

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2015/063732 A1

(43) International Publication Date
7 May 2015 (07.05.2015)

(10) International Publication Number

WO 2015/063732 A1

(51) International Patent Classification:
H04N 21/25 (2011.01)

(72) Inventors: PHILLIPS, Chris; 628 Mustang Drive, Hartwell, Georgia 30643 (US). FORSMAN, Robert; 5177 Allison Way, Sugar Hill, Georgia 30518 (US). REYNOLDS, Jennifer Ann; 3125 Abbotts Pointe Drive, Duluth, Georgia 30097 (US).

(21) International Application Number:
PCT/IB2014/065742

(74) Agents: LIU, Ronald et al.; 6300 Legacy, MS EVR 1-C-11, Plano, Texas 75024 (US).

(22) International Filing Date:
31 October 2014 (31.10.2014)

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

(25) Filing Language: English

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

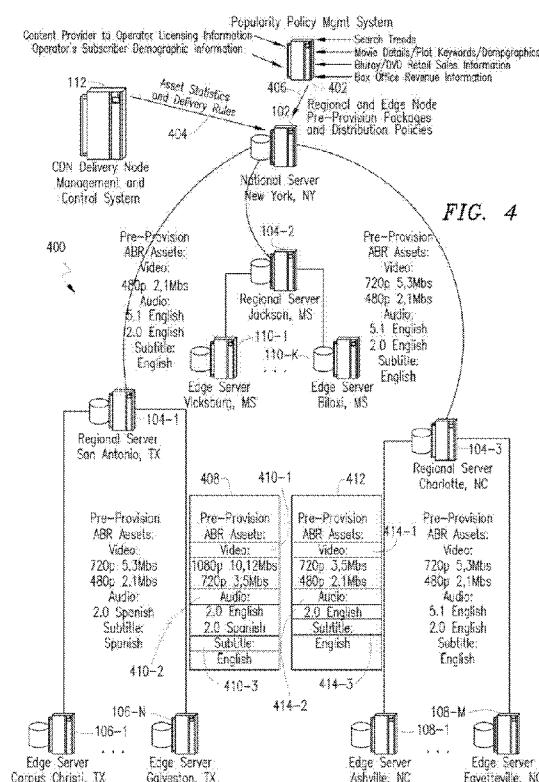
(26) Publication Language: English

(30) Priority Data:
14/069,490 1 November 2013 (01.11.2013) US

(71) Applicant: ERICSSON TELEVISION INC. [US/US];
4500 River Green Parkway, Suite 110, Duluth, Georgia
30096 (US).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR PRE-PROVISIONING ADAPTIVE BITRATE (ABR) ASSETS IN A CONTENT DELIVERY NETWORK



(57) Abstract: A system and method for pre-provisioning ABR assets in a content delivery network (CDN). In an example, historical delivery patterns of one or more ABR assets downloaded at respective content delivery nodes of the CDN are monitored. For each content delivery node, one or more delivery rules may be determined based on statistical distributions of the ABR assets delivered over a period of time. When a particular content is determined to become popular at a specific content delivery node, a pre-provisioning policy may be modulated such that only certain representations or adaptation sets of the particular content's assets that pass the delivery rules are pre-provisioned for the specific content delivery node.



GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

**SYSTEM AND METHOD FOR PRE-PROVISIONING ADAPTIVE BITRATE
(ABR) ASSETS IN A CONTENT DELIVERY NETWORK**

5 **REFERENCE TO RELATED APPLICATION(S)**

This application discloses subject matter that is related to the subject matter of the following U.S. patent application(s): (i) "BANDWIDTH POLICY MANAGEMENT IN A SELF-CORRECTED CONTENT DELIVERY NETWORK" (Ericsson Ref. No.: P40905-US1), Application No.: 13/935,381, filed July 3, 2013, in 10 the name(s) of Christopher Phillips et al., (ii) "SYSTEM AND METHOD FOR DELIVERING CONTENT IN A CONTENT DELIVERY NETWORK" (Ericsson Ref. No.: P40270-US1), Application No.: 13/935,326, filed July 3, 2013, in the name(s) of Christopher Phillips et al., (iii) "DEFRAGMENTATION OF ADAPTIVE STREAMING SEGMENT FILES IN A CONTENT DELIVERY NETWORK" 15 (Ericsson Ref. No.: P40285-US2), Application No.: 13/923,005, filed June 20, 2013, in the name(s) of Christopher Phillips et al., and (iv) "SYSTEM AND METHOD FOR OPTIMIZING DEFRAAGMENTATION OF CONTENT IN A CONTENT DELIVERY NETWORK" (Ericsson Ref. No. P41865-US1), Application No.: 14/069,565 filed November 1, 2013, in the name(s) of Christopher Phillips et al., each of which is 20 hereby incorporated by reference in its entirety.

FIELD OF THE DISCLOSURE

The present disclosure generally relates to communication networks. More particularly, and not by way of any limitation, the present disclosure is directed to 25 systems and methods for optimizing delivery of adaptive bitrate (ABR) assets in a content delivery network.

BACKGROUND

A content delivery network or CDN (sometimes also referred to as a content 30 distribution network) typically distributes content in a "best effort" fashion across the nodes throughout a network using technologies such as adaptive bitrate (ABR) streaming. It is known, however, that ABR can result in unfair and sometimes

incongruous apportionment of bandwidth in a network based upon the order multiple clients begin to use a network resource. Since the bandwidth usage is typically determined by a client requesting the content, and because the ABR clients can be opportunistic, they may become greedy. For instance, a client device that has a low 5 resolution (e.g., a cellular phone) may end up consuming a disproportionate amount of the bandwidth when compared to the bandwidth required for a higher resolution device such as, e.g., a high-definition TV (HDTV). Issues pertaining to such misappropriation of bandwidth and/or uneven distribution of bandwidth resources in a CDN are expected to become even more pronounced when higher bandwidth options become available.

10 Relatedly, when clients access streaming content via a CDN, there may exist conditions at the content delivery nodes (i.e., serving nodes) of the CDN that result in fragmented content (e.g., incomplete copies of the content) locally stored at the content delivery nodes.

15 SUMMARY

The present patent disclosure is broadly directed to systems, methods, devices, apparatuses and associated computer-readable media for pre-provisioning ABR assets in a CDN based on historical delivery trends and utilizing the historical delivery trends in efficiently managing the retrieval of content segments that may be missing from the 20 serving nodes' database caches.

In one aspect, an embodiment of a method operative at a management node associated with a CDN is disclosed. The claimed embodiment comprises, *inter alia*, obtaining historical delivery patterns of one or more ABR assets associated with one or more content delivery nodes of the CDN. For each content delivery node, one or more 25 delivery rules may be determined based on statistical distributions of the ABR assets delivered over a period of time. When a particular content is determined to become popular at a specific content delivery node, a pre-provisioning policy may be modulated such that only certain representations or adaptation sets of the particular content's assets are pre-provisioned for the specific content delivery node based on the one or more delivery rules established for the specific content delivery node. In a further embodiment of a method, historical network conditions associated with one or 30 more content delivery nodes of the CDN for delivery of one or more ABR assets may

be monitored. Based on the historical network conditions, one or more delivery rules may be determined for each content delivery node. Additionally or alternatively, when a particular content is determined to become popular at a specific content delivery node, the CDN's pre-provisioning policies may be modulated based on the delivery 5 rules such that only certain representations or adaptation sets of the particular content's assets that pass or otherwise satisfy the delivery rules are pre-provisioned for delivery at the specific content delivery node.

In another embodiment, a management node operative in association with a CDN is disclosed. The claimed embodiment may be configured as a computer-implemented data processing system and comprises, *inter alia*, one or more processors and a database for storing historical delivery patterns of one or more ABR assets downloaded at the CDN's delivery nodes. A delivery node management and control (DNMC) module may be embodied in a persistent memory coupled to the one or more processors, the DNMC module including program instructions executable by the one or 10 more processors and configured to determine, for each content delivery node, one or more delivery rules based on statistical distributions of the ABR assets delivered over a period of time. Optionally, a content pre-provisioning control module embodied in a persistent memory may be provided, which is coupled to the one or more processors, the content pre-provisioning control module having program instructions executable by 15 the one or more processors and configured to: when a particular content is determined to become popular at a specific content delivery node, apply the one or more delivery rules corresponding to the specific content delivery node against the particular content; and select or otherwise identify only certain representations of the particular content's assets for pre-populating delivery at the specific content delivery node based on the 20 application of the delivery rules established for the specific content delivery node. In a still further embodiment, a management node operative with the CDN may include a database for storing historical network conditions associated with one or more content delivery nodes of the CDN, wherein a DNMC module embodied in the persistent memory may comprise program instructions executable by the one or more processors 25 and configured to determine, for each content delivery node, one or more delivery rules based on the historical network conditions monitored for the ABR assets delivered over a period of time. When a particular content is determined to become popular at a 30

specific content delivery node, the network condition-based delivery rules may be applied to select or otherwise identify only certain representations of the particular content's assets that pass the delivery rules for pre-populating relative to the specific content delivery node.

5 In another aspect, an embodiment of a method operative at a content delivery node for optimizing defragmentation of content is disclosed. The claimed embodiment comprises, *inter alia*, analyzing a manifest available to the content delivery node with respect to a particular content as well as segment files stored in a database cache of the content delivery node to determine if any segment files referenced by the manifest are 10 absent from the database cache. If so, one or more delivery rules associated with the content delivery node may be applied to determine representations of the absent segment files that satisfy the one or more delivery rules. In one variation, the delivery rules may be established based on historical delivery patterns of one or more ABR assets downloaded at the content delivery node. In another variation, the delivery rules 15 may be established based on historical network conditions of the content delivery node. After determining the compliant representations of the absent segment files, they may be pulled from another content delivery node of the CDN (e.g., a parent node or an edge server node).

20 In a further embodiment, a content delivery node configured to optimize content defragmentation is disclosed. The claimed content delivery node may be embodied as a computer-implemented data processing system comprising, *inter alia*, one or more processors and a database cache for storing segment files of content streamed to the content delivery node. A content defragmentation module embodied in a persistent memory is coupled to the one or more processors, wherein the content defragmentation 25 module comprises program instructions executable by the one or more processors and configured to: analyze a manifest available to the content delivery node with respect to a particular content and segment files stored in the database cache in order to determine if any segment files referenced by the manifest are absent from the database; if so, apply one or more delivery rules associated with the content delivery node to determine 30 representations of the absent segment files that satisfy the one or more delivery rules; and pull from another content delivery node of the CDN only the representations of the

absent segment files satisfying the one or more delivery rules and store the pulled representations in the database cache.

In still further aspects, embodiments of a non-transitory computer-readable medium containing instructions stored thereon are disclosed for performing one or 5 more embodiments of the methods set forth above. Additional features of the various embodiments are as recited in the dependent claims.

Advantages of the present invention include, but not limited to, providing a finer-grain control over pre-provisioning of ABR assets in a CDN, thereby resulting in more efficient usage of network resources. Not only are the operator bandwidth costs 10 minimized (since the network operators no longer have to transmit assets, i.e., files, that are unlikely to be used, regardless of how popular a particular content is) but a higher quality of service may also be provided to the end users. Additionally, by optimizing content defragmentation at a CDN's delivery nodes on the basis of asset trend-based delivery rules and/or network condition-based delivery rules, further efficiencies may 15 be gained. Additional benefits and advantages of the embodiments will be apparent in view of the following description and accompanying Figures.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present disclosure are illustrated by way of example, and 20 not by way of limitation, in the Figures of the accompanying drawings in which like references indicate similar elements. It should be noted that different references to "an" or "one" embodiment in this disclosure are not necessarily to the same embodiment, and such references may mean at least one. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is 25 submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

The accompanying drawings are incorporated into and form a part of the specification to illustrate one or more exemplary embodiments of the present disclosure. Various advantages and features of the disclosure will be understood from 30 the following Detailed Description taken in connection with the appended claims and with reference to the attached drawing Figures in which:

FIG. 1 depicts an example content delivery network wherein one or more embodiments of the present patent application may be practiced;

FIG. 2 depicts an example manifest including references to one or more content representations or adaptations sets relative to video, audio and subtitle components of a 5 particular content;

FIG. 3 depicts an example Media Presentation Description (MPD) model of a manifest illustrative of further details relative to different media segments of a particular content;

FIG. 4 depicts an example content delivery network wherein content is 10 pre-provisioned based on historical delivery trends according to an embodiment of the present patent application;

FIGS. 5A and 5B depict flowcharts with blocks relative to various steps and acts that may take place according to one or more embodiments of the present patent application;

15 FIG. 6 depicts a block diagram of a computer system adapted to operate as a management node according to an embodiment of the present patent application;

FIG. 7 is a flowchart with blocks relative to various steps and acts that may take place for monitoring historical delivery trends of ABR assets according to an embodiment of the present patent application;

20 FIG. 8 is a flowchart with blocks relative to various steps and acts that may take place for pre-provisioning of content based on delivery rules according to an embodiment of the present patent application;

FIG. 9 is a flowchart with blocks relative to various steps and acts that may take 25 place for updating historical delivery trend data for one or more content delivery nodes of a CDN according to an embodiment of the present patent application;

FIGS. 10A-10C depict flowcharts with blocks relative to various steps and acts that may take place for defragmentation of media segments in a content delivery node of according to one or more embodiments of the present patent application;

30 FIG. 11 depicts a block diagram of a computer system adapted to operate as a content delivery node that may be configured to effectuate optimized content defragmentation according to an embodiment of the present patent application; and

FIG. 12 depicts an example content delivery network wherein fragmented content is replenished at a content delivery node based on historical delivery trends according to an embodiment of the present patent application.

5 DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, numerous specific details are set forth with respect to one or more embodiments of the present patent disclosure. However, it should be understood that one or more embodiments may be practiced without such specific details. In other instances, well-known circuits, subsystems, components, structures and techniques have not been shown in detail in order not to obscure the understanding of the example embodiments. Accordingly, it will be appreciated by one skilled in the art that the embodiments of the present disclosure may be practiced without such specific components-based details. It should be further recognized that those of ordinary skill in the art, with the aid of the Detailed Description set forth herein and taking reference to the accompanying drawings, will be able to make and use one or more embodiments without undue experimentation.

Additionally, terms such as “coupled” and “connected,” along with their derivatives, may be used in the following description, claims, or both. It should be understood that these terms are not necessarily intended as synonyms for each other. “Coupled” may be used to indicate that two or more elements, which may or may not be in direct physical or electrical contact with each other, co-operate or interact with each other. “Connected” may be used to indicate the establishment of communication, i.e., a communicative relationship, between two or more elements that are coupled with each other. Further, in one or more example embodiments set forth herein, generally speaking, an element, component or module may be configured to perform a function if the element is capable of performing or otherwise structurally arranged to perform that function.

As used herein, a network element or node may be comprised of one or more pieces of service network equipment, including hardware and software that communicatively interconnects other equipment on a network (e.g., other network elements, end stations, etc.), and is adapted to host one or more applications or services with respect to a plurality of subscribers. Some network elements may comprise

“multiple services network elements” that provide support for multiple network-based functions (e.g., content popularity policy management, session control, QoS policy enforcement, bandwidth scheduling management, subscriber/device policy and profile management, content provider priority policy management, streaming policy management, and the like), in addition to providing support for multiple application services (e.g., data and multimedia applications). Subscriber end stations or client devices may comprise any device configured to execute, *inter alia*, at least one streaming client application (e.g., an ABR streaming client application) for receiving content from a streaming server or content provider. Accordingly, such client devices 5 may include set-top boxes, PVR/DVRs, workstations, laptops, netbooks, palm tops, mobile phones, smartphones, multimedia phones, Voice Over Internet Protocol (VOIP) phones, mobile/wireless user equipment, high definition TV terminals, portable media players, location-aware subscriber equipment, gaming systems or consoles (such as the Wii®, Play Station 3®, Xbox 360®), etc., that may access or consume content/services 10 provided over a content delivery network in accordance with one or more embodiments 15 set forth herein. Further, the client devices may also access or consume content/services provided over broadcast networks (e.g., cable and satellite networks) as well as a packet-switched wide area public network such as the Internet via suitable service provider access networks. In a still further variation, the client devices or 20 subscriber end stations may also access or consume content/services provided on virtual private networks (VPNs) overlaid on (e.g., tunneled through) the Internet.

One or more embodiments of the present patent disclosure may be implemented using different combinations of software, firmware, and/or hardware. Thus, one or more of the techniques and blocks shown in the Figures (e.g., flowcharts) may be 25 implemented using code and data stored and executed on one or more electronic devices or nodes (e.g., a subscriber client device or end station, a network element, etc.). Such electronic devices may store and communicate (internally and/or with other electronic devices over a network) code and data using computer-readable media, such as non-transitory computer-readable storage media (e.g., magnetic disks, optical disks, 30 random access memory, read-only memory, flash memory devices, phase-change memory, etc.), transitory computer-readable transmission media (e.g., electrical, optical, acoustical or other form of propagated signals – such as carrier waves, infrared

signals, digital signals), etc. In addition, such network elements may typically include a set of one or more processors coupled to one or more other components, such as one or more storage devices (e.g., non-transitory machine-readable storage media) as well as storage database(s), user input/output devices (e.g., a keyboard, a touch screen, a 5 pointing device, and/or a display), and network connections for effectuating signaling and/or bearer media transmission. The coupling of the set of processors and other components may be typically through one or more buses and bridges (also termed as bus controllers), arranged in any known (e.g., symmetric/shared multiprocessing) or heretofore unknown architectures. Thus, the storage device or component of a given 10 electronic device or network element may be configured to store code and/or data for execution on one or more processors of that element, node or electronic device for purposes of implementing one or more techniques of the present disclosure.

Referring now to the drawings and more particularly to FIG. 1, depicted therein is an example content delivery network or CDN 100 wherein one or more embodiments 15 of the present patent application may be practiced. For purposes of the present patent application, CDN 100 may comprise an overlay network architected for high-performance streaming of a variety of digital assets or program assets as well as services (hereinafter referred to as "media content") to subscribers using one or more Internet-based infrastructures, private/dedicated infrastructures or a combination thereof. In general, the terms "media content" or "content file" (or, simply "content") 20 as used in reference to at least some embodiments of the present patent disclosure may include streaming digital assets and program assets such as any type of audio/video content or program segment, live or static (e.g., recorded over-the-air free network television (TV) shows or programs, pay TV broadcast programs via cable networks or 25 satellite networks, free-to-air satellite TV shows, IPTV programs, etc.), Over-The-Top (OTT) and video-on-demand (VOD) or movie-on-demand (MOD) shows or programs, time-shifted TV (TSTV) content, as well as other content assets provided by content publishers, owners or providers, including but not limited to software files, executable computer code or programs, online electronic games, Internet radio shows/programs, 30 entertainment programs, educational programs, movies, music video programs, and the like, that may be delivered using any known or heretofore unknown adaptive streaming technologies. Further, various programs or content files provided via streaming may be

arranged as a collection or assembly of channels that are specific to different subscribers, wherein different channels may comprise media content from one or more content sources or originators.

