

**(12) STANDARD PATENT**  
**(19) AUSTRALIAN PATENT OFFICE**

(11) Application No. **AU 2012207151 B2**

(54) Title  
**Variable bit video streams for adaptive streaming**

(51) International Patent Classification(s)  
**G06F 15/16** (2006.01)

(21) Application No: **2012207151** (22) Date of Filing: **2012.01.20**

(87) WIPO No: **WO12/100215**

(30) Priority Data

(31) Number	(32) Date	(33) Country
<b>13/011,242</b>	<b>2011.01.21</b>	<b>US</b>

(43) Publication Date: **2012.07.26**

(44) Accepted Journal Date: **2015.07.30**

(71) Applicant(s)  
**Netflix, Inc.**

(72) Inventor(s)  
**Hunt, Neil D.**

(74) Agent / Attorney  
**Freehills Patent Attorneys, Level 43 101 Collins Street, Melbourne, VIC, 3000**

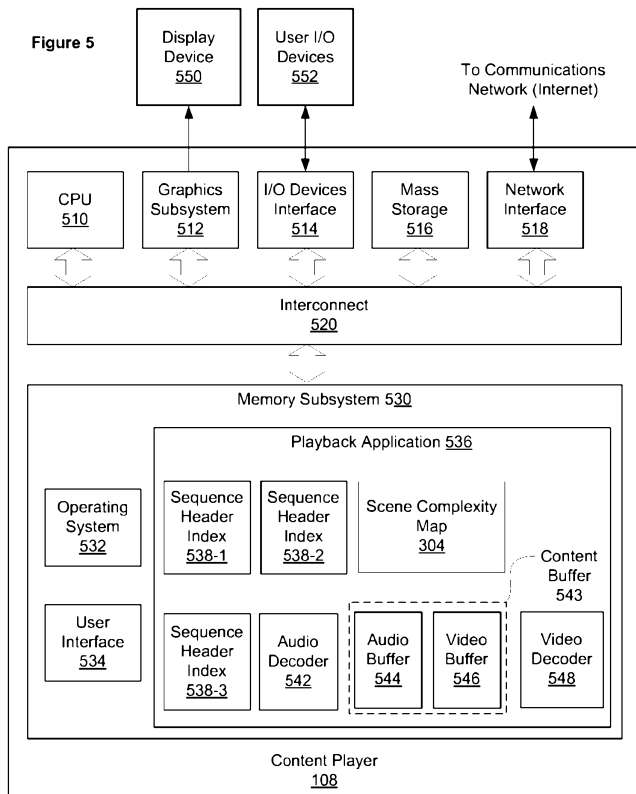
(56) Related Art  
**US 2009/0300204 A1 (ZHANG et al.) 03 December 2009**



- (51) International Patent Classification:  
*G06F 15/16* (2006.01)
- (52) International Application Number:  
PCT/US2012/022103
- (22) International Filing Date:  
20 January 2012 (20.01.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
13/011,242 21 January 2011 (21.01.2011) US
- (71) Applicant (for all designated States except US): **NET-FLIX, INC.** [US/US]; 100 Winchester Circle, Los Gatos, California 95032 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **HUNT, Neil, D.** [US/US]; 1190 Payne Drive, Los Altos, California 94024 (US).
- (74) Agents: **MCCLELLAN, Gero, G.** et al.; Patterson & Sheridan, LLP, 3040 Post Oak Blvd., Suite 1500, Houston, Texas 77056-6582 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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.

[Continued on next page]

(54) Title: VARIABLE BIT VIDEO STREAMS FOR ADAPTIVE STREAMING



(57) Abstract: One embodiment of the present invention sets forth a technique for adapting playback bit rate in a content delivery system based on scene complexity of the video content as well as network conditions and other performance factors. A scene complexity map of the video content indicates the complexity levels of different scenes within the video content. Using the scene complexity map, a content player may download scenes of lower scene complexity levels from video streams encoded to lower bit rates to manage the bandwidth consumed to download the video content and allow for the downloading of scenes of higher complexity levels from video streams encoded to higher bit rates.

WO 2012/100215 A1



**(84) Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

**Published:**

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

## VARIABLE BIT VIDEO STREAMS FOR ADAPTIVE STREAMING

### CROSS-REFERENCE TO RELATED APPLICATION

5 [0001] This application claims benefit of United States patent application serial number 13/011,242, filed January 21, 2011, which is a continuation in part of U.S. Patent Application Serial No. 12/961,375, entitled "VARIABLE BIT VIDEO STREAMS FOR ADAPTIVE STREAMING," filed December 6, 2010. This related patent application is herein incorporated by reference in its entirety.

### 10 BACKGROUND OF THE INVENTION

#### Field of the Invention

[0002] Embodiments of the present invention relate generally to digital media and, more specifically, to variable bit video streams for adaptive streaming.

#### Description of the Related Art

15 [0003] Digital content distribution systems conventionally include a content server, a content player, and a communications network connecting the content server to the content player. The content server is configured to store digital content files, which can be downloaded from the content server to the content player. Each digital content file corresponds to a specific identifying title, such as "Gone with the Wind," which is  
20 familiar to a user. The digital content file typically includes sequential content data, organized according to playback chronology, and may comprise audio data, video data, or a combination thereof.

[0004] The content player is configured to download and play a digital content file, in response to a user request selecting the title for playback. The process of playing the  
25 digital content file includes decoding and rendering audio and video data into an audio signal and a video signal, which may drive a display system having a speaker

subsystem and a video subsystem. Playback typically involves a technique known in the art as "streaming," whereby the content server sequentially transmits the digital content file to the content player, and the content player plays the digital content file while content data is received that comprises the digital content file. To account for  
5 variable latency and bandwidth within the communications network, a content buffer queues the incoming content data ahead of the content data actually being played. During moments of network congestion, which leads to lower available bandwidth, less content data is added to the content buffer, which may drain down as content data is being de-queued to support playback at a certain playback bit rate. However, during  
10 moments of high network bandwidth, the content buffer is replenished and additional buffer time is added until the content buffer is generally full again. In practical systems, the content buffer may queue content data corresponding to a time span ranging from seconds to more than a minute.

[0005] Each digital content file stored on the content server is typically encoded for a  
15 specific playback bit rate. Prior to initiating playback, the content player may measure available bandwidth from the content server and select a digital content file having an bit rate that can be supported by the measured available bandwidth. To maximize playback quality, a digital content file with the highest bit rate not exceeding the measured bandwidth is conventionally selected. To the extent the communications  
20 network can provide adequate bandwidth to download the selected digital content file while satisfying bit rate requirements, playback proceeds satisfactorily. In practice, however, available bandwidth in the communications network is constantly changing as different devices connected to the communications network perform independent tasks.

[0006] To counter the variability of network conditions, adaptive streaming may be  
25 implemented where, for each title, multiple video streams having different bit rates exist. As the network conditions vary, the content player may switch between video streams according to the network conditions. For example, video data may be downloaded from

video streams encoded to higher bit rates when the network conditions are good, and, when the network conditions deteriorate, subsequent video data may be downloaded from video streams encoded to lower bit rates.

5 [0007] A problem arises with implementing an adaptive streaming solution when the video streams are encoded using a variable bit rate (VBR) technique. In a VBR video stream, to optimize the utilization of bandwidth or space used by a file, different video scenes are encoded based on the complexity of those video scenes. A scene of low complexity is encoded to a lower bit rate to “save” bits for scenes having a higher complexity. The average bit rate across a VBR video stream is, thus, not reflective of  
10 the bit rate of a particular scene within the VBR video stream. This poses a problem when implementing adaptive streaming because the content player selects an encoded video stream based on the average bit rate, but the specific portions of video data downloaded from the encoded video stream may be encoded to a bit rate that is much higher or much lower than the average bit rate. In such a scenario, switching between  
15 encoded video streams may not be appropriate or effective, thus reducing the overall quality of the playback.

[0008] As the foregoing illustrates, what is needed in the art is an approach for downloading digital content to a content player based on the scene complexity of the digital content.

20 [0008A] It is not admitted that any background information presented or referred to in this specification forms part of the common general knowledge.

### **SUMMARY OF THE INVENTION**

[0008B] One aspect of the present invention provides a method for adaptively  
25 downloading digital video content, the method comprising: receiving a scene complexity map specifying a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content; identifying a plurality of encoded video streams associated

with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate; downloading a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to a content buffer residing within a content player device; computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content; determining a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and a relative measure of the first complexity level and the second complexity level; upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, outputting the first encoded data for playback; and upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream: selecting one of the plurality of encoded video streams that corresponds to the determined bit rate; downloading a second encoded data from the selected encoded video stream to the content buffer residing within the content player device; and outputting the second encoded data, without the first encoded data, for playback.

