

(19)



(11)

EP 3 146 446 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
26.12.2018 Bulletin 2018/52

(51) Int Cl.:
G06F 17/30 ^(2006.01) **H04H 60/37** ^(2008.01)
H04N 21/234 ^(2011.01) **G06K 9/00** ^(2006.01)
H04N 21/4402 ^(2011.01) **H04L 29/06** ^(2006.01)

(21) Application number: **15795649.1**

(86) International application number:
PCT/US2015/031167

(22) Date of filing: **15.05.2015**

(87) International publication number:
WO 2015/179245 (26.11.2015 Gazette 2015/47)

(54) MEDIA STREAM CUE POINT CREATION WITH AUTOMATED CONTENT RECOGNITION

ERZEUGUNG EINES MEDIENDATENSTROM-EINSATZPUNKTS MIT AUTOMATISCHER INHALTSERKENNUNG

CRÉATION DE POINT DE REPÈRE DANS UN FLUX MULTIMÉDIA À L'AIDE DE LA RECONNAISSANCE DE CONTENU AUTOMATISÉE

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

- **IKEZOYE, Vance**
Los Gatos, California 95032 (US)
- **AVERY, Dave**
Los Gatos, California 95032 (US)
- **WOLD, Erling**
Los Gatos, California 95032 (US)
- **EDEL, Alan**
Los Gatos, California 95032 (US)

(30) Priority: **21.05.2014 US 201462001240 P**
30.09.2014 US 201414502791

(43) Date of publication of application:
29.03.2017 Bulletin 2017/13

(74) Representative: **WP Thompson**
8th Floor
1 Mann Island
Liverpool L3 1BP (GB)

(73) Proprietor: **Audible Magic Corporation**
Los Gatos, CA 95032 (US)

(72) Inventors:

- **SCHREMPP, Jim**
Los Gatos, California 95032 (US)
- **WILLIAMS, Rob**
Los Gatos, California 95032 (US)

(56) References cited:

WO-A1-2014/004955	US-A1- 2002 097 979
US-A1- 2004 189 873	US-A1- 2006 195 860
US-A1- 2008 189 753	US-A1- 2009 300 145
US-A1- 2014 013 941	US-B1- 6 698 020

EP 3 146 446 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description**TECHNICAL FIELD**

[0001] This disclosure relates to the field of media content identification, and in particular to media stream cue point creation with automated content recognition.

BACKGROUND

[0002] A large and growing population of users enjoys entertainment through the consumption of media items, including electronic media, such as digital audio and video. Users employ various electronic devices to consume such media items. Among these electronic devices are electronic book readers, cellular telephones, personal digital assistants (PDAs), portable media players, tablet computers, electronic pads, netbooks, desktop computers, notebook computers, and the like. Media content suppliers provide media items to consumers through a variety of means. Some media content suppliers deliver media items in single units, others in bulk, others through some linear streaming experience, and others by providing playlists or manifest files of media segments or chunks. Some media content suppliers employ a variety of delivery means. Much electronic media content is offered to users free of charge and subsidized through advertising. This advertising can include advertising messages conveyed using audio, text, logos, animations, videos, photographs or other graphics. Some examples of advertisements include banner ads, frame ads, pop-up ads, floating ads, expanding ads, interstitial ads, video ads, audio ads and text ads. These advertisements can be presented to a user in a variety of ways including on a web-page, through social media applications, on a mobile device, or in an audio or video advertisement slot between or within segments of regular programming or even on top of regular programming.

[0003] Reference is made to US 2004/189873 A1, which discloses a television signal substitution system that replaces known video segments such as advertisements with selected replacement advertisements. Fingerprint data of known advertisements can be stored in a fingerprint database. When new fingerprint data is available, the fingerprint data can be automatically or manually transmitted to subscribers. Various techniques can be used to identify advertisements based on the fingerprint data.

[0004] Reference is also made to US 2006/195860 A1, which discloses a method for detecting and acting on a known video entity within a video stream. The method includes receiving a video stream and continually creating statistical parameterized representations for windows of the video stream. The statistical parameterized representation windows are continually compared to windows of a plurality of fingerprints. Each of the plurality of fingerprints includes associated statistical parameterized representations of a known video entity. A known video

entity in the video stream is detected when a particular fingerprint of the plurality of fingerprints has at least a threshold level of similarity with the video stream.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the present invention, which, however, should not be taken to limit the present invention to the specific embodiments, but are for explanation and understanding only.

Figure 1a is a block diagram illustrating a media player with a local media stream cue point creator, according to an embodiment.

Figure 1b is a block diagram illustrating a media player with a remotely connected media stream cue point creator, according to an embodiment.

Figure 1c is a block diagram illustrating a detailed view of a media stream cue point creator for use in a media player, according to an embodiment.

Figure 2a is a block diagram illustrating a media distribution point with a local media stream cue point creator, according to an embodiment.

Figure 2b is a block diagram illustrating a detailed view of a media stream cue point creator for use in a media distribution point, according to an embodiment.

Figure 2c is a block diagram illustrating a media stream cue point creator that is distributed across multiple locations, according to an embodiment.

Figure 3a is a diagram illustrating a segmented media stream, according to an embodiment.

Figure 3b is a diagram illustrating a segment information data store, according to an embodiment.

Figure 4 is a flow diagram illustrating a cue point creation method, according to an embodiment.

Figure 5 is a flow diagram illustrating a segment identification method, according to an embodiment.

Figure 6 is a flow diagram illustrating a segment replacement method, according to an embodiment.

Figure 7 is a block diagram illustrating an exemplary computer system, according to an embodiment.

DETAILED DESCRIPTION

[0006] The invention is defined by a method according to claim 1, a system according to claim 5 and a non-transitory computer-readable medium according to claim 9. Embodiments are described for media stream cue point creation with automated content recognition. In one embodiment, a media stream, including audio and/or video streams, has multiple different segments. These segments may include for example, content segments and advertisements. In this disclosure the term "content segment" is understood to mean any content in the media stream, whether advertisement, program, public service

announcement, static video image, or even a blank video or audio segment. "Content segment" may also mean a sub-segment of an advertisement, program, public service announcement, etc. In one embodiment, a party, such as a media content provider or media player manufacturer, may wish to replace one or more content segments of the media stream with different content segments. Sometimes programming content or advertisements are only permitted to be delivered in certain geographic areas. A media delivery service may wish to replace or remove certain content segments before delivery to certain geographic areas or receivers. In some cases an advertisement may have a sub-segment somewhere within it that delivers a very specific message. A media delivery service may wish to replace that sub-segment of the advertisement with another content segment which is more appropriate for the receiver. In one embodiment, a media content provider or media delivery service may include, for example, an Internet service provider, cable company, production company, web site provider, social media web site, music or video web site, telecommunications company, over-the-top content provider, cloud DVR company, or other provider. For example, a media content provider may wish to replace the advertisements in a media stream it receives with different advertisements before it provides the media stream to its users, (e.g., customers or subscribers). For example, a media content provider may wish to replace a content segment with another content segment before it provides the media stream to its customers or subscribers. In another embodiment, a party may wish to overlay some content on top of some content segment of the rendered media stream. Such overlays may be a banner ad or additional information of interest to a user. A party may wish to display sports scores over media content, but remove the overlay whenever an advertisement is playing. Alternatively, a party may wish to display additional information about an advertisement, but remove the overlay as soon as the advertisement ends. A party may wish to replace a content segment with a static video image segment. A party may wish to remove a content segment from a stream. In one embodiment, a party may wish to take actions to facilitate the augmentation of the media stream to support users with one or more disabilities. For example, the party may provide different levels of closed captioning information while content is playing. The different levels may include basic information that may be useful to a hearing person or additional information that may be useful to a deaf person. The additional information may be synchronized with the media stream, and the media stream may not have appropriate or accurate cue points.

[0007] In another embodiment, a media content provider or device may send a signal to a second device, such as a tablet or Set Top Box, which then takes some action, such as to display complementary content or advertising, present the user with a survey, or allow the user to comment on the stream. In another embodiment a par-

ty may wish to know which content segment a user is experiencing. Knowing which content segments a user is or has experienced over time, can allow a party to better select future advertising or media content for this user. Unless the media stream has cue points, such as SCTE 35 codes, embedded in the stream to designate the locations of the content segments, the media content provider may have no way of telling where or when the content segments are located in the stream. As a result, it may not be able to take action where a specific content segment is located in the media stream. In some cases, even when cue points or other indicators are present in the media stream, these may not be accurate or reliable, which can result in errors during segment replacement. In addition, embedded cue points, such as SCTE 35, are most often provided to indicate the boundary points between media segments. The media content provider may wish to trigger actions that occur inside a media segment. For instance, a media content provider may wish to trigger an action whenever a certain song is present in the audio media stream, when a certain language is detected, when a certain sound (e.g., a phone ringing, a car door closing) is detected, or when a class of sounds (e.g., dogs barking, explosions, footsteps) is detected. Or a media content provider may wish to trigger an action whenever a certain object, such as a can of soda, or a particular person is present in the video media stream. The media content provider may wish to only replace certain content segments. In many cases, embedded cue points do not supply metadata or links describing the particular content segments that would be needed to replace only certain content segments. In another embodiment a party may wish to replace a sub-segment of media content with another segment that is more appropriate for the viewer. For example, a party may wish to replace the last part of an advertisement that offers a specific discount with a different segment offering a different discount. In another example, a party may wish to completely remove some part of an advertisement that offers a specific discount or mentions a specific retailer. In another example, a party may wish to provide information about the location and contact information of a local retailer during a nationally broadcast advertisement.

[0008] In one embodiment, a media stream cue point creator, located within a media player, at a media distribution point (such as a cable operator headend, or a node of a content delivery network, or a radio or TV station rebroadcasting a syndicated program), or on a network-connected server, uses automated content recognition to identify particular content segments in a media stream that has no prior cue points or indicators, and generates cue points associated with the identified content segments. The media stream cue point creator, or other parties, may then locate a given segment in the media stream using the generated cue points and perform some operation with respect to the segment. For example, the media stream cue point creator may identify an advertisement in a media stream and allow a content provider

to replace the advertisement with a different advertisement, block the advertisement, overlay content the advertisement, replace the advertisement with other media content, or remove the advertisement.

[0009] In one embodiment, the media stream cue point creator receives an incoming media stream and generates an audio and/or video fingerprint of the media stream. The media stream cue point creator compares the fingerprint to a database of known media segments to determine if there is a match using certain techniques. One such technique is described in U.S. Patent no. 5,918,223, issued June 29, 1999. If the fingerprint matches the fingerprint of a known media segment in the database, or at least a portion of the known fingerprint, the media stream cue point creator generates a cue point entry in a segment information data store. One of skill in the art would recognize that many techniques are known for detecting specific objects within the images of a video media stream. Using one of these techniques, when an object is identified to be present in the media stream, the cue point creator generates a cue point entry in a segment information data store. The cue point entry may include information, such as, an identifier of the media stream, an identifier of the segment, a start time/value of the segment, advertising related information, a stop time/value of the segment, and other metadata describing the segment. Either prior to or during playback of the media stream, the media stream cue point creator identifies a cue point in the media stream and triggers some operation, such as a segment replacement. During segment replacement, the media stream cue point creator may identify playback context information, such as when, where, by who, etc. playback is being performed, and either provide that information to another party, or use that information itself, to select a replacement segment (e.g., new advertisement). The media stream cue point creator, or some other device, may then replace the original segment in the media stream with the new segment and continue stream playback.

[0010] In one embodiment, the media stream cue point creation operations using automated content recognition, which are described herein, allow a party to recognize and identify particular segments in a media stream which previously had no cue points or other indicators of the various segments. In other embodiments, the techniques described herein may be used on a media stream that has some cue points, but is missing others, or has cue points present that are unreliable or inaccurate. This allows the parties to perform a variety of operations on the segments, identified now by the created cue points, such as segment replacement or overlay. The ability to identify and replace and report on viewer consumption of segments in a media stream can lead to many advantages, such as improved content, increased advertising revenue, and others.

[0011] Figure 1a is a block diagram illustrating a media player 100 with a local media stream cue point creator 110, according to an embodiment. In one embodiment,

media player 100 may be any type of computing device including a server computer, gateway computer, desktop computer, laptop computer, mobile communications device, cell phone, smart phone, hand-held computer, tablet computer, set top box (STB), digital video recorder (DVR), wearable computer, or similar computing device. Media player 100 may be variously configured with different features to enable the recording and viewing of content and resources, and the execution of one or more applications. Media player 100 may be embodied, for example, by computer system 700 of Figure 7.

