



US009613042B1

(12) **United States Patent**  
**Joseph et al.**

(10) **Patent No.:** **US 9,613,042 B1**  
(45) **Date of Patent:** **Apr. 4, 2017**

(54) <b>DYNAMIC GENERATION OF VIDEO MANIFEST FILES</b>	8,677,428 B2	3/2014	Lewis et al.
	8,683,066 B2	3/2014	Hurst et al.
	8,930,991 B2	1/2015	Philpott et al.
(71) Applicant: <b>Conviva Inc.</b> , San Mateo, CA (US)	2003/0204613 A1	10/2003	Hudson et al.
	2007/0041584 A1	2/2007	O'Connor et al.
(72) Inventors: <b>Dilip Antony Joseph</b> , Union City, CA (US); <b>Aditya R. Ganjam</b> , Mountain View, CA (US); <b>Faisal Zakaria Siddiqi</b> , San Jose, CA (US)	2007/0250560 A1	10/2007	Wein et al.
	2010/0241701 A1	9/2010	Lester et al.
	2011/0058675 A1	3/2011	Bruck et al.
	2011/0296048 A1	12/2011	Knox et al.
	2011/0314130 A1	12/2011	Strasman
	2012/0047542 A1 *	2/2012	Lewis et al. .... 725/97
(73) Assignee: <b>Conviva Inc.</b> , Foster City, CA (US)	2012/0198492 A1 *	8/2012	Dhruv et al. .... 725/32
	2012/0278496 A1 *	11/2012	Hsu ..... 1104N 21/4381 709:231
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 205 days.	2013/0067109 A1 *	3/2013	Dong et al. .... 709/231
	2014/0149557 A1 *	5/2014	Lohmar ..... 1104L 65/605 709:219

\* cited by examiner

(21) Appl. No.: **13/859,706**

(22) Filed: **Apr. 9, 2013**

**Related U.S. Application Data**

(60) Provisional application No. 61/621,973, filed on Apr. 9, 2012.

(51) **Int. Cl.**  
**G06F 17/30** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 17/30115** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 17/30899; G06F 17/30115  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

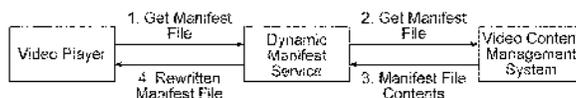
8,135,855 B2 3/2012 Sitaraman  
8,484,319 B2 7/2013 Wein et al.

*Primary Examiner* Etienne Leroux  
*Assistant Examiner* Cindy Nguyen  
(74) *Attorney, Agent, or Firm* Van Pelt, Yi & James LLP

(57) **ABSTRACT**

Some content players use manifest files to obtain and play video content. The manifest files include metadata information for obtaining video chunks that together constitute a full video, and can be used by content players to obtain the chunks, stitch them together, and render the video as a stream. Described herein are techniques for dynamically generating manifest files. A request for a manifest file is received, such as from a video player. Manifest file contents are retrieved, such as from an original manifest file. A manifest file to be used by a video player to fetch video chunks is then dynamically generated and returned.

**20 Claims, 8 Drawing Sheets**



```

#EXTM3U
#EXT-X-TARGETDURATION:10
#EXT-X-MEDIA-SEQUENCE:0
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_0.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_1.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_2.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo3.ts
#EXT-X-ENDLIST
  
```



FIG. 1A

```
#EXTM3U
#EXT-X-TARGETDURATION:10
#EXT-X-MEDIA-SEQUENCE:0
#EXTINF:10, no desc
http://cdnB.com/MyFavoriteVideo2000Kbps_0.ts
#EXTINF:10, no desc
http://cdnB.com/MyFavoriteVideo2000Kbps_1.ts
#EXTINF:10, no desc
http://cdnB.com/MyFavoriteVideo2000Kbps_2.ts
#EXTINF:10, no desc
http://cdnB.com/MyFavoriteVideo2000Kbps_3.ts
#EXT-X-ENDLIST
```

FIG. 1C

```
#EXTM3U
#EXT-X-TARGETDURATION:10
#EXT-X-MEDIA-SEQUENCE:0
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_0.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_1.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo500Kbps_2.ts
#EXTINF:10, no desc
http://cdnA.com/MyFavoriteVideo3.ts
#EXT-X-ENDLIST
```