By way of illustration, content may be delivered via CDN 100 using adaptive bit rate (ABR) streaming technologies and may be encoded to support known implementations such as Microsoft® Silverlight® Smooth Streaming, HTTP streaming (for instance, Dynamic Adaptive Streaming over HTTP or DASH, HTTP Live Streaming or HLS, HTTP Dynamic Streaming or HDS, etc.), icecast, and so on. In general, the overlay architecture of CDN 100 may include a multi-level, hierarchically-organized interconnected assembly of network servers for providing media pathways or "pipes" from one or more central distribution nodes to one or more levels of regional distribution nodes that are connected to one or more local edge servers configured to serve a plurality of end users or subscribers in respective serving location areas. It should therefore be appreciated that a regional distribution node may operate as a parent node to one or more child edge servers and a central or national distribution node may in turn operate as a parent node to one or more child regional distribution nodes. Further, in addition to such "distribution servers" (sometimes also referred to as "content delivery nodes"), CDN 100 may also include and/or interoperate with various network elements configured to effectuate request redirection or rerouting mechanisms as well as related back office systems such as subscriber management systems, bandwidth scheduling systems, account/billing systems and the like, that may be deployed as part of a streaming network back office (not specifically shown).

As illustrated, CDN 100 is exemplified with a New York-based national server 102 coupled to regional distribution servers 104-1 to 104-3 based in San Antonio (TX), Jackson (MS) and Charlotte (NC), respectively. Regional distribution server 104-1 is adapted to serve a plurality of edge servers 106-1 to 106-N, exemplified by edge server 106-1 in Corpus Christi (TX) and edge server 106-N in Galveston (TX). Likewise, regional distribution server 104-2 is adapted to serve edge servers 110-1 (in Vicksburg, MS) to 110-K (in Biloxi, MS) and regional distribution server 104-3 is adapted to server edge servers 108-1 (in Ashville, NC) to 108-M (in Fayetteville, NC).

Each local edge server is operable to serve a plurality of subscribers in a serving area for delivering content to associated end stations or client devices, as illustrated by

an example client device or user equipment (UE) device 105 that may for consume or access streaming content delivered via CDN 100 in any type or number of access technologies including broadband access via wired and/or wireless (radio) communications. For purposes of the present patent application, the terms “subscriber 5 end station” and “client device” may be used synonymously and may comprise any UE device or appliance that in one implementation not only receives digital content assets for live viewing, playback and/or decoding the content, but also operates as a command console or terminal that can accept user inputs, commands or requests to interact with a network element disposed in CDN 100 and/or associated streaming server systems for 10 requesting content that may be selectively rendered at an internal display screen and/or one or more external audio/visual (A/V) devices (not specifically shown). As such, the example client device 105 may include one or more streaming client modules (e.g., an ABR streaming client) and associated decoding functionalities depending on the 15 streaming technologies implemented, each operating in association with a processor module and suitable memory and program code (not shown) for effectuating acquisition, decoding and rendering of the streamed media content.

As alluded to previously, CDN 100 may be configured to deliver content from live sources and/or static file sources using adaptive streaming wherein the content may comprise video components, audio track components as well as subtitle language components. Further, the video components of a particular content may be transcoded or otherwise encoded with different bit rates (e.g., multi-rate transcoding) using applicable encoder(s) (e.g., a particular program content may be transcoded into five video files using variable bit rates, ranging from low to high bit rates (500 Kbs to 12 Mbs, by way of illustration). The particular content's video component is therefore 20 encoded as five different “versions” or “formats”, wherein each bit rate may be referred to as a profile or representation. Also, the audio tracks associated with particular content may include different languages and encodings to support multi-lingual 25 programming and different levels/types of audio quality experience. For instance, content may be available in multiple languages (e.g., English, French, Spanish, German, Chinese, etc.), each with several types of audio encoding: single channel 30 (mono), two-channel (stereo), 5.1 surround sound, 7.1 surround sound, 7.2 surround sound, etc. Likewise, regardless of the languages of the main audio track of a content,

various language choices for subtitles may also be provided with the content (e.g., English, Spanish, Japanese, Chinese, Korean, French, etc.).

For purposes of the present patent application, video, audio and subtitle components of the content that is distributed through CDN 100 may be referred to as "ABR assets". As the content is accessed, downloaded or otherwise consumed by the subscribers via the network of content delivery nodes of CDN 100, different consumption/downloading patterns may emerge with respect to the ABR assets. A delivery node management and control (DNMC) system 112 may be provided in association with CDN 100 as a management node adapted to monitor historical delivery trends and patterns of the ABR assets as well as network conditions associated with one or more content delivery nodes of CDN 100. As client devices 105 access various content via associated content delivery nodes over a configurable period of time, a historical delivery pattern learning and updating process is operative to obtain delivery statistics for each content delivery node on a component-by-component basis, i.e., video, audio or subtitles. In similar fashion, historical network conditions associated with each content deliver node may be obtained over a period (e.g., download rate distributions, bandwidth utilization and other performance metrics, relative performance comparisons against the performance of other content delivery nodes and/or predetermined performance thresholds, etc.), which may be utilized by DNMC system or node 112 for controlling content distribution through the network.

In certain embodiments, statistics associated with the edge server nodes may be directly provided to DNMC system 112, or indirectly via respective regional server nodes, which in turn may provide cumulative or aggregate statistics of the edge server nodes to DNMC system 112. By way of illustration, reference numeral 114-1 refers to a communication pathway for updating and/or otherwise providing delivery trends and other statistics to DNMC system 112 by a regional server node (e.g., San Antonio server 104-1). Likewise, a communication pathway 114-2 may be provided for updating and/or otherwise providing delivery trends and other statistics to DNMC system 112 by an edge server node (e.g., Galveston server 106-N). Based on the delivery statistics, a node-by-node delivery pattern may be established for each ABR asset component, which may involve determining statistical distributions, likelihood ranges, percentages, expectation values, etc. For example, an ABR asset delivery

pattern 116 associated with the Galveston edge server 106-N includes a video component delivery pattern 118-1 that shows a distribution of various video bit rates (sometimes also referred to as bitrates) delivered from that server over a certain period of time (e.g., 35% of downloaded content being high definition, 1080p video encoded 5 at 8-10 Mbs; 45% of downloaded content being high definition, 720p video encoded at 3-8 Mbs, and so on). An audio component delivery pattern 118-2 associated with the Galveston edge server 106-N likewise indicates a distribution of various languages and encodings used for all the audio tracks delivered from that server over a period of time (e.g., 10% of downloaded content being in English with 5.1 surround sound; 50% of 10 downloaded content being in English with 2.0 stereo; and 40% of downloaded content being in Spanish with 2.0 stereo). As to subtitles, a corresponding delivery pattern 118-3 associated with the Galveston edge server 106-N shows a distribution wherein 99% of downloaded subtitles being in English and the remaining 1% of downloaded subtitles being in Spanish.

15 In comparison, an ABR asset delivery pattern 120 associated with the Ashville edge server 108-1 is illustrated with a different set of video, audio and subtitle distributions. As shown in FIG.1, the ABR asset delivery pattern 120 includes a video component delivery pattern 122-1 that exemplifies a distribution of various video bit rates delivered from that server over a certain period of time, with a relatively low percentage (4%) of downloaded content being high definition, 1080p video encoded at 20 8-10 Mbs, for instance. Likewise, an audio component delivery pattern 122-2 associated with the Ashville edge server 108-1 indicates an audio track language distribution having a significantly lower percentage of Spanish language (e.g., 2% of downloaded content). A subtitle delivery pattern 120-3 is indicative of a distribution of 25 the subtitle selections downloaded from the Ashville edge server 108-1 over a period of time that is substantially similar to the subtitle delivery pattern 118-3 associated with the Galveston edge server 106-N.

30 In accordance with teachings of the present patent application, DNMC system 112 is operative to determine or otherwise establish one or more delivery rules that may be based on the differential statistical distributions of various ABR asset delivery patterns and/or historical network conditions associated with respective content delivery nodes of CDN 100. Additionally, such delivery rules may be advantageously

5 applied for modulating or otherwise controlling content pre-provisioning techniques (e.g., based on popularity determinations), defragmentation methodologies, push-based content delivery policies, etc. that may be implemented within CDN 100 for purposes of efficient resource utilization, optimized user experience, bandwidth management, and the like, as will be set forth in additional detail hereinbelow.

10 FIG. 2 depicts an example ABR asset package 200 relative to a particular content (e.g., Movie X) that may be available for distribution throughout a CDN such as CDN 100, wherein the delivery of one or more portions of the ABR asset package may be modulated or otherwise controlled according to an embodiment of the present 15 patent application. A manifest file 202 (also known as Media Presentation Description or MPD model in MPEG-DASH standard ISO/IEC 23009-1) is operative to describe available content components, including references to one or more content representations or adaptations sets relative to video, audio and subtitle components of the particular content. In respect of video asset components, the manifest 202 may 20 describe segment information such as timing, Universal Resource Locators (URLs) pointing to segments as well as media characteristics such as video resolution and bit rates. In an embodiment where the manifest 202 is provided as a master manifest, it may include a plurality of child manifests 204-1 to 204-K, each referencing a particular combination of video, audio and subtitle asset components of Movie X. By way of 25 illustration, a video component portion 206 of the ABR asset package 200 includes media segments encoded at a number of bit rates, each having certain video quality. For example, reference numeral 206-1 refers to high definition 1080p media segments encoded at 12Mbs. Likewise, reference numerals 206-2 to 206-9 refer to media segments whose video is encoded at rates all the way down to 450Kbs. An audio 30 component portion 208 of the ABR asset package 200 includes multiple languages of Movie X, each being encoded for different audio quality experiences. Reference numerals 208-1 to 208-6 refer to six asset components, English 5.1 (surround sound), English 2.0 (basic stereo), Spanish 2.0, French 2.0, Chinese 2.0 and German 2.0, respectively. A subtitles component portion 210 includes a plurality of subtitle language selections provided for Movie X, with English 210-1 to Russian 210-9 being illustrative.

FIG. 3 depicts an example Media Presentation Description model of a manifest, e.g., MPD 300 that is illustrative of further details relative to different media segments of a particular content. A plurality of arbitrarily spliced content portions 302-1 to 302-3 describe the spliced portions' respective timing and base URL parameters, each 5 including one or more adaptation sets for different asset components. As illustrated, reference numerals 304-1 and 304-2 respectively refer to a video adaptation set and an audio adaptation set of segment 302-2. The video adaptation set 304-1 in turn includes two representations 306-1 and 306-2, each having respective bit rates, display resolutions, etc, in addition to the segment information 310. By way of further 10 illustration, segment information 310 includes an initialization segment 312 as well as URLs to a plurality of media segments 314-1 to 314-L, each segment having a suitable timing parameter. Those skilled in the art should recognize that the ABR asset package 200 shown in FIG. 2 may be described in a suitable MPD model representation similar 15 to MPD 300 shown in FIG. 3 in a particular implementation that is compliant with the MPEG-DASH standard.

FIG. 4 depicts an example content delivery network 400 wherein content may be pre-provisioned based on historical delivery trends according to an embodiment of the present patent application. For purposes of illustration, CDN 400 is identical to CDN 100 shown in FIG. 1 except that a popularity policy management system 402 is 20 provided for determining or otherwise estimating or predicting where a particular content might become popular. As described in the commonly assigned patent applications (i) "BANDWIDTH POLICY MANAGEMENT IN A SELF-CORRECTED CONTENT DELIVERY NETWORK" (Ericsson Ref. No.: P40905-US1), Application No.: 13/935,381, filed July 3, 2013, and (ii) "SYSTEM AND 25 METHOD FOR DELIVERING CONTENT IN A CONTENT DELIVERY NETWORK" (Ericsson Ref. No.: P40270-US1), Application No.: 13/935,326, filed July 3, 2013, incorporated by reference hereinabove, the popularity policy management system 402 is operative to estimate or determine a particular content's popularity by location and subscriber demographics based on a plurality of information sources, e.g., 30 search trend data relative to a particular content to be distributed in a geographical region serviced by CDN 400 (for instance, Internet search engine trend data using search engines offered by Google®, Bing®, Yahoo®, and the like, as well as social

media search trend data, e.g., searches on Facebook®, Twitter®, Instagram®, Pinterest®, and the like.), in addition to trends relating to content/movie details, plot keyword search data, marketing data (e.g., box office revenue data, stored media (e.g., Blu-Ray, DVD, HD, and CD discs) sales data and related content sales data 5 (collectively, “revenue information”). Further, the popularity policy management system 402 may be configured to interface with a content policy management node (not explicitly shown) for receiving applicable licensing information an/or content-based priority levels and weights. Upon determining popularity estimations for a particular content, the popularity policy management system 402 is operative to provide pre- 10 provisioning and distribution policies to a central distribution node and/or regional server nodes, as exemplified by a communication path 406. Additional details regarding the establishment of content-specific pre-provisioning and distribution policies may be found in one or more of the foregoing commonly assigned patent applications, referenced hereinabove.

15 Further, DNMC system 112 is operative to provide one or more delivery rules established or otherwise determined for the respective content delivery nodes (e.g., regional server nodes and/or edge server nodes) based on the nodes’ ABR asset delivery/distribution patterns. In one embodiment, DNMC system 112 is further operative to apply such delivery rules (which may comprise upper/lower thresholds 20 and/or ranges for video bit rates, limits based on language/subtitle selection percentages, and the like) in conjunction with the pre-provisioning and distribution policies for a particular ABR asset package such that only certain representations or versions of the content which satisfy or otherwise pass the delivery rules may be pre- 25 provisioned. In another embodiment, applicable delivery rules may be transmitted to a central distribution node and/or regional distribution nodes, as exemplified by a communication path 404, which distribution node may then apply the delivery rules in order to select a subset of the video/audio/subtitle assets for pre-provisioning relative to a plurality of the edge server nodes where the content has been determined to be popular. In a still further embodiment, both the popularity-based pre-provisioning 30 policies as well as node-specific asset trend delivery rules may be provided to another network element or management node that is adapted to apply suitable selection/filtering logic in order for establishing filtered or otherwise modulated pre-

provisioning policies with respect a particular content's ABR assets. In such a scenario, modulated pre-provisioning policies may be transmitted to the central/regional distribution nodes of CDN 400 and a subset of the ABR assets may be appropriately pre-provisioned in accordance therewith.

5 Additionally or alternatively, DNMC system 112 may also provide one or more delivery rules established or otherwise determined for the respective content delivery nodes based on the nodes' the historical network conditions. Similar to the application of the asset trend-based delivery rules, network condition-based delivery rules may be applied for modulating or otherwise controlling the pre-provisioning policies
10 established by the popularity policy management system 402. It should be appreciated by one skilled in the art that the network condition-based delivery rules may be applied in various implementations similar to the embodiments described above, in addition to being applied in conjunction with the asset trend-based delivery rules for further refinement of a CDN's pre-provisioning policies.

15 Continuing to refer to FIG. 4, the modulated pre-provisioning of ABR assets with respect to edge server nodes 106-N and 108-1 based on respective node's delivery patterns may be further exemplified. As described previously, the ABR asset delivery pattern 116 associated with the Galveston edge server 106-N shows a video component distribution wherein more than 80% of the content is of higher bitrate category (i.e.,
20 720p content at 3-8Mbs or 1080p content at 8-12Mbps). Accordingly, a delivery rule may be established wherein only higher bitrate representations of a particular content (e.g., Movie Y) may be pre-provisioned for the Galveston node 116-N. When such a rule is applied with respect to a particular content that has been estimated to become popular at this node, a pre-provisioned ABR asset package 408 for the node will
25 include only 720p and 1080p representations of the content (encoded at, for example, 3/5 Mbs and 10/12Mbs, respectively), as illustrated by a video asset component 410-1. Likewise, given that the historical audio asset delivery pattern for the Galveston node 116-N is where only 10% of the content has been in English with surround sound (5.1), a delivery rule may be such that 5.1 English language representations of a content's
30 audio assets are excluded from pre-provisioning. In similar fashion, consistent with the distribution of historical subtitle languages downloaded at the Galveston node 106-N, a delivery rule for the subtitles may be one where any subtitles having a least likelihood

of being downloaded from that node (e.g., less than 5%) are excluded from pre-provisioning. By applying such delivery rules, it can be seen that the pre-provisioned ABR asset package 408 relative to Movie Y for the Galveston node 106-N will include an audio asset component 410-2 comprising only English 2.0 and Spanish 2.0 representations (or adaptation sets) and a subtitle asset component 410-3 comprising only English language subtitles of Movie Y, which has been estimated to become popular in the area serviced by the Galveston node 106-N.

In similar fashion, one or more delivery rules based on the learned ABR asset delivery pattern 120 associated with the Ashville edge server 108-1 may be applied to modulate the pre-provisioning of a content, e.g., Movie Y, that is estimated to become popular in the Ashville service area. As the historical video asset delivery pattern 122-1 has established that only 4% of the content delivered in that service area comprises 1080p content at 8-12Mbps, an example delivery rule may be such that 1080p representations of a content are excluded. Likewise, suitable thresholds may be established for the audio and subtitle asset components based on their respective likelihood estimates. Upon applying such delivery rules, it can be seen that the pre-provisioned ABR asset package 412 relative to Movie Y for the Ashville node 108-1 will include a video asset component 414-1 comprising only 720p and 480p representations of the content (encoded at, for example, 3/5 Mbs and 2/1 Mbs, respectively), an audio asset component 414-2 comprising only English 2.0 representations and a subtitle asset component 414-3 comprising only English language subtitles of Movie Y.

Accordingly, instead of pre-provisioning all the adaptation sets of an ABR asset package relative to a particular content that has been determined to be popular at certain nodes, the embodiments set forth herein provide finer-grain pre-provisioning of the ABR assets, resulting in more efficient deployment of network resources. Not only are the operator bandwidth costs minimized (since the network operators no longer have to transmit assets, i.e., files, that are unlikely to be used, regardless of how popular a particular content is) but a higher quality of service may also be provided to the end users. Furthermore, although the example pre-provisioning modulations described above are illustrative of application of asset trend-based delivery rules,

similar pre-provisioning modulations may be equally achieved on the basis of network condition-based delivery rules, *mutatis mutandis*.

In a further example scenario, a DNMC system may be made aware of language preferences in the area serviced by each of the CDN edge nodes, as such awareness can be informed by historical trends within the CDN itself (e.g., users in a particular service area always or most often watch the Spanish language version), or based on wider trends supplied from third parties (e.g., the particular service area has a large Spanish speaking population) if the edge node has no historical data yet. Should a particular piece of content become popular in an area serviced by that edge node, the system only pre-populates the popular audio or subtitle data (which may even be different languages), rather than all audio or subtitle options normally included with each piece of content as part of the entire ABR asset package. Should a user want access to a language other than the prevailing one (i.e., outside the filtered pre-provisioned assets), the audio/subtitle data can be sent over in the normal fashion, on demand. In a still further example, the DNMC system may be made aware of historical network conditions for each of its edge servers. Should a particular edge server experience or exhibit consistently below average edge server conditions (as may be defined for a suitable performance metric), the DNMC system is made aware of such a condition. If a particular piece of content become popular in an area serviced by that edge node, the system only pre-populates lower bitrate content, as the higher bitrates are less likely to be used (and/or less likely to be downloaded). If the network conditions improve and atypical bitrates (outside the filtered pre-provisioned assets) are requested by the subscribers, such bitrates can be sent over in the normal fashion, on demand.

