



(51) International Patent Classification:

H04L 1/00 (2006.01) H04N 21/20 (2011.01)

(21) International Application Number:

PCT/US2017/043168

(22) International Filing Date:

20 July 2017 (20.07.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/364,389	20 July 2016 (20.07.2016)	US
15/655,464	20 July 2017 (20.07.2017)	US

(71) Applicant: ARRIS ENTERPRISES LLC [US/US]; 3871 Lakefield Drive, Suwanee, Georgia 30024 (US).

(72) Inventors: MOORTHY, Praveen N.; 17324 Eagle Canyon Way, San Diego, California 92127 (US). SCHMIDT, Mark S.; 3285 Welmer Place, San Diego, California 92122 (US). LUTHRA, Ajay; 13695 Winstanley Way, San Diego, California 92130 (US). LI, Baozhou; 10944 Camino Ruiz, Apt. H, San Diego, California 92126 (US).

(74) Agent: WARD, Thomas A. et al.; 3871 Lakefield Drive, Suwanee, Georgia 30024 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: CLIENT FEEDBACK ENHANCED METHODS AND DEVICES FOR EFFICIENT ADAPTIVE BITRATE STREAMING




FIG. 10

(57) Abstract: A method for distributing video content from a server to a plurality of media devices is provided allowing adaptive bit rate encoding to better utilize bandwidth. The method includes: determining, by the server, the bandwidth to allocate to each of the plurality of media devices using a hypertext transfer protocol-based live streaming client model or a need parameter vector, refining this determination by utilizing client feedback regarding client buffer level and playback state, client hardware capabilities, and client internally measured download rate, and providing the allocated bandwidth to each of the plurality of media devices; wherein the video content is transmitted in a plurality of segments from the server, and wherein each segment is transmitted at a bitrate that may vary from segment to segment.

TITLE

Client Feedback Enhanced Methods and Devices for Efficient Adaptive Bitrate Streaming

CLAIM FOR PRIORITY

[0001] This application claims benefit of U.S. Provisional Patent Application No. 62/364,389, filed July 20, 2016, which application is hereby incorporated by reference herein.

BACKGROUND

[0002] Streaming of media over a network from a content server to a media device has been widely adopted for media consumption. Two network protocols used for media streaming include the user datagram protocol Internet protocol ("UDP IP") and the transfer control protocol ("TCP") IP. UDP IP is often used for media streaming for relatively reliable networks, such as in-home streaming over wired connections. TCP IP is often used for streaming over less reliable networks.

[0003] The hypertext transfer protocol ("HTTP") based live streaming ("HLS") protocol, used with the TCP IP, allows a content server to publish variant playlist files to media devices. A variant playlist file identifies multiple sets of video streams for a media program, such as a movie, a television program, etc., where each set of video streams has unique encoding parameters (e.g., bit rates, resolutions, etc.) for the media program. The media devices may dynamically switch between the sets of video streams identified in the variant playlist file as the sets of video streams are transmitted from the content server to the media devices. The media devices may choose to receive an initial set of video streams identified in the variant playlist file based on initial network

conditions, initial buffer conditions, etc. For example, the media devices may choose to receive a set of high definition ("HD") video streams identified in the variant playlist file if the initial network conditions, the initial buffer conditions, etc., support the streaming of the HD set of video streams. If the initial network conditions degrade, or if the initial buffer conditions degrade, etc., then the media devices may choose to receive a set of low definition video streams identified in the variant playlist file. That is, the media device may dynamically choose different sets of video streams to receive from the content server where the different sets of video streams have different encoding parameters.

[0004] Selection and transmission of the sets of video streams are driven by the media devices. In response to a selection of a set of video streams identified in the variant playlist file, the content server passively transmits the set of video streams to the media device. The media device may have limited information about the network conditions and might not select a set of video streams that is suitable for the current network conditions. Further, some types of media devices select the highest resolution and highest bit rate sets of video streams to receive. Typically the content server services multiple media devices, transmitting multiple sets of video streams to the media devices. If a media device requests a set of video streams with high resolution and high bit rate, a large portion of content server resources or network bandwidth may have to be allocated in order to service that media device. Consequently, the other media devices serviced by the content server may experience degraded performance such as intermittent interruptions in the transmission of video streams.

[0005] Server-side adaptive bitrate of HTTP streams to clients attempts to maintain a state if each client. e.g., the client's video buffer depth and subsequent video download times and rates in order to meet chunk delivery schedules by assigning more or less bandwidth to individual clients

during the schedule window. This process can be improved by providing client feedback regarding several factors, such as hardware capabilities, current buffer depths, playback state, and last-mile or last-meter measured throughput to the server.

SUMMARY

[0006] The present process and method describes a technique to improve Smart Adaptive Bit Rate (SABR) Server performance with client-based feedback. Information about a client's current video buffer size, time duration, and playback state, as well as the client's measured chunk download rate can allow a SABR server to change its deliver schedule and chunk bitrates to allow smoother playback at the client side with fewer buffer underruns. A client's hardware capabilities, such as current available memory, CPU load, screen size/resolution and/or desired caps on video bitrate and/or resolution, can allow a SABR server to set maximum and minimum bitrates and resolutions (caps). This can augment the schedule window algorithm to free up extra bandwidth when a client is expected to get a high bitrate due to need parameter values, but does not want or need it due to available resources. This excess bandwidth can be reallocated to data or peak signal-to-noise ratio gain to improve server performance.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Further details of the present method and process are explained with the help of the attached drawings in which:

[0008] FIG. 1A illustrates an example system implementation in which embodiments of the disclosure may be used;

[0009] FIG. 1B illustrates an example media-streaming system implementation in which embodiments of the disclosure may be used;

[0010] FIG. 2 illustrates a system that includes client components and server components in communication with each other and the message flows for typical adaptive streaming in accordance with embodiments of the disclosure;

[0011] FIG. 3A illustrates a video streaming system in accordance with embodiments of the disclosure;

[0012] FIG. 3B illustrates an example schedule window provided by a schedule module in accordance with embodiments of the disclosure;

[0013] FIG. 4 illustrates an example graph showing a buffering stage and a playback stage for a media-device buffer in accordance with embodiments of the disclosure;

[0014] FIG. 5 illustrates an example HLS Client Model ("HCM") block diagram in accordance with embodiments of the disclosure;

[0015] FIG. 6 illustrates an example graph of curves of bit rate vs. Need Parameter Vector ("NPV") in accordance with embodiments of the disclosure; and

[0016] FIG. 7 illustrates an example block diagram of a system implementing an HCM and schedule queue in accordance with embodiments of the disclosure.

