



US007840693B2

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

(10) **Patent No.:** **US 7,840,693 B2**  
(45) **Date of Patent:** **Nov. 23, 2010**

(54) **SERVING MEDIA ARTICLES WITH ALTERED PLAYBACK SPEED**

(75) Inventors: **Manish Gupta**, Santa Clara, CA (US);  
**John Harding**, San Francisco, CA (US);  
**Stephen Michael Lacy**, Mountain View, CA (US); **Willard Rusch, II**, Woodside, CA (US)

(73) Assignee: **Google Inc.**, Mountain View, CA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 414 days.

(21) Appl. No.: **11/620,354**

(22) Filed: **Jan. 5, 2007**

(65) **Prior Publication Data**

US 2007/0168541 A1 Jul. 19, 2007

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 11/428,319, filed on Jun. 30, 2006.

(60) Provisional application No. 60/756,787, filed on Jan. 6, 2006.

(51) **Int. Cl.**  
**G06F 13/00** (2006.01)

(52) **U.S. Cl.** ..... **709/232**; 709/219; 709/228;  
709/250

(58) **Field of Classification Search** ..... 709/217,  
709/219, 230, 232, 233, 236, 250  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

6,122,662 A \* 9/2000 Emura ..... 709/219  
6,523,027 B1 2/2003 Underwood  
6,792,468 B1 9/2004 Bloch et al.  
6,931,658 B1 \* 8/2005 Kawamura et al. .... 725/94

6,980,594 B2 \* 12/2005 Wang et al. .... 375/240.12  
7,042,844 B2 \* 5/2006 Rasanen et al. .... 370/232  
7,096,271 B1 \* 8/2006 Omoigui et al. .... 709/231  
7,136,875 B2 11/2006 Anderson et al.  
7,558,869 B2 7/2009 Leon et al.

(Continued)

**FOREIGN PATENT DOCUMENTS**

EP 1492021 A2 12/2004

(Continued)

**OTHER PUBLICATIONS**

PCT International Search Report and Written Opinion, PCT/US2007/000403, Jun. 29, 2007, 15 pages.

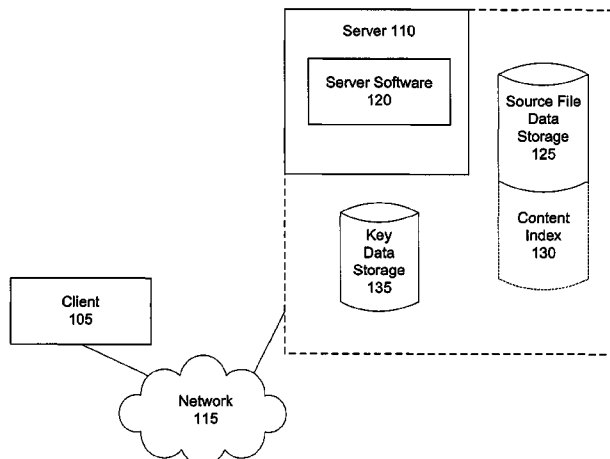
(Continued)

*Primary Examiner*—Viet Vu  
(74) *Attorney, Agent, or Firm*—Fenwick & West LLP

(57) **ABSTRACT**

Systems and methods provide for on-the-fly alteration of apparent playback rate for content sent to a client based on a request from the client to alter the playback speed of the content. The system and method can use existing, open protocols and the media files served are playable by standard media playback clients. The method includes receiving a request for media content from a client, including instructions for adjusting the number of frames of the requested content to affect an apparent change in the playback speed of the content. The adjustment to the frames may include inserting frames, removing frames, and/or inserting markers to trigger the client device to duplicate frames.

**30 Claims, 10 Drawing Sheets**



## U.S. PATENT DOCUMENTS

7,669,121 B2 2/2010 Küllerich  
 2001/0004417 A1 6/2001 Narutoshi  
 2002/0103920 A1 8/2002 Berkun et al.  
 2004/0116183 A1 6/2004 Prindle  
 2004/0187160 A1 9/2004 Cook et al.  
 2004/0267908 A1 12/2004 Doi et al.  
 2004/0267952 A1\* 12/2004 He et al. .... 709/231  
 2004/0267956 A1 12/2004 Leon et al.  
 2005/0028213 A1\* 2/2005 Adler et al. .... 725/89  
 2005/0141886 A1 6/2005 Srinivasan et al.  
 2005/0144284 A1 6/2005 Ludwig et al.  
 2005/0183017 A1 8/2005 Cain  
 2005/0254374 A1\* 11/2005 Lin ..... 369/47.36  
 2006/0002681 A1 1/2006 Spilo et al.  
 2006/0009983 A1 1/2006 Magliaro et al.  
 2006/0221788 A1\* 10/2006 Lindahl et al. .... 369/47.15  
 2006/0235654 A1 10/2006 Parnes  
 2007/0011012 A1 1/2007 Yurick et al.  
 2007/0011343 A1 1/2007 Davis et al.  
 2007/0153916 A1 7/2007 Demircin et al.  
 2007/0168254 A1 7/2007 Steelberg et al.  
 2007/0169146 A1 7/2007 Steelberg et al.  
 2007/0192789 A1 8/2007 Medford  
 2007/0271388 A1 11/2007 Bowra et al.  
 2007/0299870 A1 12/2007 Finch  
 2008/0092159 A1 4/2008 Dmitriev et al.  
 2008/0115161 A1 5/2008 Kurzion  
 2008/0167957 A1 7/2008 Steelberg et al.  
 2009/0131141 A1 5/2009 Walker et al.

## FOREIGN PATENT DOCUMENTS

WO WO 99/24987 A1 5/1999

## OTHER PUBLICATIONS

“Working Together with MXF,” Pro-MPEG Forum, Sep. 3, 2002, pp. 1-4.  
 Almeroth, K. C., et al., “On the Performance of a Multicast Delivery Video-On-Demand Service with Discontinuous VCR Actions,” CiteSeer, 2 pages, [online] [retrieved on Dec. 22, 2006] Retrieved from the Internet: <URL: <http://citeseer.ist.psu.edu/almeroth95performance.html>>.  
 Kwon, J. B., et al., “Providing VCR Functionality in Staggered Video Broadcasting,” CiteSeer, 2 pages, [online] [retrieved on Dec. 22, 2006] Retrieved from the Internet: <URL: <http://citeseer.ist.psu.edu/590569.html>>.  
 “Microsoft Advanced Streaming Format,” Multimedia Wiki, last modified May 2, 2006, 2 pages, [online] [retrieved on Dec. 22, 2006]

Retrieved from the Internet: <URL: [http://wiki.multimedia.cx/index.php?title=Microsoft\\_Advanced\\_Streaming\\_Format](http://wiki.multimedia.cx/index.php?title=Microsoft_Advanced_Streaming_Format)>.  
 Hahn, M., “Uniform Resource Locators,” EDPACS, Dec. 1995, pp. 8-13.

Office Action of the European Patent Office for European Patent Application No. 07709597.4 dated Jan. 30, 2009, 6 pages.

“Your Digital Video Repository,” YouTube, Archived on Jun. 14, 2005, pp. 1-10, [online] [retrieved on Dec. 31, 2009] Retrieved from the internet <URL: <http://web.archive.org/web/20050614234128/http://www.youtube.com>>.

First Examiner’s Report, Australian Patent Application No. 2007205046, Mar. 4, 2010, 3 Pages.

Office Action for U.S. Appl. No. 11/620,346, Mar. 31, 2010, 41 Pages.

Office Action for U.S. Appl. No. 11/620,346, Oct. 13, 2009, 28 Pages.

Office Action for U.S. Appl. No. 11/620,468, Jan. 4, 2010, 42 Pages.

Office Action for U.S. Appl. No. 11/620,386, Jan. 5, 2010, 38 Pages.

“Eyespot Releases Version 2.0 of Online Video Mixing Community at Under the Radar Conference,” Jun. 14, 2006, [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.eyespot.com/press/>>.

“Eyespot—Shoot, Mix, and Share Your Video,” Mar. 13, 2006 [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.solutionwatch.com/326/eyespot-shoot-mix-and-share-your-video/>>.

“Jumpcut Makes Movies Simple,” MiraVida Media, Inc., [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.jumpcut.com/company/>>.

Marlowe, C., “Eyespot in the Mix with Legal Video Mash-Ups,” Washingtonpost.com, Jun. 18, 2006, [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.washingtonpost.com/wp-dyn/content/article/2006/06/18/AR2006061800966.html>>.

Robinson, D., “Jumpcut: Video Editing on the Web,” FlashInsider, Apr. 6, 2006, [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.flahsinsider.com/2006/04/06/jumpcut-video-editing-on-the-web/>>.

“VideoEgg: About Us,” VideoEgg, Inc., 2006 [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://www.videoegg.com/general/about>>.

“VideoEgg Offers New Video Upload Service,” About, Inc., Sep. 19, 2005 [online] [Retrieved on Jun. 29, 2006] Retrieved from the Internet: <URL: <http://desktopvideo.about.com/b/a/203840.htm>>.

Walpole, J., et al., “A Player for Adaptive MPEG Video Streaming Over the Internet,” 12 pages, Department of Computer Science and Engineering, Oregon Graduate Institute of Science and Technology, Portland, Oregon.

