identifier module 218 receives the audio datastream from the input interface 210 and decodes it to obtain information. In preferred embodiments, the content identifier module includes a key database 216 that contains numerous keys derived from numerous songs. For each song (i.e., sound recording unit) in the database, there is a set of keys. The set of keys provides a means of identifying a song or a segment of a song. A section of a song will have a corresponding subset of keys that allow the section to be identified, thus retaining the ability to identify the presence of only a segment of a song.

[0027] In one exemplary embodiment, the key database takes the form of a key table. Each row of the keytable includes a key sequence in a first column, a time group value in a second column, and a song ID (e.g., title) in a third column. The key database is constructed by applying a key generation program to known songs, and associating each key obtained by the key generation program with the title in the key database. The time group is a time (measured in units of a predetermined interval relative to the start of the song) at which an audio feature from which a key is derived occurred. Each time group includes events (explained further below) that occurred during a period of time equal to the predetermined interval. Two different recordings of the same song can have slightly different start times (e.g., the recordings may start a few seconds before or after the actual beginning of the song) or can include only segments of the song. Such variances present a difficulty that preferred audio identification methods surmount.

[0028] FIG. 2B is a block diagram of software for a system for selling a product according to another embodiment of the present invention. In this embodiment, the audio identifying information is coupled through communication link 208 to the input interface 210. For example, the content player can contain an identifying information generator and can supply the generated identifying information for the audio content, instead of the audio content itself. The input

interface **210** supplies the identifying information to a product presentation application **212** that includes a content identifier module **218**.

[0029] In one exemplary embodiment, a user's computer system generates audio content identifying information for audio content or an audio content image, and then uploads the generated audio content identifying information (as opposed to the audio content image itself) to a retailer's system (e.g., web site or kiosk). In such an embodiment, the software for generating the audio content identifying information is executed on the user's computer system. Thus, the identifying information generating function and the matching function can be separated, and these functions can be performed by the same or different computer systems.

[0030] FIG. 3 is a flow diagram of a process **300** that is performed by the audio identification and product presentation application to automatically identify audio content for selling purposes according to a preferred embodiment of the present invention. In step **302**, audio content from products available for purchase (e.g., songs from albums or concert videos) is registered along with identifying information in an audio content database. In preferred embodiments, the identifying information is produced by a feature generation system and is based on the audio content itself. For example, the process described in detail below can be used to generate a unique signature (in the form of a set of keys) for each piece of audio content that is to be searched for.

[0031] In further embodiments, the identifying information can be any other type of audio signature or audio fingerprint that is computed based on the audio data itself, an identifying watermark, embedded identification information such as an identification number, or any other type of information that allows individual pieces of audio content to be uniquely (or substantially uniquely) identified. Further, the audio content that is registered can be a single selection from each product or multiple selections (e.g., all songs on an album or in a movie).

[0032] In step **304**, a listener hearing audio content records an image of at least a portion of the audio content. Then, in step **305**, the listener supplies the audio content image to a retailer's audio identification system. As described above with respect to FIG. 1, the steps of recording an image and supplying it to the retailer's system can be achieved in any available manner depending on the practicalities of a particular application. For example, the listener can record an audio content image on a machine-readable medium, and then insert the medium into a kiosk at the retailer's store. Alternatively, the listener can record an audio content image on a computer as an audio file (or transfer a previously-recorded image to such an audio file format), and then upload the audio file to the retailer's web site. In yet another embodiment, the user generates the audio content identifying information and then supplies the generated audio content identifying information (as opposed to the audio content image itself) to the retailer's system. For example, the user could upload the identifying information to the retailer's system.

[0033] In step **306**, the content identifier module **218** analyzes the audio content image that was supplied in step **304** using an algorithm that generates audio content identifying information like the information stored in the audio content database. For example, in preferred embodiments, a

unique feature signature (in the form of a set of keys) is generated for each audio content image that is supplied. Next, in step **308**, the unique signature computed for the audio content image (or audio content image portion) is compared with the information stored in the audio content database. The content identifier module **218** determines whether or not the audio content image matches any of the audio content in the audio content database. For example, a best match algorithm employing some minimum threshold could be used to make such a determination. One exemplary match determination algorithm is discussed in detail below.

