



US007818444B2

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

(10) **Patent No.:** **US 7,818,444 B2**
(45) **Date of Patent:** **Oct. 19, 2010**

(54) **APPARATUS, SYSTEM, AND METHOD FOR
MULTI-BITRATE CONTENT STREAMING**

6,490,627 B1 12/2002 Kalra et al.

(75) Inventors: **David F. Brueck**, Saratoga Springs, UT
(US); **Mark B. Hurst**, Cedar Hills, UT
(US)

(Continued)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **Move Networks, Inc.**, American Fork,
UT (US)

CA 2466482 5/2003

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 319 days.

(Continued)

OTHER PUBLICATIONS

(21) Appl. No.: **11/673,483**

Pathan, Al-Mukaddim, et al., "A Taxonomy and Survey of Content
Delivery Networks", Australia, Feb. 2007. Available at <http://www.gridbus.org/reports/CDN-Taxonomy.pdf>.

(22) Filed: **Feb. 9, 2007**

(65) **Prior Publication Data**

(Continued)

US 2008/0195743 A1 Aug. 14, 2008

Related U.S. Application Data

(63) Continuation-in-part of application No. 11/116,783,
filed on Apr. 28, 2005.

(60) Provisional application No. 60/566,831, filed on Apr.
30, 2004.

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/231**

(58) **Field of Classification Search** 709/231,
709/232; 375/240.02

See application file for complete search history.

(56) **References Cited**


U.S. PATENT DOCUMENTS

4,535,355 A 8/1985 Arn et al.
5,953,506 A 9/1999 Kalra et al.
6,091,775 A 7/2000 Hibi et al.
6,091,777 A 7/2000 Guetz et al.
6,195,680 B1 2/2001 Goldszmidt et al. 709/203
6,366,614 B1 4/2002 Pian et al. 375/240.02
6,486,803 B1 * 11/2002 Luby et al. 341/50

(57) **ABSTRACT**

An apparatus for multi-bitrate content streaming includes a receiving module configured to capture media content, a streamlet module configured to segment the media content and generate a plurality of streamlets, and an encoding module configured to generate a set of streamlets. The system includes the apparatus, wherein the set of streamlets comprises a plurality of streamlets having identical time indices and durations, and each streamlet of the set of streamlets having a unique bitrate, and wherein the encoding module comprises a master module configured to assign an encoding job to one of a plurality of host computing modules in response to an encoding job completion bid. A method includes receiving media content, segmenting the media content and generating a plurality of streamlets, and generating a set of streamlets.

55 Claims, 12 Drawing Sheets



U.S. PATENT DOCUMENTS

6,510,553	B1	1/2003	Hazra	
6,574,591	B1	6/2003	Kleiman et al.	
6,604,118	B2	8/2003	Kleiman et al.	
6,731,600	B1	5/2004	Patel et al.	
6,732,183	B1	5/2004	Graham	
6,968,387	B2	11/2005	Lanphear	
7,075,986	B2	7/2006	Girod et al.	
7,096,271	B1	8/2006	Omoigui et al.	
7,349,976	B1	3/2008	Glaser et al.	
7,408,984	B2 *	8/2008	Lu et al.	375/240.02
2002/0073167	A1	6/2002	Powell et al.	
2002/0144276	A1	10/2002	Radford et al.	
2002/0174434	A1	11/2002	Lee et al.	
2002/0188745	A1	12/2002	Hughes et al.	
2003/0005455	A1 *	1/2003	Bowers	725/90
2003/0067872	A1	4/2003	Harrell et al.	
2003/0081582	A1	5/2003	Jain et al.	
2003/0107994	A1	6/2003	Jacobs et al.	
2003/0135631	A1	7/2003	Li et al.	
2003/0151753	A1	8/2003	Li et al.	358/1.9
2003/0152036	A1	8/2003	Quigg Brown et al.	
2003/0204519	A1 *	10/2003	Sirivara et al.	707/101
2004/0003101	A1	1/2004	Roth et al.	
2004/0030797	A1	2/2004	Akinlar et al.	
2004/0031054	A1	2/2004	Dankworth et al.	
2004/0054551	A1 *	3/2004	Ausubel et al.	705/1
2004/0071209	A1	4/2004	Burg et al.	
2004/0093420	A1	5/2004	Gamble	
2004/0143672	A1 *	7/2004	Padmanabham et al.	709/231
2004/0170392	A1 *	9/2004	Lu et al.	386/96
2005/0108414	A1	5/2005	Taylor et al.	709/231
2005/0123058	A1	6/2005	Greenbaum et al.	
2007/0024705	A1	2/2007	Richter et al.	
2008/0028428	A1	1/2008	Jeong et al.	
2008/0056373	A1	3/2008	Newlin et al.	
2008/0162713	A1 *	7/2008	Bowra et al.	709/231
2008/0195744	A1	8/2008	Bowra et al.	
2008/0195745	A1	8/2008	Bowra et al.	