\* cited by examiner

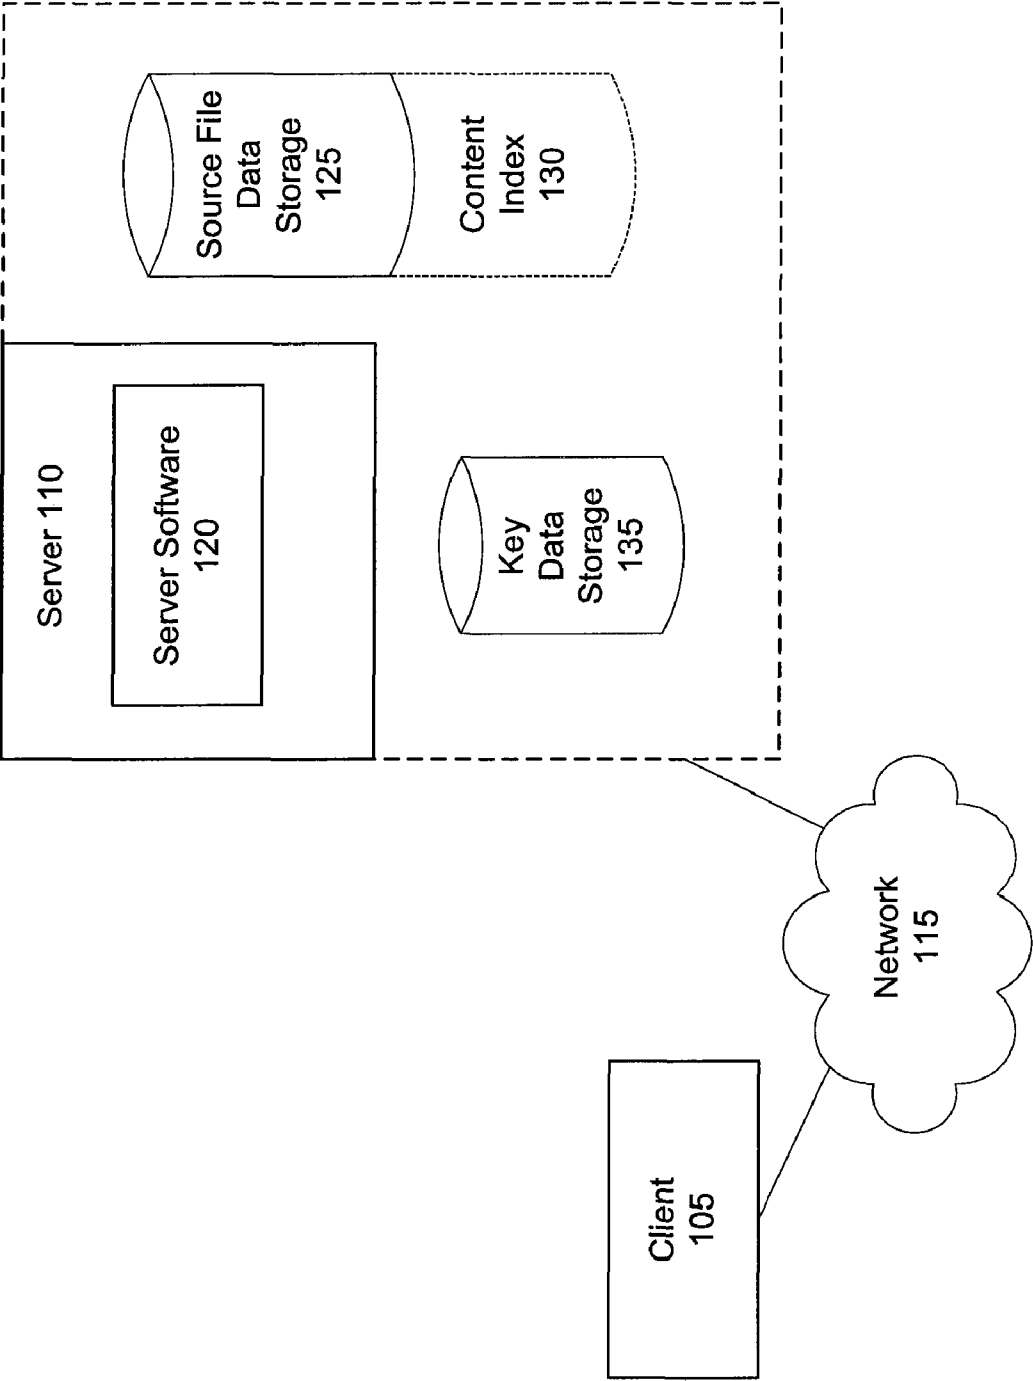


FIG. 1

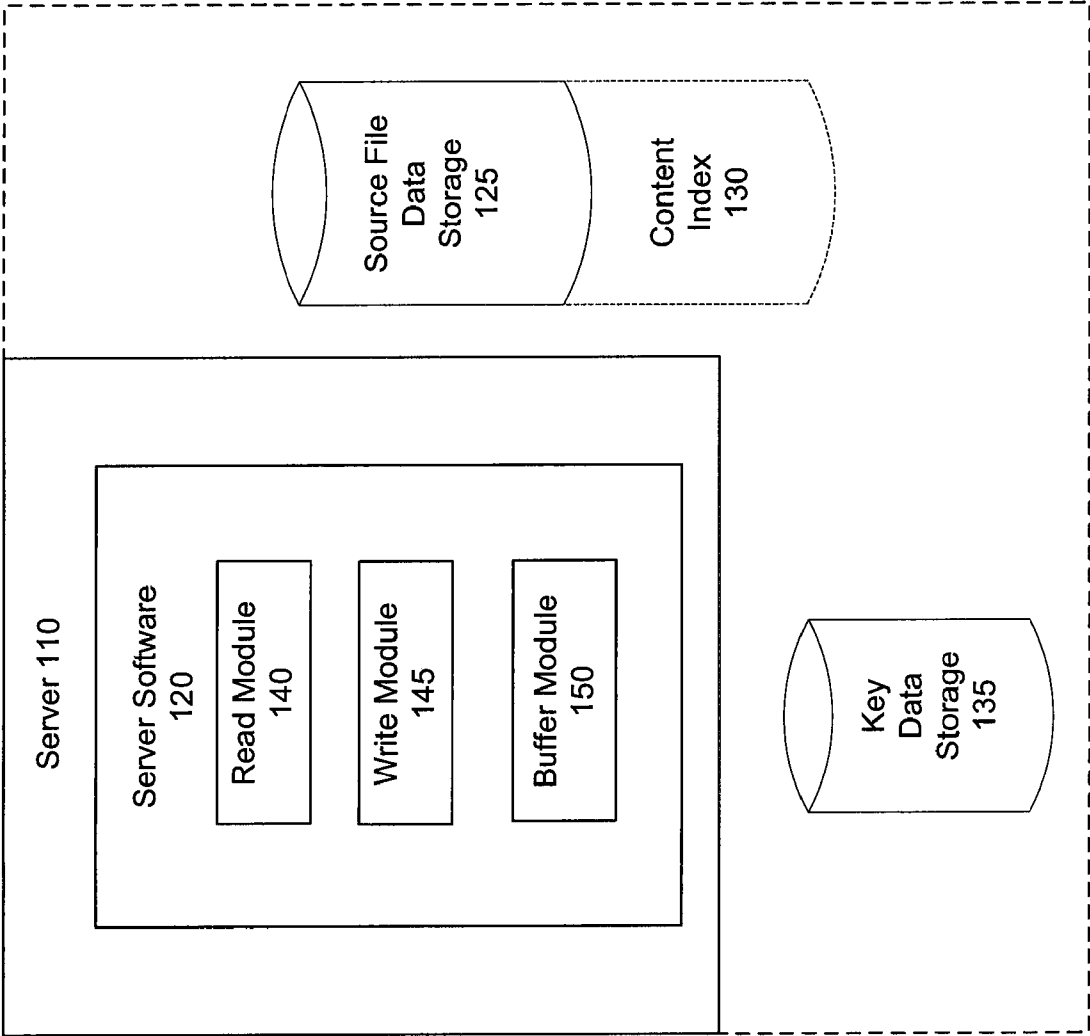


FIG. 2

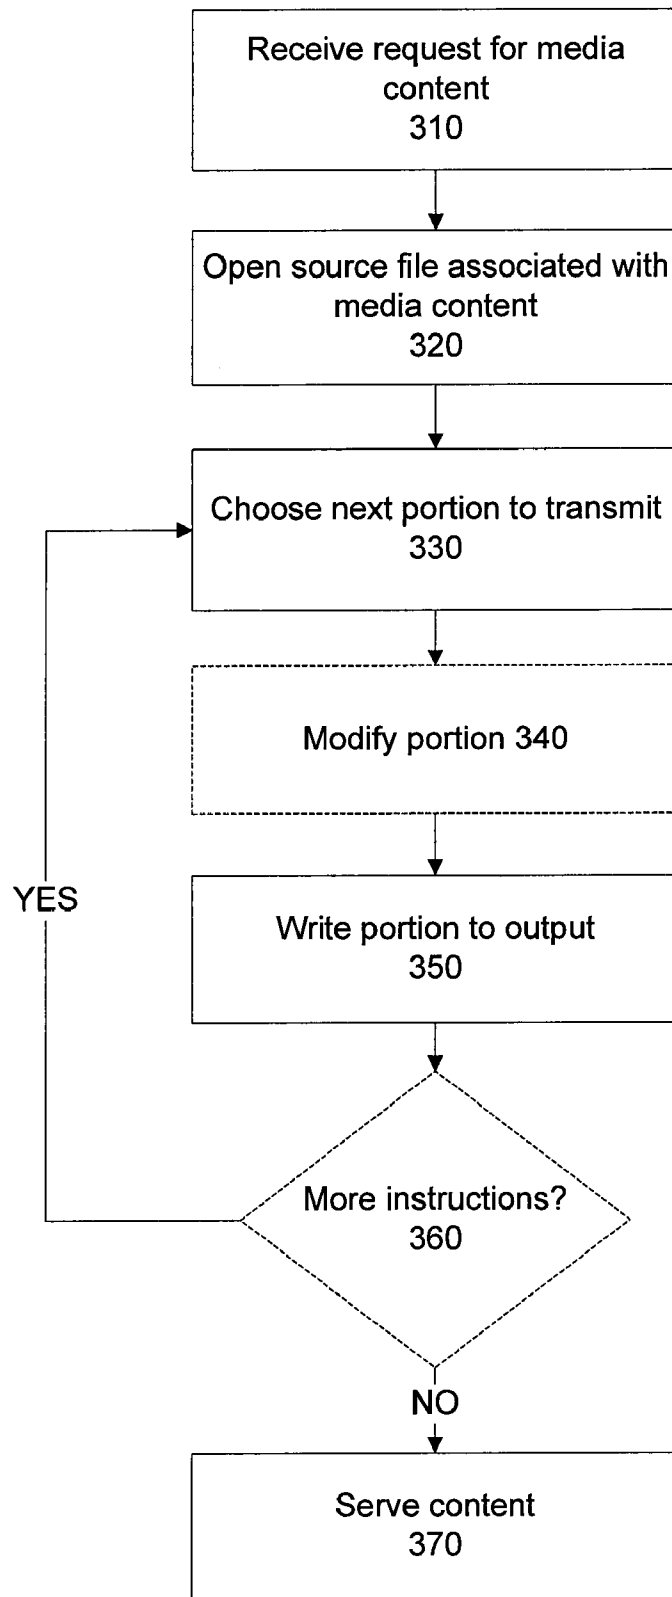


FIG. 3

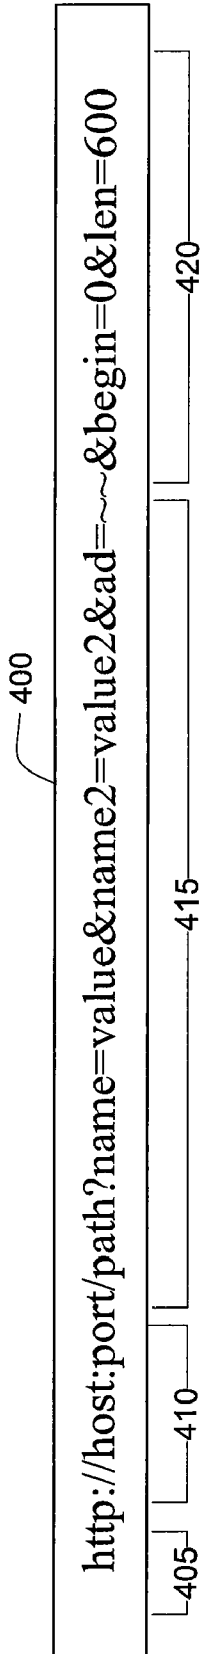


FIG. 4A




FIG. 4B

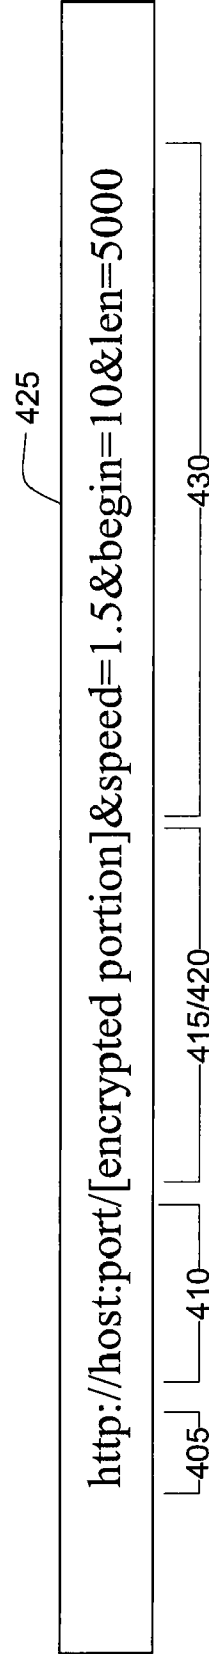


FIG. 4C

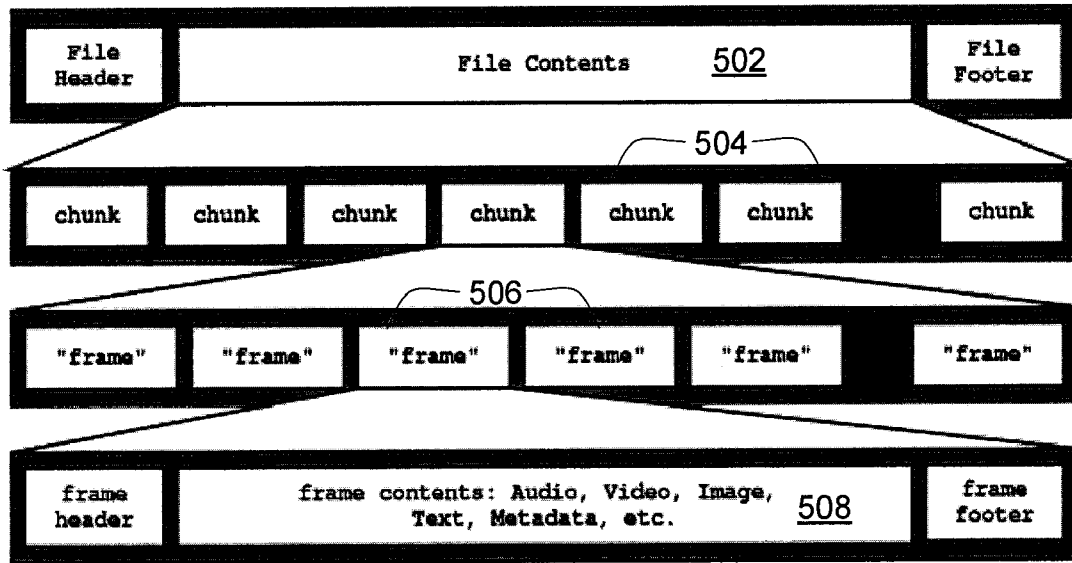


FIG. 5A

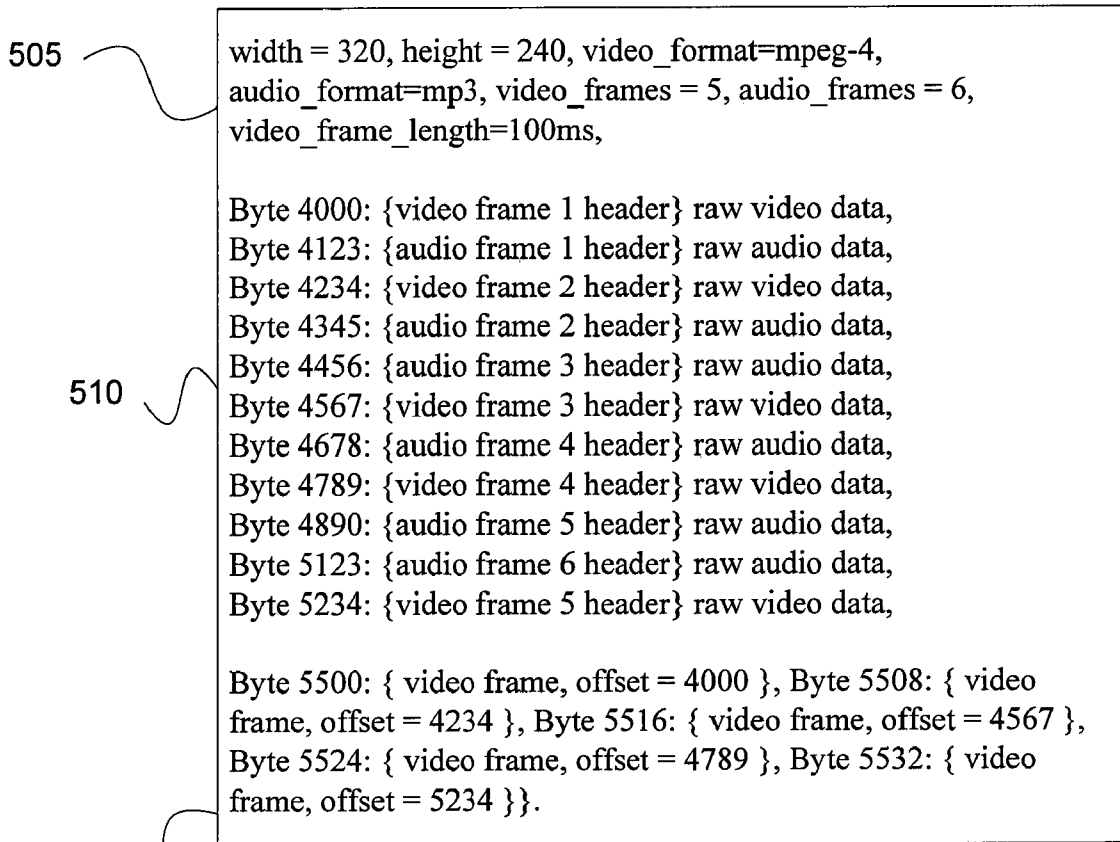


FIG. 5B

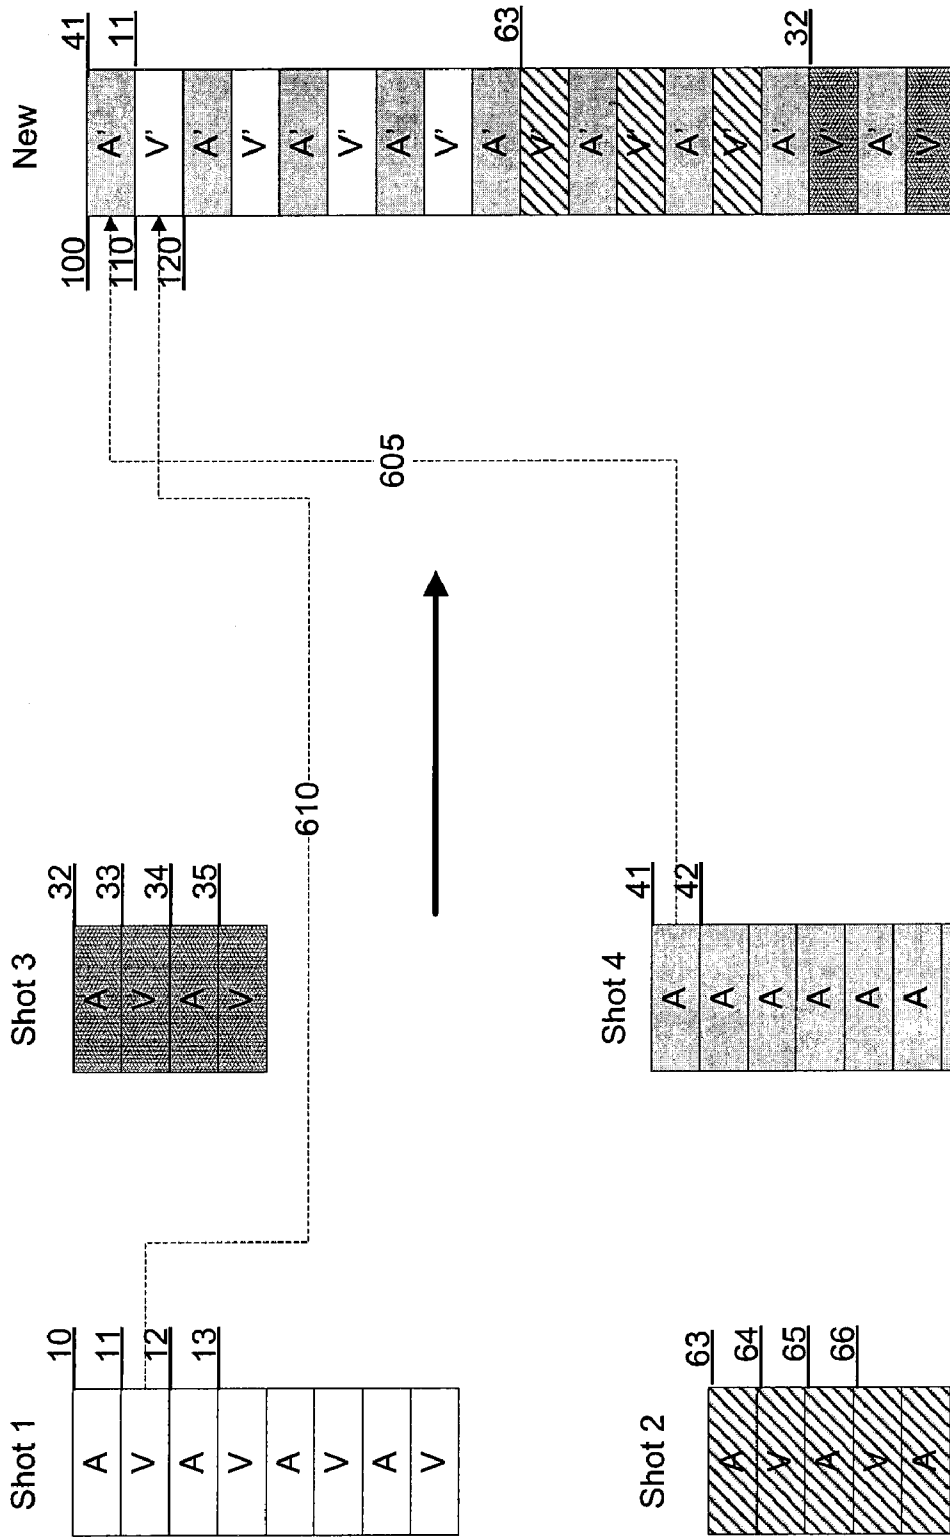


FIG. 6



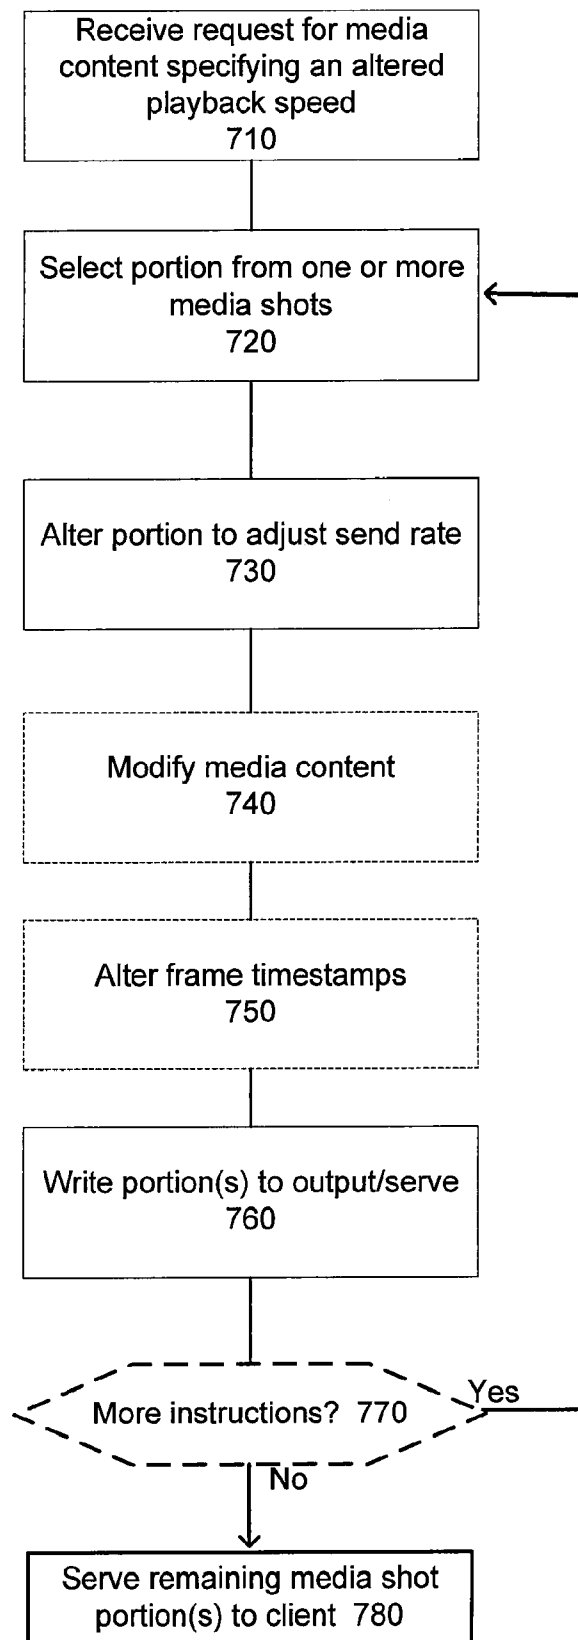


FIG. 7

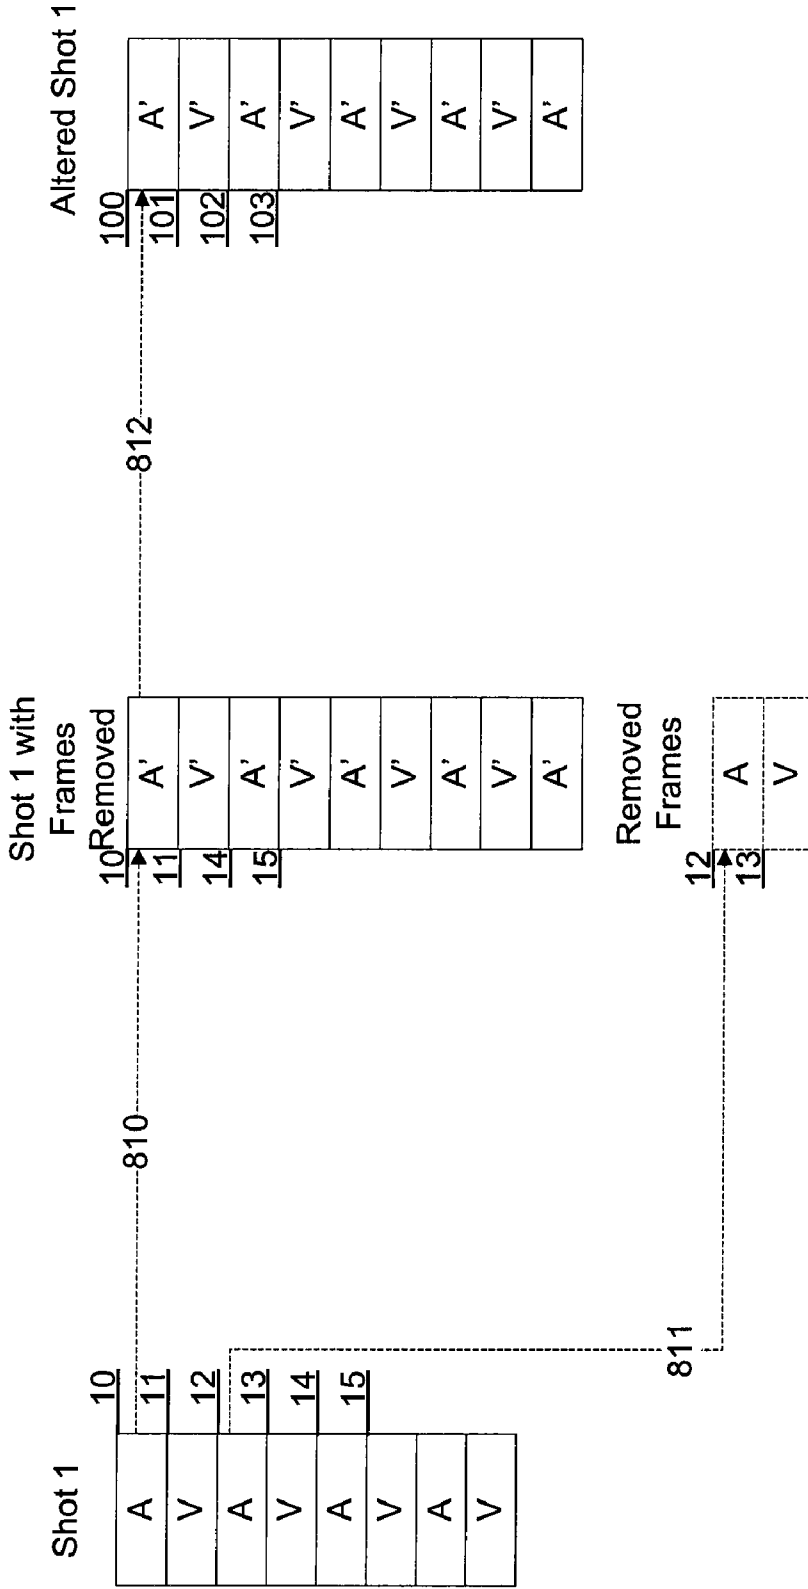


FIG. 8

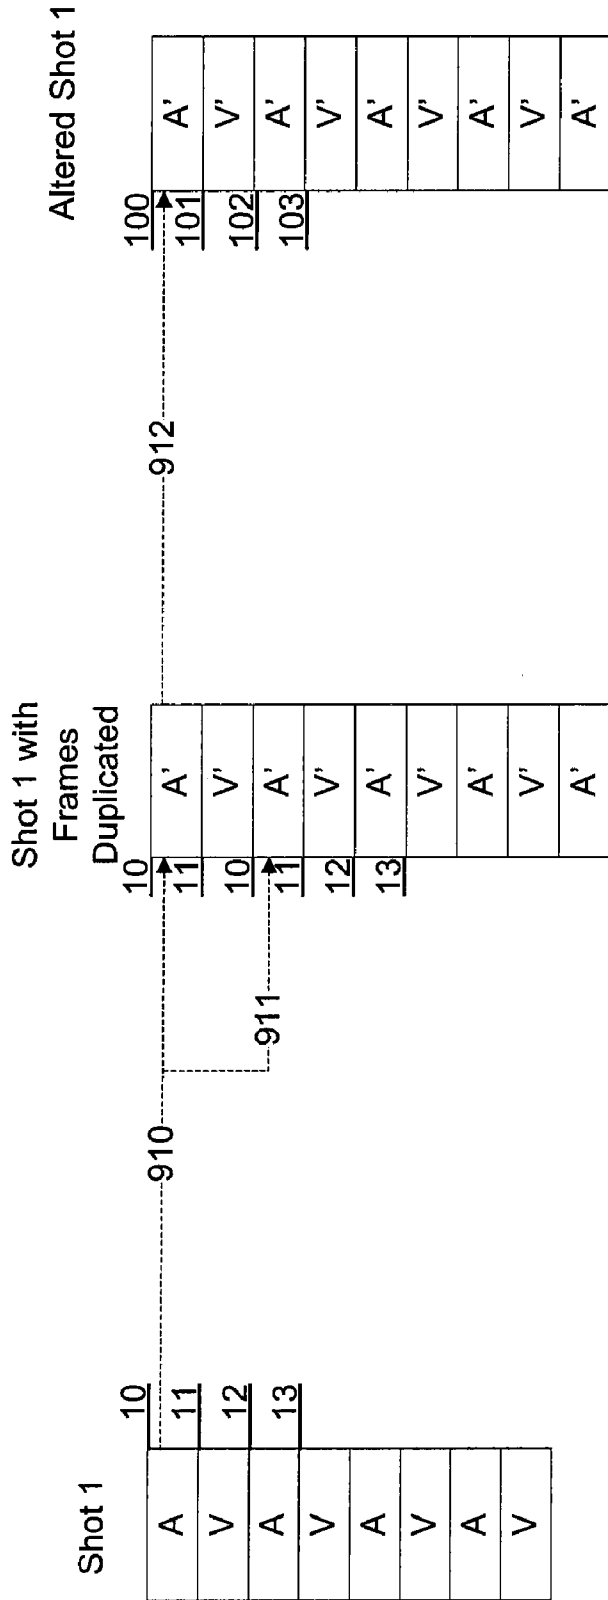


FIG. 9

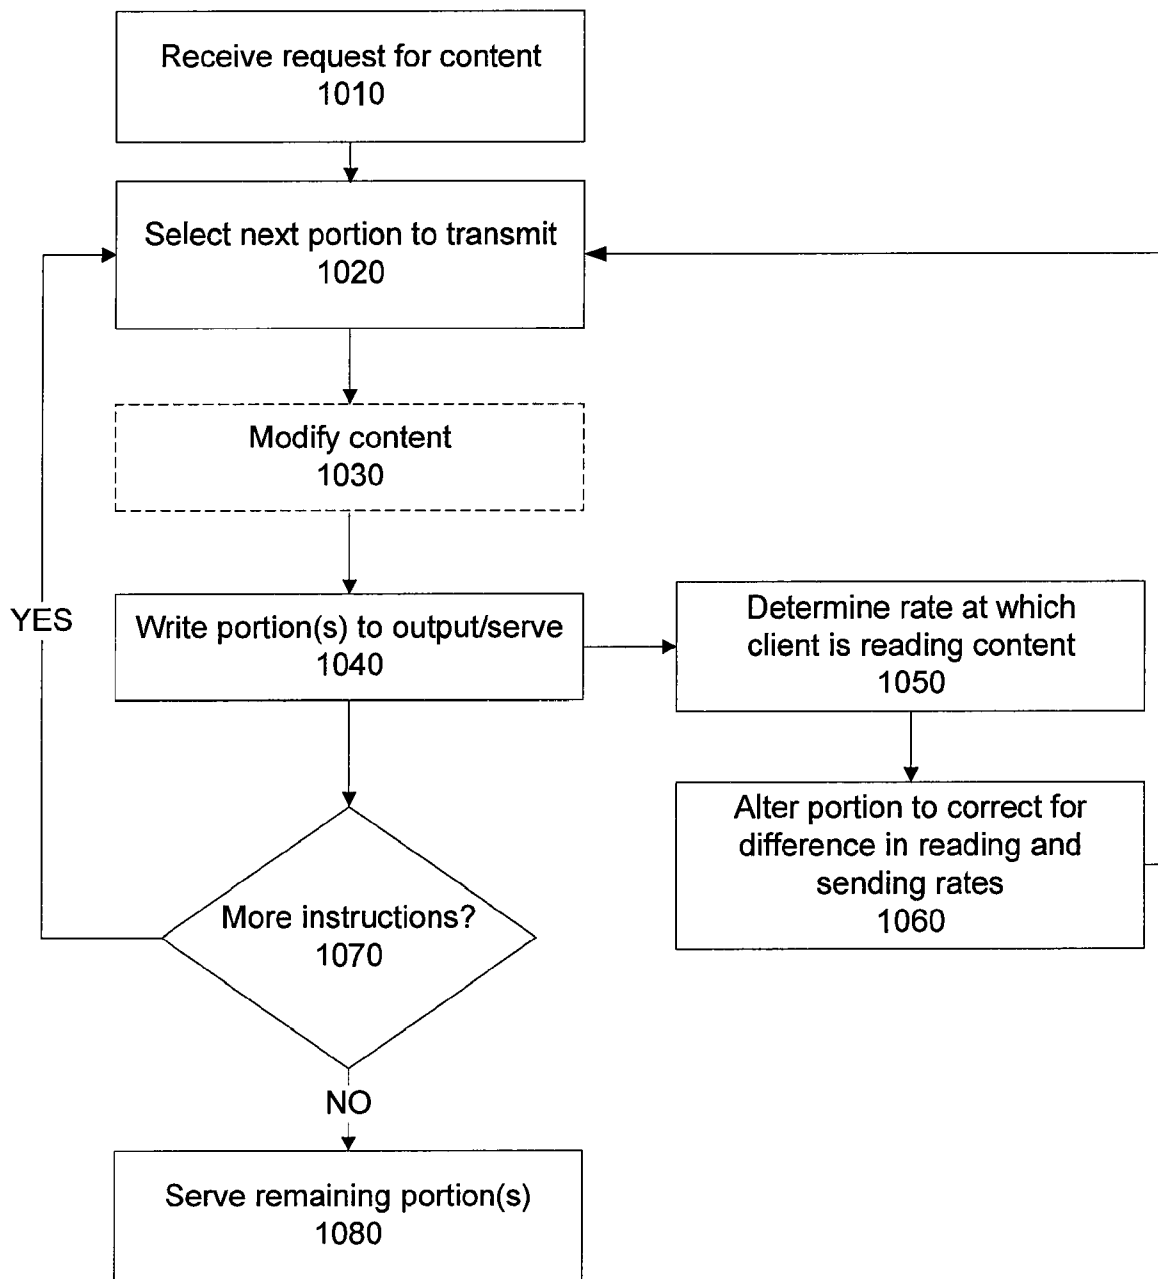


FIG. 10

## SERVING MEDIA ARTICLES WITH ALTERED PLAYBACK SPEED

### RELATED APPLICATIONS

This application claims the benefit under 35 U.S.C. §119 (e) of Provisional Patent Application Ser. No. 60/756,787, entitled "Discontinuous Download of Video Articles," filed Jan. 6, 2006. This application is a continuation-in-part under 35 U.S.C. §120 of U.S. patent application Ser. No. 11/428,319, entitled "Dynamic Media Serving Infrastructure," filed Jun. 30, 2006. This application is related to U.S. patent application Ser. No. 11/620,346, entitled "Combining and Serving Media Content", filed Jan. 5, 2007; U.S. patent application Ser. No. 11/620,386, entitled "Media Article Adaptation to Client Device", filed Jan. 5, 2007; and U.S. patent application Ser. No. 11/620,468, entitled "Discontinuous Download of Media Articles", filed Jan. 5, 2007. All of the foregoing applications are incorporated herein in their entirety by reference for all purposes.

