



- (51) **International Patent Classification:**
H04H 60/37 (2008.01)
- (21) **International Application Number:**
PCT/US2016/034401
- (22) **International Filing Date:**
26 May 2016 (26.05.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/175,145 12 June 2015 (12.06.2015) US
14/815,256 31 July 2015 (31.07.2015) US
- (71) **Applicant:** SORENSON MEDIA, INC. [US/US]; 13961 Minuteman Dr. Suite 100, Draper, Utah 84020 (US).
- (72) **Inventor:** CHEN, Juikun; 6869 Running Springs Drive, West Jordan, Utah 84084 (US).
- (74) **Agents:** GRANGE, Kevin O. et al.; Lowenstein Sandler LLP, 65 Livingston Avenue, Roseland, New Jersey 07068 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,

DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

(54) **Title:** DETECTING CHANNEL CHANGE IN AUTOMATIC CONTENT RECOGNITION FINGERPRINT MATCHING




FIG. 5

(57) **Abstract:** The disclosed system and method receives, from a media device, a query fingerprint that includes an ordered sequence of frames of content being consumed, and matches the query fingerprint with frame fingerprints by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames. The matching may be performed by populating a binary tree structure with respective distances, chronologically across leaf nodes of the binary tree structure, and then noting a value of the root node. If below a maximum threshold, then there is a match. The system and method may detect a channel change at the media device by traversing a sub-tree of the binary tree structure, and locating a low value (or a zero) at the first or second nodes of the sub-tree. This channel change may be useful in targeting media to content being streamed to the media device.

WO 2016/200622 A1

DETECTING CHANNEL CHANGE IN AUTOMATIC CONTENT RECOGNITION FINGERPRINT MATCHING

BACKGROUND

[0001] Media consumption devices, such as smart televisions (TVs), may access broadcast digital content and receive data, such as streaming media, from data networks (such as the Internet). Streaming media refers to a service in which media content such as movies or news may be provided to an end user over a telephone line, cable, Internet, and so forth upon request. For example, a user may view a movie without having to leave their residence. Also, users may access various types of educational content, such as video lectures, without having to physically attend at a school or educational institution.

[0002] As the number of media consumption devices continues to increase, video content generation and delivery may similarly increase. With an increase in use of media consuming devices (such as smartphones, tablets, and smart televisions) to access streaming media, content or network providers (such as local broadcasters, multi-channel networks, and other content owners/distributors) may distribute contextually-relevant material to viewers that are consuming streaming media (e.g., media programs). For example, local broadcasters may include contextually-relevant advertisements and interactive content with streaming media.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The disclosure will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the disclosure. The drawings, however, should not be taken to limit the disclosure to the specific embodiments, but are for explanation and understanding only.

[0004] FIG. 1 illustrates a system diagram of a content distribution network according to one embodiment.

[0005] FIG. 2 illustrates a content manager to provide overlay content to a client device according to one embodiment.

[0006] FIG. 3 illustrates a system diagram of an automatic content recognition (ACR) engine used to fingerprint media content for the content manager of FIG. 2.

[0007] FIG. 4 illustrates a flowchart of a method of automatic content recognition (ACR) that matches a sequence of frames of an input (or query) fingerprint to identify a corresponding television program, according to one embodiment.

[0008] FIG. 5 is an example diagram of a populated binary tree structure indicating a channel change of a media device in the middle of a 64-frame query fingerprint.

[0009] FIG. 6 is an example diagram of a populated binary tree structure indicating a channel change after 12 frames into a 64-frame query fingerprint.

[0010] FIG. 7 is an example diagram of a populated binary tree structure indicating a channel change 16 frames from an end of a 64-frame query fingerprint.

[0011] FIG. 8 illustrates a flowchart of a method of ACR that employs a binary tree structure to match fingerprints and to detect a channel change of a media device.

[0012] FIG. 9 illustrates a diagrammatic representation of a machine in the example form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

DESCRIPTION OF EMBODIMENTS

[0013] Media content broadcasting or streaming, such as television (TV) or internet show broadcasting, can be an engaging venue to advertise products and services, provide information to viewers, or any combination thereof. Accordingly, broadcasters want to know what programs individual viewers are watching, and when, so that subject matter of those programs may be used to accurately target advertising and other useful, optionally non-commercial information to the viewers. Non-commercial may include, for example, news alerts, announcements or educational information. It would therefore be advantageous to determine a program a user is watching or is about to watch, and to send an identification of the program information to an advertising (“ad”) targeter or server for use in such targeting actions. In furthering such targeting actions, it is advantageous to detect a channel change quickly so as not to miss targeting opportunities that may follow closely after the channel change.

[0014] The disclosed system and method receives, from a media device, a query fingerprint that includes an ordered sequence of frames of content being consumed by a user. The system and method matches the query fingerprint with frame fingerprints, which are stored by the system for individual frames of programs streaming to a plurality of television channels, by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames. A distance may include, for example, a numerical value indicating how different the two frames are from each other or a difference between channels on which the two respective frames are being streamed. In another embodiment, the distance is derived by summing the absolute distance of each frame pair between two fingerprints.

[0015] Matching a query fingerprint with a fingerprint from a known program or channel may be performed by populating a binary tree structure with respective distances (between respective frames) chronologically across leaf nodes of the binary tree structure, and then noting a value of the root node. When the value is below a preset maximum threshold, then there is a match. In one example, the maximum threshold is a small number or zero, e.g., indicating very little difference and a statistically close match.

[0016] Populating the binary tree structure may occur by populating the respective distances between the respective frame fingerprints and the ordered sequence of frames from left to right (or right to left) across the leaf nodes (or first children) of the binary tree structure. Parents then take on the sum or addition of the value of their two children as the binary tree structure is populated going up to the root node, as will be discussed in more detail.

[0017] In some cases, a television channel may be changed partway through the ordered sequence of frames of the query fingerprint, e.g., during a transition of changing from a first to a second channel. This may result in a medium value at the root node, thus indicating a mismatch, although the latter part of the ordered sequence of frames of the query fingerprint may actually match the frame fingerprints of the originating content. The system and method may detect a channel change at the media device by traversing a right (or a left) sub-tree of the binary tree structure, and locating a threshold low value (or a zero) at the first or second node of the sub-tree. A channel streaming media content matching the ordered sequence of frames of the query fingerprint at the end of the ordered sequence of frames may then be considered the new (or second) channel. Detecting the channel change may be useful in targeting media to content being streamed to the media device as quickly as possible and continuously despite channel changes.

[0018] The system and method may send an identification of the channel to an advertising server and receive, from the advertising server, an advertisement contextually-relevant to a subject matter of a media program being played on that channel. The system and method may then deliver the advertisement (or other content) to the media device for display as an overlay or as a replacement advertisement (or informational segment) during a break (e.g., a certain slot or position) in the media program.

[0019] To provide viewers with media content, an individual or an organization may stream the media content to viewers, such as by delivering the media content over the Internet to the viewers. The media content used by the individual or the organization may be media content (such as video data) acquired from one or more live broadcast media feeds. For

example, a media content provider may provide a user with a linear media channel (e.g., media provided from a live media feed source to a viewer) over the Internet.

[0020] The word “content” may be used to refer to media or multimedia. The word “content” may also be a specific term that means the subject matter of the medium rather than the medium itself. Likewise, the word “media” and some compound words that include “media” (e.g. multimedia, hypermedia) are instead referring to content, rather than to the channel through which the information is delivered to the end user/audience. Media or media content may include graphical representations, such as: videos, films, television shows, commercials, streaming video, and so forth; text; graphics; animations; still images; interactivity content forms; and so forth. An example of a type of content commonly referred to as a type of media is a “motion picture” referred to as “a film.”

[0021] In one embodiment, a content overlay system or a content overlay device may enable combining media content with specific, timely, and/or targeted overlay content such as advertising. The content overlay system or content overlay device may enable overlay content providers to engage with viewers by inviting the viewers to respond to a call to action within the content overlays (e.g., an invitation to engage the content overlay). One advantage of inviting the viewers to a call to action may be to provide a return path or follow-up path for the viewers to request additional information, ask questions, provide input, contact a provider of a service or product advertised, and so forth. Another advantage of inviting the viewer to a call to action may be to provide a return path or follow up path for the advertisers to provide additional information, further engage the viewers, gather additional information about the viewers, answer viewer questions about the product or service advertised, and so forth. In another example, the content overlay system or the content overlay device may enable an advertiser to use cross platform retargeting campaigns once a viewer has viewed and/or interacted with an overlay content of a media program.

[0022] FIG. 1 illustrates a system diagram of a content distribution network 100 according to one example. In the content distribution network 100, a content provider 102 may broadcast a content feed to a local provider 106. The local provider 106 may include a headend 104 and an automatic content recognition (ACR) fingerprinter server 105. The content feed from the content provider 102 may be received at the headend 104 of the local provider 106. The headend 104 may generate a local content feed based on the received content feed. For example, the headend 104 may be a local affiliate broadcaster receiving a network channel with programming and advertisements from a national or global broadcaster.

[0023] The headend 104 may communicate the local content feed to the ACR fingerprinter server 105, an over-the-air (OTA) broadcaster 108, and/or a multichannel video programming distributor (MVPD) 110. The OTA broadcaster 108 and/or the MVPD 110 may communicate the local content feed to a media device 115. Some examples of the media devices 115 include client devices 118 and 120, a set top box 114 that streams provider content to the client devices 118 and 120, as well as other devices 116 through which the user may stream the local content feed, e.g., wirelessly.

[0024] In one example, the OTA broadcaster 108 may broadcast the local content feed using traditional local television or radio channels. In this example, the client devices 118 and 120 may include antennas (such as TV or radio antennas) and receive the local content feed. In another example, the MVPD 110 (such as cable or satellite broadcaster) may communicate the local content feed to a set top box 114. In this example, the set top box 114 may format the content feed for the client devices 118 and 120 and may communicate the formatted content feed to the client devices 118 and 120. The client devices 118 and 120 may include a display device, such as a television screen or a touch screen, to display the local content to a viewer. Various components of the content distribution network 100 may be integrated or coupled to the client devices 118 and 120. For example, a smart television may include the antennas, the set top box 114, and a display device in a single unit.