FOREIGN PATENT DOCUMENTS

EP	0 711 077	A2	5/1996
EP	0 919 952	A1	6/1999
EP	1 641 271	A2	3/2006

EP	1 670 256	A2	6/2006
EP	1 777 969	A1	4/2007
WO	00/67469		11/2000

OTHER PUBLICATIONS

On2 Technologies, Inc., "TrueMotion VP7 Video Codec", White Paper, Document Version 1.0, Jan. 10, 2005, (13 pages).
 Supplemental European Search Report, Sep. 22, 2008, (3 pages).
 Office Action mailed May 14, 2008 for U.S. Appl. No. 11/116,783, filed on Apr. 28, 2005.
 Final Office Action mailed Feb. 20, 2009 for U.S. Appl. No. 11/116,783, filed on Apr. 28, 2005.
 Inteligent Streaming, Bill Birney, May 2003, Microsoft, pp. 7 total.
 Advisory Action for U.S. Appl. No. 11/116,783, Mailed Apr. 02, 2009, 3 pages.
 Advisory Action for U.S. Appl. No. 11/116,783, Mailed May 12, 2009, 4 pages.
 Wicker, Stephen B., "Error Control Systems for Digital Communication and Storage", Prentice-Hall, Inc., New Jersey, USA, 1995 (Book: see NPL's Parts 1-6).
 PCT Notification of Transmittal of the International Search Report and Written Opinion of the International Searching Authority, for PCT/US05/15091, Oct. 29, 2007, 8 pages.
 PCT Notification of Transmittal of International Preliminary Report on Patentability, for PCT/US05/15091, Oct. 29, 2007, 6 pages.
 Office Action for U.S. Appl. No. 11/116,783, Aug. 20, 2009, 14 pages.
 Albanese, Andres, et al. "Priority Encoding Transmission", TR-94-039, Aug. 1994, 36 pages, International Computer Science Institute, Berkeley, California.
 Puri, Rohit, et al. "Multiple Description Source Coding Using Forward Error Correction Codes", Oct. 1999, 5 pages, Department of Electrical Engineering and Computer Science, University of California, Berkeley, California.
 Goyal, Vivek K., "Multiple Description Coding: Compression Meets the Network", Sep. 2001, pp. 74-93, IEEE Signal Processing Magazine.
 Advisory Action for U.S. Appl. No. 11/116,783, Mailed Mar. 23, 2010, 3 pages.
 Advisory Action for U.S. Appl. No. 11/116,783, Mailed May 17, 2010, 3 pages.
 Final Office Action for U.S. Appl. No. 11/116,783, Mailed Feb. 22, 2010, 16 pages.

* cited by examiner




FIG. 1




FIG. 2A




FIG. 2B




FIG. 3A




FIG. 3B




FIG. 4


800
↓

FIG. 5




FIG. 5a




FIG. 5B




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 7

800




FIG. 8




FIG. 9




FIG. 10

1

APPARATUS, SYSTEM, AND METHOD FOR MULTI-BITRATE CONTENT STREAMING

CROSS-REFERENCES TO RELATED APPLICATIONS

This application is a continuation-in-part of application Ser. No. 11/116,783, filed on Apr. 28, 2005, which claims the benefit of U.S. Provisional Application No. 60/566,831, filed on Apr. 30, 2004, both of which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention relates to video streaming over packet switched networks such as the Internet, and more particularly relates to adaptive-rate shifting of streaming content over such networks.

2. Description of the Related Art

The Internet is fast becoming a preferred method for distributing media files to end users. It is currently possible to download music or video to computers, cell phones, or practically any network capable device. Many portable media players are equipped with network connections and enabled to play music or videos. The music or video files (hereinafter "media files") can be stored locally on the media player or computer, or streamed or downloaded from a server.

"Streaming media" refers to technology that delivers content at a rate sufficient for presenting the media to a user in real time as the data is received. The data may be stored in memory temporarily until played and then subsequently deleted. The user has the immediate satisfaction of viewing the requested content without waiting for the media file to be completely downloaded. Unfortunately, the audio/video quality that can be received for real time presentation is constrained by the available bandwidth of the user's network connection. Streaming may be used to deliver content on demand (previously recorded) or from live broadcasts.

Alternatively, media files may be downloaded and stored on persistent storage devices, such as hard drives or optical storage, for later presentation. Downloading complete media files can take large amounts of time depending on the network connection. Once downloaded, however, the content can be viewed repeatedly anytime or anywhere. Media files prepared for downloading usually are encoded with a higher quality audio/video than can be delivered in real time. Users generally dislike this option, as they tend to want to see or hear the media file instantaneously.