### BACKGROUND

The present invention relates to display and modification of streaming and downloading media content, and more specifically, to serving playback speed-adjusted media content.

Methods are known for streaming and downloading media content, for example, across the Internet from a server to a client device in response to a client request for media content. Existing technologies use a traditional static file serving interface, in which a complete file or client-requested byte ranges of files is served to the client. The file is usually stored or cached on the server for playback. Typically, if a user wants to view media content in this manner, a specialized client player application for playing the media content must be downloaded and installed on a client system. The server and client player application then use specialized protocols, file formats, and video encodings to transmit, decode and playback the media content.

### SUMMARY

The present invention provides various embodiments of methods and systems for serving media content. The system provides for on-the-fly alteration of apparent playback rate for content sent to a client based on a request from the client to alter the playback speed of the content.

A client requests media content from a media server, the request including an list of media source files and an edit list. The edit list comprises a set of instructions for various methods of combining and/or modifying one or more media shots. The server opens one or more source media files, parses these files, and selects frames or other portions to transmit based on the edit list, and sequentially writes those portions to an output file for serving to the client. According to the instructions, the server adjusts the number of frames of the requested content to affect an apparent change in the playback speed of the content. The adjustment to the frames may include inserting frames, removing frames, and/or inserting markers to trigger the client device to duplicate frames.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a client-server architecture according to one embodiment of the present invention.

FIG. 2 is a block diagram illustrating the server in greater detail.

FIG. 3 is a flowchart illustrating a method for serving media content.

FIGS. 4A-4C illustrate possible uniform resource locators.

FIG. 5A is an illustration of a file format.

FIG. 5B is an illustration of a sample AVI file.

FIG. 6 shows an example of a use case for the method described herein.

FIG. 7 is a flowchart illustrating a method for serving media content with an altered playback speed.

FIG. 8 shows an example of a use case for the method described herein wherein apparent playback speed is increased.