[0034] In preferred embodiments, the content identifier module can differentiate among different recordings of a song. There can be several recordings of the same song, including different recordings by the same artist and recordings by different artists. While conventional techniques such as manual review may result in the identification of a different recording of the song than the one heard by a listener, the content identifier module of preferred embodiments is able to use the unique signatures to differentiate among the different recordings of the song and identify the recording that was heard by the listener.

[0035] In step **310**, if the audio content image is determined to match audio content in the database, the retailer's system displays an identification of the matching audio content (e.g., the artist and title) for the listener, along with one or more products containing (or relating to) the matching audio content. For example, if the audio content image is a portion of a song, the list displayed to the listener could include products such as a cassette single of the song, an album containing the song, a music video for the song, a concert video including a live performance of the song, a movie in which the song appears, and a movie soundtrack including the song.

[0036] Preferably, the user can then select one of more of the products on the list and purchase or order these selections directly from the retailer. The operator of the audio identification system could instead (or additionally) charge a fee for identifying the audio content from the audio content image supplied by the listener. Further, the audio identification system could instead (or additionally) create a product containing the matching audio content (e.g., a cassette, CD, or digital audio file containing a recording of the matching song) for purchase by the listener. In some embodiments, the listener is asked to verify the audio content of interest by selecting it from among one or more audio content titles in the database that match or most closely match the audio content image that was supplied by the listener.

[0037] FIGS. 4A through 8 illustrate in detail relevant portions of one exemplary process for identifying audio content. FIG. 4A is a first part of a flow diagram of a process **400** for generating an event stream for an audio segment (e.g., song) according to a preferred embodiment of the present invention. The process **400** accepts an audio signal as input and outputs a sequence of "events". In some embodiments of the present invention, the audio signal is reproduced from an MP3 file. In step **402**, an audio signal is sampled. In one embodiment, the audio signal is sampled at about 22050 Hz or lower. This allows frequency components up to 11 KHz to be accurately determined. It is advantageous to use an upper frequency limit of about 11 KHz because 11

KHz is about the frequency cutoff for FM broadcast radio, and it is desirable to be able to generate the same set of keys for a song regardless of whether the song recording was at one point transmitted through FM radio or obtained directly from a high quality source (e.g., a CD).

[0038] In step 404, for each successive test period the spectrum of the audio signal is computed. The duration of the test period preferably ranges from about $\frac{1}{43}$ of a second to about $\frac{1}{10.75}$ of a second, and more preferably the test period is about $\frac{1}{21.5}$ of a second. The spectrum of the audio signal is preferably analyzed using a fast Fourier transform (FFT) algorithm. The accuracy of spectrum information obtained using an FFT algorithm can be improved by averaging together the results obtained by applying the FFT to several successive periods (sample sets). In preferred embodiments of the present invention, spectrum information is improved by averaging together the results obtained by applying the FFT to two or more successive periods, and preferably 3 or more successive periods, and even more preferably 4 successive periods. According to one exemplary embodiment of the present invention, the spectrum associated with a given test period having a duration of $\frac{1}{21.5}$ of a second is obtained by sampling an audio signal at a rate of 22050 Hz and averaging together the results obtained by applying an FFT algorithm to four successive periods, each of which has a duration of $\frac{1}{86}$ seconds and includes 2048 samples.

[0039] Step 404 can be accomplished by using an FFT algorithm run on the microprocessor of the AIC 120. Alternatively, the AIC could be provided with FFT hardware for performing step 404. Other spectrum analyzers, such as a filter bank, can alternatively be used for carrying out step 404. Additionally, in process 404, successive sets of samples can alternatively be projected onto another type of basis besides a Fourier basis. One particular alternative to the Fourier basis is a wavelet basis. Like Fourier basis functions, wavelets are also localized in the frequency domain (although to a lesser degree). Wavelets have the added property that they are localized in the time domain as well. This opens up the possibility of projecting the audio signal as a whole, rather than successive sample sets of the audio signal onto a wavelet basis, and obtaining time dependent frequency information about the signal.

[0040] One common set of frequencies used in composing music are the notes of the even-tempered scale. The even tempered scale includes notes that are equally spaced on a logarithmic scale. Each note covers a frequency band called a "semitone". The inventors have determined that improved signatures can be obtained by collecting spectral power in discrete semitone bands as opposed to the evenly spaced frequency bands output by an FFT algorithm. In step 406, the spectrum information (e.g., Fourier frequency components) obtained in step 404 are collected into a number of semitone frequency bands or channels.