[0017] FIG. 8 illustrates an example HLS Client Model ("HCM") block diagram incorporating client feedback in accordance with embodiments of the disclosure.

[0018] FIG. 9 illustrates an example block diagram of a system implementing an HCM and schedule queue in accordance with embodiments of the disclosure.

[0019] FIG. 10 illustrates an example block diagram of a system implementing an HCM and schedule queue utilizing client feedback in accordance with embodiments of the disclosure.

DETAILED DESCRIPTION

[0020] Turning to the drawings, wherein like reference numerals refer to like elements, techniques of the present disclosure are illustrated as being implemented in a suitable environment. The following description is based on embodiments of the claims and should not be taken as limiting the claims with regard to alternative embodiments that are not explicitly described herein.

[0021] With reference to FIG. 1A, an example media device 100 includes a processing unit 120 and a system bus 110 that couples various system components including the system memory 130 such as read-only memory ("ROM") 140 and random-access memory ("RAM") 150 to the processor 120. The media device 100 can include a cache 122 of high-speed memory connected directly with, in close proximity to, or integrated as part of the processor 120. The media device 100 may be configured to copy data from the memory 130 or a storage device 160 to the cache 122 for quick access by the processor 120. In this way, the cache 122 provides a performance boost that avoids processor delays while waiting for data. These and other modules can control or be configured to control the processor 120 to perform various actions. Other system memory 130 may be available for use as well. The memory 130 can include multiple different types of memory with different performance characteristics. It can be appreciated that the disclosure may operate on a media device 100 with more than one processor 120 or on a group or cluster of computing devices networked together to provide greater processing capability. The processor 120 can include any general-purpose processor and a hardware module or software module, such as module 1 162, module 2 164, and module 3 166 stored in the storage device 160 and configured to control the processor 120 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The processor 120 may essentially be a completely self-contained

computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0022] The system bus 110 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. A basic input/output system stored in ROM 140 or the like may provide the basic routine that helps to transfer information between elements within the media device 100, such as during start-up. The media device 100 further includes storage devices 160, such as a hard disk drive, a magnetic disk drive, an optical disk drive, tape drive, or the like. The storage device 160 is connected to the system bus 110 by a drive interface. The drives and the associated computer-readable storage media provide nonvolatile storage of computer-readable instructions, data structures, program modules, and other data for the media device 100. In some embodiments, a hardware module that performs a particular function includes the software component stored in a non-transitory computer-readable medium in connection with the necessary hardware components, such as the processor 120, bus 110, display 170, and so forth, to carry out the function. The basic components are known to those of skill in the art and appropriate variations are contemplated depending on the type of device, such as whether the device 100 is a small, handheld computing device, a desktop computer, a computer server, or the like.

[0023] Although some implementations employ the hard disk 160, it should be appreciated by those skilled in the art that other types of computer-readable media which can store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, digital versatile disks, cartridges, RAM 150, ROM 140, a cable or wireless signal containing a bit stream and the like, may also be used in the exemplary operating environment. Non-transitory computer-readable

storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0024] Media device 100 also includes a receive buffer 105 that includes three buffer sections 105a, 105b, and 105c. A first buffer section 105a may be for video packets that media device 100 has received from a content server but has not consumed for media play. Media device 100 may have acknowledged receipt of the video packets in the first buffer section 105a to the content server via an acknowledgment. A buffer management module (not shown) may monitor the rate at which video packets in the first buffer section 105a are retrieved for consumption by media device 100.

[0025] A second buffer section 105b may be for video packets that media device 100 has received from a content server but has not consumed for media play. Media device 100 may not have sent acknowledgments to the content server for the video packets in the second buffer section 105b. Portions of the second buffer section 105b may be categorized as a portion of the first buffer section 105a as acknowledgments for video packets in the second buffer section 105b are transmitted to the content server from media device 100. A buffer management module (not shown) may track the portions of the second buffer section 105b that are categorized as a portion of the first video buffer 105a when media device 100 sends an acknowledgment to the content server for acknowledging receipt of the video packets in the second buffer section 105b.

[0026] A third buffer section 105c may be available for receipt of video packets. A buffer management module (not shown) may monitor the third buffer section 105c to determine when the third buffer section 105c receives video packets and is categorized as a portion of the second buffer section 105b. Portions of the first buffer section 105a may be categorized as a portion of the third buffer section 105c as video packets from the first buffer section 105a are consumed. That

is, the portion of the first buffer section 105a for which video packets are consumed may receive new video packets from the content server.

[0027] The sizes of the first, second, and third buffer sections 105a-105c together define the maximum buffer size for video-packet buffering according to some embodiments. The maximum buffer size may be allocated by the media device 100 when opening an initial connection with a content server. The maximum buffer size typically remains unchanged after the allocation.

[0028] To enable user interaction with the media device 100, an input device 190 represents any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, speech and so forth. An output device 170 can also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems enable a user to provide multiple types of input to communicate with the media device 100. The communications interface 180 generally governs and manages the user input and system output. There is no restriction on operating on any particular hardware arrangement, and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0029] For clarity of explanation, the illustrative system embodiment is presented as including individual functional blocks, including functional blocks labeled as a "processor" or processor 120. The functions these blocks represent may be provided through the use of either shared or dedicated hardware, including, but not limited to, hardware capable of executing software and hardware, such as a processor 120, that is purpose-built to operate as equivalent to software executing on a general-purpose processor. For example the functions of one or more processors presented in FIG. 1A may be provided by a single shared processor or multiple processors. Illustrative embodiments may include microprocessor or digital signal processor ("DSP") hardware, ROM 140 for storing

software performing the operations discussed below, and RAM 150 for storing results. Very large scale integration ("VLSI") hardware embodiments, as well as custom VLSI circuitry, in combination with a general purpose DSP circuit, may also be provided.