FIG. 9 shows an example of a use case for the method described herein wherein apparent playback speed is decreased.

FIG. 10 is a flowchart illustrating a method for serving content adapted to the read rate of the client.

The figures depict various embodiments of the present invention for purposes of illustration only. One skilled in the art will readily recognize from the following description that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles of the invention described herein.

### DETAILED DESCRIPTION OF THE EMBODIMENTS

Described herein is a method and system for serving media. Media content, as used herein, refers to any type of video and audio content formats, including movie files, i.e., files including both video and audio content, as well as audio-only formats, and may include any form of audio, video, metadata, or combination thereof. A common example is a video article including two content streams, one video stream and one audio stream. However, the techniques described herein can be used with any number of file portions or streams, and may include metadata.

The various embodiments of the invention can be implemented in the context of a standard-client server architecture. FIG. 1 is a block diagram illustrating a client server architecture suitable for this purpose. Such a system comprises a client **105** and a server **110**, communicatively coupled, e.g., by a network **115**. The client **105** can be any type of client computing device, for example, a device executing a browser application or other application adapted to communicate over Internet related protocols (e.g., TCP/IP and HTTP) and/or display a user interface through which media content can be output. According to one embodiment, the user interface of the client **105** allows a user to view, manipulate, and select media content portions and sequence them to form the basis for an edit list as described herein. The client **105** includes a processor, an addressable memory, and other features (not illustrated) such as a display adapted to display video content, local memory, input/output ports, and a network interface. The network interface and a network communication protocol provide access to a network **115** and other computers, such as server **110** or third party computers, along with access to the Internet, via a TCP/IP type connection, or to other network embodiments, such as a LAN, a WAN, a MAN, a wired or wireless network, a private network, a virtual private network, via other networks, or other systems allowing for data communication between two or more computing systems, as well as through a combination of networks (e.g., accessing the server via a bridge from a cellular or wireless networking system to a TCP/IP system). In various embodiments the client **105** may be implemented on a computer running a Microsoft Corp. operating system, an Apple Com-

puter Inc., Mac OS, various flavors of Linux, UNIX, Palm OS, and/or other operating systems. While only a single client **105** is shown, the system can support a large number of concurrent sessions with many clients **105**.

In one implementation, the system operates on high performance server class computers **110**. The details of the hardware aspects of servers **110** are well known to those of skill in the art and are not further described herein. The server **110** is depicted in FIG. 1 as a single computer system, however, the system may be implemented as a network of computer processors and associated storage devices. Examples of server devices **110** are servers, mainframe computers, networked computers, a processor-based device, and similar types of systems and devices. The server **110** and its associated storage devices need not be physically co-located, and there need not be any predefined mapping between storage devices and server computers. The server processor can be any of a number of well known computer processors, such as processors from Intel Corporation of Santa Clara, Calif. and Motorola Corporation of Schaumburg, Ill. In one embodiment, the server **110** has sufficient processing power and bandwidth to perform modifications of content prior to serving as described herein. The network **115** may comprise the Internet alone or in combination with other networks, wired and wireless, such as an intranet, a local area network, a wide area network, or a broadcast network according to various embodiments.

FIG. 2 is a block diagram illustrating one embodiment of the server **110** in greater detail. Here, the server **110** includes server software **120**, source file data storage **125**, a content index **130**, and a key data storage **135**. The source file data storage **125**, content index **130**, and key data storage **135** may also be separate from, but accessible by, the server **110**, and/or form a single data store. The server software **120** is comprised of a number of executable code portions and data files. These include code for enabling the methods described herein. The server software **120** is responsible for orchestrating the processes performed according to the methods of the present invention. The server software **120** includes a read module **140**, a write module **145**, and a buffer module **150**. The read module **140** is one means for enabling the system to use one or more read threads to read portions of a stored file. The read module **140** controls the source media file access, processing, header formation, content reading, content modification, footer construction, and writing to the output as described herein. The write module **145** is one means for enabling the server **110** to control the writing of data from the output (network transmit buffer) to the network **115**. The buffer module **150** is one means for enabling the system to prepare data for writing to the network by placing it in a network transfer buffer. The buffer module **150** controls the functioning of the network transfer buffer. The above software portions **140-150** need not be discrete software modules. The software configuration shown is meant only by way of example; other configurations are contemplated by and within the scope of the present invention. These aspects of the present invention allow the server **110** to serve media content directly from stored content, without the need to duplicate content for serving. In addition, the media files served are playable by standard media playback clients.

Stored media content can be previously received from a client **105** as a result of being uploaded by a user. In another embodiment, the media content is from one or more other sources, for example television content. Once the media content is received at the server **110**, the media file is transformed into a serving format. In this example, transforming into the serving file format includes processing the media file, identifying the indexable portions within the media file, enumer-

ating and assigning each frame a timestamp, parsing the media file headers and footers, and computing format-specific "chunk" boundary points based on the indexable file portions, or chunks. The media content is stored in a source file data storage **125**, and data about the media content is stored out-of-band from the media content in a content index **130**, which can be queried using a content identifier. A content identifier is an arbitrary binary identifier that uniquely identifies a single piece of stored media content and data about the media content, e.g., a 64-bit hash value computed from attributes of the stored media content, any unique string of characters, or a unique number assigned to this content at the time it is created. The content identifier is used to anonymously reference media content instead of using file names, e.g., to provide greater privacy and security. The size of a portion of a file, or chunk, varies from the very small to the very large. The size of the portions of a file determines the granularity available for the editing of the file. The number of chunks stored per media stream is limited only by the size and speed of the storage mechanism used, and by the amount of acceptable metadata storage size per media file. Systems with smaller sized storage would choose a longer interval for chunks, thus reducing the storage overhead for each media file. Systems with large amounts of fast storage could choose chunk sizes down to the individual frame granularity. Preferably, the ratio of the size of the metadata to the size of the media file is small, e.g., less than 1%. Alternatively, a static size for the metadata, such as 64 kB, could be used.

The source file data storage **125** may be a relational database or any other type of database that stores the media content data according to one embodiment. Media content on the server may be received from various sources in various (e.g., compressed or uncompressed) forms. A media article can be stored in various ways and formats. For example, a media article can be stored as a file or other data in a database or other storage medium. A media article can comprise, for example, video content stored in any format, such as Moving Picture Experts Group (MPEG-2, MPEG-4), Windows Media 9 (Windows Media High Definition Video), Video for Windows (AVI), QuickTime, Indeo, etc., as well as in any resolution (e.g., HD or SD), or signal format (e.g., National Television Standards Committee (NTSC), Phase Alternating Line (PAL), System Electronique pour Couleur avec Mémoire (SECAM)), or any other suitable video content of any format and standard whatsoever. The media articles may be suitable for broadcast and/or made available on a network (such as the Internet), a personal computer, or other computing or storage means, and may include other data associated with the video content such as program guide information, metadata, and closed captioning text. The embodiments described herein are described generally in relation to video content, and thus are referred to as video articles, but embodiments may operate on any suitable type of content, including audio-only content.

A generic container file format is used to encapsulate the underlying video or audio data. An example container file format according to one embodiment is shown in FIG. 5A. The example generic file format includes an optional file header followed by file contents **502** and an optional file footer. The file contents **502** comprise a sequence of zero or more chunks **504**, and each chunk is a sequence of "frames" **506** (or groups, as described herein). Each frame **506** includes an optional frame header followed by frame contents **508** and an optional frame footer. A frame **506** (or group) can be of any type, either audio, video, or metadata according to one embodiment. In other embodiments, this model may be used for other file types, for example, Adobe Post Script™, PDF,

mp3, or other custom file types. For temporal media, e.g., most audio or video, frames are defined by a specific (e.g., chronological) timestamp. For non-temporal media, frames are defined by any monotonically increasing value, e.g., page numbers of a PDF document. A frame or group **506** includes one or more streams of content. A stream refers to a set of frames identified by the containing file format. For example, a file may contain zero or more video streams, and zero or more audio streams. Each frame **506** has an identifier that marks it as being part of at least one stream. Some formats implicitly define streams based on frame format types.

Each frame **506** may be a keyframe according to one embodiment, as defined by the video encoding format and encoding parameters used. The media file must include one or more keyframes, and chunks may contain zero or more keyframes. Key frames are used, in part, to define segments of media content from which playback can begin. However, many audio formats do not have this feature, and can be segmented on individual audio frame boundaries. For each frame **506**, a timestamp can be computed. A timestamp need not necessarily correspond to a physical time, and should be thought of as an arbitrary monotonically increasing value that is assigned to each frame of each stream in the file. For example, as referenced above, the monotonically increasing value as applied to a PDF file would correspond to a page number. Thus, the use of the term “timestamp” is used herein in this broader context. If a timestamp is not directly available, the timestamp can be synthesized through interpolation according to the parameters of the video file. Each frame **506** is composed of data, typically compressed audio, compressed video, text meta data, binary meta data, or of any other arbitrary type of compressed or uncompressed data. Various container file formats may be used. An AVI file is described below as one possible container file format, however, other container formats, e.g., ASF, etc., may be used.

In one embodiment, an AVI file is used to store the media article. FIG. 5B shows sample AVI file contents according to one embodiment. In this example, the AVI file includes a header **505** including basic information about the segment, e.g., width, height, video format, audio format, number of video frames, number of audio frames, length of video frame, etc. The bulk of the file lists byte content information **510**. In addition, the file ends with a footer **515** including information about byte offsets, etc. In this example every frame is a key frame.

The content index **130** may be a relational database or any other type of database that stores the indexed data corresponding to the source file media according to one embodiment. Here, the content index **130** includes a content identifier, the number of file portions or streams and stream types within the media file, the file content offset in bytes, the file footer offset in bytes, stream-specific information (e.g., frame rate, format, width, height, duration, etc.), file-format specific flags, and a file format identifier. In this example, the indexed data is available as a result of the above-referenced transforming operations. In addition, other information from the received video file (e.g., width, height, audio and video frame rate, audio and video duration, file format, audio and video codec, etc.) may also be stored in the content index **130**. Media content is stored in source file(s) in a plurality of separately indexed portions. “Portions,” as used herein refer to any segments or fragments of a file, and includes frames, chunks, groups, segments, or other portions defined herein.

The content index **130** is a file-format specific data structure, that includes relevant information from the source media file headers and footers, as well as overall file information, such as described above. The content index **130** may be

structured as a set of sorted lists, tables, trees, or other implementations. The content index **130** stores the media segmented into chunks, which are logical units representing indexed portions of files and identifying the location of the frames in the chunk. In other words, each media file is associated with some number of chunks, which can then be arbitrarily and independently accessed and manipulated. Similar to the file structure depicted in FIG. 5A, each chunk has its own file name for internal data access, and includes a set of frames and/or groups. Each chunk is associated with a row or portion thereof in the content index **130** that represents a single media file. Each row has a maximum size of 64 k bytes, and is identified by a row identifier, which is a hash of the content identifier and an identifier for the video encoding format and resolution. In this example, the row identifier includes multiple columns to specify the content index **130** data as discussed herein. Alternatively, all data for all chunks of a media file is stored in a single database row, using one or more columns. For each chunk, a chunk index stores a chunk identifier number, chunk byte offset (from the start of the file associated with the chunk), chunk length in bytes, chunk start timestamp, chunk end timestamp, and for each stream in the chunk, a stream frame number, stream frame count, stream starting timestamp, and stream ending timestamp. This aspect of the present invention allows a given media file to be represented as a number of individual chunks, which can then be stored separately or together, so as to optimize read and write performance. Further separation of chunks into frames, groups, or other file portions allows the various file operations (transcoding, resolution, etc.) to be performed on the frame or group level, allowing the system to make changes at this level of granularity.

The content index **130** maps from content identifier and time to file name, byte offset, and (video, audio) frame number and/or group number. The mapping is created when media content is received by the server; the mapping is stored in the content index **130** in association with the row for a given media file. The mapping maps a specific time index point in both the video and audio feeds to a specific byte in the source file. Byte offsets are relative to the start of an audio or video frame according to one embodiment, and the content index **130** itself is stored in a format that is file-agnostic with the specific file format and codec stored as metadata in the index **130**. Thus, the content index **130** provides an indication of the relative location of a time index point in the original file, and exactly what content (e.g., audio and video) is stored at that byte.

Time to byte mappings exist for every video key frame in one embodiment. In other embodiments, time to byte mappings exist for every audio frame and video frame in the video article, for every N frames in the video article, or for every N video key frames, without limitation as to size (number of items in the mapping) of the information. The mapping from time to offset preferably maps to the nearest group. Most compressed video formats have the notion of a group of pictures (“GOP,” “group”), which can be thought of an atomic set of video frames. This set of video frames cannot be cut at any point without decoding and re-encoding the video. Therefore, the media content may take advantage of this design, and be addressed at the group level rather than the frame level. Typically, audio formats do not have groups, and can be cut on individual audio frame boundaries. Each frame or group is sorted by time, and includes a start time, audio and/or video frame number, and byte offset. In other words, for any given time point (or range of time) in a media file, the content index **130** can be used to determine the specific chunk(s) that are

associated with that time, and the specific frames thereof, and thereby retrieve them as needed.