[0041] In step 408, a first average of the power in each semitone frequency channel is taken over the last T1 seconds. In step 410, a second average of the power in each semitone frequency channel is taken over the last T2 seconds, where T2 is greater than T1. T1 is preferably from about $\frac{1}{10}$ to about 1 second. T2 is preferably larger than T1 by a factor of from 2 to 8. According to a one exemplary embodiment of the present invention, T2 is equal to one

second, and T1 is equal to one-quarter of a second. The "events" mentioned above occur when the value of the first average crosses the second average.

[0042] In step 412, the values of the first and second averages are recorded for each semitone channel. Recording is done so that it can be determined during the following test period whether the first average crossed the second average. In step 414, for each semitone channel it is determined if the first average crossed the second average. This is done by comparing the inequality relation between the first and second averages during the current test period to the inequality relation for the last period. Although comparison between only two averages has been discussed above, it is possible according to alternative embodiments of the present invention to use more than two averages, and identify events as the crossing points between different sub-combinations of the more than two averages.

[0043] In the vicinity of an extremum (local maximum or minimum) in a semitone frequency channel, the two averages will cross. Rather than looking for the crossing point of two running averages with different averaging periods, another type of peak detector (e.g., an electronic circuit) could be used. Such could advantageously be used in combination with an FFT in an implementation of the present invention that is implemented predominately in hardware, as opposed to software.

[0044] Rather than looking for a peak in the signal in a frequency channel, another type of curve characteristic such as an inflection point could be used as a trigger event. An inflection point can be found by calculating a second derivative of a frequency channel by operating on three successive values of the power in a given frequency channel, and identifying a time at which the second derivative changes from positive to negative or vice versa. The second derivative can be approximated using function (time dependent frequency component) values for three successive points in the following formula.

$$(F(N+2)-2F(N+1)+F(N))/\Delta T$$

[0045] where F(I) is the value of the function at the i^{th} time (e.g., at the i^{th} test period), and ΔT is the interval between successive function values (e.g., the duration of the test period).

[0046] At an extremum of a time dependent frequency component, its first derivative is equal to zero. At an inflection point of a time dependent frequency component, its second derivative is equal to zero. Extrema and inflection points are both kinds of events. More generally events can be defined as points (i.e., points in time) at which an equation involving a time dependent frequency component derivative of one or more orders of the time dependent frequency components, and/or integrals involving the time dependent frequency components is satisfied. To allow their use in identifying different audio content, an essential part of the definition of "events" is that they occur at a subset of test periods, not at each test period.

[0047] Step 416 is a decision block, the outcome of which depends on whether averages for a semitone channel crossed. Step 416 is tested for each semitone channel. If averages for a semitone channel were not crossed during the current test period, then in step 418 it is determined if the audio signal is over. If the audio stream is finished, then the

process 400 terminates. If the audio signal is not finished, then the process 400 is advanced to the next test period and the process continues with step 404. If on the other hand, averages did cross during the last test period then the process 400 continues with step 422 in which each event is assigned to the current time group and information related to the average crossing event is generated.

[0048] Event information preferably includes the time group for the event, the test period for the event, the semitone frequency band of the event, and the value of the fast average (average over T1) at the time of crossing. Event information can be recorded in a memory or storage device associated with the AIC. Each time group covers a period of time that is longer than a test period, and preferably time groups cover successive periods of time equal to from ¼ to 2 seconds, and more preferably each time group covers a period of from one-half to three-quarters of a second. Grouping events into successive time groups has the advantage that keys obtained by processing two recordings of the same song will tend to match more completely despite the fact that one or both of the recordings may have some distortions (e.g., distortions that arise in the course of recording on magnetic tape).

[0049] In step 424, the process 400 is incremented to the next test period. In step 426, it is determined if the audio segment (e.g., song) is finished. If the audio segment is finished then the process 400 terminates. If the audio segment is not finished, then the test period is incremented and the process loops back to step 404.

[0050] Thus, the result of the process is to take an audio signal and produce a plurality of events. Each event is assigned to a semitone frequency band in which it occurred and a time group (interval) within which it occurred. The events can be stored in a memory (e.g., RAM in the AIC 120). The events can be stored in a buffer from which they are successively read by one or more key generator processes. The events output by the process could be in the form of an event stream, which is to say that after each time group, all the events occurring within the time group could be written to memory and thereby made available for further processing. An alternative is to write all the events for a song to memory or storage at one time.