[0030] The logical operations of the various embodiments may be implemented as: (1) a sequence of computer-implemented steps, operations, or procedures (generally "instructions") running on a programmable circuit within a general-use computer, (2) a sequence of computer-implemented steps, operations, or procedures running on a specific-use programmable circuit, or (3) interconnected machine modules or program engines within the programmable circuits. The media device 100 shown in FIG. 1A can practice all or part of the disclosed methods, or can be a part of the disclosed systems, or can operate according to instructions in the disclosed computer-readable storage media. Such logical operations can be implemented as modules configured to control the processor 120 to perform particular functions according to the programming of the module. For example, FIG. 1A illustrates three modules Mod1 162, Mod2 164, and Mod3 166 which are modules configured to control the processor 120. These modules may be stored on the storage device 160 and loaded into RAM 150 or memory 130 at runtime or may be stored as would be known in the art in other computer-readable memory locations.

[0031] Content delivery describes the delivery of media "content" such as audio or video or computer software and games over a delivery medium such as broadcasting or the Internet. Content delivery generally has two parts: delivery of finished content for digital distribution, with its accompanying metadata; and delivery of the end product to an end-user.

[0032] As used herein, "streaming media" are media received by and presented to an end-user while being delivered by a streaming provider using Adaptive Bit Rate ("ABR") streaming methods. The name refers to the delivery method of the medium rather than to the medium itself.

The distinction is usually applied to media that are distributed over telecommunications networks, e.g., "on-line," as most other delivery systems are either inherently streaming (e.g., radio, television) or inherently non-streaming (e.g., books, video cassettes, audio CDs). Hereinafter, on-line media and on-line streaming using ABR methods are referred to as "media" and "streaming."

[0033] ABR streaming is a technology that works by breaking the overall media stream or media file into a sequence of small HTTP-based file downloads, each download loading one short segment of an overall potentially unbounded transport stream or media elementary streams. As the stream is played, the client (e.g., the media player) may select from a number of different alternate streams containing the same material encoded at a variety of data rates, allowing the streaming session to adapt to the available data rate. At the start of the streaming session, the player downloads a manifest containing the metadata for the various sub-streams which are available. Since its requests use only standard HTTP transactions, ABR streaming is capable of traversing a firewall or proxy server that lets through standard HTTP traffic, unlike UDP-based protocols such as Real-Time Transport Protocol. This also allows a content delivery network to readily be implemented for any given stream. ABR-streaming methods have been implemented in proprietary formats including HTTP Live Streaming by Apple, Inc., and HTTP Smooth Streaming by Microsoft, Inc. ABR streaming has been standardized as ISO/IEC 23009-1, Information Technology--Dynamic adaptive streaming over HTTP: Part 1: Media presentation description and segment formats.

[0034] An increasing number of video playback devices, such as the Apple iPad, prefer video content to be delivered via ABR streaming rather than streamed continuously. The iPad, using Apple's HTTP Live Streaming format, receives the manifest as an m3u8 file that contains links, media uniform resource identifiers (URIs), to each of the segments or "chunks" of video content,

and processes the manifest file to retrieve and play back each media segment in turn. In this disclosure, "HLS" represents the range of protocols that segment media content and employ a playlist to manage playback.

[0035] Having disclosed some components of a computing system, the disclosure now turns to FIG. 1B, which illustrates an example media-streaming system embodiment 1000. The communications between the entities depicted in FIG. 1B can occur via one or more wired or wireless networks. Further, the devices can communicate directly, via the World Wide Web, or via an application-programming interface. A playback device 1002, such as a tablet device, smartphone, desktop or portable computer, set-top box, Internet-enabled television, media center PC, or any other suitable device, first makes a request to a media server 1004 for playback of media content. Typically, the media server 1004 resides in a network, such as the Internet or a third-party content-distribution network.

[0036] In HLS, the media server 1004 receives the request and generates or fetches a manifest file 1006 to send to the media device 1002 in response to the request. Example formats for the manifest file 1006 include the m3u and m3u8 formats. An m3u8 file is a specific variation of an m3u encoded using UTF-8 Unicode characters. The m3u file format was initially used in the WINAMP Media Player for audio-only files but has since become a de facto playlist standard on many media devices for local or streaming media, including music and other media types. Many media devices employ variations of the m3u file format, any of which can be used according to the present disclosure. A manifest file can include links to media files as relative or absolute paths to a location on a local file system or as a network address, such as a URI path. The m3u8 format is used herein as a non-limiting example to illustrate the principles of manifest files including non-standard variants.

[0037] The manifest file 1006 includes a list of Uniform Resource Locators ("URLs") to different representations of the requested segmented media content. Before or at the time of the request, the media server 1004 generates or identifies the media segments of the requested media content as streaming media content 1010. The media segments of the streaming media content 1010 are generated, either by the media server 1004, by the content producer, or by some other entity, by splitting, transcoding, or transrating the original media content 1008. Upon receiving the manifest file 1006, the media device 1002 can fetch a first media segment for playback from the streaming media content 1010, then, during playback of that media segment, fetch a next media segment for playback after the first media segment, and so on until the end of the media content.

[0038] Referring to FIG. 2, a system 200 is shown that includes client components 210 and server components 250 in communication with each other and the message flows for typical adaptive streaming. Flows related to security between the client components 210 and server components 250 have been omitted for clarity.

[0039] Client components 210 may include an application graphic user interface ("App GUI") 220 and ABR player 230. Server components 250 may include a content server 260, which may be configured to store or produce multi-bitrate media steams and manifest files.

[0040] In a first step 205, a user navigates through movie listing and selects an audiovisual media asset for viewing. In some embodiments, the audiovisual media asset is linked to a URL pointing to a high-level playlist.

[0041] In a next step 215, the ABR player 230 requests a high-level manifest file for the audiovisual media asset that includes information about the ABR profiles and links to the manifests corresponding to each media bandwidth.

[0042] In a next step 225, the ABR player 230 looks at the high-level manifest or master playlist and either starts by requesting the first manifest file, the lowest bandwidth manifest file, or optionally may do some bandwidth availability estimation and select the corresponding bandwidth manifest file.

[0043] In a next step 235, ABR player 230 requests a 2nd-level manifest for the corresponding bandwidth. In a next step 245, ABR player 230 determines the media segment file in the 2nd-level manifest.

[0044] In a next step 255, ABR player 230 requests media segment files in succession. In a next step 265, ABR player 230 continuously monitors the media buffer fullness to determine if it is necessary to request lower or higher bandwidth media-segment representations. For example, if the bandwidth conditions change, the player selects the corresponding bandwidth manifest file and selects media segments in succession.

