



(51) International Patent Classification:

<i>H04N 21/2187</i> (2011.01)	<i>H04N 21/84</i> (2011.01)
<i>H04N 21/24</i> (2011.01)	<i>H04N 21/845</i> (2011.01)
<i>H04N 21/44</i> (2011.01)	<i>H04N 21/8547</i> (2011.01)
<i>H04N 21/4402</i> (2011.01)	<i>H04N 21/414</i> (2011.01)
<i>H04N 21/442</i> (2011.01)	<i>H04N 21/2343</i> (2011.01)
<i>H04N 21/6332</i> (2011.01)	<i>H04N 21/234</i> (2011.01)

(72) **Inventors:** MCLEOD, Euan; 410 Terry Avenue N., Seattle, Washington 98109-5210 (US). JOLIVEAU, Marc; 410 Terry Avenue N., Seattle, Washington 98109-5210 (US). RICHTER, Stefan Christian; 410 Terry Avenue N., Seattle, Washington 98109-5210 (US).

(21) International Application Number:

PCT/US2017/034066

(22) International Filing Date:

23 May 2017 (23.05.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/170,164	01 June 2016 (01.06.2016)	US
15/170,169	01 June 2016 (01.06.2016)	US

(74) **Agent:** VILLENEUVE, Joseph M. et al.; P.O. Box 70250, Oakland, California 94612-0250 (US).

(71) **Applicant:** AMAZON TECHNOLOGIES, INC. [US/US]; P.O. Box 81226, Seattle, Washington 98108-1226 (US).

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

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,

(54) Title: CATCHING UP TO THE LIVE PLAYHEAD IN LIVE STREAMING




FIG. 1

(57) **Abstract:** Techniques are described for reducing the delay between the live playhead of live streaming content and the client playhead of a client device consuming the live stream. In one technique, an increased playback speed is used by the media player on the client device so that the delay is gradually reduced. In another technique, the media player jumps forward in the stream, skipping content identified as expendable.



UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

CATCHING UP TO THE LIVE PLAYHEAD IN LIVE STREAMING

RELATED APPLICATION DATA

[0001] The present application claims priority to U.S. Patent Application No. 15/170,164
5 entitled *Catching Up to the Live Playhead in Live Streaming* filed on June 1, 2016 (Attorney
Docket No. AMZ1P062/P32499-US), and to U.S. Patent Application No. 15/170,169
entitled *Catching Up to the Live Playhead in Live Streaming* filed on June 1, 2016 (Attorney
Docket No. AMZ1P069/PM13188-US), the entire disclosures of both of which are
incorporated herein by reference for all purposes.

BACKGROUND

[0002] Live streaming content includes channels or feeds with scheduled content (e.g.,
premium movie channels) and live broadcasts (e.g., sporting events, news, etc.). Unlike
video-on-demand (VOD) content, live streaming content often does not have a distinct end
point and may continue indefinitely. In addition, VOD content may be buffered in client
15 devices well in advance of the client playhead (i.e., the content fragment currently being
rendered by the client). This is typically not the case for live content because of the
constraint that the delay between the live playhead (i.e., the latest content fragment
available) and the client playhead be as low as possible.

[0003] Because content services that provide live content often prioritize minimizing the
20 delay between the live and client playheads, this has the potential to result in buffering
events in which client devices run out of content fragments to play back. This is particularly
the case for clients that maintain short buffers. Each buffering event for a given client
device introduces additional delay between the live and client playheads (beyond the
inherent latency associated with making the content available via a streaming protocol). For
25 live content having a relatively long duration, e.g., a live sporting event, some client devices
may fall several minutes behind the live playhead. Not only might this negatively affect
viewer experience, it might also result in lack of compliance with the provisions of service
level agreements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] FIG. 1 is an illustration of a delay between the live and client playheads of a live
30 content stream.

[0005] FIG. 2 is a simplified diagram of a computing environment in which implementations enabled by the present disclosure may be practiced.

[0006] FIG. 3 is a simplified diagram of an example of a client device that may be used with implementations enabled by the present disclosure.

5 [0007] FIG. 4 is a flowchart illustrating operation of a particular implementation.

[0008] FIG. 5 is a flowchart illustrating operation of a particular implementation.

DETAILED DESCRIPTION

[0009] This disclosure describes techniques for reducing the delay between the live playhead of live streaming content and the client playhead of a client device consuming the live stream. In some cases, an increased playback speed is used by the media player on the client device so that the delay is gradually reduced. The increase in the playback speed is preferably small enough so that the faster playback is not perceptible to a human viewer. In other cases, the media player jumps forward in the stream, skipping content that is considered expendable, e.g., black frames, slate frames (e.g., an image with “please stand
10 by” or “we’ll be right back), low-value advertising content, etc. Using these mechanisms separately or in combination, the client playhead may be brought closer in time to the live playhead. An example will be instructive.
15

[0010] Suppose two friends who are avid soccer fans are both watching live streams of the World Cup final match on their tablets while messaging each other. One of the friends on device 102 is using cellular data in a crowded public location to connect with the stream while the other on device 104 is using wifi on his home network. Because device 102 is competing for bandwidth with everyone else at its location it experiences multiple buffering events, and the client playhead falls behind that of his friend’s device 104 which is consuming the stream with considerably more available bandwidth and is therefore much
20 closer in time to the live playhead. When the viewer associated with device 104 triumphantly texts his friend as the ball hits the back of the net, the viewer associated with device 102 is understandably confused and frustrated. Using techniques enabled by the present disclosure, the delay resulting in this frustration can be significantly reduced or
25 substantially eliminated.

[0011] Live streaming content is sometimes annotated with metadata in real time by human operators as the content is being generated. As part of this annotation, segments of the content may be identified by annotators as being expendable. For example, when the cameras at the event focus on the crowd or an aerial view of the surrounding geography rather than the pitch for a few seconds, or when play is stopped for an injury, such segments of the content can be identified by a human annotator as expendable. When the media player on client device 102 reaches such a point in the live stream, it can skip ahead to the next fragments in the stream that are not identified as expendable. In this way, the delay between the live playhead and the client playhead of device 102 can be shortened, reducing the likelihood of viewer frustration.

[0012] Alternatively or in conjunction with the skipping of content, the playback speed on device 102 can be increased for a period of time in a way that is not noticeable to the viewer to allow for a more gradual reduction in the delay between the live and client playheads. And if both client devices 102 and 104 are attempting to minimize this delay in similar fashion, it is much more likely that the respective viewers will be having viewer experiences that are more closely synchronized in time.

[0013] FIG. 2 illustrates an example of a computing environment in which a video content service 202 provides live streaming content (e.g., audio or video) via network 204 to a variety of client devices (206-1 through 206-5) in accordance with the techniques described herein. Content service 202 may conform to any of a wide variety of architectures such as, for example, a services platform deployed at one or more co-locations, each implemented with one or more servers 203. Network 204 represents any subset or combination of a wide variety of network environments including, for example, TCP/IP-based networks, telecommunications networks, wireless networks, satellite networks, cable networks, public networks, private networks, wide area networks, local area networks, the Internet, the World Wide Web, intranets, extranets, etc. Client devices 206 may be any suitable device capable of connecting to network 204 and consuming live streaming content provided by service 202. Such devices may include, for example, mobile devices (e.g., cell phones, smart phones, and tablets), personal computers (e.g., laptops and desktops), set top boxes (e.g., for cable and satellite systems), smart televisions, gaming consoles, wearable computing devices (e.g., smart watches or smart glasses), etc.

[0014] At least some of the examples described herein contemplate implementations based on computing models that enable ubiquitous, convenient, on-demand network access to a shared pool of computing resources (e.g., networks, servers, storage, applications, and services). As will be understood, such computing resources may be integrated with and/or
5 under the control of the same entity controlling content service 202. Alternatively, such resources may be independent of content service 202, e.g., on a platform under control of a separate provider of computing resources with which content service 202 connects to consume computing resources as needed.

[0015] It should also be noted that, despite any references to particular computing
10 paradigms and software tools herein, the computer program instructions on which various implementations are based may correspond to any of a wide variety of programming languages, software tools and data formats, may be stored in any type of non-transitory computer-readable storage media or memory device(s), and may be executed according to a variety of computing models including, for example, a client/server model, a peer-to-peer
15 model, on a stand-alone computing device, or according to a distributed computing model in which various functionalities may be effected or employed at different locations.