[0012] In one embodiment, media player 100 includes various modules and data stores to enable the recording and playback of media content. For example, media player 100 may include DVR record controller 102, media data store 104 and DVR playback controller 106. In one embodiment, DVR record controller 102 receives an incoming media stream and may optionally, at the instruction of a user or program running in media player 100, record a copy of the media stream in media data store 104. A copy of the media stream may be stored temporarily (while media stream cue point creator 110 performs associated processing) before playback, or may be stored for a longer period of time for playback at a later time or date. Media data store 104 may include one or more mass storage devices which can include, for example, flash memory, magnetic or optical disks, or tape drives; read-only memory (ROM); random-access memory (RAM); erasable programmable memory (e.g., EPROM and EEPROM); flash memory; or any other type of storage medium. The recorded media stream may be stored in the MPEG format or in any other suitable format. DVR playback controller 106 manages playback of recorded media from media data store 104 on an internal or externally connected display device, such as a directly connected display or a display at a remote terminal connected over a network (e.g., a cloud DVR). DVR playback controller 106 may start, stop, pause, or continue playback of recorded media at the instruction of a user or other program running in media player 100.

[0013] In one embodiment, media player 100 further includes media stream cue point creator 110 and associated data store 120. Data store 120 may be a standalone data store or may share a storage device with media data store 104. Media stream cue point creator 110 uses automated content recognition to identify particular segments in a media stream stored in media data store 104, and generate cue points associated with the identified segments. The media stream cue point creator 110, or another module, may then locate a given segment in the media stream using the generated cue points and perform some operation on the segment. For example, the media stream cue point creator 110 may identify an original advertisement in a stored media stream and replace the original advertisement with a new advertisement. This replacement may be done as the media stream is being viewed. In the case of recorded media streams this replacement may also be done before the

media stream is viewed. In another embodiment, an input media stream is received and the data is sent to media stream cue point creator 110 and to an output device, such as screen or speaker, at the same time. Media stream cue point creator 110 may take a short period of time to perform its processing (i.e., identifying a content segment and creating a cue point). In the meantime, the content segment may have already started playing on the output device. Media stream cue point creator 110 may trigger an overlay to be put up during the content segment. Media stream cue point creator 110 may predict when the content segment will end and trigger the removal of the overlay, right as the content segment ends. In this case, media data store 104 may not be involved. In another embodiment, media data store 104 may keep a small buffer of the media stream to delay the playback, perhaps to give the user a "skip ahead" or "skip back" function. Additional details of media stream cue point creator 110 are provided below.

[0014] Figure 1b is a block diagram illustrating a media player 100 with a remotely connected media stream cue point creator 110, according to an embodiment. In one embodiment, media player 100 may not have an internal media stream cue point creator, but rather may be connected to a remote server 150 over a network 140. Network 140 can be a local area network (LAN), a wireless network, a telephone network, a mobile communications network, a wide area network (WAN), such as the Internet, near field communication link, Bluetooth, or similar communication system. Server 150 may include a network-accessible server-based functionality, various data stores, and/or other data processing equipment. Server 150 may be implemented by a single machine or a cluster of machines. Server 150 may be hosted, for example, by computer system 700 of Figure 7.

[0015] In one embodiment, server 150 hosts media stream cue point creator 110 and associated data store 120. When running on remote server 150, media stream cue point creator 110 and associated data store 120 may perform substantially the same as when they are hosted locally within media player 100. In this embodiment, communication between media stream cue point inserter 110 and the other components of media player 100 (e.g., DVR record controller 102, media data store 104, DVR playback controller 106) may occur across network 140. One of skill in the art would understand that the description of media stream cue point creator 110 herein should apply equally whether media stream cue point creator 110 is located locally within media player 100, remotely at server 150 or any other location. In other embodiments, the functionality of media stream cue point creator 110 may be divided across multiple devices. For example, a portion within media player 100 may create fingerprints of the media stream and send them to another portion located on server 150, where additional analysis is performed.

[0016] Figure 1c is a block diagram illustrating a detailed view of a media stream cue point creator 110 for

use in a media player 100, according to an embodiment. In one embodiment, media stream cue point creator 110 includes segment identifier module 112, segment cue point alert module 114 and segment pre-fetch module 116. This arrangement of modules and components may be a logical separation, and in other embodiments, these modules or other components can be combined together or separated in further components, according to a particular embodiment. In one embodiment, a data store connected to media stream cue point creator 110 includes segment reference data store 122, segment information data store 124, and replacement segment data store 126. In one embodiment, media player 100 may include both media stream cue point creator 110 and data store 120. In another embodiment, the data store 120 may be external to media stream cue point creator 110 and may be connected over a network or other connection. In other embodiments, media player 100 may include different and/or additional components that are not shown to simplify the description.

[0017] In one embodiment, segment identifier module 112 identifies segments of a media stream stored in media data store 104 using automated content recognition. In one embodiment, there may be a minor delay when automated content recognition is performed on a live data stream. This delay, however, may be inconsequential to the end viewer and may run concurrently with normal broadcast delays. Segment identifier module 112 generates an audio and/or video fingerprint of at least a portion of the media stream. For example, segment identifier module 112 may divide a stored or buffered audio stream into overlapping frames (e.g., approximately 0.1 seconds in length with an overlap of approximately 0.025 seconds) and extract acoustic or perceptual features of each frame. Perceptual features may include characteristics of sounds in terms of commonly understood physical characteristics, such as brightness, pitch and loudness. Other characteristics may include bass, bandwidth, and Mel-frequency Cepstral Coefficients (MFCCs) at periodic intervals over the length of the frame. Segment identifier module 112 may generate an n-dimensional feature vector for each frame and compare the feature vector to stored feature vectors corresponding to known segments from segment reference data store 122. Segment reference data store 122 may store feature vectors, fingerprints, complete copies, or other representations of known segments of media content. Data store 122 may store metadata used to identify the content represented by the feature vector. For example, segment reference data store 122 may include representations of known advertisements which could be identified and potentially removed from the media streams in media data store 104. In one embodiment, segment identifier module 112 may extract perceptual features of the video image stream. These perceptual features may include commonly understood characteristics such as brightness, color, luminosity, etc. Other characteristics may include lines, shapes, depth features, occlusion, and other features

used in object detection. Segment reference data store 122 may contain video fingerprints of known content segments, as well as models and classes of objects. In one embodiment, segment identifier module 112 may compare characteristics of the video image stream to these fingerprints in order to identify media segments within the video image stream. In another embodiment, segment identifier module 112 may use object models and classes to identify objects within the video image stream or the audio stream. When a given object is identified, this may signify a portion of interest in the media stream. For example, a certain object may appear at the beginning of a content segment, and detection of that object may be used to identify the segment and trigger segment replacement or some other action. Segment reference data store 122 may be populated with information obtained directly from advertisers, other content creators, or a content distributor, or with information obtained from monitoring previous media streams. Segment reference data store 122 may be populated with models or classes describing specific objects or sounds (e.g. a certain person, a certain model of automobile, a certain building, or a certain corporate logo) or generic objects or sounds (e.g. a soda can, an automobile, a building). These classes and models may be created specifically for this application or obtained from a library of object classes and models. In one embodiment, segment identifier module 112 or some other module may monitor live media streams and locate repeating segments of media using acoustic, visual or other perceptual features or other data, such as closed captioning text, etc. Segment identifier module 112 may flag repeating segments of a specified length (e.g., 10 seconds, 30 seconds, 45 seconds, 60 seconds) for further investigation. In one embodiment, a human curator may verify that the flagged segments are advertisements or other interesting pieces of content and may add them to segment reference data store 122. Additional details about identifying known segments for segment reference data store 122 may be found in U.S. Patent No. 7,877,438 issued January 25, 2011.

[0018] In one embodiment, upon identifying a segment that matches a known segment or object in segment reference data store 122, segment identifier module 112 creates a cue point associated with the identified segment. In one embodiment, segment identifier module 112 creates a corresponding entry in segment information data store 124. In one embodiment, the cue point entry may include an identifier of the media stream from which the segment was identified, an identifier of the segment itself, a start time/value of the segment within the media stream, a stop time/value of the segment, and other information or metadata associated with the segment, such as an indication of what product the advertisement is for, the subject matter of the advertisement, a recommended audience for the advertisement, the name of a person, the location of a building, a description of the object, etc. One example of segment information data store 124 is described below with respect to Figure 3b. In another

embodiment, segment identifier module 112 may insert an indication of the cue point directly into the media stream itself. For example, segment identifier module 112 could send instructions to segment cue point alert module 114, or some other module, which could embed SCTE 35 codes in the media stream or media file.

[0019] In one embodiment, segment cue point alert module 114 generates an alert signaling replacement of a segment in the media stream identified by the cue point entry in segment information data store 124. Segment cue point alert module 114 identifies a cue point for the media stream. This may occur during playback of the media stream by DVR playback controller 106 or prior to playback when the stream is stored as a file in media data store 104. In one embodiment, segment cue point alert module 114 receives an indication of playback from DVR playback controller 106 and compares the location to the cue points in segment information data store 124. When a cue point is reached, segment cue point alert module 114 may trigger an operation, such as segment replacement. In other embodiments, segment cue point alert module 114 may trigger an operation several seconds (or some other period of time) prior to the cue point being reached. In other embodiments, other operations that may be triggered include the overlay of complementary advertisement data, the overlay of complementary program data (such as sports scores), the removal of an overlay that was put in place earlier, such as a persistent ad overlay that taken down as soon as recognized programming begins or a program overlay that is taken down as soon as an ad segment begins, the removal of a segment, or the insertion of a new segment. In other embodiments, the segment cue point alert module 114 may send a trigger to another device such as a tablet or smartphone which then takes action. In other embodiments, segment cue point module 114 may send an alert to a data store which monitors viewing habits.

[0020] In one embodiment, segment cue point alert module 114 identifies playback context information for the media stream. The playback context information may include, for example, a time when playback occurs, a location where playback occurs, an indication of who requested playback, or other information. Using the playback context information, segment cue point module 114 may request a new segment from replacement segment data store 126. In one embodiment, segment pre-fetch module 116 may request and receive content segments (e.g., advertisements, online videos, overlays, information) from an external source, such as media replacement server 160. Examples of replacement media server 160 may include an ad server, an online content provider, a social network, etc. Segment pre-fetch module 116 may populate replacement segment data store 126 with the received replacement segments. In one embodiment, segment pre-fetch module 116 may request segments based on the playback context information and media replacement server 160 may determine the appropriate segments to provide. In another embodiment, segment

pre-fetch module 116 may request generic segments from media replacement server 160 and segment cue point alert module 114 may make a selection of segments from replacement segment data store 126 based on the playback context information. In another embodiment, segment cue point alert module 114 may instruct DVR playback controller 106 to insert a media stream that is being streamed in real time from replacement media server 160. In one embodiment, segment cue point alert module 114 provides an instruction to DVR playback controller 106 to replace the original segment in the playing media stream with the replacement segment from replacement segment data store 126. In another embodiment, rather than replacing a segment, segment cue point alert module 114 may provide an instruction to DVR playback controller 106 to block a certain segment, remove a certain segment, or to perform some other operation. In another embodiment, DVR playback controller 106 may modify a media file in media data store 104.

[0021] Figure 2a is a block diagram illustrating a media distribution point 200 with a local media stream cue point creator 210, according to an embodiment. Media distribution point 200 may be embodied, for example, by computer system 700 of Figure 7. In one embodiment, media distribution point 200 may be a retransmission facility or a cable television headend, and may include a master facility for receiving television or other signals for processing and distribution over a cable television system or other communication system. The media distribution point may include electronic equipment used to receive and retransmit, audio, video or other signals over the local communication infrastructure.

[0022] In one embodiment, media distribution point 200 includes various modules and data stores to enable the broadcast or other distribution of media content. For example, media distribution point 200 may include encoder 202, origin server 204 and splicer 206. In one embodiment, encoder 202 receives an incoming media stream and encodes the media stream into various distribution formats. The incoming media stream, or input stream, may be received, for example, over a satellite downlink, high speed network link, or other communications system. The formats may include for example, standard definition, high definition, web format, or other formats, aspect ratios and/or pixel resolutions. Encoder 202 provides the various encoded media streams to origin server 204. In one embodiment, origin server 204 serves as the source of the media stream to the rest of the downstream network. For example, the origin server might host several different encodings of the same stream. The media stream might not be stored for long on the origin server 204, which may be serving as a buffer between the encoder 202 and the distribution network. In other embodiments, origin server 204 might store the streams for a longer period of time (e.g., a week or month). In other embodiments, origin server 204 might store content forever, as in the case of a Video On Demand origin server. Splicer 206 is responsible for per-

forming segment replacement of certain segments in the various media streams. For example, splicer 206 may be configured to remove a certain segment and insert a new segment, stop playback of a first stream, play a segment from a different stream, and resume playback of the first stream, or other operations.