[0008C] Another aspect of the present invention provides a non-transitory computer-readable storage medium storing instructions that, when executed by a processor, cause a computer system to adaptively download digital video content, by performing an operation comprising: receiving a scene complexity map specifying a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content; identifying a plurality of encoded video streams associated with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate; downloading a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to a content buffer residing within a content player device; computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content; determining a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and

a relative measure of the first complexity level and the second complexity level; upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, outputting the first encoded data for playback; and upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream: selecting one of the plurality of encoded video streams that corresponds to the determined bit rate; downloading a second encoded data from the selected encoded video stream to the content buffer residing within the content player device; and outputting the second encoded data, without the first encoded data, for playback.

[0008D] Another aspect of the present invention provides a system for adaptively downloading digital video content, the system comprising: a memory configured to store a playback application, a content buffer, and a scene complexity map associated with the digital video content, wherein the scene complexity map specifies a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content; and a processor that is coupled to the memory and configured to execute the playback application that is configured to: identify a plurality of encoded video streams associated with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate; download a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to the content buffer; compute an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content; determine a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and a relative measure of the first complexity level and the second complexity level; upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, output the first encoded data for playback; and upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream: select one of the plurality of encoded video streams that corresponds to the determined bit rate; download a second encoded data from the selected encoded video stream to the content buffer residing

within the content player device; and output the second encoded data, without the first encoded data, for playback.

[0009] One embodiment of the present disclosure sets forth a method for adaptively downloading digital video content. The method comprises the steps of receiving a scene complexity map associated with the digital video content and specifying a complexity level associated with one or more portions of the digital video content. A plurality of encoded video streams associated with the digital video content are identified, where each encoded video stream is associated with a different bit rate and

includes an encoded portion encoded to the different bit rate for each of the one or more portions of the digital video content. An estimated bandwidth for a first time window is computed based on the actual bandwidth for downloading the digital video content over one or more time windows. The complexity level associated with a first portion of the digital video content and a second portion of the digital video content is determined based on the scene complexity map. During playback of a different portion of the digital video content, a first encoded video stream included in the plurality of encoded video streams from which to download a first encoded portion corresponding to the first portion of the digital video content is selected based on the estimated bandwidth, the first complexity level, and the second complexity level. The first encoded portion is downloaded for playback from the first encoded video stream to a content buffer residing within a content player device.

[0010] One advantage of the disclosed technique is that a variable bit encoded stream is dynamically generated by the content player device at playback time by selecting portions of video data from different constant bit rate encoded streams based on the complexity levels of the portions of video data. Such a technique allows for optimizing the playback of the video data and generating the highest playback quality video stream based on the current conditions and scene complexities.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0011] So that the manner in which the above recited features of the present invention can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0012] Figure 1 illustrates a content distribution system configured to implement one or more aspects of the present invention;

[0013] Figure 2 is a more detailed view of the encoding server of Figure 1, according to one embodiment of the invention;

5 [0014] Figure 3 is an illustration of a scene complexity map generated by the complexity map generator of Figure 2, according to one embodiment of the invention;

[0015] Figure 4 is an illustration of different video streams generated by the video stream encoder, according to one embodiment of the invention;

10 [0016] Figure 5 is a more detailed view of the content player of Figure 1, according to one embodiment of the invention;

[0017] Figure 6 is a flow diagram of method steps for selecting a next scene for playback from one of a plurality of video streams based on scene complexity, according to one embodiment of the invention;

15 [0018] Figure 7 is an illustration of a scene complexity map generated by the complexity map generator of Figure 2, estimated bandwidth, and selected encoded video streams, according to one embodiment of the invention; and

[0019] Figure 8 is a flow diagram of method steps for selecting encoded video streams for playback from one of a plurality of video streams based on scene complexity and estimated bandwidth, according to one embodiment of the invention.

20 **DETAILED DESCRIPTION**

[0020] In the following description, numerous specific details are set forth to provide a more thorough understanding of the present invention. However, it will be apparent to one of skill in the art that the present invention may be practiced without one or more

of these specific details. In other instances, well-known features have not been described in order to avoid obscuring the present invention.

[0021] Figure 1 illustrates a content distribution system 100 configured to implement one or more aspects of the invention. As shown, the content distribution system 100  
5 includes an encoding server 102, a communications network 104, a content distribution network (CDN) 106 and a content player 108.

[0022] The communications network 104 includes a plurality of network communications systems, such as routers and switches, configured to facilitate data communication between the encoding server 102, the CDN 106 and the content player  
10 108. Persons skilled in the art will recognize that many technically feasible techniques exist for building the communications network 104, including technologies practiced in deploying the well-known internet communications network.

[0023] The encoding server 102 is a computer system configured to encode video streams associated with digital content files for adaptive streaming. The encoding  
15 workflow for encoding the video streams for adaptive streaming is described in greater detail below with respect to Figures 2 and 3. The content distribution system 100 maybe include one or more encoding servers 102, where each encoding server 102 is configured to perform all the functions needed to encode the video streams or where  
20 each encoding server 102 is configured to perform a particular function needed to encode the video streams. The digital content files including the encoded video streams are retrieved by the CDN 106 via the communications network 104 for distribution to the content player 108.

[0024] The CDN 106 comprises one or more computer systems configured to serve download requests for digital content files from the content player 108. The digital  
25 content files may reside on a mass storage system accessible to the computer system. The mass storage system may include, without limitation, direct attached storage,

network attached file storage, or network attached block-level storage. The digital content files may be formatted and stored on the mass storage system using any technically feasible technique. A data transfer protocol, such as the well-known hyper-text transfer protocol (HTTP), may be used to download digital content files from the content server 106 to the content player 108.

[0025] The content player 108 may comprise a computer system, a set top box, a mobile device such as a mobile phone, or any other technically feasible computing platform that has network connectivity and is coupled to or includes a display device and speaker device for presenting video frames, and generating acoustic output, respectively. The content player 108 is configured for adaptive streaming, i.e., to download units of a video stream encoded to a specific playback bit rate, and switch to downloading subsequent units of a video stream encoded to a different playback bit rate based on prevailing bandwidth conditions within the communications network 104. As bandwidth available within the communications network 104 becomes limited, the content player 108 may select a video stream encoded to a lower playback bit rate. As the bandwidth increases, a video stream encoded to a higher playback bit rate may be selected.

[0026] Although, in the above description, the content distribution system 100 is shown with one content player 108 and one CDNs 106, persons skilled in the art will recognize that the architecture of Figure 1 contemplates only an exemplary embodiment of the invention. Other embodiments may include any number of content players 108 and/or CDNs 106. Thus, Figure 1 is in no way intended to limit the scope of the present invention in any way.

[0027] Figure 2 is a more detailed illustration of the encoding server 102 of Figure 1, according to one embodiment of the invention. As shown, the encoding server 102 includes a central processing unit (CPU) 202, a system disk 204, an input/output (I/O)

devices interface 206, a network interface 208, an interconnect 210 and a system memory 212.

[0028] The CPU 202 is configured to retrieve and execute programming instructions stored in the system memory 212. Similarly, the CPU 202 is configured to store  
5 application data and retrieve application data from the system memory 212. The interconnect 210 is configured to facilitate transmission of data, such as programming instructions and application data, between the CPU 202, the system disk 204, I/O devices interface 206, the network interface 208, and the system memory 212. The I/O devices interface 206 is configured to receive input data from I/O devices 222 and  
10 transmit the input data to the CPU 202 via the interconnect 210. For example, I/O devices 222 may comprise one or more buttons, a keyboard, and a mouse or other pointing device. The I/O devices interface 206 is also configured to receive output data from the CPU 202 via the interconnect 210 and transmit the output data to the I/O devices 222. The system disk 204, such as a hard disk drive or flash memory storage  
15 drive or the like, is configured to store non-volatile data such as encoded video streams. The encoded video streams can then be retrieved by the CDN 106 via the communications network 104. The network interface 218 is coupled to the CPU 202 via the interconnect 210 and is configured to transmit and receive packets of data via the communications network 104. In one embodiment, the network interface 208 is  
20 configured to operate in compliance with the well-known Ethernet standard.

[0029] The system memory 212 includes software components that include instructions for encoding one or more video streams associated with a specific content title for adaptive streaming. As shown, these software components include a complexity map generator 214, a video stream encoder 216, and a sequence header  
25 index (SHI) generator 218.

[0030] For a particular video stream, the complexity map generator 214 analyzes the

video stream to determine the complexity of the video content within different portions of the video stream (referred to herein as “scenes”). A complex scene is, typically, a scene that changes significantly frame-over-frame, for example a car crash scene in an action movie. Conversely, a simple scene is, typically, a scene with not many changes frame-over-frame, for example a scene of a still body of water at night. The complexity generator 214 may analyze the video stream based on pre-determined heuristics information. Based on the analysis, the complexity map generator 214 generates a scene complexity map that is described in greater detail below with respect to Figure 3.