[0051] FIG. 5 is a flow diagram of a key generator process for generating keys from the events produced by a process such as that shown in FIGS. 4A and 4B. The events output by process 400 are processed by a plurality of key generator processes 500. Each of the plurality of key generator processes is assigned to one semitone frequency band that is designated as its main frequency. However, each key generator also uses events that occur in other semitone frequency bands near its main frequency. Preferably each key generator monitors from 5 to 15 semitone frequency bands. If the number of frequency bands monitored is too few, the resulting keys will not be as strongly characteristic of the particular audio segment. On the other hand, a higher number of frequency bands will result in higher computational expense for computing and comparing keys, greater memory requirements for storing keys, and potential performance loss due to key saturation in the key table from the increased number of keys. According to one embodiment of the present invention, each key generator monitors its main

semitone frequency band and four other semitone frequency bands, two on each side of the main semitone frequency band.

[0052] Referring now to FIG. 5, in step 502 each successive time group of events output by process 400 is monitored for events occurring within the semitone frequency bands assigned to this key generator. Step 504 is a decision block, the outcome of which depends on whether the key generator detected (e.g., by reading from memory) any new events in step 502. If not, then in step 514, the process 500 is incremented to the next time group and loops back to step 502. If, on the other hand, new events did occur in the time group and semitone frequency bands checked, then in step 506 the new events are written to a key generator buffer for the key generator under consideration, and the events for the oldest time group that were stored in the key generator buffer are deleted. In one exemplary embodiment, the buffer can be seen as an array in which the rows correspond to time groups and the columns to frequency bands. Thus, in the embodiment of the present invention mentioned above, there would be five columns for each of the semitone frequency bands monitored by each key generator.

[0053] The key generator buffer preferably includes events from 3 to 7 time groups. More preferably, events from five or six time groups are maintained in each key buffer array. Note that in this embodiment not all time groups are represented in the key generator buffer. As shown in FIG. 5, if no events occur in the semitone frequency bands for a key generator in a certain time group, then no change will be made to the key generator buffer. In other words, a blank row will not be recorded. Therefore, each time group recorded in the key generator buffer includes at least one event.

[0054] Step 508 is decision block whose outcome depends on whether an event that occurred in the current time group (e.g., current pass through program loop) is a trigger event. According to a preferred embodiment of the present invention, a trigger event is an event that occurs at the main frequency assigned to this key generator. If a trigger event did not occur, then the process loops back to step 514. If a trigger event did occur, then the process continues with step 510 in which keys are generated from the contents of the key generator buffer. The process 500 continues until all of the events produced by process 400 have been processed.

[0055] FIG. 6 is a flow diagram of a process for generating keys from the contents of a key generator buffer according to one embodiment of the present invention. In particular, the process 600 shows in detail one embodiment of the implementation of step 510 of FIG. 5. In step 602, for each key generator (as explained above there are a plurality of key generators carrying out process 500) and for each trigger event for the key generator under consideration, one or more different combinations of events from the key generator buffer are selected. Each combination includes only one event from each time group. (There may be more than one event for each time group in each key generator buffer.) According to a preferred embodiment of the present invention, not all possible combinations are selected, rather only combinations for which a power associated with each event changes monotonically from one event to the next in the combination are selected.

[0056] In this embodiment, the order of events within a combination corresponds to the time group order. The power associated with each event is preferably the magnitude of the fast (first) average at the test period at which the event occurred. In this embodiment, less than all of the possible combinations of keys will be taken, so that the total number of keys for a given audio segment will tend to be reduced which leads to lower memory and processing power requirements. On the other hand, there will be enough keys that the identity of the song will be well characterized by (i.e., strongly correlated to) the set of keys generated from the song. According to an alternative embodiment, only a single combination is selected from the contents of the key generator buffer. The single combination includes the event associated with the highest fast average power from each time group. According to another alternative embodiment, all the different combinations of events taking one event from each time group are taken.

[0057] In step 604, for each selected combination of events a key sequence is composed that includes a sequence of numerical values of frequency offsets (relative to the main key generator frequency) for the sequence of events from each combination formed in step 602. Each frequency offset is the difference between the frequency of the semitone band in which the event occurred and the main frequency of the key generator. In step 606, test period information (e.g., a sequence number for the test period of the trigger event, where the sequence number for the first test period for each song is designated by the number one) for the trigger event is associated with the key sequence.