[0023] In one embodiment, media distribution point 200 further includes media stream cue point creator 210 and associated data store 220. Data store 220 may include one or more mass storage devices which can include, for example, flash memory, magnetic or optical disks, or tape drives; read-only memory (ROM); random-access memory (RAM); erasable programmable memory (e.g., EPROM and EEPROM); flash memory; or any other type of storage medium. Media stream cue point creator 210 may be the same as or similar to media stream cue point creator 110 described above, such that it uses automated content recognition to identify particular segments and objects in a media stream that has no prior cue points or indicators, or inaccurate or unreliable cue points, and generates cue points associated with the identified segments. The media stream cue point creator 210, or another module, may locate a given segment in the media stream using the generated cue points and either perform, or alert some other module to perform, an operation on the segment. For example, the media stream cue point creator 210 may identify an original advertisement in a media stream and signal splicer 206 to replace the original advertisement with a new advertisement. For example, the media stream cue point creator 210 may identify a content segment and alert encoder 202 to place the appropriate cue markers, such as SCTE 35, into the encoded streams. Additional details of media stream cue point creator 210 are provided below.

[0024] Figure 2b is a block diagram illustrating a detailed view of a media stream cue point creator 210 for use in a media distribution point 200, according to an embodiment. In one embodiment, media stream cue point creator 210 is the same as or similar to media stream cue point creator 110 described above, and includes segment identifier module 112, segment cue point alert module 114 and segment pre-fetch module 116. In one embodiment, data store 220 is connected to media stream cue point creator 210 and includes segment reference data store 122, segment information data store 124, and replacement segment data store 126. In one embodiment, media distribution point 200 may include both media stream cue point creator 210 and data store 220. In another embodiment, the data store 220 may be external to media stream cue point creator 210 and may be connected over a network or other connection. In other embodiments, media distribution point 200 may include different and/or additional components which are not shown to simplify the description. One of skill in the art would understand that the description of media stream cue point creator 110 herein should apply equally to media stream cue point creator 210, as their operation is substantially the same. Accordingly, a description of the

operation of media stream cue point creator 210 is not repeated here so as not to obscure the present disclosure.

[0025] Figure 2c is a block diagram illustrating a media stream cue point creator that is distributed across multiple locations, according to an embodiment. In one embodiment, a network environment includes two separate entities, such as media content provider 250 and media player 260. Media content provider 250 may include, for example, a cable, satellite or broadcast television provider, an Internet service provider, an on-line media provider, or some other content provider. Media player 260 may be part of a user device configured to receive and playback media content received from media content provider, including for example, a set top box, an over-the-top box, etc. In one embodiment, media content provider 250 and media player 260 are located in different physical locations but are connected via a network, such as the Internet. In one embodiment, different components of the media stream cue point creator described above are distributed across media content provider 250 and media player 260. For example, segment identifier 112, segment reference data store 122, segment information data store 124 and segment cue point alert module 114 may be located at media content provider 250. These components may recognize content segments in an input media stream provided by encoder 202 and may add cue points to the input media stream. In one embodiment, segment prefetcher 116 and replacement segment data store 126 are located at media player 260. In one embodiment, action module 206 asks segment cue point alert module 114 if there are any upcoming cue points in the current stream. In response, segment cue point alert module 114 may provide an indication of the upcoming cue points. Action module 206 may detect a cue point and perform an action with respect to a certain segment (e.g., segment replacement using segment prefetcher 116 and replacement segment data store 126). In another embodiment, some other distribution of the components of the media stream cue point creator may be used.

[0026] Figure 3a is a diagram illustrating a media stream 300, according to an embodiment. The media stream 300 may be, for example, an audio stream, a video stream, a combination audio and video stream, or some other type of media stream. According to the processes described herein, segment identifier module 112 may identify particular content segments of the media stream 300. In one embodiment, the media stream 300, or stored media file, may begin at time 302 and have regular content segments between times 302-304, 306-308, and 310-312. Segment identifier module 112 may identify advertisement segments between times 304-306, 308-310, 312-314, and 314-316. In other embodiments, the identified segments may have some other lengths or be some other segment types. In other embodiments, some other unit of measurement, besides time, may be used to designate the start and end points of the various segments in media stream 300.

[0027] Figure 3b is a diagram illustrating a segment information data store 124, according to an embodiment. In one embodiment, upon identifying one or more segments in media stream 300 that match known segments in segment reference data store 122, segment identifier module 112 creates cue points associated with the identified segments. In one embodiment, segment identifier module 112 creates corresponding entries in segment information data store 124. In one embodiment, the cue point entries may include an identifier of the media stream 300 from which the segment was identified, an identifier of the segment itself 01, 02, 03, 04, a start time/value of the segment 304, 308, 312, 314 within the media stream, a stop time/value of the segment 306, 310, 314, 316, and other information or metadata associated with the segment, such as an indication of what product the advertisement is for, the subject matter of the advertisement, a recommended audience for the advertisement, an object found to be present, the name of a specific person, etc. In one embodiment, segment identifier module 112 may create an entry with a start time/value for the segment and then update the entry later with a stop time/value. In other embodiments, the cue point information may be stored in segment information data store 124 in some other format and each entry may have additional, fewer or different fields.

[0028] Figure 4 is a flow diagram illustrating a cue point creation method, according to an embodiment. The method 400 may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof. The processing logic is configured to identify segments in a media stream, generate corresponding cue points and facilitate the performance of various segment operations using the generated cue points. In one embodiment, method 400 may be performed by media stream cue point creator 110, as shown in Figures 1a-2b. In another embodiment, method 400 may be performed by segment identifier 112, segment cue point alert module 114 and action module 206, as shown in Figure 2c.

[0029] Referring to Figure 4, at block 410, method 400 identifies a segment of a media stream or file using automated content recognition. In one embodiment, segment identifier module uses fingerprinting technology and the audio and/or video perceptual characteristics of the stream identify particular segments in the stream. In one embodiment, a segment may include a series of frames in the stream that have the same perceptual characteristics as a known segment in segment reference data store 122. Additional details of the segment identification process can be found in U.S. Patent No. 5,918,223 issued June 29, 1999, U.S. Patent No. 6,834,308 issued December 21, 2004, U.S. Patent No. 8,031,947 issued October 4, 2011, and U.S. Patent No. 8,112,818 issued February 7, 2012. In another embodiment, segment identifier 112 uses object recognition and

classification technology to identify objects found within the video images of the media stream. Additional details of segment identification are also described below with respect to Figure 5.

[0030] At block 420, method 400 creates a cue point associated with the segment identified at block 410. In one embodiment, upon identifying a segment that matches a known segment in segment reference data store 122, segment identifier module 112 creates a cue point associated with the identified segment. In one embodiment, segment identifier module 112 creates a corresponding entry in segment information data store 124. In one embodiment, the cue point entry may include an identifier of the media stream from which the segment was identified, an identifier of the segment itself, a start time/value of the segment within the media stream, a stop time/value of the segment, and other information or metadata associated with the identification made of the segment.

[0031] At block 430, method 400 generates an alert to signal replacement of the segment found at the cue point created at block 420, or some other operation. In one embodiment, segment cue point alert module 114 generates an alert signaling replacement of a segment in the media stream identified by the cue point entry in segment information data store 124. Segment cue point alert module 114 identifies a cue point in the media stream. This may occur during playback of the media stream by DVR playback controller 106 or prior to playback when the stream is stored as a file in media data store 104. In one embodiment, segment cue point alert module 114 receives an indication of playback from DVR playback controller 106 and compares the location to the cue points in segment information data store 124. When a cue point is reached, segment cue point alert module 114 may trigger an operation, such as segment replacement. Additional details of segment replacement are described below with respect to Figure 6.

[0032] Figure 5 is a flow diagram illustrating a segment identification method, according to an embodiment. The method 500 may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof. The processing logic is configured to identify content segments in a media stream on which various operations may be performed. In one embodiment, method 500 may be performed by media stream cue point creator 110, as shown in Figures 1a-2b. In another embodiment, method 500 may be performed by segment identifier 112, segment cue point alert module 114 and action module 206, as shown in Figure 2c.

[0033] Referring to Figure 5, at block 510, method 500 generates a fingerprint of the media stream. In one embodiment, segment identifier module 112 identifies segments of a media stream stored in media data store 104 using automated content recognition. Segment identifier

module 112 generates an audio and/or video fingerprint of at least a portion of the media stream. For example, segment identifier module 112 may divide a stored audio stream into frames and extract acoustic or perceptual features of each frame. Segment identifier module 112 may generate an n-dimensional feature vector for each frame for later comparison. For example, segment identifier 112 may extract features from a frame of video, or a group of frames of video. Segment identifier 112 may generate an n-dimensional feature vector for each frame or group of frames.

[0034] At block 520, method 500 compares the fingerprint to a database of known media segments, objects, and classes of objects. In one embodiment, segment identifier module 112 compares the feature vector to stored feature vectors corresponding to known segments from segment reference data store 122. Segment reference data store 122 may store feature vectors, fingerprints, complete copies, or other representations of known segments of media content or objects. For example, segment reference data store 122 may include indication of known advertisements which could be identified and removed from the media streams in media data store 104. For example, segment reference data store 122 may include indication of certain types of objects, and even specific objects. Segment reference data store 122 may be populated with information obtained directly from advertisers, other content creators, or a content distributor, or with information obtained from monitoring previous media streams, or with models and classifiers of objects.

[0035] At block 530, method 500 determines whether the fingerprint matches a known segment or object. In one embodiment, a segment may match if the feature vectors are within an acceptable threshold, thereby indicating that the segment from the stream and the known segment have the same or similar perceptual characteristics. In another embodiment, a segment may match if the computed characteristics match closely to a model or classifier for a particular object or class of objects.

[0036] If the fingerprint matches, at block 540, method 500 creates a cue point entry in segment information data store 124. In one embodiment, segment identifier module 112 creates an entry associated with the identified segment in segment information data store 124. In one embodiment, the cue point entry may include an identifier of the media stream from which the segment was identified, an identifier of the segment itself, a start time/value of the segment within the media stream, a stop time/value of the segment, and other information or metadata associated with the identified segment, as shown in Figure 3b. At block 550, method 500 optionally updates the cue point entry created at block 540. For example, in a live data stream segment identifier module 112 may create an entry with a start time/value for the segment and then update the entry later with a stop time/value, once the stop time/value is known. In another embodiment, there may be separate entries for start time/value and the stop time/value.

[0037] Figure 6 is a flow diagram illustrating a segment replacement method, according to embodiment. The method 600 may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof. The processing logic is configured to replace an identified segment in a media stream with a new or different segment. In one embodiment, method 600 may be performed by media stream cue point creator 110, as shown in Figures 1a-2b. In another embodiment, method 600 may be performed by segment identifier 112, segment cue point alert module 114 and action module 206, as shown in Figure 2c.

[0038] Referring to Figure 6, at block 610, method 600 identifies a cue point in the media stream. In one embodiment, segment cue point alert module 114 identifies a cue point in the media stream. This may occur during playback of the media stream by DVR playback controller 106 or prior to playback when the stream is stored as a file in media data store 104. In one embodiment, segment cue point alert module 114 receives an indication of playback from DVR playback controller 106 and compares the location to the cue points in segment information data store 124. When a cue point is reached, at block 620, segment cue point alert module 114 may trigger an operation, such as segment replacement, display of an overlay, etc. In one embodiment, segment cue point alert module 114 generates an alert signaling replacement of, or other action with respect to, the segment in the media stream identified by the cue point entry in segment information data store 124.

[0039] At block 630, method 600 identifies playback context information for the media stream. The playback context information may include, for example, a time when playback occurs, a location where playback occurs, an indication of who requested playback, or other information. Using the playback context information, segment cue point module 114 may request a new segment from replacement segment data store 126, or determine what other appropriate action to take.