Turning now to FIGS. 5A and 5B, depicted therein are flowcharts with blocks relative to various steps and acts that may take place according to one or more embodiments of the present patent application for modulating pre-provisioning of ABR assets. Flowchart 500A is representative of a collection of processes that may be practiced either independently or in one or more combinations. In one example process, blocks 502, 504 and 510 illustrate an application of asset trend-based delivery rules with respect to pre-provisioning of popular content. In another example process, blocks 506, 508 and 510 illustrate an application of network condition-based delivery rules with respect to pre-provisioning. In a still further example process, blocks 502-

-20-

510 may be combined in a number of ways where both asset-based delivery rules as well as network condition-based delivery rules may be applied for purposes of pre-provisioning policy modulation.

At block 502, historical delivery patterns and trends of one or more ABR assets 5 may be obtained for one or more content delivery nodes of a CDN, for example, over a period of time such as a configurable moving/sliding time window. For each content delivery node, appropriate statistical distributions of the ABR assets may be determined, e.g., percentage distributions, likelihood estimates, multi-level ranking distributions, etc., over the relevant sample time base. Based on the ABR distributions, 10 one or more delivery rules may be established or otherwise determined, which may be provided as part of an adaptive learning process, e.g., a knowledge-based or rules-based expert system (block 504). When a particular content is determined or otherwise estimated to become popular at different nodes in the CDN, only certain representations or adaptation sets of the particular content's video/audio/subtitle 15 components that satisfy or otherwise meet the delivery rules may be pre-provisioned for delivery with respect to the nodes where the content has been determined to be popular (block 510).

Alternatively or additionally, historical network conditions (e.g., download rates, bandwidth usage, etc.) associated with respective content delivery nodes of the 20 CDN may be monitored (block 506), which may then be used for determining suitable delivery rules. Again, such delivery rules may be static/deterministic or varied responsive to pre-configured control inputs from CDN management nodes (i.e., self-adaptive learning). Accordingly, the network condition-based rules may be established in a number of ways, depending on a specific implementation (block 508). As with the 25 asset trend-based delivery rules, the network condition-based delivery rules may be used for filtering only certain adaptation sets of popular content for pre-provisioning (block 510).

Additional steps or acts relative to block 510 are illustrated as a separate flowchart 500B in FIG. 5B. At block 512, an application of the delivery rules may 30 comprise comparisons of ABR assets against a range of video bitrates, thresholds, language/subtitle likelihood estimates, network connection performance/conditions, etc., as well as other metrics relating thereto. Responsive to the comparisons,

video/audio adaptation sets and subtitles of a particular content may be selected or otherwise identified for pre-provisioning as set forth at block 514.

FIG. 6 depicts a block diagram of a computer-implemented data processing system adapted to operate as a management node 600 according to an embodiment of the present patent application. Taking reference to FIGS. 4 and 6, the management node 600 may be configured as an example implementation of DNMC system 112 that may inter-operate with one or more CDN management systems such as a bandwidth allocation/scheduling system, a distribution pipe policy management and optimization system, a content policy and QoS management system, or a request rerouting/redirecting system, or any combination thereof, in addition to the popularity policy management system 402 associated with CDN 400 as well as the CDN nodes. One or more processors or processor boards 602-1 to 602-N coupled to a bus structure 464 provide for the overall control of the node 600, including executing one or more functionalities set forth in the present patent disclosure. Bus 604 may be coupled to one or more bus bridges or bus controllers 606-1, 606-2 for extending to or coupling to additional or peripheral buses 608, 610, which in turn may support a plurality of interfaces (I/F). By way of example, one or more of such interfaces 612-1, 612-2 may effectuate interfacing with various CDN nodes such as the national server node(s) 102, regional server nodes 104-1 to 104-3, and edge server nodes, 106-1 to 106-N; 108-1 to 108-M; 110-1 to 110-K. Additional interfaces 614-1, 614-2 may effectuate interfaces to one or more CDN management systems mentioned above.

A memory subsystem 619 provided with the management node 600 may include one or more memory controllers 616 for controlling memory operations relative to a plurality of memory modules 618-1 to 618-M as well as nonvolatile memory modules such as persistent memory module 617. Program instructions or logic, which may be organized and/or executed as one or more software processes, modules, blocks, routines, threads, etc. may be stored in or uploaded/downloaded into the persistent memory 617 for effectuating one or more embodiments of the present disclosure. For example, a delivery node and management control module 621 may be configured to effectuate or otherwise establish delivery rules based on historical ABR asset delivery trends and/or historical network connection conditions, which may be arranged in a node-by-node basis and stored in non-volatile databases 623, 627, respectively. In one

implementation, the delivery rules may be established pursuant to statistical and mathematical modeling techniques involving, e.g., cluster analysis, multivariate analysis, principal component analysis, factor analysis, correlation analysis, adaptive neural networks, regression analysis, etc. to not only learn but also forecast ABR asset 5 distributions on a node-by-node basis. As pointed out previously, the delivery rules may be distributed to appropriate CDN nodes (e.g., the national server 102) for modulating content pre-provisioning policies. Alternatively or additionally, yet another module 625 may be provided for applying the delivery rules in conjunction with content pre-provisioning policies as determined by the popularity policy management 10 system 402, whereby modulated pre-provisioning policies may be provided to the CDN nodes. In a still further variation, the delivery rules may also be applied in conjunction with a request redirector system that is operative to redirect subscribers' requests to appropriate edge server nodes based on popularity, network bandwidth conditions, etc. for purposes of the present patent disclosure.

15 FIG. 7 is a flowchart of an example process 700 with blocks relative to various steps and acts that may take place for monitoring historical delivery trends of ABR assets according to an embodiment of the present patent application. When a client requests to watch ABR content (block 702), appropriate data such as subtitle/language selection, etc. may be saved for the delivery session (blocks 704, 706). Depending on 20 whether the ABR content is in separate adaptation sets or in interleaved content stream (i.e., multiplexed or muxed content stream provided in certain legacy streaming implementations) as set forth in decision block 708, further data relative to the delivery session may be stored, saved or otherwise monitored. For the content in separate adaptation sets, a loop sub-process 710 may be effectuated with respect to each audio 25 and video segment of the ABR asset package, wherein segment video bitrate(s), audio bitrate(s) and delivery format (e.g., audio formats such as MP3, AAC, AAC+, eAAC+, FLAC, WMA, WAV, AMR, OGG, DTS, AC3, LPCM and MIDI as well as video formats such as, e.g., MPEG4, H.263, H.264, DivX, XviD, WMV, AVI, 3GO, Flash 30 Video), may be stored, as set forth at blocks 712, 714. For muxed content, a loop sub-process 716 is operative to store the segment bitrate(s) for each interleaved segment (block 718). Thereafter, when the client delivery session is terminated (block 720, 724), suitable data updating procedures may take place relative to the delivery statistics

of video/audio/subtitle assets (block 728) or the muxed set bitrate(s) and subtitle selection (block 726). In one implementation, the foregoing sub-processes 710, 716 and subsequent updating procedures may take place relative to each content delivery node of a CDN over a respective period of time, which nodes may be configured to 5 provide (block 728) the delivery statistics to a management node such as DNMC system 102 shown in FIG. 1, for example. Further, transmission of delivery statistics between the DNMC system and CDN nodes may be effectuated by means of request/response query mechanisms, push mechanisms, pull mechanisms, in addition being triggered based on timers, user/network settings, and other conditionalities.

10 FIG. 8 is a flowchart an example process 800 with blocks relative to various steps and acts that may take place for modulating pre-provisioning policies of content based on delivery rules according to an embodiment of the present patent application. When pre-provisioning of content in a CDN is commenced (block 802), a loop sub-process 805 may be effectuated with respect to each content delivery node, which is initiated at a decision block 803 by determining whether a particular content is 15 estimated or other determined to be popular in a service location of the delivery node. If the content is not determined to be popular in a delivery node's area, that delivery node may be skipped (block 804) and the content's popularity relative to a next delivery node may be analyzed. If the content is determined to be popular with respect to a specific content delivery node, appropriate historical video/audio/subtitle delivery statistics and associated asset delivery rules for that node may be obtained and/or utilized (block 806). The particular popular content's manifest may be analyzed relative to the ABR package assets available, including the corresponding adaptation sets if applicable (block 808). For each asset in the ABR package, a loop sub-process 20 810 is effectuated that commences with a determination as to whether the ABR asset (and its adaptation set(s)) meets the corresponding delivery rule(s) based on the distribution statistics (decision block 814). If so, that ABR asset (and its adaptation set(s)) may be identified for pre-provisioning (block 816) with respect to the specific content delivery node identified in decision block 803. Otherwise, the ABR asset is 25 skipped and the next asset in the ABR package of the particular popular content is analyzed (block 818). After applying the delivery rules to all the ABR assets relative to each delivery node where the particular content has been determined to be popular,

appropriate ABR assets (and corresponding adaptation set(s)) are pre-provisioned for the “popular” nodes (block 820), whereupon the process is completed (block 822). As illustrated in FIGS. 1 and 4, not all popular nodes may be pre-provisioned with the same ABR assets/adaptation sets because of the varying historical delivery trends, 5 network conditions, and the like.

FIG. 9 is a flowchart of an example process 900 with blocks relative to various steps and acts that may take place for updating historical delivery trends and/or network conditions data for one or more content delivery nodes of a CDN according to an embodiment of the present patent application. Pursuant to invoking an updating 10 routine (block 902), a timer-related determination may be made as to whether it is appropriate to update trend/conditions statistics in the CDN (decision block 904). If so, a complete list of content delivery nodes may be obtained (e.g., by the DNMC system) from a national origin server of the CDN (block 908). For each delivery node of the 15 list, a loop sub-process 910 may be effectuated for obtaining video/audio/subtitle delivery statistics and updating a database record or storage associated with the node based thereon (block 912). Additionally or alternatively, appropriate network conditions may also be monitored on a node-by-node basis. If the timer-related (or, other trigger-based) determination indicates that it is not appropriate to update the 20 statistics yet, the updating process 900 may wait for a pre-configured amount of time to elapse (block 906). It should be appreciated that although a single timing determination is illustrated (i.e., decision block 904) with respect updating the statistics for all CDN nodes, multiple timing determinations and/or other trigger-based thresholds may also be implemented for initiating separate updating processes 25 corresponding to the respective CDN nodes.

25 In addition to facilitating modulation of content pre-provisioning policies in a CDN, the historical delivery trends and network conditions described hereinabove may also be utilized for optimizing defragmentation of content in the CDN in a further aspect. As pointed out previously, when clients access ABR content via a CDN, it is possible that servers at the edge may become fragmented because of the inherent 30 features of adaptive streaming architecture. Typically, edge server nodes of the CDN may serve one or many individual segment files of a particular ABR asset package to a client and when the client requests content from a second edge server (e.g., due to

redirection), both the first edge server and the second edge server may end up holding an incomplete collection of the segment files referenced by the associated manifest.

When edge servers do not hold a copy of a requested segment file, they may be required to retrieve or pull the requested segment file from a parent node such as a regional server. While this generally causes the parent node to hold a complete copy of the requested segment files, there is no guarantee that a complete segment collection referenced in a manifest will reside on the parent node for all advertised bitrates. Accordingly, fragmentation may occur when clients request some, but not all, of the segment files for a manifest from the same edge server resulting in an incomplete copy of the entire manifest segment collection. Clients changing bitrates within an adaptive streaming architecture can compound this fragmentation on an edge server even further. As a solution, particular embodiments described in the commonly assigned patent application entitled "DEFragmentATION OF ADAPTIVE STREAMING SEGMENT FILES IN A CONTENT DELIVERY NETWORK" (Ericsson Ref. No.: P40285-US2), Application No.: 13/923,005, filed June 20, 2013, in the name(s) of Christopher Phillips et al., incorporated by reference hereinabove, provide a "defragmentation" system and method within a CDN whereby edge servers are ensured to hold a complete copy of the entire segment collection referenced in a manifest. In a further enhancement, embodiments of the present disclosure described below provide an optimized defragmentation system and method wherein the fragmented content at edge server nodes is replenished based on historical delivery trends and/or network conditions associated with the nodes. Furthermore, the disclosed embodiments may be controlled to occur at times of low network utilization such that conditions of bandwidth congestion may be avoided or otherwise ameliorated.

FIGS. 10A-10C depict flowcharts with blocks relative to various steps and acts that may take place for optimizing defragmentation of media segments in a content delivery node of according to one or more embodiments of the present patent application. Turning first to FIG. 10C, an embodiment of optimized defragmentation process 1000C may take place at a CDN node with respect to a particular content. At block 1080, the CDN node is configured to analyze the content's manifest as well as the ABR segment files that are stored in a database cache of the CDN node in order to determine if any segment files referenced by/in the manifest are absent thereat. If so,

one or more ABR asset delivery rules associated with the CDN node may be applied to determine representations or adaptation sets of the absent/missing segment files that satisfy the delivery rules (block 1082). Responsive thereto, only compliant representations/adaptation sets of the absent segments are pulled or retrieved from 5 another CDN node (e.g., another edge server node or a parent node such as a regional server node), which are then stored locally in the database cache of the CDN node (block 1084).

FIG. 10A depicts a flowchart of an optimized defragmentation process 1000A in further detail, which may be performed by a child node or edge node such as edge 10 server nodes 106-N, 108-1 of CDN 100. Similar to the flowcharts described hereinabove, blocks enclosing a portion of steps of the process indicate a loop sub- process wherein the enclosed operations may be iteratively/repeatedly performed until the list of entities for which the loop is effectuated is exhausted. In one implementation, each node within the CDN may determine at various time intervals 15 whether a current time falls within a period that is characterized by the node's low network utilization (e.g., during late night), which may be configured as the node's Minimum Usage Timeframe parameter. Nodes that do not provide self-calculating timeframes may send a request to a parent node, a redirector server, or another server to determine its Minimum Usage Timeframe. Thus, at block 1002, the edge node checks 20 a defragmentation wakeup timer, which causes the edge node to determine at decision block 1004 whether the current time is within the period of minimum usage of network resources within the CDN. The edge node may self-calculate the period of minimum usage of network resources if able, and if not able, the edge node receives information identifying the period of minimum usage of network resources within the CDN from its 25 parent node, the redirector server, or another server in the CDN. If the current time is not within the period of minimum usage of network resources within the CDN, the process moves to block 1008 where defragmentation is skipped for that particular edge node, with control returning to block 1002. When it is determined at block 1004 that the current time is within the period of minimum usage of network resources within the 30 CDN, the process moves to block 1006 where the edge node retrieves all content popularity timeframe windows, e.g., from the redirector server operating in conjunction with a popularity policy management server. Process 1000A thereafter enters a loop

sub-process 1010 that encompasses a collection of actions performed for each adaptive streaming manifest associated with segment files stored in the edge node's local cache memory. At block 1012 of loop 1010, the edge node determines whether the current time is within the popularity time window threshold for a particular manifest/package being analyzed. It should be appreciated that such a determination may be made in order to ensure the complete piece of electronic content is available (and potentially within a licensing window) when clients are likely to request the segment files. If the current time is not within the popularity time window threshold, the process moves to block 1014 where the edge node skips that particular manifest/package. When it is determined at block 1012 that the current time is within the popularity time window threshold for the particular package being analyzed, the process moves to block 1016 where the edge node reads the master manifest file(s) associated therewith. The process thereafter enters another loop sub-process 1018 that encompasses a series of actions performed for each child manifest in the master manifest. At block 1020, the edge node reads the URLs of the segment files from a particular child manifest (such as, e.g., the URLs illustrated in MPD model 300 shown in FIG. 3). A determination is made as to whether the ABR asset segment referenced by the URL satisfies the applicable delivery rules associated with the node (decision block 1022). If not, the URL is skipped and the next URL referenced in the child manifest is analyzed. For the URLs that refer to ABR asset segments complying with applicable delivery rules, yet another loop sub-process 1026 is entered for determining if the respective asset segments exist in the edge node's cache memory (decision block 1028). If so, the edge node skips the URL for that segment file at block 1030 and moves to the next segment file. This sub-loop process continues for each of the segment files referenced by the child manifest, and when a referenced segment file is determined to be absent from the edge node's local cache memory, the process moves to block 1032 where the edge node utilizes the URL of the absent segment file to pull the absent segment file from another CDN node such as the edge node's parent node or another edge node. Subsequently, the edge node caches the segment file pulled from the other CDN node. When all the loops have been performed for each adaptive streaming manifest in cache, for each child manifest in the master manifest, and for each URL in each child manifest, all segment files should exist in the edge node's local cache memory. Upon verification

that all segment files for all manifests exist in the edge node's local cache memory at block 1034, the process moves to block 1036 where the edge node reports to its parent node that defragmentation is complete. It will be apparent to one skilled in the art that the sequence of the loops and sub-loops illustrated in FIG. 10A may be interchanged in 5 some implementations, and some of the operations from one loop may be moved to a sub-loop within.

FIG. 10B depicts a flowchart of another embodiment of an optimized defragmentation process 1000B in further detail, which may be performed by a content delivery node of the CDN. Similar to the flowcharts described hereinabove, blocks 10 enclosing a portion of steps of the process indicate a loop sub-process wherein the enclosed actions may be iteratively/repeatedly performed until the list of entities for which the loop is effectuated is exhausted. When asset defragmentation is commenced in the CDN, e.g., responsive to a defragmentation timer mechanism or some other trigger-based mechanism (block 1052), a loop sub-process 1054 may be effectuated 15 with respect to each ABR package within the content delivery node, which is initiated at a decision block 1056 by determining whether a particular ABR package is within a popularity time window or threshold. If not, the particular ABR package is skipped and a next ABR package is analyzed (block 1058). If the ABR package is within the popularity time window, the manifest of the ABR package is analyzed (block 1060), 20 whereupon a loop sub-process 1062 is entered that is performed for each asset component of the ABR package. A determination is made as to whether the asset contains fragmented segments (block 1064). If so, a further determination is made as to whether the fragmented segments are within the historical audio/video/subtitle delivery trend rules and thresholds (decision block 1068). If the missing audio, video, 25 and/or subtitle asset segments are within the historical delivery trends/thresholds, such segments may be pulled from a regional or national server node or another edge server node (block 1072). If the ABR asset does not contain any fragmented segments, the asset is skipped and a next asset of the ABR package is analyzed (block 1066). Further, if the fragmented segments are not compliant with the applied delivery rules, 30 they are skipped from segment retrieval (block 1070). When all the loops have been performed for each delivery node, each ABR package within a delivery node, etc., the optimized defragmentation process 1000B is completed (block 1074).