[0016] In the following examples and for the sake of simplicity, content service 202 is described as if it were integrated with the platform(s) that provides the live streaming content to client devices. However, it will be understood that content service 202 may
20 provide access to live streaming content in conjunction with one or more content delivery networks (e.g., CDN 214) that may or may not be independent of content service 202. In addition, the source of the live content may or may not be independent of content service 202 (e.g., as represented by content provider server 216). The range of variations known to those of skill in the art are contemplated to be within the scope of this disclosure.

[0017] Some of the implementations enabled by the present disclosure contemplate logic
25 resident on the client devices consuming live streaming content from content service 202; such logic being configured to make decisions in conjunction with consuming the video content such as, for example, monitoring the delay between playheads, increasing playback speed, and/or skipping expendable content. The logic might be part of an existing algorithm
30 or module on the client device or implemented to work in conjunction with such an algorithm or module. The logic might be implemented, for example, in a media player on the client device or as a separate application or module resident on the client device.

[0018] It should also be noted that implementations are contemplated in which, in addition to content delivery logic 210 (which facilitates various aspects of content delivery to client devices 206), content service 202 may include logic that facilitates at least some aspects of monitoring and reducing the delay between playheads as described herein (e.g., as represented by playhead catchup logic 211). For example, such logic might be used to associate metadata with fragments or segments of the content that identify expendable content that can potentially be skipped to allow for reduction of the delay between the live playhead and client playheads. As discussed below, this may be done manually (e.g., by human operators), using existing content metadata, or by real-time analysis of the frames or fragments of the content. Such logic might also be configured to determine whether and how far a particular client is behind the live playhead and/or take steps or send instructions to the client (e.g., to initiate higher-speed playback or skipping of content) to support getting the client's playhead closer to the live playhead.

[0019] In addition to providing access to the live streaming content, content service 202 may also include a variety of information related to the live streaming content (e.g., other associated metadata and manifests in data store 212 to which service 202 provides access. Alternatively, such information about the live streaming content, as well as the live streaming content itself may be provided and/or hosted by one or more separate platforms, e.g., CDN 214. It should be noted that, while logic 210 and 211, and data store 212 are shown as integrated with content service 202, implementations are contemplated in which some or all of these operate remotely from the associated content service, and/or are under the control of an independent entity. From these examples, those of skill in the art will understand the diversity of use cases to which the techniques described herein are applicable.

[0020] A block diagram of an example of a client device 300 suitable for use with various implementations is shown in FIG. 3. Device 300 includes one or more single or multi-core processors 302 configured to execute stored instructions (e.g., in device memory 320). Device 300 may also include one or more input/output (I/O) interface(s) 304 to allow the device to communicate with other devices. I/O interfaces 304 may include, for example, an inter-integrated circuit (I2C) interface, a serial peripheral interface (SPI) bus, a universal serial bus (USB), an RS-232 interface, a media device interface, and so forth. I/O interface(s) 304 is coupled to one or more I/O devices 306.

[0021] Device 300 may also include one or more communication interfaces 308 configured to provide communications between the device and other devices. Such communication interface(s) 308 may be used to connect to cellular networks, personal area networks (PANs), local area networks (LANs), wide area networks (WANs), and so forth.

5 For example, communications interfaces 308 may include radio frequency modules for a 3G or 4G cellular network, a WiFi LAN and a Bluetooth PAN. Device 300 also includes one or more buses or other internal communications hardware or software (not shown) that allow for the transfer of data and instructions between the various modules and components of the device.

10 **[0022]** Device 300 also includes one or more memories (e.g., memory 310). Memory 310 includes non-transitory computer-readable storage media that may be any of a wide variety of types of volatile and non-volatile storage media including, for example, electronic storage media, magnetic storage media, optical storage media, quantum storage media, mechanical storage media, and so forth. Memory 310 provides storage for computer
15 readable instructions, data structures, program modules and other data for the operation of device 300. As used herein, the term “module” when used in connection with software or firmware functionality may refer to code or computer program instructions that are integrated to varying degrees with the code or computer program instructions of other such “modules.” The distinct nature of the different modules described and depicted herein is
20 used for explanatory purposes and should not be used to limit the scope of this disclosure.

[0023] Memory 310 includes at least one operating system (OS) module 312 configured to manage hardware resources such as I/O interfaces 304 and provide various services to applications or modules executing on processor(s) 302. Memory 310 also includes a user interface module 316, a content rendering module 318, and other modules. Memory 310
25 also includes device memory 320 to store a wide variety of instructions and information using any of a variety of formats including, for example, flat files, databases, linked lists, trees, or other data structures. Such information includes content for rendering and display on display 306(1) including, for example, any type of video content. In some implementations, a portion of device memory 320 may be distributed across one or more
30 other devices including servers, network attached storage devices, and so forth.

[0024] The logic or computer program instructions used to support reducing the delay between live and client playheads as described herein (represented by playback speed

module 319 and content skipping module 321) may be implemented in a variety of ways. For example, at least some of this functionality may be implemented as part of the code of a media player operating on device 300. Alternatively, modules 319 and 321 may be implemented separately from and interact with the device's media player, web browser, mobile app, decoder, etc.

[0025] And as mentioned above, implementations are contemplated in which at least a portion of the logic or computer program instructions may reside on a separate platform, e.g., service 202, CDN 214, etc.; potentially working in conjunction with the client-side logic to reduce the delay between the respective playheads. Suitable variations and alternatives will be apparent to those of skill in the art. It will also be understood that device 300 of FIG. 3 is merely an example of a device with which various implementations enabled by the present disclosure may be practiced, and that a wide variety of other devices types may also be used (e.g., devices 206-1 to 206-5). The scope of this disclosure should therefore not be limited by reference to device-specific details.

[0026] The delivery of live streaming content to a client device according to a particular implementation is illustrated in the flow chart of FIG. 4. This and other examples described herein assume the use of H.265 encoding (also commonly referred to as HEVC) for video content. However, it will be understood that the basic principles described herein may be employed with any of a variety of video and audio codecs including, for example, MPEG-1, MPEG-2, MPEG-4 Part 2, VC-1, H.263, VP8, VP9, Daala, and H.264. The example illustrated in FIG. 4 also assumes a media player on the client device that includes logic (e.g., modules 319 and 321) configured to manage at least some aspects of reducing the delay between the live and client playheads as described herein. Again, these details are merely presented by way of example.

[0027] When a user wants to connect with a content service using a client device, the connection is typically achieved through some kind of login process to the service in a user interface presented on the client device. Content playback is provided, for example, via a resident media player, web browser, or mobile app. Access to content over the Internet is typically governed by a DRM system such as Google's Widevine, Microsoft's PlayReady, Apple's FairPlay, or Sony's OpenMG to name a few representative examples. Live streaming content is typically delivered in an encrypted stream using any of a variety of encryption technologies including, for example, various Advanced Encryption Standard

(AES) and Elliptic Curve Cryptography (ECC) encryption techniques. The live stream may also be delivered using an adaptive bit rate streaming technique such as, for example, MPEG-DASH (Dynamic Adaptive Streaming over HTTP), Apple's HLS (HTTP Live Streaming), or Microsoft's Smooth Streaming, to name a few representative examples. It should be noted that the techniques described herein are compatible with a wide range of content services, media players, DRM systems, encryption technologies, and streaming technologies, the details of which are known to those of skill in the art. The nature and operation of these technologies will therefore not be described in detail to promote clarity.

[0028] Referring now to FIG. 4, when live content is selected in a user interface on a client device (402), a request for the content is sent to the corresponding content service (404). The content service provides the client device with the information the client device needs to acquire a stream of the content (406). This may include, for example, DRM licenses, a decryption key, content metadata, and information about where the client can request the fragments of the selected content at various resolutions (e.g., a manifest). The client device then acquires a stream of the live content using the information received from the content service (408).

[0029] As the client device is consuming the live content, the delay between the live playhead and the client playhead is tracked (410). This may be done in a variety of ways. For example, logic on the client device can count the cumulative amount of time required to recover from rebuffering events. In a simpler approach, a fixed amount of time could be added to the delay for each rebuffering event. In another example, time stamps associated with the recently requested fragments and representative of or close in time to the live playhead could be compared to a local clock on the client device.