[0040] At block 640, method 600 receives a second segment identified based on the playback information. In one embodiment, segment pre-fetch module 116 may request and receive content segments (e.g., advertisements) from an external source, such as media replacement server 160. Segment pre-fetch module 116 may populate replacement segment data store 126 with the received replacement segments. In one embodiment, segment pre-fetch module 116 may request segments based on the playback context information and media replacement server 160 may determine the appropriate segments to provide. In another embodiment, segment pre-fetch module 116 may request generic segments from media replacement server 160 and segment cue point alert module 114 may make a selection of segments from replacement segment data store 126 based on the

playback context information.

[0041] At block 650, method 600 replaces the original first segment in the media stream with the new second segment. In one embodiment, segment cue point alert module 114 provides an instruction to DVR playback controller 106 to replace the original segment in the playing media stream with the replacement segment from replacement segment data store 126. In one embodiment, replacement of the segment may include modifying the media stream such that the old segment is removed and the new segment is inserted. In another embodiment, replacement may refer to halting playback of the media stream, playing the replacement segment, and then resuming playback of the original media stream. In another embodiment, rather than replacing a segment, segment cue point alert module 114 may provide an instruction to DVR playback controller 106 to block a certain segment, remove a certain segment, insert an overlay of complementary information, insert a clickable banner advertisement, or to perform some other operation. In another embodiment, segment cue point alert module 114 may provide an instruction to another device (e.g. a computer tablet, mobile phone, laptop, game console, lighting controller, set top box, robot actuator) to perform some other operation.

[0042] Figure 7 illustrates a diagrammatic representation of a machine in the exemplary form of a computer system 700 within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. The system 700 may be in the 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. 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 machine in client-server network environment. The machine may be a personal computer (PC), a set-top box (STB), 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. In one embodiment, computer system 700 may represent media player 100 of Figure 1a, server 150 of Figure 1b or media distribution point 200 of Figure 2a.

[0043] The exemplary computer system 700 includes a processing device (processor) 702, a main memory 704 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) such as synchronous DRAM (SDRAM)), a static memory 706 (e.g., flash memory, static random access memory (SRAM)), and a data storage device 718, which communicate with each other via a bus 730.

[0044] Processing device 702 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device 702 may be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing device 702 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. The processing device 702 is configured to execute the notification manager 210 for performing the operations and steps discussed herein.

[0045] The computer system 700 may further include a network interface device 708. The computer system 700 also may include a video display unit 710 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device 712 (e.g., a keyboard), a cursor control device 714 (e.g., a mouse), and a signal generation device 716 (e.g., a speaker).

[0046] The data storage device 718 may include a computer-readable medium 728 on which is stored one or more sets of instructions 722 (e.g., instructions of media stream cue point creator 110 or 210) embodying any one or more of the methodologies or functions described herein. The instructions 722 may also reside, completely or at least partially, within the main memory 704 and/or within processing logic 726 of the processing device 702 during execution thereof by the computer system 700, the main memory 704 and the processing device 702 also constituting computer-readable media. The instructions may further be transmitted or received over a network 720 via the network interface device 708.

[0047] While the computer-readable storage medium 728 is shown in an exemplary embodiment 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 for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present invention. 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.

[0048] The preceding description sets forth numerous specific details such as examples of specific systems, components, methods, and so forth, in order to provide a good understanding of several embodiments of the present invention. It will be apparent to one skilled in the art, however, that at least some embodiments of the present invention may be practiced without these specific

details. In other instances, well-known components or methods are not described in detail or are presented in simple block diagram format in order to avoid unnecessarily obscuring the present invention. Thus, the specific details set forth are merely exemplary. Particular implementations may vary from these exemplary details and still be contemplated to be within the scope of the present invention.

[0049] In the above description, numerous details are set forth. It will be apparent, however, to one of ordinary skill in the art having the benefit of this disclosure, that embodiments of the invention may be practiced without these specific details. In some instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the description.

[0050] 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 steps leading to a desired result. The steps 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.

[0051] 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 "determining", "identifying", "adding", "selecting" or the like, refer to the actions and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (e.g., electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0052] Embodiments of the invention also relate to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-op-

tical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions.

[0053] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0054] It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

Claims

1. A method comprising:

receiving a media stream (300) comprising at least one of an audio portion or a video portion, wherein the media stream (300) comprises a plurality of content segments that lack a cue point;
storing the media stream (300) in a media data store (104) for playback at a later time;
upon identifying, by a processing device (702), a first content segment of the plurality of content segments of the media stream (300), wherein the first content segment corresponds to a known content segment in a segment reference data store (122),
creating a first cue point entry in a segment information data store (124); wherein the first cue point entry is associated with the first content segment, and wherein the first cue point entry in the segment information data store (124) comprises an identifier of the media stream (300), an identifier of the first content segment, a value indicating a start time of the first content segment, and a value indicating an end time of the first content segment;
prior to playback of the media stream (300) at the later time, and upon receiving an indication of playback comparing at least a portion of the plurality of content segments to one or more cue

point entries in the segment information data store (124),
and upon reaching the first cue point entry in the segments information data store (124), triggering operation on the first segment of the media stream identified by the first cue point entry while the media stream is stored in the media data store (104).

2. The method of claim 1, further comprising: generating an alert signaling the operation to be performed on the first content segment in the media stream (300).

3. The method of claim 2, wherein the operation comprises replacing the first content segment in the media stream (300) with a second content segment and, optionally, further comprising:

identifying playback context information for playback of the media stream (300); and
identifying the second content segment in view of the playback context information and, optionally,
wherein the playback context information comprises at least one of a time when playback of the media stream is requested, a location where playback of the media stream is requested, or an indication of a requestor of playback of the media stream.

4. The method of claim 1, wherein identifying the first content segment of the plurality of content segments of the media stream (300) comprises:

generating a fingerprint of at least a portion of the first content segment;
comparing the fingerprint to a plurality of known content segments in the segment reference data store (122); and
determining whether the fingerprint matches at least one of the plurality of known content segments.

5. A system comprising:

a processing device (702);
a memory coupled to the processing device (702), the memory storing a media stream cue point creator (110), executable by the processing device (702) from the memory, to:

receive a media stream (300) comprising at least one of an audio portion or a video portion, wherein the media stream (300) comprises a plurality of content segments that lack a cue point;

store the media stream (300) in a media data store for playback at a later time; upon identifying a first content segment of the plurality of content segments of the media stream (300), wherein the content segment corresponds to a known content segment in a segment reference data store (122);
 create a first cue point entry in a segment information data store (124), wherein the first cue point entry is associated with the first content segment, and wherein the first cue point entry in a segment information data store (124) comprises an identifier of the media stream (300), an identifier of the first content segment, value indicating a start time of the first content segment, and a value indicating an end time of the first content segment;
 prior to playback of the media stream (300) at the later time, and upon receiving an indication of playback, compare at least a portion of the plurality of content segments to one or more cue point entries in the segment information data store (124),

and upon reaching the cue point entry in the segment information data store (124), trigger an operation on the first segment of the media stream (300) identified by the first cue point entry while the media stream (300) is stored in the media data store (104).

- 6. The system of claim 5, wherein the media stream cue point creator (110) is further to:
 generate an alert signaling an operation to be performed on the first segment in the media stream (300).
- 7. The system of claim 6, wherein the operation comprises replacing the first content segment in the media stream (300) with a second content segment and, optionally, wherein the media stream cue point creator (110) is further to:
 identify playback context information for playback of the media stream (300); and
 identify the second content segment in view of the playback context information and, optionally, wherein the playback context information comprises at least one of a time when playback of the media stream (300) is requested, a location where playback of the media stream (300) is requested, or an indication of a requestor of playback of the media stream (300).
- 8. The system of claim 5, wherein to identify the first

content segment of the plurality of content segments of the media stream (300), the media stream cue point creator (110) is further to:

- 5 generate a fingerprint of at least a portion of the first content segment;
 compare the fingerprint to a plurality of known content segments in the segment reference data store (122); and
 10 determine whether the fingerprint matches at least one of the plurality of known content segments.

- 9. A non-transitory computer-readable storage medium storing instructions which, when executed by a processing device (702), cause the processing device (702) to perform operations comprising:

receiving a media stream (300) comprising at least one of an audio portion or a video portion, wherein the media stream (300) comprises a plurality of content segments that lack a cue point;
 storing the media stream (300) in a media data store (104) for playback at a later time;
 upon identifying, by the processing device (702), a first content segment of the plurality of content segments of the media stream (300), wherein the first content segment corresponds to a known content segment in a segment reference data store (122),
 creating a first cue point entry in a segment information data store (124), wherein the first cue point entry is associated with the first content segment, and wherein the first cue point entry in the segment information data store (124) comprises an identifier of the media stream (300), an identifier of the first content segment, a value indicating a start time of the first content segment, and a value indicating an end time of the first content segment;
 prior to playback of the media stream (300) at the later time, and upon receiving an indication of playback, comparing at least a portion of the plurality of content segments to one or more cue point entries in the segment information data store (124),
 and upon reaching the first cue point entry in the segments information data store (124), triggering an operation on the first segment of the media stream (300) identified by the first cue point entry while the media stream (300) is stored in the media data store (104).

- 10. The non-transitory computer-readable storage medium of claim 9, wherein the operations further comprise:
 generating an alert signaling the operation to be per-

formed on the first content segment in the media stream (300).

11. The non-transitory computer-readable storage medium of claim 10, wherein the operation is performed by a first device that is different from a second device that created the cue point as is performed in response to an instruction received from the second device and, optionally, wherein the operations further comprise:

identifying playback context information for playback of the media stream (300); and identifying the second content segment in view of the playback context information and, optionally, wherein the playback context information comprises at least one of a time when playback of the media stream (300) is requested, a location where playback of the media stream (300) is requested, or an indication of a requestor of playback of the media stream (300).

12. The non-transitory computer-readable storage medium of claim 9, wherein identifying the first content segment of the plurality of content segments of the media stream (300) comprises:

generating a fingerprint of at least a portion of the first content segment; comparing the fingerprint to a plurality of known content segments in the segment reference data store; and determining whether the fingerprint matches at least one of the plurality of known content segments.

Patentansprüche

1. Verfahren, umfassend:

Empfangen eines Medienstroms (300), umfassend mindestens eines von einem Audioteil oder einem Videoteil, wobei der Medienstrom (300) eine Vielzahl von Inhaltssegmenten, denen ein Cue-Point fehlt, umfasst;
Speichern des Medienstroms (300) in einem Mediendatenspeicher (104) zur Wiedergabe zu einem späteren Zeitpunkt;
nach dem Identifizieren, durch eine Verarbeitungsvorrichtung (702), eines ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300), wobei das erste Inhaltssegment einem bekannten Inhaltssegment in einem Segmentreferenzdatenspeicher (122) entspricht,
Erstellen eines ersten Cue-Point-Eintrags in ei-

nem Segmentinformationsdatenspeicher (124), wobei der erste Cue-Point-Eintrag mit dem ersten Inhaltssegment assoziiert ist und wobei der erste Cue-Point-Eintrag in dem Segmentinformationsdatenspeicher (124) einen Identifikator des Medienstroms (300), einen Identifikator des ersten Inhaltssegments, einen Wert, der einen Startzeitpunkt des ersten Inhaltssegments angibt, und einen Wert, der einen Endzeitpunkt des ersten Inhaltssegments angibt, umfasst;
vor der Wiedergabe des Medienstroms (300) zu dem späteren Zeitpunkt und nach dem Empfangen einer Wiedergabeangabe, Vergleichen mindestens eines Teils der Vielzahl von Inhaltssegmenten mit einem oder mehreren Cue-Point-Einträgen in dem Segmentinformationsdatenspeicher (124) und nach dem Erreichen des ersten Cue-Point-Eintrags in dem Segmentinformationsdatenspeicher (124), Auslösen einer Operation an dem ersten Segment des Medienstroms, identifiziert von dem ersten Cue-Point-Eintrag, während der Medienstrom in dem Mediendatenspeicher (104) gespeichert ist.

2. Verfahren nach Anspruch 1, ferner umfassend: Erzeugen einer Benachrichtigung, die die an dem ersten Inhaltssegment in dem Medienstrom (300) durchzuführende Operation signalisiert.
3. Verfahren nach Anspruch 2, wobei die Operation das Ersetzen des ersten Inhaltssegments in dem Medienstrom (300) durch ein zweites Inhaltssegment umfasst und optional ferner Folgendes umfasst:

Identifizieren von Wiedergabekontextinformationen zur Wiedergabe des Medienstroms (300); und
Identifizieren des zweiten Inhaltssegments angesichts der Wiedergabekontextinformationen, und wobei optional die Wiedergabekontextinformationen mindestens eines von einem Zeitpunkt, wenn die Wiedergabe des Medienstroms angefordert wird, einem Ort, an dem die Wiedergabe des Medienstroms angefordert wird, oder einer Angabe eines Anforderers der Wiedergabe des Medienstroms umfassen.