FIG. 1B

110 →

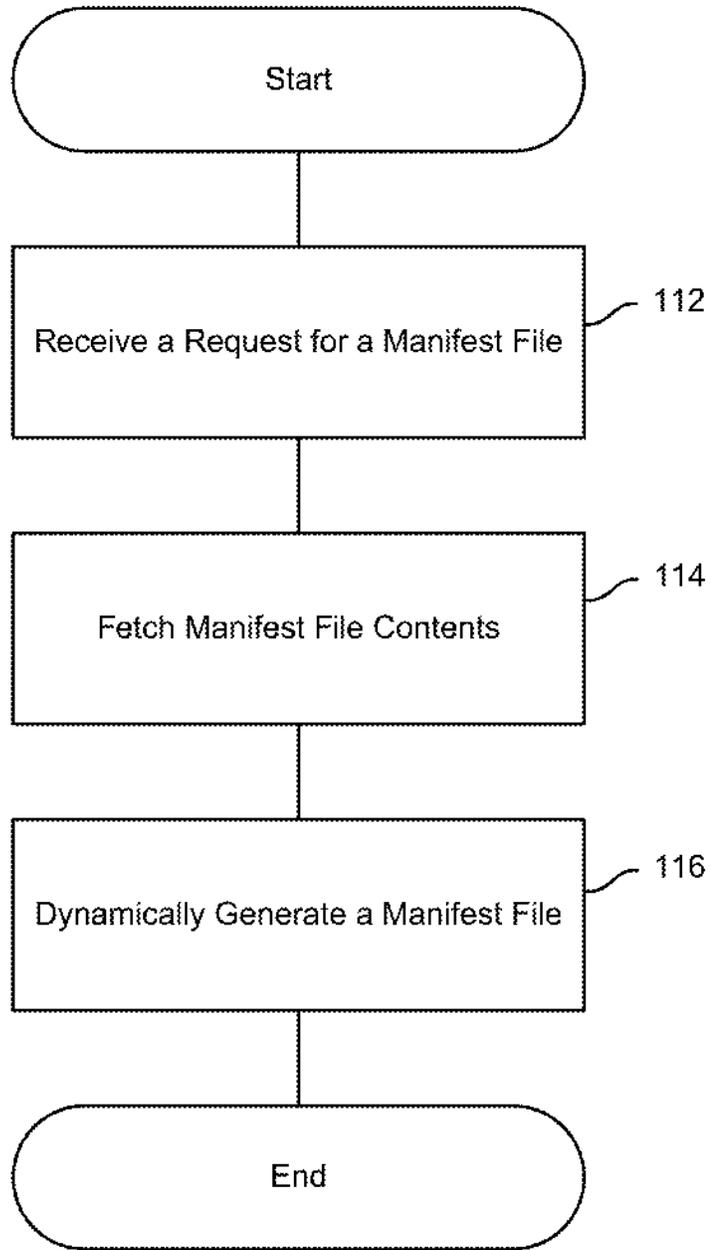


FIG. 1D

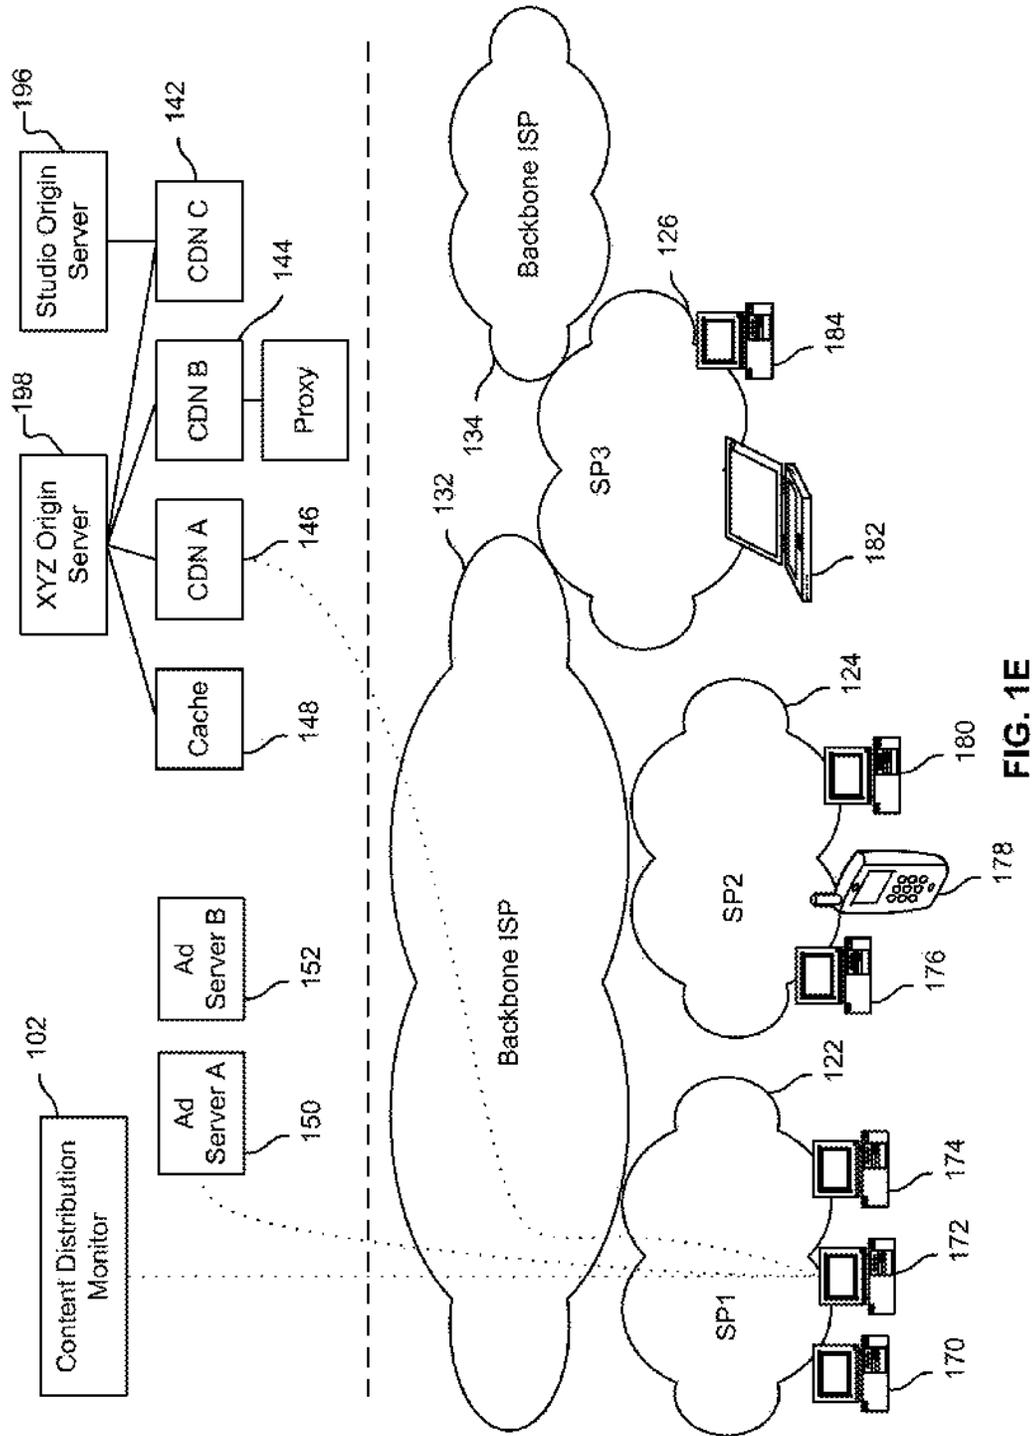


FIG. 1E

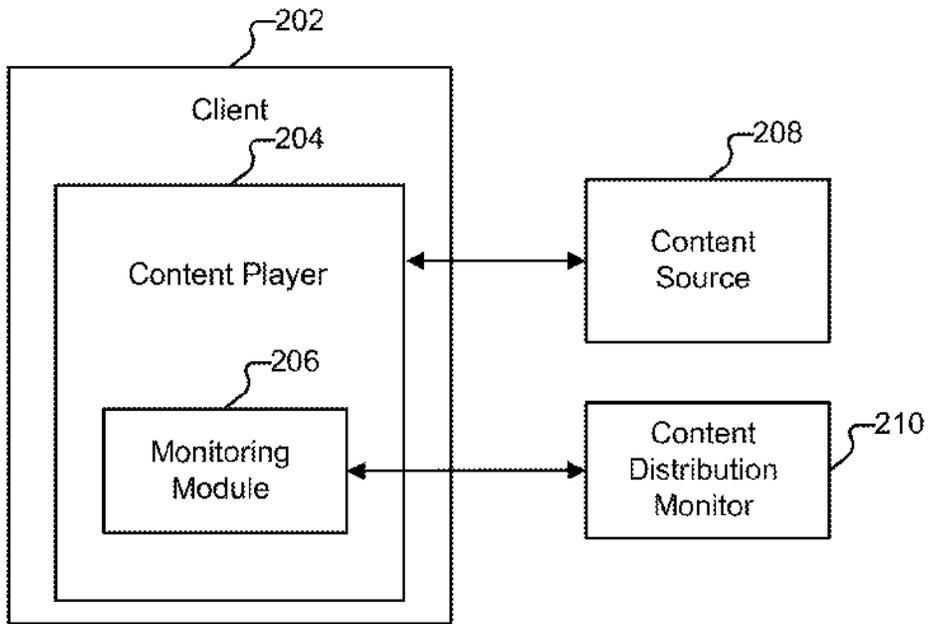


FIG. 2A

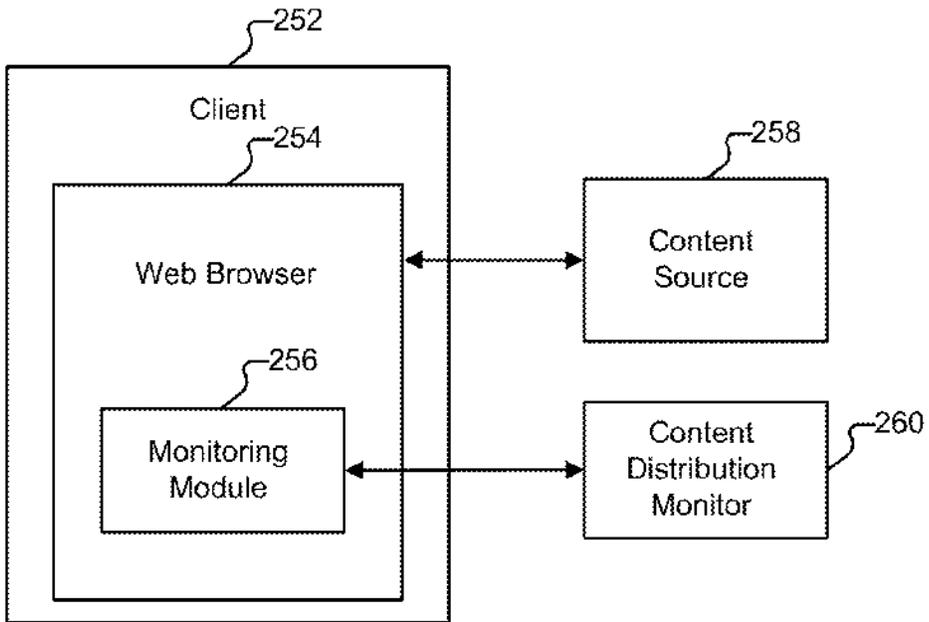


FIG. 2B

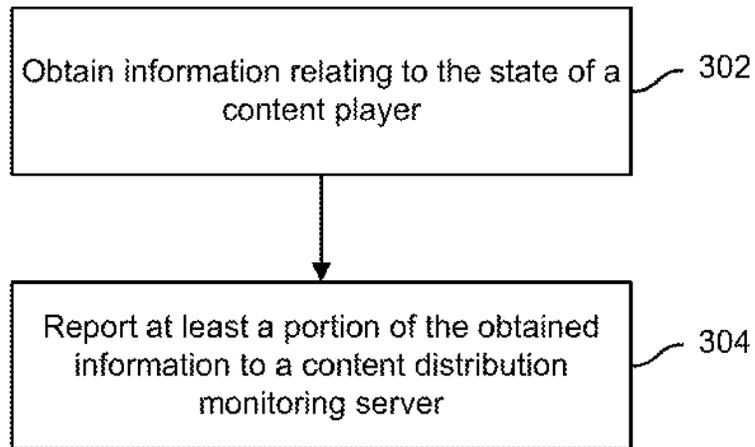


FIG. 3

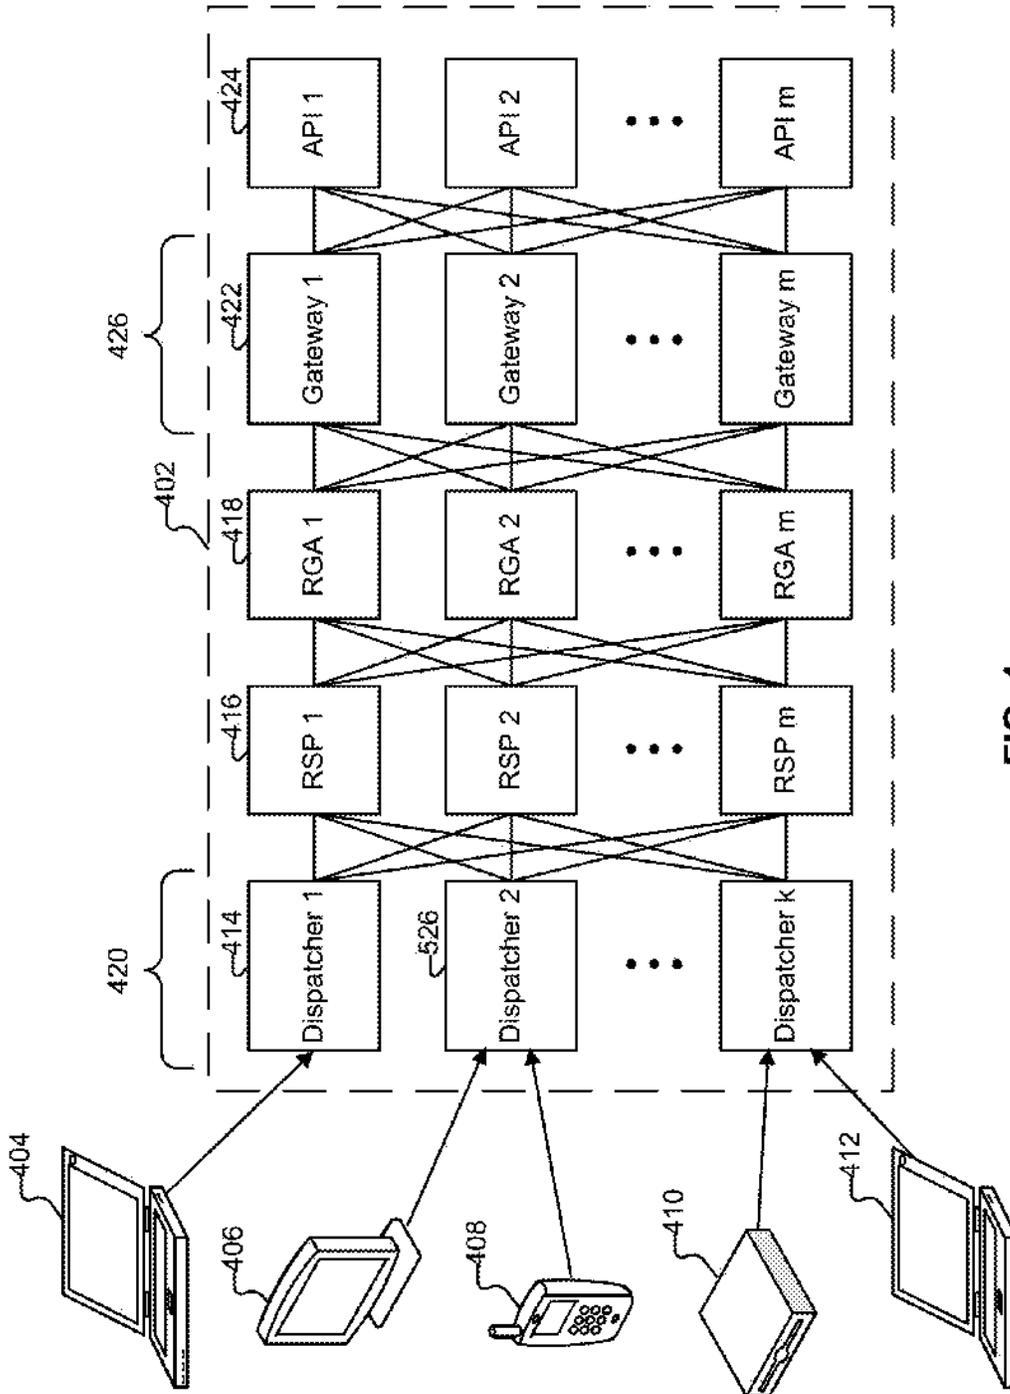


FIG. 4

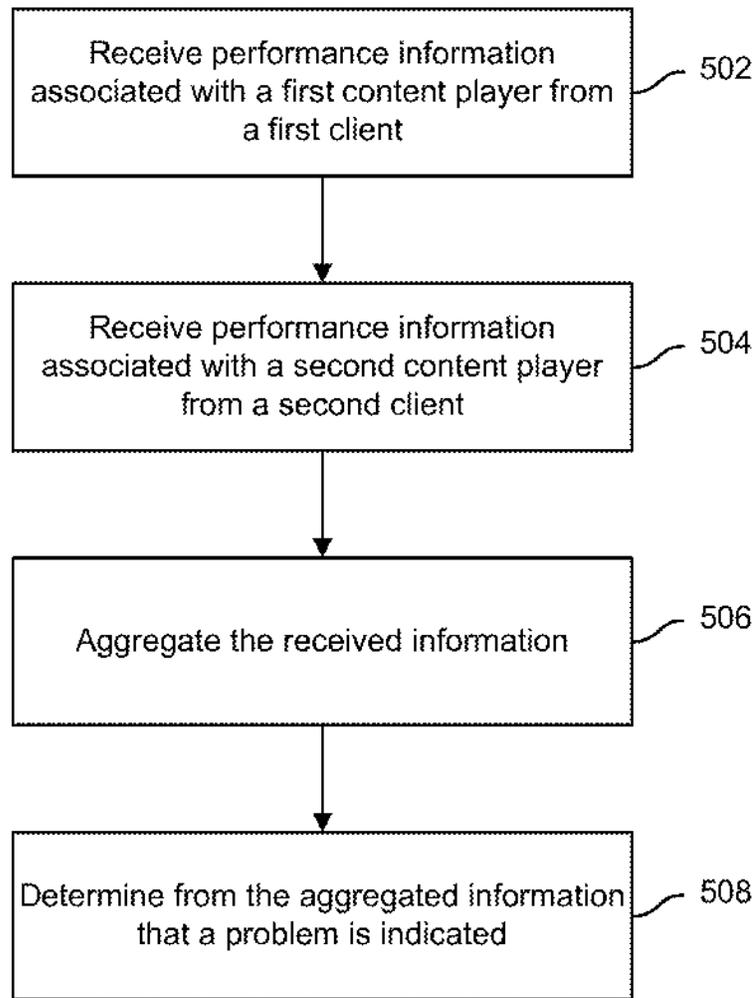


FIG. 5

1

## DYNAMIC GENERATION OF VIDEO MANIFEST FILES

### CROSS REFERENCE TO OTHER APPLICATIONS

This application claims priority to U.S. Provisional Patent Application No. 61/621,973 entitled DYNAMIC GENERATION OF VIDEO MANIFEST FILES filed Apr. 9, 2012 which is incorporated herein by reference for all purposes.

### BACKGROUND OF THE INVENTION

Chunk based protocols are used in online video delivery to various devices. As the bandwidth available to a video player and the performance of video sources involved in the delivery of video chunks can vary significantly while video is playing, there are challenges in maintaining a good viewing experience

### BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

FIG. 1A is an example illustration of a protocol to use dynamic manifest files.

FIG. 1B is an example illustration of a static manifest file.

FIG. 1C is an example illustration of a dynamically rewritten manifest file.

FIG. 1D is a flow diagram illustrating an embodiment of a process for dynamically generating video manifest files.

FIG. 1E is an illustration of an environment in which content is distributed.

FIG. 2A illustrates an embodiment of a client

FIG. 2B illustrates an embodiment of a client.

FIG. 3 illustrates an example of a process for monitoring the performance of a content player.

FIG. 4 is an illustration of an environment in which status information is received and processed.

FIG. 5 illustrates an example of a process for detecting a problem in a content distribution.

### DETAILED DESCRIPTION

The invention can be implemented in numerous ways, including as a process; an apparatus; a system; a composition of matter; a computer program product embodied on a computer readable storage medium; and/or a processor, such as a processor configured to execute instructions stored on and/or provided by a memory coupled to the processor. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention. Unless stated otherwise, a component such as a processor or a memory described as being configured to perform a task may be implemented as a general component that is temporarily configured to perform the task at a given time or a specific component that is manufactured to perform the task. As used herein, the term 'processor' refers to one or more devices, circuits, and/or processing cores configured to process data, such as computer program instructions.

A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodi-

2

ments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

#### Introduction

Chunk-based protocols dominate online video delivery in mobile devices, game consoles and set top boxes. In a chunk-based protocol, the video player first fetches a manifest (i.e., metadata) file that lists a sequence of URIs pointing to the chunks (typically 2-10 seconds long) that together constitute a full video. The player then fetches the chunks one by one, stitches them together, and renders the video as one smooth stream.

The bandwidth available to the video player and the performance of the video servers/Content Delivery Networks (CDNs) involved in the delivery of the chunk bytes can vary significantly while the video is playing. By dynamically adjusting the video bitrate and the CDN, a good viewing experience for the user can be better maintained. Unfortunately, current manifest files are static, and restrict such dynamic adjustments.

Described herein are techniques for maintaining a good viewing experience by dynamically generating manifest files.

#### Static Manifest Files

A static manifest file fetched by a video player can be of two types:

1. The manifest file contains a list of URIs to chunks at one particular bitrate on one particular server/CDN.
2. The manifest file contains a list of URLs to different sub-manifest files, each corresponding to one particular bitrate and one particular server/CDN. The sub-manifest file for bitrate B on CDN C contains URIs to bitrate B chunks on CDN C.

In (1), the video player cannot change bitrate/CDN to improve video quality, as it does not know about the other bitrate/CDN choices it can use to play the same video.

In (2), the video player can change bitrate/CDN by downloading chunks from a different sub-manifest file. The decision to change the bitrate/CDN is made solely based on local information at the video player. For example, the video player can decide to lower the bitrate if it detects that its available bandwidth has decreased. It can decide to download chunks from a different CDN if it experiences a connection error when talking to the current CDN. One limitation of this approach is that the video player can rely only on local information, which by itself is not likely sufficient to make the best decision.

#### Dynamic Generation of Manifest Files

FIG. 1A is a diagram illustrating an embodiment of a protocol to use dynamic manifest files.