5 [0031] The video stream encoder 216 executes encoding operations for encoding a video stream to a specific playback bit rate such that the encoded video stream complies with a particular video codec standard, such as VC1, and is configured for adaptive streaming. In an alternative embodiment, the video stream can be encoded to comply with a different video codec standard such as MPEG or H.264. In operation, for a particular video stream, the video stream encoder 216 encodes the video stream to 10 different constant bitrates to generate multiple encoded video streams, each encoded video stream associated with a different constant bitrate and, thus, having a different quality. An encoded video stream generated by the video stream encoder 216 includes a sequence of groups of pictures (GOPs), each GOP comprising multiple frames of video data.

15 [0032] The SHI generator 218 generates a sequence header index associated with each encoded video stream. To generate the sequence header index, the SHI generator 218 first searches the encoded video stream for the key frames associated with the different GOPs included in the encoded video stream. The key frames can be located by the SHI generator 218 based on the sequence start codes specified in the 20 sequence headers included in the key frames. For the GOP associated with each of the identified key frames, the SHI generator 218 defines a switch point within the sequence header index that stores (i) a data packet number that identifies the data 25

packet that includes the key frame associated with the GOP and (ii) the playback offset associated with the GOP. Again, the playback offset associated with the GOP is determined based on the location of the GOP in the sequence of GOPs included in the encoded video stream.

- 5 [0033] The encoding server 102 may generate multiple encoded video streams associated with the same content title and encoded to different playback bit rates in the manner described above. The encoding process described herein ensures that, across the different encoded video streams the GOPs are associated with the same playback time interval and that corresponding GOPs across the different encoded video streams
- 10 are associated with the same playback offsets. Therefore, each switch point defined in a sequence header included in one of the encoded video stream associated with a specific content title has a corresponding switch point defined in a sequence header included in each of the other encoded video stream associated with the same content title.
- 15 [0034] Based on the sequence header indices included in two encoded video streams associated with the same content title, a content player can efficiently switch between the encoded video streams by identifying the appropriate switch points in the sequence header indices. When switching between a currently playing encoded video stream and a new encoded video stream, a content player, such as the content player
- 20 108, searches the sequence header index included in the new encoded video stream to locate the particular switch point specifying the playback offset associated with the next GOP to be played. The content player can then switch to the new encoded video stream and download the GOP stored in the data packet specified at the particular switch point for playback. For example, for encoded video streams where each GOP
- 25 were associated with a playback time interval of three seconds, if the first GOP associated with the playback offset of zero seconds were currently being played, then the next GOP to be played would be associated with the playback offset of three

seconds. In such a scenario, the content player searches the sequence header associated with the new encoded stream for the particular switch point specifying a playback offset of three seconds. Once locating the particular switch point, the content player would download the GOP stored in the data packet specified in the switch point  
5 for playback.

[0035] In practice, a GOP may include multiple scenes or portions of a scene. For the sake of simplicity, with respect to the current invention, the discussion set forth below focuses on particular scenes within an encoded video stream rather than the GOPs within the encoded video stream. While a content player, such as content player  
10 108, may switch between different encoded video streams based on the GOP boundaries defined by the corresponding sequence header indices, the switching process takes into account the complexities of the different scenes included within the GOP. This switching process is described in further detail below.

[0036] Figure 3 is an illustration of a scene complexity map 302 generated by the complexity map generator 214 of Figure 2, according to one embodiment of the  
15 invention. As shown, the scene complexity map 302 specifies the complexity level of different scenes 304 within a video stream. For example, scene 304(0) and 304(4) each has a medium scene complexity, scene 304(1) has a low scene complexity and scene 304(2) and 304(3) each has a high scene complexity. Other embodiments of the  
20 scene complexity map 302 are also contemplated by this invention. In alternate embodiments, the scene complexity levels may be numeric-based and/or be more granular. For the purposes of this invention, a scene complexity map 302 specifies a scene complexity level for each scene in a video stream, where a particular scene corresponds to a specific set of frames within the video stream.

25 [0037] Figure 4 is an illustration of different encoded video streams 404 generated by the video stream encoder 216, according to one embodiment of the invention. As

shown, each encoded video stream 404 is associated with the same title and includes the scenes 304 illustrated in Figure 3. Further, each encoded video stream 404 is encoded to a different bit rate. Encoded video stream 404(0) is encoded to bit rate that is lower than the bit rate of encoded video stream 404(1). Similarly, encoded video stream 404(1) is encoded to bit rate that is lower than the bit rate of encoded video stream 404(2). As also shown, the bit allocation every scene 304 in each encoded video stream 404 is constant. For example, each scene 304 within encoded video stream 404(0) has a bit allocation identified by bit allocation 406. Similarly, each scene 304 within encoded video stream 404(1) has a bit allocation identified by bit allocation 408, and each scene 304 within encoded video stream 404(2) has a bit allocation identified by bit allocation 410. Importantly, the bit allocations 406, 408 and 410 vary according to the bit rate associated with the corresponding encoded video stream 404, where bit allocation 406 is smaller than bit allocation 408 and bit allocation 408 is smaller than bit allocation 410.

**[0038]** Figure 5 is a more detailed view of the content player 108 of Figure 1, according to one embodiment of the invention. As shown, the content player 108 includes, without limitation, a central processing unit (CPU) 510, a graphics subsystem 512, an input/output (I/O) device interface 514, a network interface 518, an interconnect 520, and a memory subsystem 530. The content player 108 may also include a mass storage unit 516.

**[0039]** The CPU 510 is configured to retrieve and execute programming instructions stored in the memory subsystem 530. Similarly, the CPU 510 is configured to store and retrieve application data residing in the memory subsystem 530. The interconnect 520 is configured to facilitate transmission of data, such as programming instructions and application data, between the CPU 510, graphics subsystem 512, I/O devices interface 514, mass storage 516, network interface 518, and memory subsystem 530.

[0040] The graphics subsystem 512 is configured to generate frames of video data and transmit the frames of video data to display device 550. In one embodiment, the graphics subsystem 512 may be integrated into an integrated circuit, along with the CPU 510. The display device 550 may comprise any technically feasible means for  
5 generating an image for display. For example, the display device 550 may be fabricated using liquid crystal display (LCD) technology, cathode-ray technology, and light-emitting diode (LED) display technology (either organic or inorganic). An input/output (I/O) device interface 514 is configured to receive input data from user I/O devices 552 and transmit the input data to the CPU 510 via the interconnect 520. For  
10 example, user I/O devices 552 may comprise one or more buttons, a keyboard, and a mouse or other pointing device. The I/O device interface 514 also includes an audio output unit configured to generate an electrical audio output signal. User I/O devices 552 includes a speaker configured to generate an acoustic output in response to the electrical audio output signal. In alternative embodiments, the display device 550 may  
15 include the speaker. A television is an example of a device known in the art that can display video frames and generate an acoustic output. A mass storage unit 516, such as a hard disk drive or flash memory storage drive, is configured to store non-volatile data. A network interface 518 is configured to transmit and receive packets of data via the communications network 150. In one embodiment, the network interface 518 is  
20 configured to communicate using the well-known Ethernet standard. The network interface 518 is coupled to the CPU 510 via the interconnect 520.

[0041] The memory subsystem 530 includes programming instructions and data that comprise an operating system 532, user interface 534, and playback application 536. The operating system 532 performs system management functions such as managing  
25 hardware devices including the network interface 518, mass storage unit 516, I/O device interface 514, and graphics subsystem 512. The operating system 532 also provides process and memory management models for the user interface 534 and the playback application 536. The user interface 534 provides a specific structure, such as

a window and object metaphor, for user interaction with content player 108. Persons skilled in the art will recognize the various operating systems and user interfaces that are well-known in the art and suitable for incorporation into the content player 108.

[0042] The playback application 536 is configured to retrieve digital content from the  
5 CDN 106 via the network interface 518 and play the digital content through the graphics subsystem 512. The graphics subsystem 512 is configured to transmit a rendered video signal to the display device 550. In normal operation, the playback application 536 receives a request from a user to play a specific title. The playback application 536 then identifies the different encoded video streams associated with the requested title,  
10 wherein each encoded video stream is encoded to a different playback bit rate. After the playback application 536 has located the encoded video streams associated with the requested title, the playback application downloads sequence header indices associated with each encoded video stream associated with the requested title from the CDN 106. As previously described herein, a sequence header index associated with  
15 an encoded video stream includes information related to the encoded sequence included in the digital content file.