[0025] The ACR fingerprint server 105 may analyze the local content feed and determine fingerprint information (e.g., fingerprints). The ACR fingerprint server 105 may communicate the fingerprints to ACR systems 124 and/or 126. The ACR systems 124 and 126 may be different ACR systems selected by device manufacturers, such as smart TV manufacturers. In some embodiments, as described below, the ACR system 124 includes an ACR fingerprint sequence matcher 125. The ACR fingerprint sequence matcher 125 may match frame fingerprints to the original video content from which the corresponding video frames originate. The ACR system 126 may or may not have the ACR fingerprint sequence matcher 125. Details regarding the ACR fingerprint sequence matcher 125 are described below with respect to FIG. 3.

[0026] The ACR fingerprint server 105 may analyze the local content feed and capture fingerprints, which may include an ordered sequence of frames from the local content feed. The ACR fingerprint server 105 may communicate the fingerprints to the ACR systems 124 and/or 126. The ACR systems 124 and 126 may be different ACR systems selected by device manufacturers, such as smart TV manufacturers. In one example, the ACR fingerprint server 105 may format fingerprints for the different ACR systems 124 and 126, e.g., that include

different types of fingerprinting technology such as different fingerprinting algorithms. The ACR systems 124 and 126 may establish communication connections with the different media devices 115, including the client devices 118 and 120, respectively. The client devices 118 and 120 may communicate fingerprint information to the ACR systems 124 and 126, respectively. When the ACR system 124 or 126 receives ACR fingerprint information from the client devices 118 and/or 120, the ACR system 124 or 126 may match the received fingerprints with those generated by the ACR fingerprint server 105 and when a match occurs and the content has been identified, may communicate ACR events to a content manager 122.

[0027] In another example, the ACR systems 124 and 126 may receive ACR fingerprint information from the client devices 118 and 120 and may match the received fingerprints with those generated by the ACR fingerprint server 105. When a match occurs and the content has been identified, the ACR systems 124 and/or 126 may notify the client device 118 or 120 of the ACR events and then the client device 118 or 120 may communicate those ACR events to a content manager 122. Alternatively, or additionally, the ACR systems 124 and/or 126 may directly communicate the ACR events to the content manager 122. The ACR fingerprint information may include: a display of advertisements in the local content feed to a viewer, a display of selected or flagged content in the local content feed to a viewer, a change of content channel at the client device 118 or 120, and so forth.

[0028] The event information from the different ACR systems 124 and 126 may be in different formats and the content manager 122 may normalize the data into a common format before storing the data into a database 123. For example, the content manager 122 may receive disparate data sets from the ACR systems 124 and 126 that include similar but not identical data, such as data with the same content but formatted differently. The content manager 122 may process and reformat the disparate data sets to create a single data model or format (e.g., reformatted data sets) and the reformatted data sets may be populated into the database 123 in the content manager 122.

[0029] In one embodiment, to normalize disparate data sets from ACR systems 124 and 126, the content manager 122 may cleanse or filter data in the data sets. For example, some data sets may contain fields or data that may be irrelevant to the content manager 122. In this example, the content manager 122 may cleanse or filter the irrelevant data (e.g., the data may be removed or ignored). In another example, some data sets may include instances of incomplete or incorrect data or data sets and the content manager 122 may cleanse or filter the incomplete or incorrect data or data sets. In another embodiment, to normalize the disparate data sets from ACR systems 124 and 126, the content manager 122 may map fields of the data

sets. For example, when the content manager 122 receives a first data set from the ACR system 124 and a second data set from the ACR system 126, at least some of the data fields of the first data set and the second data set may be common to both the first and second data set. However, the common data fields may be located at different places in the first and second data sets. In this example, the content manager 122 may map the different data fields of the first and second data sets to normalized fields and have the same data fields in the same data field locations in the database 123.

[0030] In another embodiment, to normalize disparate data sets from the ACR systems 124 and 126, the content manager 122 may derive data from the data sets. For example, data from the ACR systems 124 and/or 126 may not contain all of the fields that are needed to fill the data fields in the database. However, the content manager 122 may use other fields in the data sets from the ACR systems 124 and 126 to derive data for these data fields.

[0031] In one example, the database 123 may include data fields for such as a state in a country field, a designated market area (DMA), and a county and/or city field but the data sets from the ACR systems 124 and 126 may only include zone improvement plan (ZIP) codes. In this example, the content manager 122 may use the ZIP codes to derive data for the fields in the database. In another example, the data set may not contain any geographic location information, but may include an internet protocol (IP) address of the ACR systems 124 and 126. In this example, the content manager 122 may use a geo-IP lookup service to derive the state, DMA, county, city and ZIP code information.

[0032] In another example, the database 123 may include demographic fields such as an age field, a gender field, a household income field, and so forth. However, the data sets from the ACR systems 124 and 126 may not include the demographic fields or demographic data. In this example, the ACR systems 124 and 126 may provide the content manager 122 with the IP address of the client devices 118 and 120. The content manager 122 may use the IP addresses to determine the demographic data to populate the data fields in the database.

[0033] In another example, a field in a first data set from the ACR system 124 may include local time zone information, such as a mountain daylight time (MDT) zone, and a second data set from the ACR system 126 may include information from another time zone, such as a coordinated universal time (UTC) zone. The database may store all data using the UTC and the content manager 122 may convert the local time to UTC before storing the data in the database 123.

[0034] In one embodiment, the content manager 122 may use the normalized data to generate reports or data (viewing data) about user viewing behavior across different ACR

technology vendors and smart TV or other Internet-connected video devices. The content manager 122 and the media devices 115 may include communications interfaces to communicate information, such as overlay content, between the media devices 115 and the content manager 122.

[0035] In one example, the communication interface may communicate the information using a cellular network and/or a wireless network. In one example, the communications network may be a cellular network that may be a third generation partnership project (3GPP) release 8, 9, 10, 11, or 12 or Institute of Electronics and Electrical Engineers (IEEE) 802.16p, 802.16n, 802.16m-2011, 802.16h-2010, 802.16j-2009, 802.16-2009. In another embodiment, the communications network may be a wireless network (such as a network using the Wi-Fi® technology developed by the Wi-Fi Alliance) that may follow an IEEE® standard developed by the Institute of Electrical and Electronics Engineers, Inc., such as the IEEE 802.11-2012, IEEE 802.11ac, or IEEE 802.11ad standards. In another embodiment, the communications network may be a Bluetooth® connection developed by Bluetooth Special Interest Group (SIG) such as Bluetooth v1.0, Bluetooth v2.0, Bluetooth v3.0, or Bluetooth v4.0. In another embodiment, the communications network may be a Zigbee® connection developed by the ZigBee Alliance such as IEEE 802.15.4-2003 (Zigbee 2003), IEEE 802.15.4-2006 (Zigbee 2006), IEEE 802.15.4-2007 (Zigbee Pro).

[0036] In one example, the content manager 122 may also instruct the media devices 115 to replace portions of the local content feed received from the OTA broadcaster 108 or the MVPD 110 with overlay content. In another example, the content manager 122 may instruct the media devices 115 to overlay or superimpose overlay content onto portions of the local content feed. The content manager 122 may aggregate ACR information across multiple ACR systems 124 and 126 and may communicate overlay content to different client devices 118 and 120, where the client devices 118 and 120 may be from different device manufacturers.

[0037] The content manager 122 may also establish communication connections with other devices 116 categorized generally media devices 115. In one example, the other device 116 may communicate with the client devices 118 or 120 and provide an additional screen (e.g., a second screen) to display overlay content. For example, the client devices 118 and 120 may receive the local content feed from the OTA broadcaster 108 or the MVPD 110 and display the local content feed to the user. The other devices 116 may also communicate ACR event information to the ACR systems 124 and 126 when an ACR event occurs, as discussed in the preceding paragraphs. When the content manager 122 receives the ACR event

information, the content manager 122 may communicate overlay content to the other devices 116.

[0038] In one example, the client devices 118 and 120 may continue to display the local content feed while the other devices 116 display the overlay content. In another example, the client devices 118 and 120 and the other devices 116 may both display the overlay content. In another example, the client devices 118 and 120 and the other devices 116 may display a portion of the overlay content and a portion of the local content feed. In another example, the client devices 118 and 120 and the other devices 116 may display different local content feeds and/or overlay content.

[0039] In one example, the client devices 118 and 120 and/or the other devices 116 may display the overlay content at the time the overlay content is received. In another example, the client devices 118 and 120 and/or the other devices 116 may delay displaying the overlay content for a threshold period of time. The threshold period of time may be a predefined period of time or the content manager 122 may select a period of time for the client devices 118 and 120 and/or the other devices 116 to delay displaying the overlay content.

[0040] FIG. 2 illustrates a content manager 222 to provide overlay content to a media device 115 such as to client devices 218 and/or 220 according to one embodiment. A content provider 202 may stream media content to the media device 115 over a network 219, which streaming may be intercepted by the content manager 222 before, or simultaneously with the streaming of the media content to the media device 115. The content manager 222 may also communicate with an advertisement server (or “ad” server) 230 such as to send the ad server 230 an advertising call that requests an advertisement be served with (or as an overlay to) the media content to target the subject matter of the media content and/or interests of a user as will explained in more detail. The ad server 230 may be a third party or external server that provides advertising or other overlay content to the content manager 222 for later delivery to the media devices 115 or may provide the content directly to the media devices 115 for delivery as overlay or replacement advertisements.

[0041] The content manager 222 may include an ACR engine 204, a look-up server 206, an overlay decision engine 210, an overlay database 211 in which to store overlay content, and an ad targeter 212. The content provider 202 may upload media content to the ACR engine 204.

[0042] The ACR engine 204 may fingerprint the media content. In one example, fingerprints may be generated by fingerprinting every frame of a feed, every other frame of the feed, a sequence of frames and so forth. For example, the ACR engine 204 may generate a

fingerprint for a frame of the feed by performing a discrete cosine transform (DCT) of the frame and designating a subset of the resulting coefficients (e.g., the low-frequency coefficients) as the fingerprint. Additionally, or alternatively, the ACR engine 204 may generate (or identify) a target block of fingerprints that correspond to an ordered sequence of frames of content being streamed to a media device 115. Furthermore, the look-up server 206 may find a matching source block of frame fingerprints stored in a fingerprint database 207 for the originating media programs (e.g., originating content) for which an identity is known.