[0058] In step 608, the key which includes the key sequence and the test period information is associated with a song (or other audio) identifier or ID (e.g., title). Process 600 includes step 608 in the case that known songs are being used to construct a song database against which unknown songs will be compared. In comparing two songs, both the key sequence and test period information will be used, as described further below with reference to FIG. 8. The song database can take the form of a table including three columns and a plurality of rows. The first column includes key sequences, the next column includes corresponding test periods associated with the key sequences, and the final column includes an identification of the song from which the keys in the row were obtained.

[0059] While the processes described above can be used to identify audio content, it is advantageous to filter percussion events. More specifically, percussion sounds in a song, if not filtered, typically account for high percentage of the events output by process 400. In the interest of saving computer resources (e.g., memory and processing power) and obtaining a more characteristic set of keys, it is desirable to reduce the number of percussion events such as by eliminating some percussion events before events are processed by the key generator process 500. It has been recognized by the inventors that percussion sounds lead to events being triggered during the same test period in adjacent semitone frequency bands. For example, percussion sounds can lead to events occurring in a sequence of 2 or more adjacent semitone frequency bands.

[0060] FIG. 7 is a flow diagram of a process used in a preferred embodiment to filter percussion events from the events produced by the process of FIGS. 4A and 4B. In step

702, for each successive test period it is determined if multiple events occurred in a sequence of two or more adjacent semitone frequency bands. Step 704 is a decision block, the outcome of which depends on whether multiple events in adjacent frequency bands occurred. A threshold of some predetermined number of events occurring in adjacent frequency bands used in the process. Preferably, a lower limit on the number of adjacent frequency bands in which events must be found (in order to consider that the events were produced by a percussion sound) is set at three or more. According to an exemplary embodiment of the present invention, events must occur in three successive semitone frequency bands for the outcome of step 704 to be positive.

[0061] If the outcome of step 704 is negative, then the process continues with step 708 in which the process increments to the next test period and loops back to step 702. If, on the other hand, the outcome of step 704 is positive, then the process 700 continues with step 706 in which each sequence of events that occurred during the same test period in adjacent frequency bands is pared down to a single event. All of the events except for the event in the sequence that has the highest fast average value are deleted from the event stream produced by process 400. Alternatively, instead of deleting all but one, up to a certain predetermined number of events can be retained.

[0062] The processes described above produce keys for a sound recording based on the features (i.e., events) contained in the sound recording. Thus, the processes can be run on known audio content to construct a feature database of the known audio content during a storage phase. After the database is created, during a retrieval phase the above processes can be used to extract features from unknown audio content and then the database can be accessed to identify the audio content based on the features that are extracted. For example, the same processes can be run on the unknown audio content to extract features in real time (or even faster), and then the audio content is identified with the best match in the database. In one embodiment, a best match can be reported for each predetermined interval (e.g., 10 to 30 seconds) of the audio content.

[0063] FIG. 8 is a flow diagram of a song identification process that uses the keys generated in the processes of FIGS. 5 and 6 to identify an audio segments. A song database (such as that described above) is used to identify an unknown song such as a song downloaded from a web site in step 804 of process 300. The key sequence field (column) of the song database can be used as a database key. The records (rows) of the song database are preferably stored in a hash table for direct lookup. The identification process 800 is an exemplary implementation of step 308 of FIG. 3.

[0064] In step 802, keys are generated from a song to be identified (for example, by carrying out the processes shown in FIGS. 5 and 6). In step 804, each key in the set of keys generated in step 802 is looked up in a song database that includes keys for a plurality of songs. The key sequence part (as opposed to the test period part) of each key is used as a database key. In other words, the song database is searched for any entries that have the same key sequence as a key sequence belonging to a key obtained from the song to be identified. More than one key in the song database can have the same key sequence, and furthermore by happenstance more than one song in the song database can share the same

key sequence. In step **806**, for each key in the database that matched (by key sequence) one or more keys in the song database, an offset is calculated by taking the difference between a test period associated with the key being looked up and a test period associated with each matching key in the song database.

[**0065**] In step **808**, the offsets are collected into offset time groups. The offset time groups for the offsets are distinct from the time groups used in key generation. According to a preferred embodiment, an offset time group will be equal to from 2 to 10 test periods. By way of illustration, if each offset time group were 5, then any pair of keys for which the difference determined in step **806** was between 0 and 5 would be assigned to a first offset time group, and any pair of keys for which the difference was between 6 and 10 would be assigned to a second offset time group. According to an exemplary embodiment of the present invention, each offset time group is equal to 5 test periods.