Streaming offers the advantage of immediate access to the content but currently sacrifices quality compared with downloading a file of the same content. Streaming also provides the opportunity for a user to select different content for viewing on an ad hoc basis, while downloading is by definition restricted to receiving a specific content selection in its entirety or not at all. Downloading also supports rewind, fast forward, and direct seek operations, while streaming is unable to fully support these functions. Streaming is also vulnerable to network failures or congestion.

Another technology, known as "progressive downloads," attempts to combine the strengths of the above two technologies. When a progressive download is initiated, the media file download begins, and the media player waits to begin playback until there is enough of the file downloaded that playback can begin with the hope that the remainder of the file will be completely downloaded before playback "catches up." This waiting period before playback can be substantial

2

depending on network conditions, and therefore is not a complete or fully acceptable solution to the problem of media presentation over a network.

Generally, three basic challenges exist with regard to data transport streaming over a network such as the Internet that has a varying amount of data loss. The first challenge is reliability. Most streaming solutions use a TCP connection, or "virtual circuit," for transmitting data. A TCP connection provides a guaranteed delivery mechanism so that data sent from one endpoint will be delivered to the destination, even if portions are lost and retransmitted. A break in the continuity of a TCP connection can have serious consequences when the data must be delivered in real-time. When a network adapter detects delays or losses in a TCP connection, the adapter "backs off" from transmission attempts for a moment and then slowly resumes the original transmission pace. This behavior is an attempt to alleviate the perceived congestion. Such a slowdown is detrimental to the viewing or listening experience of the user and therefore is not acceptable.

The second challenge to data transport is efficiency. Efficiency refers to how well the user's available bandwidth is used for delivery of the content stream. This measure is directly related to the reliability of the TCP connection. When the TCP connection is suffering reliability problems, a loss of bandwidth utilization results. The measure of efficiency sometimes varies suddenly, and can greatly impact the viewing experience.

The third challenge is latency. Latency is the time measure from the client's point-of-view, of the interval between when a request is issued and the response data begins to arrive. This value is affected by the network connection's reliability and efficiency, and the processing time required by the origin to prepare the response. A busy or overloaded server, for example, will take more time to process a request. As well as affecting the start time of a particular request, latency has a significant impact on the network throughput of TCP.

From the foregoing discussion, it should be apparent that a need exists for an apparatus, system, and method that alleviate the problems of reliability, efficiency, and latency. Additionally, such an apparatus, system, and method would offer instantaneous viewing along with the ability to fast forward, rewind, direct seek, and browse multiple streams. Beneficially, such an apparatus, system, and method would utilize multiple connections between a source and destination, requesting varying bitrate streams depending upon network conditions.

SUMMARY OF THE INVENTION

The present invention has been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available content streaming systems. Accordingly, the present invention has been developed to provide an apparatus, system, and method for adaptive-rate content streaming that overcome many or all of the above-discussed shortcomings in the art.

The apparatus for adaptive-rate content streaming is provided with a logic unit containing a plurality of modules configured to functionally execute the necessary steps. These modules in the described embodiments include a receiving module configured to receive media content, a streamlet module configured to segment the media content and generate a plurality of sequential streamlets, and an encoding module configured to encode each streamlet as a separate content file.

The encoding module is further configured to generate a set of streamlets for each of the sequential streamlets. Each

streamlet may comprise a portion of the media content having a predetermined length of time. The predetermined length of time may be in the range of between about 0.1 and 5 seconds.

In one embodiment, a set of streamlets comprises a plurality of streamlets having identical time indices, and each streamlet of the set of streamlets has a unique bitrate. The receiving module is configured to convert the media content to raw audio or raw video. The encoding module may include a master module configured to assign an encoding job to one of a plurality of host computing modules in response to an encoding job completion bid. The job completion bid may be based on a plurality of computing variables selected from a group consisting of current encoding job completion percentage, average encoding job completion time, processor speed, and physical memory capacity.

A system of the present invention is also presented for adaptive-rate content streaming. In particular, the system, in one embodiment, includes a receiving module configured to receive media content, a streamlet module configured to segment the media content and generate a plurality of sequential streamlets, each streamlet comprising a portion of the media content having a predetermined length of time, and an encoding module configured to encode each streamlet as a separate content file and generate a set of streamlets.

The system also includes a plurality of streamlets having identical time indices and each streamlet of the set of streamlets having a unique bitrate. The encoding module comprises a master module configured to assign an encoding job to one of a plurality of host computing modules in response to an encoding job completion bid.

A method of the present invention is also presented for adaptive-rate content streaming. In one embodiment, the method includes receiving media content, segmenting the media content and generating a plurality of sequential streamlets, and encoding each streamlet as a separate content file.