4. Verfahren nach Anspruch 1, wobei das Identifizieren des ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300) Folgendes umfasst:

Erzeugen eines Fingerabdrucks mindestens eines Teils des ersten Inhaltssegments;
Vergleichen des Fingerabdrucks mit einer Viel-

zahl von bekannten Inhaltssegmenten in dem Segmentreferenzdatenspeicher (122); und Bestimmen, ob der Fingerabdruck zu mindestens einem der Vielzahl von bekannten Inhaltssegmenten passt.

5. System, umfassend:

eine Verarbeitungsvorrichtung (702); einen Speicher, der mit der Verarbeitungsvorrichtung (702) gekoppelt ist, wobei der Speicher einen Medienstrom-Cue-Point-Ersteller (110) speichert, der von der Verarbeitungsvorrichtung (702) aus dem Speicher für Folgendes ausgeführt werden kann:

Empfangen eines Medienstroms (300), umfassend mindestens eines von einem Audioteil oder einem Videoteil, wobei der Medienstrom (300) eine Vielzahl von Inhaltssegmenten, denen ein Cue-Point fehlt, umfasst;

Speichern des Medienstroms (300) in einem Mediendatenspeicher zur Wiedergabe zu einem späteren Zeitpunkt;

nach dem Identifizieren eines ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300), wobei das Inhaltssegment einem bekannten Inhaltssegment in einem Segmentreferenzdatenspeicher (122) entspricht;

Erstellen eines ersten Cue-Point-Eintrags in einem Segmentinformationsdatenspeicher (124), wobei der erste Cue-Point-Eintrag mit dem ersten Inhaltssegment assoziiert ist und wobei der erste Cue-Point-Eintrag in einem Segmentinformationsdatenspeicher (124) einen Identifikator des Medienstroms (300), einen Identifikator des ersten Inhaltssegments, einen Wert, der einen Startzeitpunkt des ersten Inhaltssegments angibt, und einen Wert, der einen Endzeitpunkt des ersten Inhaltssegments angibt, umfasst;

vor der Wiedergabe des Medienstroms (300) zu dem späteren Zeitpunkt und nach dem Empfangen einer Wiedergabeangabe, Vergleichen mindestens eines Teils der Vielzahl von Inhaltssegmenten mit einem oder mehreren Cue-Point-Einträgen in dem Segmentinformationsdatenspeicher (124) und nach dem Erreichen des ersten Cue-Point-Eintrags in dem Segmentinformationsdatenspeicher (124), Auslösen einer Operation an dem ersten Segment des Medienstroms (300), identifiziert von dem ersten Cue-Point-Eintrag, während der Medienstrom (300) in dem Mediendatenspei-

cher (104) gespeichert ist.

6. System nach Anspruch 5, wobei der Medienstrom-Cue-Point-Ersteller (110) ferner Folgendes tut:
Erzeugen einer Benachrichtigung, die eine an dem ersten Segment in dem Medienstrom (300) durchzuführende Operation signalisiert.

7. System nach Anspruch 6, wobei die Operation das Ersetzen des ersten Inhaltssegments in dem Medienstrom (300) durch ein zweites Inhaltssegment umfasst und wobei optional der Medienstrom-Cue-Point-Ersteller (110) ferner Folgendes tut:

Identifizieren von Wiedergabekontextinformationen zur Wiedergabe des Medienstroms (300); und

Identifizieren des zweiten Inhaltssegments angesichts der Wiedergabekontextinformationen, und wobei optional

die Wiedergabekontextinformationen mindestens eines von einem Zeitpunkt, wenn die Wiedergabe des Medienstroms (300) angefordert wird, einem Ort, an dem die Wiedergabe des Medienstroms (300) angefordert wird, oder einer Angabe eines Anforderers der Wiedergabe des Medienstroms (300) umfassen.

8. System nach Anspruch 5, wobei zum Identifizieren des ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300) der Medienstrom-Cue-Point-Ersteller (110) ferner Folgendes tut:

Erzeugen eines Fingerabdrucks mindestens eines Teils des ersten Inhaltssegments;

Vergleichen des Fingerabdrucks mit einer Vielzahl von bekannten Inhaltssegmenten in dem Segmentreferenzdatenspeicher (122); und Bestimmen, ob der Fingerabdruck zu mindestens einem der Vielzahl von bekannten Inhaltssegmenten passt.

9. Nicht flüchtiges computerlesbares Speichermedium, das Anweisungen speichert, die, wenn sie von einer Verarbeitungsvorrichtung (702) ausgeführt werden, bewirken, dass die Verarbeitungsvorrichtung (702) Operationen durchführt, umfassend:

Empfangen eines Medienstroms (300), umfassend mindestens eines von einem Audioteil oder einem Videoteil, wobei der Medienstrom (300) eine Vielzahl von Inhaltssegmenten, denen ein Cue-Point fehlt, umfasst;

Speichern des Medienstroms (300) in einem Mediendatenspeicher (104) zur Wiedergabe zu einem späteren Zeitpunkt;

nach dem Identifizieren, durch die Verarbei-

- tungsvorrichtung (702), eines ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300), wobei das erste Inhaltssegment einem bekannten Inhaltssegment in einem Segmentreferenzdatenspeicher (122) entspricht,
- Erstellen eines ersten Cue-Point-Eintrags in einem Segmentinformationsdatenspeicher (124), wobei der erste Cue-Point-Eintrag mit dem ersten Inhaltssegment assoziiert ist und wobei der erste Cue-Point-Eintrag in dem Segmentinformationsdatenspeicher (124) einen Identifikator des Medienstroms (300), einen Identifikator des ersten Inhaltssegments, einen Wert, der einen Startzeitpunkt des ersten Inhaltssegments angibt, und einen Wert, der einen Endzeitpunkt des ersten Inhaltssegments angibt, umfasst;
- vor der Wiedergabe des Medienstroms (300) zu dem späteren Zeitpunkt und nach dem Empfangen einer Wiedergabeangabe, Vergleichen mindestens eines Teils der Vielzahl von Inhaltssegmenten mit einem oder mehreren Cue-Point-Einträgen in dem Segmentinformationsdatenspeicher (124)
- und nach dem Erreichen des ersten Cue-Point-Eintrags in dem Segmentinformationsdatenspeicher (124), Auslösen einer Operation an dem ersten Segment des Medienstroms (300), identifiziert von dem ersten Cue-Point-Eintrag, während der Medienstrom (300) in dem Medien- datenspeicher (104) gespeichert ist.
- 10.** Nicht flüchtiges computerlesbares Speichermedium nach Anspruch 9, wobei die Operationen ferner Folgendes umfassen:
- Erzeugen einer Benachrichtigung, die die an dem ersten Inhaltssegment in dem Medienstrom (300) durchzuführende Operation signalisiert.
- 11.** Nicht flüchtiges computerlesbares Speichermedium nach Anspruch 10, wobei die Operation von einer ersten Vorrichtung durchgeführt wird, die sich von einer zweiten Vorrichtung unterscheidet, die den Cue-Point erstellt hat, wie es als Reaktion auf eine von der zweiten Vorrichtung empfangene Anweisung durchgeführt wird, und wobei optional die Operationen ferner Folgendes umfassen:
- Identifizieren von Wiedergabekontextinformationen zur Wiedergabe des Medienstroms (300); und
- Identifizieren des zweiten Inhaltssegments angesichts der Wiedergabekontextinformationen, und wobei optional
- die Wiedergabekontextinformationen mindestens eines von einem Zeitpunkt, wenn die Wiedergabe des Medienstroms (300) angefordert wird, einem Ort, an dem die Wiedergabe des
- Medienstroms (300) angefordert wird, oder einer Angabe eines Anforderers der Wiedergabe des Medienstroms (300) umfassen.
- 12.** Nicht flüchtiges computerlesbares Speichermedium nach Anspruch 9, wobei das Identifizieren des ersten Inhaltssegments der Vielzahl von Inhaltssegmenten des Medienstroms (300) Folgendes umfasst:
- Erzeugen eines Fingerabdrucks mindestens eines Teils des ersten Inhaltssegments;
- Vergleichen des Fingerabdrucks mit einer Vielzahl von bekannten Inhaltssegmenten in dem Segmentreferenzdatenspeicher; und
- Bestimmen, ob der Fingerabdruck zu mindestens einem der Vielzahl von bekannten Inhaltssegmenten passt.
- Revendications**
- 1.** Un procédé comprenant :
- la réception d'un flux multimédia (300) contenant au moins un élément parmi une partie audio ou une partie vidéo, où le flux multimédia (300) comprend une pluralité de segments de contenu auxquels il manque un point de repère,
- la conservation en mémoire du flux multimédia (300) dans un espace mémoire de données multimédia (104) à des fins de lecture à un instant ultérieur,
- après l'identification, par un dispositif de traitement (702), d'un premier segment de contenu de la pluralité de segments de contenu du flux multimédia (300), où le premier segment de contenu correspond à un segment de contenu contenu dans un espace mémoire de données de référence de segments (122),
- la création d'une première entrée de point de repère dans un espace mémoire de données d'informations de segments (124), où la première entrée de point de repère est associée au premier segment de contenu, et où la première entrée de point de repère dans l'espace mémoire de données d'informations de segments (124) comprend un identifiant du flux multimédia (300), un identifiant du premier segment de contenu, une valeur indiquant une heure de début du premier segment de contenu et une valeur indiquant une heure de fin du premier segment de contenu,
- avant la lecture du flux multimédia (300) à l'instant ultérieur, et après la réception d'une indication de lecture, la comparaison d'au moins une partie de la pluralité de segments de contenu à une ou plusieurs entrées de point de repère dans l'espace mémoire de données d'informa-

- tions de segments (124),
 et, après l'atteinte de la première entrée de point
 de repère dans l'espace mémoire de données
 d'informations de segments (124), le déclenche-
 ment d'une opération sur le premier segment du
 flux multimédia identifié par la première entrée
 de point de repère pendant que le flux multimé-
 dia est conservé en mémoire dans l'espace mé-
 moire de données multimédia (104). 5
2. Le procédé selon la Revendication 1, comprenant
 en outre :
 la génération d'une alerte signalant l'opération à exé-
 cuter sur le premier segment de contenu dans le flux
 multimédia (300). 10
3. Le procédé selon la Revendication 2, où l'opération
 comprend le remplacement du premier segment de
 contenu dans le flux multimédia (300) par un deuxiè-
 me segment de contenu et, éventuellement,
 comprenant en outre : 15
- l'identification d'informations de contexte de lec-
 ture destinées à la lecture du flux multimédia
 (300), et 20
- l'identification du deuxième segment de contenu
 en considération des informations de contexte
 de lecture, et éventuellement,
 où les informations de contexte de lecture com-
 prennent au moins un élément parmi une heure
 à laquelle la lecture du flux multimédia est de-
 mandée, un emplacement où la lecture du flux
 multimédia est demandée ou une indication d'un
 demandeur de la lecture du flux multimédia. 25
4. Le procédé selon la Revendication 1, où l'identifica-
 tion du premier segment de contenu de la pluralité
 de segments de contenu du flux multimédia (300)
 comprend : 30
- la génération d'une empreinte d'au moins une
 partie du premier segment de contenu,
 la comparaison de l'empreinte à une pluralité de
 segments de contenu connus dans l'espace mé-
 moire de données de référence de segments
 (122), et 35
- la détermination si l'empreinte correspond à au
 moins un segment de contenu de la pluralité de
 segments de contenu connus. 40
5. Un système comprenant : 45
- un dispositif de traitement (702),
 une mémoire couplée au dispositif de traitement
 (702), la mémoire conservant en mémoire un
 créateur de point de repère de flux multimédia
 (110), exécutable par le dispositif de traitement
 (702) à partir de la mémoire, de façon à : 50
- recevoir un flux multimédia (300) contenant
 au moins un élément parmi une partie audio
 ou une partie vidéo, où le flux multimédia
 (300) comprend une pluralité de segments
 de contenu auxquels il manque un point de
 repère,
 conserver en mémoire le flux multimédia
 (300) dans un espace mémoire de données
 multimédia à des fins de lecture à un instant
 ultérieur,
 après l'identification d'un premier segment
 de contenu de la pluralité de segments de
 contenu du flux multimédia (300), où le seg-
 ment de contenu correspond à un segment
 de contenu connu dans un espace mémoire
 de données de référence de segments
 (122),
 créer une première entrée de point de re-
 père dans un espace mémoire de données
 d'informations de segments (124), où la
 première entrée de point de repère est as-
 sociée au premier segment de contenu, et
 où la première entrée de point de repère
 dans l'espace mémoire de données d'infor-
 mations de segments (124) comprend un
 identifiant du flux multimédia (300), un iden-
 tifiant du premier segment de contenu, une
 valeur indiquant une heure de début du pre-
 mier segment de contenu et une valeur in-
 diquant une heure de fin du premier seg-
 ment de contenu,
 avant la lecture du flux multimédia (300) à
 l'instant ultérieur, et après la réception
 d'une indication de lecture, comparer au
 moins une partie de la pluralité de segments
 de contenu à une ou plusieurs entrées de
 point de repère dans l'espace mémoire de
 données d'informations de segments (124),
 et, après l'atteinte de la première entrée de
 point de repère dans l'espace mémoire de
 données d'informations de segments (124),
 déclencher une opération sur le premier
 segment du flux multimédia (300) identifié
 par la première entrée de point de repère
 pendant que le flux multimédia (300) est
 conservé en mémoire dans l'espace mé-
 moire de données multimédia (104). 55
6. Le système selon la Revendication 5, où le créateur
 de point de repère de flux multimédia (110) est conçu
 en outre de façon à :
 générer une alerte signalant une opération à exé-
 cuter sur le premier segment dans le flux multimédia
 (300). 60
7. Le système selon la Revendication 6, où l'opération
 comprend le remplacement du premier segment de
 contenu dans le flux multimédia (300) par un deuxiè- 65

me segment de contenu et, éventuellement, où le créateur de point de repère de flux multimédia (110) est conçu en outre de façon à :