With respect to the optimized defragmentation embodiments described hereinabove, it should be appreciated that the historical trend-based delivery rules and/or network condition-based delivery rules may be provided by the DNMC system to each CDN node as it commences its defragmentation process relative to the contents 5 cached thereat. In another variation, because each CDN node may also be configured to maintain its own trends (at least for a short term), there would be no need to retrieve the trend data from a management node in such a scenario. In a still further variation, locally maintained short-term based trends as well as long-term trends maintained at the DNMC system may be applied for purposes of determining which of the missing 10 segments should be pulled from another node in the CDN. Additionally, the foregoing optimized defragmentation embodiments may be practiced with respect to parent nodes as well as child nodes of a CDN.

FIG. 11 depicts a block diagram of a computer-implemented data processing system adapted to operate as a content delivery node 1100 that may be configured to 15 effectuate optimized content defragmentation according to an embodiment of the present patent application. The functions of the CDN node 1100 may be controlled, for example, by one or more processors 1102 executing computer program instructions stored on a memory 1110 including a persistent module 1111. A defragmentation wakeup timer 1104 provides the current time to a comparison unit 1112, which 20 compares the current time with a Minimum Usage Timeframe parameter 1106. If the current time is within the Minimum Usage Timeframe, the comparison unit 1112 further compares the current time with the manifest popularity time window/threshold 1108. If the current time is within the manifest popularity time window threshold, the comparison unit 1112 notifies a segment file reader 1114 that it is appropriate to start 25 the defragmentation process. The segment file reader 1114 is operative to read the ABR asset segment files from a local database cache 1118. A URL reader 1116 is operative to read the URLs of each of the segment files referenced by the associated manifest and provide them to an absent file identifier 1122. The absent file identifier 1122 is operative to identify one or more referenced segment files that comply with 30 applicable ABR asset delivery rules 1120 but are absent from the local cache memory 1118. URL information relative to the compliant missing segments files is provided to a communication unit 1124 that is operative to send one or more requests for the

-30-

segment files by sending the requested URL(s) to another CDN node 1126, which may be, for example, a parent node. In response, the other CDN node 1126 returns the requested segment file(s) to the communication unit 1124, which sends them to the cache 1118 for storage. Subsequently, the segment file reader 1114 and file identifier 1122 are operative to determine that all segment files referenced by the associated manifest document now exist in the cache 1118. The file identifier 1122 may thereafter notify the communication unit 1124 that the defragmentation process is complete, which in turn may notify the other CDN node 1126 that the defragmentation process is complete. It should be appreciated that dashed lines in FIG. 11 indicate additional functionality when the CDN node 1100 is also a parent node. As set forth in additional detail in the commonly assigned patent application (Appl. No.: 13/923,005, filed June 20, 2013) incorporated by reference hereinabove, the parent node functionality may involve determining whether all of its child node(s) 1128 have reported that their defragmentation process is complete. Once all the child node(s) have reported that their defragmentation is complete, the comparison unit 1112 may continue by determining whether the current time is within the manifest popularity time window threshold 1108. If so, the defragmentation process for the parent node may be initiated.

FIG. 12 depicts an example content delivery network 1200 wherein fragmented content is replenished at a content delivery node (e.g., edge server nodes 106-N, 108-1) based on historical delivery trends according to an embodiment of the present patent application. For purposes of illustration, CDN 1200 is identical to CDN 100 shown in FIG. 1 except that a CDN redirector node 1202 is shown to interface with DNMC system 112 for receiving suitable delivery rules as discussed hereinabove. Further, the CDN redirector node 1202 may also receive content popularity windows from a content popularity policy management node (not specifically shown) for distributing the data to the various nodes of CDN 1200, as illustrated by example request/response paths 1204-1, 1204-2 and 1204-3 relative to nodes 104-1, 106-N and 108-1, respectively. As described in detail in the foregoing sections with respect to FIGS. 10A-10C and 11, an edge server node is operative to apply the delivery rules in order to identify a subset of representations of the missing segments of a fragmented content file. By way of example, the Galveston edge server 106-N is operative to apply suitable delivery rules based on its ABR asset delivery pattern 116, whereby only certain representations of

the missing segments that satisfy the rules may be retrieved from another CDN node, as shown by a missing segments identifier 1208 that comprises a video asset component 1210-1 (for pulling only 720p and 1080p representations of the content encoded at, for example, 3/5 Mbs and 10/12Mbs, respectively), an audio asset component 1210-2 (for pulling only English 2.0 and Spanish 2.0 representations), and a subtitle asset component 1210-3 (for pulling only English language subtitles of the content).

In similar fashion, one or more delivery rules based on the learned ABR asset delivery pattern 120 associated with the Ashville edge server 108-1 may be applied to optimize the server's missing segment retrieval. Reference numeral 1212 refers to the 10 Ashville server's missing segments identifier that illustrates a video asset component 1214-1 (for pulling only 720p and 480p representations of the content encoded at, for example, 3/5 Mbs and 2/1 Mbs, respectively), an audio asset component 1214-2 (for pulling only English 2.0 representations), and a subtitle asset component 1214-3 (for pulling only English language subtitles of the content).

15 One skilled in the art will recognize upon reference hereto that if an entire file becomes necessary (for example, additional languages are needed), the defragmentation system is aware of which files are already transferred to the edge of the network and only transfers those sections of the file which are additionally needed. Similarly, should an edge node decide to delete files due to memory constraints, the system will 20 be aware of which fragments of the file are no longer needed (for example, due to the language(s) or bitrate(s) being no longer required).

Based upon the foregoing Detailed Description, it should be appreciated that the 25 embodiments of the present disclosure can be advantageously implemented that allow CDNs to function more efficiently, by making available what is needed where it is needed before it is needed. In addition to savings in operator bandwidth costs, the disclosed embodiments also provide higher QoS levels with respect to end user experience.

30 In the above-description of various embodiments of the present disclosure, it is to be understood that the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. Unless otherwise defined, all terms (including technical and scientific terms) used herein have the same meaning as commonly understood by one of ordinary skill in the art to which

this invention belongs. It will be further understood that terms, such as those defined in commonly used dictionaries, should be interpreted as having a meaning that is consistent with their meaning in the context of this specification and the relevant art and may not be interpreted in an idealized or overly formal sense expressly so defined 5 herein.

At least some example embodiments are described herein with reference to block diagrams and/or flowchart illustrations of computer-implemented methods, apparatus (systems and/or devices) and/or computer program products. It is understood that a block of the block diagrams and/or flowchart illustrations, and combinations of 10 blocks in the block diagrams and/or flowchart illustrations, can be implemented by computer program instructions that are performed by one or more computer circuits. Such computer program instructions may be provided to a processor circuit of a general purpose computer circuit, special purpose computer circuit, and/or other programmable data processing circuit to produce a machine, so that the instructions, which execute via 15 the processor of the computer and/or other programmable data processing apparatus, transform and control transistors, values stored in memory locations, and other hardware components within such circuitry to implement the functions/acts specified in the block diagrams and/or flowchart block or blocks, and thereby create means (functionality) and/or structure for implementing the functions/acts specified in the 20 block diagrams and/or flowchart block(s). Additionally, the computer program instructions may also be stored in a tangible computer-readable medium that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable medium produce an article of manufacture including instructions which implement the functions/acts 25 specified in the block diagrams and/or flowchart block or blocks.

As alluded to previously, tangible, non-transitory computer-readable medium may include an electronic, magnetic, optical, electromagnetic, or semiconductor data storage system, apparatus, or device. More specific examples of the computer-readable medium would include the following: a portable computer diskette, a random access 30 memory (RAM) circuit, a read-only memory (ROM) circuit, an erasable programmable read-only memory (EPROM or Flash memory) circuit, a portable compact disc read-only memory (CD-ROM), and a portable digital video disc read-only memory

(DVD/Blu-ray). The computer program instructions may also be loaded onto or otherwise downloaded to a computer and/or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer and/or other programmable apparatus to produce a computer-implemented process such that

5 the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts specified in the block diagrams and/or flowchart block or blocks. Accordingly, embodiments of the present invention may be embodied in hardware and/or in software (including firmware, resident software, micro-code, etc.) that runs on a processor such as a digital signal processor,

10 which may collectively be referred to as "circuitry," "a module" or variants thereof.

Further, in at least some additional or alternative implementations, the functions/acts described in the blocks may occur out of the order shown in the flowcharts. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse

15 order, depending upon the functionality/acts involved. Moreover, the functionality of a given block of the flowcharts and/or block diagrams may be separated into multiple blocks and/or the functionality of two or more blocks of the flowcharts and/or block diagrams may be at least partially integrated. Moreover, the acts, steps, functions, components or blocks illustrated in a particular flowchart may be inter-mixed or

20 otherwise inter-arranged with the acts, steps, functions, components or blocks illustrated in another flowchart in order to effectuate additional variations, modifications and configurations with respect to one or more implementations of pre-provisioning policy modulation and/or optimized content defragmentation techniques, either in combination or separately, for purposes of the present patent disclosure.

25 Finally, other blocks may be added/inserted between the blocks that are illustrated. Moreover, although some of the diagrams include arrows on communication paths to show a primary direction of communication, it is to be understood that communication may occur in the opposite direction relative to the depicted arrows.

Although various embodiments have been shown and described in detail, the

30 claims are not limited to any particular embodiment or example. None of the above Detailed Description should be read as implying that any particular component, element, step, act, or function is essential such that it must be included in the scope of

-34-

the claims. Reference to an element in the singular is not intended to mean "one and only one" unless explicitly so stated, but rather "one or more." All structural and functional equivalents to the elements of the above-described embodiments that are known to those of ordinary skill in the art are expressly incorporated herein by reference and are intended to be encompassed by the present claims. Accordingly, those skilled in the art will recognize that the exemplary embodiments described herein can be practiced with various modifications and alterations within the spirit and scope of the claims appended below.

WHAT IS CLAIMED IS:

1. A method operative at a management node (112) associated with a content delivery network (CDN) (100, 400) configured to deliver adaptive streaming content, the method comprising:
 - 5 obtaining (502) historical delivery patterns of one or more adaptive bitrate (ABR) assets associated with one or more content delivery nodes of the CDN (100, 400);
 - 10 determining (504), for each content delivery node, one or more delivery rules based on statistical distributions of the ABR assets delivered over a period of time; and when a particular content is determined to become popular at a specific content delivery node (106-N), pre-provisioning (510) only certain representations of the particular content for delivery at the specific content delivery node (106-N) based on the one or more delivery rules established for the specific content delivery node (106-N).
2. The method as recited in claim 1, wherein the ABR assets comprise at least one of video components, audio track components and one or more subtitle languages.
- 20 3. The method as recited in claim 2, wherein the historical delivery patterns are obtained for all the video components, all the audio track components, and all subtitle languages downloaded over a period of time at respective content delivery nodes of the CDN (100, 400).
- 25 4. The method as recited in claim 3, wherein the one or more delivery rules for a respective content delivery node (106-N) comprise at least one of a video delivery rule based on bit rates of video components downloaded from the respective content delivery node (106-N), an audio delivery rule based on audio language tracks downloaded from the respective content delivery node (106-N), and a subtitle delivery rule based on subtitle language selections downloaded from the respective content delivery node (106-N).
- 30

5. The method as recited in claim 1, wherein the particular content is estimated to become popular at a specific content delivery node (106-N) based on at least one of revenue data, search trend data, subscriber demographic data, and social media data obtained for the particular content.

5

6. The method as recited in claim 1, wherein the content delivery nodes comprise one or more edge server nodes (106-1 to 106-N, 108-1 to 108-M, 110-1 to 110-K) and one or more regional server nodes (104-1 to 104-3) coupled to a national server node (102).

10

7. A management node (600) operative with a content delivery network (CDN) (100, 400) configured to deliver adaptive streaming content, the management node (600) comprising:

one or more processors (602-1 to 602-N);

15

a database (623) for storing historical delivery patterns of one or more adaptive bitrate (ABR) assets associated with one or more content delivery nodes of the CDN (100, 400);

20

a delivery node management and control (DNMC) module (621) embodied in a persistent memory (617) coupled to the one or more processors (602-1 to 602-N), the DNMC module (621) having instructions executable by the one or more processors (602-1 to 602-N) and configured to determine (504), for each content delivery node, one or more delivery rules based on statistical distributions of the ABR assets delivered over a period of time; and

25

a content pre-provisioning control module (625) embodied in a persistent memory (617) coupled to the one or more processors (602-1 to 602-N), the content pre-provisioning control module (625) having instructions executable by the one or more processors (602-1 to 602-N) and configured to:

30

when a particular content is determined to become popular at a specific content delivery node (106-N), apply (512) the one or more delivery rules corresponding to the specific content delivery node (106-N) against the particular content; and

select (514) only certain representations of the particular content for pre-populating delivery at the specific content delivery node (106-N) based on the application of the delivery rules established for the specific content delivery node (106-N).

5

8. The management node (600) as recited in claim 7, wherein the ABR assets comprise at least one of video components, audio track components and one or more subtitle languages.

10

9. The management node (600) as recited in claim 8, wherein the historical delivery patterns are obtained and stored in the database (623) for all the video components, all the audio track components, and all subtitle languages downloaded over a period of time at respective content delivery nodes of the CDN (100, 400).

15

10. The management node (600) as recited in claim 9, wherein the DNMC module (621) is operative to determine the one or more delivery rules for a respective content delivery node (106-N) comprising at least one of a video delivery rule based on bit rates of video components downloaded from the respective content delivery node (106-N), an audio delivery rule based on audio language tracks downloaded from the respective content delivery node (106-N), and a subtitle delivery rule based on subtitle language selections downloaded from the respective content delivery node (106-N).

20

11. The management node (600) as recited in claim 7, wherein the particular content is estimated to become popular at a specific content delivery node (106-N) based on at least one of revenue data, search trend data, subscriber demographic data, and social media data obtained for the particular content.

25

12. A method operative at a management node (112) associated with a content delivery network (CDN) (100, 400) configured to deliver adaptive streaming content, the method comprising:

monitoring (506) historical network conditions associated with one or more content delivery nodes of the CDN (100, 400) for delivery of one or more adaptive bitrate (ABR) assets;

5 determining (508), for each content delivery node, one or more delivery rules based on the historical network conditions monitored for the ABR assets delivered over a period of time; and

10 when a particular content is determined to become popular at a specific content delivery node (106-N), pre-provisioning (510) only certain representations of the particular content for delivery at the specific content delivery node (106-N) based on the one or more delivery rules established for the specific content delivery node (106-N).

13. The method as recited in claim 12, wherein the ABR assets comprise at least one of video components, audio track components and one or more subtitle languages.

15

14. The method as recited in claim 12, wherein the historical network conditions comprise at least one of content download rates sustained by respective content delivery nodes over a period of time, temporal performance variations of respective content delivery nodes compared against a threshold level, and relative performance of respective content delivery nodes in relation to an average performance level of a sample of the content delivery nodes monitored over a period of time.

15. The method as recited in claim 14, wherein the one or more delivery rules for a respective content delivery node (106-N) comprise allowing for delivery of high bitrate content only if the respective content delivery node (106-N) has sustained a high content download rate over a period of time.

16. The method as recited in claim 14, wherein the one or more delivery rules for a respective content delivery node (106-N) comprise allowing for delivery of high bitrate content only if the respective content delivery node (106-N) has sustained a relative performance that is higher than the average performance level of a sample of the content delivery nodes monitored over a period of time.

17. The method as recited in claim 12, wherein the content delivery nodes comprise one or more edge server nodes (106-1 to 106-N, 108-1 to 108-M, 110-1 to 110-K) and one or more regional server nodes (104-1 to 104-3) coupled to a national server node 5 (102).

18. A management node (600) operative with a content delivery network (CDN) (100, 400) configured to deliver adaptive streaming content, the management node comprising:

10 one or more processors (602-1 to 602-N);

15 a database (627) for storing historical network conditions associated with one or more content delivery nodes of the CDN (100, 400) for delivery of one or more adaptive bitrate (ABR) assets;

20 a delivery node management and control (DNMC) module (621) embodied in a persistent memory (617) coupled to the one or more processors (602-1 to 602-N), the DNMC module (621) having instructions executable by the one or more processors (602-1 to 602-N) and configured to determine (506), for each content delivery node, one or more delivery rules based on the historical network conditions monitored for the ABR assets delivered over a period of time; and

25 a content pre-provisioning control module (625) embodied in a persistent memory (617) coupled to the one or more processors (602-1 to 602-N), the content pre-provisioning control module (625) having instructions executable by the one or more processors (602-1 to 602-N) and configured to:

30 when a particular content is determined to become popular at a specific content delivery node (106-N), apply (512) the one or more delivery rules corresponding to the specific content delivery node (106-N) against the particular content; and

35 select (514) only certain representations of the particular content for pre-populating delivery at the specific content delivery node (106-N) based on the application of the one or more delivery rules established for the specific content delivery node (106-N).