[0045] In a next step 275, when the end of the manifest file is reached, ABR player 230 signals the App GUI 220 that playback of the audiovisual media asset is complete. The signaling that the stream playback is complete is shown as step 285.

[0046] As explained above, the content server 260 services multiple media devices or ABR players 230, transmitting multiple sets of video streams to the media devices. If a media device requests a set of video streams with high resolution and high bit rate, a large portion of content-server resources or network bandwidth may have to be allocated in order to service that media device. Consequently, the other media devices serviced by the content server may experience degraded performance such as intermittent interruptions in the transmission of video streams.

[0047] Consequently, a server-side approach that is able to efficiently utilize multiplexing techniques with ABR streaming for multiple clients is highly desirable. For example, a system and

method that can determine the resources to allocate to each client, e.g., by not giving the client multiple options, while maximizing the performance for each client, offers many advantages over the current client-driven model.

[0048] FIG. 3A depicts a video streaming system 300 according to some embodiments. Video streaming system 300 includes a content server 305, a network 315, a set of media devices 320, and a transcoder element 322. Content server 305 may transmit video content or sets of video streams to media devices 320 via network 315. A set of video streams may be for a media program, such as a movie, a television program, etc. Each video stream in a set of video streams may be a short segment of video (e.g., two second, ten seconds, etc.). A set of video streams may include thousands of video streams for a media program, such as a two-hour movie. As used herein, encoded content such as video transport or elementary stream may be divided into fixed-duration segments (e.g., chunks). The segments or chunks are typically between two and 10 seconds in duration, although they may be longer or shorter. In some embodiments, shorter segments reduce coding efficiency while larger segments impact speed to adapt to changes in network throughput. In some embodiments, the video and audio transport stream is composed of 188-byte transport packets which are grouped together into HLS chunks or segments. For Microsoft HTTP Smooth Streaming, however, the video and audio elementary streams are grouped into separate data blocks, chunked into file fragments, and indexed with the MP4 or ISO-MBFF "boxes" or "atoms" to hint to the player how to find samples (coded video and audio frames) in these containers.

[0049] The sets of video streams may be provided to content server 305 from transcoder element 322. Transcoder element 322 includes a number of transcoder resources 323 where each transcoder resource provides a set of video streams having unique encoding parameters (e.g., a bit rate, a resolution, etc.). Network 315 may include the Internet, various intranets, etc. Network 315

may include wired links and wireless links. It will be understood that the various references made herein to "media" and "video" include both video content and audio content.

[0050] Content server 305 includes a set of processors 305a and a non-transitory computer-readable storage medium (memory) 305b. Memory 305b may store instructions, which the set of processors 305a may execute to carry out various embodiments described herein. Content server 305 includes a number of computer devices that share a domain. Content server 305 also includes a schedule module 305c, described in further detail in reference to FIG. 3B.

[0051] Referring now to FIG. 3B, an example schedule window 310 provided by schedule module 305c is shown. Ideally, content server 305 maximizes network utilization and maintains comparable video quality levels across the media streams sent to media devices 320. Content server 305 may accomplish this goal in part by utilizing schedule module 305c. Schedule module 305c may construct and implement schedule window 310 based on one or more factors it receives, determines, or assumes from one or more transcoders, content-management services, content-delivery servers, or media devices in communication with content server 305. In some embodiments, schedule module 305c uses metadata in the form of Need Vectors corresponding to the transcoded video (which may come from upstream transcoders or content servers) as well as information such as the client's subscription parameters (e.g., service level agreement for the streaming service, e.g., did the client pay for SD or HD service or gold/silver/bronze, etc., which may come from operator-managed content-directory services or fulfillment managers), client device information reported directly from the client media device (screen size, A/V decoder capabilities), etc.

[0052] As shown, schedule window 310 includes a plurality of client or media device individual schedules 320. For example, client #1 is assigned schedule 320a, client #2 is assigned

schedule 320b, and client #N is assigned schedule 320n. Within each individual schedule 320 is the order of segments to be delivered to the client. For example, client #1 shows that segments chunk $S_1 + 1$, chunk $S_1 + 2$, chunk $S_1 + 3$, . . . , chunk $S_1 + K_1$ are to be delivered over the schedule-window time period T_{SW} .

[0053] As is known, a media-device receive buffer is a necessary component to accommodate network jitter. In other words, timing constraints on the arrival of video data for a media device may be relaxed based on the buffer fullness. For example, if media device has T amount of video data segments in its buffer, the arrival of next video segment can be delayed by $(T - T_{chunk})$ seconds before the media device runs out of media data or underflows. In some embodiments, the player has T seconds of media presentation data in its buffer. If it does not download any new data over the next T seconds, then when T seconds expire, the player's buffer will underflow, and it will stop playing audio and video. In some embodiments, the player must download new segments of T_{chunk} seconds long at a rate that must equal the rate at which video is decoded and rendered, on average, using the T-second buffer to absorb delays.

[0054] Additionally, the media-device receive buffer also provides the opportunity to relax the timing constraints on segment or chunk downloads without affecting media-device user experience. Thus, a solution that schedules chunks and determines their bitrates for each client for a period of time defined by a schedule window (such as schedule window 310) is desirable. Under the Schedule-Window model, statistical multiplexing may be described as the following: Given a predefined period of time called the Schedule Window with window size TSW, and N clients playing different HLS programs, determine optimal rates for each HLS client that maximizes the utilization of network bandwidth while maintaining good video quality for each client.

[0055] As indicated in FIG. 3B, the problem for selecting different rates and number of segments for media devices becomes a scheduling problem. In the time period defined as T_{SW} , K_i number of chunks are scheduled for media device i . K_i is dependent on the media device status and on the value of the size of the media-device buffer minus the fullness of the media-device buffer, e.g., the available buffer or $(T_{MAX} - T_{buffer})$. To select the optimal bitrates for scheduled chunks, content server 305 determines which overall quality level Q will be supported based on network bandwidth.

[0056] The media-device status and available buffer may be determined using an HLS client model.