5 identifier des informations de contexte de lecture destinées à la lecture du flux multimédia (300), et
 10 identifier le deuxième segment de contenu en considération des informations de contexte de lecture, et éventuellement,
 15 où les informations de contexte de lecture comprennent au moins un élément parmi une heure à laquelle la lecture du flux multimédia (300) est demandée, un emplacement où la lecture du flux multimédia (300) est demandée ou une indication d'un demandeur de la lecture du flux multimédia (300).

8. Le système selon la Revendication 5, où, de façon à identifier le premier segment de contenu de la pluralité de segments de contenu du flux multimédia (300), le créateur de point de repère de flux multimédia (110) est conçu en outre de façon à :

25 générer une empreinte d'au moins une partie du premier segment de contenu,
 30 comparer l'empreinte à une pluralité de segments de contenu connus dans l'espace mémoire de données de référence de segments (122), et
 35 déterminer si l'empreinte correspond à au moins un segment de contenu de la pluralité de segments de contenu connus.

9. Un support à mémoire lisible par ordinateur non transitoire conservant en mémoire des instructions qui, lorsqu'elles sont exécutées par un dispositif de traitement (702), amènent le dispositif de traitement (702) à exécuter des opérations comprenant :

40 la réception d'un flux multimédia (300) contenant au moins un élément parmi une partie audio ou une partie vidéo, où le flux multimédia (300) comprend une pluralité de segments de contenu auxquels il manque un point de repère,
 45 la conservation en mémoire du flux multimédia (300) dans un espace mémoire de données multimédia (104) à des fins de lecture à un instant ultérieur,
 50 après l'identification, par le dispositif de traitement (702), d'un premier segment de contenu de la pluralité de segments de contenu du flux multimédia (300), où le premier segment de contenu correspond à un segment de contenu connu dans un espace mémoire de données de référence de segments (122),
 55 la création d'une première entrée de point de repère dans un espace mémoire de données

d'informations de segments (124), où la première entrée de point de repère est associée au premier segment de contenu, et où la première entrée de point de repère dans l'espace mémoire de données d'informations de segments (124) comprend un identifiant du flux multimédia (300), un identifiant du premier segment de contenu, une valeur indiquant une heure de début du premier segment de contenu et une valeur indiquant une heure de fin du premier segment de contenu,
 avant la lecture du flux multimédia (300) à l'instant ultérieur, et après la réception d'une indication de lecture, la comparaison d'au moins une partie de la pluralité de segments de contenu à une ou plusieurs entrées de point de repère dans l'espace mémoire de données d'informations de segments (124),
 et, après l'atteinte de la première entrée de point de repère dans l'espace mémoire de données d'informations de segments (124), le déclenchement d'une opération sur le premier segment du flux multimédia (300) identifié par la première entrée de point de repère pendant que le flux multimédia (300) est conservé en mémoire dans l'espace mémoire de données multimédia (104).

10. Le support à mémoire lisible par ordinateur non transitoire selon la Revendication 9, où les opérations comprennent en outre :
 la génération d'une alerte signalant l'opération à exécuter sur le premier segment de contenu dans le flux multimédia (300).

11. Le support à mémoire lisible par ordinateur non transitoire selon la Revendication 10, où l'opération est exécutée par un premier dispositif qui est différent d'un deuxième dispositif qui a créé le point de repère comme exécutée en réponse à une instruction reçue du deuxième dispositif et, éventuellement,
 où les opérations comprennent en outre :

l'identification d'informations de contexte de lecture destinées à la lecture du flux multimédia (300), et
 l'identification du deuxième segment de contenu en considération des informations de contexte de lecture et, éventuellement
 où les informations de contexte de lecture comprennent au moins un élément parmi une heure à laquelle la lecture du flux multimédia (300) est demandée, un emplacement où la lecture du flux multimédia (300) est demandée ou une indication d'un demandeur de la lecture du flux multimédia (300).

12. Le support à mémoire lisible par ordinateur non transitoire selon la Revendication 9, où l'identification du

premier segment de contenu de la pluralité de segments de contenu du flux multimédia (300) comprend :

la génération d'une empreinte d'au moins une partie du premier segment de contenu, 5
la comparaison de l'empreinte à une pluralité de segments de contenu connus dans l'espace mémoire de données de référence de segments, et 10
la détermination si l'empreinte correspond à au moins un segment de contenu de la pluralité de segments de contenu connus.