[0043] In one embodiment, the playback application 536 begins downloading the digital content file associated with the requested title comprising the encoded sequence encoded to the lowest playback bit rate to minimize startup time for playback. For the  
20 purposes of discussion only, the digital content file is associated with the requested title and comprises the encoded sequence encoded to the lowest playback bit rate. The requested digital content file is downloaded into the content buffer 543, configured to serve as a first-in, first-out queue. In one embodiment, each unit of downloaded data comprises a unit of video data or a unit of audio data. As units of video data associated  
25 with the requested digital content file are downloaded to the content player 108, the units of video data are pushed into the content buffer 543. Similarly, as units of audio data associated with the requested digital content file are downloaded to the content

player 108, the units of audio data are pushed into the content buffer 543. In one embodiment the units of video data are stored in video buffer 546 within the content buffer 543, and units of audio data are stored in audio buffer 544, also within the content buffer 543.

5 [0044] A video decoder 548 reads units of video data from the video buffer 546, and renders the units of video data into a sequence of video frames corresponding in duration to the fixed span of playback time. Reading a unit of video data from the video buffer 546 effectively de-queues the unit of video data from the video buffer 546 (and from the content buffer 543). The sequence of video frames is processed by graphics  
10 subsystem 512 and transmitted to the display device 550.

[0045] An audio decoder 542 reads units of audio data from the audio buffer 544, and renders the units of audio data into a sequence of audio samples, generally synchronized in time with the sequence of video frames. In one embodiment, the sequence of audio samples is transmitted to the I/O device interface 514, which  
15 converts the sequence of audio samples into the electrical audio signal. The electrical audio signal is transmitted to the speaker within the user I/O devices 552, which, in response, generates an acoustic output.

[0046] Given the bandwidth limitations of the communications network 150, the playback application 536 may download consecutive portions of video data from  
20 different constant bit rate encoded video streams based on scene complexities. In operation, when playback is initiated, the playback application 536 receives the scene complexity map 302 associated with the digital video to be played. As previously described, the scene complexity map 302 specifies the complexity level of different scenes of the digital video. When selecting a next portion of video data for download,  
25 the playback application 536 determines the complexity level of the scene(s) included in the portion of the video data based on the scene complexity map 302. Based on the

complexity level of the scene(s) and one or more performance factors, the playback application 536 then determines the particular encoded video stream from which to download the portion of the video data. For example, in a scenario where the bandwidth is low, if the scene(s) are of low complexity, then the playback application 536 downloads the portion of video data including the scenes from an encoded video stream encoded to low bit rate. In this manner, bandwidth of the communications network 150 may be effectively managed by the playback application 536 to download subsequent portions of the higher bit rate encoded video stream for higher complexity scenes. In such a scenario, less bandwidth is used to download low complexity scenes compared with medium complexity scenes, and bandwidth is advantageously conserved in order to download portions of medium or high bit rate encoded video streams for portions of the video data including scenes of high complexity. In contrast, a conventional content player simply selects one of the variable bit rate encoded video streams based on the available bandwidth, without accounting for the complexity of the scene that is encoded in that particular portion of the variable bit rate video stream.

**[0047]** Other performance factors, apart from the complexity levels of the scenes included in a portion of video data, that may influence the specific encoded stream from which to download the portion of video data include complexity levels of subsequent scenes of the video data, the buffer size of the video buffer 546, the behavior of the end-user viewing the video content, the type of display being generated (high-definition, standard-definition, etc) and the available lead time. These factors combined with the bandwidth limitations of the communications network 150 may be used to determine a specific encoded video stream from which to download each portion of the video data based on the complexity levels of the scenes included in the portion of the video data. In such a manner, a variable bit rate video stream is generated from different constant bit rate encoded video streams.

**[0048]** In an alternate embodiment, only portions of video data that include scenes

of high complexity are encoded to a high bit rate. Similarly, only portions of video data that include scenes of medium or high complexity are encoded to a medium bit rate. Portions of video data that only include scenes of low complexity are only encoded to a low bit rate. Referring back to Figure 4, the medium bit rate encoded level of the video stream, encoded video stream 404(1) would not include scene 304(1) and the high bit rate encoded level of the video stream, encoded video stream 404(2) would not include scene 304(0), 304(1), and 304(4). In such an embodiment, the playback application 536 may download only portions of video data including scenes of high complexity of encoded video streams encoded to high bit rates and all other portions of video data from encoded video streams encoded to lower bit rates.

[0049] Figure 6 is a flow diagram of method steps for selecting a next scene for playback from one of a plurality of video streams based on scene complexity, according to one embodiment of the invention. Although the method steps are described in conjunction with the systems for Figures 1-5, persons skilled in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the invention.

[0050] At step 602, the playback application 536 receives the scene complexity map 302 associated with the digital video for which playback has been initiated. As previously described, the scene complexity map 302 specifies the complexity level of different scenes of the digital video. At step 604, the playback application 536 identifies a set of encoded video streams associated with the digital video to be played. Each encoded video stream is encoded to a different bit rate, as previously described in conjunction with Figure 3.

[0051] At step 606, the playback application 536 determines, for a next portion of the video data, the complexity level associated with the scene(s) included in the next portion. The complexity level is determined based on the scene complexity map 302

received at step 602. At step 608, the playback application 536 then selects a specific encoded video stream for downloading the next portion of the video data based on the determined complexity level as well as one or more performance factors. As previously described, the performance factors may include bandwidth limitations and the size of the content buffer 543. In order to select the specific encoded video stream, the playback application 536 executing on the content player 108 dynamically determines the encoding level (high, medium, or low bit rate) of the video stream for the next portion of the video data to be downloaded during playback of a different (previous) portion of the digital video content.

10 [0052] At step 610, the playback application 536 determines if another timestep occurs during playback of the video data, and, if so, then the playback application 536 repeats steps 606 and 608 for another portion of the video stream. When another timestep does not occur during the playback of the video data, i.e., playback of the video content is done. The timestep may occur at a constant rate (in seconds or  
15 frames) or be triggered based on a fullness or emptiness of the content buffer 543.

[0053] One advantage of the disclosed technique is that a variable bit encoded stream is dynamically generated at playback time by selecting portions of video data from different constant bit rate encoded streams based on the complexity levels of the portions of video data. Such a technique allows for optimizing the playback of the video data and generating the highest playback quality video stream based on the current  
20 conditions and scene complexities.

#### **Video Stream Selection Based on Bandwidth and Scene Complexity**

[0054] The selection of a video stream to be downloaded may vary based on bandwidth limitations between the content player 108 and the encoding server 102. In some cases, the scene complexity map is used by the content player 108 to download  
25 scenes of lower scene complexity levels from video streams encoded to lower constant

bit rates to allow for the subsequent downloading of scenes of higher complexity levels from video streams encoded to higher constant bit rates. In other words, the quality of the encoded video downloaded for some scenes may be compromised (downloaded at a lower bit rate) in order to download a higher quality of encoded video for subsequent scenes that are encoded at a higher bit rate. The objective is to balance the quality level over several scenes so that scenes of higher complexity are downloaded at a higher bit rate compared with scenes of lower complexity. A performance factor that is used to select the bit rate of the encoded video for download may include a maximum tolerable quality degradation. For example, the maximum tolerable quality degradation may be a single level, so that a scene of high complexity may be downloaded at a high or medium bit rate, but not at a low bit rate.

[0055] The content player 108 may determine the estimated bandwidth by sampling the actual bandwidth between the content player 108 and the encoding server 102. The estimated bandwidth is an estimate of the actual bandwidth for a future time window. For example, the estimated bandwidth may be determined based on a history of the actual bandwidth for a portion of the video data or a time window. The duration of the time window may be fixed or programmed to be as small as a single frame or as large as multiple GOPs. The estimated bandwidth may be computed as the average of samples of the actual bandwidth during one or more previous time windows. Alternatively, the estimated bandwidth may be computed using a function, where the actual bandwidth samples during the time window are an input to the function. The actual bandwidth between the content player 108 and the encoding server 102 may be sampled hundreds, thousands, or any number of times during a time window.

[0056] In addition to computing the estimate bandwidth, the bandwidth variability may also be determined as the variability of the actual bandwidth relative to the estimated bandwidth and represented as a deviation value, e.g., number of standard deviations from the estimated bandwidth. At various points in time during a time

window, the actual bandwidth may vary from at or near zero to the high bit rate or even higher than the high bit rate. In particular video content streamed via satellite may have a large amount of bandwidth variability and a correspondingly high deviation value. In other cases, the actual bandwidth may vary by small amounts and the bandwidth variability will have a correspondingly low deviation value. At the beginning of a new streaming session, the estimated bandwidth and bandwidth variability may be initialized to a value based on the behavior from a previous session. The computed estimated bandwidth and one or more of the scene complexity and bandwidth variability may be used to select an encoded video stream at a particular fixed bit rate for download to optimize the quality of the digital content received by the content player 108.

[0057] Figure 7 is an illustration of a scene complexity map 700 generated by the complexity map generator 214 of Figure 2, according to one embodiment of the invention. The estimated bandwidth 710 and the selected stream 720, i.e., the portions of fixed rate video streams that are selected for each scene to generate the variable bit rate video stream, are also shown in Figure 7.