[0030] In another example, the time reference used by logic on the client device to determine the delay could be a time stamp associated with one or more fragments acquired at the beginning of the session. For example, logic on the client could compare the difference between such a time stamp and that of a later fragment with actual time elapsed on the client device (e.g., as determined by a local clock) to determine the extent to which the delay has grown over time. In another example, server-side logic could determine the delay for a particular client by comparing the time stamp for a recently requested fragment by that client with the time stamp for the fragment most recently made available by the content service, or the latest fragment requested by any client consuming the live stream. Where server-side

logic determines the delay for a particular client device, the server-side logic could also determine whether the delay exceeds a threshold and, when that occurs, transmit a message or an instruction to the client to initiate use of one or both of the catch-up mechanisms. Alternatively, the server-side logic could periodically transmit the delay to the client device
5 for decision making on the client.

[0031] It should be noted that the delay value being tracked may only be an approximation of the actual delay between the live and client playheads. That is, both client and server-side logic might use time references that are suitable proxies for one or both of the playheads without departing from the scope of this disclosure. For example, the time
10 stamp associated with a fragment most recently requested by a client device will likely differ from the time the fragment is actually rendered and displayed by the client device, but may otherwise be suitable for purposes of determining a reliable approximation of the actual delay. In another example, the time stamp associated with the fragment most recently made
15 available by the content service might be earlier than the time at which the fragment becomes available to some client devices. More generally, and as will be appreciated by the diversity of the foregoing examples, there are many ways in which a delay between the live playhead and the client playhead for a particular client may be tracked, determined, or approximated for use as described herein. The scope of the present disclosure should therefore not be limited by reference to such examples.

[0032] Referring again to FIG. 4, if the delay between the live playhead and the client
20 playhead exceeds a predetermined threshold (412), the operation of logic configured to reduce the delay is initiated (414). The threshold may be selected to keep the client device acceptably close to the live playhead while ensuring a particular level of content quality. The threshold might also be selected, at least in part, to ensure compliance with any
25 applicable service level agreement(s). And as mentioned above, the logic initiated may be configured to increase the playback speed of the client device's media player, skip playback of expendable content, or use a combination of these "catch-up mechanisms." Once the delay is reduced to an acceptable level (416) (which may or may not be the same as the threshold used to trigger use of the catch up mechanism(s)), use of the catch-up
30 mechanism(s) is/are terminated (418), and tracking of the delay continues (410).

[0033] An example of the operation of an implementation that uses both of these mechanisms is illustrated in the flowchart of FIG. 5. According to the implementation

depicted, when the delay between playheads exceeds the threshold (e.g., as determined in 412), the playback speed of the media player on the client device is increased (502). The amount by which the playback speed is increased may be relatively small, e.g., 2-5%, and potentially as much about 15%. For example, for a normal playback frame rate of 30 frames
5 per second, the playback speed could be increased to 32 or 33 frames per second. Ideally, the increase in playback speed is imperceptible to most human viewers and may be empirically determined, e.g., using viewer assessment by human subjects.

[0034] It should be noted that the increase in playback speed may be at least partially dependent on the type of content. For example, humans may more readily distinguish an
10 increase in playback speed for musical content than for content that is primarily visual in nature. So, for content that includes musical content (e.g., pure audio content, or video with significant musical content), the increase in playback speed may be lower than for some video content. And although it is preferable that the increase in playback speed be imperceptible to some, most, or all human viewers, implementations are contemplated in
15 which this does not need to be the case.

[0035] According to some implementations, the increased playback speed may be a constant speed. Alternatively, implementations are contemplated in which the playback speed may vary dynamically. For example, if the increased playback speed is not successful
20 in reducing the delay between the live and the client playheads after some programmable period of time, the playback speed might be further increased. In another alternative, the playback speed for different types of content may be different. For example, segments of the content that include musical content could be identified or detected (e.g., from content metadata) and the playback speed could be reduced for those segments while playback of segments not including musical content could be at a higher rate. In another alternative,
25 different playback speeds might be used for different ranges of delay so as to enable faster catch-up for larger delays.

[0036] Increased playback speed or content skipping might also affect or interact with the operation of other logic on the client device such as, for example, adaptive bit rate selection logic. For example, such logic might be configured to request fragments at
30 reduced quality when the media player is operating at a higher playback speed so that the high-speed playback can continue even if available bandwidth is low. Alternatively, increased playback speed and/or content skipping might be disabled where the content

quality attainable by the adaptive bit rate selection logic is or would be negatively affected by the operation of the catch-up mechanisms (e.g., the video quality drops below a threshold). In some cases, the available bandwidth may be checked before initiating use of a catch-up mechanism to ensure playback quality. For example, if available bandwidth is
5 below a certain level, increased playback speed may be disallowed or, if already started, suspended.

[0037] Referring back to FIG. 5, and in parallel with the increased playback speed, the identification of expendable content in the live stream is initiated (504). Each time one or more expendable content frames, fragments, or segments are encountered in the stream
10 (506), playback of the expendable content is skipped (508). Such expendable content may correspond to any of a variety of breaks between the most relevant or interesting segments of the content such as, for example, black frames, slate frames, credits, opening or closing montages, commercial breaks, etc.

[0038] As mentioned above, such segments of content may be identified with reference
15 to metadata that is introduced into the live stream (e.g., as metadata tags) by human operators in substantially real time. That is, human operators may view and annotate the live content (e.g., as it is received from the live content source) for a wide variety of purposes such as, for example, dynamic insertion of advertisements, providing additional descriptive content (e.g., sports play-by-play), rating of content for different viewing audiences, etc.
20 According to some implementations, human operators annotate the live content by identifying expendable content, i.e., content for which playback may be skipped on client devices that are sufficiently behind the live playhead. So, for example, as a human operator is viewing the stream of a live sporting event close in time to the actual creation of the content, she can identify segments that correspond, for example, to views of the crowd,
25 commercial breaks, replay reviews, etc., by inserting metadata tags into the stream that are associated with individual fragments, groups of fragments, or even longer segments. These tags may then be used (e.g., by content skipping module 321) to skip playback of the corresponding fragments.

[0039] In addition or as an alternative to being tagged by human operators, expendable
30 content may be identified in a variety of other ways. For example and as mentioned above, expendable content might be identified using information about the content that is provided by the content provider. For example, content providers often provide information (e.g.,

content stream metadata) about events or breaks in content (e.g., commercial breaks, breaks between blocks of content, the beginning or end of scheduled content, the beginning of important live content, etc.) that may present opportunities for content skipping. Such events or breaks might include a fade to black, a few black frames, or content that is less important to viewers (e.g., commercial breaks, credits, etc.). And using such information for catching up to the live playhead may be advantageous in that there is a relatively high degree of reliability in the timing of such events as they are explicitly identified by the content provider. Further, for some types of live streams (e.g., streams of scheduled content), such events or breaks may be relatively far out into the future and thus may be communicated to the client well in advance.

[0040] In some implementations, the identification of expendable content may be based on real-time or near-real-time video inspection and analysis. For example, video fragments, GOPs, and individual video frames can be analyzed to determine whether they are black frames, or correspond to scenes in which the display images do not appreciably change for an extended period of time. As should be appreciated, such an approach may be particularly important for live streams that do not follow a strict schedule, e.g., live sporting events in which commercial breaks or the end of the program is determined by play on the field.

[0041] Identification of expendable content may be done by the client (e.g., content skipping module 321) with reference to either or both of information from the content provider (e.g., in stream metadata), or by inspection of the fragments or frames of the current stream as they are received. For example, the client might be configured to identify low-complexity or static content (e.g., by virtue of the relationships or dependencies among frames in a GOP). This might be done instead of or in addition to identification of expendable content on the server side (e.g., by playhead catchup logic 211).