Instead of requesting the manifest file from the actual "URLA" provided by the video content management system (CMS), the video player requests the manifest file from a dynamic manifest service. The request includes the following information:

- a. The actual URL/A
- b. Device Type
- c. Content Type

3

- d. Current and historical bandwidth estimate
- e. Bitrate history from prior videos
- f. CDN history from prior videos
- g. Geographical information like country, state, city, ASN
- h. Device resources (e.g., available memory, CPU)
- i. Content length (e.g., short form and long form VoD, live)

The dynamic manifest service fetches the contents of the manifest file from URI A. It then overwrites the contents of the manifest file to specify the bitrate and CDN/server that can provide the best video quality to the requesting video player. The overwritten manifest file is returned to the video player.

The video player fetches the chunks of the specified bitrate from the CDN/server specified in the dynamically generated manifest file.

The techniques described herein do not require the video player to expose a new API to control it, thus allowing the video player software to remain unmodified.

#### Example

##### Start-Time Bitrate and CDN Selection for Apple HTTP Live Streaming (HLS)

Apple HTTP Live Streaming, or HLS in short, is a video delivery mechanism widely used in iPhones, iPads, Android phones, and set-top boxes like Roku. In this example, we describe how our dynamic manifest service invention helps the video player start playing at the best available bitrate/CDN combination.

The video player first requests the manifest file MyFavoriteVideo.m3u8 from the dynamic manifest service using the standard HTTP protocol as follows:

```
HTTP GET
http://BestQualityVideoCorp.com/DynamicManifest?
originalUrl=http%3A%2F%2FcdnA.com%
2FMyFavoriteVideo.m3u8&city=SanFrancisco&
lastBandwidthMeasurement 3000Kbps&deviceType
iPad&contentType VoD&isp Comcast
```

The dynamic manifest service at BestQualityVideo Corp.com fetches the static HLS manifest file (shown in FIG. 1B) from `http://cdnA.com/MyFavoriteVideo.m3u8`. It then calculates the best bitrate/CDN to be used for this particular video session. For example, it decides to use bitrate 2000 Kbps since the request indicated that the video player has previously experienced 3000 Kbps bandwidth. It decides to use CDN B instead of CDN A, since it knows that CDN B is currently the best performing CDN for iPads in San Francisco on the Comcast ISP network, based on global knowledge aggregated from a large number of video players. Once the decision is made, the original manifest file is rewritten to include URLs pointing to 2000 Kbps chunks from CDN B. The rewritten manifest file is shown in FIG. 1C.

When the video player receives this rewritten manifest file, it directly starts playing 2000 Kbps bitrate video from CDN B. Unlike existing solutions, the viewer does not have to endure watching video at the lowest quality for the first few tens of seconds.

##### Example Deployment Options for the Dynamic Manifest Service

The following are four different example options for deploying the Dynamic Manifest Service:

1. On the same device as the video player
2. On the video Content Management System

4

3. On the CDN origin or edge servers
4. On a 3rdParty service provider

When the dynamic manifest service runs as a local proxy on the same device as the video player, it can overwrite all chunk URLs to point to its local address. It then transparently translates each chunk request to the appropriate bitrate and CDN.

##### Selecting the Best Bitrate and CDN

The dynamic manifest service can rely on the following information to decide the best CDN and bitrate to be used for a particular video player:

1. The attributes sent by the video player in the request, listed earlier. For example, the video player is running on an iPad in San Francisco on the Comcast ISP network.
2. Global information about current network conditions aggregated from other video players. For example, a large number of iPad video players on CDN A have been experiencing high buffering (and hence poor video quality) on Comcast and ATT ISP networks in San Francisco during the last 35 minutes. Hence, based on global network information, CDN B is a better choice for iPad video players in San Francisco on Comcast or ATT.
3. Traffic and business policies. For example, the customer or CDN may dictate that all viewers should currently play a lower bitrate than what they are normally capable of because the network is currently undergoing maintenance or high congestion. In another example, the content provider may dictate that certain devices (e.g., iOS devices) only stream from a particular CDN, as that CDN provides better support for those devices.
4. Local information from the video player gathered via a separate communication channel. The video player can periodically update the dynamic manifest with information about its local state like current bandwidth, whether the user is playing in full screen mode, etc through an independent communication channel.