[0058] As shown, the scene complexity map 700 specifies the complexity level of different scenes 704 within a video stream. For example, scene 704(0), 704(4), and 704(5) each has a medium scene complexity, scene 704(1) and 704(9) each has a low scene complexity, and scene 704(2), 704(3), 704(6), 704(7), and 704(8) each has a high scene complexity. Other embodiments of the scene complexity map 700 are also contemplated by this invention. In alternate embodiments, the scene complexity levels may be numeric-based and/or be more granular. For the purposes of the embodiments of the invention described herein, a scene complexity map 700 specifies a scene complexity level for each scene in a video stream, where a particular scene corresponds to a specific set of frames within the video stream.

[0059] As described above, different encoded video streams are generated by the

video stream encoder 216, according to one embodiment of the invention. Each encoded video stream is associated with the same title and includes the scenes 704 illustrated in Figure 7. Further, each encoded video stream is encoded to a different fixed bit rate, e.g., low, medium, and high fixed bit rates. In an alternate embodiment, only portions of video data that include scenes of high complexity are encoded to a high bit rate. Similarly, only portions of video data that include scenes of medium or high complexity are encoded to a medium bit rate. Portions of video data that only include scenes of low complexity are only encoded to a low bit rate. Therefore, the high and medium encoded video streams may not include all of the scenes 704 that are illustrated in Figure 7. For example, the high and medium encoded video streams may not include scenes 704 (1) and 704(9).

[0060] The estimated bandwidth 710 between the content player 108 and the encoding server 102 is determined by the content player 107 for each time window, where a time window corresponds to a scene 704. The estimated bandwidth for each scene may be computed to be low, medium, or high corresponding to the scene complexities of low, medium, and high, respectively. The estimated bandwidth associated with scene 704(0) is medium and the complexity of scene 704(0) is also medium, so scene 704(0) can be downloaded at the medium fixed bit rate.

[0061] The estimated bandwidth 710 associated with scene 704(1) is high and the complexity of scene 704(1) is low, so scene 704(1) can be downloaded at the low, medium, or high fixed bit rate. Since the complexity of scene 704(1) is low, scene 704(1) is downloaded at the low fixed bit rate (as shown in selected stream 720). During the time window corresponding to the scene 704(1) additional encoded video data may be downloaded, so scene 704(2) having a complexity that is high is also downloaded. Scene 704(2) is downloaded at a high fixed bit rate during the time window corresponding to scene 704(1).

[0062] The estimated bandwidth 710 associated with scene 704(2) and 704(3) is medium, so scene 704(3) may be downloaded at a high fixed bit rate during the time windows corresponding to scenes 704(2) and 704(3). The estimated bandwidth associated with scenes 704(4), 704(5), 704(6), 704(7), 704(8), and 704(9) is medium.

5 In order to download at least scenes 704(6) and 704(7) at a high fixed bit rate, the low fixed bit rate encoded video streams are selected for scenes 704(4) and 704(5) (see the selected stream 720). The low fixed bit rate encoded video streams are downloaded for scenes 704(4) and 704(5) during the time window corresponding to scene 704(4). The high fixed bit rate encoded video streams are then selected and  
10 downloaded for scenes 704(6) and 704(7) during the time windows corresponding to scene 704(5), scene 704(6), and 704(7). Because scenes 704(4) and 704(5) are downloaded at a lower bit rate, scenes 704(6) and 704(7) may be downloaded at a higher bit rate.

[0063] Finally, the medium fixed bit rate encoded video stream is selected for  
15 downloading scene 704(8) since the estimated bandwidth is medium. The low fixed bit rate encoded video stream is selected for downloading scene 704(9) since scene 704(9) has a low complexity. Any surplus bandwidth during the time window associated with scene 704(9) may be used to download a video stream for a subsequent scene.

20 [0064] Figure 8 is a flow diagram of method steps 800 for selecting encoded video streams for playback from one of a plurality of video streams based on scene complexity and estimated bandwidth, according to one embodiment of the invention. Although the method steps 800 are described in conjunction with the systems for Figures 1-7, persons skilled in the art will understand that any system configured to  
25 perform the method steps, in any order, is within the scope of the invention.

[0065] At step 802, the playback application 536 receives the scene complexity map

302 associated with the digital video for which playback has been initiated. As previously described, the scene complexity map 700 specifies the complexity level of different scenes of the digital video. At step 804, the playback application 536 identifies a set of encoded video streams associated with the digital video to be played. Each  
5 encoded video stream is encoded to a different fixed bit rate, as previously described in conjunction with Figures 3 and 7.

[0066] At step 806, the playback application 536 computes, for a next time window, the estimated bandwidth for downloading the encoded video streams from the encoding server 102 to the content player 108. The estimated bandwidth is an estimation that is  
10 computed based on the actual bandwidth over one or more previous time windows. At step 807, the playback application 536 computes, for a next time window, the bandwidth variability of the actual bandwidth relative to the estimated bandwidth. The bandwidth variability is computed based on the actual bandwidth over one or more previous time windows and the estimated bandwidth.

15 [0067] At step 808, the playback application 536 determines, for the next one or more time windows, the complexity level associated with the scene(s). Looking ahead by more than one time window allows the playback application 536 to balance the quality level of the video over several scenes based on the estimated bandwidth so that scenes of higher complexity are downloaded at a higher bit rate compared with scenes  
20 of lower complexity. The complexity level is determined based on the scene complexity map 700 received at step 802. At step 810, the playback application 536 then selects a specific encoded video stream for downloading the next portion of the video data based on the computed estimated bandwidth, the computed bandwidth variability, and the determined complexity level as well as one or more performance factors. The  
25 performance factors may include the size of the content buffer 543 and a maximum tolerable quality degradation. In order to select the specific encoded video stream, the playback application 536 executing on the content player 108 dynamically determines

the encoding level (high, medium, or low bit rate) of the video stream for the next portion of the video data to be downloaded during playback of a different (previous) portion of the digital video content.

5 [0068] At step 812, the playback application 536 determines whether another time window occurs during playback of the video data, and, if so, then the playback application 536 repeats steps 806, 808, and 810 for another portion of the video stream. When another time window does not occur during the playback of the video data, i.e., playback of the video content terminates. The time window may occur at a constant rate (in seconds, frames, or GOPs) or be triggered based on a fullness or  
10 emptiness of the content buffer 543.

[0069] In one embodiment, the playback application 536 downloads the video stream at the low bit rate for one or more time windows to ensure that at least the low bit rate encoded video stream is available in a first portion of the video buffer 546. Whenever the low bit rate encoded video stream is selected for a time window at step  
15 810, the playback application 536 does not download the low bit rate encoded video stream for the time window since that data is already stored in the first portion of the video buffer 546. When, at step 810 the medium or high bit rate encoded video stream is selected for a time window, the playback application 546 download the selected encoded video stream to a second portion of the video buffer 546.

20 [0070] One advantage of the disclosed technique is that a variable bit encoded stream is dynamically generated at playback time by selecting portions of video data from different constant bit rate encoded streams based on the estimated bandwidth for downloading the encoded video data and complexity levels of the portions of video data. Such a technique allows for optimizing the playback of the video data based on  
25 the available download bandwidth, bandwidth variability, and scene complexities.

[0071] In one embodiment of the invention the content player 108 comprises an

embedded computer platform such as a set top box. An alternative embodiment of the invention may be implemented as a program product that is downloaded to a memory within a computer system, for example as executable instructions embedded within an internet web site. In this embodiment, the content player 108 comprises the computer system.

[0072] While the foregoing is directed to embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof. For example, aspects of the present invention may be implemented in hardware or software or in a combination of hardware and software.

10 One embodiment of the invention may be implemented as a program product for use with a computer system. The program(s) of the program product define functions of the embodiments (including the methods described herein) and can be contained on a variety of computer-readable storage media. Illustrative computer-readable storage media include, but are not limited to: (i) non-writable storage media (e.g., read-only memory devices within a computer such as CD-ROM disks readable by a CD-ROM drive, flash memory, ROM chips or any type of solid-state non-volatile semiconductor memory) on which information is permanently stored; and (ii) writable storage media (e.g., floppy disks within a diskette drive or hard-disk drive or any type of solid-state random-access semiconductor memory) on which alterable information is stored. Such

15

20 computer-readable storage media, when carrying computer-readable instructions that direct the functions of the present invention, are embodiments of the present invention.

[0073] In view of the foregoing, the scope of the present invention is determined by the claims that follow.

The claims defining the invention are as follows:

1. A method for adaptively downloading digital video content, the method comprising:

receiving a scene complexity map specifying a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content;

identifying a plurality of encoded video streams associated with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate;

downloading a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to a content buffer residing within a content player device;

computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content;

determining a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and a relative measure of the first complexity level and the second complexity level;

upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, outputting the first encoded data for playback; and

upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream:

selecting one of the plurality of encoded video streams that corresponds to the determined bit rate;

downloading a second encoded data from the selected encoded video stream to the content buffer residing within the content player device; and

outputting the second encoded data, without the first encoded data, for playback.