The method also includes segmenting the media content into a plurality of streamlets, each streamlet comprising a portion of the media content having a predetermined length of time. In one embodiment, the method includes generating a set of streamlets comprising a plurality of streamlets having identical time indices, and each streamlet of the set of streamlets having a unique bitrate.

Furthermore, the method may include converting the media content to raw audio or raw video, and segmenting the content media into a plurality of sequential streamlets. The method further comprises assigning an encoding job to one of a plurality of host computing modules in response to an encoding job completion bid, and submitting an encoding job completion bid based on a plurality of computing variables.

Reference throughout this specification to features, advantages, or similar language does not imply that all of the features and advantages that may be realized with the present invention should be or are in any single embodiment of the invention. Rather, language referring to the features and advantages is understood to mean that a specific feature, advantage, or characteristic described in connection with an embodiment is included in at least one embodiment of the present invention. Thus, discussion of the features and advantages, and similar language, throughout this specification may, but do not necessarily, refer to the same embodiment.

Furthermore, the described features, advantages, and characteristics of the invention may be combined in any suitable manner in one or more embodiments. One skilled in the relevant art will recognize that the invention may be practiced without one or more of the specific features or advantages of a particular embodiment. In other instances, additional fea-

tures and advantages may be recognized in certain embodiments that may not be present in all embodiments of the invention.

These features and advantages of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the advantages of the invention will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings, in which:

FIG. 1 is a schematic block diagram illustrating one embodiment of a system for dynamic rate shifting of streaming content in accordance with the present invention;

FIG. 2a is a schematic block diagram graphically illustrating one embodiment of a media content file;

FIG. 2b is a schematic block diagram illustrating one embodiment of a plurality of streams having varying degrees of quality and bandwidth;

FIG. 3a is a schematic block diagram illustrating one embodiment of a stream divided into a plurality of source streamlets;

FIG. 3b is a schematic block diagram illustrating one embodiment of sets of streamlets in accordance with the present invention;

FIG. 4 is a schematic block diagram illustrating in greater detail one embodiment of the content module in accordance with the present invention;

FIG. 5a is a schematic block diagram illustrating one embodiment of an encoder module in accordance with the present invention;

FIG. 5b is a schematic block diagram illustrating one embodiment of parallel encoding of streamlets in accordance with the present invention;

FIG. 6a is a schematic block diagram illustrating one embodiment of a virtual timeline in accordance with the present invention;

FIG. 6b is a schematic block diagram illustrating an alternative embodiment of a VT in accordance with the present invention;

FIG. 6c is a schematic block diagram illustrating one embodiment of a QMX in accordance with the present invention;

FIG. 7 is a schematic block diagram graphically illustrating one embodiment of a client module in accordance with the present invention;

FIG. 8 is a schematic flow chart diagram illustrating one embodiment of a method for processing content in accordance with the present invention;

FIG. 9 is a schematic flow chart diagram illustrating one embodiment of a method for viewing a plurality of streamlets in accordance with the present invention; and

FIG. 10 is a schematic flow chart diagram illustrating one embodiment of a method for requesting streamlets within a adaptive-rate shifting content streaming environment in accordance with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Many of the functional units described in this specification have been labeled as modules, in order to more particularly emphasize their implementation independence. For example, a module may be implemented as a hardware circuit comprising custom VLSI circuits or gate arrays, off-the-shelf semi-conductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices or the like.

Modules may also be implemented in software for execution by various types of processors. An identified module of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions which may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together, but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module.

Indeed, a module of executable code may be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices. Similarly, operational data may be identified and illustrated herein within modules, and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set, or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network.

Reference throughout this specification to “one embodiment,” “an embodiment,” or similar language means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases “in one embodiment,” “in an embodiment,” and similar language throughout this specification may, but do not necessarily, all refer to the same embodiment.

Reference to a signal bearing medium may take any form capable of generating a signal, causing a signal to be generated, or causing execution of a program of machine-readable instructions on a digital processing apparatus. A signal bearing medium may be embodied by a transmission line, a compact disk, digital-video disk, a magnetic tape, a Bernoulli drive, a magnetic disk, a punch card, flash memory, integrated circuits, or other digital processing apparatus memory device.

Furthermore, the described features, structures, or characteristics of the invention may be combined in any suitable manner in one or more embodiments. In the following description, numerous specific details are provided, such as examples of programming, software modules, user selections, network transactions, database queries, database structures, hardware modules, hardware circuits, hardware chips, etc., to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention may be practiced without one or more of the specific details, or with other methods, components, materials, and so forth. In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

FIG. 1 is a schematic block diagram illustrating one embodiment of a system 100 for dynamic rate shifting of streaming content in accordance with the present invention. In one embodiment, the system 100 comprises a content

server 102 and an end user station 104. The content server 102 and the end user station 104 may be coupled by a data communications network. The data communications network may include the Internet 106 and connections 108 to the Internet 106. Alternatively, the content server 102 and the end user 104 may be located on a common local area network, wireless area network, cellular network, virtual local area network, or the like. The end user station 104 may comprise a personal computer (PC), an entertainment system configured to communicate over a network, or a portable electronic device configured to present content. For example, portable electronic devices may include, but are not limited to, cellular phones, portable gaming systems, and portable computing devices.