[0043] In another or related embodiment, the ACR engine 204 may also analyze ACR event information to determine what event may have occurred, e.g., a positive match between a sequence-of-frames query fingerprint and frame fingerprints of originating content stored in the ACR system 124 or 126. When there is a positive match, the ACR engine 204 may send a positive match indicator to the requesting media device 115 that may include a media content identifier (ID) that identifies the content for which a positive match results. The media device 115 may send an overlay request to the overlay decision engine 210 requesting a media content overlay. In one example, the overlay request may include the media content ID. In another example, the overlay request may include overlay information or overlay parameters.

[0044] With further reference to FIG. 2, the ACR engine 204 may communicate fingerprints to the look-up server 206, which may look up and determine a television program (or channel, for example) and location within the television program corresponding to an ACR event received from the ACR system 124 or 126. Each fingerprint of a segment of the feed may be associated with a time stamp. The time stamp may belong to individual frames of the segment of the feed when received by the ACR engine 204. The time stamp may be a frame number within the feed from an arbitrary starting point. The look-up server 206 may store the fingerprints in association with their respective time stamps (e.g., in the fingerprint database 207), and aid the ad targeter 212 and the overlay decision engine 210 in timing and content targeting within the media content of the feed that the user is viewing.

[0045] In one embodiment, the ACR engine 204 interacts with an ACR client 215 at various media devices 115. The ACR client 215 may locally match fingerprints and confirm whether or not the user has changed a channel to watch a different television program, and to report the channel change to the content manager 222. Accordingly, matching of fingerprints may occur locally at the media devices 115 in some cases.

[0046] Furthermore, the ACR client 215 may periodically, continuously, or semi-continuously communicate user fingerprint information to the look-up server 206, e.g., in the form of query fingerprints, including target blocks, requesting to confirm the television

program or channel being watched on the media device 115. The look-up server 206 may determine when there is a match between the query fingerprint(s) (or target blocks) and a multitude of frame fingerprints, which may be stored as source blocks of fingerprints in the fingerprints database 207 or across the network 219 from the look-up server. The query fingerprint may be an ordered sequence of frames, respective ones of which may be matched with individual ones of the frame fingerprints until finding a sufficient match to be associated with the television program or channel streamed the television program. When there is a positive match, the look-up server 206 may communicate a positive match indicator to the ACR client 215. The ACR client 215 may send an overlay request to an overlay decision engine 210 requesting a media content overlay.

[0047] In one example, the overlay request may include a media content identifier (ID). In another example, the overlay request may include overlay information or overlay parameters. In one example, the overlay decision engine 210 may use the content ID, overlay information, and/or overlay parameters to identify targeted overlay content. In another example, the overlay decision engine 210 may use the content ID, overlay information, and/or overlay parameters to identify an overlay format. The overlay decision engine 210 may compare the content ID, overlay information, and/or overlay parameters with an overlay database 211 to identify the targeted overlay content and the overlay format. The overlay database may be updated, by a content provider or an advertiser (e.g., the ad server 230), with new overlay content and overlay formats on a periodic or continuous basis. The overlay content may populate the overlay format (such as an overlay template or the like) before or after being delivered to an overlay position of the streamed media content of the television program of channel.

[0048] The ad targeter 212 may track and analyze user interaction with and behavior regarding advertisements and other overlay content delivered to the media devices 115 by the overlay decision engine. The ad targeter 212 may also receive and incorporate user profile information with the analysis of user behavior on a per-media-device basis, to determine subject matter of interest to users. This information and data gathered on a user or group of users may extend to preferred viewing times and typical viewing habits with regards to television programs and the channels typically watched, and when. The ad targeter 212 may then inform the overlay decision engine 210, e.g., in the form of parameters, different subject matters of interest and viewing habits that the overlay decision engine 210 may use in deciding what overlay content to select for delivery to respective users, how to format it and when best to deliver for maximum return on investment of campaign budgets.

[0049] When the overlay decision engine 210 identifies the targeted overlay content with the help of the ad targeter 212, the overlay decision engine may return targeted overlay content to the media device 115. In one example, the overlay decision engine 210 may communicate the targeted overlay content directly to the media device 115, such as via a wireless communications network. In another example, the overlay decision engine 210 may communicate the targeted overlay content to the media device 115 via a universal resource locator (URL). In one example, when multiple targeted overlay contents match the content ID, overlay information, and/or overlay parameters, the overlay decision engine 210 may select the targeted content overlay that meets a greatest number parameters or other information. In another example, when multiple targeted overlay contents match the content ID, overlay information, and other overlay parameters, the overlay decision engine 210 may randomly select an overlay content that meets the parameters and other information. In another example, when multiple targeted overlay contents match the content ID, overlay information, and overlay parameters, the overlay decision engine 210 may select predetermined overlay content that matches the content ID, overlay information, and/or overlay parameters. The overlay content may be populated with dynamic content (e.g., content that may be updated or refreshed at periodic intervals). The dynamic content may be stored in a local database or an external system.

[0050] The ACR client 215 of the media device 115 may superimpose overlay content over the content feed when the ACR fingerprint information matches the user fingerprint information. In one example, the media device 115 may superimpose overlay content over the content feed in a hypertext markup language (HTML) browser. In another example, the media device 115 may superimpose overlay content over a content feed from an OTA broadcaster or a cable broadcaster. When the overlay content is placed over the content feed, the overlay content may be displayed to the user via a display of the media device 115. In one example, the overlay content may include one or more call-to-action options that may be displayed to a user. In this example, the user may interact with the overlay content using an input device (such as a TV remote, keyboard, a smartphone, or a tablet) to create feedback information. The ACR client 215 may communicate the feedback information to an ad targeter 212. Another individual, such as an advertiser, may access the feedback information and analyze the feedback information to determine desired information, such as user interest in the overlay content.

[0051] The ACR client 215 may monitor the content feed to determine when the overlay content and content feed match ceases and/or a threshold period of time expires. In one

example, when the overlay content and content feed match ceases or a threshold period of time expires, the media device may cease to superimpose the overlay content for display on the media device 115.

[0052] FIG. 3 is a system diagram of the ACR engine 204 used to fingerprint media content for the content manager of FIG. 2. The ACR engine 204 may receive content frames of the media content to be fingerprinted, a fingerprinter 305, a fingerprint sequence matcher 325 and a database 327 in which are stored frame fingerprints, according to one embodiment. The content provider 202 may generate multimedia content that is streamed to the media devices 115, including the client devices 218 and/or 220.

[0053] The fingerprinter 305 may detect or select a number of content frames 302 from the multimedia content as a fingerprint. In one embodiment, the number of content frames 302 may be sequentially ordered (e.g., as a target block), and thus include sequential time stamps from a beginning to an end of the fingerprint. In one example, the content may be audio data, video data, or both. In this example, video content may be raw video frames.

[0054] For example, when the fingerprinter 305 receives the content frames 302, the fingerprinter 305 may determine how to process the content frames 302, such as the raw video and/or audio frames to generate the fingerprint. In one example, the frames may be fingerprinted individually. In another example, the frames may be fingerprinted in collections or sequences. The fingerprinter 305 may determine when to fingerprint the frames individually or sequentially based on an ACR algorithm that the fingerprinter 305 executes during fingerprinting, examples of which will be discussed.

[0055] Additionally, the fingerprinter 305 may fingerprint the content frames 302 differently for different broadcasters or users. In this example, the fingerprinter 305 may include different ACR fingerprinting algorithms for different ACR vendors. In one example, the different ACR fingerprinting algorithms may be predetermined and stored in memory of the fingerprinter 305.

[0056] In yet another example, the different ACR fingerprinting algorithms may be provided by third party ACR vendors. When the different ACR fingerprinting algorithms are provided by such vendors, the fingerprinter 305 may aggregate the different ACR fingerprinting algorithms. In one example, ACR fingerprinting may use raw video with in the YUV 4:2:2 colorspace and at high resolutions or other levels of resolution. When video content is received at the local provider 106 or other content provider 202, the fingerprinter 305 may convert the video content to YUV 4:2:0 colorspace and scale it down in resolution to

a threshold resolution level for encoding by a broadcaster or distributor before being sent to the media devices 115.

[0057] In some embodiments, the fingerprinter 305 may include, or may communicate with the fingerprint sequence matcher 325. The fingerprint sequence matcher 325 may match a sequence of fingerprints to an original video content from which a given set of individual frame fingerprints originated as described in more detail below.

[0058] When the fingerprinter 305 has fingerprinted the content frames 302, the fingerprinter 305 may send fingerprints (including channel information, time codes, and fingerprint information) to the overlay database 211 and/or to a look-up server 206. The look-up server 206 may retrieve the fingerprints and related information from the overlay database 211. The look-up server 206 may also be in communication or coupled with the overlay decision engine 210 and the ad targeter 212, to send overlay and subject matter matching information to the overlay decision engine 210 and the ad targeter 212 with which to contextually target users on the client device 218 or 220.

[0059] In one example, the different ACR fingerprinting algorithms may be used on the same content to provide different fingerprint information to look-up servers of different ACR vendors. An advantage of fingerprinting the same content (e.g., content frames) 302 differently may be to provide contextually-relevant advertisements and interactive content to different viewers of media consumption devices. In another example, the content frames 302 may include media content from different feeds. In this example, the different ACR fingerprinting algorithms may be used on the content of the different feeds of the content frames 302 to provide different fingerprinting information to the look-up servers of different ACR vendors.