##### Mid-Stream CDN/Bitrate Switching

The dynamic manifest service can instruct the video player to switch the bitrate/CDN being after the video has started playing as follows:

When responding to the initial manifest file request from the video player, the dynamic manifest service only includes URLs to the first few video chunks. It also omits the End-Of-File marker. Chunk-based video players are already designed to refetch a manifest file which does not have an End-Of-File marker when it has finished playing all the chunks listed in the partial manifest file.

When the video player requests the manifest file again, the dynamic manifest service looks up internal per video player session state to determine the next chunk to give this video player. It calculates the best bitrate and CDN to be used currently and writes out a manifest file with the next few chunks of the selected bitrate from the selected CDN. It then updates the internal per video player session state counter that tracks the chunks already provided to this video player.

The video player fetches and plays the chunks from the URLs specified in the latest manifest file. This protocol is applicable to both live and on demand video.

FIG. 1D is a flow diagram illustrating an embodiment of a process for dynamically generating video manifest files. In some embodiments, process 110 is executed by the dynamic manifest service described herein. The process begins at 112 when a request for a manifest file is received. The request can be received from a video player. The request can include information such a URL provided by a video content management system (CMS), device type, content type, current

and historical bandwidth estimate, bitrate history from prior videos, CDN history from prior videos, geographical information (e.g., country, state, city, ASN), device resources (e.g., available memory, CPU), content length (e.g., short form and long form VoD, live), or any other appropriate information. At 114, contents (e.g., sequence of URLs pointing to video chunks) of a manifest file (e.g., original manifest specified in a URI, included in a request) are fetched/retrieved. In some embodiments, a manifest file (e.g., original manifest file) is fetched. At 116, a manifest file is dynamically generated. For example, the contents of an original, static, manifest file specified in a URI, included in a request can be overwritten/rewritten to specify the bitrate and CDN/server that can provide the best video quality to the requesting video player. In some embodiments, the best bitrate/CDN to be used for a particular video session are calculated, for example, based on global knowledge aggregated from a large number of video players. The best bitrate/CDN can be selected, as described above, based on the attributes sent by the video player in the request, listed earlier, global information about current network conditions aggregated from other video players, traffic and business policies, local information from the video player gathered via a separate communication channel, etc. The manifest file can be written with chunks of the selected/calculated best bitrate and selected/calculated best CDN. The dynamically generated manifest file can then be returned to the video player. The video player can then fetch the chunks of the specified bitrate from the CDN/server specified in the dynamically generated manifest file.

#### Example Environment

Additional detail regarding an example environment in which the techniques described above can be deployed are provided below.

FIG. 11 is an illustration of an environment in which content is distributed. In the example shown, clients 170-184 are used to access content, such as audiovisual content (e.g., movies, songs, television shows, sporting events, games, images, etc.) that is owned by content owners. The content is stored (or captured) at origin servers 196-198, then distributed via other servers, caches, content distribution networks (CDNs), proxies, etc. (collectively, "content sources"). Content sources employ a variety of technologies and include HTTP, Adobe Flash Media, and Microsoft Internet Information Service servers. In some embodiments content is also distributed by clients (e.g., using peer-to-peer techniques).

Examples of clients include personal computers (170), laptops (182), cellular phones/personal digital assistants (178), and other types of information appliances (not shown) such as set-top boxes, game consoles, broadband routers, file servers, video servers, and digital video recorders, as applicable. The clients shown are used by subscribers to various Internet service providers (ISPs). For example, clients 170, 172, and 174 are subscribed to SP1 (122), while clients 176, 178, and 180 are subscribed to SP2 (124), and clients 182 and 184 are subscribed to SP3 (126).

In the example shown, a movie studio ("Studio") has contracted with content distributor 142 to provide downloadable copies of its films in exchange for a fee. Similarly, a television network ("XYZ") has contracted with content distributors 142-146 to provide viewers with access to live streams of its broadcasts as well as streams of television show episodes and sporting events. In some cases, the content distributor is owned/operated by the content owner.

Content distributor 142 has a data center that is provided with network access by backbone ISP 132. Though represented here by a single node (also referred to herein as a "CDN node"), content distributor 142 may typically have multiple data centers (not shown) and may make use of multiple backbone or other ISPs. Content distributor 144 has a data center that is provided with network access by backbone ISP 134. Advertisements are served to various clients via ad servers 150-152.

Suppose a user of client 172 (hereinafter "Alice") would like to watch a live soccer game owned by XYZ. Client 172 includes a web browser application. Alice uses the web browser application to navigate to a portal owned by XYZ, such as "http://xyztvnetwork.com/livegames." Her request for the game is directed to a CDN node that is closest to her. In this case, CDN 146 is the fewest hops away from her client. Her client then begins streaming the content from CDN 146, which is in turn rendered in her browser (e.g., via a Flash or Silverlight player). Advertisements, associated with the portal, are served to her by ad server 150.

In addition to CDN 146 and ad server 150, Alice's client is also in communication with content distribution monitor 102. As will be described in more detail below, client 172 provides status information (also referred to herein as a "heartbeat"), on a recurring basis, to content distribution monitor 102.

The status information includes a variety of telemetry data such as information that captures the quality of the user experience (e.g., video stream quality), and information pertaining to user behavior. Examples of quality metrics include: the length of time it takes for the soccer game video to start playing, the number of buffering events (if any), the length of buffering events, and the number of frames per second rendered by the video player. Examples of user behavior include: starting and stopping playing a video or audio stream, seeking within the stream, switching the player to full screen mode, minimizing/restoring the player, a change in the volume level of the player, and clicking on an advertisement.

As other users of clients 170-184 request content, their respective players similarly obtain content from content sources such as CDN 144 and also communicate status information (also referred to herein as telemetry information) to content distribution monitor 102. Such players may be browser-based as with Alice's, or they may be standalone applications, as applicable. In various embodiments, all clients in the environment provide status information to content distribution monitor 102.

As will be described in more detail below, different clients may provide content distribution monitor 102 with different levels of detail, and may also do so with differing frequency. For example, client 178 is a smartphone with less powerful hardware than client 172 and more limited bandwidth. It is configured to provide less information to content distribution monitor 102 than client 172 and also does so less frequently than client 172.

Content distribution monitor 102 collects and processes the information received from Alice's client along with other clients. The collected information is made available in real-time to control entities/operators and can be used to detect and remedy problems in the content distribution. Examples of such problems include excessive buffering, freezing, and frame skipping. Additional information pertaining to delivery resources (e.g., CDN 142) and network providers (e.g., ISP 126) is also made available, as is other information pertaining to clients such as demographic information.

In the example shown in FIG. 1E, a single content distribution monitor 102 is used. Portions of content distribution monitor 102 may be provided by and/or replicated across various other modules or infrastructure depending, for example, on factors such as scalability and availability (reducing the likelihood of having a single point of failure), and the techniques described herein may be adapted accordingly. In some embodiments content distribution monitor 102 is implemented across a set of machines distributed among several data centers. A Resilience Service Layer (RSL) can also be used to ensure that the monitoring service is not disrupted when a subset of machines fail or a subset of data centers hosting the content distribution monitor are disconnected from the Internet.

#### Examples of Client Architecture

In various embodiments, the collection of status information and the reporting of that information to the content distribution manager are performed by a “monitoring module” included in the client.

FIG. 2A illustrates an embodiment of a client. In the example shown, client 202 includes a content player application 204 which in turn incorporates monitoring module 206. Content is received by the player from content source 208. Monitoring module 206 is in communication with content distribution monitor 210.

In the example shown, module 206 is implemented in ActionScript and deployed as an SWF library that is dynamically loaded by player 204. In Flash, the NetStream() class is mainly responsible for streaming, buffering, and playing the video. The monitor is implemented as an element in the content player which wraps the video element, and provides the same interface as the video element, as described in more detail below.

Module 206 can also be implemented using other approaches, such as in the .NET platform for Silverlight and deployed as a DLL library that is dynamically loaded by player application 204. In Silverlight, the MediaElement() class is mainly responsible for streaming, buffering, and playing the video. The monitor is implemented as an element in the content player which wraps the video element, and provides the same interface as the video element, as described in more detail below.

Examples of some of the information collected by monitoring module 206 include the following, expressed in key-value pairs:

- (player\_state, “buffering”): The stream is buffering.
- (buffer\_length, 5 s): The current buffer length is five seconds.
- (join\_time, 3 s): The join time was 3 sec.
- (frame\_per\_second, 30): The number of frames per second is 30.
- (player-mode, “Full Screen”): The player is running in full-screen mode.

FIG. 2B illustrates an embodiment of a client. In the example shown, client 252 includes a web browser application which in turn incorporates monitoring module 256. Content is received by the player from content source 258. Monitoring module 256 is in communication with content distribution monitor 260. In the example shown, module 256 is implemented in JavaScript. The monitor periodically collects information about the current web page and the browser. Examples of some of the information collected by monitoring module 256 include the following, expressed in key, value pairs:

- (browser\_minimized, “yes”): The browser window is minimized.

(tab\_visible, “no”): The tab containing the web page is not visible.

(pointer\_pos, “x, y”): The position of the pointer on the web page.

(banner\_display, “ACMECars”): The banner ad on the web page is for ACME Cars.

As explained above, the monitoring module can be implemented in a variety of ways. For example, the monitoring module can be included in the client’s player by the author of the player. The monitoring module can also extend the functionality of an existing player, such as by being implemented as a plugin, library, or standalone application (e.g., a helper application). Various examples of techniques for integrating a monitoring module with an existing player will now be described.

In one embodiment, the monitoring module is a wrapper around the lowest level streamer module (e.g., “NetStream” in Flash, or “MediaElement” in Silverlight). The player uses the wrapper, which provides the same or enhanced version of the API provided by the streamer module, to stream video. Logic in the wrapper captures the appropriate telemetry data.

#### Example

“ConvivaNetStream” Extends and Wraps “NetStream”

```
var ns:NetStream new ConvivaNetStream( );
ns.play(<stream>);
```

In a second embodiment, the monitoring module exists as a side attachment and is not present in the code path for streaming to occur. The player passes the streamer module to the monitoring module either directly or through a proxy module. The monitoring module reads properties and listens for events from the streamer module or proxy module to collect data. The proxy module method prevents the monitoring module from interfering with the streaming.