2. The method of claim 1, further comprising selecting a second encoded video stream included in the plurality of encoded video streams from which to download a second encoded portion corresponding to the second portion of the digital video content based on the estimated bandwidth, the first complexity level, and the second complexity level.
3. The method of claim 1, wherein computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content further comprises:
  - determining a measure of actual bandwidth that is indicative of an amount of network throughput currently available for the content player device;
  - determining a measure of estimated bandwidth that is indicative of an amount of network throughput available on the content player device during one or more historical windows of time; and
  - calculating the estimated available bandwidth as a measure of variability of the measure of actual bandwidth to the measure of estimated bandwidth, such that the estimated available bandwidth is expressed as a deviation value.
4. The method of claim 1, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on the size of the content buffer, and wherein selecting occurs during playback of a different portion of the digital video content.
5. The method of claim 1, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on a maximum tolerable quality degradation for each one of the one or more portions of the digital video content.
6. The method of claim 1, wherein selecting the first encoded video stream of the plurality of encoded video streams further comprises:
  - upon determining that the second complexity level is more complex than the first complexity level, selecting the first encoded video stream, wherein the first encoded

video stream is a lowest bit rate encoded video stream of the plurality of encoded video streams, thereby increasing an amount of available bandwidth for use in downloading the second portion of digital video content.

7. The method of claim 1, further comprising computing a bandwidth variability value for the first time window based on actual bandwidth available for downloading the digital video content over one or more time windows and the estimated bandwidth.

8. The method of claim 7, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on the bandwidth variability.

9. The method of claim 1, wherein selecting one of the plurality of encoded video streams from which to download encoded data further comprises:

upon determining that the relative measure of the first complexity level and the second complexity level indicates that the first complexity level is more complex than the second complexity level, selecting one of the plurality of encoded video streams encoded at a higher bit rate, thereby increasing an amount of bandwidth used in downloading the first encoded portion from the first encoded video stream, relative to an anticipated amount of bandwidth; and

upon determining that the relative measure of the first complexity level and the second complexity level indicates that the first complexity level is less complex than the second complexity level, selecting one of the plurality of encoded video streams encoded at a lower bit rate, thereby decreasing an amount of bandwidth used in downloading the first encoded portion from the first encoded video stream, relative to the anticipated amount of bandwidth.

10. The method of claim 1, wherein the second portion is subsequent to the first portion with respect to a temporal direction of playback of the digital video content.

11. A non-transitory computer-readable storage medium storing instructions that, when executed by a processor, cause a computer system to adaptively download digital video content, by performing an operation comprising:

receiving a scene complexity map specifying a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content;

identifying a plurality of encoded video streams associated with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate;

downloading a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to a content buffer residing within a content player device;

computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content;

determining a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and a relative measure of the first complexity level and the second complexity level;

upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, outputting the first encoded data for playback; and

upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream:

selecting one of the plurality of encoded video streams that corresponds to the determined bit rate;

downloading a second encoded data from the selected encoded video stream to the content buffer residing within the content player device; and

outputting the second encoded data, without the first encoded data, for playback.

12. The non-transitory computer-readable storage medium of claim 11, further comprising selecting a second encoded video stream included in the plurality of

encoded video streams from which to download a second encoded portion corresponding to the second portion of the digital video content based on the estimated bandwidth, the first complexity level, and the second complexity level.

13. The non-transitory computer-readable storage medium of claim 11, wherein computing an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content further comprises:

determining a measure of actual bandwidth that is indicative of an amount of network throughput currently available for the content player device;

determining a measure of estimated bandwidth that is indicative of an amount of network throughput available on the content player device during one or more historical windows of time; and

calculating the estimated available bandwidth as a measure of variability of the measure of actual bandwidth to the measure of estimated bandwidth, such that the estimated available bandwidth is expressed as a deviation value.

14. The non-transitory computer-readable storage medium of claim 11, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on the size of the content buffer, and wherein selecting occurs during playback of a different portion of the digital video content.

15. The non-transitory computer-readable storage medium of claim 11, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on a maximum tolerable quality degradation for each one of the one or more portions of the digital video content.

16. The non-transitory computer-readable storage medium of claim 11, further comprising computing a bandwidth variability value for the first time window based on actual bandwidth available for downloading the digital video content over one or more time windows and the estimated bandwidth.

17. The non-transitory computer-readable storage medium of claim 16, wherein selecting the first encoded video stream from which to download the first encoded portion is further based on the bandwidth variability.

18. A system for adaptively downloading digital video content, the system comprising:

a memory configured to store a playback application, a content buffer, and a scene complexity map associated with the digital video content, wherein the scene complexity map specifies a first complexity level associated with a first portion of the digital video content and a second complexity level associated with a second portion of the digital video content, wherein the second portion is subsequent to the first portion in playback of the digital video content; and

a processor that is coupled to the memory and configured to execute the playback application that is configured to:

identify a plurality of encoded video streams associated with the digital video content, wherein each of the plurality of encoded video streams comprises an encoding of the digital video content at a respective bit rate;

download a first encoded data corresponding to the first portion from a lowest bit rate stream of the plurality of encoded video streams to the content buffer;

compute an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content;

determine a bit rate at which to playback the first portion of the digital video content based on the estimated available bandwidth and a relative measure of the first complexity level and the second complexity level;

upon determining that the determined bit rate does not exceed a first bit rate of the lowest bit rate stream, output the first encoded data for playback; and upon determining that the determined bit rate exceeds the first bit rate of the lowest bit rate stream:

select one of the plurality of encoded video streams that corresponds to the determined bit rate;

download a second encoded data from the selected encoded video stream to the content buffer residing within the content player device; and  
output the second encoded data, without the first encoded data, for playback.

19. The system of claim 18, wherein the playback application is configured to compute an estimated available bandwidth for use in downloading data relating to the first portion of the digital video content by:

determining a measure of actual bandwidth that is indicative of an amount of network throughput currently available for the content player device;

determining a measure of estimated bandwidth that is indicative of an amount of network throughput available on the content player device during one or more historical windows of time; and

calculating the estimated available bandwidth as a measure of variability of the measure of actual bandwidth to the measure of estimated bandwidth, such that the estimated available bandwidth is expressed as a deviation value.

20. The system of claim 19, wherein the playback application is further configured to select the first encoded video stream from which to download the first encoded portion further based on the bandwidth variability.