[0042] Referring again to FIG. 5, the delay between the live and client playheads is monitored to determine whether operation of one or both of the catch-up mechanisms should be terminated. That is, for example, if the delay drops below a threshold (510), normal speed playback without content skipping may be resumed (512). The threshold used may be chosen to get the client playhead as close as possible to the live playhead without negatively affecting the user experience in terms of content quality and/or an unacceptably high rate of rebuffering events. The threshold might be the same as the one used to initiate operation of the catch-up mechanisms (e.g., 412 of FIG. 4). Alternatively, some level of hysteresis might

be built into the system, using a lower threshold for termination of higher-speed playback or content skipping to ensure that the client isn't rapidly switching between normal playback and use of the catch-up mechanisms. Either or both thresholds might be dynamic in nature, depending, for example, on available bandwidth or a current state of an adaptive bit rate
5 algorithm.

[0043] According to some implementations in which both catch-up mechanisms are employed, different thresholds may apply to initiation and termination of each mechanism. That is, increasing playback speed may allow for a finer control of the reduction of the delay between the live and client playheads as compared to the cruder but faster control
10 represented by the skipping of content. Implementations are therefore contemplated in which the threshold(s) associated with increased playback speed is lower than the threshold(s) associated with content skipping. For example, the increase in playback speed might be initiated when the delay reaches 30 seconds, but content skipping might not be initiated until the delay reaches a minute or more. This allows for a more brute force
15 approach (represented by content skipping) for longer delays, while allowing for a more fine-grained and precise approach (represented by higher-speed playback) for shorter delays.

[0044] And while the different catch-up mechanisms may be used in combination, they may or may not be used simultaneously. For example, the different mechanisms might be used alternatively, e.g., with content skipping being used until the delay has decreased
20 sufficiently to allow for further reduction using higher-speed playback. Variations on this theme within the scope of the present disclosure will be understood by those of skill in the art.

[0045] Embodiments disclosed herein may include a computer program product, having one or more non-transitory computer-readable media having computer program instructions
25 stored therein, the computer program instructions being configured such that, when executed by one or more computing devices, the computer program instructions cause the one or more computing devices to at least one of, acquire a stream of live video content for playback on a client device, determine that a delay between a live playhead of the live video content and a client playhead associated with the playback of the live video content on the client device at
30 a first frame rate exceeds a threshold, increase a playback speed of the live video content on the client device to a second frame rate in a manner that is substantially imperceptible to a human viewer, thereby reducing the delay, identify one or more expendable portions of the

live video content, skip playback of at least one of the one or more expendable portions of the live video content, thereby reducing the delay, determine that the delay is below the threshold, and/or decrease the playback speed of the live video content on the client device to the first frame rate.

5 [0046] Optionally, the one or more processors may be configured to identify the one or more expendable portions of the live video content (1) using first content metadata associated with the live video content by a content provider in conjunction with generation of the live video content, (2) using second content metadata associated with the live video content by a human operator after generation of the live video content, or (3) by analyzing
10 the live video content substantially in real time. Optionally, the one or more processors may further be configured to determine the delay between the live playhead and the client playhead by determining a cumulative time for recovering from rebuffering events occurring on the client device.

[0047] Embodiments disclosed herein may include a client device including a memory,
15 an output device, and one or more processors configured, in conjunction with the memory, to one or more of acquire a stream of content for playback on the output device, determine that a delay between a live playhead of the content and a client playhead associated with the playback of the content exceeds a first threshold, and/or increase a playback speed of the content.

20 [0048] Optionally, the one or more processors may be further configured to determine that the delay is below a second threshold, the second threshold being lower than the first threshold, and/or decrease the playback speed of the content. Optionally, the one or more processors may be further configured to at least one of determine that the delay exceeds a second threshold, the second threshold being higher than the first threshold, and/or further
25 increase the playback speed of the content.

[0049] Optionally, the one or more processors may be further configured to identify one or more expendable portions of the content, and/or skip playback of at least one of the one or more expendable portions of the content. Optionally, the one or more processors may be further configured to determine that the delay exceeds a second threshold, the second
30 threshold being different from the first threshold, and the one or more processors may be configured to skip playback of at least one of the one or more expendable portions of the content in response to determining that the delay exceeds the second threshold.

[0050] Optionally, the one or more processors may be further configured to determine that there is sufficient available bandwidth for increasing the playback speed of the content. Optionally, the one or more processors may be further configured to determine that available bandwidth or a playback quality of the content has dropped below a corresponding
5 threshold, and/or decrease the playback speed of the content. Optionally, the one or more processors may be further configured to determine the delay based on one or more rebuffering events or using time stamps associated with fragments or frames of the content.

[0051] Embodiments disclosed herein may include a computer-implemented method including at least one of, acquiring a stream of content for playback, determining that a delay
10 between a live playhead of the content and a client playhead associated with the playback of the content exceeds a delay threshold, and/or increasing a playback speed of the content. Optionally, the method may further include determining that the delay is below a second threshold, the second threshold being lower than the first threshold, and/or decreasing the playback speed of the content. Optionally, the method may further include determining that
15 the delay exceeds a second threshold, the second threshold being higher than the first threshold, and/or further increasing the playback speed of the content.

[0052] Optionally, the method may further include identifying one or more expendable portions of the content, and/or skipping playback of at least one of the one or more expendable portions of the content. Optionally, the method may further include determining
20 that the delay exceeds a second threshold, the second threshold being different from the first threshold, and wherein skipping playback of at least one of the one or more expendable portions of the content may occur in response to determining that the delay exceeds the second threshold.

[0053] Optionally, the method may further include determining that there is sufficient
25 available bandwidth for increasing the playback speed of the content. Optionally, the method may further include determining that available bandwidth or a playback quality of the content has dropped below a corresponding threshold, and/or decreasing the playback speed of the content. Optionally, the method may further include determining the delay based on one or more rebuffering events or using time stamps associated with fragments or
30 frames of the content.

[0054] Embodiments disclosed herein may include a computer-implemented method including at least one of receiving content from a live content source, encoding the content

for streaming to client devices, streaming the content to the client devices, determining a delay between a live playhead of the content and a client playhead associated with the playback of the content on a first client device, and/or instructing the first client device to increase a playback speed of the content.

- 5 [0055] Optionally, determining the delay may include comparing a first time stamp associated with a first fragment of the content requested by the first client device with a second time stamp associated with a second fragment of the content available for streaming to the client devices, or with a third time stamp associated with a third fragment requested by one or more other client devices.
- 10 [0056] Embodiments disclosed herein may include a computer program product, including one or more non-transitory computer-readable media having computer program instructions stored therein, the computer program instructions being configured such that, when executed by one or more computing devices, the computer program instructions cause the one or more computing devices to at least one of, acquire a stream of live video content
15 for playback on a client device, the stream of the live video content including a plurality of video fragments identify one or more rebuffering events on the client device, determine a delay between a live playhead of the live video content and a client playhead associated with the playback of the live video content on the client device based on the one or more rebuffering events, determine that the delay exceeds a threshold, identify one or more
20 expendable portions of the live video content using content metadata associated with at least some of the video fragments of the live video content, and/or skip playback of at least one of the one or more expendable portions of the live video content, thereby reducing the delay.

[0057] Optionally, the content metadata may be associated with the live video content by a human operator after generation of the live video content. Optionally, the one or more
25 processors may be configured to determine the delay between the live playhead and the client playhead by determining a cumulative time for recovering from the one or more rebuffering events.

[0058] Embodiments disclosed herein may include a client device, including memory, an output device, and one or more processors configured, in conjunction with the memory, to
30 one or more of, acquire a stream of content for playback on the output device, determine that a delay between a live playhead of the content and a client playhead associated with the

playback of the content exceeds a first threshold, identify an expendable portion of the content, and/or skip playback of the expendable portion of the content.

[0059] Optionally, the one or more processors may be configured to identify the expendable portion of the content (1) using first content metadata associated with one or more fragments of the content by a content provider of the content, (2) using second content metadata associated with one or more fragments of the content by a human operator after generation of the content, and/or (3) by analyzing one or more frames or fragments of the content substantially in real time. Optionally, the one or more processors may be further configured to increase a playback speed of the content. Optionally, the one or more processors may be configured to increase the playback speed of the content in response to the delay exceeding a second threshold, the second threshold being lower than the first threshold.

[0060] Optionally, the one or more processors may be further configured to one or more of determine that the delay has dropped below a second threshold, the second threshold being lower than the first threshold, and/or terminate skipping of expendable content. Optionally, the one or more processors may be further configured to determine the delay based on one or more rebuffering events or using time stamps associated with fragments or frames of the content.