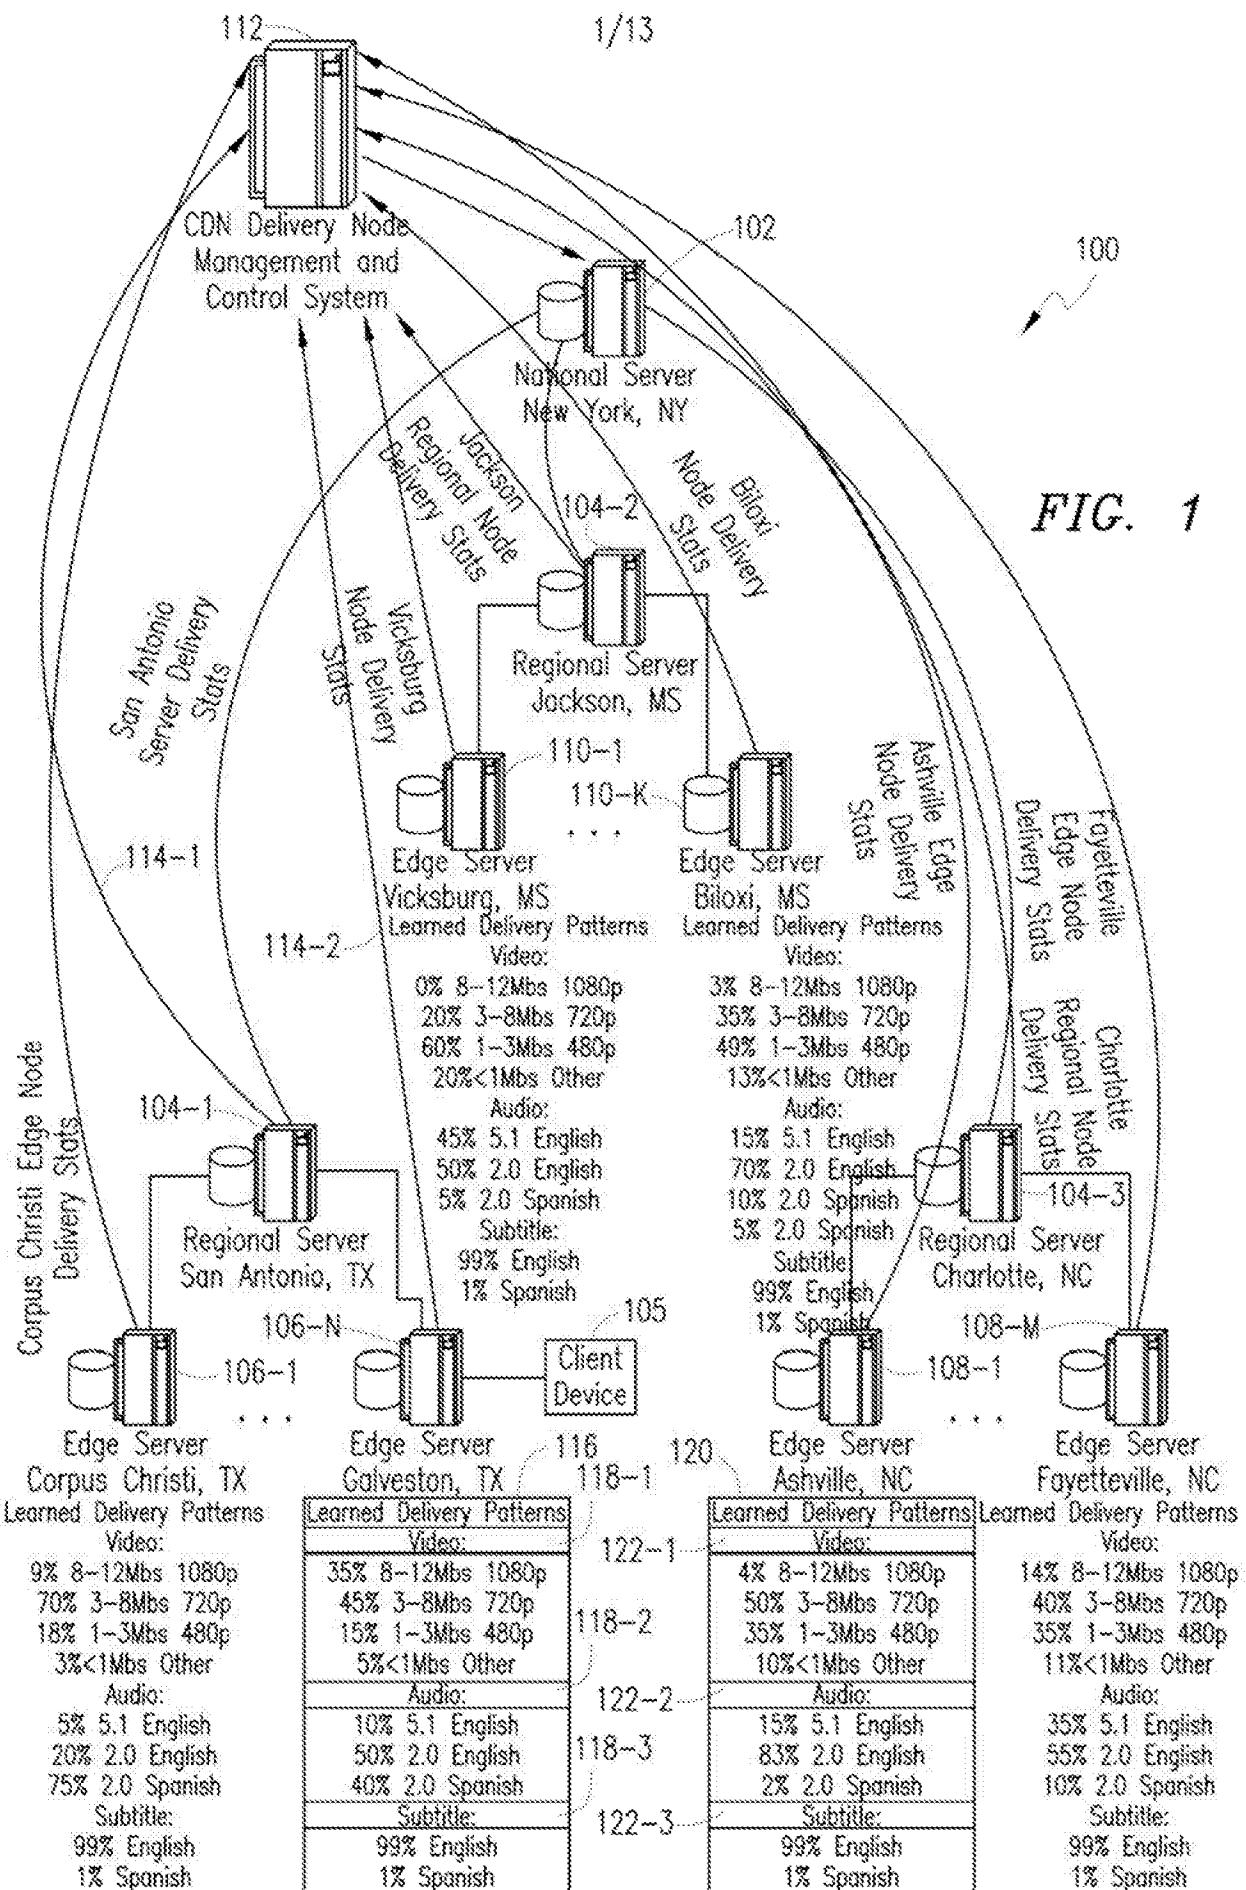
-40-

19. The management node (600) as recited in claim 18, wherein the ABR assets comprise at least one of video components, audio track components and one or more subtitle languages.

5 20. The management node (600) as recited in claim 12, wherein the historical network conditions comprise at least one of content download rates sustained by respective content delivery nodes over a period of time, temporal performance variations of respective content delivery nodes compared against a threshold level, and relative performance of respective content delivery nodes in relation to an average
10 performance level of a sample of the content delivery nodes monitored over a period of time.

21. The management node (600) as recited in claim 14, wherein the DNMC module (621) is operative to determine the one or more delivery rules for a respective content delivery node (106-N) that allow for delivery of high bitrate content only if the
15 respective content delivery node (106-N) has sustained a high content download rate over a period of time.

22. The management node (600) as recited in claim 14, wherein the DNMC module (621) is operative to determine the one or more delivery rules for a respective content delivery node (106-N) that allow for delivery of high bitrate content only if the respective content delivery node (106-N) has sustained a relative performance that is higher than the average performance level of a sample of the content delivery nodes monitored over a period of time.