#### Example 1

```
var ns:NetStream new NetStream( );
ns.play(<stream>);
LivePass.createMonitoringSession(ns);
```

#### Example 2

“NetStreamProxy” wraps “NetStream,” but prohibits any calls that may adversely impact “NetStream”:

```
var ns:NetStream=new NetStream( );
var nsProxy:NetStreamProxy new NetStreamProxy(ns);
ns.play(<stream>);
LivePass.createMonitoringSession(nsProxy);
```

In a third embodiment, the monitoring module is built into the streamer module. The player uses the streamer module. The monitoring module within collects data.

#### Example

```
var ns:NetStream new NetStream( )
ns.play(<stream>);
```

#### Additional Status Information Examples

As mentioned above, clients downloading/streaming content from content providers are also in communication with content distribution monitor 102 on a recurring, time-driven basis. Each heartbeat contains a snapshot of the session state at the time the heartbeat was sent out and a summary of the events and measurements since the last heartbeat or since the

start of the session, as applicable. Examples of information that can be included in a heartbeat (though need not be present in every heartbeat) include the following:

version: The version of the heartbeat.  
 clientID: A unique identifier associated with the client's monitoring module (described in more detail below) when it is downloaded/installed on the client.  
 clientVersion: The version of the monitoring module.  
 customerID: An identifier associated with an entity such as XYZ or Studio.  
 sessionID: A unique identifier associated with the content viewing session.  
 objectID: A unique identifier associated with the content being streamed. In Alice's case, this is the soccer game.  
 currentResource: The source from which the content is currently being obtained. In Alice's case, this is CDN 146.  
 candidateResourceList: A list of potential sources from which the player can obtain the content. In Alice's case, this could include CDN 144, CDN 142, cache 148, etc. The "currentResource" is one resource in the "candidateResourceList."  
 resourceUsage: The number of bytes loaded/streamed since the start of the session.  
 currentBufferSize: The current buffer size.  
 minBufferSize: The minimum buffer size during the previous heartbeat interval.  
 maxBufferSize: The maximum buffer size during the previous heartbeat interval.  
 numBufferEmptyEvents: The number of buffering empty events since previous heartbeat.  
 currentBitrate: The bitrate at which the content is currently streamed.  
 playheadTime: For video-on-demand content (i.e., not a live event), this is the time offset of the stream file.  
 currentlyPlaying: For live content, this is the time offset of the stream since the player started.  
 joinTime: The amount of time that it took the player to enter the player state for the session.  
 averageFPS: The average rendering frame per second (FPS) at which the content was rendered since the last heartbeat.  
 encodedFPS: The FPS at which the content was encoded.  
 averageFPS: Can be lower than the "encodedFPS" if the client's CPU is unable to render the content at the "encodedFPS" or if the client cannot stream the content fast enough.  
 totalPlayTime: The total play time since the beginning of the session.  
 totalBufferingTime: The total buffering time since the beginning of the session.  
 totalPauseTime: The total pause time since the beginning of the session.  
 totalSleepTime: The total sleep time since the beginning of the session. The client is in a sleep state if the client is suspended.  
 sessionTime: The total time elapsed since the session started.  
 currentState: The state of the player when the heartbeat was sent. Examples of player states include: play, pause, buffering, joining, and seeking  
 numRateSwitches: The number of bitrate switches since the beginning of the session.  
 numResourceSwitches: The number of resource switches since the beginning of the session.

rateSwitchingEvent.time: The time at which a bitrate switch was attempted, measured from the start of the session.  
 rateSwitchingEvent.from: The bitrate of the stream at the time the switch was attempted.  
 rateSwitchingEvent.to: The target bitrate of the switch.  
 rateSwitchingEvent.result: Indicates whether the switch was successful or not. For example, a switch may not succeed if the client cannot sustain the "rateSwitchingEvent.result" bitrate.  
 resourceSwitchingEvent.time: The time at which a resource switch was attempted, measured from the start of the session.  
 resourceSwitchingEvent.from: The resource from which the content was streamed at the time the switch was attempted.  
 resourceSwitchingEvent.to: The target resource for the switch.  
 resourceSwitchingEvent.results: Indicates whether the switch was successful or not.  
 errorList: A list of errors encountered from the start of the session.

FIG. 3 illustrates an example of a process for monitoring the performance of a content player. In some embodiments the process shown in FIG. 3 is performed by a client such as client 172. The process begins at 302 when information relating to the state of a content player is obtained. For example, at 302, client 172 collects various information pertaining to the video player that Alice uses to watch the streaming soccer game. Information such as the amount of time an advertisement took to load and which ad server supplied the advertisement is also obtained. In some embodiments, additional processing is performed in the collected information. Examples include computing averages, minimums, maximums, and the  $x^{th}$  percentile over the last few samples.

At 304, at least a portion of the obtained information is reported to a content distribution monitoring server. For example, at 304, client 172 sends a heartbeat to content distribution monitor 102, including some or all of the information collected at 302. Information sent at 304 can also include processed results such as averages and minimums, instead of or in addition to raw data.

In various embodiments, the process shown in FIG. 3 repeats throughout a content playing session. The process may repeat regularly (e.g., once every second), and may also repeat with varying frequency.

The amount of data reported by the client in a heartbeat to content distribution monitor 102 can affect the scalability of the content distribution monitor. The amount of data collected and reported by the client can also potentially impact the performance of the client. If the client is otherwise resource constrained (e.g., due to other processes, or due to hardware limitations), adjustments can be made to how much information is collected and with what frequency. For example, suppose the buffer size is expected to be of size three seconds. The time period employed by the monitoring module can be set to one second or a few hundred milliseconds. The adjustment can be dynamically adjusted as needed. The period can be decreased if the buffer grows and increased if the buffer shrinks, thus minimizing overhead while still being able to detect a low buffer size which can impact the video quality.

The amount of data sent can be adjusted according to two parameters: the heartbeat "period," and the size of heartbeats, both of which can be adaptively changed by either the client or the content distribution monitor as needed. For

example, if the quality experienced by the client is acceptable, the client may reduce the heartbeat frequency. This can entail collecting state information less frequently and can also entail sending collected information less frequently. If the quality degrades, the client can increase the heartbeat frequency accordingly. As one example, client 172 can employ a rule that, as long as the buffering ratio is less than or equal to 0.5%, the heartbeat period is 30 seconds. If the buffering ratio is between 0.5% and 1% the heartbeat period is adjusted to 20 seconds. If the buffering ratio is equal to or greater than 1%, the heartbeat period is adjusted to 10 seconds.

Content distribution monitor 102 can also direct client 172 to increase or decrease the heartbeat frequency based on factors such as the current load on the content distribution monitor (or applicable component thereof). For example, if the load on content distribution monitor 102 exceeds a predefined threshold, clients are instructed to increase the heartbeat interval and/or reduce the detail of the information sent in a heartbeat.

As one example, content distribution monitor 102 can employ a rule that, if its resource utilization exceeds 80%, clients are instructed to double their respective heartbeat intervals. Monitor 102 can also selectively send such instructions to clients. For example, clients with periods under 30 seconds can be instructed to double their periods, while clients with periods above 30 seconds are not sent such instructions.

As another example, when monitor 102's utilization exceeds 80%, it can instruct clients to reduce heartbeat size by sending less information. This can be less detailed information (e.g., sending only the number of resource switches, instead of the detailed information about these switches), or can include aggregate measurements about certain events instead of sending individual details. If the client is already aggregating certain information, the measurements can be aggregated over a longer time interval (e.g., 10 seconds instead of 2 seconds). Additional examples follow.

Instead of sending every buffering event, send the aggregate time the session spent in the buffering state since the last heartbeat.

Instead of sending every measurement of the rendering rate, send the average over the entire heartbeat interval.

Instead of sending information about a switch event at the time when the switch event happened, store the event on the client and send it with the next heartbeat.

Inferring State Information

In some cases, the monitoring and reporting functionality described herein as being provided by a client is wholly integrated with the player itself. For example, Studio might make available a custom player application that can be installed on a client by users that want to watch Studio films. The customer player includes all of the necessary logic to provide complete heartbeat information, and also includes logic for communicating with the content distribution monitor, adjusting the frequency with which status information is communicated to the content distribution monitor, and other functionality described herein. However, depending on the player deployed on a client, not all of the example heartbeat information described above may be directly accessible for collection and transmission. For example, the Flash player plugin used by Alice to view the soccer game does not provide direct access to its state. Using the techniques described herein, content distribution monitor 102 can nonetheless be provided with the state of the player.

In some embodiments, a monitoring module is embedded in the content player deployed on the client. For example, when Alice directs her browser to XYZ's website, a custom Flash-based player is loaded. The custom player has direct access to the API of the video plugin. Possible states of the player are as follows:

**Playing:** The player's buffer is not empty, and the player renders frames on the screen.

**Buffering:** The player's buffer is either empty or is less than a predefined threshold, making it impossible for the player to render new frames.

**Joining:** The player has just connected and has not yet accumulated enough data in its buffer to start rendering frames.

**Pause:** The player is paused, usually, as a result of a user action. During the pause state, no new frames are rendered even if the buffer is full.

**Sleep:** The client is suspended.

In the case of the aforementioned Flash player, NetStatus events provided by the player can be used to infer the state, as described below:

**Buffering:** When one of a NetStatus.Buffer.Empty, NetStream.Play.Start, or NetStream.Play.Reset events is received, an inference is made that the player is buffering data and the video is not playing. For live RTMP streams, the mentioned events might not be available. In that case, if the playhead time has stopped moving and there is no data in the content player buffer, an inference is made that the player is buffering data and the video is not playing.

**Playing:** When a NetStatus.Buffer.Full event is received, an inference is made that the player is playing video. If the player is currently in a Paused state and the playhead time starts to move, an inference is made that the player is playing video.

**Paused:** If the playhead time does not progress for 1.6 seconds, an inference is made that the player is paused.

**Sleep:** If the amount of time that has elapsed between two firings of a periodic one second timer is greater than 30 seconds, an inference is made that the client has been in a sleep mode since the last firing of the timer.

**Stopped:** For HTTP streams, when a NetStatus.Play.Stop event is received, an inference is made that the stream has ended. For RTMP streams, when a NetStatus.Play.Complete event is received, an inference is made that the stream has ended. When a NetStatus.Play.Stop is received, an inference is made that the download has finished, but the stream is still playing.

**Error:** If the stream stays in stopped state for fifteen seconds at the start of a session, an inference is made that it has failed to connect.