In the depicted embodiment, the system 100 also includes a publisher 110, and a web server 116. The publisher 110 may be a creator or distributor of content. For example, if the content to be streamed were a broadcast of a television program, the publisher 110 may be a television or cable network channel such as NBC®, or MTV®. Content may be transferred over the Internet 106 to the content server 102, where the content is received by a content module 112. The content module 112 may be configured to receive, process, and store content. In one embodiment, processed content is accessed by a client module 114 configured to play the content on the end user station 104. In a further embodiment, the client module 114 is configured to receive different portions of a content stream from a plurality of locations simultaneously. For example, the client module 114 may request and receive content from any of the plurality of web servers 116.

Content from the content server 102 may be replicated to other web servers 116 or alternatively to proxy cache servers 118. Replicating may occur by deliberate forwarding from the content server 102, or by a web, cache, or proxy server outside of the content server 102 asking for content on behalf of the client module 114. In a further embodiment, content may be forwarded directly to web 116 or proxy 118 servers through direct communication channels 120 without the need to traverse the Internet 106.

FIG. 2a is a schematic block diagram graphically illustrating one embodiment of a media content (hereinafter “content”) file 200. In one embodiment, the content file 200 is distributed by the publisher 110. The content file 200 may comprise a television broadcast, sports event, movie, music, concert, etc. The content file 200 may also be live or archived content. The content file 200 may comprise uncompressed video and audio, or alternatively, video or audio. Alternatively, the content file 200 may be compressed using standard or proprietary encoding schemes. Examples of encoding schemes capable of use with the present invention include, but are not limited to, DivX®, Windows Media Video®, Quicktime Sorenson 3®, On2, OGG Vorbis, MP3, or Quicktime 6.5/MPEG-4® encoded content.

FIG. 2b is a schematic block diagram illustrating one embodiment of a plurality of streams 202 having varying degrees of quality and bandwidth. In one embodiment, the plurality of streams 202 comprises a low quality stream 204, a medium quality stream 206, and a high quality stream 208. Each of the streams 204, 206, 208 is a copy of the content file 200 encoded and compressed to varying bit rates. For example, the low quality stream 204 may be encoded and compressed to a bit rate of 100 kilobits per second (kbps), the medium quality stream 206 may be encoded and compressed to a bit rate of 200 kbps, and the high quality stream 208 may be encoded and compressed to 600 kbps.

FIG. 3a is a schematic block diagram illustrating one embodiment of a stream 302 divided into a plurality of source

7

streamlets **303**. As used herein, streamlet refers to any sized portion of the content file **200**. Each streamlet **303** may comprise a portion of the content contained in stream **302**, encapsulated as an independent media object. The content in a streamlet **303** may have a unique time index in relation to the beginning of the content contained in stream **302**. In one embodiment, the content contained in each streamlet **303** may have a duration of two seconds. For example, streamlet **0** may have a time index of 00:00 representing the beginning of content playback, and streamlet **1** may have a time index of 00:02, and so on. Alternatively, the time duration of the streamlets **304** may be any duration smaller than the entire playback duration of the content in stream **302**. In a further embodiment, the streamlets **303** may be divided according to file size instead of a time index and duration.

FIG. **3b** is a schematic block diagram illustrating one embodiment of sets **306** of streamlets in accordance with the present invention. As used herein, the term "set" refers to a group of streamlets having identical time indices and durations but varying bitrates. In the depicted embodiment, the set **306a** encompasses all streamlets having a time index of 00:00. The set **306a** includes encoded streamlets **304** having low, medium, and high **204**, **206**, **208** bitrates. Of course each set **306** may include more than the depicted three bitrates which are given by way of example only. One skilled in the art will recognize that any number of streams having different bitrates may be generated from the original content **200**.

As described above, the duration of one streamlet **304** may be approximately two seconds. Likewise each set **306** may comprise a plurality of streamlets **304** where each streamlet **304** has a playable duration of two seconds. Alternatively, the duration of the streamlet **304** may be predetermined or dynamically variable depending upon a variety of factors including, but not limited to, network congestion, system specifications, playback resolution and quality, etc. In the depicted embodiment, the content **200** may be formed of the plurality of sets **306**. The number of sets **306** may depend on the length of the content **200** and the length or duration of each streamlet **304**.