[0061] Embodiments disclosed herein may include a computer-implemented method including one or more of acquiring a stream of content for playback, determining that a delay between a live playhead of the content and a client playhead associated with the playback of the content exceeds a first threshold, identifying an expendable portion of the content, and/or skipping playback of the expendable portion of the content.

[0062] Optionally, identifying the expendable portion of the content may include (1) using first content metadata associated with one or more fragments of the content by a content provider of the content, (2) using second content metadata associated with one or more fragments of the content by a human operator after generation of the content, and/or (3) analyzing one or more frames or fragments of the content substantially in real time. Optionally, the expendable portion of the content may include one or more of black frames, slate frames, credits, an opening montage, a closing montage, a commercial break, a break in action, a replay review, a time out, and/or substantially static content. Optionally, the method may further include increasing a playback speed of the content. Optionally,

increasing the playback speed of the content may occur in response to the delay exceeding a second threshold, the second threshold being different than the first threshold.

5 [0063] Optionally, the method may further include one or more of determining that the delay has dropped below a second threshold, the second threshold being lower than the first threshold, and/or terminating skipping of expendable content. Optionally, the method may further include determining the delay based on one or more rebuffering events or using time stamps associated with fragments or frames of the content.

10 [0064] Embodiments disclosed herein may include a computer-implemented method, including one or more of receiving content from a live content source, encoding the content for streaming to client devices, including associating metadata with portions of the content, the metadata identifying the portions of the content with which the metadata are associated as expendable content, and/or streaming the content to the client devices, at least some of the client devices being configured to use the metadata to skip playback of the portions of the content identified as expendable content.

15 [0065] Optionally, associating the metadata with portions of the content may include (1) using first content metadata associated with the content before the content was received, (2) using second content metadata associated with the content by a human operator after the content was received, and/or (3) analyzing the content substantially in real time. Optionally, the method may further include determining a delay between a live playhead of the content and a client playhead associated with the playback of the content on a first client device. Optionally, determining the delay may include comparing a first time stamp associated with a first fragment of the content requested by the first client device with a second time stamp associated with a second fragment of the content available for streaming to the client devices, or with a third time stamp associated with a third fragment requested by one or
20
25 more other client devices.

[0066] While the subject matter of this application has been particularly shown and described with reference to specific implementations thereof, it will be understood by those skilled in the art that changes in the form and details of the disclosed implementations may be made without departing from the spirit or scope of the invention. Examples of some of
30 these implementations are illustrated in the accompanying drawings, and specific details are set forth in order to provide a thorough understanding thereof. It should be noted that implementations may be practiced without some or all of these specific details. In addition,

well known features may not have been described in detail to promote clarity. Finally, although various advantages have been discussed herein with reference to various implementations, it will be understood that the scope of the invention should not be limited by reference to such advantages. Rather, the scope of the invention should be determined

5 with reference to the appended claims.

CLAIMS

What is claimed is:

1. A client device, comprising:
memory;
5 an output device; and
one or more processors configured, in conjunction with the memory, to:
acquire a stream of content for playback on the output device;
determine that a delay between a live playhead of the content and
a client playhead associated with the playback of the content exceeds a
10 first threshold; and
increase a playback speed of the content.
2. The client device of claim 1, wherein the one or more processors are further
configured to:
determine that the delay is below a second threshold, the second threshold being
15 lower than the first threshold; and
decrease the playback speed of the content.
3. The client device of any of claims 1 or 2, wherein the one or more processors
are further configured to:
determine that the delay exceeds a second threshold, the second threshold being
20 higher than the first threshold; and
further increase the playback speed of the content
4. The client device of any of claims 1, 2, or 3, wherein the one or more
processors are further configured to:
identify one or more expendable portions of the content; and
25 skip playback of at least one of the one or more expendable portions of the content.
5. The client device of claim 4, wherein the one or more processors are further
configured to determine that the delay exceeds a second threshold, the second threshold
being different from the first threshold, and wherein the one or more processors are

configured to skip playback of at least one of the one or more expendable portions of the content in response to determining that the delay exceeds the second threshold.

6. The client device of any of claims 1, 2, 3, 4, or 5, wherein the one or more processors are further configured to determine that there is sufficient available bandwidth for
5 increasing the playback speed of the content.

7. The client device of any of claims 1, 2, 3, 4, 5, or 6, wherein the one or more processors are further configured to:

determine that available bandwidth or a playback quality of the content has dropped below a corresponding threshold; and

10 decrease the playback speed of the content.

8. The client device of any of claims 1, 2, 3, 4, 5, 6, or 7, wherein the one or more processors are further configured to determine the delay based on one or more rebuffering events or using time stamps associated with fragments or frames of the content.

9. A client device, comprising:

15 memory;

an output device; and

one or more processors configured, in conjunction with the memory, to:

acquire a stream of content for playback on the output device;

20 determine that a delay between a live playhead of the content and a client playhead associated with the playback of the content exceeds a first threshold;

identify an expendable portion of the content; and

skip playback of the expendable portion of the content.

10. The client device of claim 9, wherein the one or more processors are
25 configured to identify the expendable portion of the content (1) using first content metadata associated with one or more fragments of the content by a content provider of the content, (2) using second content metadata associated with one or more fragments of the content by a human operator after generation of the content, or (3) by analyzing one or more frames or fragments of the content substantially in real time.

11. The client device of any of claims 9 or 10, wherein the one or more processors are further configured to increase a playback speed of the content.

12. The client device of any of claims 9, 10, or 11, wherein the one or more processors are configured to increase the playback speed of the content in response to the
5 delay exceeding a second threshold, the second threshold being lower than the first threshold.

13. The client device of any of claims 9, 10, 11, or 12, wherein the one or more processors are further configured to:

determine that the delay has dropped below a second threshold, the second
10 thresholding being lower than the first threshold; and

terminate skipping of expendable content.

14. The client device of any of claims 9, 10, 11, 12, or 13, wherein the one or more processors are further configured to determine the delay based on one or more rebuffering events or using time stamps associated with fragments or frames of the content.