In addition to the player states described above, an additional state can also be inferred:

**Zombie state:** An inference is made that the player is in a zombie state (a non-operating state with no user actively using it) if it persists in a non-playing state for more than a specified time. Once the player is in a zombie state, the transmission of heartbeats to content distribution monitor 102 is halted. Reporting will resume if the player returns to a playing state. Non-playing states include: "Buffering," "Paused," "Stopped," and "Error."

In addition to the player states, the monitoring module monitors various metrics, including the number of bytes downloaded/streamed since the session started, the bitrate of the stream, and the rendering rate.

13

Some players, such as the Flash 9 plugin, do not provide a direct API from which the number of bytes streamed/downloaded or the stream bitrate. Nonetheless, a variety of techniques can be used to obtain or estimate the bitrate. As one example, the developer of the player might expose the

Another technique for obtaining a bitrate is to use metadata associated with the content. For example, suppose metadata is made available through a configuration file associated to the content. The configuration resides on an origin server such as origin server 198 and also includes other information such as the title of the content, its genre, and a list of CDNs where the content is available.

Yet another technique for obtaining a bitrate is to examine the content URI. For example, a movie might be accessible at the URI: `http://www.CDN-C.com/Studio/JanetJrc300` Kbps. In this case, the bitrate of the content is likely to be 300 Kbps.

An estimation of the stream bitrate can also be made. The estimation, along with the amount of time the player was in the playing state, and the size of the buffer, are used to estimate the number of bytes downloaded/streamed. As one example, suppose a player streamed data at 400 Kbps, was in a playing state for 315 seconds, and at the time of taking the measurement the buffer size contained data for playing another 15 seconds. The total number of bytes downloaded is:  $400 \text{ Kbps} \times (315 \text{ seconds} + 15 \text{ seconds}) = 16 \text{ MB}$ .

In the event that the bitrate is available through multiple of the above techniques, in various embodiments each applicable approach is used, in a specified order, until a valid bitrate (e.g., greater than 0 and less than 10,000 kbps) is obtained.

The number of bytes downloaded can be estimated as  $(\text{totalPlayingTime} + \text{bufferLength}) \times \text{bitrate}$ . If multiple bitrates are used, in some embodiments the bytes downloaded is estimated as the sum of  $\text{totalPlayingTime}[\text{bitrate}] \times \text{bit-rate} + \text{bufferLength} \times \text{currentBitrate}$ .

Rendering quality is another example metric that can be obtained and provided to content distribution coordinator 102. Rendering quality is the ratio between the frames per second (FPS) rendered by the player and the FPS at which the stream was encoded.

Some players do not directly provide the information needed to compute the rendering quality. For example, Flash 9 has an API that exposes rendered frames per second, but not encoded frames per second. One technique to compute the encoded FPS, in such a scenario, is as follows. The rendering FPS is measured and an assumption is made that the encoded FPS is the maximum FPS observed over the course of the session. Another technique is to estimate the encoded FPS as the maximum of a three second sliding average of rendered FPS sampled five times a second.

In order to ensure accurate accounting of the player status over the course of the session, it is important to determine the time intervals when the system is in sleep or suspend mode. These intervals can be detected using notification APIs provided by the system for this purpose. For systems that do not provide such API (e.g., Flash Player, or Silverlight), the intervals can be estimated by setting a timer to fire at periodic intervals (e.g., every 1 second). At each firing of the timer the current system time is recorded, and the time elapsed since the last firing is computed. If the time elapsed is greater than a given threshold (e.g., 10 seconds), an inference is made that the given elapsed time interval was spent in sleep or suspend mode.

14

#### Inferring Player Capability Information

The following are example player capabilities that impact the quality of the video experience: processing speed, video rendering capabilities, amount of memory available, and amount of bandwidth available.

**Processing speed:** The monitoring module can use the API of the content player, or the underlying platform to read the processing speed of the player. For example, if the content player is implemented on the Silverlight platform the monitor can use the `Environment.ProcessorCount` to obtain the count of the processors available.

For platforms whose API does not provide direct reading of processing speed capabilities, the monitoring module can derive it by using a timer to measure the time required to perform a fixed CPU-intensive computation.

**Video rendering capabilities:** Some content players have the ability to use hardware acceleration to render video. The monitoring module can determine the speed at which the content player can render video, either by using the API of the content player or the underlying platform. For example, a monitoring module using the Flash platform can use the `flash.system.Capabilities.ScreenResolutionX` to determine the screen resolution.

**Available memory:** The amount of available memory has a direct influence on how well the content player can support rendering video while performing the computation required for monitoring and other user-interaction tasks. The monitoring module can obtain the amount of available memory by using the API of the underlying platform. For example, a monitoring module using the Flash platform can use the `API flash.system.System.totalMemory` to determine how much memory is available.

**Available download bandwidth:** The available download bandwidth has a direct impact on the quality of the video stream that can be downloaded and played. The monitoring module can infer the available bandwidth by measuring the time interval it takes to download a fixed size file from the server. Alternatively, the underlying platform can provide API that can be used to determine the download bandwidth. For example, a player using the Flash platform can use the `API flash.net.NetStreamInfo.currentBytesPerSecond`.

FIG. 4 is an illustration of an environment in which status information is received and processed. In various embodiments, the services provided by content distribution monitor 102 are implemented across a scalable infrastructure, particularly in embodiments where telemetry data is received from all clients. In the example shown, the elements contained within dashed region 402 collectively provide the functionality of content distribution monitor 102. Each of the layers (e.g., dispatcher layer 420) is horizontally scalable and their respective components can be implemented on standard commercially available server hardware (e.g., having a multi-core processor, 4G+ of RAM, and Gigabit network interface adaptors) running a typical server-class operating system (e.g., Linux).

Clients 404-412 each include a monitoring module that collects status information. When the monitoring module on a client is activated, the client is mapped to a dispatcher server. As one example, when the monitoring module starts, it reads a configuration file that includes a list of dispatcher servers. The monitoring module selects a dispatcher server at random from the list.

A dispatcher server (414) includes two conceptual modules. The first module implements a communication interface for receiving status information from clients. In some

embodiments the module is implemented using an off-the-shelf web server, and allows clients to connect over the HTTP protocol (and also allows clients to securely communicate via SSL). Data received by the first module is passed to the second module. The second module normalizes the data (to a format suitable for further processing) and passes the normalized data to a real-time stream processing component (416).

The real-time stream processing (RSP) layer includes an optimized software component that processes the telemetry data that it receives from the dispatcher in real-time. A dispatcher sends all heartbeats belonging to the same session to the same RSP component.

In some embodiments the RSP component is implemented as a continuously running service that reads and processes the telemetry data received from dispatchers via the network over TCP. The telemetry data stream comprises individual records, each of which represents the telemetry data sent by the monitoring module. The RSP component reads network data one record at a time and parses the data into a local data representation. The data received by the RSP component can be stored as in-memory hash tables of records allowing fast execution, and very high throughputs. Since the RSP component does not require old information, it can periodically purge the in-memory hash tables and increase scalability accordingly. In other embodiments, optimized in-memory databases are used.

A mapping function to map heartbeats having a session identifier "ID" to a particular RSP component "i" as follows:

$$i = \text{hash}(\text{ID}) \bmod m,$$

where hash( ) is a hash function and "m" is the total number of RSP components.

Once an RSP component parses the data records, it performs two main tasks.

First, it performs data filtering. A filter is a logical expression and is installed at each RSP component instance. As one example, the following filter would identify viewers located in San Francisco, connected to ISP SP1, streaming from CDN A, one of two particular shows:

```
(city="San Francisco" AND ISP="SP1" AND
CDN "CDN A" AND ((show "NewsAt10") OR
(show "SundayMagazine")))
```

For each message of incoming telemetry data, the (key, value) pairs in the record are matched against the filter. If the filter is matched, the data is associated with the filter.

The second task performed is to compute snapshots and on-line statistics over the telemetry data matching each filter. One example of a snapshot is the number of players that are in a particular state (e.g., "playing"). The RSP component generates sequences of these snapshots (e.g., one every second). Examples of statistics computed by the RSP component include: the average number of bytes played over all video streams matching a filter over a given time interval (e.g., 10 seconds) and the minimum frames per second experienced by a stream matching a filter over a time interval. Snapshots and statistics are updated continuously, from new telemetry data received from clients.

The RSP component provides its computed snapshots and statistics to a real-time global aggregation component (418). The real-time global aggregation (RTGA) component aggregates the information provided by the RSP component for each filter specified by a user (described in more detail below).

As explained above, each RSP component (416) receives (via a dispatcher) telemetry data from a subset of the monitoring modules and calculates snapshots and statistics

for all filters. Each RGA component instance is in turn responsible for a subset of the filters. Based on the identifier of the filter, all RSP components send data for that filter to a single RGA component. The RGA component combines the data from all RSP components for the filters that it is responsible for, resulting in a global snapshot and statistics based on information from all monitoring modules. Examples of aggregation operations performed by an RGA component include: counting the total number of viewers that match a particular filter, determining the percentage of viewers matching a given filter that are in buffering state, the joint time distribution experienced by viewers joining a given stream, the current number of viewers in a given city, the rank of the most popular live events, and so on.

In some embodiments an RGA component's functionality is implemented as a continuously running service. It reads records sent by the RSPs asynchronously, thus achieving a high throughput. The RGA component stores the records it receives in in-memory hash tables allowing optimized access for real-time processing. Old information is periodically purged from the hash tables to improve efficiency.

As shown in FIG. 4, gateway 422 provides a web service API 424 for accessing data. Through the API, RGAs data is available on a per-filter basis. In addition, it also exposes APIs to edit and install new filters and aggregation functions. In some embodiments gateway 422 is implemented using an off-the shelf web server (such as Apache) with customized code to handle the various web service API calls. The handlers return data for a given web API call in various data formats including XML, JSON, SOAP, and HTML, as applicable. The handlers serve as middleware for querying and interactively controlling the RSPs and RGAs.

Gateway 422 also provides access controls and persistently stores the information regarding the API requests, data access and presentation policies, and filter descriptions. The information is maintained in a persistent database, such as MySQL or Oracle database.

Automatically Detecting and Resolving Content Distribution Problems

As explained above, content distribution monitor 102 aggregates telemetry information from all clients, processes the information, and, as will be explained in more detail below, allows users (e.g., via a user interface) to view the multi-dimensional results in realtime. Examples of dimensions include:

- Client properties: Including browser type and version, player type and version, operating system, CPU speed, connection type, IP address, geographic location, language, autonomous system, and ISP.

- Content properties: Including category, title, show, episode number, duration, encoding format, encoding quality, and language.