FIG. **4** is a schematic block diagram illustrating in greater detail one embodiment of the content module **112** in accordance with the present invention. The content module **112** may comprise a capture module **402**, a streamlet module **404**, an encoder module **406**, a streamlet database **408**, and the web server **116**. In one embodiment, the capture module **402** is configured to receive the content file **200** from the publisher **110**. The capture module **402** may be configured to "decompress" the content file **200**. For example, if the content file **200** arrives having been encoded with one of the above described encoding schemes, the capture module **402** may convert the content file **200** into raw audio and/or video. Alternatively, the content file **200** may be transmitted by the publisher in a format **110** that does not require decompression.

The capture module **402** may comprise a capture card configured for TV and/or video capture. One example of a capture card suitable for use in the present invention is the DRC-2500 by Digital Rapids of Ontario, Canada. Alternatively, any capture card capable of capturing audio and video may be utilized with the present invention. In a further embodiment, the capture module **402** is configured to pass the content file to the streamlet module **404**.

The streamlet module **404**, in one embodiment, is configured to segment the content file **200** and generate source streamlets **303** that are not encoded. As used herein, the term "segment" refers to an operation to generate a streamlet of the content file **200** having a duration or size equal to or less than the duration or size of the content file **200**. The streamlet

8

module **404** may be configured to segment the content file **200** into streamlets **303** each having an equal duration. Alternatively, the streamlet module **404** may be configured to segment the content file **200** into streamlets **303** having equal file sizes.

The encoding module **406** is configured to receive the source streamlets **303** and generate the plurality of streams **202** of varying qualities. The original content file **200** from the publisher may be digital in form and may comprise content having a high bit rate such as, for example, 2 mbps. The content may be transferred from the publisher **110** to the content module **112** over the Internet **106**. Such transfers of data are well known in the art and do not require further discussion herein. Alternatively, the content may comprise a captured broadcast.

In a further embodiment, the encoding module **406** is configured to generate a plurality of sets **306** of streamlets **304**. The sets **306**, as described above with reference to FIG. **3b**, may comprise streamlets having an identical time index and duration, and a unique bitrate. As with FIG. **3b**, the sets **306** and subsequently the plurality of streams **202** may comprise the low quality stream **204**, the medium quality stream **206**, and the high quality stream **208**. Alternatively, the plurality of streams **202** may comprise any number of streams deemed necessary to accommodate end user bandwidth.

The encoder module **406** is further configured to encode each source streamlet **303** into the plurality of streams **202** and streamlet sets **306** and store the streamlets in the streamlet database **408**. The encoding module **406** may utilize encoding schemes such as DivX®, Windows Media Video 9®, Quicktime 6.5 Sorenson 3®, or Quicktime 6.5/MPEG-4®. Alternatively, a custom encoding scheme may be employed.

The content module **112** may also include a metadata module **412** and a metadata database **414**. In one embodiment, metadata comprises static searchable content information. For example, metadata includes, but is not limited to, air date of the content, title, actresses, actors, length, and episode name. Metadata is generated by the publisher **110**, and may be configured to define an end user environment. In one embodiment, the publisher **100** may define an end user navigational environment for the content including menus, thumbnails, sidebars, advertising, etc. Additionally, the publisher **110** may define functions such as fast forward, rewind, pause, and play that may be used with the content file **200**. The metadata module **412** is configured to receive the metadata from the publisher **110** and store the metadata in the metadata database **414**. In a further embodiment, the metadata module **412** is configured to interface with the client module **114**, allowing the client module **114** to search for content based upon at least one of a plurality of metadata criteria. Additionally, metadata may be generated by the content module **112** through automated process(es) or manual definition.

Once the streamlets **304** have been received and processed, the client module **114** may request streamlets **304** using HTTP from the web server **116**. Using a standard protocol such as HTTP eliminates the need for network administrators to configure firewalls to recognize and pass through network traffic for a new, specialized protocol. Additionally, since the client module **114** initiates the request, the web server **116** is only required to retrieve and serve the requested streamlet **304**. In a further embodiment, the client module **114** may be configured to retrieve streamlets **304** from a plurality of web servers **116**.

Each web server **116** may be located in various locations across the Internet **106**. The streamlets **304** may essentially be static files. As such, no specialized media server or server-side intelligence is required for a client module **114** to retrieve

streamlets **304**. Streamlets **304** may be served by the web server **116** or cached by cache servers of Internet Service Providers (ISPs), or any other network infrastructure operators, and served by the cache server. Use of cache servers is well known to those skilled in the art, and will not be discussed further herein. Thus, a highly scalable solution is provided that is not hindered by massive amounts of client module **114** requests to the web server **116** at any specific location, especially the web server **116** most closely associated with or within the content module **112**.