[0057] Without wishing to be bound by any particular theory, it has been discovered that HLS client behavior can be characterized and thus is predictable. For example, when the client starts to play a stored HLS program, it first reads a manifest file (playlist file) from an HLS server with specified URI, parses the content of the file, and starts to request HLS chunks sequentially starting with the chunk of lowest sequence number (as described in reference to FIG. 2). Throughout the playing of the HLS program, two playback stages are observed and defined as Buffering stage and Playback stage. FIG. 4 is an example graph 400 showing the two stages: Buffering stage 410 and Playback stage 420. When the client starts, it will be in the Buffering stage 410 with the HLS client buffer empty, and it tries to fill up its HLS Client Buffer by requesting segments or chunks immediately after it finishes downloading the previous chunk. When the client fills up its HLS Client Buffer, it moves into the Playback stage 420. In this stage, the client will fetch one HLS chunk during chunk duration. In other words, the overall download chunk speed matches its real-time playback speed.

[0058] To predict the client behavior and learn its status, content server 305 builds a state-based HCM for each client or media device. The HCM provides information on whether a client is in Buffering stage or Playback stage. HCM also offers a means to estimate the fullness level of HLS client buffer. FIG. 5 illustrates an example HCM block diagram 500.

[0059] HCM block diagram 500 includes a player state 510. Player state 510 provides the current state of HCM, e.g., play, seek, pause, resume. Also included with player state 510 is media time 520. Media time 520 is the media timestamp of current playing media frame.

[0060] HCM block diagram 500 also includes buffer fullness 530, buffer size 540, and current request chunk 550. Buffer fullness 530 is the number of chunks downloaded into HLS client buffer and not yet consumed by the media device. Buffer size 540 is the total size of HLS Client buffer, and current request chunk 550 is the sequence number of current requesting chunk by HLS client.

[0061] Referring now to FIGS. 4 and 5, a few variables which are used to construct the HCM are described herein. For example, T_{start} indicates the system time when client finishes downloading the first media chunk and starts to play video. $T_{current}$ indicates the system time when client requests current request chunk 550. T_{start_pts} represents the presentation time stamp ("PTS") of first video frame in the first media chunk. $T_{current_pts}$ represents the PTS of first video frame in current request chunk 550.

[0062] Buffer fullness 530 is represented by the filling of the buffer (relative elapsed time) minus the draining of the buffer (actual elapsed time). For example, let T_{buffer} represent Client buffer fullness 530 measured in time. The video data measured in time (seconds) streamed to HCM, $T_{filling}$, can be written as:

$$T_{filling} = T_{current_pts} - T_{start_pts} \quad (\text{Equation 1})$$

While the video data measured in time (seconds) consumed by HCM, $T_{draining}$, can be written as:

$$T_{draining} = T_{current} - T_{start} \quad (\text{Equation 2})$$

T_{buffer} can then be calculated as the following:

$$T_{buffer} = T_{filling} - T_{draining} = (T_{current_{pts}} - T_{start_{pts}}) - (T_{current} - T_{start}) \quad (\text{Equation 3})$$

[0063] Let T_{MAX} represent the size of HLS client buffer size 540, then HLS client is in Buffering Stage if $T_{buffer} < T_{MAX} - T_{chunk}$, otherwise it operates in normal Playback Stage.

[0064] As explained above with reference to FIG. 3, to select the optimal bitrates for scheduled chunks, content server 305 determines which overall quality level Q will be supported based on network bandwidth. To achieve a quality level Q , the encoded bitrate can be calculated based on its Need Parameter Vector ("NPV").

[0065] The NPV is a composite of several factors including Video Complexity ("VC"), Device Profile, Service Priority Level, Codec Profile, etc. VC is derived from video content as an estimation of the complexity level of the video content. Given an NPV, content server 305 calculates what bit rate is needed to obtain a targeted level of quality. This information can be provided by a family of curves of Bit Rate vs. NPV for constant quality, such as shown in FIG. 6. Device Profile can include the screen size, codec profiles (e.g., MPEG2 and AVC for video or Dolby AC-3 or AAC for audio) that are supported, software/hardware/firmware versions, OS versions, player-application version, etc. Service Priority Level can include parameters such as those included in service level agreements such as guaranteed bandwidths supported or access to high definition videos versus standard definition associated with higher and lower cost subscription services, respectively.

[0066] NPV can be computed based on complexity of content for each segment, with the curve of Bit Rate vs. NPV being linear for any given quality level. This means if NPV for program A is

twice the NPV for program B, it will take program A twice the bit rate as program B to maintain a similar video quality.

[0067] An encoded bitrate can be calculated based on its NPV (in bytes per second) as shown in the following equation:

$$\text{byterate}(\text{NPV}) = \alpha \times \text{NPV} \quad (\text{Equation 4})$$

where α represents the quality level of the encoded video and is a scalar and when normalized it is in the range of 0 to 1. Higher α indicates higher video quality. The total scheduling budget in bytes for a single client over K chunks then is:

$$\text{budget} = \sum_{j=1}^K \alpha \times \text{NPV}(j) \times T_{\text{chunk}} \quad (\text{Equation 5})$$

[0068] For any given budget, content server 305 may determine α for a single client, thus the achievable video quality under the given budget. Next, content server 305 expands the computation of α over multiple HLS clients. In order to maintain comparable quality for each HLS client, the same α value may be selected. Therefore, the total bytes for all clients during the scheduling window ("SW") can be calculated as:

$$\begin{aligned} \text{SUM}_{\text{bytes}} &= \sum_{i=1}^{N_i} \sum_{j=1}^{K_j} \text{byterate}(i, j) \times T_{\text{chunk}} \\ &= \sum_{i=1}^{N_i} \sum_{j=1}^{K_j} \alpha \times \text{NPV}(i, j) \times T_{\text{chunk}} \end{aligned} \quad (\text{Equation 6})$$

where N is the number of HLS clients, S_i is the sequence number of last downloaded chunk for client i, and K_i is the number of chunks scheduled for HLS client i. For simplification, all chunks may be assumed to have the same duration T_{chunk} .