2/13

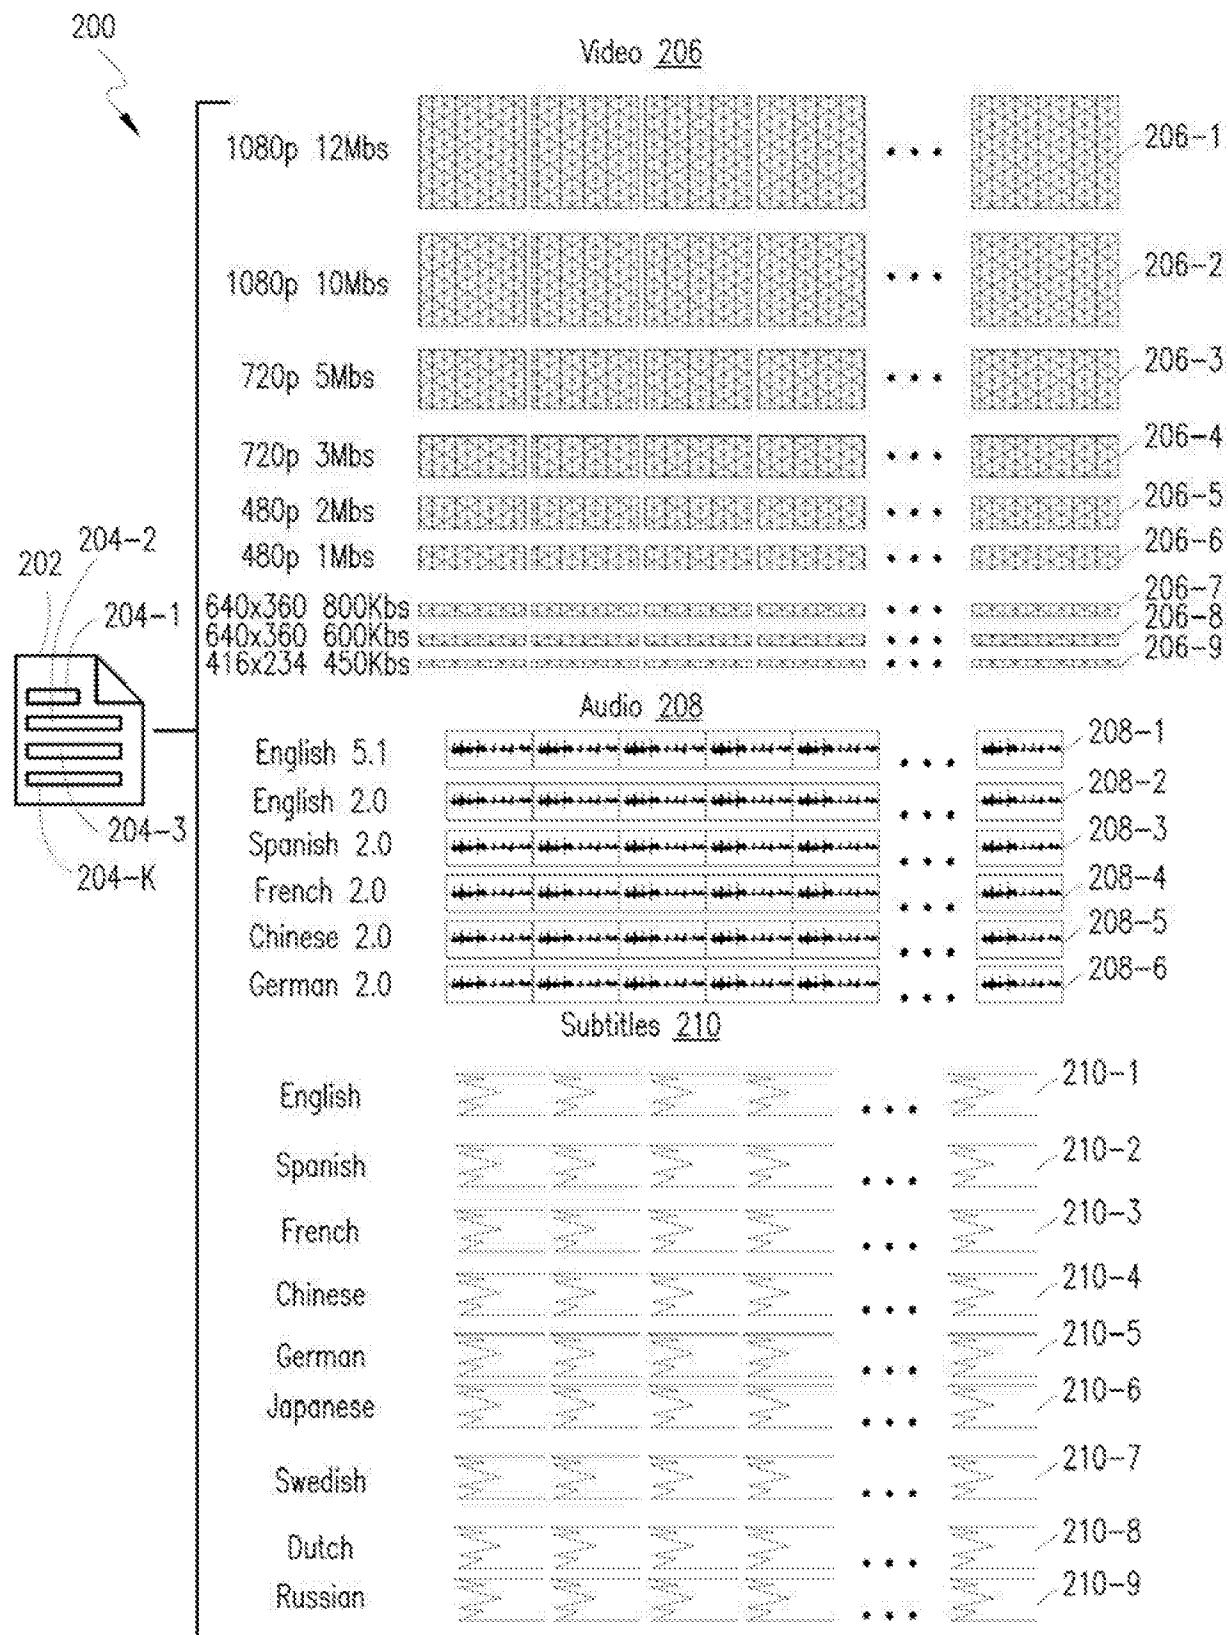


FIG. 2

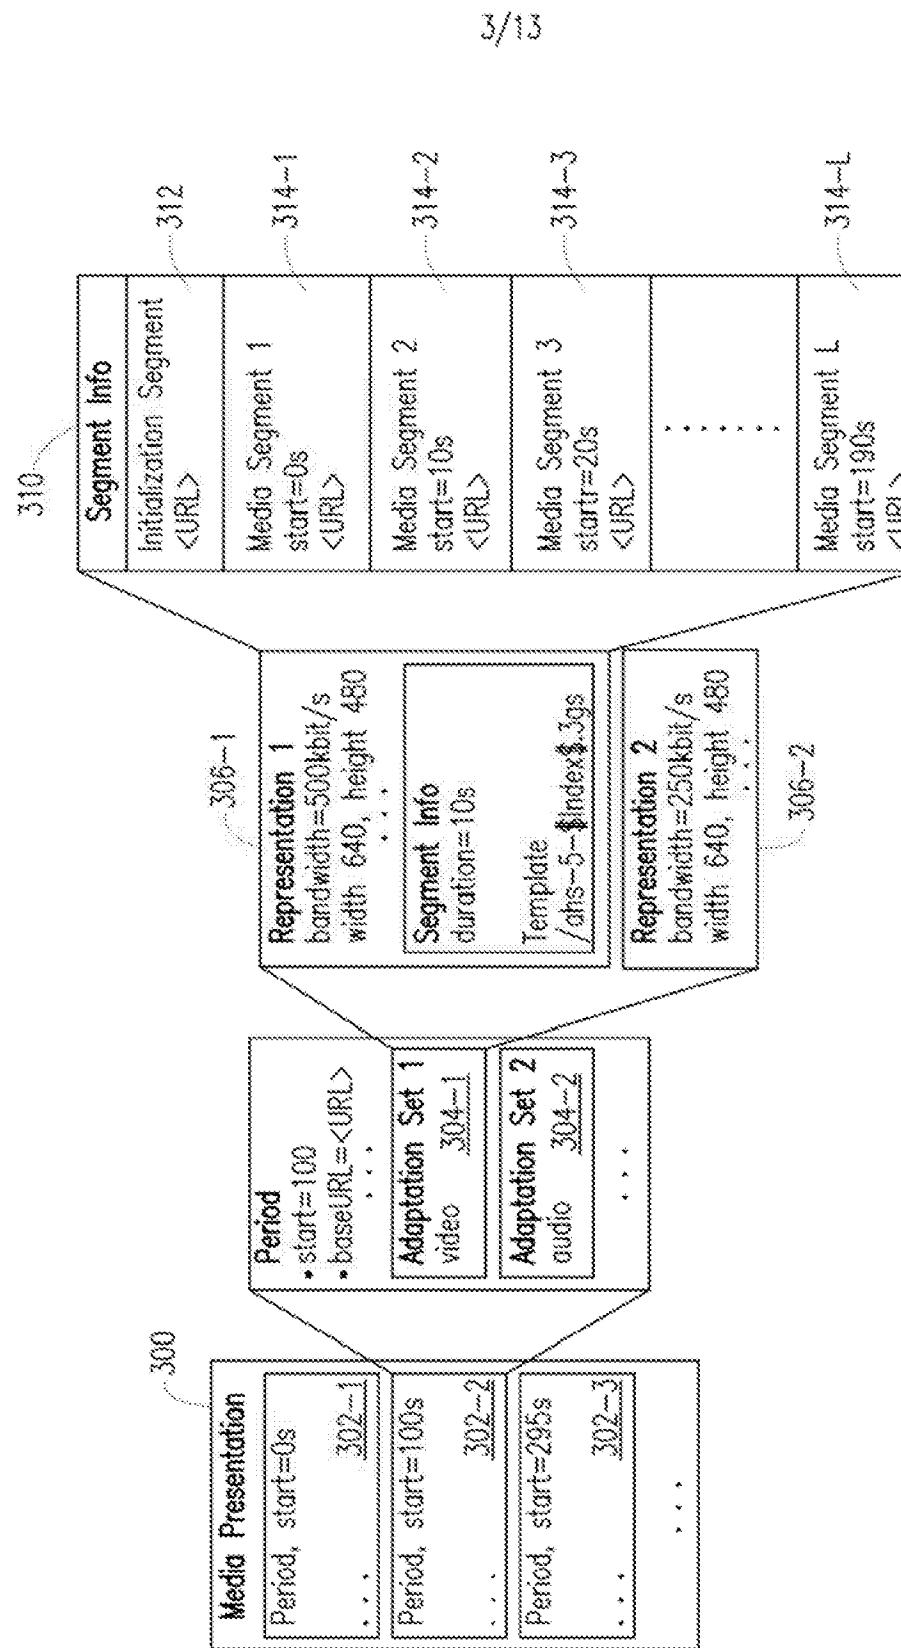


FIG. 3

4/3

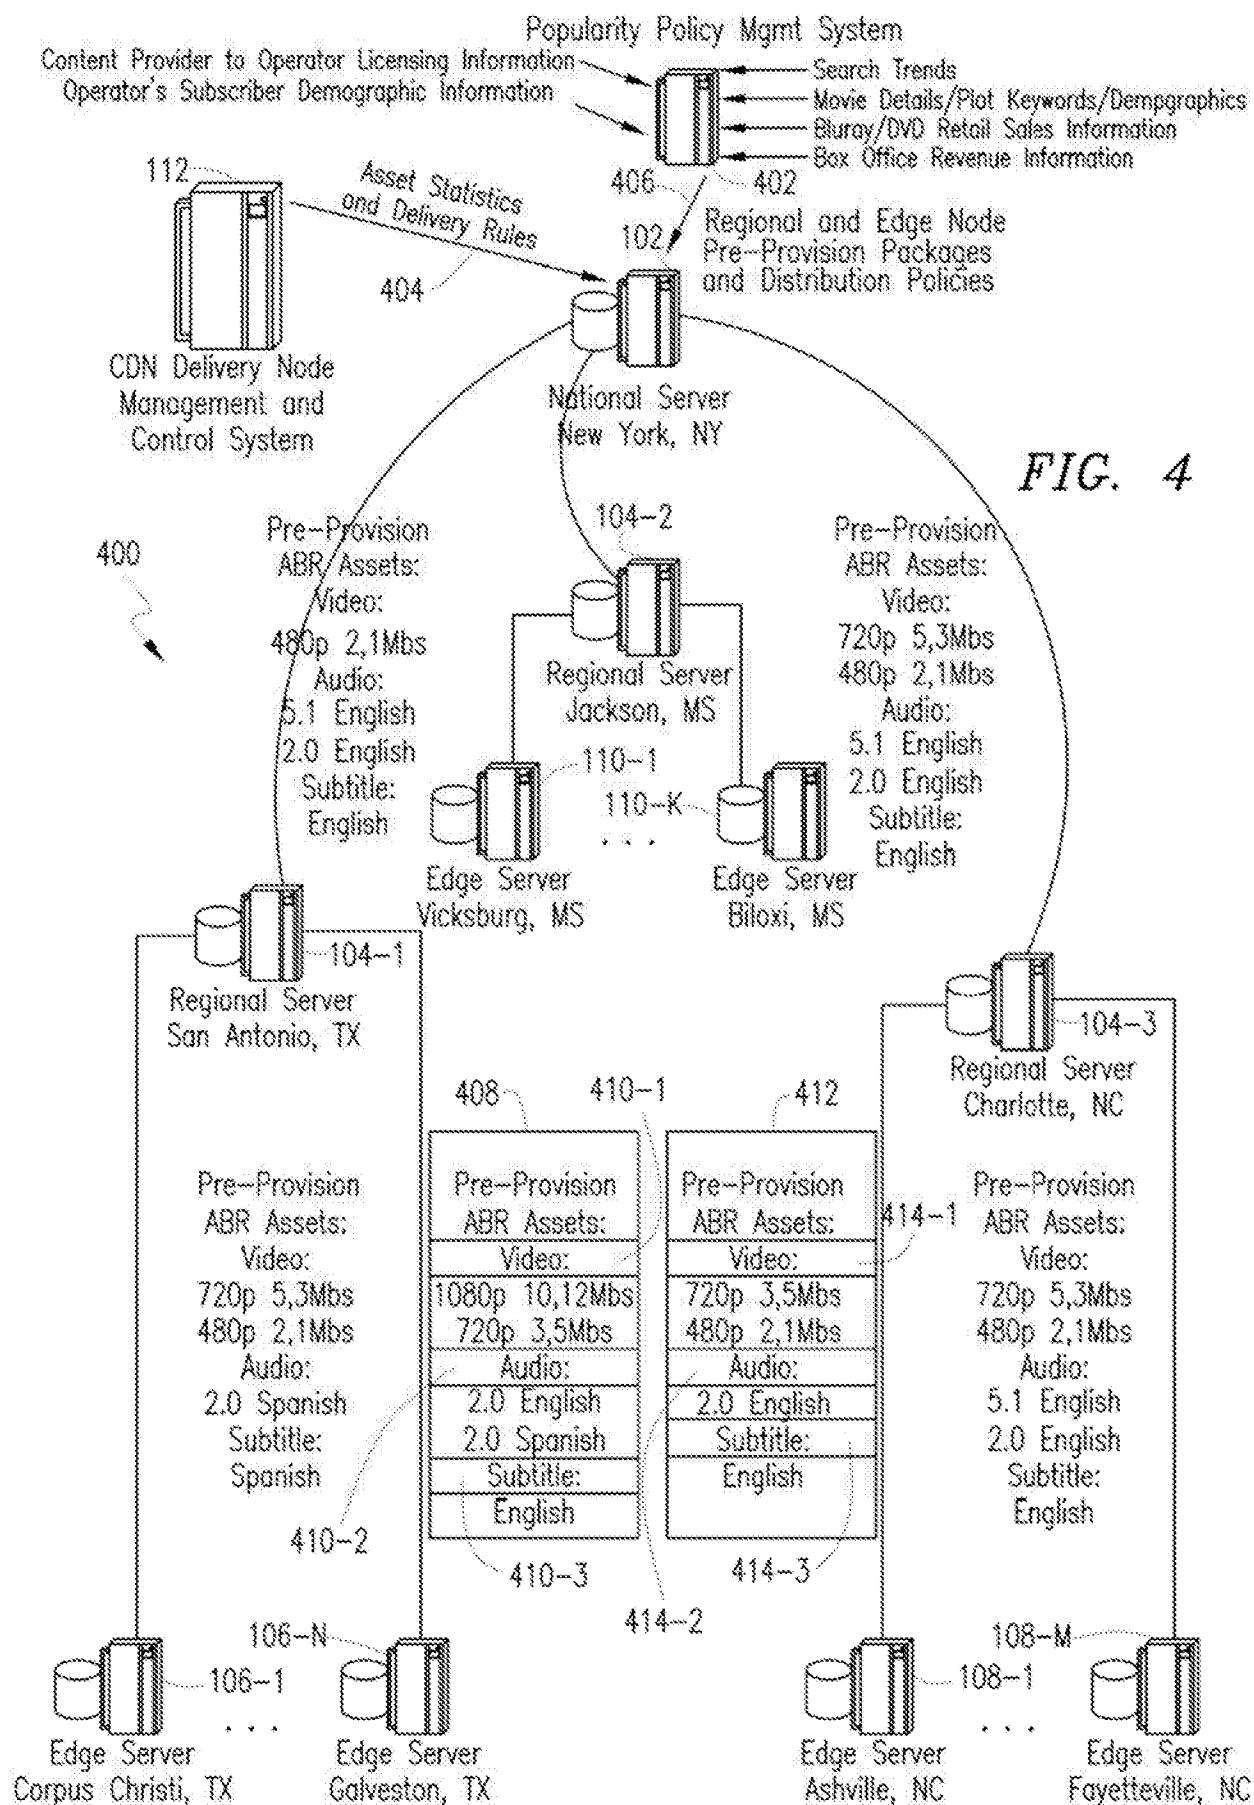
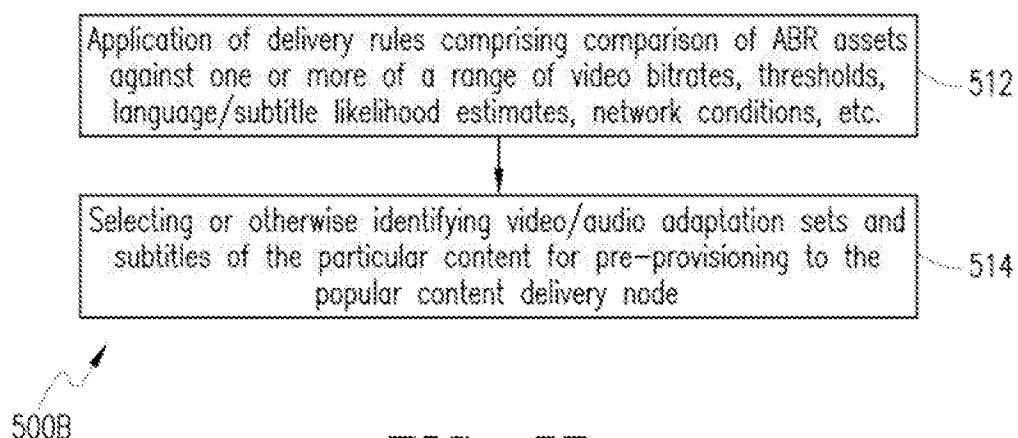
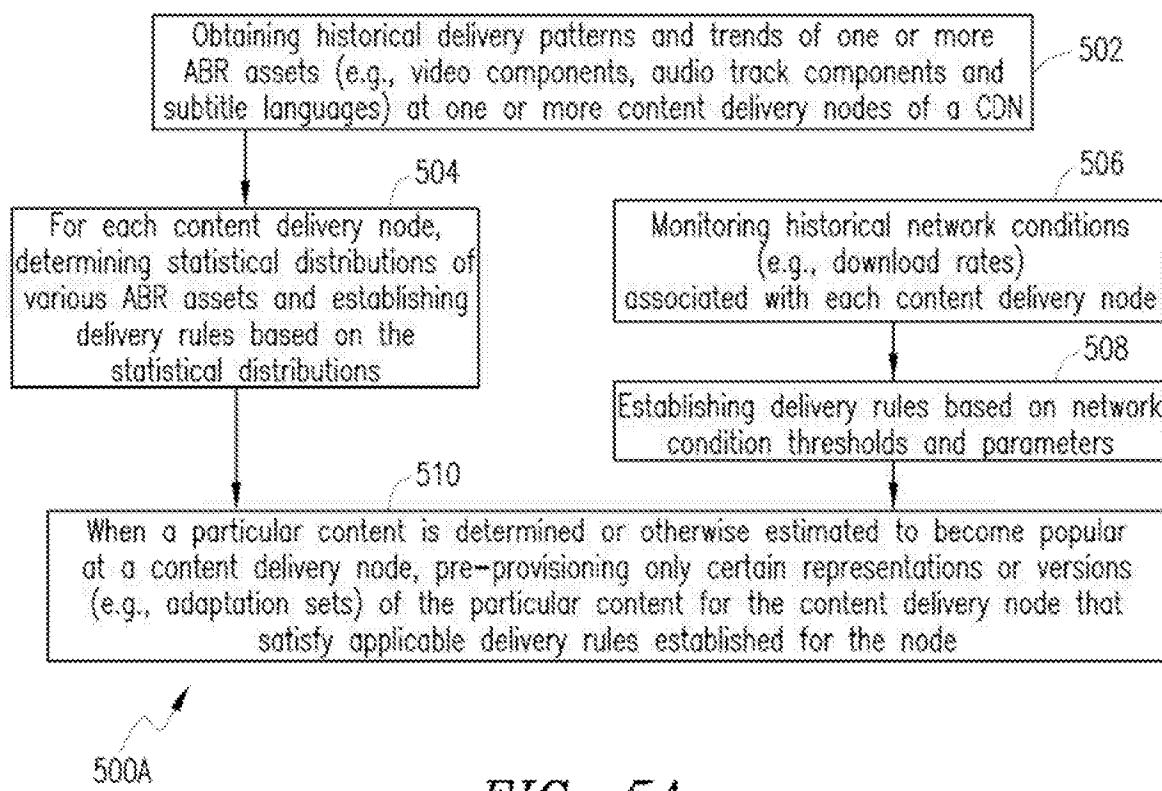