FIG. **5a** is a schematic block diagram illustrating one embodiment of an encoder module **406** in accordance with the present invention. In one embodiment, the encoder module **406** may include a master module **502** and a plurality of host computing modules (hereinafter "host") **504**. The hosts **504** may comprise personal computers, servers, etc. In a further embodiment, the hosts **504** may be dedicated hardware, for example, cards plugged into a single computer.

The master module (hereinafter "master") **502** is configured to receive streamlets **303** from the streamlet module **404** and stage the streamlet **303** for processing. In one embodiment, the master **502** may decompress each source streamlet **303** to produce a raw streamlet. As used herein, the term "raw streamlet" refers to a streamlet **303** that is uncompressed or lightly compressed to substantially reduce size with no significant loss in quality. A lightly compressed raw streamlet can be transmitted more quickly and to more hosts. Each host **504** is coupled with the master **502** and configured to receive a raw streamlet from the master **502** for encoding. The hosts **504**, in one example, generate a plurality of streamlets **304** having identical time indices and durations, and varying bitrates. Essentially each host **504** may be configured to generate a set **306** from the raw streamlet **503** sent from the master **502**. Alternatively, each host **504** may be dedicated to producing a single bitrate in order to reduce the time required for encoding.

Upon encoding completion, the host **504** returns the set **306** to the master **502** so that the encoding module **406** may store the set **306** in the streamlet database **408**. The master **502** is further configured to assign encoding jobs to the hosts **504**. Each host is configured to submit an encoding job completion bid (hereinafter "bid"). The master **502** assigns encoding jobs depending on the bids from the hosts **504**. Each host **504** generates a bid depending upon a plurality of computing variables which may include, but are not limited to, current encoding job completion percentage, average job completion time, processor speed and physical memory capacity.

For example, a host **504** may submit a bid that indicates that based on past performance history the host **504** would be able to complete the encoding job in 15 seconds. The master **502** is configured to select from among a plurality of bids the best bid and subsequently submit the encoding job to the host **504** with the best bid. As such, the described encoding system does not require that each host **504** have identical hardware but beneficially takes advantage of the available computing power of the hosts **504**. Alternatively, the master **502** selects the host **504** based on a first come first serve basis, or some other algorithm deemed suitable for a particular encoding job.

The time required to encode one streamlet **304** is dependent upon the computing power of the host **504**, and the encoding requirements of the content file **200**. Examples of encoding requirements may include, but are not limited to, two or multi-pass encoding, and multiple streams of different bitrates. One benefit of the present invention is the ability to perform two-pass encoding on a live content file **200**. Typi-

cally, in order to perform two-pass encoding prior art systems must wait for the content file to be completed before encoding.

The present invention, however, segments the content file **200** into source streamlets **303** and the two-pass encoding to a plurality of streams **202** may be performed on each corresponding raw streamlet without waiting for a TV show to end, for example. As such, the content module **112** is capable of streaming the streamlets over the Internet shortly after the content module **112** begins capture of the content file **200**. The delay between a live broadcast transmitted from the publisher **110** and the availability of the content depends on the computing power of the hosts **504**.

FIG. **5b** is a schematic block diagram illustrating one embodiment of parallel encoding of streamlets in accordance with the present invention. In one example, the capture module **402** (of FIG. **4**) begins to capture the content file and the streamlet module **404** generates a first streamlet **303a** and passes the streamlet to the encoding module **406**. The encoding module **406** may take 10 seconds, for example, to generate the first set **306a** of streamlets **304a** (**304a₁**, **304a₂**, **304a₃**, etc. represent streamlets **304** of different bitrates). FIG. **5b** illustrates the encoding process generically as block **502** to graphically illustrate the time duration required to process a raw or lightly encoded streamlet **303** as described above with reference to the encoding module **406**. The encoding module **406** may simultaneously process more than one streamlet **303**, and processing of streamlets will begin upon arrival of the streamlet from the capture module **402**.

During the 10 seconds required to encode the first streamlet **303a**, the streamlet module **404** has generated five additional 2-second streamlets **303b**, **303c**, **303d**, **303e**, **303f**, for encoding and the master **502** has prepared and staged the corresponding raw streamlets. Two seconds after the first set **306a** is available the next set **306b** is available, and so on. As such, the content file **200** is encoded for streaming over the Internet and appears live. The 10 second delay is given herein by way of example only. Multiple hosts **504** may be added to the encoding module **406** in order to increase the processing capacity of the encoding module **406**. The delay may be shortened to an almost unperceivable level by the addition of high CPU powered systems, or alternatively multiple low powered systems.