[0060] The different fingerprinting information may be uploaded to the look-up servers of the different ACR vendors, respectively. Different ACR vendors may be integrated on viewing devices manufactured by different contract equipment manufacturers (CEMs). For example, Toshiba televisions may utilize Samba® ACR fingerprinting and Samsung® televisions may use Enswer® ACR fingerprinting. An advantage of the fingerprinter 305 including ACR fingerprinting algorithms for different ACR vendors, may be to fingerprint content provided to viewers via different ACR vendors regardless of the manufacturer of the media consumption device. In one example, the ACR fingerprinting information may be used for digital advertisement replacement (DAR). In another example, the ACR fingerprinting information may be used for advertisement or content augmentation and data collection. The overlay decision engine 210 and the ad targeter 212 (FIG. 2) may use the fingerprinting

information to match the encoded content with contextually-relevant advertisements, informational and/or interactive content. The matched content and/or advertisements may then be provided to the media devices 115 for streaming display thereon. In another example, information about matching events may be communicated to a broadcast cloud for analysis of content matching.

[0061] In one example, the fingerprinter 305 may perform ACR fingerprinting upstream of broadcast distribution in the system. In another example, the fingerprinting may occur in the broadcast chain where broadcast feeds may be delivered for encoding/uploading. When a single broadcast feed is used, a number of devices and/or applications that may need to be purchased, installed, monitored, and maintained for ACR fingerprinting and encoding/uploading may be reduced. For example, capital expenses (CAPEX) and/or operational expenses (OPEX) may be reduced, such as a reduced system infrastructure (rack space, power, and Ethernet connectivity). Additionally, when a single broadcast feed is used, a number of broadcast feeds generated from a broadcaster's distribution amplifier may be reduced.

[0062] As discussed, the fingerprinter 305 may generate individual fingerprints from multimedia content, such as may pass through a broadcasting server, a headend, a switch, and/or a set-top box, in route to being displayed on one of the media devices 115. The fingerprints may include one frame or a number of frames. The frames may be taken in a sequential order with chronological time stamps, or may be taken at some interval such as every other (or every third frame), for example, still with corresponding time stamps.

[0063] The fingerprint sequence matcher 325 may process a set or sequence of these individual frames (as a query fingerprint) to match the frames to frame fingerprints (stored in the database 327) of original video content from which the corresponding video frames represented by these fingerprints originated. The frame fingerprints may be single frames and each include a corresponding time slot (or some kind of index such as a frame number or an amount of time from the beginning of the media program). Finding a match may result in determining the media program the user is watching on the media device 115, a channel or a channel change, a time slot of a schedule in terms of beginning and ending times of the program, commercial time slots, and such information of the matching media program or television channel.

[0064] To have confidence in the matching results, the fingerprint sequence matcher 325 may take advantage of the fact that the sequence of the fingerprint is ordered in time, and so the matching fingerprints not only belong to the same video, but are similarly ordered in time.

The fingerprint sequence matcher 325 exploits this property to populate a binary tree structure chronologically across leaf nodes of the binary tree structure, and then to traverse a fully populated binary tree structure to detect a channel change, as will be explained in more detail.

[0065] The fingerprint sequence matcher 325 may also use such temporal properties of sequences of frames in a fingerprint to detect how a matching fingerprint sequence was played out (faster, slower, or in reverse). An algorithm in the ACR fingerprint sequence matcher 325 may be implemented to detect certain playback scenarios, such as, for example, normal speed at full frame rate, normal speed at 1/2 frame rate, normal speed at 1/3 frame rate, or the like.

[0066] FIG. 4 illustrates a flowchart 400 of a method of automatic content recognition (ACR) that matches a sequence of frames of an input (or query) fingerprint to identify a corresponding television program, according to one embodiment. The method may be at least partially performed by processing logic that may include hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions executed by a processing device), firmware or a combination thereof. The method may be performed by processing logic of the client device such as client device 218 or 220, by a server system such as the ACR system 124 or 126 of FIG. 1 or the ACR system 224 or 226 of FIG. 2. The content manager 122 or 222 (FIGS. 1 and 2) may also play a part in executing the method. Alternatively, the method may be performed by other processing devices in various types of user device, portable devices, televisions, projectors, or other media devices.

[0067] Referring to FIG. 4, the processing logic begins with receiving media content from a content feed at a content (or media) device (402). The logic may continue to perform fingerprinting on the media content to generate an input (or query) fingerprint containing a sequence of frames and a corresponding time-based index (such as corresponding time stamps) (404). The logic may continue to match the query fingerprint, according to the time-based index, with a plurality of frame fingerprints from original media content, to identify a media program corresponding to the media content (406).

[0068] FIG. 5 is an example diagram of a populated binary tree structure 500 indicating a channel change of a media device in the middle of a 64-frame query fingerprint. While the query fingerprint (or target block) and source block of frame fingerprints are shown herein as spanning 64 frames of source content, different sizes of target blocks and source blocks may be used, such as 32 frames or 128 frames, and other sizes being the power of two.

[0069] Traditional methods in which media frames are matched frame by frame may be by calculating a distance between the frames, and a match is found only when the distances

are all zero between respective frames of a query fingerprint. As illustrated in FIG. 5, a binary tree structure may provide an alternative to this tradition method. For example, the ACR engine may populate leaf nodes 504 of the binary tree structure with the distances between respective frames of the query fingerprint and frame fingerprints of the source block from originating content. This population of the leaf nodes 504 may occur chronologically from left to right. In an alternative embodiment, the leaf nodes may be populated from right to left. The time stamps may aid the ACR engine (or other server) in making sure to match up the proper frames of target and source blocks and to populate the leaf nodes chronologically. A plurality of parent nodes 506 may then be populated with values from adding (or summing) the values of their respective children nodes. When a root node 508 contains a value within a threshold maximum number (such as a low number such as 8, 10, 12 or even zero), then the fingerprints may be said to match. Otherwise, there is no match.

[0070] More specifically, the example of FIG. 5 illustrates populating the binary tree structure 500 with distances where the target block indicates a channel change half-way through the 64 frames (e.g., at frame 33) from channel 8 to channel 9. The source block is channel 9 the whole time because it has already been streaming from channel 9 and the actual content arriving at the media device 115 has not yet caught up (thus catching the media device mid-way through a channel change). Although there is not room to show all 64 frames at the leaf nodes, the binary tree structure includes a left sub-tree 512 and a right sub-tree 516. The distances of “1” between all channels 8 and 9 through the first 32 frames get propagated up the binary tree structure, thus ending in a value of 32 at the root node of the left sub-tree and at the root node 508 of the entire binary tree structure. This creates the appearance of a mismatch, particularly with a threshold maximum number being relatively small.

[0071] But, because the channel was changed part-way through, the second half of the query fingerprint matches the frame fingerprints, resulting in distances of zero all the way along the right sub-tree 516. Because of this behavior within the binary tree structure, the ACR engine (or other server) may traverse the right sub-tree 516 and, upon identifying a zero at the first or second node of the right sub-tree, detect a channel change. (Note, if the binary tree was populated from right to left as indicated in an alternative embodiment, then traversal would be of the left sub-tree.) In another embodiment, the ACR engine may detect a low value below a threshold maximum distance (something greater than zero) to also signal a channel change. The new channel may be identified as the channel with media content matching the frames at the end of the query fingerprint, e.g., the ordered sequence of frames. This new channel may be associated with a content ID, as discussed, and be immediately used

for content and ad targeting by the content manager 222. This advantageously avoids missing any overlay or replacement advertising opportunity following soon after the point the media device 115 starts to receive media content from the new channel (e.g., channel 9 in FIG. 5).

[0072] FIG. 6 is an example diagram of a populated binary tree structure 600 indicating a channel change after 12 frames into a 64-frame query fingerprint. Within a plurality of leaf nodes 604, the channel is changed at the 13th node into the ordered sequence of frames of a current query fingerprint. A plurality of parent nodes 606 propagate distances up a left sub-tree 612 of the binary tree structure, ending in a smaller distance value of 12 at a root node 608. A value of 12, in one embodiment, may be low enough below a threshold maximum distance at the root node to be considered a match. In this case, the match will be with the current channel or program, designated as “9” in FIG. 6.

[0073] In another embodiment, however, the value of 12 at the root node may not be considered a match, and therefore, the ACR engine may continue traversing the binary tree structure 600 down a right sub-tree 616 of the binary tree structure, to determine whether there is a small enough distance (e.g., a threshold maximum value or below) at the first or second node of the right sub-tree to detect a channel change. The threshold maximum value may be zero or a relatively small value such as 8, 10, 12 or 16, for example.

[0074] FIG. 7 is an example diagram of a populated binary tree structure 700 indicating a channel change 16 frames from an end of a 64-frame query fingerprint. In this embodiment, the channel is changed at the media device from which the query fingerprint is generated at 17 frames before the end, providing less time in which to detect the channel change.

Accordingly, a plurality of leaf nodes 706 are populated with distances of “1” until frame 16 from the end, at which point streaming to the media device catches up to the channel change, resulting in a distance of zero. A plurality of parent nodes 706 propagate increasingly higher values of distances in the respective frames until a root node 708 takes on a distance value of 48.

[0075] A distance value of 48 reflects a mismatch between the query fingerprint and the originating content from the new television channel. But, according to the disclosed process, the ACR engine (server) may traverse a right sub-tree 716 of the binary tree structure and find a zero at not the first node, but at a second node down the right sub-tree 716. As discussed, instead of zero it may be detection of a low value (e.g., a threshold maximum value or below) to trigger detection of a channel change. So, although the root node 708 may indicate a mismatch, finding the zero (or a threshold low value) at the second node down the right sub-tree signals to the ACR engine a channel change. Advantageously, with knowing the new

channel (or media program associated with the new channel), the ad targeter 212 and/or the ad server 230 may continue to target overlay or replacement media content to the media device 115 at which the channel changed, with no interruption or waiting for a channel change determination at the next query fingerprint.

[0076] FIG. 8 illustrates a flowchart 800 of a method of ACR that employs a binary tree structure to match fingerprints and to detect a channel change of a media device. The method may be at least partially performed by processing logic that may include hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions executed by a processing device), firmware or a combination thereof. The method may be performed by processing logic of the client device such as client device 218 or 220, by a server system such as the ACR system 124 or 126 of FIG. 1 or the ACR system 224 or 226 of FIG. 2. The content manager 122 or 222 (FIGS. 1 and 2) may also play a part in executing the method. Alternatively, the method may be performed by other processing devices in various types of user device, portable devices, televisions, projectors, or other media devices.