For a fixed-rate channel, the total available network bandwidth is defined as "BW" in bytes/second. So the total budget for all HLS clients is $\text{BW} \times T_{\text{SW}}$. Therefore, in order to send

SUM_{bytes} during SW with window size defined as T_{SW}, it must satisfy the equation below:

$$SUM_{bytes} \leq BW \times T_{SW} \quad (\text{Equation 7})$$

[0069] With the known NPV values for scheduled chunks, content server 305 can calculate α as follows: (Equation 8)

$$SUM_{bytes} = \sum_{i=1}^{N_i} \sum_{j=1}^{K_j} \alpha \times NPV(i,j) \times T_{chunk} \leq BW \times T_{SW}$$

then

$$\alpha \leq \frac{BW \times T_{SW}}{T_{chunk} \times \sum_{i=1}^{N_i} \sum_{j=1}^{K_j} NPV(i,j)} \quad (\text{Equation 9})$$

[0070] While not wishing to be bound by any particular theory, it has been discovered that when a media device or client is in normal Playback Stage, it will only try to download T_{SW}/T_{chunk} number of chunks. However, during its initial Pre-Buffering Stage, it will try to download as many chunks as possible until its client buffer is full. To determine the number of chunks for a client during SW, it is important to estimate client buffer fullness and figure out the number of chunks to download during SW. To estimate the client buffer fullness, content server 305 utilizes Equation 3. Based on the value of T_{buffer}, K (e.g., scheduled number of chunks to download) is determined according to:

$$K = (T_{SW} + \beta \times (T_{MAX} - T_{buffer})) / T_{chunk} \quad (\text{Equation 10})$$

where β is the weight factor for buffering stage. The larger β indicates shorter time for filling up client buffer, thus shorter buffering stage.

[0071] With the known NPV values and scheduled number of chunks for each client, content server 305 determines α with Equation 9 and determines the achievable optimal quality for all HLS

clients. Subsequently, content server 305 selects the bitrate for scheduled chunk that is closest to $\alpha \times \text{NPV}$.

[0072] FIG. 7 is an example block diagram of a system 700 implementing an HCM 710 and schedule queue 720. In some embodiments, HCM 710 indicates a program ID 730 (e.g., what show it is requesting) and the corresponding program 745 in the content server 740. HCM 710 tracks the associated HLS client state 735 (e.g., running, seeking, pausing, resuming, etc.). HCM 710 also estimates the buffer time 750 for its associated HLS client. In some embodiments, schedule queue 720 stores the latest schedule results from schedule window (not shown) and maintains a pointer to a current serving chunk number or request chunk number 755.

[0073] In some embodiments, a client (e.g., associated with HCM 710) will request a program 745 from content server 305. Rather than allow the client to select the bit rate at which the client receives the program content, the content server 305 publishes a single available bit rate for the client to receive. Thus, in FIG. 7, the schedule queue 720 indicates that the client will only receive chunk #1 of bitrate #2 stream for that particular segment request. As content server 305 determines a different assignment for a particular client, the schedule queue 720 reflects the assignment change almost immediately, e.g., with little or no delay. In some embodiments, the assignment from content server 305 may change from segment to segment. In some embodiments, this flexibility in changing the bit rate the client receives can prevent or minimize the number of segments that a client would re-buffer.

[0074] Still referring to FIG. 7, in addition to normal playback operations, the HLS client may also perform some common operations which can affect the scheduling decisions. For example, the HLS client may switch to different video content (e.g., channel change). The HLS client may pause the content and after a while resume the playback. The HLS client may seek to a specific

position and start to watch from the new position. The HLS client may complete playback and leave. All these client events indicate a transition period, requiring a modification to the bandwidth assigned to the HLS client.

[0075] As used herein, a channel change is when the same client asks for program different from the program it is currently playing. Pause and Resume is when the client does not request any chunk download for a prolonged period of time. However, it then requests the chunk with sequence number following the chunk it downloads before the prolonged period of time. Seek is when a client is requesting a chunk with a sequence number not following the previous chunk it downloads. Complete is when a client downloads the last chunk of a program.

[0076] In some embodiments, additional refinements in the form of an adjustment factor may be used for assigning bandwidth to clients. For example, for certain video programs, they may not offer multiple streams with the dynamic range of bit rates that satisfy $\alpha \times NP \subset \{\text{bitrate}\}$. In this case, content server 305 must handle cases for either over budget or under budget bit-rate assignments. As used herein, over budget is when the calculated $\alpha \times NP$ is much lower than the available lowest bitrate (e.g., the bit rate of a chunk is greater than the allocated bandwidth for the client). Similarly, under budget is when the calculated $\alpha \times NP$ is much higher than the available higher bitrate (e.g., the bit rate of a chunk is less than the allocated bandwidth for a client).

[0077] In some embodiments, content server 305 may remove either under budget or over budget client out of the Equation 6. For example, assuming Client P is over/under budget, the budget for Client P is calculated as:

$$\text{budget}(P) = \sum_{i=1}^P \text{ChunkSize}_{\text{lowest/highest}}(i) \quad (\text{Equation 11})$$

And

$$SUM_{\text{bytes}} = \sum_{i=1, i \neq P}^N \sum_{j=1}^{K_j} \alpha \times NPV(i, j) \times T_{\text{chunk}} \leq \quad (\text{Equation 12})$$

$$(BW \times T_{SW} - \text{budget}(P))$$

[0078] In some embodiments, if the bit streams are pre-encoded and stored, then it may happen that the available bit rates are over or under budget. In this case, a transcoder may be used to generate the bit streams close to $\alpha \times \text{NPV}$. Content server 305 can work closely with the transcoder by computing α and adjusting the transcoder's output bitrate as $\alpha \times \text{NPV}$ at each chunk interval.

[0079] As provided above, a transition period occurs when a client playback approaches the end of program. This is because when client playback approaches the end of the program (e.g., the number of available chunks to schedule is smaller than K), there may not be enough chunks for downloading to fill the schedule window. In this case, Equation 13 may be used to compensate for the gap in the schedule window.

[0080] In Equation 13, M represents the number of total remaining chunks for download, and K represents the number of scheduling chunks for Schedule Window. In this case, $M < K$. To assure scheduled chunks arrive at client in time, a modification is made for Equation 5. Instead of using sum of NPVs of M chunks, content server 305 determines the sum as follows:

$$\frac{K}{M} \times \sum_{i=1}^M \text{NPV}(i) \quad (\text{Equation 13})$$

This assigns a higher bitrate to the client such that it completes the download of $M < K$ chunks in less than the T_{SW} schedule window time. The client may terminate the session at this time which triggers the calculation of a new Schedule Window applied to the remaining clients communicating with the content server.