- Content source properties: Including CDN, data center, and clients served.

- User properties: Including whether the user (of the client) is a premium or free user, and returning or first-time visitor.

Accordingly, using the techniques herein, one is able to track, in realtime, granular information such as the quality experienced by viewers located in Denver, connected to SP2, and streaming videos from CDN B, using Flash Player 10. Further, by aggregating and correlating the data it receives from all clients, content distribution monitor 102 exposes, in realtime, the performance of content sources and network providers. (ISPs). Problems in content delivery can be automatically detected by examining the results.

17

FIG. 5 illustrates an example of a process for detecting a problem in a content distribution. In some embodiments the process shown in FIG. 5 is performed by content distribution monitor 102. The process begins when information associated with a first content player and a second content player is received from a first client (502) and second client (504), respectively. For example, at 502, telemetry information is received from client 404 by dispatcher 414. At 504, telemetry information is received from client 406 at dispatcher 426. At 506, the received information is aggregated. For example, at 506, distributed processing of the received data is performed by the RSP and RGA layers. Finally, at 508, a determination is made from the aggregate information that a content distribution problem is indicated. For example, at 508, gateway 422 determines that the CDN from which client 404 is obtaining content is experiencing a problem. Specific examples of analysis and diagnosis will now be given.

#### Example

##### Diagnosing Client Problems (Single Client, Single CDN)

Suppose a content owner, such as Studio-A, distributes its content via a single CDN, such as CDN C. Suppose that the content is encoded at multiple bitrates B1 and B2 where  $B1 < B2$ . If a client A is able to sustain the download rate, but detects problems rendering the frames, an inference can be made that the client A likely has problems with its CPU utilization. An example of performing such an inference follows.

Observe the client buffer size over a period of time T (e.g., T 10 seconds) when the client is streaming at bitrate B2. Observe the rendering quality measured using dropped frames reported by the player. If the client buffer size is greater than B\_threshold (e.g.,  $B\_threshold = 0.5 * B\_max$ ) at all times during the period of observation, and the rendering quality is less than R\_threshold\_1 (say,  $R\_threshold\_1 = 0.6$ ), then a conclusion can be made that the client cannot sustain displaying bit-rate B2 due to CPU issues on the client.

Perform the same observation with the client playing at bit-rate B1. If the client buffer size is greater than B\_threshold (e.g.,  $B\_threshold = 0.5 * B\_max$ ) at all times during the period of observation, and the rendering quality is greater than R\_threshold\_2 (e.g.,  $R\_threshold\_2 = 0.75$ ), then a conclusion can be made that the client can sustain displaying bit-rate B1.

A numerical example is now provided. Consider a setting where a client streams a 1 Mbps stream for 10 seconds with buffer size varies between 9 and 10 seconds. Suppose that B\_max is 10 seconds and B\_threshold is 5 seconds. Assume that the rendering quality at 1 Mbps is 0.55 over this interval. When the client plays a 500 kbps stream instead, the buffer size is in the same range (9-10 seconds), but the rendering quality is 0.9 instead. Then, assuming that R\_threshold\_1 0.6 and R\_threshold\_1 0.75, a conclusion can be made that the client can sustain 500 kbps but not 1 Mbps.

#### Example

##### Diagnosing CDN Problems (Single Client, Two CDNs)

Suppose a content owner, such as Studio-A, distributes its content via two CDNs, such as CDN D and CDN E. If a

18

client A is able to sustain the download rate from CDN D, but not CDN E, then the client can conclude that CDN E has problems streaming to client A. An example of performing such an inference follows.

5 Compute an aggregate quality metric Q using various measurements from the client over a time T (e.g., T 60 seconds). Examples of measurements which can be directly used as the metric Q across a group of clients over a time T include:

10 buffering ratio: The total number of seconds the client experienced buffering divided by the total playing time of the clients during interval of time T.

join time: The average joining time over all clients during time T over a group of clients.

15 join failures: The fraction of cases (across all attempts) where the join failed during time T.

If the quality Q is greater than a threshold T1 (e.g.,  $T1 = 0.95$ ) for CDN D, but lower than another threshold T2 (e.g.,  $T2 = 0.50$ ) for CDN E, then a conclusion can be made that CDN E has problems streaming to client A.

One example remedial action that can be taken at the client is to select CDN D for streaming the current stream, or use it first for future streams.

25 A numerical example is provided. Consider a setting where a client M streams from CDN D and experiences a buffering ratio of 0.5 over K1 attempts and time T1. Also, client M streams from CDN E, and over K2 attempts and time T2 experiences a buffering ratio of 0.02. In this case, given that K1 and K2 are above a threshold (e.g., both greater than 2), and T1 and T2 are above a threshold (e.g., both are greater than 10 minutes), a conclusion can be made that CDN D has problems streaming to the client. A further refinement can be performed by computing the buffering ratio for each of the K1 (or K2) streaming attempts from the CDNs to the client. In this case, it can be stated that CDN D has problems streaming to client M if more than 75% of the attempts from CDN D (out of the K1 attempts) have buffering ratio greater than 0.5, and more than 75% of the attempts from CDN D have buffering ratio less than 0.02.

#### Example

##### Two Clients in the Same Location Using Different ISPs Both Streaming from a Single CDN

Suppose a content owner, such as Studio XYZ, distributes its content via a single CDN, such as CDN C. The CDN comprises several geographically distributed servers. As clients request content, they are assigned to a particular server based on their geographical location. If many clients at the same geographical location, but using different ISPs, experience quality problems, an inference can be made that the CDN servicing that geographical location is having a problem. Note that since XYZ also uses the services of CDN C, the experiences of clients obtaining XYZ's content from the same location can be included in the analysis. One way this analysis can be implemented on content distribution monitor 102 is as follows.

60 First, obtain location and network information for all clients. Examples of location information include city, state, country, DMZ code, and geographic coordinates. The location information can be obtained in a variety of ways, such as by providing the client's IP address to a geo-location service such as Quova. Also obtain network information such as Autonomous System Number (ASN) information and ISP information for the client.

19

Second, group clients based on their ASN and one geographic attribute, such as DMZ code, city, or state. Let  $G(\text{ASN1}, \text{Geo1})$  denote the set of clients connected to a specific ASN, ASN1, which have the same geographic attribute value, Geo1.

Third, compute an aggregate quality  $Q$  for each group  $G(\text{ASN}, \text{Geo})$  over last  $T$  seconds. Denote this aggregate quality by  $Q(G(\text{ASN}, \text{Geo}))$ . Examples of quality metrics are provided above.

Finally, for each Geo attribute value Geo1, check whether there are at least two group of clients connected via different ASNs. If the aggregate quality of a certain fraction  $F$  of the groups is greater than a threshold  $\text{BufferingThreshold1}$ , then conclude that the CDN is experiencing problems serving the clients at location Geo1. In some embodiments a minimum group size is defined so that only groups having a number of clients larger than a given threshold are considered. As one example, the fraction  $F$  is selected between 0.75 and 1.

A numerical example is now provided. Consider a system setting where  $\text{BufferingThreshold1}=0.1$ ,  $\text{MinGroupSize}=200$ , and  $F=1.0$ . Based on an analysis of their IP addresses, 20,000 viewers are determined to be located in San Francisco, connected to CDN C, and watching a new release by Studio. Assume 5,000 clients are connected to ASN1 belonging to SP1; 5,000 clients are connected to ASN2 belonging to SP2; 9,900 clients are connected to ASN3 belonging to SP3; and 100 clients are connected to ASN4 which belongs to a wireless company. Four groups exist:  $G(\text{SF}, \text{ASN1})$ ,  $G(\text{SF}, \text{ASN2})$ ,  $G(\text{SF}, \text{ASN4})$ , and  $G(\text{SF}, \text{ASN4})$ , respectively.

Each client reports for each session (a) the total time it spent in buffering during every 10 second interval; and (b) the total playing time during the same 10 second interval. For example, suppose a user Charlie has the IP address 1.2.3.4 which belong to ASN1, and watches the stream with the URL: `rtmp://www.CDN.C.example.com/12346/video.flv`. Charlie reports in a particular heartbeat that he has experienced 1 second of buffering and 9 seconds of playing time. Another user, Bob, has the IP address 2.3.4.5 which belongs to ASN3, and watches the stream with the same URL. Bob reports in a particular heartbeat that he has experienced 0.5 seconds of buffering and 5 seconds of playing time. Bob's player spent the remaining 5 seconds in pause state. The pause time is not used in the computation of the buffering ratio.

Content distribution monitor 102 computes over each interval of time  $T$  5 minutes, the aggregate average buffering ratio, defined as the (1) the total buffering time experienced by all sessions in a group during  $T$  divided by (2) the total playing time over all sessions in the group during the same time interval  $T$ . The aggregate quality ratio for each of the four groups is as follows:

$Q(G(\text{SF}, \text{ASN1}))=0.14$ ,  $Q(G(\text{SF}, \text{ASN2}))=0.21$ ,  $Q(G(\text{SF}, \text{ASN3}))=0.18$ ,  $Q(G(\text{SF}, \text{ASN4}))=0.04$ .

Since the number of clients in  $G(\text{SF}, \text{ASN4})$  is less than  $\text{MinGroupSize}$ , the diagnosis analysis ignores this group. Also, since the buffering ratio of all remaining groups is greater than  $\text{BufferingThreshold1}$ , a conclusion is made that CDN C experiences quality issues serving clients in San Francisco.

Suppose the above example is modified such that ASN4 has 1000 clients, and  $Q(G(\text{SF}, \text{ASN1}))=0.02$ ,  $Q(G(\text{SF}, \text{ASN2}))=0.21$ ,  $Q(G(\text{SF}, \text{ASN3}))=0.18$ ,  $Q(G(\text{SF}, \text{ASN4}))=0.04$ . In this case, an inference can be made that CDN C does not experience quality issues serving clients in San Francisco in general. The problem can then be narrowed down

20

further and smaller or different sets of clients investigated, such as clients that are connected to ASN1 served by CDN C.

#### Example

##### Inferring Problems in a Particular ASN (Single CDN)

Suppose a content owner, such as Studio XYZ, distributes its content via a single CDN, such as CDN C. If the quality of clients at the same location but using different ISPs is very different, an inference can be made that the quality problems experienced by the clients are due to the ISP (and not due to the CDN). One way this analysis can be implemented on content distribution monitor 102 is as follows.