FIGURE 1

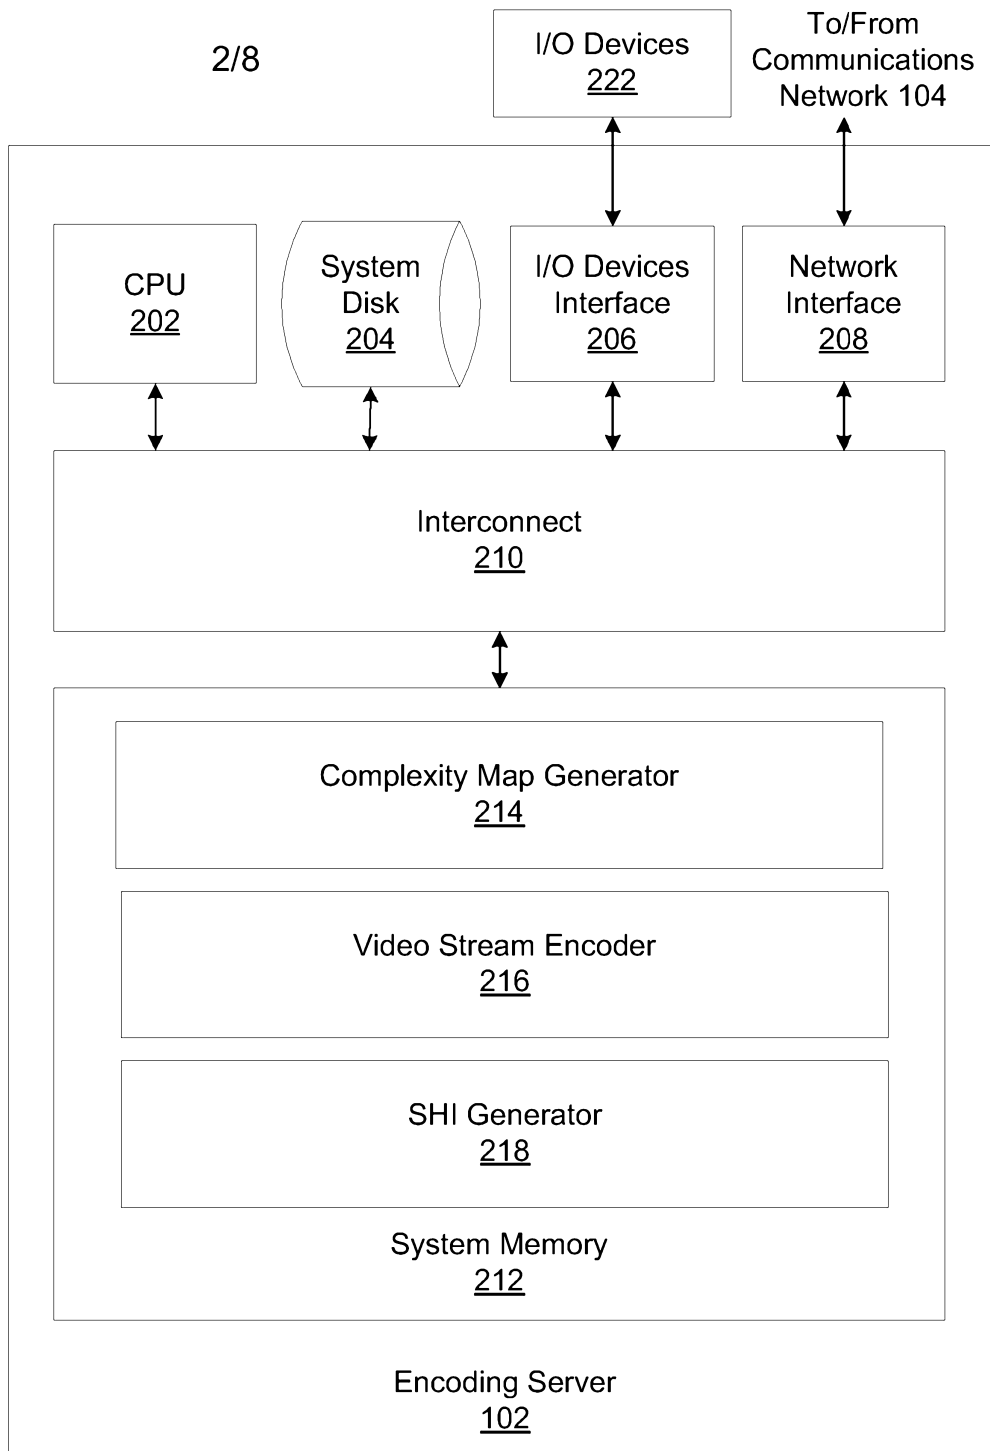


FIGURE 2

Scene Complexity Map  
302

	High Complexity	Medium Complexity	Low Complexity
Scene <u>304(0)</u>		×	
Scene <u>304(1)</u>			×
Scene <u>304(2)</u>	×		
Scene <u>304(3)</u>	×		
Scene <u>304(4)</u>		×	