[0077] Referring to FIG. 4, the processing logic begins with receiving, from a media device, a query fingerprint of content being consumed, the query fingerprint including an ordered sequence of frames (810). The processing logic may further match the query fingerprint with the plurality of frame fingerprints stored in the database by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames (820). Matching may include a number of steps, but in one embodiment, matching may include to populate leaf nodes of the binary tree structure with values of the respective distances chronologically from left to right (822) and to populate parent nodes with values derived from adding value of children nodes connected to the parent nodes (824).

[0078] The processing logic may further detect a channel change from a first channel to a second channel responsive to locating a zero value within a threshold number of children nodes from a root node when traversing down the binary tree structure (830). The processing logic may further determine an identity of the second channel by identifying a channel streaming content with frames matching those at an end of the ordered sequence of frames of the query fingerprint (840). The processing logic may further notify an advertisement (“ad”) targeter of the identity of the second channel being watched in response to detecting the channel change so that the ad targeter can adjust targeting media to content streamed to the second channel (850).

[0079] For example, the logic may send an identification of the media program to an advertising server and receive, from the advertising server, an advertisement contextually-

relevant to a subject matter of the media program. The logic may then deliver the advertisement (or other content) to the media device for display as an overlay or as a replacement advertisement (or informational segment) during a commercial break in the media program.

[0080] FIG. 9 illustrates a diagrammatic representation of a machine in the example form of a computer system 900 within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. In alternative embodiments, the machine may be connected (e.g., networked) to other machines in a LAN, an intranet, an extranet, or the Internet. The machine may operate in the capacity of a server or a client device in a client-server network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine may be a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a smartphone, a web appliance, a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0081] The computer system 900 may correspond to the ACR system 124 or 126 of FIG. 1, to the ACR system 224 or 226 of FIGS. 2 and 3, or to the content manager 122 of FIG. 1 or the content manager 222 of FIG. 2. The computer system 900 may correspond to any media device 115 such as the client device 118 or 120 of FIG. 1. The computer system 900 may correspond to at least a portion of a cloud-based computer system.

[0082] The computer system 900 includes a processing device 902, a main memory 904 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) (such as synchronous DRAM (SDRAM) or DRAM (RDRAM), etc.), a static memory 506 (e.g., flash memory, static random access memory (SRAM), etc.), and a data storage device 918, which communicate with each other via a bus 930.

[0083] Processing device 902 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device may be complex instruction set computing (CISC) microprocessor, reduced instruction set computer (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processing device 902 may also be one or more special-

purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. In one embodiment, processing device 902 may include one or more processing cores. The processing device 902 may execute the instructions 926 of a mirroring logic for performing the operations discussed herein.

[0084] The computer system 900 may further include a network interface device 908 communicably coupled to a network 920. The computer system 900 also may include a video display unit 910 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device 912 (e.g., a keyboard), a cursor control device 914 (e.g., a mouse), a signal generation device 916 (e.g., a speaker), or other peripheral devices. Furthermore, computer system 900 may include a graphics processing unit 922, a video processing unit 928, and an audio processing unit 932. In another embodiment, the computer system 900 may include a chipset (not illustrated), which refers to a group of integrated circuits, or chips, that are designed to work with the processing device 902 and controls communications between the processing device 902 and external devices. For example, the chipset may be a set of chips on a motherboard that links the processing device 902 to very high-speed devices, such as main memory 904 and graphic controllers, as well as linking the processing device 902 to lower-speed peripheral buses of peripherals, such as USB, PCI or ISA buses.

[0085] The data storage device 918 may include a computer-readable storage medium 925 on which is stored instructions 926 embodying any one or more of the methodologies of functions described herein. The instructions 926 may also reside, completely or at least partially, within the main memory 904 and/or within the processing device 902 during execution thereof by the computer system 900; the main memory 904 and the processing device 902 also constituting computer-readable storage media.

[0086] The computer-readable storage medium 925 may also be used to store instructions 926 utilizing logic and/or a software library containing methods that call the above applications. While the computer-readable storage medium 925 is shown in an exemplary implementation to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions 926 for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present disclosure. The term “computer-readable storage medium” shall accordingly be taken

to include, but not be limited to, solid-state memories, optical media, and magnetic media.

The following examples pertain to further embodiments.

[0087] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

[0088] In the description herein, numerous specific details are set forth, such as examples of specific types of processors and system configurations, specific hardware structures, specific architectural and micro architectural details, specific register configurations, specific instruction types, specific system components, specific measurements/heights, specific processor pipeline stages and operation etc. in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that these specific details need not be employed to practice the present invention. In other instances, well known components or methods, such as specific and alternative processor architectures, specific logic circuits/code for described algorithms, specific firmware code, specific interconnect operation, specific logic configurations, specific manufacturing techniques and materials, specific compiler implementations, specific expression of algorithms in code, specific power down and gating techniques/logic and other specific operational details of computer system have not been described in detail in order to avoid unnecessarily obscuring the present invention.

[0089] The embodiments are described with reference to secure memory repartitioning in specific integrated circuits, such as in computing platforms or microprocessors. The embodiments may also be applicable to other types of integrated circuits and programmable logic devices. For example, the disclosed embodiments are not limited to desktop computer systems or portable computers, such as the Intel® Ultrabooks™ computers. And may be also used in other devices, such as handheld devices, tablets, other thin notebooks, systems on a chip (SoC) devices, and embedded applications. Some examples of handheld devices include cellular phones, Internet protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications typically include a microcontroller, a digital signal processor (DSP), a system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that may perform the functions and operations taught below. It is described that the system may be any kind of computer or embedded system. The disclosed embodiments may especially be used for low-end devices, like wearable devices (e.g., watches), electronic implants, sensory and control infrastructure

devices, controllers, supervisory control and data acquisition (SCADA) systems, or the like. Moreover, the apparatuses, methods, and systems described herein are not limited to physical computing devices, but may also relate to software optimizations for energy conservation and efficiency. As will become readily apparent in the description below, the embodiments of methods, apparatuses, and systems described herein (whether in reference to hardware, firmware, software, or a combination thereof) are vital to a 'green technology' future balanced with performance considerations.

[0090] Although the embodiments herein are described with reference to a processor, other embodiments are applicable to other types of integrated circuits and logic devices. Similar techniques and teachings of embodiments of the present invention may be applied to other types of circuits or semiconductor devices that may benefit from higher pipeline throughput and improved performance. The teachings of embodiments of the present invention are applicable to any processor or machine that performs data manipulations. However, the present invention is not limited to processors or machines that perform 512 bit, 256 bit, 128 bit, 64 bit, 32 bit, or 16 bit data operations and may be applied to any processor and machine in which manipulation or management of data is performed. In addition, the description herein provides examples, and the accompanying drawings show various examples for the purposes of illustration. However, these examples should not be construed in a limiting sense as they are merely intended to provide examples of embodiments of the present invention rather than to provide an exhaustive list of all possible implementations of embodiments of the present invention.

[0091] Although the below examples describe instruction handling and distribution in the context of execution units and logic circuits, other embodiments of the present invention may be accomplished by way of a data or instructions stored on a machine-readable, tangible medium, which when performed by a machine cause the machine to perform functions consistent with at least one embodiment of the invention. In one embodiment, functions associated with embodiments of the present invention are embodied in machine-executable instructions. The instructions may be used to cause a general-purpose or special-purpose processor that is programmed with the instructions to perform the steps of the present invention. Embodiments of the present invention may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform one or more operations according to embodiments of the present invention. Alternatively, operations of embodiments of the present invention might be

performed by specific hardware components that contain fixed-function logic for performing the operations, or by any combination of programmed computer components and fixed-function hardware components.

[0092] Instructions used to program logic to perform embodiments of the invention may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions may be distributed via a network or by way of other computer readable media. Thus a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium includes any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

[0093] A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine readable medium. A memory or a magnetic or optical storage such as a disc may be the machine readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may store on a tangible,

machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present invention.

[0094] A module as used herein refers to any combination of hardware, software, and/or firmware. As an example, a module includes hardware, such as a micro-controller, associated with a non-transitory medium to store code adapted to be executed by the micro-controller. Therefore, reference to a module, in one embodiment, refers to the hardware, which is specifically configured to recognize and/or execute the code to be held on a non-transitory medium. Furthermore, in another embodiment, use of a module refers to the non-transitory medium including the code, which is specifically adapted to be executed by the microcontroller to perform predetermined operations. And as may be inferred, in yet another embodiment, the term module (in this example) may refer to the combination of the microcontroller and the non-transitory medium. Often module boundaries that are illustrated as separate commonly vary and potentially overlap. For example, a first and a second module may share hardware, software, firmware, or a combination thereof, while potentially retaining some independent hardware, software, or firmware. In one embodiment, use of the term logic includes hardware, such as transistors, registers, or other hardware, such as programmable logic devices.

[0095] Use of the phrase ‘configured to,’ in one embodiment, refers to arranging, putting together, manufacturing, offering to sell, importing and/or designing an apparatus, hardware, logic, or element to perform a designated or determined task. In this example, an apparatus or element thereof that is not operating is still ‘configured to’ perform a designated task if it is designed, coupled, and/or interconnected to perform said designated task. As a purely illustrative example, a logic gate may provide a 0 or a 1 during operation. But a logic gate ‘configured to’ provide an enable signal to a clock does not include every potential logic gate that may provide a 1 or 0. Instead, the logic gate is one coupled in some manner that during operation the 1 or 0 output is to enable the clock. Note once again that use of the term ‘configured to’ does not require operation, but instead focus on the latent state of an apparatus, hardware, and/or element, where in the latent state the apparatus, hardware, and/or element is designed to perform a particular task when the apparatus, hardware, and/or element is operating.

[0096] Furthermore, use of the phrases ‘to,’ ‘capable of/to,’ and or ‘operable to,’ in one embodiment, refers to some apparatus, logic, hardware, and/or element designed in such a way to enable use of the apparatus, logic, hardware, and/or element in a specified manner. Note as above that use of to, capable to, or operable to, in one embodiment, refers to the latent

state of an apparatus, logic, hardware, and/or element, where the apparatus, logic, hardware, and/or element is not operating but is designed in such a manner to enable use of an apparatus in a specified manner.