Perform the first three steps of the previous diagnosis (describing the diagnosis of CDN quality issues). As a final step, let  $G(\text{ASN1}, \text{Geo1})$  be a group experiencing low quality. If there is at least another group,  $G(\text{ASN2}, \text{Geo1})$  at the same location experiencing high quality, then conclude that ASN1 is the cause of quality issues. In particular, a conclusion that ASN2 has quality issues is reached if, and only if,  $Q(G(\text{ASN2}, \text{Geo1})) - Q(G(\text{ASN1}, \text{Geo1})) > \text{BufferingRatioDiscrepancy}$ .

A numerical example is now provided. Consider a system setting where  $\text{BufferingRatioDiscrepancy}=0.1$ . Suppose 5,000 viewers are determined to be located in New York, connected to ASN1, and streaming a baseball game from CDN C. Another 4,000 viewers in New York are connected to ASN2 and streaming the same game from CDN C.

Each client reports for each session (a) the total time it spent in buffering during every 10 second interval; and (b) the total playing time during the same 10 second interval. The aggregate quality ratio for each group as computed by content distribution monitor 102 is:  $Q(G(\text{NY}, \text{ASN1}))=0.14$ , and  $Q(G(\text{NY}, \text{ASN2}))=0.03$ , respectively.

Since  $Q(G(\text{SF}, \text{ASN1})) - Q(G(\text{SF}, \text{ASN2})) > \text{BufferRatioDiscrepancy}$ , the quality issues are pinpointed as being ASN1.

#### Example

##### Inferring Problems with a Particular CDN (Multiple CDNs)

Suppose a content owner, such as XYZ, distributes its content via multiple CDNs (e.g., CDN A, CDN B, and CDN C). If clients at the same location and using the same ISP experience significantly different quality when connected to different CDNs, the difference in quality can be attributed to CDNs. One way this analysis can be implemented on content distribution monitor 102 is as follows.

First, obtain location and network information for all clients. Second, classify each client based on its ASN, the CDN it gets data from, and one of its geographic attributes (e.g., DMZ code, city, or state). Let  $G(\text{ASN}_i, \text{CDN}_i, \text{Geo}_i)$  denote the set of clients connected to a specific ASN and a geographic region (e.g., for  $i=1, 2, 3, 4, 5$ ), getting data from same CDN, CDN<sub>*i*</sub>. One way of determining the particular CDN from which a client is receiving content is to extract it from the URL used by the client to obtain the content.

Third, compute an aggregate quality  $Q$  for each group  $G(\text{ASN}, \text{CDN}, \text{Geo})$  over the last  $T$  seconds. Denote this aggregate quality by  $Q(G(\text{ASN}, \text{CDN}, \text{Geo}))$ .

Finally, check whether there are at least  $K$  pairs of groups that share the same ASN and geo-location, but get their data

from different CDNs, and which experience different quality. In particular, let  $G(\text{ASNi}, \text{CDNi}, \text{Geo1})$  and  $G(\text{ASNj}, \text{CDNj}, \text{Geo1})$  be one such pair of groups. Then if  $Q(G(\text{ASNi}, \text{CDN1}, \text{Geo1})) - Q(G(\text{ASNi}, \text{CDN2}, \text{Geo1})) > \text{QualityThreshold2}$  for greater than a fraction  $F$  of the  $K$  pairs, a conclusion is made that CDN2 has problems serving clients in general across a large set of regions. Alternatively, if this happens for smaller than a fraction  $F$  but a non-zero set, then a conclusion can be made that CDN2 has problems serving those (ASN, Geo) combinations for which the difference in quality exceeds the quality threshold.

In some embodiments, a minimum group size is defined so that only groups having a number of clients larger than a given threshold are considered.

A numerical example is provided. Content owner Charlie's Studio uses CDNs A and B for delivering content. At some point, there are 1100 streams from ISP-1 in San Francisco, 1200 streams from ISP-2 in Los Angeles, 1500 streams from ISP-3 in New York, 1050 streams from ISP-4 in Boston, 600 streams from ISP-1 in Denver, with exactly half the users in each group streaming from CDN A and the other half from CDN B. A minimum group size of 500 users streaming from a particular (ASN, CDN, Geo) group is used, and hence there are four groups to consider. A requirement of at least four pairs ( $K=4$ ) exists. For CDN A, the buffering ratios for users in San Francisco, Los Angeles, and Boston is 0.3, and for users in Denver it is 0.1. For CDN B, the buffering ratio of all user groups is 0.05. The quality threshold to separate good and bad quality is 0.1, and the fraction  $F$  required is 75%. This condition is satisfied, since three out of the four groups have a quality difference between CDN A and CDN B as 0.25. Hence, a conclusion is made that CDN A is having problems streaming to users.

### Example

#### Multiple CDNs with CDN Optimization

Suppose a content owner, such as XYZ, distributes its content via multiple CDNs (e.g., CDN A, CDN B, and CDN C). Assume that if a client connected to CDN A experiences buffering beyond a threshold, it switches to another CDN (and stays there for at least a threshold interval of time). Based on the amount of switching observed from a CDN, the relative extent of quality problems the CDNs have can be quantified. One way this analysis can be implemented on content distribution monitor 102 is as follows.

For each  $\text{CDNi}$  for a period  $T$ , determine the number of clients  $N_i$  that start with  $\text{CDNi}$ , and the number of clients  $M_i$  that start with  $\text{CDNi}$  and switch away from  $\text{CDNi}$  (i.e. count only the first switch of a client). Compute the switch fraction  $\text{SW}_i = M_i/N_i$  for all CDNs. If  $(\text{SW}_i - \text{avg}(\text{SW})) > \text{SW\_Threshold}$ . A conclusion can be made that  $\text{CDNi}$  has quality problems in general. Similar analysis can also be performed with respect to a geographic region or an ASN restriction.

A numerical example is provided. A content owner Charlie's Studios uses CDN A and CDN B to deliver content. Over a 10-minute period, 10,000 users watch content starting streaming from CDN A, and in the course of viewing, 1000 users switch to CDN B based on client-based quality detection algorithm. In the same period, 10,000 users watch the same content starting streaming from CDN B, and out of these, 2,500 switch to CDN B. The switch fraction for CDN A is 0.1 and that for CDN B is 0.2. Switching threshold  $\text{SW\_Threshold}$  is 0.1, and a conclusion is made that CDN A has quality problems in general.

Other content distribution problems can also be detected using the techniques described herein. For example, a problem with a content item itself (e.g., a particular movie) can be inferred if several clients, using different CDNs and different ISPs, experience quality issues with the same content. As another example, a problem with ad server 150 can be inferred if several clients report excessively long advertisement load times or timeouts.

A numerical example is provided. A content owner publishes 100 clips daily using two CDNs: CDN A and CDN B. One of the clips, clip X, has 100% join failures on both CDNs. Using this information, an inference can be made that there is a problem in publishing clip X. Now assume that there is 100% join failure for clients that join CDN A, but <2% of the users have problems with CDN B. Then an inference can be made that CDN A's publishing path has problems for that clip (but CDN B's does not).

Other content distribution problems can also be detected using the techniques described herein. For example, a problem with a content item itself (e.g., a particular movie) can be inferred if several clients, using different CDNs and different ISPs, experience quality issues with the same content. As another example, a problem with ad server 150 can be inferred if several clients report excessively long advertisement load times or timeouts.

In various embodiments, when problems are detected by content distribution monitor 102, corrective actions are automatically taken, either by a control module included in content distribution monitor 102, or by a separate control system, configured to work in conjunction with content distribution monitor 102.

Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

What is claimed is:

1. A system, comprising:

a set of one or more processors configured to, using a set of one or more interfaces:

receive, from a video player, a request for a manifest file, wherein the request includes location information associated with the manifest file;

fetch, from a remote source, the manifest file using the location information included in the request;

rewrite the fetched manifest file, wherein the rewriting of the fetched manifest file is based at least in part on global information aggregated from a plurality of video players; and

return the rewritten manifest file to the video player, wherein the video player is configured to obtain video data according to the rewritten manifest file; and

a memory coupled to the set of one or more processors and configured to provide the set of one or more processors with instructions.

2. The system of claim 1, wherein the set of one or more processors is further configured to select at least one of a bitrate and a content delivery network.

3. The system of claim 2, wherein the selection is based at least in part on at least one of a traffic and a business policy.

4. The system of claim 2, wherein the at least one of the bitrate and the content delivery network is selected based at least in part on an evaluation of the global information aggregated from a plurality of video players.

23

5. The system of claim 2, wherein the fetched manifest file is rewritten to include the selected at least one of the bitrate and the content delivery network.

6. The system of claim 2, wherein the selection is based at least in part on the request.

7. The system of claim 6, wherein the request includes a device type associated with the video player.

8. The system of claim 6, wherein the request includes a content type associated with the video player.

9. The system of claim 6, wherein the request includes historical information associated with the video player.

10. The system of claim 1, wherein the rewriting includes overwriting at least a portion of the fetched manifest file.

11. The system of claim 1, wherein the rewritten manifest file, when followed by the video player, causes a bitrate switch.

12. The system of claim 1, wherein the rewritten manifest file, when followed by the video player, causes a content delivery network switch.

13. The system of claim 1, wherein the fetched manifest file comprises a static manifest file.

14. The system of claim 1, wherein the location information comprises a universal resource locator associated with the manifest file.

15. A method, comprising:

receiving, from a video player, request for a manifest file, wherein the request includes location information associated with the manifest file;

fetching, from a remote source, the manifest file using the location information included in the request;

rewriting, using a set of one or more processors, the fetched manifest file, wherein the rewriting of the

24

fetched manifest file is based at least in part on global information aggregated from a plurality of video players; and

returning the rewritten manifest file to the video player, wherein the video player is configured to obtain video data according to the rewritten manifest file.

16. The method of claim 15, further comprising selecting at least one of a bitrate and a content delivery network.

17. The method of claim 16, wherein the fetched manifest file is rewritten to include the selected at least one of the bitrate and the content delivery network.

18. The method of claim 15, wherein the rewritten manifest file, when followed by the video player, causes a bitrate switch.

19. The method of claim 15, wherein the rewritten manifest file, when followed by the video player, causes a content delivery network switch.

20. A computer program product embodied in a non-transitory tangible computer readable storage medium and comprising computer instructions for:

receiving, from a video player, a request for a manifest file, wherein the request includes location information associated with the manifest file;

fetching, from a remote source, the manifest file using the location information included in the request;

rewriting the fetched manifest file, wherein the rewriting of the fetched manifest file is based at least in part on global information aggregated from a plurality of video players; and

returning the rewritten manifest file to the video player, wherein the video player is configured to obtain video data according to the rewritten manifest file.

\* \* \* \* \*