15

20

25

30

35

40

45

50

55




Fig. 1a

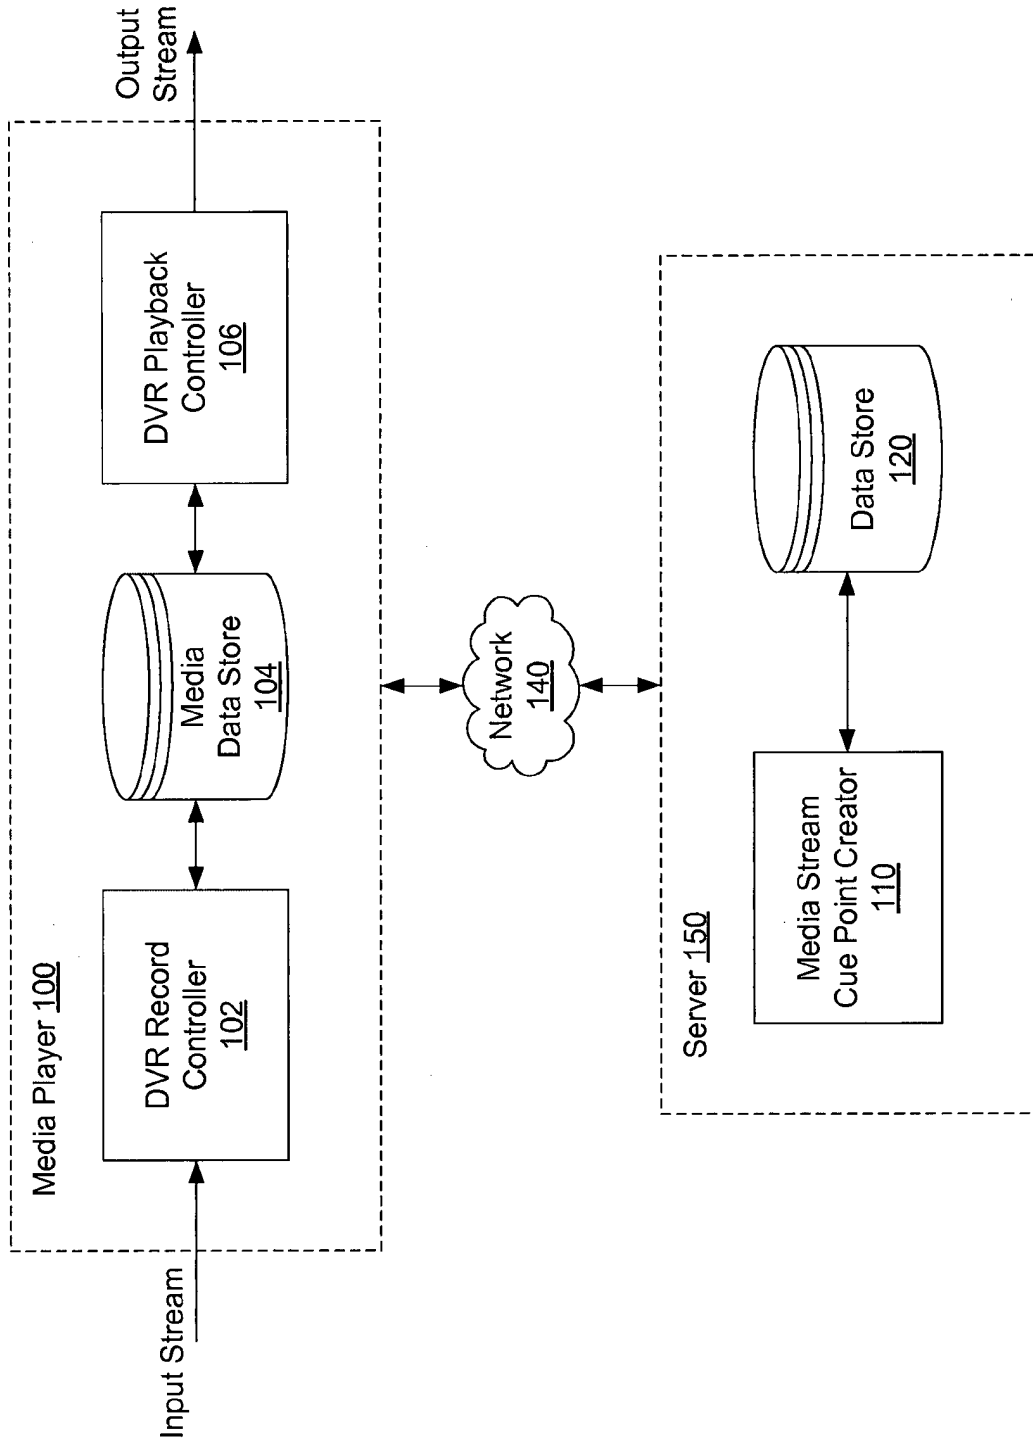


Fig. 1b

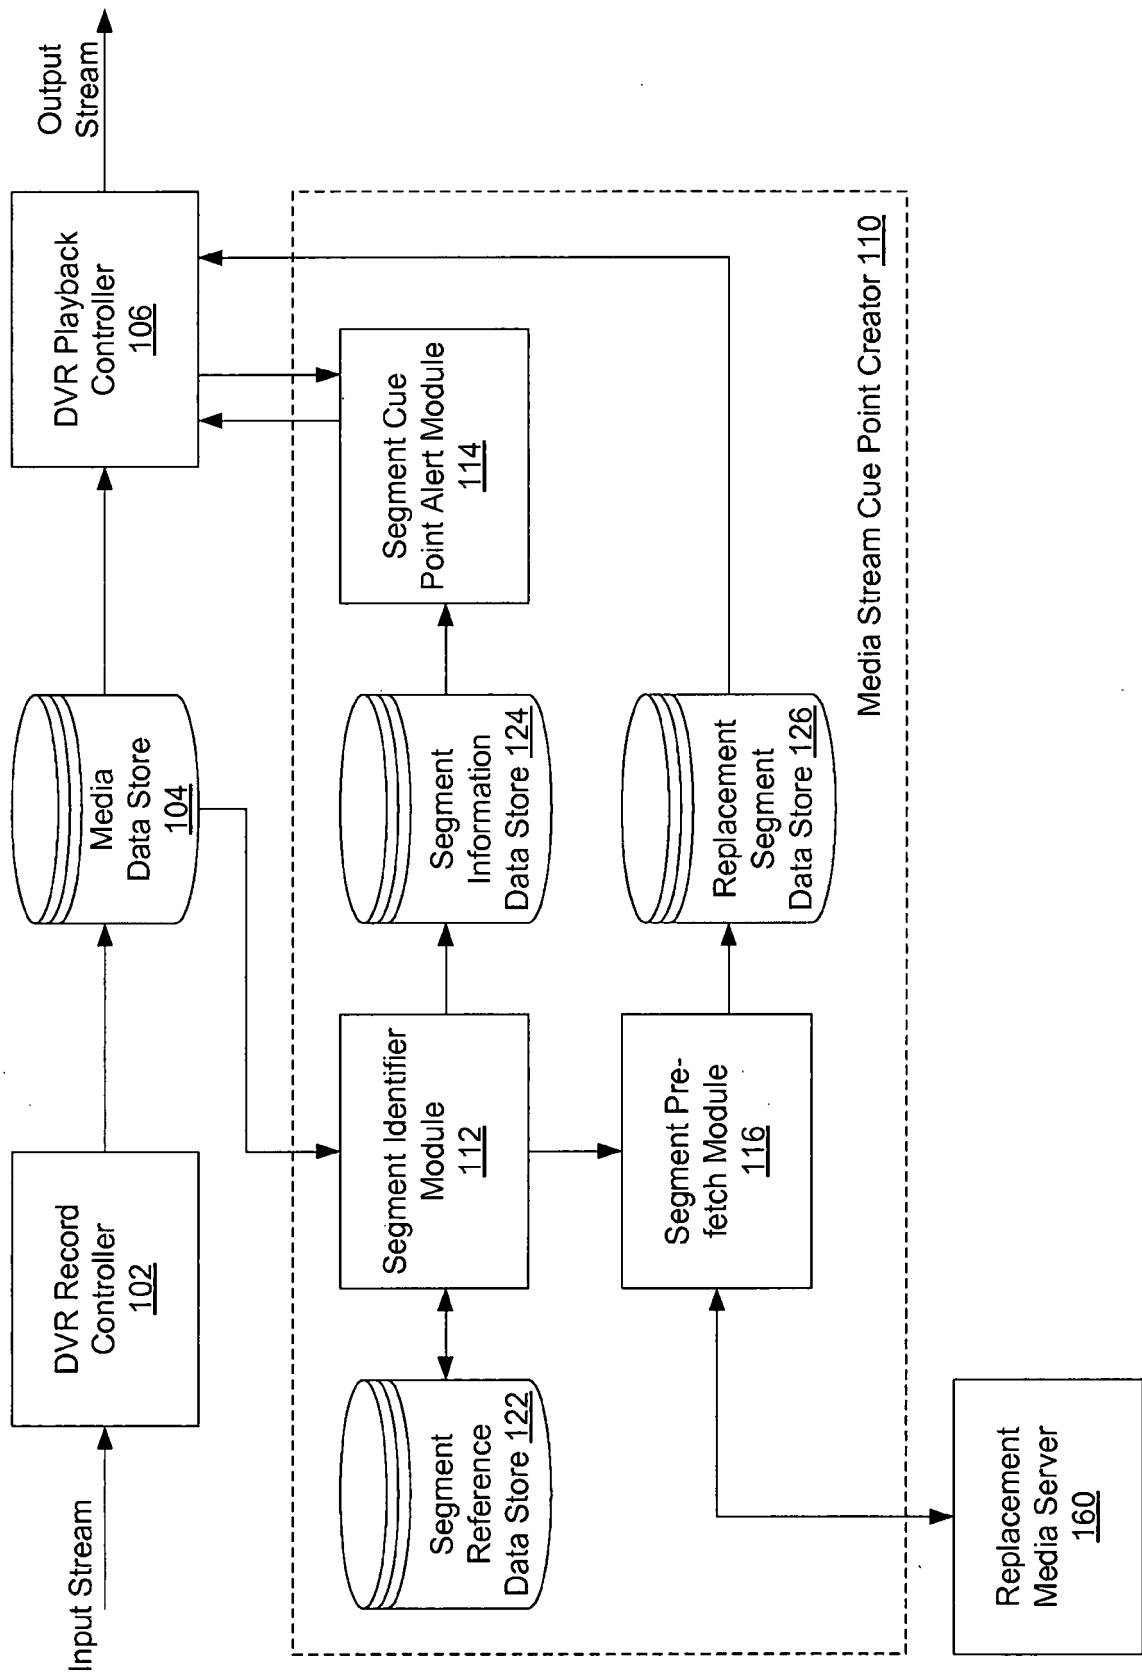


Fig. 1c




Fig. 2a

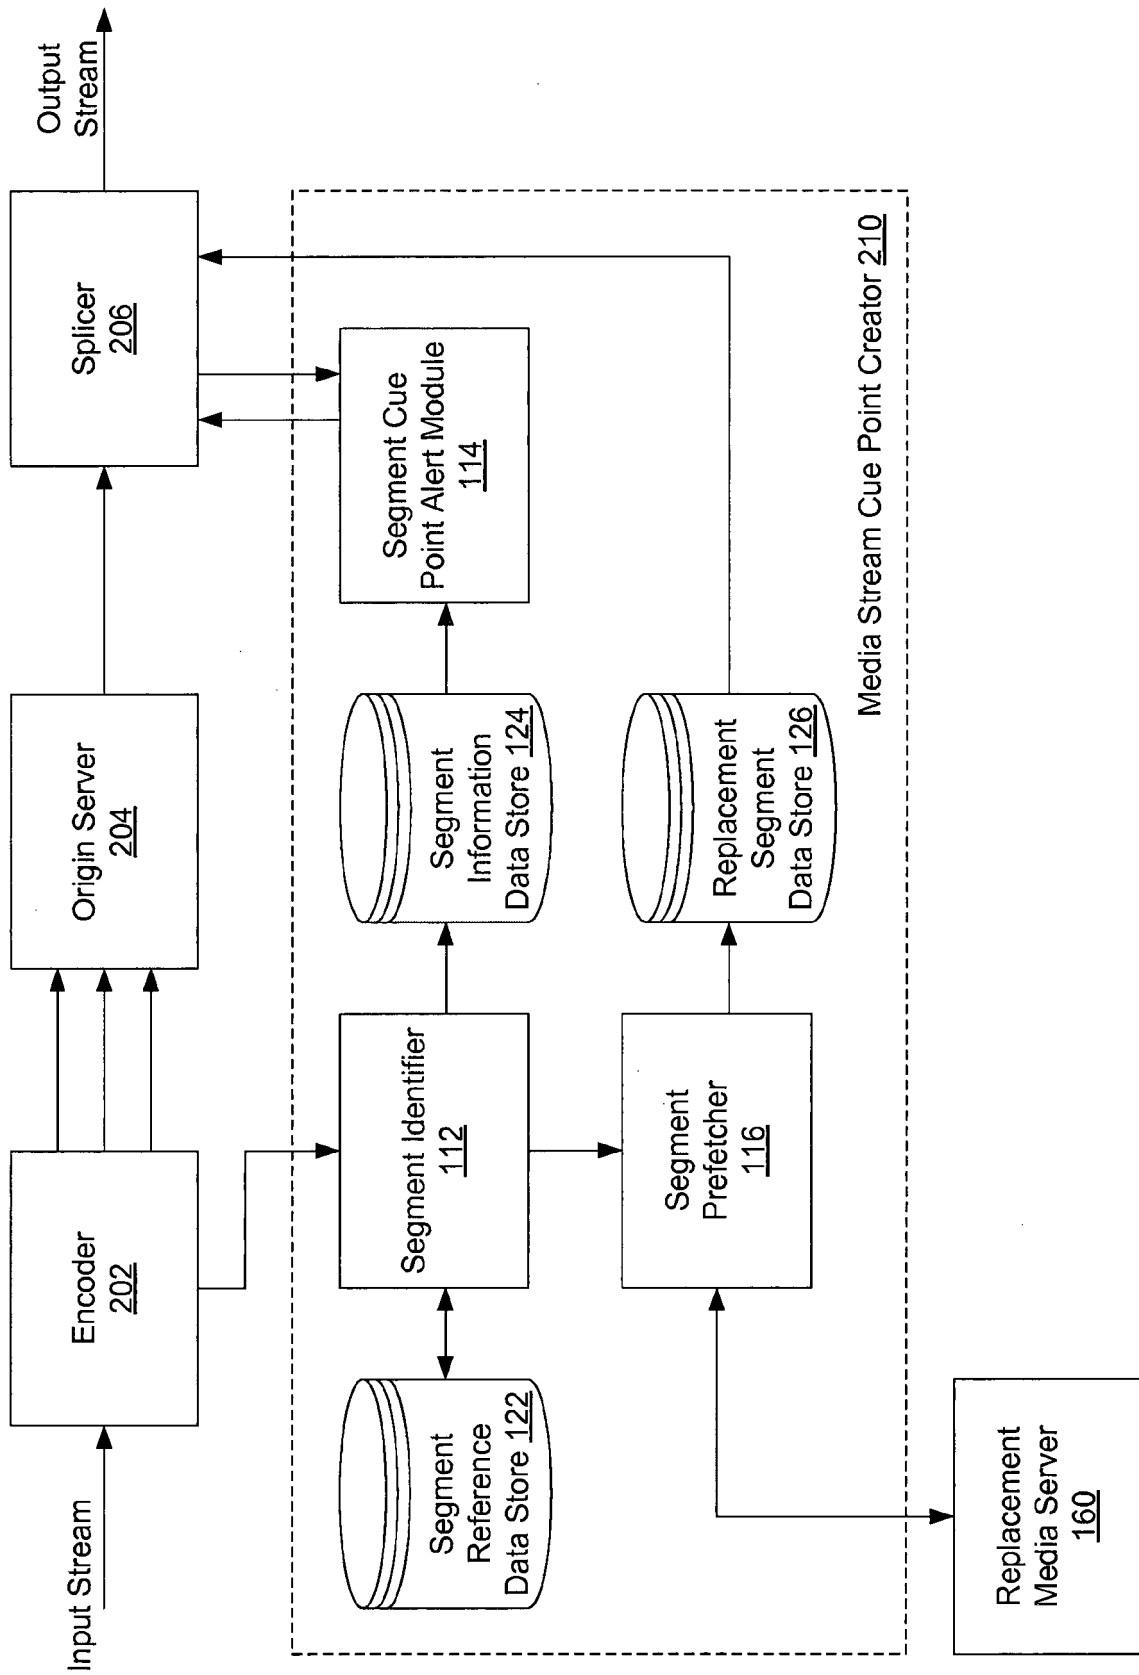


Fig. 2b




Fig. 2c




Fig. 3a

Segment Information Data Store <u>124</u>				
Stream ID	Segment ID	Start Cue Pt.	End Cue Pt.	Metadata
300	01	304	306	...
300	02	308	310	...
300	03	312	314	...
300	04	314	316	...