The key data storage **135** may be a relational database or any other type of database that stores the data corresponding to keys for encrypting data within a URL. The key data storage **135** may be accessible by the server software **120** through a user interface.

The server software **120**, the source file data storage **125**, content index **130**, and key data storage **135** may be stored and operated on a single computer or on separate computer systems communicating with each other through a network.

The server **110** can store the media content either locally, or on remotely accessible systems, or a combination of both.

One skilled in the art will recognize that the system architecture illustrated in FIGS. **1** and **2** is merely exemplary, and that the invention may be practiced and implemented using many other architectures and environments.

As mentioned above, the present invention operates in the context of an information retrieval system for processing queries for video content. FIG. **3** is a flowchart illustrating one method for serving media content. The method begins when a request from a client **105** is received **310** for media content stored on a server **110**. The request is received at a network interface to the server **110**. The request to the server **110** includes the content identifier for the content that the user desires to play back, the resolution at which to play the media content, and a time range, specified as the starting time (in milliseconds, or another time unit, into the content) and length (in milliseconds or other units). The content identifier is an arbitrary binary identifier that uniquely identifies a single piece of stored media content and data about the media content, e.g., a 64-bit hash value computed from attributes of the stored media content, any unique string of characters, or a unique number assigned to this content at the time it is created. The content identifier is used to anonymously reference media content instead of using file names, e.g., to provide greater privacy and security. This aspect of the present invention is advantageous because it allows for time-based requests from clients, rather than traditional byte-based requests.

The request also includes an edit list. An edit list, as used herein, refers to a list of dynamic or static instructions for serving and editing stored media content. Edit list instructions can be added either by the client **105**, the server **110**, or both. For example, given an edit list received from the client **105**, the server **110** can pre-pend, append, or insert additional instructions into the client-requested content, or can modify the length, rate, number of streams, file wrapper format of the content. The edit list instructions can range from the very simple serving of an entire source file, to more complicated modifications, including serving a portion of a source file, serving more than one portion of one or more source files, changing the rate at which a source file or portion is served (speeding up or slowing down), adapting the rate at which a source file or portion is served to the client capabilities, encrypting a file or portion, adding or removing streams from the file or portion, interleaving streams from different files or portions, adding a data stream to a file or portion, transforming a file wrapper format, transforming from a proprietary file format to a standards-compliant format, inserting metadata into the file or portion, and various other modifications according to one embodiment. Other possible modifications include changing stereo audio into monophonic audio, changing color video into black and white video, reducing the resolution of a video stream, reducing the resolution of a still image stream, performing various visual algorithms on video or still images streams (sharpening, blurring, distorting, crossfade, etc.), and performing various audio algorithms

(changing sample rate, pitch shifting, rate adjusting, echo, flange, etc.) according to various embodiments. Examples of ways to accomplish these modifications are included in the following description. The edit list that is provided in the request also may include an identifier which indicates that further instructions should be generated and included. These extra instructions may be generated by the server, or received from any machine connected to the server. Thus, the complete edit list need not be specified in the request, but can be generated internally at the time the request is received. This process allows every user to be provided with a unique media stream in response to the unique request.

The details of the request can be generated automatically or programmatically. The user is not required to specifically input the request details.

A request may be for content from a single media source file, an altered (e.g., truncated) version of the media content, or a combination of pieces from different media source files. For a request that includes more than one portion of one or more files, the request includes more than one time range, and the server **110** queries the content index to find the chunks that lie nearest to the requested timestamps according to one embodiment.

The request is a hyper text transfer protocol (HTTP) request (or similar transport protocol), and takes the form of a uniform resource locator (URL) **400**, as shown in FIG. **4A**. The URL **400** is in the typical format, including a protocol **405**, a host (and port, if applicable) **410**, and a resource path and/or query string **415**. In addition, an edit list **420**, in the form of instructions for performing operations on one or more files, is included. FIG. **4A** shows a simple example, in which a time range is the only parameter. In some embodiments, part of the URL **400** is encrypted as described herein. FIG. **4B** shows an example of an encrypted URL **425**. In embodiments in which the URL is encrypted, the client cannot modify the encrypted portion, but can add modifications on to the URL, for example, to speed up the rate or to request a smaller subset of the media content. FIG. **4C** shows an example of an encrypted URL with additional edit list instructions **430** appended. In this example, an instruction has been added to speed up the rate for the content (to 1.5 times normal speed) and to select the time range 10 to 5000. The modification permitted by the client **105** does not alter the encrypted portion according to one embodiment. For example, if the file size is 1000, the request for 10 to 5000 may automatically be adjusted by the server to the available time range.

Once a request is received **310**, at least one source media file corresponding to the media content requested is opened **320**. To do this, the request is passed to a read thread, which does asynchronous lookups to the content index **130** using the content identifier, computes starting byte offsets and frame lengths from the specified time range, and then opens **320** the source media file. In one embodiment, if more than one portion will be combined or more than one stream interleaved, the server **110** sorts the content items in the order in which they will be transmitted prior to reading the content index for all content items associated with all portions. In addition, the server **110** stores the relevant header and footer information, such as video resolution and frame rate, byte range of the portion of the file that is being composed, and other file format specific information. If a data stream is being added, the server **110** also reads content index information for a data stream, which includes a list of time-stamped data objects to be inserted in the stream. The time-stamped data objects can be arbitrary binary data, text data, or any other format.

In addition, a new file header is created from the information stored in the content index **130** according to one embodi-



ment, which is written to the output. In one embodiment, the server **110** looks at the content index **130**, reads the header information, and constructs a segment media header representing the requested time range. The server **110** computes the number of audio and video frames that will be transmitted during this request. If more than one portion is being used, the server **110** computes the total number as the sum of frames from all content items that will be transmitted. If the content is being sped up or slowed down, the number of frames that will be transmitted will be different from the number of frames in the source media file(s). The server **110** also appends data specifying the original byte range and (video and audio) frame numbers according to one embodiment. If the data is being encrypted, the header may contain encryption keys, user identifying marks, or other information that allows the client to decrypt the encrypted data that follows. If the file wrapper is being transformed, the server **110** composes a valid file header, according to the “new” wrapper format (e.g., the format into which the wrapper is being transformed). If the wrapper transformation is from a proprietary format to a standards-compliant format, e.g., the original wrapper format is a format that is not standards compliant, and the new wrapper format is a standards compliant one.

Once the file is opened **320**, some or all of the file contents are read.

To reconstruct a media file in the same file format as the original source, the byte offsets of the first and last video and audio frames (or groups) of interest, as well as the number of video and audio frames in the segment, as well as any other metadata that may be useful to clients, are located. In addition, a new header (and footer, if applicable for the format) is created and spliced together to construct the output file.

A separately indexed portion of the media content is then selected **330** for transmitting based upon an instruction in the edit list. In one embodiment, the selecting **330** is part of a loop of reading from the source media file(s), and writing that data to the output. Once selected **330**, the server **110** finds the byte range in the original media source file and reads that region. In one embodiment, the output is a network transmit buffer. The network transmit buffer stores data ready for output to the network, for serving to the client according to one embodiment.

Optionally, the server **110** next modifies **340** the portion. Various modifications are possible; each modification is listed as an instruction in the edit list. Each modification step may or may not be length preserving, and thus the frames and other portions may be of any length after each step of the modification process. One possible data modification is data encryption. The contents of each frame or portion are encrypted with a per-request generated key. To accomplish this, the server **110** looks at each frame of this content, applies a standard encryption algorithm to some or all of these frames, and modifies the frames in a way that clearly delineates them as being encrypted before writing the encrypted data to the output. The keys are stored, e.g., in key data storage **135**, such that they can be retrieved later for decryption purposes. The use of encryption allows for insertion of additional data, e.g., ad targeting information. Another possible data modification is a rate change to a rate different from the original encoding rate. For example, a user can request a “fast forward” stream or portion that sends a file that will play back faster than the original content. To accomplish fast forward, in one embodiment frames are selected to be “dropped” from the source file, while the remaining frames are written to the output as normal, or with their timestamps adjusted accordingly. For example, for each video frame, the frame is either discarded or written to the output. The requested speedup

factor and the specifics of the video encoding determine how many frames need to be discarded, and which frames may be discarded. The server **110** also needs to construct a valid audio stream that matches the rate of the video stream. The audio stream is constructed, e.g., by using a precomputed “silence” audio track in place of the original audio, doing more sophisticated audio processing to adjust audio rate to match the new video rate, or discarding a fraction of the original audio frames.

Another example is a “slow motion” stream or portion that will send data to be played back at a lower rate than the original source. To accomplish slow motion, frames are selected for duplication, or “padding” frames are added to automatically duplicate frames. For example, for each video frame, the frame is either transmitted intact, or a format-specific marker is placed in the file to duplicate the previously transmitted frame. The requested slow down factor and the specifics of the video encoding format determine how many frames need to be duplicated, and which frames may be duplicated, as well as the bytes that should be sent to mark duplicate frames. As with speeding up the content, the server **110** also needs to construct a valid audio stream that matches the rate of the video stream, e.g., as described above. Another possible rate change modification is an adaptive bit rate. To accomplish this modification, the server **110** determines whether the client **105** is able to receive the video at a rate sufficient to be decoded in real time. For example, the rate at which the client reads from the output is monitored throughout the request. If the rate at which the client is reading is less than the encoded rate of the video, the server **110** may replace the original encoded video frame with a significantly smaller, dynamically created frame that tells the video decoder in the client to duplicate the previous frame onscreen. In another embodiment, the server **110** applies an algorithm to modify the original encoded video frame to create a new frame of a smaller size. This algorithm may include a full or partial decode of the original video frame, as well as a full or partial re-encode of the video frame. This process reduces the overall transmitted audio and video rate.

Another possible data modification is selection of data streams or other file portions from within a file. For a data file that includes multiple streams of content, the server **110** can dynamically choose not to send some of the streams. For example, a data file may contain one video stream and three audio streams. At the time the file is sent, one of the three audio streams is chosen for sending with the video stream. For example, the server **110** reads correct byte range from the first file, and writes this to the output, and then repeats this process for subsequent files until each content item has been read and written to the output. Yet another possible data modification is stream interleaving or data interleaving. The server **110** interleaves data from one or more files together to create a new, dynamic piece of content. For example, the server could choose to overlay an audio stream from one file with video stream from another file. In this example, the server **110** reads chunks simultaneously from each file, then takes each video and audio stream pair, and applies a file format specific stream marker to each matched pair of audio and video. It then sorts all of the frames by their representative timestamps, and writes this data to the output. If the number of audio and video frames is different between the streams, then the server **110** may pad the video with empty video frames, black video frames, or some other file format specific construct. A similar padding operation may be used for audio streams. This process ensures that the total duration of video and audio content is identical in each stream. Another example of data interleaving is mixing a text or metadata

stream, e.g., closed captioning or promotional data, with an existing audio and/or video stream. To add a data stream, the index information for the data stream, e.g., the byte range, is gathered and transformed into a format suitable for inclusion in the resultant media file. The video frames and data stream frames are then sorted by their internal timestamps, and each frame is subsequently written to the output. This process results in an effective interleaving of the data and video streams, and the data stream will be synchronized correctly with the video and audio streams. In one embodiment, the added data is metadata dynamically computed by the server. The metadata may be data that is not useful to the user, and is consumed by the playback client.

Another possible data modification is data wrapper transformation. Different computers and operating systems support different file wrapper and frame wrapper formats. For each source file, the frame header and footer specific to the original format are discarded, and generate new headers and footers that are specific to a desired file format. For each frame contained in this data, the server **110** parses the data according to the original wrapper format, extracts the compressed video or audio data, and then encodes this data using the “new” wrapper format, and writes this byte range to the output. Yet another possible data modification is calculated data insertion, wherein data is generated on the fly and inserted into the served file according to one embodiment. For example, a new “stream” is created and specified in the file header, and this stream is transmitted by the server. In this example, the new stream includes data transmission statistics, popularity statistics, or other computed metadata according to various embodiments.