15

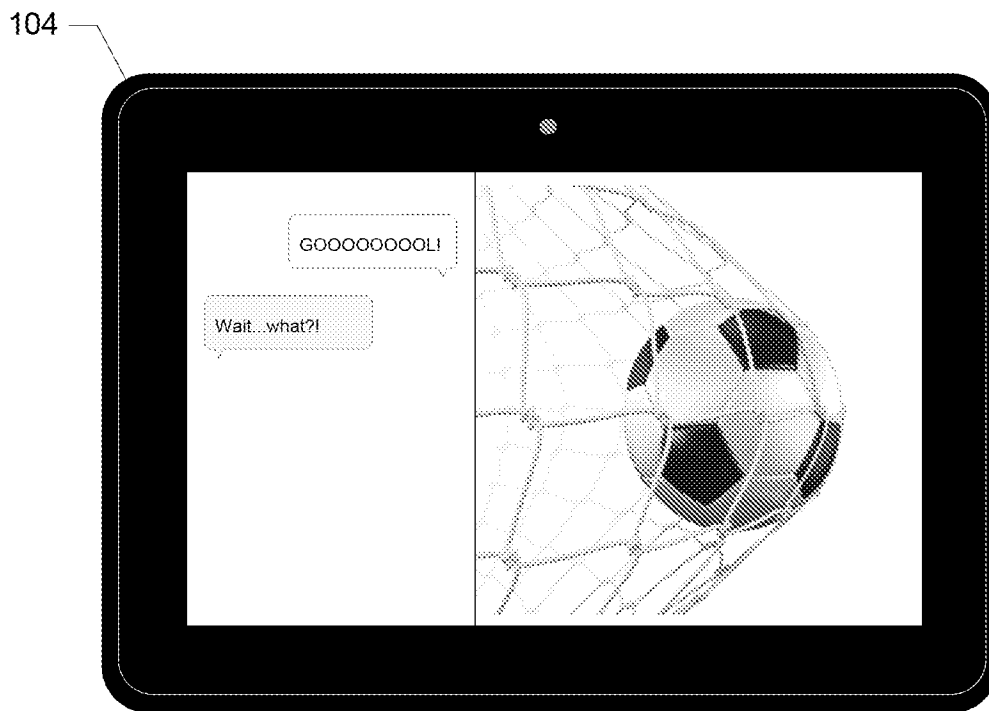
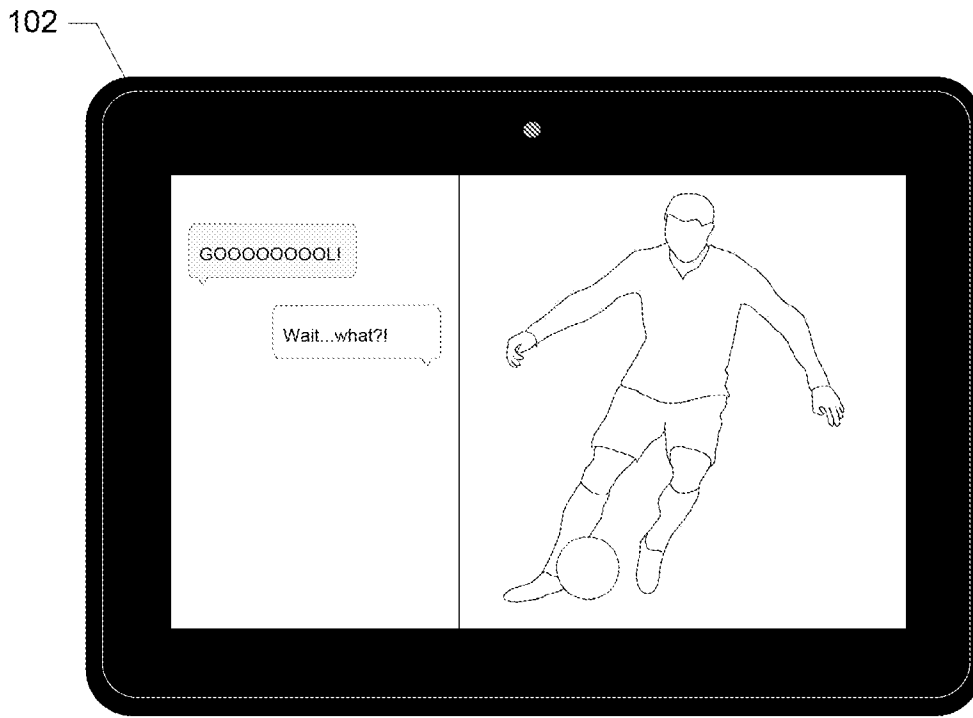