Figure 3

4/8




Figure 4

+

5/8

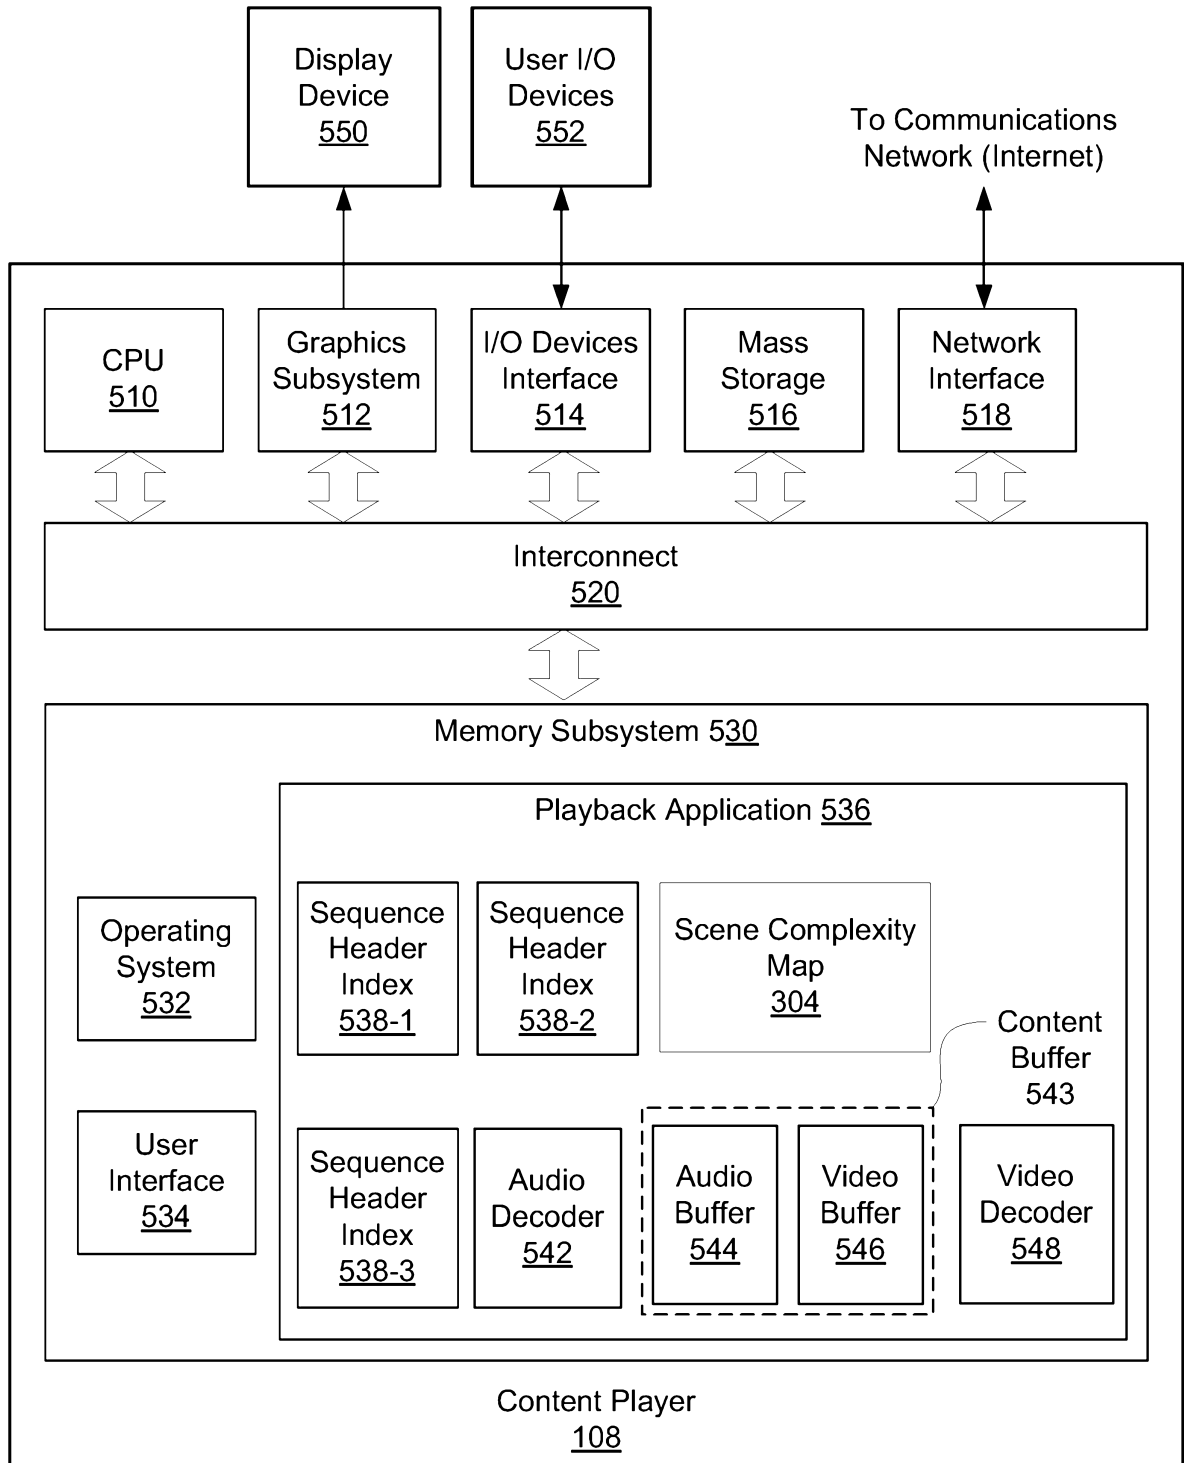


Figure 5

+

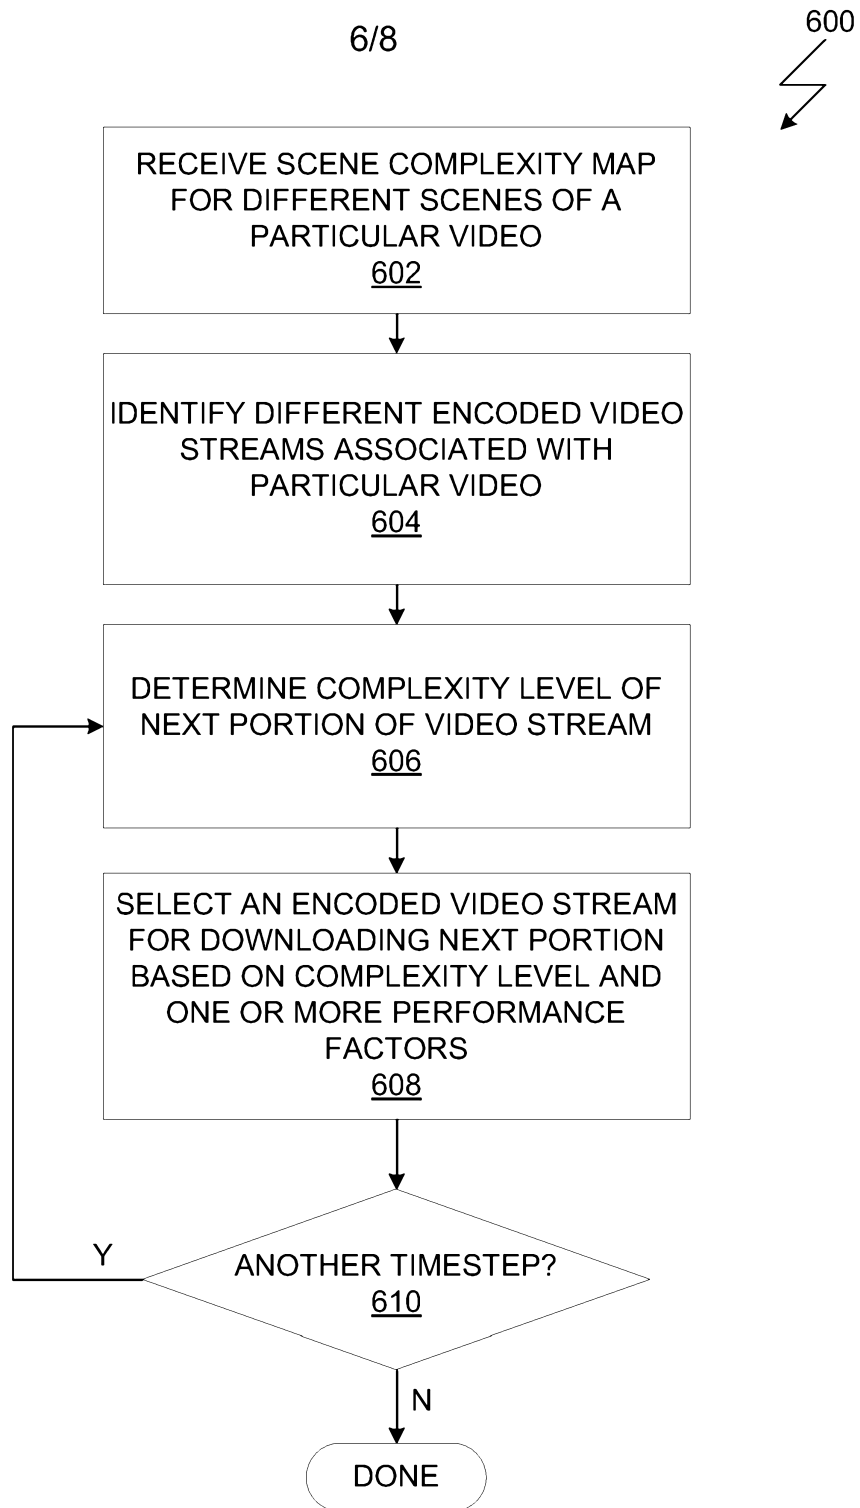


Figure 6

Scene Complexity Map 700			Estimated Bandwidth 710	Selected Stream 720
	High Complexity	Medium Complexity	Estimated Bandwidth	Selected Stream
Scene <u>704(0)</u>		×	Medium	Medium
Scene <u>704(1)</u>			High	Low
Scene <u>704(2)</u>	×		Medium	High
Scene <u>704(3)</u>	×		Medium	High
Scene <u>704(4)</u>		×	Medium	Low
Scene <u>704(5)</u>		×	Medium	Low
Scene <u>704(6)</u>	×		Medium	High
Scene <u>704(7)</u>	×		Medium	High
Scene <u>704(8)</u>	×		Medium	Medium
Scene <u>704(9)</u>			Medium	Low

Figure 7

8/8

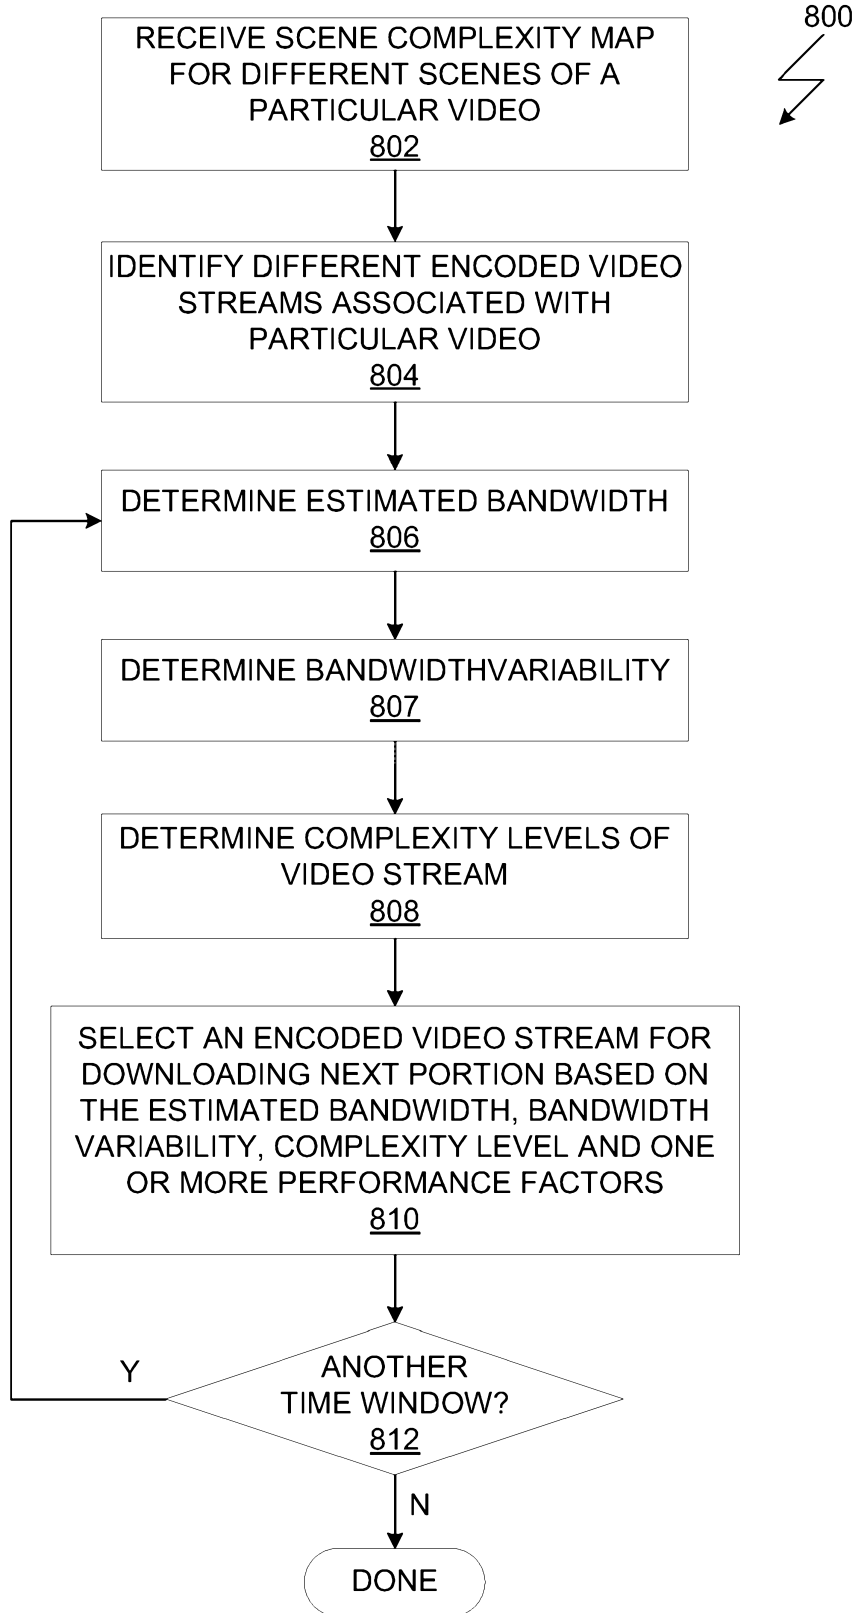


Figure 8