FIG. 4



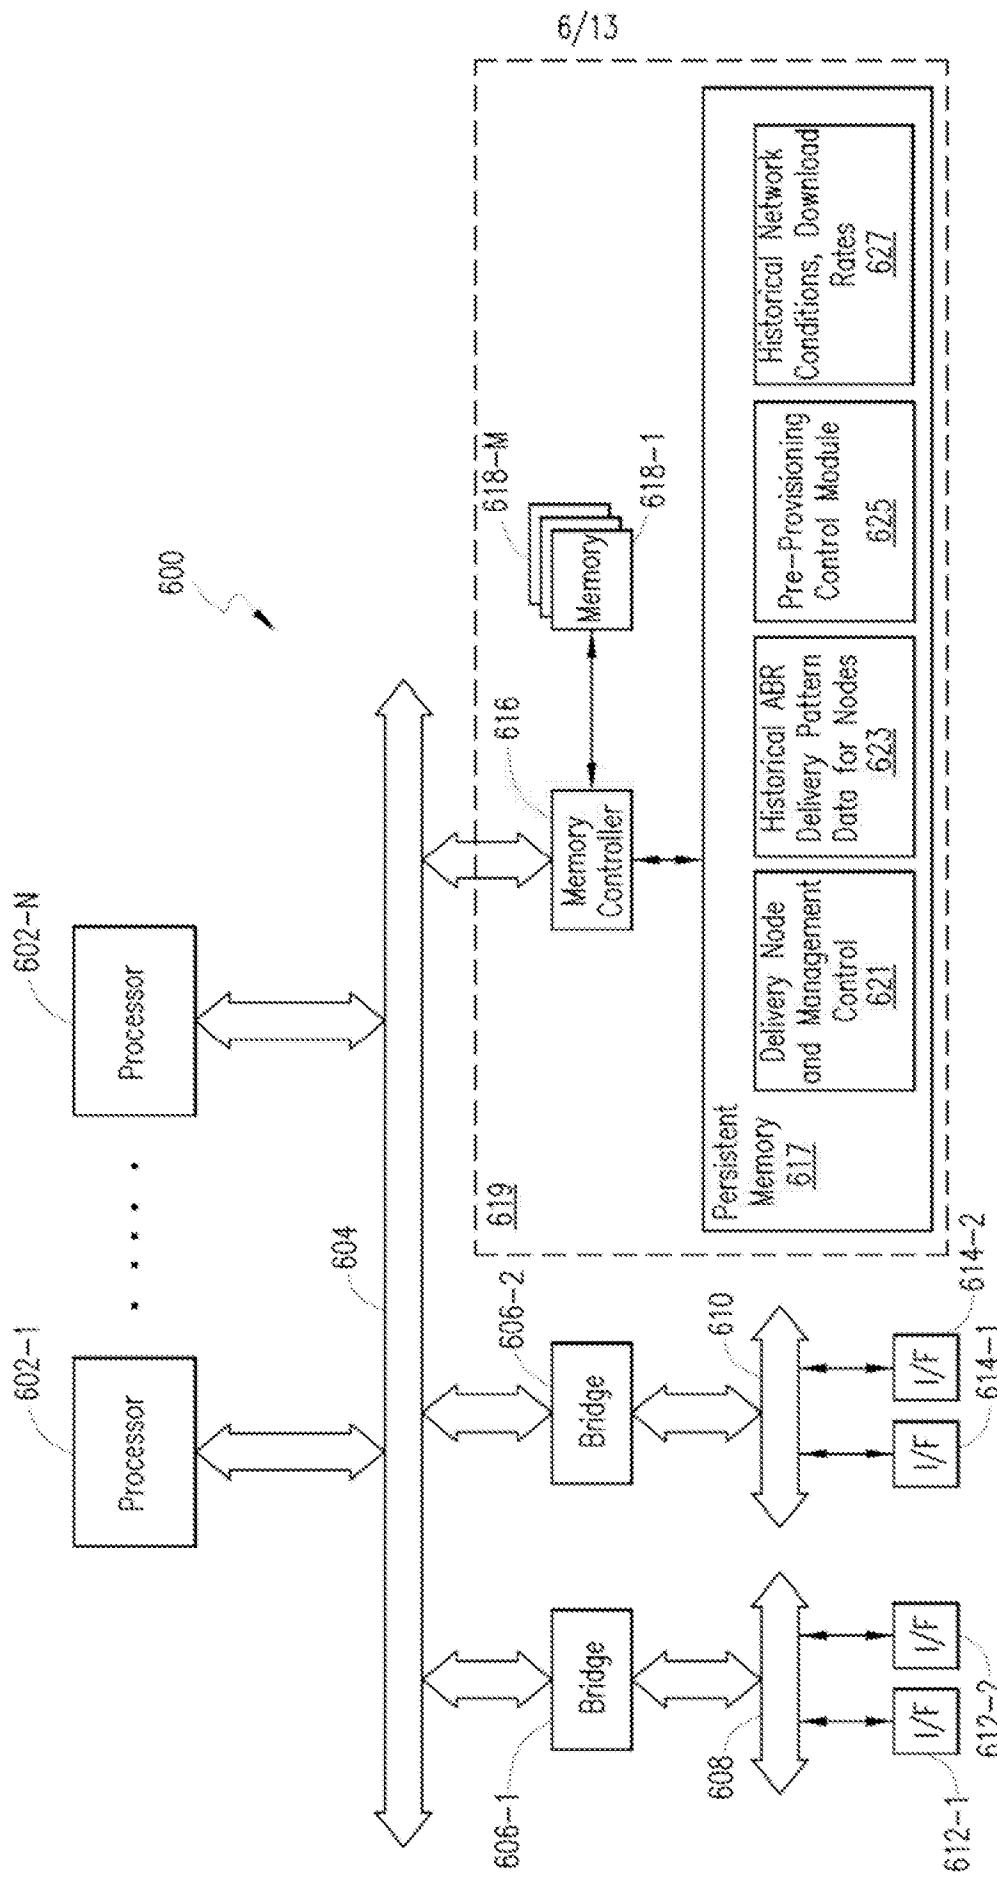


FIG. 6

7/13

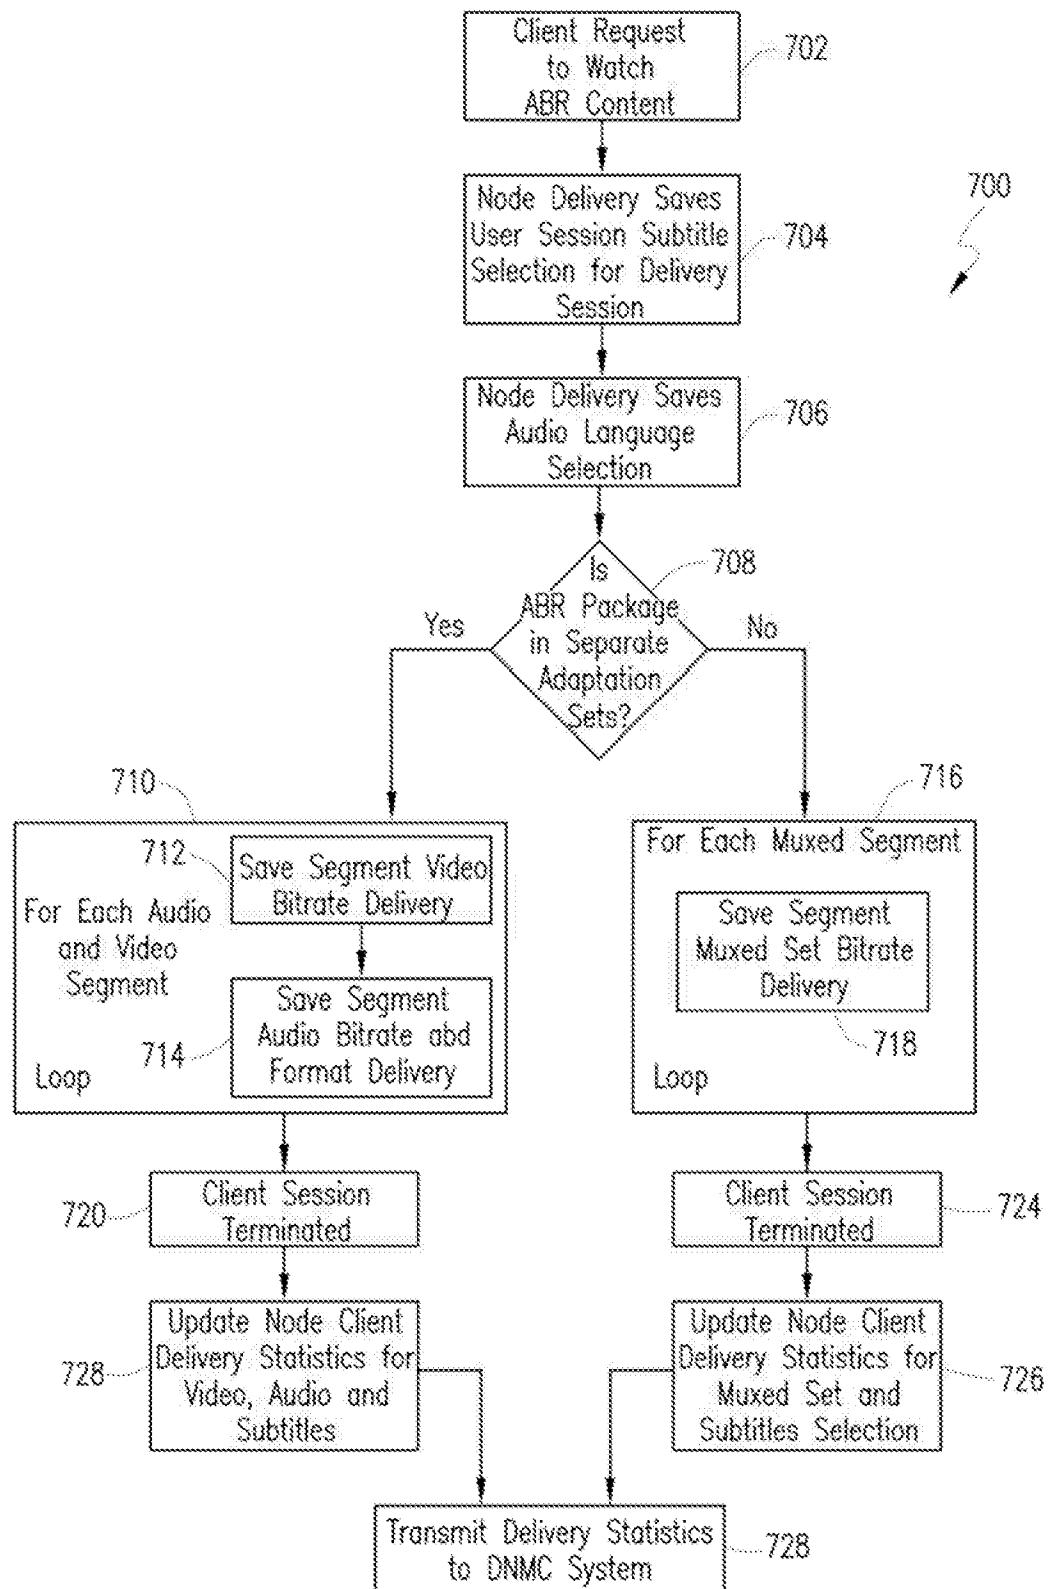


FIG. 7

8/13

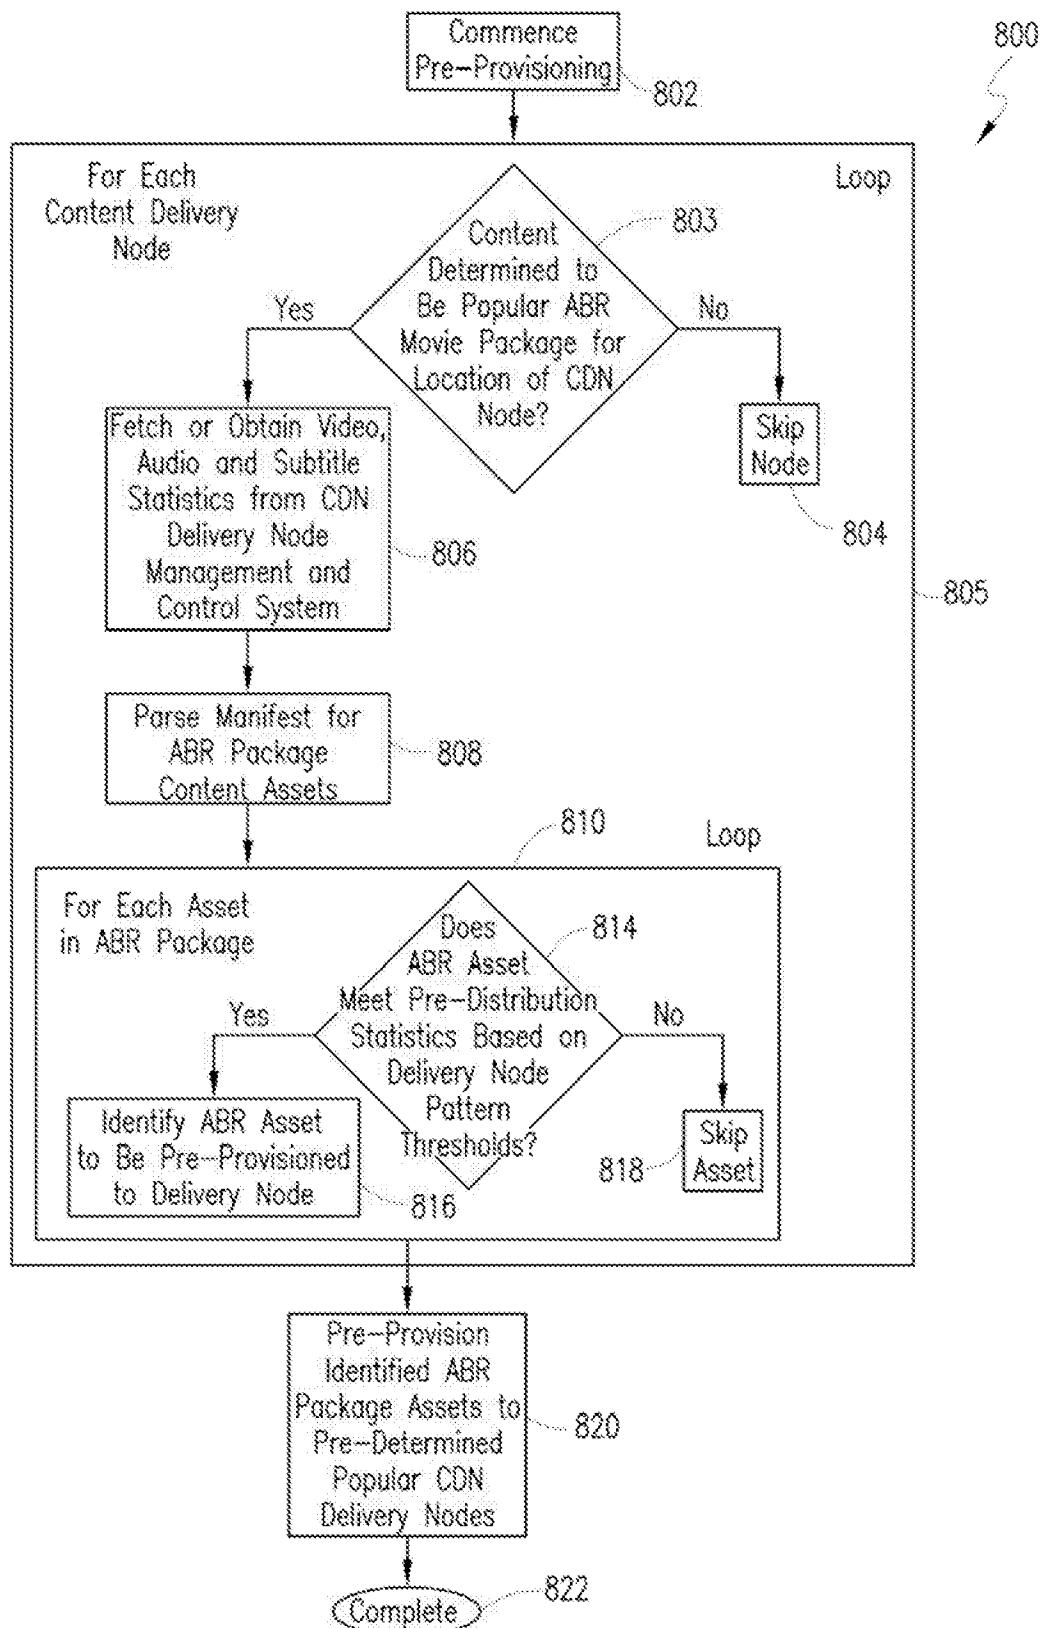


FIG. 8

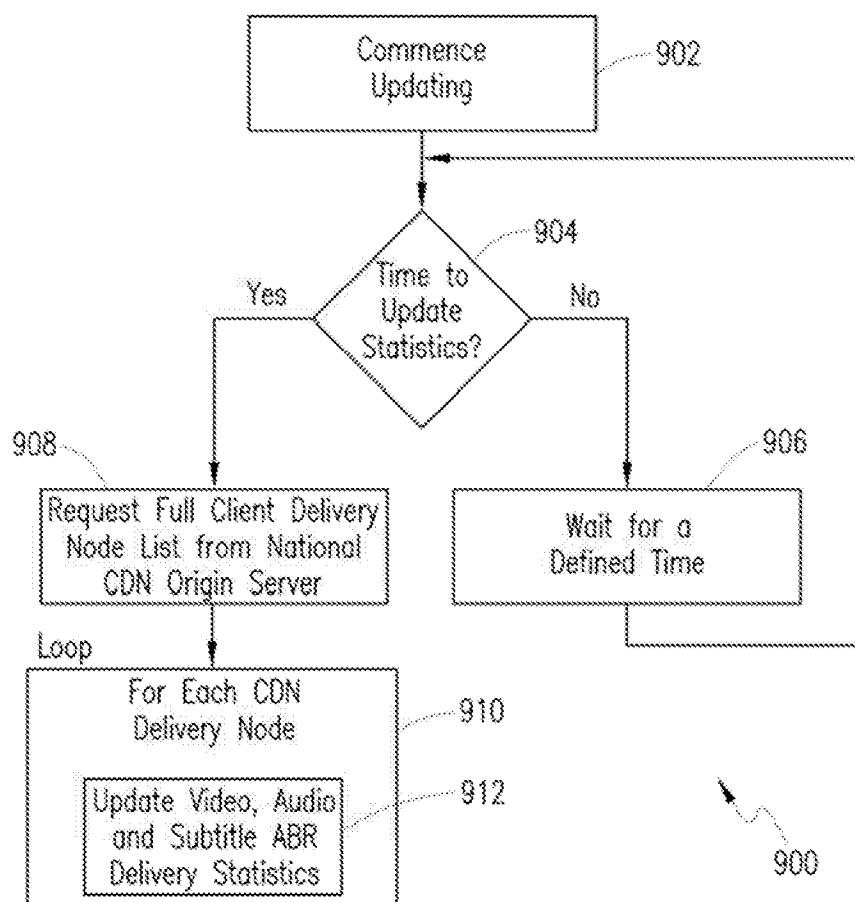


FIG. 9

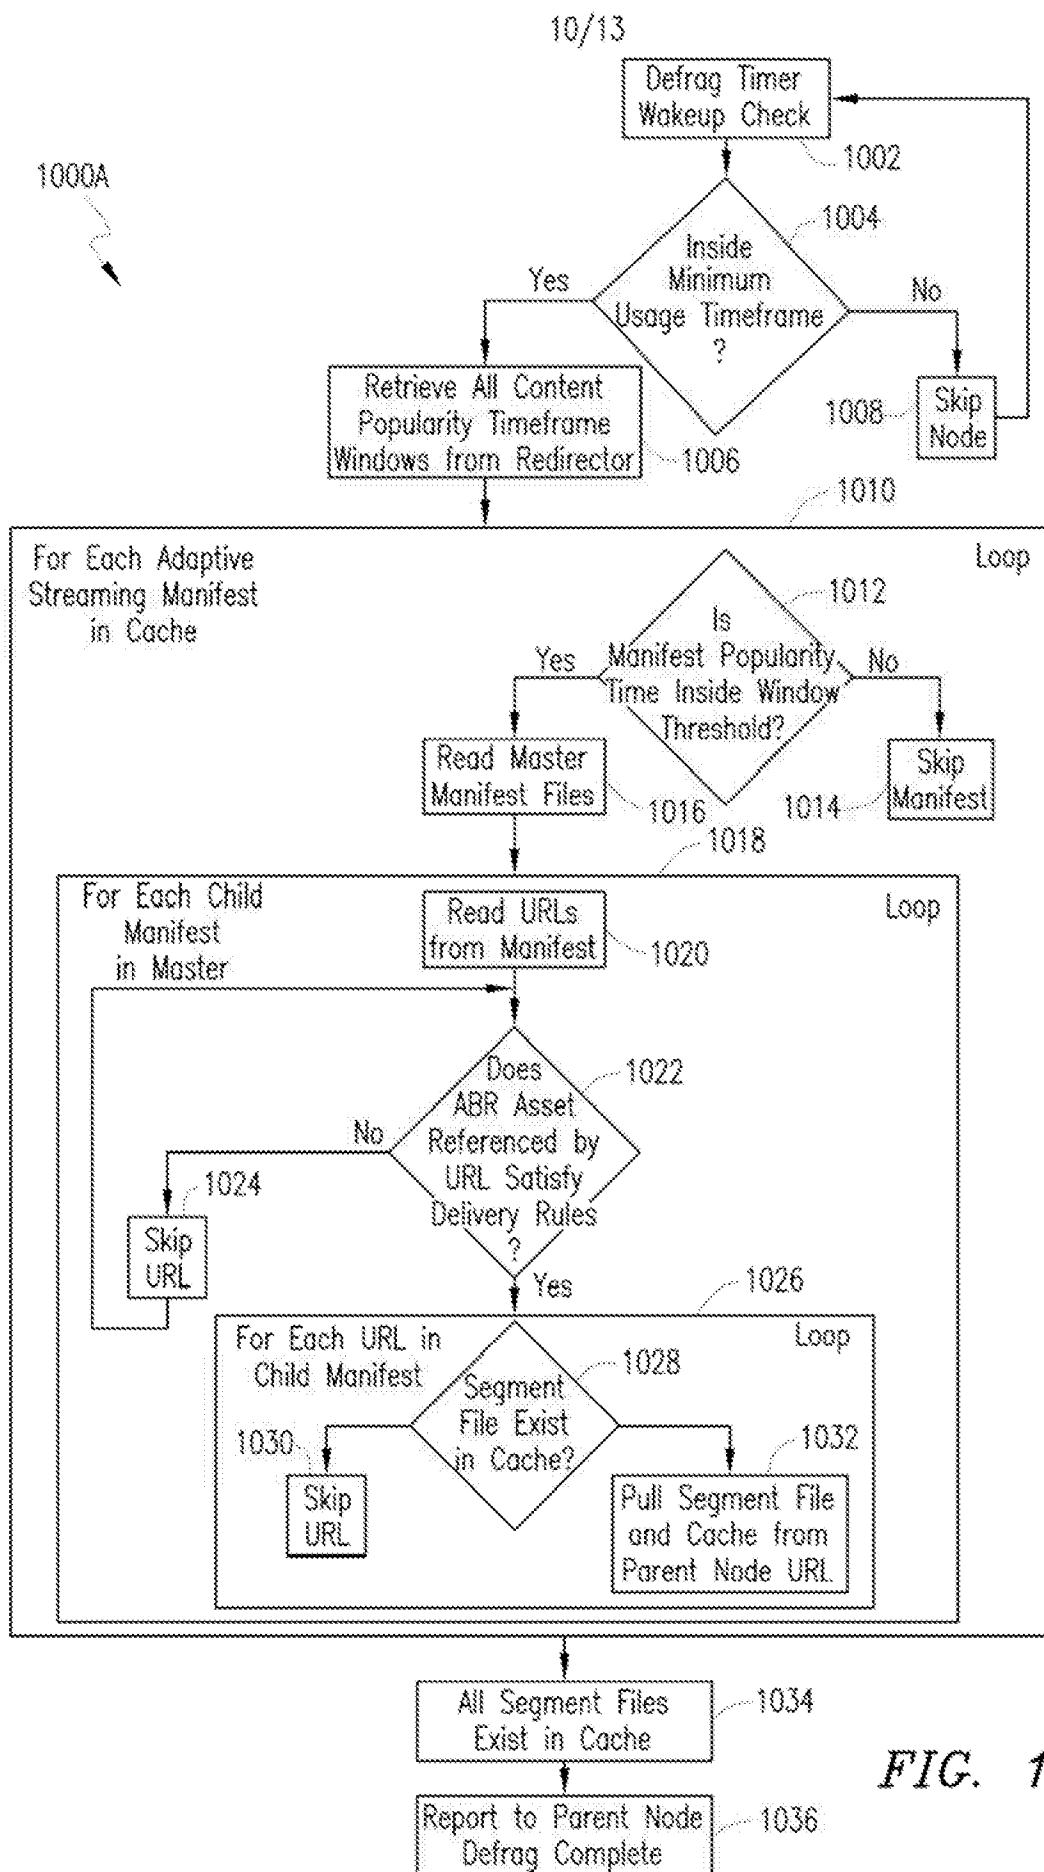


FIG. 10A

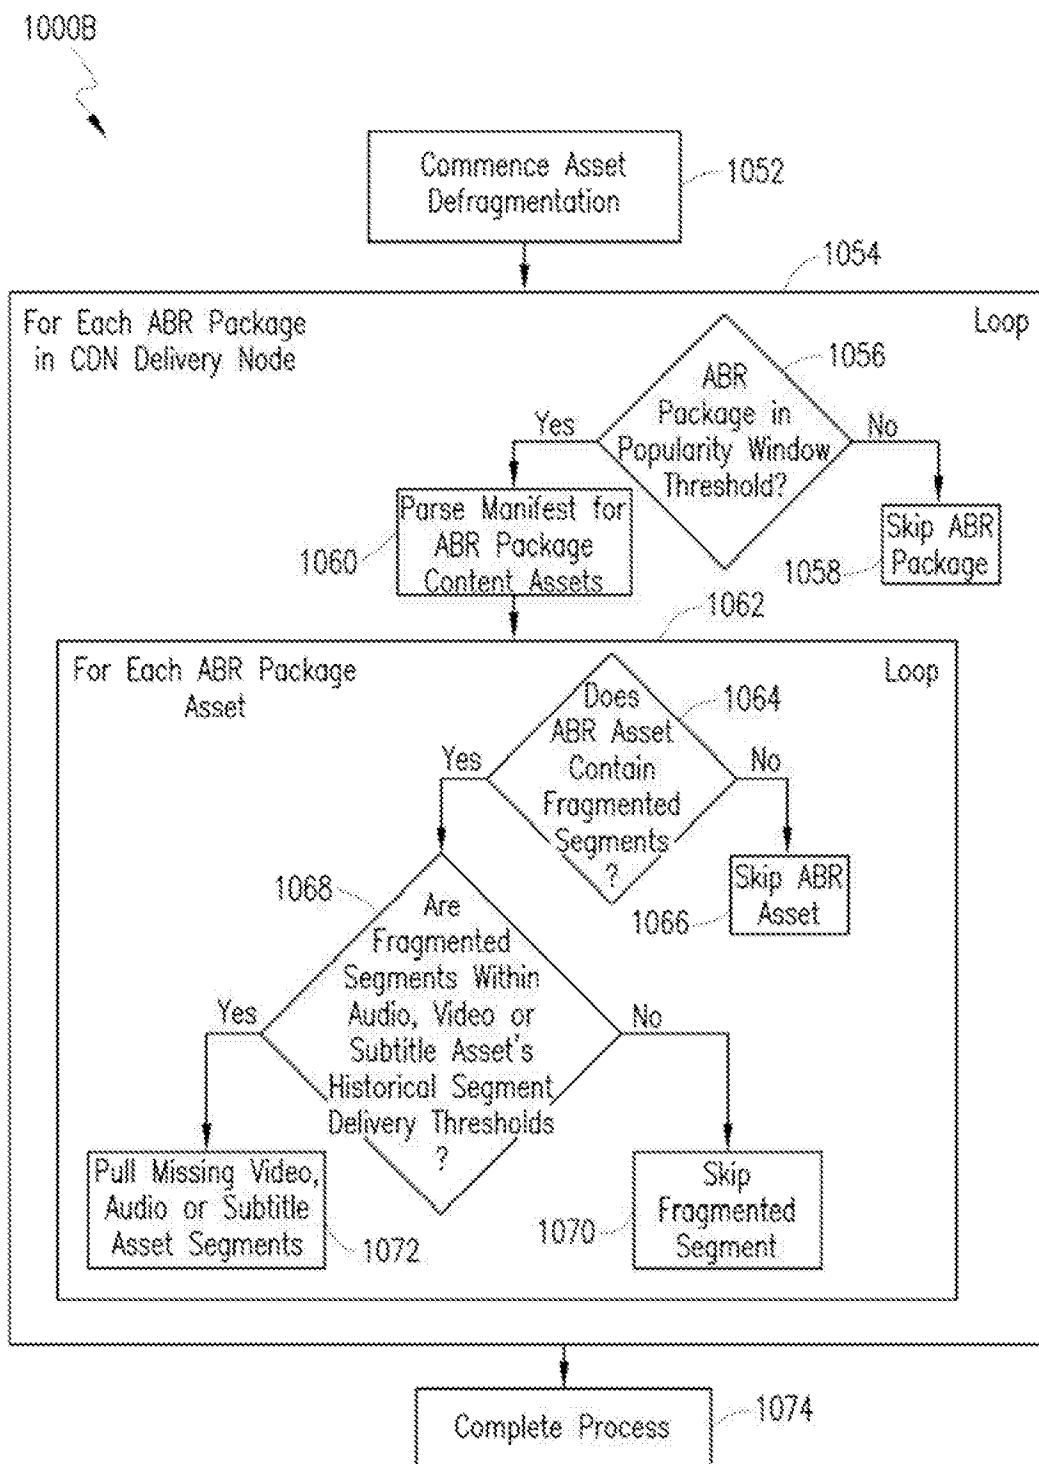


FIG. 10B

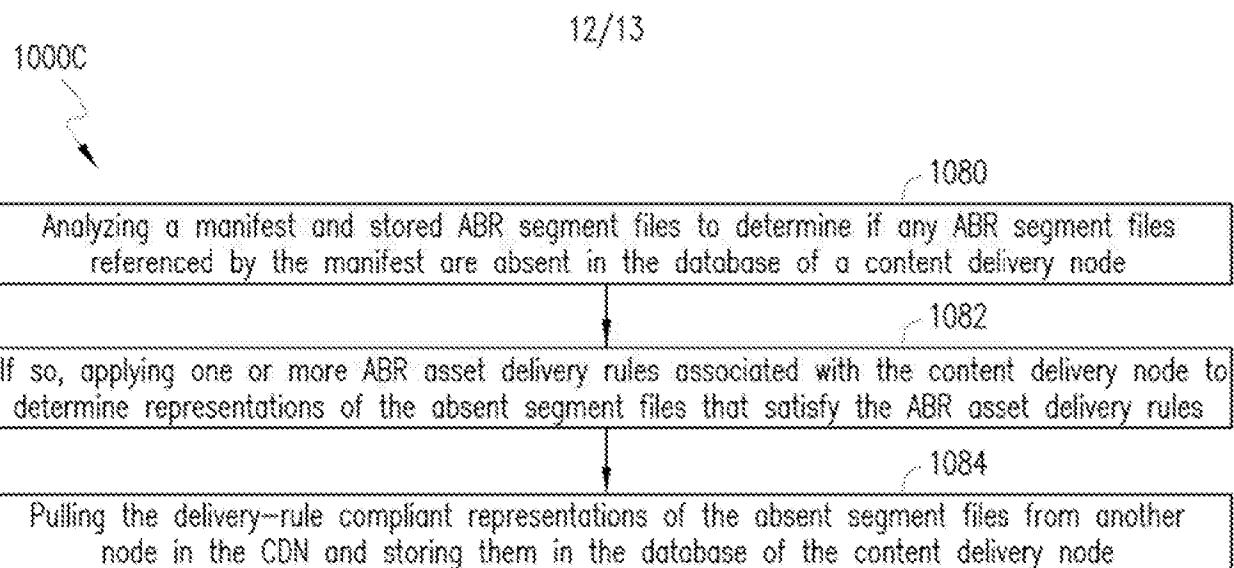
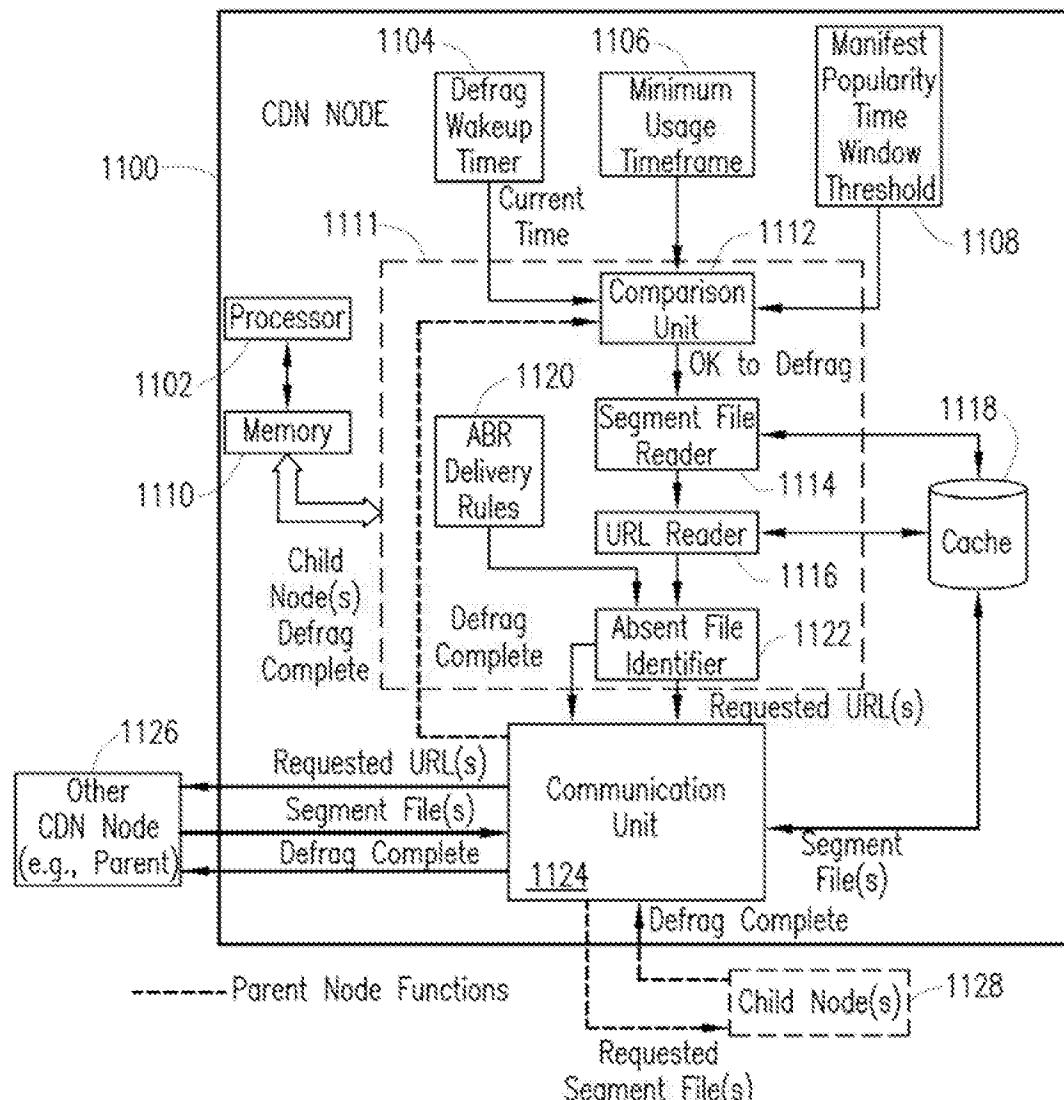
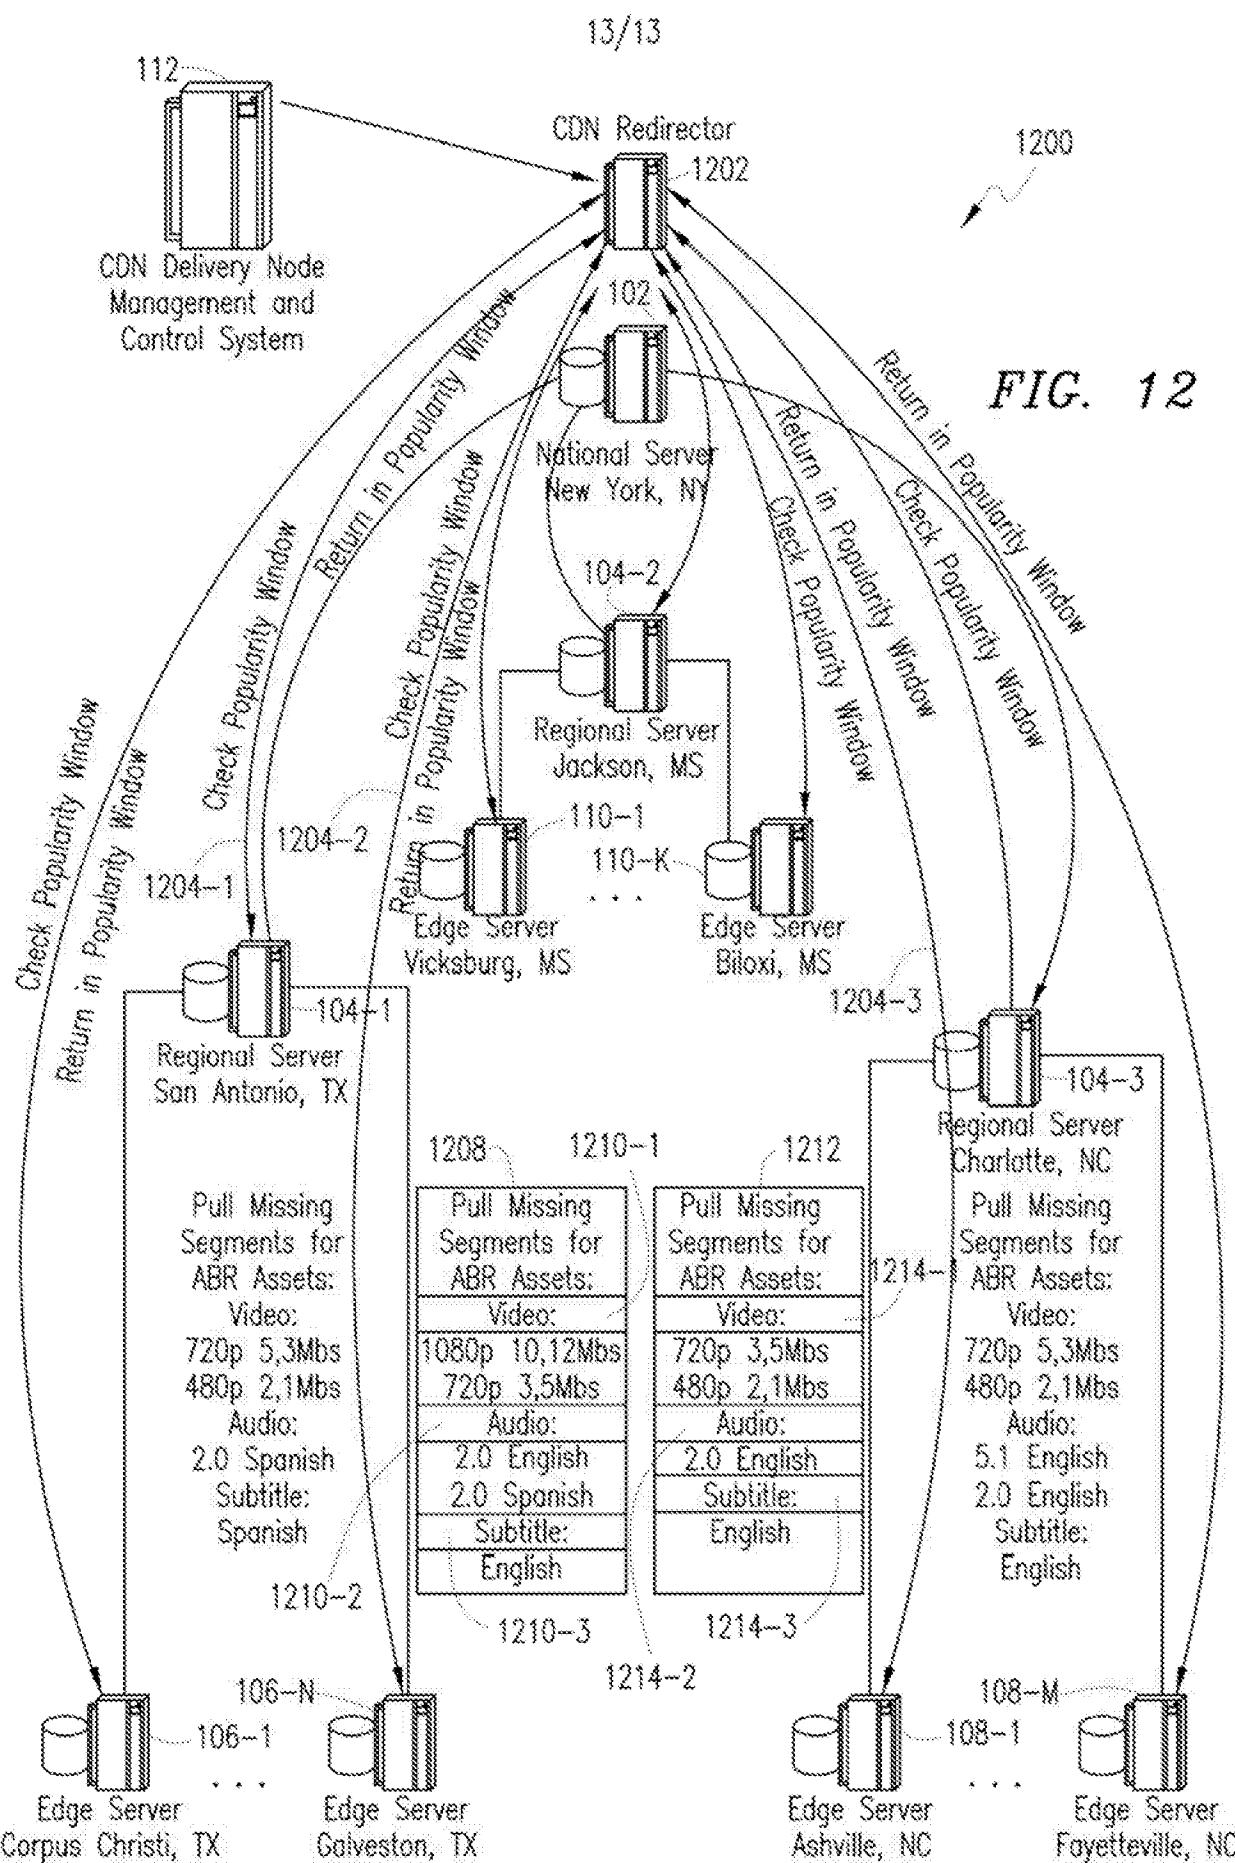


FIG. 10C





INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB2014/065742**A. CLASSIFICATION OF SUBJECT MATTER****H04N 21/25(2011.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
H04N 21/25; G06F 15/16; G06F 15/173; H04L 12/26; H04L 29/06Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: content delivery network (CDN), adaptive bitrate (ABR), historical delivery pattern, delivery rule, popular, pre-provisioning, edge server, historical network condition**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 2615777 A1 (TEKTRONIX, INC.) 17 July 2013 See paragraphs 28-32, 72; and figure 2.	1-22
A	US 2012-0124179 A1 (ADAM CAPPIO et al.) 17 May 2012 See paragraphs 50-58; and figure 3.	1-22
A	US 2007-0250560 A1 (JOEL M. WEIN et al.) 25 October 2007 See paragraphs 115-117; claim 1; and figure 4.	1-22
A	US 2013-0144979 A1 (VIVEK KANSAL et al.) 6 June 2013 See paragraphs 45-64; and figures 4-8.	1-22
A	US 2012-0198041 A1 (BRYAN BLACK et al.) 2 August 2012 See paragraphs 54-72; claims 1, 8; and figures 5, 6.	1-22

 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
28 January 2015 (28.01.2015)Date of mailing of the international search report
29 January 2015 (29.01.2015)Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea
Facsimile No. +82 42 472 3473

Authorized officer

LEE, Jin Ick

Telephone No. +82-42-481-5770