In addition, for each frame, various types of audio and/or video processing may be applied, e.g., using one or more processing algorithms. The algorithm(s) may operate directly on the encoded video to create the desired video effect. Or the server **110** may do a partial or full decode of each video frame, apply an arbitrary video transformation on the partially or fully decoded video and then do a partial or full encode to produce the resultant frame. Each audio frame may be processed in a similar manner. Video processing and enhancement algorithms are well known, and include such things as color to black and white transformation, sharpening, blurring, arbitrary distortions, resolution reduction, and blending effects between one or more video streams. Audio processing steps include such things as sample rate changes, pitch shifting, bitrate adjustments, echo, flange, crossfade, converting to monaural, and dynamic range compression.

In some embodiments, a per-frame transformation is applied to make the file appear consistent. The timestamp of the output video article is updated so that the first frame is timestamp **0** and increases over time according to some embodiments. Although the time durations of video and audio frames are often unequal, because separate video and audio time offsets are stored in the content index **130**, there is as little offset as possible between video and audio frames. Thus, no re-mixing of video and audio frames is required (but is possible, if desired). According to one embodiment, the server **110** also has the ability to dynamically modify data by chaining together multiple modifications for the same piece of media content. These aspects allow insertions, deletions, and other modifications to be made at the time of serving the media file portion.

Once any modifications **340** are made to the portion, the portion is written **350** to the output. Then, optionally it is determined **360** whether more instructions apply for the

requested media. If the edit list includes additional instructions, steps **320-350** repeat according to the next instruction in the edit list.

In one embodiment, the server **110** next constructs a footer that includes an index of relevant information from the original media source file and writes the footer to the output. In one embodiment, a region of the source file footer is read, modified, and placed in the network transmit buffer. If greater than one file portion is to be included in the served content, a region of the first source media file footer is read first, then subsequent files in series. If the data wrapper is being transformed, the server **110** parses the footer according to the original wrapper format, and modifies the footer to be consistent with the new wrapper format. As with the header, if the transformation is from a proprietary format to a standards-compliant format, the original wrapper format is a format that is not standards compliant, and the new wrapper format is a standards-compliant one.

Up to this point, the steps at the server **110** have been performed by a single read thread according to one embodiment. A separate thread, a write thread, controls the writing of data from the output (network transmit buffer) to the network **115** for serving to the client. The write thread determines whether there is data in the network transfer buffer. If there is, the data is written to the network. If there is no data in the network transfer buffer, or the buffer is nearly empty, the write thread asks for more data, i.e., from the above process by the read thread.

FIG. **6** shows an example of a use case for the method described herein according to one embodiment. In this example, streams from different file portions are interleaved. As a preliminary step, a user uploads one or more pieces of media content (“shots”) to a server, e.g., via a user interface. In this example, the user uploads media content including shots **1-4**. The media content is transcoded on the server into the serving format. Then, using a user interface, the user selects, edits, and arranges the content into the format desired for viewing. In this example, the user wants to create a new piece of media content that includes a portion of the video from the content associated with shots **1-3**, i.e., the portion shown as shots **1-3**, with a portion of the audio associated with shot **4**, i.e., the portion shown as shot **4**. The user interface software converts the instructions into an edit list as part of a request from the client for the edited content, which is sent to the server. After the steps detailed below, the client will receive from the server a single piece of content edited per the request.

The server receives the request from the client, in the form of a URL as described in conjunction with FIGS. **4A-C**, which includes the content identifier for each piece of content that the user desires to play back, in this example shots **1-4**, and the resolution at which to play the content. The request also includes a time range for each shot, specified as the starting time (e.g., in seconds into the content) and length (e.g., in seconds). For example, for shot **1**, the starting time is 10 seconds (begin=10) and the length is 8 seconds (len=8). The request also includes an edit list that refers to a list of instructions for interleaving the four streams from shots **1-4**, such that the new piece of content includes the video of each of shots **1-3** and the audio of shot **4**.

As discussed in conjunction with FIG. **3**, the source media file corresponding to the media content requested, in this example, corresponding to each of shots **1-4**, is opened. To accomplish this step, the request is passed to a read thread, which does asynchronous lookups to the content index using the content identifiers, computes starting byte offsets and

frame lengths from the specified time ranges, and opening each file. The read thread also creates a new header for the new content.

Following the instructions in the edit list sequentially, a next frame of the media content is selected for transmitting. In this example, the server interleaves data from four files together to create a new, dynamic piece of content. Specifically, the server selects an audio stream from one file (shot 4) to interleave with the video streams from three other files (shots 1-3).

In this example, the audio frame of shot 4 with timestamp "41" is selected 605 for transmitting and is written to the network transfer buffer, as illustrated in FIG. 6. Next, the video frame of shot 1 with timestamp "11" is selected 610 and written to the output buffer. The read thread continues to select the next frame in sequence until completing the last instruction in the edit list. The frames are renumbered such that the timestamps associated with the newly created piece of content increase monotonically from the new starting time. The frames of the new shot may include any of the various modifications described herein as compared with the frames from the original shots, as indicated by the fact that the audio and video frames respectively are designated A' and V' rather than A and V. The read buffer then constructs the new footer. The new header, content, and footer are spliced together to construct the new output file for serving to the client. When the write thread detects that there is data in the buffer, the data is transmitted to the network for serving to the client.

FIG. 7 is a flowchart illustrating a method for serving media content with an altered playback speed. The method begins when a request for a playback speed-altered media shot is received 710 from a client at a server, for example, client 105 and server 110. The request includes a content identifier for the content that the user desires to play back and an altered playback speed at which the media content should be played. For example, the request may indicate that the playback speed be twice the standard speed (e.g., "speed=2" in the request parameters) Similarly, for a request indicating "speed=0.5" the content will appear to be played at half the standard speed. The altered playback speed may specify that the media content should be played either faster or slower, and can specify a change of any magnitude. The request may contain additional information that identifies other transformations that the server 110 will make to the media content. Examples of these transformations are as described in conjunction with FIG. 3.

Once a request is received 710, at least one source file corresponding to the content requested is opened, and a portion of the content corresponding to the request is selected 720 and read. The selection is made based on a content identifier, a time range, and other information included in the request from the client 105.

Next, the server 110 alters 730 the portion to adjust the send rate of the selected content. This alters the apparent playback speed of the content, by altering the number of frames in the media content, e.g., by adding or removing frames. Thus when the client 105 plays the content, at the same number of frames per second as if the content had not been altered, the fact that frames are removed or added will increase or decrease, respectively, the playback speed.

For increasing the playback speed, the server 110 removes frames from the media content. The number of frames removed depends on the playback speed specified. For instance, if the client requests 710 that the media be played at 1.5 times the normal speed, then the server 110 would remove every third frame, so that the file size is  $\frac{2}{3}$  of the original size. When played back, the media will appear to be playing at

1.5x speed. The file format may dictate that certain frames cannot be removed. In this case, the server 110 will adapt the removal process accordingly by removing one of the frames adjacent to the frame that cannot be removed. As a result of the removal of  $\frac{1}{3}$  of the frames, the total time to play the altered media content will be  $\frac{2}{3}$  of the total time to play the original media content, appearing to the user as if the content has been sped up.

For decreasing the playback speed, the server 110 inserts duplicate frames into the media content. Alternatively, if the file format permits, the server 110 adds a format-specific marker to frames that directs the client 105 to play the previous frame a second time rather than inserting a duplicate frame. The number of frames inserted depends on the playback speed specified. For instance, for a request 710 that the media be played at 0.5 times the normal speed, the server 110 duplicates every frame. As a result of the duplicated frames, the altered media content will take twice the amount of time to play as it would to play the original media content, appearing to the user that the content has been slowed down.

If the media content contains audio, then the audio content's playback speed can also be altered by removing or inserting frames as described above. Alternatively, the audio content is altered by processing it to maintain the pitch of the original audio content while altering the number of frames. Or, the audio content can be replaced with a silence track that is either pre-computed or computed during the processing.

If the request also includes instructions for modifying the media content, the selected content may be modified 740 according to the instructions in the request. The instructions can be in the form of an edit list as described in conjunction with FIG. 3. Possible modifications may include data encryption, length modification, encoding rate modification, stream quantity modification, file wrapper format modification, video transformation processing, audio transformation processing, and a text or metadata insertion. Examples of these modifications are discussed in greater detail in conjunction with FIG. 3. The modifications can be applied to a portion or all of a media shot.

The server 110 next alters 740 the timestamps of the frames in the media content. Because frames have been removed or inserted, the timestamps will no longer be consecutive and possibly will no longer be monotonically increasing. This process will change the timestamps so that they are monotonically increasing and/or numbered consecutively. Some video formats may not have timestamps or do not require consecutively numbered timestamps, and the server 110 instead may not alter the timestamps for those video formats.

The altered media content, or a portion thereof, then is written 760 to an output medium and is served to the client 105. Portions of the media content are served to the client 105 when each portion is ready for serving. There is no need to wait for the entire altered media shot to be formed to serve it to the client 105. Then, it is determined 770 whether more instructions apply to the requested media content. If the request includes more instructions, steps 720-760 repeat according to the next instruction. Once all instructions are processed and the altered media shot is formed from the selected frames, the server 110 serves 780 any remaining portion of the combined media shot to the client 105.

FIG. 8 shows an example of a use case for the method described herein wherein apparent playback speed is increased. In this example, the apparent playback speed of media content is increased by removing frames from the original content on the server 110. As a preliminary step, a user selects a piece of media content to view, e.g., via a user interface. The user also selects an increased playback speed

15

for the media content. In this example, the user wants to view shot **1** at an increased playback speed. The instructions are converted into a request from the client for the edited content, for example in the form of a URL as described in conjunction with FIGS. 4A-C, which is sent to the server **110**. After the steps detailed below, the client will receive from the server a single piece of content with an apparent increased playback speed as per the request.

In this example, the user requests shot **1** to be played at an increased speed. The server accomplishes this alteration by removing frames from shot **1**. First, the first audio frame of shot **1** (timestamp “**10**”) and the first video frame of shot **1** (timestamp “**11**”) are selected **810** to be included in the altered shot **1**. The server then removes **811** the frames with timestamps “**12**” and “**13**.” These frames will not be included in the altered shot **1**. The server then selects **810** the frames from shot **1** with timestamps “**14**” and “**15**” to be included in the altered shot **1**.

The frames of the new shot may include any of the various modifications described herein as compared with the frames of the original shot, as indicated by the fact that the audio and video frames respectively are designated A' and V' rather than the original A and V. For instance, the audio frames may be processed to preserve the pitch of shot **1**.

The server may alter the timestamps of the altered shot **1**. In this example, the server modifies **812** the timestamps of the altered shot **1** so that the frames are numbered consecutively. The audio frame with timestamp “**10**” becomes timestamp “**100**” in the altered shot **1**, the frame with timestamp “**11**” becomes timestamp “**101**” in the altered shot **1**, the frame with timestamp “**14**” is changed to timestamp “**102**”, and the frame with timestamp “**15**” is changed to timestamp “**103**.” The server **110** piecemeal serves the altered shot **1** to the client **105** on-the-fly as described herein, maintaining the order designated by the timestamps.

FIG. 9 shows an example of a use case for the method described herein wherein apparent playback speed is decreased. In this example, the playback speed of media content is decreased by inserting duplicate frames into the content on the server **110**. As described in conjunction with FIG. 8, a user selects a piece of media content, selects in this case a decreased playback speed, and the instructions are converted into a request from the client **105** for the edited content, which is sent to the server. After the steps detailed below, the client will receive from the server a single piece of content with a decreased playback speed as per the request.

In this example, the user requests shot **1** to be played at a decreased speed. The server then processes shot **1** to decrease the playback speed as follows. In this example, the first audio frame of shot **1** (timestamp “**10**”) and the first video frame of shot **1** (timestamp “**11**”) are selected **910** for to be included in the altered shot **1**. The server then copies **911** the frames with timestamps “**10**” and “**11**” again, for duplicate inclusion in altered shot **1**. The output thus includes two copies of the frames with original timestamps “**10**” and “**11**.” The server selects **912** the frames with timestamps “**12**” and “**13**” to be included in the altered shot **1**. Because the client will display the media at a constant number of frames per second, the insertion of duplicate frames will decrease the apparent playback speed. As described in FIG. 8, the frames of the altered shot **1** may include any of the various modifications described herein.

The server **110** then alters the timestamps of the altered shot **1**. In this example, the server modifies the timestamps **912** of the altered shot so that the frames are numbered consecutively and are monotonically increasing. The first audio frame with timestamp “**10**” becomes timestamp “**100**,” the

16

first frame with timestamp “**11**” becomes timestamp “**101**,” the second frame with timestamp “**10**” is changed to timestamp “**102**,” and the second frame with timestamp “**11**” is changed to timestamp “**103**.” The server **110** piecemeal serves the altered shot **1** to the client **105** on-the-fly as described herein, maintaining the order designated by the timestamps.