FIG. 1

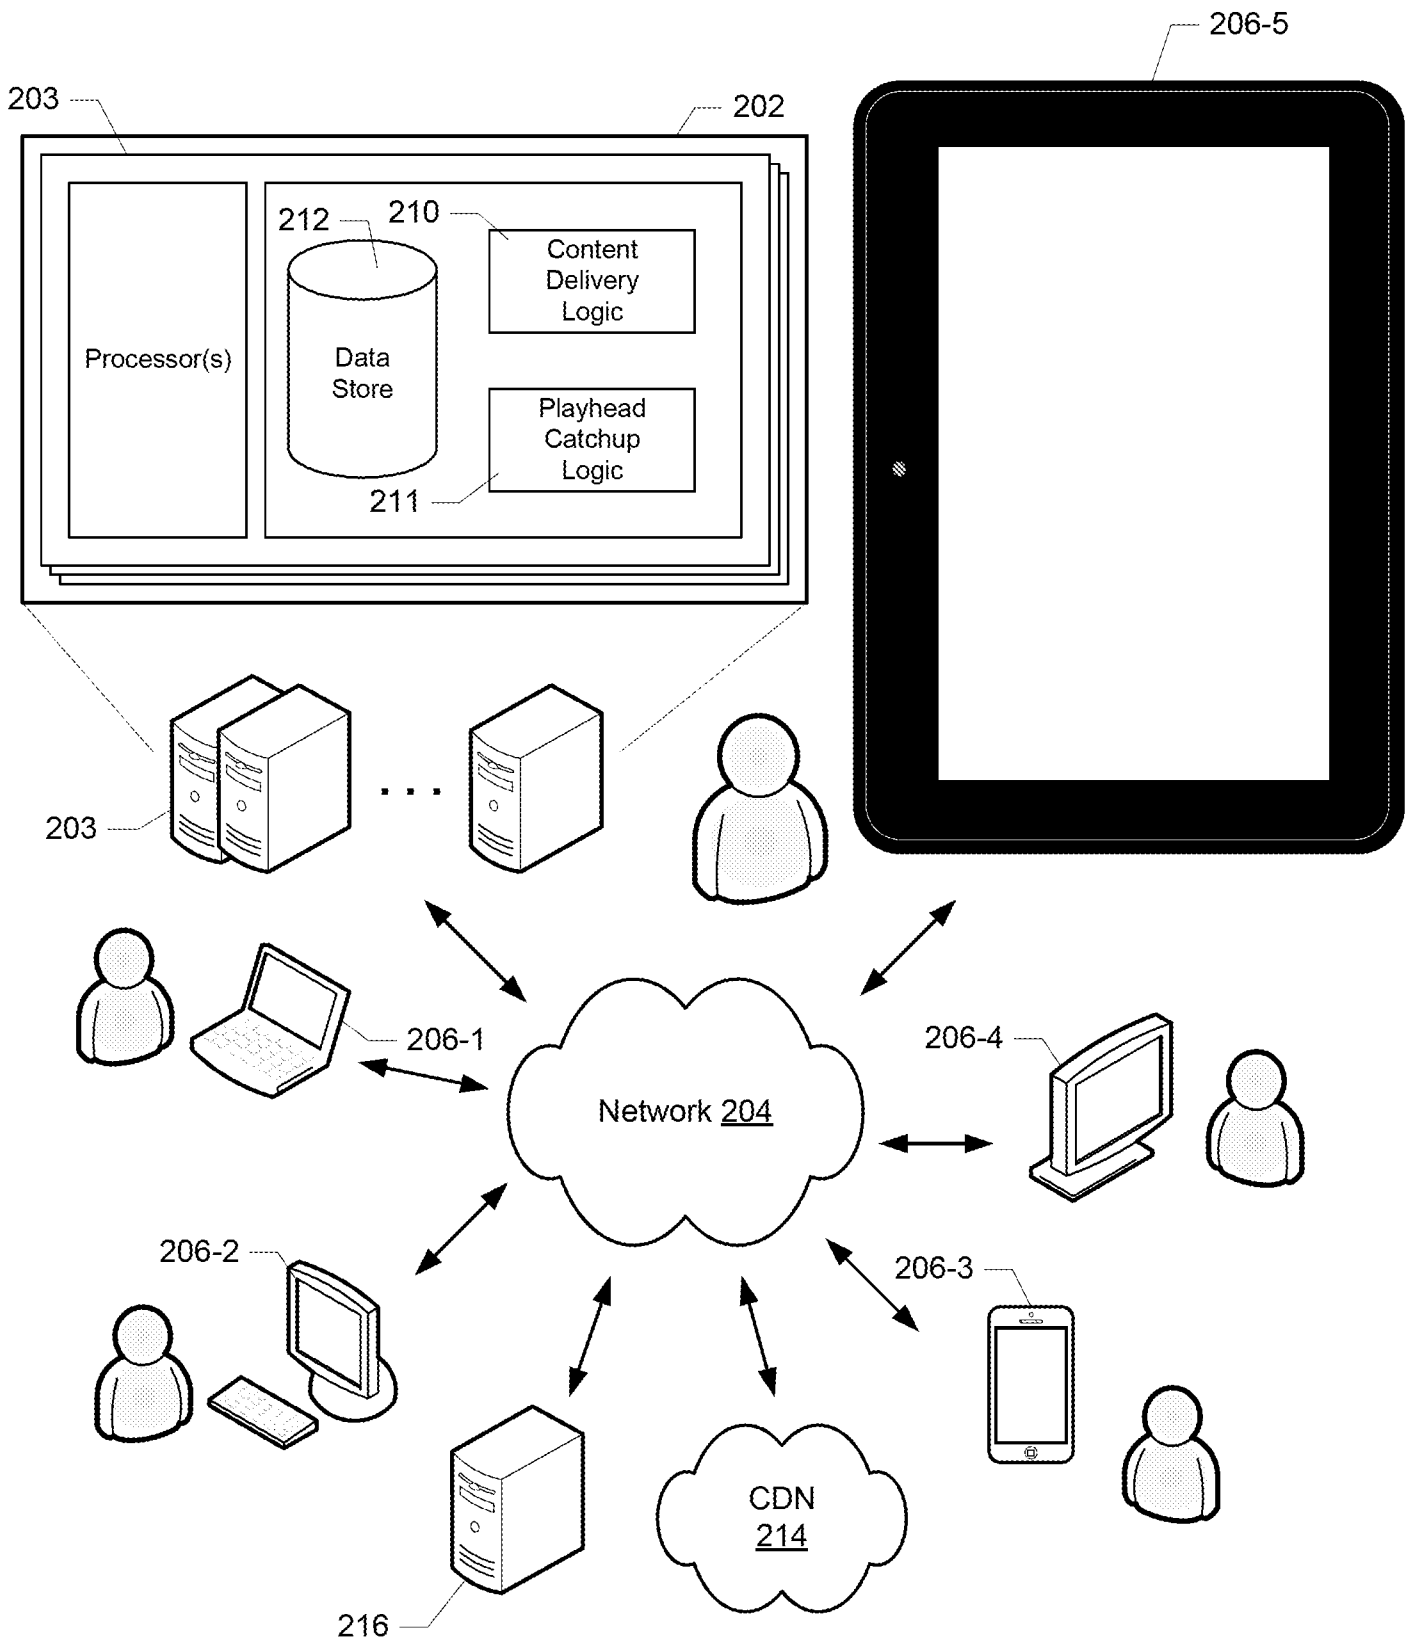


FIG. 2

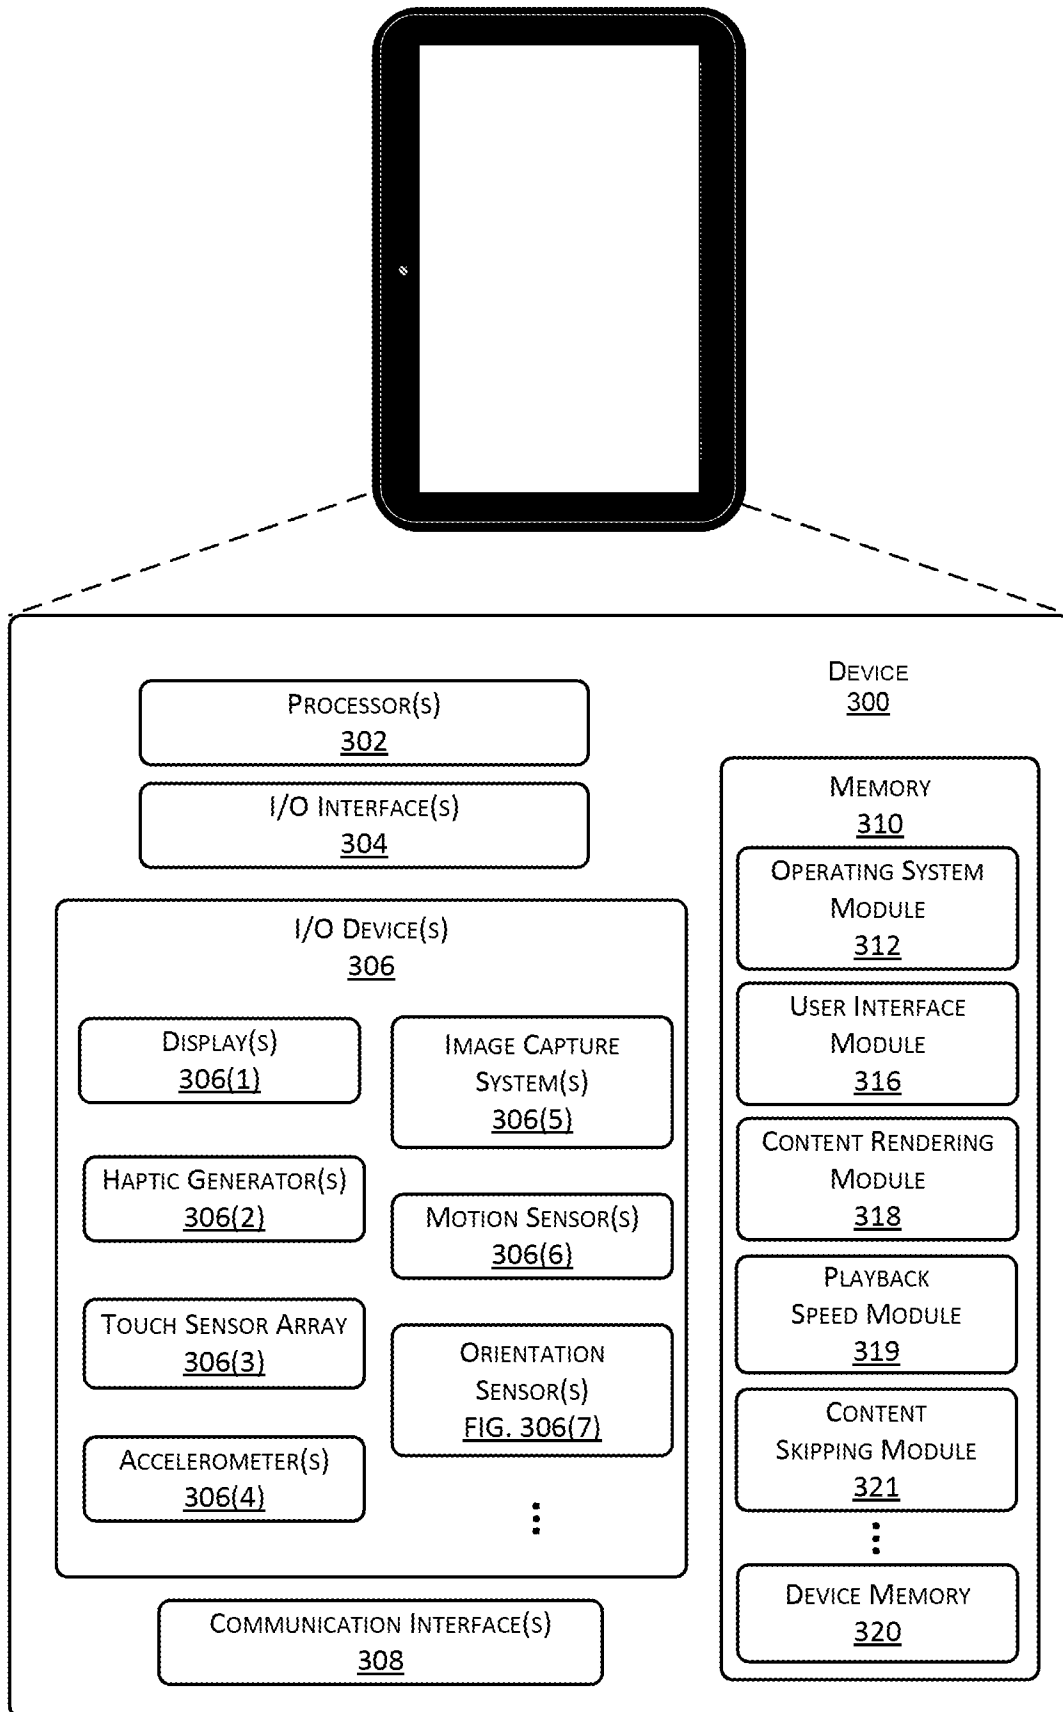
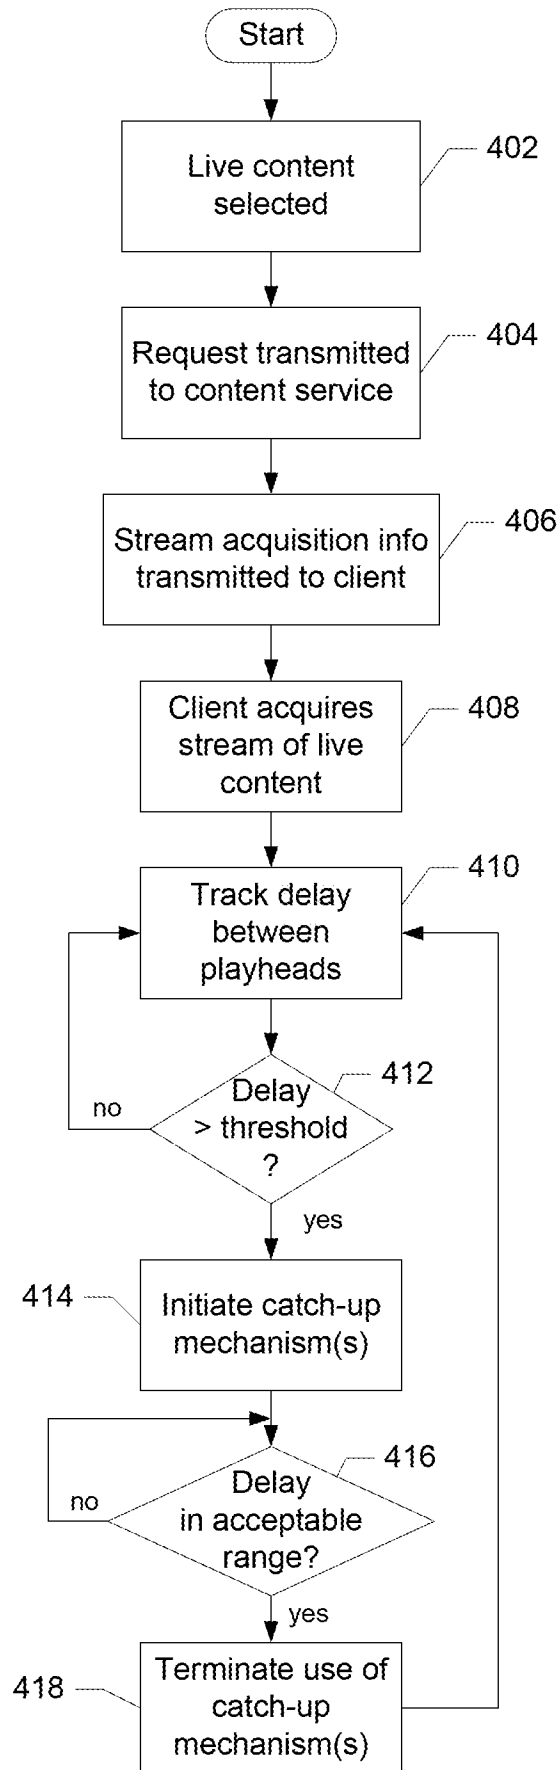