[**0066**] In step **810**, for each song that has keys that match keys in the song to be identified, and for each offset time group value that was determined in step **808** and involved keys for a given song in the song database, a count is made of the number of matching keys that had the same time group offset value. One can visualize step **810** in the following way, which may also be used as a basis for an implementation approach. A temporary table is constructed where each row corresponds to a song from the song database that had one or more key matches with the song to be identified. The first column includes names of the songs. In the second column, adjacent each song name there is a value of the offset time group that was found between keys found for the named song in the song database, and matching keys from the song to be identified. After completing step **810**, the third column will include counts of the number of key matches corresponding to a particular song identified in the first column, that had the same offset time group as identified in the second column. The table might appear as follows.

TABLE 1

SONG TITLE	OFFSET VALUE (UNITS OF TIME GROUP INTERVAL)	COUNT OF KEY SEQUENCE MATCHES FOR THIS SONG AND WITH THIS OFFSET VALUE
Title1	3	1
Title1	4	1
Title2	2	2
Title2	3	107
Title3	5	1
Title2	8	1

[**0067**] If the song to be identified is in the database, then one particular time group offset value will accumulate a high count. In other words, a high number of matching pairs of keys will be found to have some particular value of offset time group. In the example above, the song entitled Title2 has a count of 107 for an offset time group of 3. For example, the time group offset may arise because the specific recording that is being identified started a few seconds after the recording of the song used to generate keys for the song database, or because a small segment of the song is being identified.

[**0068**] In step **812**, the song from the song database that has the highest count of matching keys with the same offset is identified. In decision block **814**, the count is compared to a threshold value. The threshold can be set based on the particular application or through a determination of the minimum value for the highest counts that are found when songs actually match, and the maximum value of the highest counts when songs tested do not match any songs in the database. The value of the threshold used also depends on the specific embodiment chosen for step **602** discussed above, as this determines the total number of keys.

[**0069**] Rather than comparing the count to a threshold, it is possible instead in step **812** to compare a threshold with the ratio of the highest count to the total number of keys generated from the song to be identified. Another alternative is to compare a threshold with the ratio of the highest count to the average of the remaining counts. These latter two alternatives can also be viewed as comparing the highest count to a threshold, although in these cases the threshold is not fixed. If, as would be the case when the song to be identified is not in the database, the count does not meet the threshold criteria, then the song identification process **800** terminates. Additional steps maybe provided for reporting (e.g., to a user) that the song to be identified could not be identified. If on the other hand the count does meet the threshold criteria, then in step **814** information identifying the song that had the highest count (which met the threshold criteria) is output. In further embodiments, the processes of the present invention are used to identify segments of songs.

[**0070**] The process for identifying songs described above with reference to **FIGS. 4A through 8** is robust in terms of its ability to handle distortion and alteration. Furthermore, the process is also efficient in terms of computational complexity and memory requirements. The processes for generating an event stream, filtering percussion events, generating keys, and looking up the keys in a song database can also be conducted in real time (or faster). The computational expense of the process is low enough to allow it to run in real time on a common personal computer.

[**0071**] Accordingly, the present invention provides systems and methods in which audio content is automatically identified and then one or more products containing (or relating to) the identified audio content are presented for purchase. Thus, a prospective buyer can quickly and easily purchase audio content that is heard but whose identity is not known. This has the potential to dramatically increase sales.

[**0072**] The above description frequently describes a process and system for identifying musical content from a radio broadcast. However, this is for illustrative purposes only and the present invention is not so limited. The present invention can similarly be applied to audio content derived from any source (including a television broadcast, web broadcast, recorded video, film, and even live performance), because the initial source of the audio content is unimportant to the system of the present invention.

[**0073**] Further, while the embodiments of the present invention described above relate to musical content, the system of the present invention could easily be adapted by one of ordinary skill in the art to automatically identify any other type of unidentifiable audio content that was recorded by a user. In such further embodiments, an image of the unidentifiable audio content is recorded, the recorded image

is supplied to the audio identification system, identifying information is generated for the supplied image, and then the generated identification information is compared with a database of identifying information for content of interest. For example, in one embodiment, an audio content database for the sounds of different birds (or other biological sounds) is created. An avid bird watcher can then record an image of the sounds made by an unknown type of bird, and then use the audio identification system to identify the type of bird that was heard. In such a case, the system could also be used to identify products relating to the recorded image (e.g., a book on the type of bird identified or some other product relating to that type of bird).