[0081] Lastly, in some embodiments, a network bottleneck issue on the client end may occur. This may occur when the end network on client side (e.g., 802.11g/n Wi-Fi home network) significantly slows down, and its bandwidth drops below the scheduled bit rate. To address this issue, a history of download rate of previous HLS chunks can be stored. Thus, in some

embodiments, the content server 305 checks if a client's previous download rate is higher than the scheduled weighted network bandwidth before sending the chunk requested by HLS client. If not, a new lower rate is calculated based on previous download rate history and HLS client information.

[0082] In further embodiments, client feedback can further enhance the performance of server 305. In some embodiments, as shown in FIG.8, a server 305 can detect pause playback executed by a user. Actions by a user watching a client video, such as pause playback or fast-forwarding, can cause buffer depth to change such that an adaptive server 305 can be unaware of the changing depth. In such embodiments, a client can inform an adaptive server 305 that a client buffer is empty. In the embodiment of FIG.8, a client buffering state 812 can provide information of user action, e.g., pause playback or fast-forwarding.

[0083] To aid the adaptive server schedule window calculations and chunk delivery, a client media playback application can send a message to a server 305, such as, but not limited to HTTP REST messaging, that can inform the server that the client user has paused media playback or performed a “trick-play” operation (e.g., fast-forward or rewind). Pausing playback can result a client to cease downloading of media chunks for the pause duration. If a server 305 is unaware of this client action, a server 305 can perceive such an action as a loss of network connectivity to the client as opposed to paused playback. A server 305 can continue to schedule chunks for a client during a schedule window 310 if a network outage has occurred, whereas a server 305 can stop scheduling chunks to a client if a client is in a paused state. During fast-forward or rewind, a client can either attempt to download full media segments faster than real-time in order to playback portions of the segments faster than real-time or download special media segments, such as, but not limited to HLS I-FRAMES-ONLY segments, that have only a KEY-FRAME or Instantaneous Decoder Refresh (IDR) video frame.

[0084] If a client buffering state 812 indicates pause playback, as detected at a server 305 by reception of a message from a client, a server 305 can terminate the current schedule window 310, reduce the number N of HLS clients by one to $N-1$, and recalculate a new schedule window 310 for the remaining $N-1$ clients to reallocate bandwidth. When a client exits the pause state to resume playback, a client can send a message to a server 305 announcing it will be downloading more chunks, or make a contiguous media segment request from where it last paused, and a server 305 can halt a current schedule window 310, and increase the number of clients from $N-1$ back to N , recalculate a new schedule window bandwidth allocation for the N clients. This feedback from the client can prevent a server 305 from needlessly allocating bandwidth to a paused client that is unused during its paused state, resulting in possible underutilization of the network bandwidth by the non-paused $N-1$ clients.

[0085] If a client buffering state 812 indicates fast-forwarding, a client can send a message to a server 305 indicating the transition to fast-forward, and a server 305 can prepare a new schedule window 310 that can either allocate more bandwidth to that client to allow for chunk downloads at faster than real-time playback speed or schedules for downloading of smaller I-FRAMES-ONLY or IDR video-only chunks. When a client exits the fast-forward mode, a client can send a message to a server 305 indicating return to real-time playback and a server 305 can respond by recalculating a schedule window 310 to allow for real-time playback chunk download rates. This feedback can be required to facilitate proper playback during fast-forward by a client since a server 305 can schedule chunk downloads at real-time rate during SW phase. A client might play out content in its buffer very rapidly in this state and eventually underrun if a server 305 does not adjust its schedule window chunk delivery to keep up with this pace.

[0086] In some embodiments, feedback of client capabilities can further refine the complexity as calculated by a content server 740. As shown in FIG. 9, an HCM 710 can attempt to schedule A/V chunks in a SW time period to be downloaded by a particular client based on the VC of a content stream. For content having a high VC, a server 740 can attempt to download a higher quality (e.g., higher bitrate and video resolution) set of content variant A/V chunks to a client.

[0087] In the embodiment of FIG. 9, a client can *a priori* inform an adaptive server 740 of client hardware and/or software limitations 937, e.g., that a client cannot perform video decoding of video above a certain profile having a given resolution, frame rate, and/or bitrate. A server 740 can then configure its scheduling algorithm to cap a video stream at an A/V profile level that a client supports. As an example, a client app provisioned by a service operator may constrain or cap the media delivery to its clients to a qHD mode (960x540 at 30 frames/second) while an adaptive server 305 has variants for HD formats (720p60 or 1080p30).

[0088] When an adaptive server 305 is made aware of such limitations through reception of a message from the client during, by non-limiting example, client application start-up, then during schedule window calculations for that client, an HCM server will not schedule content at resolutions and bit rates greater than qHD even if a Need Vector dictated that the content should be delivered at higher bitrates. This can free up bandwidth in a schedule window 310 to assign to the remaining clients.

[0089] By way of non-limiting example, a client can be constrained in its playback capabilities by its screen-size, CPU load memory space, supported A/V Codecs, Wi-Fi download bandwidth, and other resource limitations. In other embodiments, an adaptive server 305 can monitor the chunk download rate as it delivers content to a client to estimate the bandwidth available to the client. In some embodiments, as shown in FIG. 10, a client can provide feedback of a client's

internal download rate measurements to refine an available bandwidth estimate. However, there can be situations where a server 305 measurement technique yields a different bandwidth value than that experienced by a client. By way of non-limiting example, if there is a TCP proxy between a server 305 and a client that performs TCP fragment ACK directly to a server 305 on one side of the proxy while delivering chunks with different TCP ACK timings on the client side, the measurements will disagree.

[0090] In some embodiments, as shown in FIG.10, a client can send internal download rate measurements 1060 to a server 740, which can then correctly assign media variant bitrates during a schedule window period. By way of non-limiting example, if a schedule window calculation dictates a certain client should receive a high bitrate/resolution A/V variant, and a download bandwidth measurement by a server 740 is greater than this bitrate, the high-rate content download can be scheduled. However, if a client cannot complete download in the given chunk playback duration due to a lower bandwidth availability on its side of the proxy, a client can inform a server 305 and a server 305 can schedule a lower bit rate variant 1070 for download to a client. As depicted in FIG. 10 and in some embodiments, a chunk download can be delivered to a client media player and a chunk download and bandwidth measurement 1060 can be performed. A chunk bandwidth report message 1080 containing data based at least in part on the chunk download and bandwidth measurement 1060 can be delivered to an HCM 710. Based at least in part on the chunk bandwidth report message 1080 and/or chunk download and bandwidth measurement 1060, a desired bitrate for chunk delivery can be selected. In the embodiment depicted in FIG. 10, a lower bitrate chunk can be selected if the client chunk bandwidth is less than a scheduled bitrate.