FIG. 3

FIG. 4



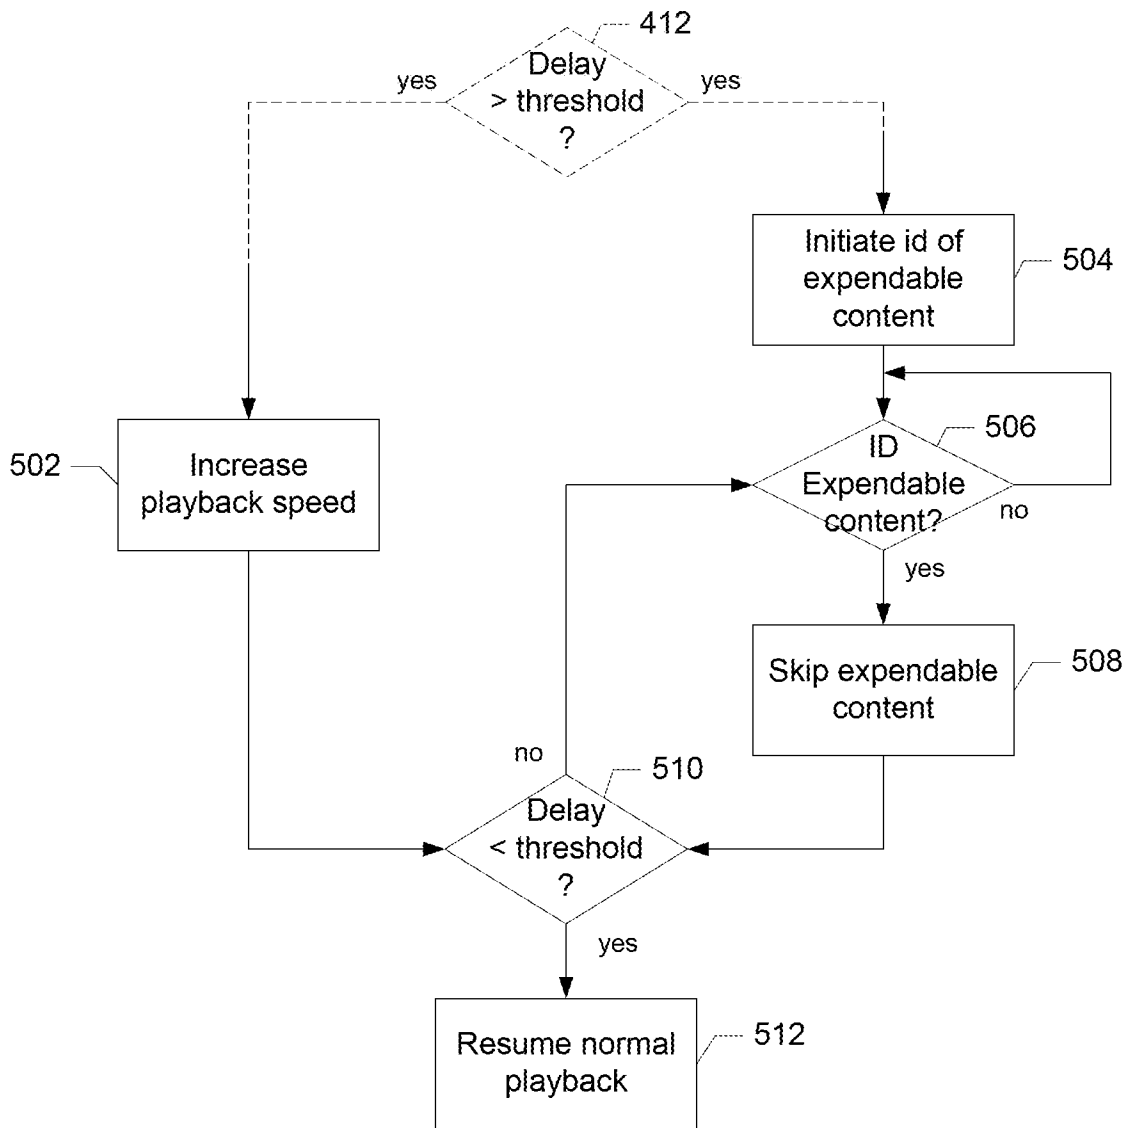


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No PCT/US2017/034066

A. CLASSIFICATION OF SUBJECT MATTER

INV.	H04N21/2187	H04N21/24	H04N21/44	H04N21/4402	H04N21/442
	H04N21/6332	H04N21/84	H04N21/845	H04N21/8547	H04N21/414
	H04N21/2343	H04N21/234			

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	US 2011/246657 A1 (GLOW ANDY [US]) 6 October 2011 (2011-10-06) figure 2 claim 15 paragraph [0048] paragraph [0051] paragraph [0056] paragraph [0087]	1-5,9, 11-13 6-8,10, 14
Y A	----- US 2011/080940 A1 (BOCHAROV JOHN A [US] ET AL) 7 April 2011 (2011-04-07) paragraph [0043] paragraph [0031] paragraph [0033] -----	6-8,10, 14 1-5,9, 11-13

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
 "E" earlier application or patent but published on or after the international filing date
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
 "O" document referring to an oral disclosure, use, exhibition or other means
 "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
 "&" document member of the same patent family

Date of the actual completion of the international search

19 June 2017

Date of mailing of the international search report

26/06/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

Ong, Bee

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2017/034066

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011246657	A1	06-10-2011	NONE

US 2011080940	A1	07-04-2011	CN 102577272 A 11-07-2012
			EP 2486705 A2 15-08-2012
			US 2011080940 A1 07-04-2011
			WO 2011044285 A2 14-04-2011