[0074] The present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system—or other apparatus adapted for carrying out the methods described herein—is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein.

[0075] The present invention can also be embedded in a computer program product, which includes all the features enabling the implementation of the methods described herein, and which—when loaded in a computer system—is able to carry out these methods. Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or, notation; and b) reproduction in a different material form. Each computer system may include, inter alia, one or more computers and at least a computer readable medium allowing a computer to read data, instructions, messages or message packets, and other computer readable information from the computer readable medium. The computer readable medium may include non-volatile memory, such as ROM, Flash memory, disk drive memory, CD-ROM, and other permanent storage. Additionally, a computer medium may include, for example, volatile storage such as RAM, buffers, cache memory, and network circuits. Furthermore, the computer readable medium may include computer readable information in a transitory state medium such as a network link and/or a network interface, including a wired network or a wireless network, that allow a computer to read such computer readable information.

[0076] While there has been illustrated and described what are presently considered to be the preferred embodiments of the present invention, it will be understood by those skilled in the art that various other modifications may be made, and equivalents may be substituted, without departing from the true scope of the invention. Additionally, many modifications may be made to adapt a particular situation to the teachings of the present invention without departing from the central inventive concept described herein. Therefore, it is intended that the present invention not be limited to the particular embodiments disclosed, but that the invention include all embodiments falling within the scope of the appended claims.

What is claimed is:

1. A method for selling products containing or relating to audio content, said method comprising the steps of:

receiving a recorded audio content image;
generating audio identifying information for the audio content image;

determining whether the audio identifying information generated for the audio content image matches audio identifying information in an audio content database; and

if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, charging a fee for at least one product containing or relating to audio content that corresponds to the matching audio identifying information.

2. The method according to claim 1,

wherein the receiving step includes the sub-step of receiving a transmitted audio content image of at least a portion of a song, and

the charging step includes the sub-step of transmitting a recording of at least a song that corresponds to the matching audio identifying information.

3. A method for selling products containing audio content, said method comprising the steps of:

receiving a recorded audio content image;
generating audio identifying information for the audio content image;

determining whether the audio identifying information generated for the audio content image matches audio identifying information in an audio content database; and

if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, generating a product containing audio content that corresponds to the matching audio identifying information.

4. The method according to claim 3, further comprising the step of charging a fee for the generated product.

5. A method for selling products containing or relating to audio content, said method comprising the steps of:

determining whether audio identifying information generated for an audio content image matches audio identifying information in an audio content database; and

if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, identifying at least one product containing or relating to audio content that corresponds to the matching audio identifying information.

6. The method according to claim 5, further comprising the steps of:

receiving the audio content image; and
generating the audio identifying information for the audio content image.

7. The method according to claim 5, further comprising the step of receiving the audio identifying information for the audio content image.

8. The method according to claim 5, further comprising the step of charging a fee for the identified product.

9. The method according to claim 5, wherein the audio identifying information is an audio feature signature that is based on audio content.

10. The method according to claim 9, wherein the determining step includes the sub-step of comparing the audio feature signature generated for the audio content image with the audio feature signatures stored in the audio content database.

11. The method according to claim 5, further comprising the steps of:

generating audio identifying information corresponding to predetermined audio content; and

storing the audio identifying information corresponding to the predetermined audio content in the audio content database.

12. The method according to claim 5, further comprising the step of charging a fee for identifying the audio content that corresponds to the matching audio identifying information.

13. The method according to claim 6, wherein the generating step includes the sub-steps of:

obtaining an audio signal characterized by a time dependent power spectrum;

analyzing the spectrum to obtain a plurality of time dependent frequency components;

detecting a plurality of events in the plurality of time dependent frequency components; and

producing the audio identifying information for the audio content from the audio datastream based on the plurality of events.

14. The method according to claim 13, wherein the sub-step of analyzing the spectrum includes:

sampling the audio signal to obtain a plurality of audio signal samples;

taking a plurality of subsets from the plurality of audio signal samples; and

performing a Fourier transform on each of the plurality of subsets to obtain a set of Fourier frequency components.

15. The method according to claim 6, wherein the generating step includes the sub-steps of:

performing a Fourier transformation of the audio signals of the portion into a time series of audio power dissipated over a first plurality of frequencies;

grouping the frequencies into a smaller second plurality of bands that each include a range of neighboring frequencies;

detecting power dissipation events in each of the bands; and

grouping together the power dissipation events from mutually adjacent bands at a selected moment so as to form the identifying feature.