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/IB2014/065742

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 2615777 A1	17/07/2013	CN 103281594 A US 2013-0067109 A1	04/09/2013 14/03/2013
US 2012-0124179 A1	17/05/2012	EP 2638682 A2 EP 2638682 A4 WO 2012-065186 A2 WO 2012-065186 A3	18/09/2013 23/07/2014 18/05/2012 30/08/2012
US 2007-0250560 A1	25/10/2007	US 2007-0288588 A1 US 2012-151016 A1 US 2013-0297735 A1 US 7240100 B1 US 8122102 B2 US 8484319 B2	13/12/2007 14/06/2012 07/11/2013 03/07/2007 21/02/2012 09/07/2013
US 2013-0144979 A1	06/06/2013	US 8903955 B2	02/12/2014
US 2012-0198041 A1	02/08/2012	AU 2010-202034 B1 CN 102203758 A CN 102640135 A CN 103370919 A CN 103477335 A EP 2329395 A2 EP 2483793 A1 EP 2659388 A1 EP 2671165 A1 US 2010-0077056 A1 US 2011-0082982 A1 US 2011-0252100 A1 US 2011-0307586 A1 US 2012-0166574 A1 US 2012-0198022 A1 US 2012-0198042 A1 US 2012-0198069 A1 US 2012-0198070 A1 US 2012-0198071 A1 US 2012-0297033 A1 US 2013-0212208 A1 US 2013-0246570 A1 US 2013-0262627 A1 US 8090863 B2 US 8219645 B2 US 8255557 B2 US 8291083 B2 US 8370452 B2 US 8396970 B2 US 8458290 B2 US 8463876 B2 US 8478858 B2	23/12/2010 28/09/2011 15/08/2012 23/10/2013 25/12/2013 08/06/2011 08/08/2012 06/11/2013 11/12/2013 25/03/2010 07/04/2011 13/10/2011 15/12/2011 28/06/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 02/08/2012 03/01/2012 10/07/2012 28/08/2012 16/10/2012 05/02/2013 12/03/2013 04/06/2013 11/06/2013 02/07/2013

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/IB2014/065742

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 8521813 B2	27/08/2013
		US 8615577 B2	24/12/2013
		US 8856329 B2	07/10/2014
		WO 2010-033938 A2	25/03/2010
		WO 2010-033938 A3	01/07/2010
		WO 2011-041709 A1	07/04/2011
		WO 2012-091693 A1	05/07/2012
		WO 2012-105967 A1	09/08/2012