[0097] A value, as used herein, includes any known representation of a number, a state, a logical state, or a binary logical state. Often, the use of logic levels, logic values, or logical values is also referred to as 1's and 0's, which simply represents binary logic states. For example, a 1 refers to a high logic level and 0 refers to a low logic level. In one embodiment, a storage cell, such as a transistor or flash cell, may be capable of holding a single logical value or multiple logical values. However, other representations of values in computer systems have been used. For example the decimal number ten may also be represented as a binary value of 1010 and a hexadecimal letter A. Therefore, a value includes any representation of information capable of being held in a computer system.

[0098] Moreover, states may be represented by values or portions of values. As an example, a first value, such as a logical one, may represent a default or initial state, while a second value, such as a logical zero, may represent a non-default state. In addition, the terms reset and set, in one embodiment, refer to a default and an updated value or state, respectively. For example, a default value potentially includes a high logical value, i.e. reset, while an updated value potentially includes a low logical value, i.e. set. Note that any combination of values may be utilized to represent any number of states.

[0099] The embodiments of methods, hardware, software, firmware or code set forth above may be implemented via instructions or code stored on a machine-accessible, machine readable, computer accessible, or computer readable medium which are executable by a processing element. A non-transitory machine-accessible/readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine, such as a computer or electronic system. For example, a non-transitory machine-accessible medium includes random-access memory (RAM), such as static RAM (SRAM) or dynamic RAM (DRAM); ROM; magnetic or optical storage medium; flash memory devices; electrical storage devices; optical storage devices; acoustical storage devices; other form of storage devices for holding information received from transitory (propagated) signals (e.g., carrier waves, infrared signals, digital signals); etc., which are to be distinguished from the non-transitory mediums that may receive information there from.

[0100] Instructions used to program logic to perform embodiments of the invention may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions may be distributed via a network or by way of other

computer readable media. Thus a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium includes any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer)

[0101] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0102] In the foregoing specification, a detailed description has been given with reference to specific exemplary embodiments. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense. Furthermore, the foregoing use of embodiment and other exemplarily language does not necessarily refer to the same embodiment or the same example, but may refer to different and distinct embodiments, as well as potentially the same embodiment.

[0103] Some portions of the detailed description are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined,

compared and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers or the like. The blocks described herein may be hardware, software, firmware or a combination thereof.

[0104] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as “defining,” “receiving,” “determining,” “issuing,” “linking,” “associating,” “obtaining,” “authenticating,” “prohibiting,” “executing,” “requesting,” “communicating,” or the like, refer to the actions and processes of a computing system, or similar electronic computing device, that manipulates and transforms data represented as physical (e.g., electronic) quantities within the computing system's registers and memories into other data similarly represented as physical quantities within the computing system memories or registers or other such information storage, transmission or display devices.

[0105] The words “example” or “exemplary” are used herein to mean serving as an example, instance or illustration. Any aspect or design described herein as “example” or “exemplary” is not necessarily to be construed as preferred or advantageous over other aspects or designs. Rather, use of the words “example” or “exemplary” is intended to present concepts in a concrete fashion. As used in this application, the term “or” is intended to mean an inclusive “or” rather than an exclusive “or.” That is, unless specified otherwise, or clear from context, “X includes A or B” is intended to mean any of the natural inclusive permutations. That is, if X includes A; X includes B; or X includes both A and B, then “X includes A or B” is satisfied under any of the foregoing instances. In addition, the articles “a” and “an” as used in this application and the appended claims should generally be construed to mean “one or more” unless specified otherwise or clear from context to be directed to a singular form. Moreover, use of the term “an embodiment” or “one embodiment” or “an implementation” or “one implementation” throughout is not intended to mean the same embodiment or implementation unless described as such. Also, the terms “first,” “second,” “third,” “fourth,” etc. as used herein are meant as labels to distinguish among different elements and may not necessarily have an ordinal meaning according to their numerical designation.

Claims

What is claimed is:

1. A system comprising:
 - computer-readable storage medium storing instructions and a database having a
 - 5 plurality of frame fingerprints associated with media programs being streamed to a plurality of channels;
 - at least one processing device to execute the instructions to:
 - receive, from a media device, a query fingerprint of content being consumed,
 - the query fingerprint including an ordered sequence of frames;
 - 10 match the query fingerprint with the plurality of frame fingerprints stored in the database by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames, wherein to match comprises to populate a binary tree structure with respective distances;
 - detect a channel change from a first channel to a second channel responsive to
 - 15 locating a zero value within a threshold number of children nodes from a root node when traversing down the binary tree structure; and
 - determine an identity of the second channel by identifying a channel streaming content with frames matching those at an end of the ordered sequence of frames of the query fingerprint.
- 20 2. The system of claim 1, wherein to detect the channel change, the processing device to locate the zero value at a first node of a right or a left sub-tree positioned below the root node.
3. The system of claim 1, wherein to detect the channel change, the processing device to locate the zero value at a second node of a right or a left sub-tree positioned below the root node.
- 25 4. The system of claim 1, wherein to populate the binary tree structure with respective distances, the at least one processing device further to:
 - populate leaf nodes of the binary tree structure with values of the respective distances chronologically from one side to an opposite side; and
 - populate parent nodes with values derived from adding values of children nodes
 - 30 connected to the parent nodes.

5. The system of claim 1, where the at least one processing device further to notify an advertisement (“ad”) targeter of the identity of the second channel being watched in response to detecting the channel change so that the ad targeter can adjust targeting media to content streamed to the second channel.
- 5 6. The system of claim 5, wherein to adjust targeting media to the content streamed to the second channel comprises to update an overlay or replacement advertisement to, respectively, overlay on top of or to replace an advertisement streamed to the second channel.
7. The system of claim 5, wherein the ad targeter is further to:
sends a notification of the channel change and the identity of the second channel to an
10 ad server located across a network; and
receive, from the ad server, a replacement ad or overlay content to target commercial slots of the second channel.
8. A method comprising:
receiving, from a media device, a query fingerprint of content being consumed, the
15 query fingerprint including an ordered sequence of frames;
matching, using at least one processing device, the query fingerprint with a plurality of frame fingerprints stored in a database in memory by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames, comprising:
20 populating leaf nodes of the binary tree structure with values of the respective distances chronologically from one side to an opposite side of the binary tree structure;
populating parent nodes with values derived from adding values of children nodes connected to the parent nodes of the binary tree structure; and
25 detecting a value of or lower than a preset maximum threshold at a root node of the binary tree structure as indicating a match; and
notifying content manager logic of a channel corresponding to the matching query fingerprint so that the content manager logic can target media to content streamed to the channel.
- 30 9. The method of claim 8, wherein the preset maximum threshold is zero, further comprising determining an identity of the second channel by identifying a channel streaming

content with frames matching those at an end of the ordered sequence of frames of the query fingerprint.

10. The method of claim 8, wherein the determined channel is a first channel, the method further comprising:

5 detecting a channel change from the first channel to a second channel responsive to locating a low value within a threshold number of children nodes from the root node when traversing down the binary tree structure, wherein the low value is at or below the preset maximum threshold; and

notifying the content manager logic of the channel change to the second channel.

10 11. The method of claim 10, wherein the preset maximum threshold is zero.

12. The method of claim 10, wherein detecting the channel change comprises locating the low value at a first node of a right or a left sub-tree positioned below the root node.

13. The method of claim 10, wherein detecting the channel change comprises locating the low value at a second node of a right or a left sub-tree positioned below the root node.

15 14. The method of claim 10, further comprising:

sending a notification of the second channel to an ad server located across a network;

and

receiving, from the ad server, a replacement ad or overlay content to target commercial slots of the second channel.

20 15. A server comprising:

computer-readable storage medium storing instructions and a database having a plurality of frame fingerprints associated with media programs being streamed to a plurality of channels;

at least one processing device to execute the instructions to:

25 match the query fingerprint with a plurality of frame fingerprints stored in the database by determining a distance between respective frame fingerprints and corresponding ones of the ordered sequence of frames, comprising to:

populate leaf nodes of the binary tree structure with values of the respective distances chronologically from one side to an opposite side of the

30 binary tree structure;

populate parent nodes with values derived from adding values of children nodes connected to the parent nodes of the binary tree structure; and detect a value of or lower than a preset maximum threshold at a root node of the binary tree structure as indicating a match; and

5 notify content manager logic of a channel corresponding to the matching query fingerprint so that the content manager logic can target media to content streamed to the channel.

16. The server of claim 15, wherein the determined channel is a first channel, the at least one processing device further to:

10 detect a channel change from the first channel to a second channel responsive to locating a low value within a threshold number of children nodes from the root node when traversing down the binary tree structure, wherein the low value is at or below the preset maximum threshold;

15 determine an identity of the second channel by identifying a channel streaming content with frames matching those at an end of the ordered sequence of frames of the query fingerprint; and

notify the content manager logic of the channel change and the identity of the second channel.

17. The server of claim 16, wherein the preset maximum threshold is zero.

20 18. The server of claim 17, wherein to detect the channel change comprises to locate the zero value at a first node of a right or a left sub-tree positioned below the root node.

19. The server of claim 17, wherein to detect the channel change comprises to locate the zero value at a second node of a right or a left sub-tree positioned below the root node.

25 20. The server of claim 16, wherein the at least one processing device further to:
send a notification of the second channel to an ad server located across a network; and receive, from the ad server, a replacement ad or overlay content to target commercial slots of the second channel.