Fig. 3b

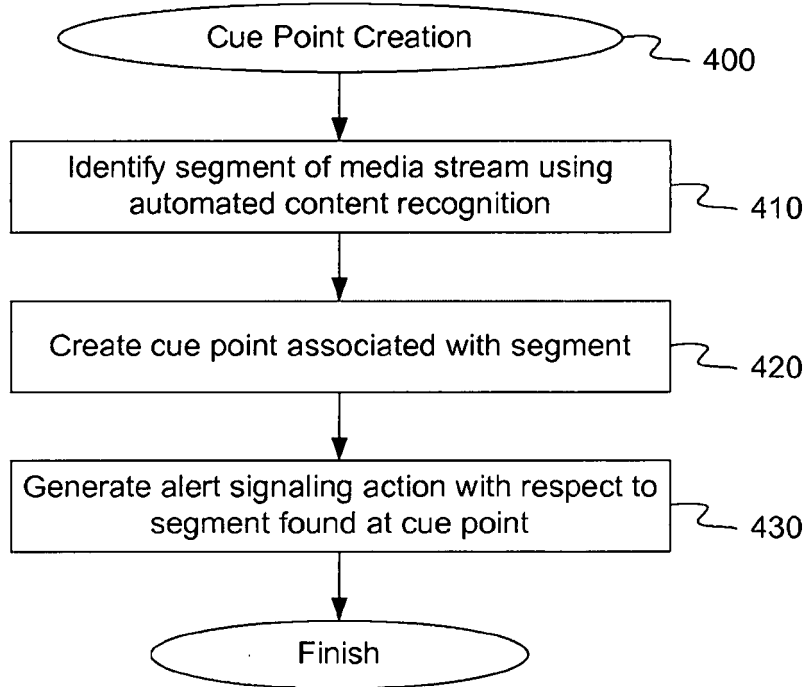


Fig. 4

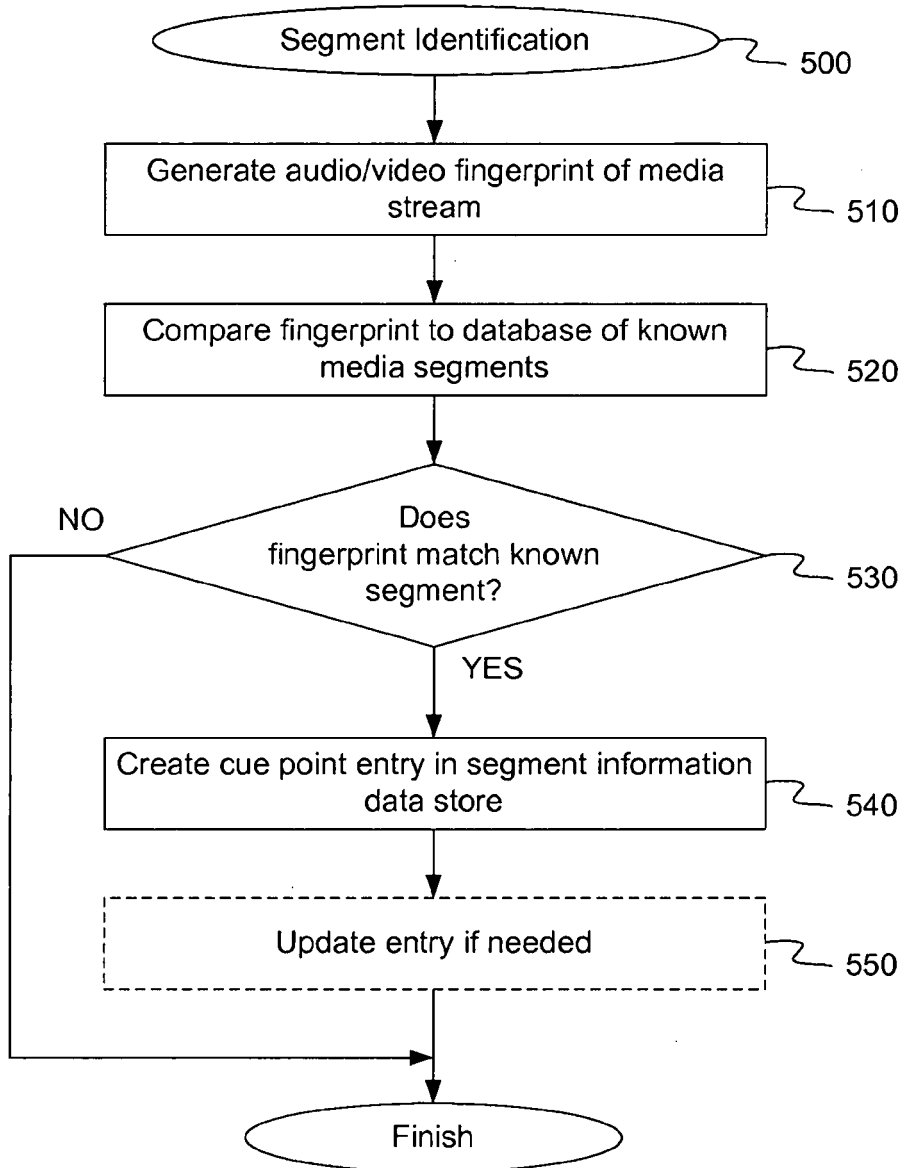


Fig. 5

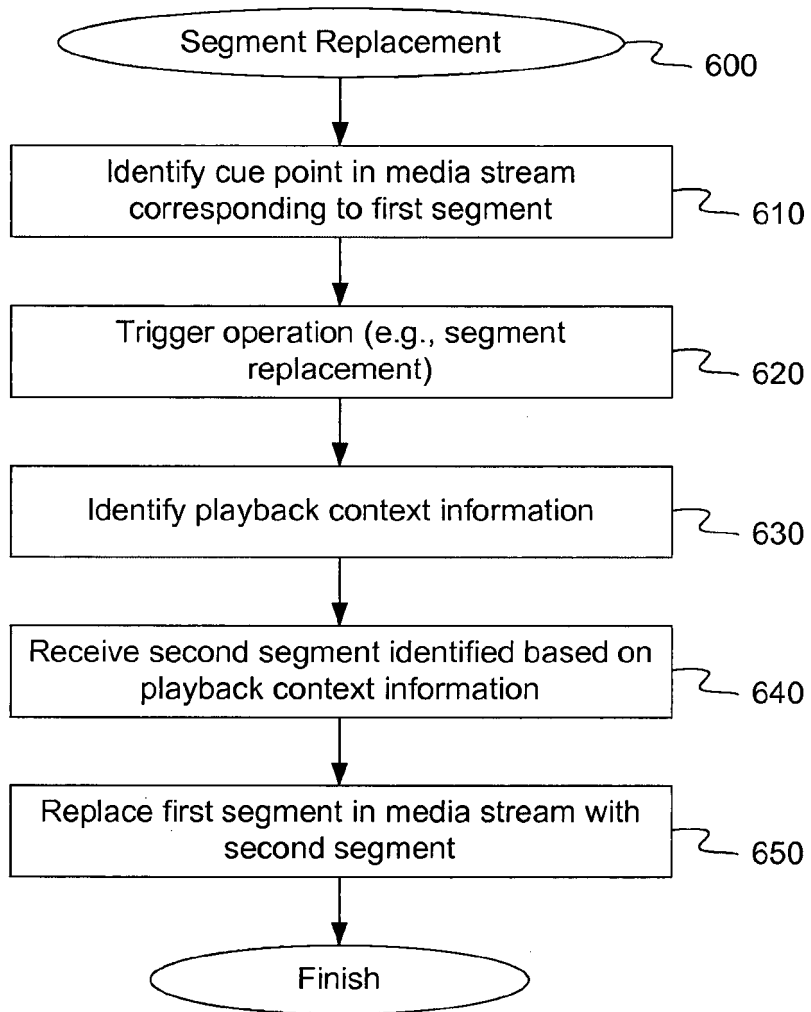


Fig. 6

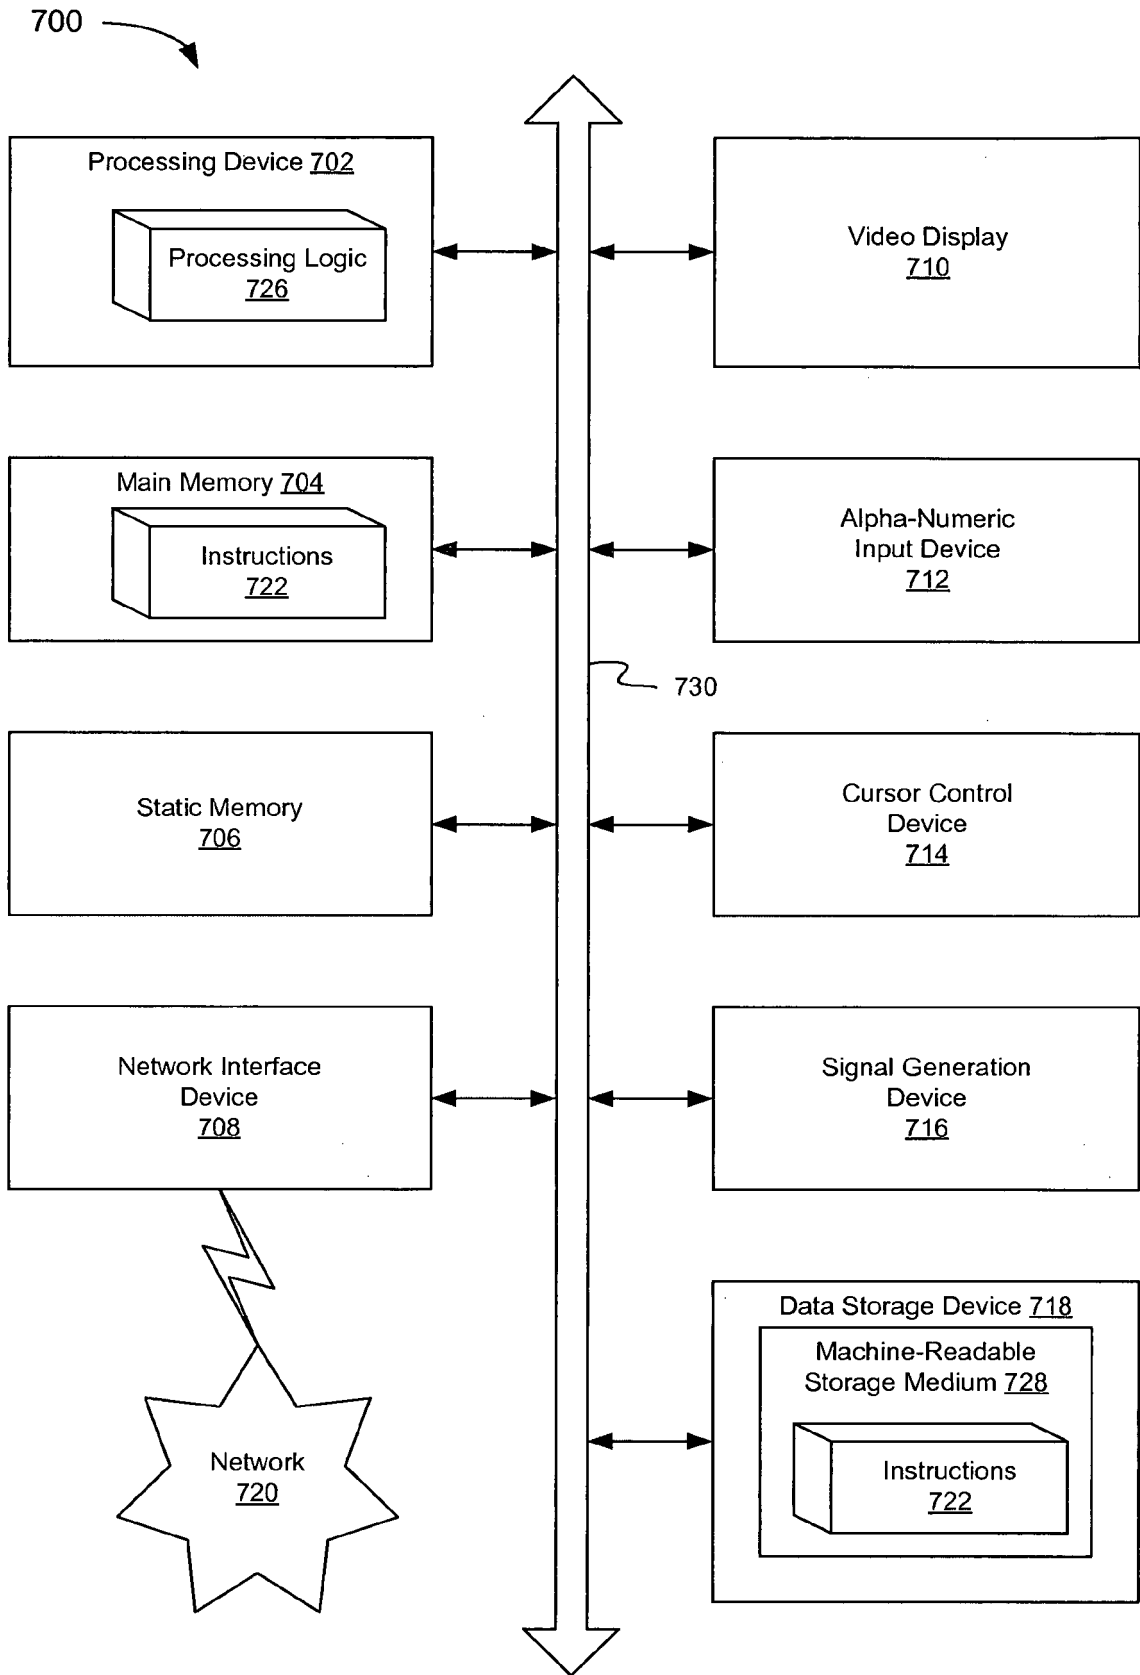


Fig. 7

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2004189873 A1 [0003]
- US 2006195860 A1 [0004]
- US 5918223 A [0009] [0029]
- US 7877438 B [0017]
- US 6834308 B [0029]
- US 8031947 B [0029]
- US 8112818 B [0029]