16. A computer-readable medium encoded with a program for selling products containing audio content, said program containing instructions for performing the steps of:

receiving a recorded audio content image;

generating audio identifying information for the audio content image;

determining whether the audio identifying information generated for the audio content image matches audio identifying information in an audio content database; and

if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, generating a product containing audio content that corresponds to the matching audio identifying information.

17. The computer-readable medium according to claim 16, wherein said program further contains instructions for performing the step of charging a fee for the generated product.

18. A computer-readable medium encoded with a program for selling products containing or relating to audio content, said program containing instructions for performing the steps of:

determining whether audio identifying information generated for an audio content image matches audio identifying information in an audio content database; and

if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database, identifying at least one product containing or relating to audio content that corresponds to the matching audio identifying information.

19. The computer-readable medium according to claim 18, wherein said program further contains instructions for performing the steps of:

receiving the audio content image; and

generating the audio identifying information for the audio content image.

20. The computer-readable medium according to claim 18, wherein said program further contains instructions for performing the step of receiving the audio identifying information for the audio content image.

21. The computer-readable medium according to claim 18, further comprising:

receiving a transmitted audio content image of at least a portion of a song; and

transmitting a recording of at least a song that corresponds to the matching audio identifying information.

22. The computer-readable medium according to claim 18,

wherein the audio identifying information is an audio feature signature that is based on audio content, and

the determining step includes the sub-step of comparing the audio feature signature generated for the audio content image with the audio feature signatures stored in the audio content database.

23. The computer-readable medium according to claim 19, wherein said program further contains instructions for performing the steps of:

generating audio identifying information corresponding to predetermined audio content; and

storing the audio identifying information corresponding to the predetermined audio content in the audio content database.

24. The computer-readable medium according to claim 19, wherein the generating step includes the sub-steps of:

obtaining an audio signal characterized by a time dependent power spectrum;

analyzing the spectrum to obtain a plurality of time dependent frequency components;

detecting a plurality of events in the plurality of time dependent frequency components; and

producing the audio identifying information for the audio content from the audio datastream based on the plurality of events.

25. The computer-readable medium according to claim 24, wherein the sub-step of analyzing the spectrum includes:

sampling the audio signal to obtain a plurality of audio signal samples;

taking a plurality of subsets from the plurality of audio signal samples; and

performing a Fourier transform on each of the plurality of subsets to obtain a set of Fourier frequency components.

26. The computer-readable medium according to claim 19, wherein the generating step includes the sub-steps of:

performing a Fourier transformation of the audio signals of the portion into a time series of audio power dissipated over a first plurality of frequencies;

grouping the frequencies into a smaller second plurality of bands that each include a range of neighboring frequencies;

detecting power dissipation events in each of the bands; and

grouping together the power dissipation events from mutually adjacent bands at a selected moment so as to form the identifying feature.

27. A system comprising:

an input interface for receiving a recorded audio content image;

an identifying information generator for generating audio identifying information for the audio content image;

a match detector for determining whether the audio identifying information generated for the audio content image matches audio identifying information in an audio content database; and

a product generator for generating a product containing audio content that corresponds to the matching audio identifying information, if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database.

28. A system comprising:

a match detector for determining whether audio identifying information generated for an audio content image matches audio identifying information in an audio content database; and

a product identifier for identifying at least one product containing or relating to audio content that corresponds to the matching audio identifying information, if the audio identifying information generated for the audio content image matches audio identifying information in the audio content database.

29. The system according to claim 28, further comprising:

an input interface for receiving the audio content image; and

an identifying information generator for generating the audio identifying information for the audio content image.

30. The system according to claim 28, further comprising an input interface for receiving the audio identifying information for the audio content image.

31. The system according to claim 28, further comprising:

an output interface for transmitting a recording of at least a song that corresponds to the matching audio identifying information,

wherein the input interface is adapted for receiving a transmitted audio content image of at least a portion of a song.

32. The system according to claim 28, wherein the audio identifying information is an audio feature signature that is based on audio content.

33. The system according to claim 32, wherein the match detector compares the audio feature signature generated for the audio content image with the audio feature signatures stored in the audio content database.

34. The system according to claim 28, wherein the audio content database stores audio identifying information for predetermined audio content.

* * * * *