FIG. 10 is a flowchart illustrating a method for serving content adapted to the read rate of the client. The method begins when a request for content is received **1010** from a client at a server, for example, client **105** and server **110**. The request includes similar characteristics to the request described in conjunction with FIG. 3. The instructions also can contain ways to modify the media content.

Once a request is received **1010**, at least one source file corresponding to the content requested is opened, and a portion of the content corresponding to the request is selected **1020** and read. The selection is made based on a content identifier, a time range, and other information included in the request from the client **105**.

If the request also includes instructions for modifying the media content, the selected content may be modified **1030** according to the instructions in the request. The instructions can be in the form of an edit list as described in conjunction with FIG. 3. Possible modifications may include data encryption, length modification, encoding rate modification, stream quantity modification, file wrapper format modification, video transformation processing, audio transformation processing, and a text or metadata insertion. Examples of these modifications are discussed in greater detail in conjunction with FIG. 3. The modifications can be applied to a portion or all of a media shot.

The requested content, or a portion thereof, then is written **1040** to an output medium and is served to the client **105**. Portions of the media content are served to the client **105** when each portion is ready for serving. There is no need to wait for the entire altered media shot to be formed to serve it to the client **105**. The content may be sent to the client **105** in small pieces, e.g., 64 kB each.

Once at least a portion of the content is served **1040**, the server **110** determines **1050** the rate at which the client **105** is reading content. To make this determination, the server **110** compares the rate at which the client **105** receives, and is therefore able to decode, the content (“the reading rate”) to the rate, in frames per second, at which the content is being served from the server **110** (“the sending rate”).

More specifically, the time between when the data transmission begins for a portion of content, and the time at which the data transmission ends for the portion of content can be measured, and the average rate at which the client **105** is reading the content can be determined. For example, the server **110** can measure the time between client reads of successive 64 kB chunks of content. This is a reasonable measurement of client **105** throughput because the standard internet networking architecture (routers, switches, etc.) have very small buffers, and thus, the rate at which the client is reading can be estimated in this way. The amount of network buffering may also be able to be determined by keeping statistics and probing each client’s destination network or sub-net. Other factors that affect client networking speed (such as the required retransmission of TCP/IP network packets) can be known, measured, and accounted for as well, to give a more accurate estimation of client read speed. This technique is useful for a gross estimation of client read speed. For example, from this process it easily could be determined whether a client is using a dial-up connection (<56 kbps), low-speed DSL (~768 kbps) or high-speed DSL (>3 Mbps).

Alternatively, one or more test packets of non-requested content could be sent to make the above determination.

Alternatively or in conjunction with the above, aggregate statistics also could be measured for the client's destination host domain and kept in a persistent storage database on the server **110** (not shown), such that a reasonable estimate could be made for unknown clients, even if the server **110** has not sent any data to that client yet, or even seen the client's specific destination IP address.

After making this comparison, the server **110** alters **1060** the content to correct for any difference in sending rate and reading rate. In one embodiment, this correction **1060** is accomplished by adding or removing frames from the content, in a manner similar to that described in conjunction with FIG. 7. In order to determine whether to add or remove and how many frames to add or remove, the server **110** does the following calculation to determine the altered frame count (altered number of total frames):

$$\left[ \frac{\text{(reading rate)}}{\text{(sending rate)}} \right] * [\text{original number of frames in content}]$$

wherein the reading rate and the sending rate are in frames per second.

Once the sending rate is adjusted **1060**, the server **110** returns to selecting **1020** the next portion of content, modifying **1030** the content if necessary, and writing **1040** the portion to output. Following the writing **1040** step, the server **110** next determines **1070** whether more instructions apply to the requested media content. If the request includes more instructions, steps **1020-1040** repeat according to the next instruction. Once all instructions are processed and the altered media shot is formed from the selected frames, the server **110** serves **1080** any remaining portion of the combined media shot to the client **105**.

Thus, the present invention provides for the dynamic increase in playback speed of media content by a server, on-the-fly from a corpus of indexed data source files. The media content served is playable by standard media playback clients, and uses simple, existing, open protocols such as HTTP, and standards compliant video encoding such as MPEG-4 and h.263. The method uses an increased playback speed specification received as part of a request for media content from a client, or via a determination of client read speed, and the server dynamically modifies the content to increase the playback speed before serving it.

The present invention has been described in particular detail with respect to one possible embodiment. Those of skill in the art will appreciate that the invention may be practiced in other embodiments. First, the particular naming of the components, capitalization of terms, the attributes, data structures, or any other programming or structural aspect is not mandatory or significant, and the mechanisms that implement the invention or its features may have different names, formats, or protocols. Further, the system may be implemented via a combination of hardware and software, as described, or entirely in hardware elements. Also, the particular division of functionality between the various system components described herein is merely exemplary, and not mandatory; functions performed by a single system component may instead be performed by multiple components, and functions performed by multiple components may instead performed by a single component.

Some portions of above description present the features of the present invention in terms of algorithms and symbolic representations of operations on information. These algorithmic descriptions and representations are the means used by

those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. These operations, while described functionally or logically, are understood to be implemented by computer programs. Furthermore, it has also proven convenient at times, to refer to these arrangements of operations as modules or by functional names, without loss of generality.

Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system memories or registers or other such information storage, transmission or display devices.

Certain aspects of the present invention include process steps and instructions described herein in the form of an algorithm. It should be noted that the process steps and instructions of the present invention could be embodied in software, firmware or hardware, and when embodied in software, could be downloaded to reside on and be operated from different platforms used by real time network operating systems.

The present invention also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored on a computer readable medium that can be accessed by the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, application specific integrated circuits (ASICs), or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus. Furthermore, the computers referred to in the specification may include a single processor or may be architectures employing multiple processor designs for increased computing capability.

The algorithms and operations presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may also be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will be apparent to those of skill in the art, along with equivalent variations. In addition, the present invention is not described with reference to any particular programming language. It is appreciated that a variety of programming languages may be used to implement the teachings of the present invention as described herein, and any references to specific languages are provided for invention of enablement and best mode of the present invention.

The present invention is well suited to a wide variety of computer network systems over numerous topologies. Within this field, the configuration and management of large networks comprise storage devices and computers that are communicatively coupled to dissimilar computers and storage devices over a network, such as the Internet.

Finally, it should be noted that the language used in the specification has been principally selected for readability and instructional purposes, and may not have been selected to delineate or circumscribe the inventive subject matter. Accordingly, the disclosure of the present invention is

intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the following claims.

What is claimed is:

1. A method of altering apparent media playback speed on a server, comprising:

receiving at a server from a client a request for media content specifying an altered playback speed, wherein the media content comprises video content;  
 altering the number of frames of the media content according to the altered playback speed by inserting duplicate frames into the media content without overlapping the duplicate frames;  
 serving at least a portion of the altered media content from the server to the client; and  
 storing a reading rate for a client host domain.

2. The method of claim 1, wherein altering the number of frames of the media content further comprises calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content.

3. The method of claim 1, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving.

4. A method of altering apparent media playback speed on a server, comprising:

receiving at a server from a client a request for media content specifying an altered playback speed, wherein the media content comprises video content;  
 altering the number of frames of the media content according to the altered playback speed by:  
 inserting duplicate frames into the media content without overlapping the duplicate frames;  
 inserting format-specific markers into the media content that direct the client to duplicate the previous frame; and  
 serving at least a portion of the altered media content from the server to the client.

5. The method of claim 1, wherein the media content comprises at least audio content, the method further comprising replacing the audio content with silence.

6. The method of claim 1, wherein the media content comprises at least audio content, the method further comprising processing the audio content to alter the playback speed while preserving pitch characteristics of the audio content.

7. The method of claim 1, wherein the media content comprises at least audio content, the method further comprising inserting duplicate audio frames.

8. The method of claim 1, further comprising storing a reading rate for a client host domain.

9. The method of claim 1, wherein the request comprises at least one content identifier and instructions corresponding to the at least one content identifier.

10. The method of claim 9, wherein the instructions further comprise a content transformation instruction.

11. A method of altering apparent media playback speed on a server, comprising:

receiving at a server from a client a request for media content specifying an altered playback speed, wherein the content comprises a frame selected from a group consisting of a video stream, an audio stream, a text stream, and a metadata stream and the request comprises at least one content identifier and instructions corresponding to the at least one content identifier, the instructions further comprising a content transformation instruction;  
 altering the number of frames of the media content according to the altered playback speed by inserting duplicate

frames into the media content without overlapping the duplicate frames, wherein altering the number of frames of the media content further comprises calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content;

serving at least a portion of the altered media content from the server to the client, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving; and

storing a reading rate for a client host domain.

12. A computer program product for altering apparent media playback speed on a server, the computer program product comprising:

a computer-readable medium; and

computer program code, coded on the medium, for:

receiving from a client at a server a request for media content specifying an altered playback speed, wherein the media content comprises video content;  
 altering the number of frames of the media content according to the altered playback speed by inserting duplicate frames into the media content without overlapping the duplicate frames;

serving at least a portion of the altered media content from the server to the client; and

storing a reading rate for a client host domain.

13. The computer program product of claim 12, wherein altering the number of frames of the media content further comprises calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content.

14. The computer program product of claim 12, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving.

15. A computer program product for altering apparent media playback speed on a server, the computer program product comprising:

a computer-readable medium; and

computer program code, coded on the medium, for:

receiving from a client at a server a request for media content specifying an altered playback speed, wherein the media content comprises video content;  
 altering the number of frames of the media content according to the altered playback speed by:

inserting duplicate frames into the media content without overlapping the duplicate frames;

inserting format-specific markers into the media content that direct the client to duplicate the previous frame; and

serving at least a portion of the altered media content from the server to the client.

16. The computer program product of claim 12, wherein the media content comprises at least audio content, the method further comprising processing the audio content to alter the playback speed while preserving pitch characteristics of the audio content.

17. The computer program product of claim 12, wherein the request comprises at least one content identifier and instructions corresponding to the at least one content identifier.

18. The computer program product of claim 17, wherein the instructions further comprise a content transformation instruction.

## 21

19. The computer program product of claim 12, wherein the content comprises a frame selected from a group consisting of a video stream, an audio stream, a text stream, and a metadata stream.

20. A system for altering apparent media playback speed on a server, the system comprising:

means for receiving from a client at a server a request for media content specifying an altered playback speed, wherein the media content comprises video content;

means for altering the number of frames of the media content according to the altered playback speed by inserting duplicate frames into the media content without overlapping the duplicate frames;

means for serving at least a portion of the altered media content from the server to the client;

means for storing a reading rate for a client host domain.

21. The system of claim 20, wherein the means for altering the number of frames of the media content further comprises means for calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content.

22. The system of claim 20, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving.

23. The system of claim 20, further comprising means for storing a reading rate for a client host domain.

24. The system of claim 20, wherein the content comprises a frame selected from a group consisting of a video stream, an audio stream, a text stream, and a metadata stream.

25. A computer program product for altering apparent media playback speed on a server, the computer program product comprising:

a computer-readable medium; and

computer program code, coded on the medium, for:

receiving at a server from a client a request for media content specifying an altered playback speed, wherein the content comprises a frame selected from a group consisting of a video stream, an audio stream, a text stream, and a metadata stream and the request comprises at least one content identifier and instructions corresponding to the at least one content identifier, the instructions further comprising a content transformation instruction;

altering the number of frames of the media content according to the altered playback speed by inserting duplicate frames into the media content without overlapping the duplicate frames, wherein altering the number of frames

## 22

of the media content further comprises calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content;

serving at least a portion of the altered media content from the server to the client, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving; and

storing a reading rate for a client host domain.

26. The method of claim 4, wherein the request comprises at least one content identifier and instructions corresponding to the at least one content identifier.

27. The method of claim 26, wherein the instructions further comprise a content transformation instruction.

28. The computer program product of claim 15, wherein the request comprises at least one content identifier and instructions corresponding to the at least one content identifier.

29. The computer program product of claim 28, wherein the instructions further comprise a content transformation instruction.

30. A system for altering apparent media playback speed on a server, the system comprising:

means for receiving at a server from a client a request for media content specifying an altered playback speed, wherein the content comprises a frame selected from a group consisting of a video stream, an audio stream, a text stream, and a metadata stream and the request comprises at least one content identifier and instructions corresponding to the at least one content identifier, the instructions further comprising a content transformation instruction;

means for altering the number of frames of the media content according to the altered playback speed by inserting duplicate frames into the media content without overlapping the duplicate frames, wherein altering the number of frames of the media content further comprises calculating alterations to the number of frames of the requested media content to conform to the altered playback speed to produce the altered media content;

means for serving at least a portion of the altered media content from the server to the client, wherein the portion of the altered media content is capable of being served prior to the remainder of the altered media content being available for serving; and

means for storing a reading rate for a client host domain.

\* \* \* \* \*