[0091] By way of non-limiting example, variants available on a content server 740 can consist of bitrate streams 1 through 4 at respective bitrates of 1, 2, 3, and 4 Mbps. Based on an NPV for a

given program during a given schedule window period, scheduled chunks to a client can be chosen from the 4 Mbps stream. However, if a client receiving these chunks can only download at a 2.5 Mbps rate due to Wi-Fi or cellular bandwidth limitations, then the scheduled 4 Mbps chunks will not complete their downloads in real-time and cause the client media buffer to shrink during real-time playback. A client can measure its chunk download rate through any exemplary method, such as, but not limited to, chunk size divided by the HTTP download time, and provide this bandwidth estimate back to a server 740 in an HTTP REST message. An adaptive server 740 HCM 710 can adjust its schedule window calculations so as to not schedule chunk downloads that exceed this bandwidth measurement to that specific client. This can allow that bandwidth-stressed client to maintain its media buffer at a more constant level.

[0092] In view of the many possible embodiments to which the principles of the present discussion may be applied, it should be recognized that the embodiments described herein with respect to the drawing figures are meant to be illustrative only and should not be taken as limiting the scope of the claims. Therefore, the techniques as described herein contemplate all such embodiments as may come within the scope of the following claims and equivalents thereof.

CLAIMS

What Is Claimed:

1. A method comprising:

determining, by a server, bandwidth to allocate to each of a plurality of media devices configured to provide video content using a HyperText Transfer Protocol-based live streaming client model ("HCM"), a bitrate based on a corresponding need parameter vector ("NPV") varied by a scalar quality value for each of the plurality of media devices and client feedback regarding client buffer level and playback state, client hardware capabilities, and client internally measured download rate; and

providing the determined bandwidth to allocate to each of the plurality of media devices, wherein the video content is transmitted in a plurality of segments from the server, and

wherein each segment is transmitted using a variable bitrate from segment to segment.

2. The method of claim 1 wherein the server constructs a state-based HCM for each of the plurality of media devices.

3. The method of claim 2 wherein the HCM provides one or both of the following pieces of information: if a media device is in a buffering state or playback state, and an estimate of a fullness of a media-device buffer.

4. The method of claim 2 wherein the determined bandwidth to allocate to each of the plurality of media devices prevents a media device from buffering content already received from the server.
5. The method of claim 1 wherein the server or a proxy constructs a NPV for each of the plurality of media devices.
6. The method of claim 5 wherein the NPV is based on one or more of the following: video complexity, device profile, service priority level, and codec profile.
7. The method of claim 6 wherein the video complexity is derived from video content as an estimation of complexity of the video content.
8. The method of claim 7 wherein the NPV is computed based on a complexity of video content for each segment.
9. The method of claim 7 wherein the NPV has a linear relationship with the bitrate for a given segment.
10. The method of claim 1 further comprising: transmitting the video content from the server to one or more media devices at a bit rate that is within the confines of the determined bandwidth to allocate for each media device.

11. A method for providing video content using a HyperText Transfer Protocol-based live streaming client model ("HCM"), the method comprising:

- determining a bandwidth to allocate to a plurality of media devices;
- determining a number of active media devices associated with the plurality of media devices to allocate the determined bandwidth;
- determining a need parameter vector ("NPV") for each of the active media devices; and
- assigning a fraction of the bandwidth to each of the active media devices based on the NPV varied by a scalar quality value for each of the active media devices.

12. The method of claim 11 wherein the NPV is based on one or more of the following: video complexity, device profile, service priority level, and codec profile.

13. The method of claim 12 wherein video complexity is derived from video content as an estimation of a complexity of the video content.

14. The method of claim 11 wherein the NPV is computed based on a complexity of video content for each segment.

15. The method of claim 12 wherein the NPV is based on device profile.

16. The method of claim 15 wherein the device profile indicates that an active device is undergoing a transition period requiring a modification to the bandwidth assigned to the active device.

17. The method of claim 16 wherein the transition period is one or more of the following: a channel change, a pause or resume, a seek, a complete, and a join.

18. The method of claim 11 further comprising: determining an adjustment factor for assigning the fraction of the total bandwidth to each active media device.

19. The method of claim 18 wherein the adjustment factor is based on or more of the following: the NPV for an active media device is over budget, the NPV for an active media device is under budget, an active media device completes playback, and a bottleneck occurs at the active media device.

20. A server configured to provide video content, the server comprising:

a set of processors; and

a computer-readable storage medium including instructions for controlling the set of processors to be configured for:

determining bandwidth to allocate to each of a plurality of media devices using a HyperText Transfer Protocol-based live streaming client model ("HCM") and a bitrate based on a corresponding need parameter vector ("NPV") varied by a scalar quality value for each of the plurality of media devices; and

providing the determined bandwidth to allocate to each of the plurality of media devices,

wherein the video content is transmitted in a plurality of segments from the server,

and

wherein each segment is transmitted using a variable bitrate from segment to segment.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4




FIG. 5




FIG. 6





FIG. 7

**FIG. 8**




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/043168

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L1/00 H04N21/20
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L H04N

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/013376 A1 (XU HAIFENG [US] ET AL) 9 January 2014 (2014-01-09)	11-20
Y	claims 1-22 Equation 4,5,6 paragraph [0062] - paragraph [0064] -----	1-10
Y	US 2013/326024 A1 (CHEN MING [US] ET AL) 5 December 2013 (2013-12-05) paragraph [0001] - paragraph [0006] paragraph [0044] - paragraph [0047] paragraph [0058] -----	1-10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

27 September 2017

Date of mailing of the international search report

09/10/2017

Name and mailing address of the ISA/

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

Authorized officer

Nölle, Markus

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/043168

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014013376 A1	09-01-2014	CN 104471955 A	25-03-2015
		EP 2870776 A1	13-05-2015
		KR 20150042191 A	20-04-2015
		US 2014013376 A1	09-01-2014
		WO 2014008200 A1	09-01-2014

US 2013326024 A1	05-12-2013	NONE	