FIG. 1




FIG. 2




FIG. 3

4/9

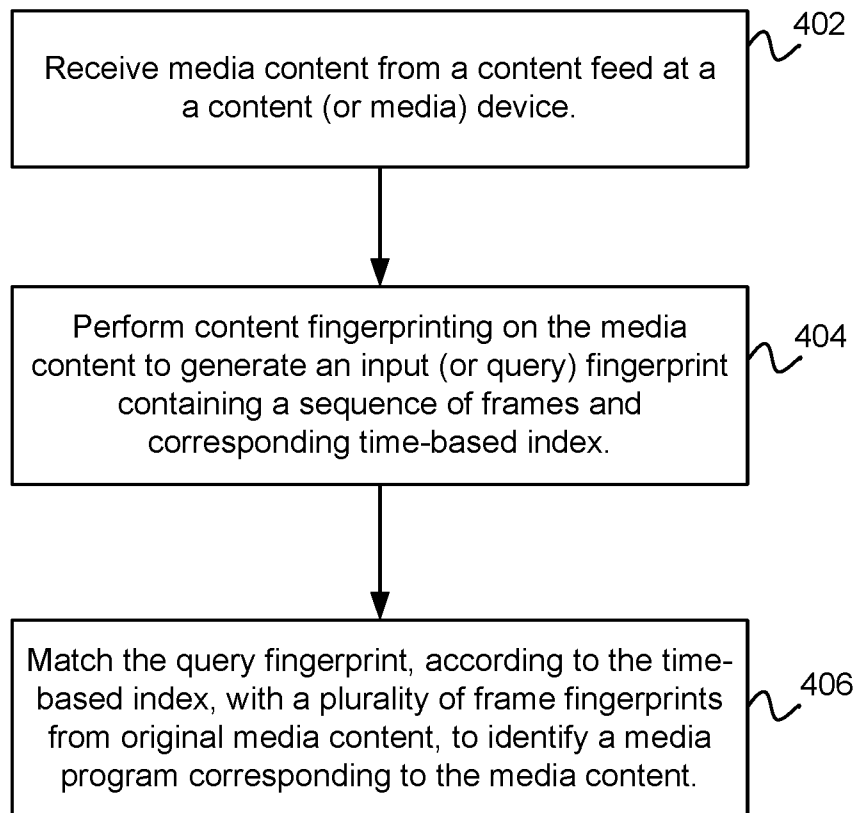
400

FIG. 4

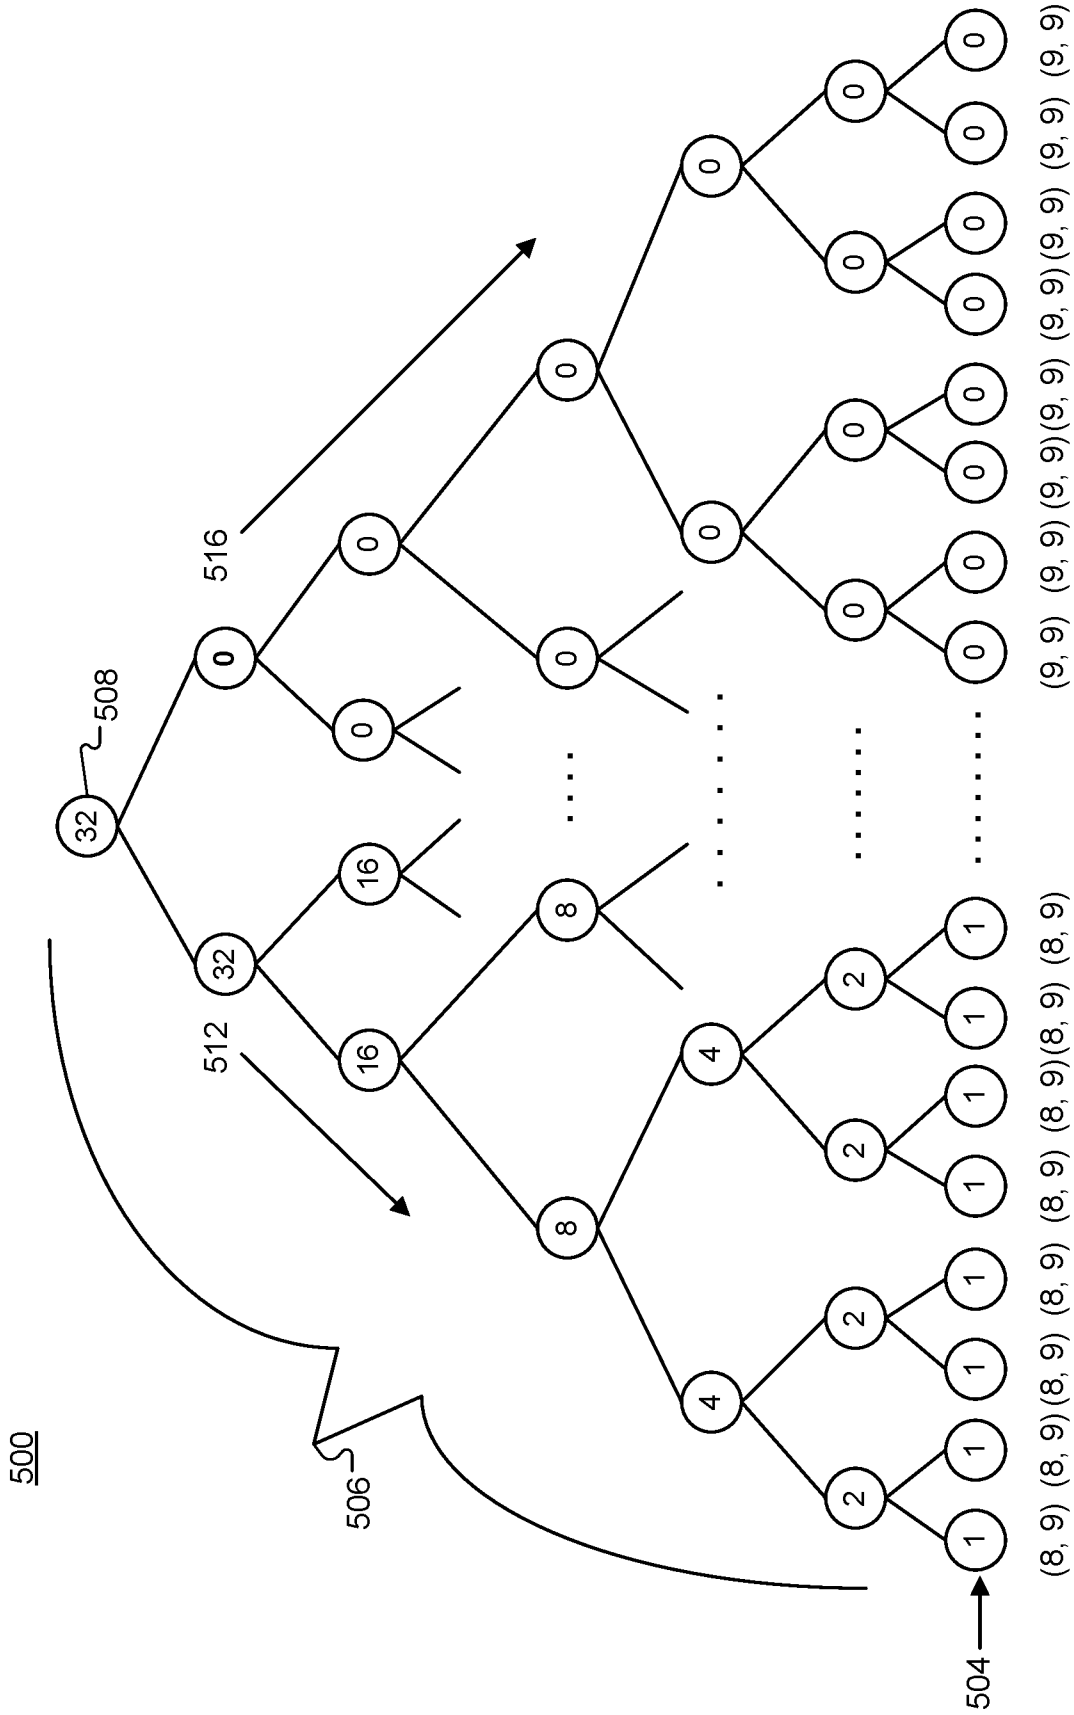


FIG. 5

8/9

800

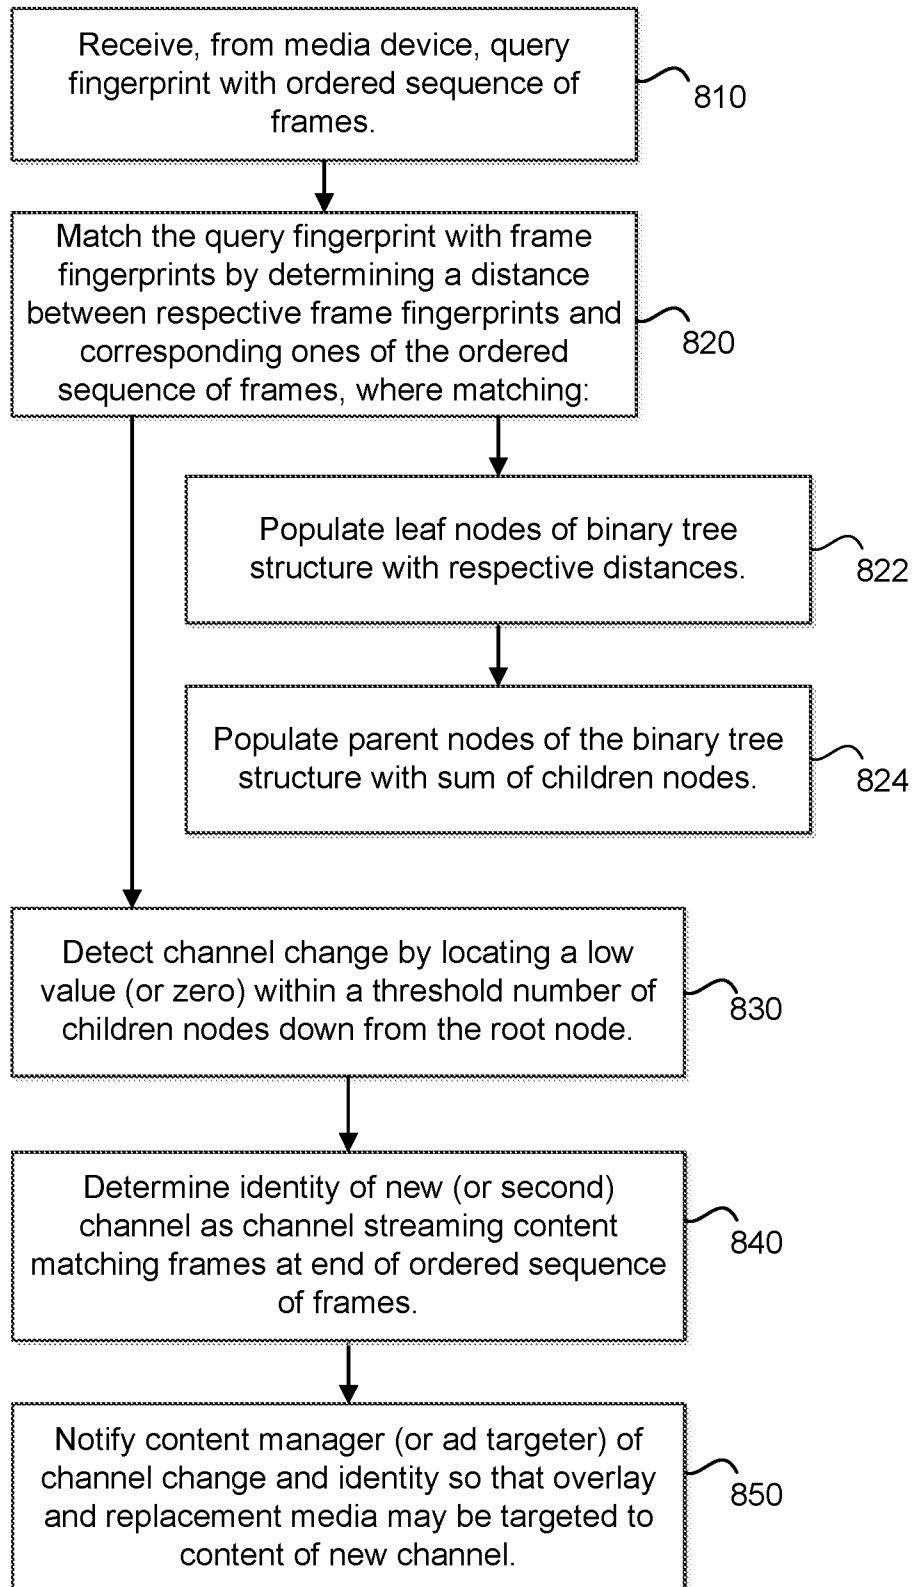


FIG. 8

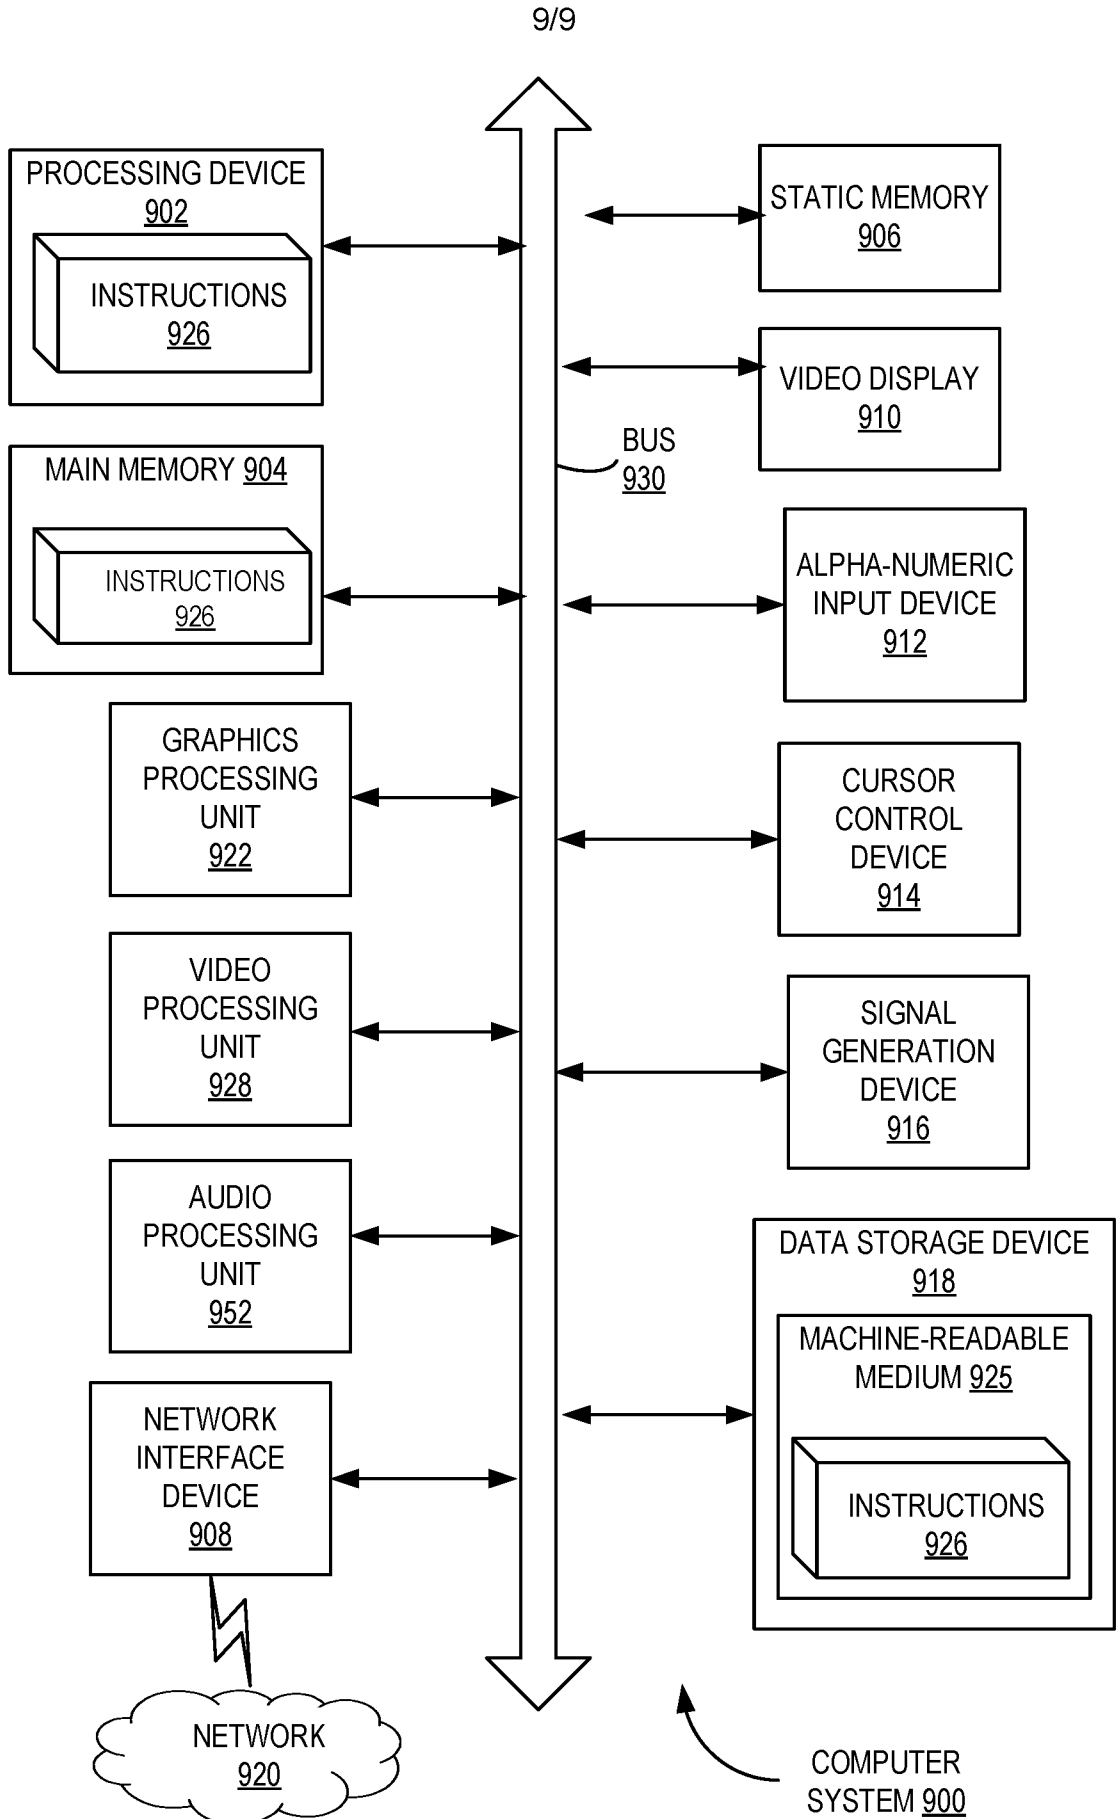


FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 16/34401

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - H04H 60/37 (2016.01)

CPC - H04N21/44008

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
CPC: H04N21/44008; IPC(8): H04H 60/37 (2016.01); USPC: 725/19Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
CPC: H04N21/44008, H04H60/37, H04H60/56, H04H60/59, H04N21/8352; USPC: 725/19, 725/20, 725/18 (keyword limited; terms below)Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PatBase; Google Scholar
Search Terms: frame fingerprint media program stream channel query match database binary tree distance root child zero

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2013/0104179 A1 (SHAH) 25 April 2013 (25.04.2013), entire document, especially abstract and para [0022], [0025], [0034], [0036], [0039], [0043], [0045]-[0046], [0049]-[0050], [0052]-[0053], [0063], [0075], [0107]-[0108], [0130]-[0134], [0143]-[0144], [0166], [0174]-[0176], Fig. 1, Fig. 3.	1-20
A	US 2005/0249080 A1 (FOOTE et al.) 10 November 2005 (10.11.2005), entire document.	1-20
A	US 2015/0073967 A1 (KATSUYAMA et al.) 12 March 2015 (12.03.2015), entire document.	1-20

 Further documents are listed in the continuation of Box C.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier application or patent but published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

20 July 2016 (20.07.2016)

Date of mailing of the international search report

22 AUG 2016

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774