A system as described above beneficially enables multi-pass encoding of live events. Multi-pass encoding systems of the prior art require that the entire content be captured (or be complete) because in order to perform multi-pass encoding the entire content must be scanned and processed more than once. This is impossible with prior art systems because content from a live event is not complete until the event is over. As such, with prior art systems, multi-pass encoding can only be performed once the event is over. Streamlets, however, may be encoded as many times as is deemed necessary. Because the streamlet is an encapsulated media object of 2 seconds (for example), multi-pass encoding may begin on a live event once the first streamlet is captured. Shortly after multi-pass encoding of the first streamlet **303a** is finished, multi-pass encoding of the second streamlet **303b** finishes, and as such multi-pass encoding is performed on a live event and appears live to a viewer.

Any specific encoding scheme applied to a streamlet may take longer to complete than the time duration of the streamlet itself, for example, a very high quality encoding of a 2-second streamlet may take 5 seconds to finish. Alternatively, the processing time required for each streamlet may be less than the time duration of a streamlet. However, because the offset parallel encoding of successive streamlets are encoded by the

11

encoding module at regular intervals (matching the intervals at which the those streamlets are submitted to the encoding module **406**, for example 2 seconds) the output timing of the encoding module **406** does not fall behind the real-time submission rate of the unencoded streamlets. Conversely, prior art encoding systems rely on the very fastest computing hardware and software because the systems must generate the output immediately in lock-step with the input. A prior art system that takes 2.1 seconds to encode 2 seconds worth of content is considered a failure. The present invention allows for slower than real-time encoding processes yet still achieves a real-time encoding effect due to the parallel offset pipes.

The parallel offset pipeline approach described with reference to FIG. **5b** beneficially allows for long or short encoding times without “falling behind” the live event. Additionally, arbitrarily complex encoding of streamlets to multiple profiles and optimizations only lengthens the encoding time **502** without a perceptible difference to a user because the sets **306** of streamlets **304** are encoded in a time-selective manner so that streamlets are processed at regular time intervals and transmitted at these time intervals.

Returning now to FIG. **5a**, as depicted, the master **502** and the hosts **504** may be located within a single local area network, or in other terms, the hosts **504** may be in close physical proximity to the master **502**. Alternatively, the hosts **504** may receive encoding jobs from the master **502** over the Internet or other communications network. For example, consider a live sports event in a remote location where it would be difficult to setup multiple hosts. In this example, a master performs no encoding or alternatively light encoding before publishing the streamlets online. The hosts **504** would then retrieve those streamlets and encode the streamlets into the multiple bitrate sets **306** as described above.

Furthermore, hosts **504** may be dynamically added or removed from the encoding module without restarting the encoding job and/or interrupting the publishing of streamlets. If a host **504** experiences a crash or some failure, its encoding work is simply reassigned to another host.

The encoding module **406**, in one embodiment, may also be configured to produce streamlets that are specific to a particular playback platform. For example, for a single raw streamlet, a single host **504** may produce streamlets for different quality levels for personal computer playback, streamlets for playback on cell phones with a different, proprietary codec, a small video-only streamlet for use when playing just a thumbnail view of the stream (like in a programming guide), and a very high quality streamlet for use in archiving.

FIG. **6a** is a schematic block diagram illustrating one embodiment of a virtual timeline **600** in accordance with the present invention. In one embodiment, the virtual timeline **600** comprises at least one quantum media extension **602**. The quantum media extension (hereinafter “QMX”) **602** describes an entire content file **200**. Therefore, the virtual timeline (hereinafter “VT”) **600** may comprise a file that is configured to define a playlist for a user to view. For example, the VT may indicate that the publisher desires a user to watch a first show QMX **602a** followed by QMX **602b** and QMX **602c**. As such, the publisher may define a broadcast schedule in a manner similar to a television station.

FIG. **6b** is a schematic block diagram illustrating an alternative embodiment of a VT **600** in accordance with the present invention. In the depicted embodiment, the VT **600** may include a single QMX **602** which indicates that the publisher desires the same content to be looped over and over again. For example, the publisher may wish to broadcast a never-ending infomercial on a website.

12

FIG. **6c** is a schematic block diagram illustrating one embodiment of a QMX **602** in accordance with the present invention. In one embodiment, the QMX **602** contains a multitude of information generated by the content module **112** configured to describe the content file **200**. Examples of information include, but are not limited to, start index **604**, end index **606**, whether the content is live **608**, proprietary publisher data **610**, encryption level **612**, content duration **614** and bitrate values **616**. The bitrate values **616** may include frame size **618**, audio channel **620** information, codecs **622** used, sample rate **624**, and frames parser **626**.

A publisher may utilize the QVT **600** together with the QMX **602** in order to prescribe a playback order for users, or alternatively selectively edit content. For example, a publisher may indicate in the QMX **602** that audio should be muted at time index 10:42 or video should be skipped for 3 seconds at time index 18:35. As such, the publisher may selectively skip offensive content without the processing requirements of editing the content.

FIG. **7** is a schematic block diagram graphically illustrating one embodiment of a client module **114** in accordance with the present invention. The client module **114** may comprise an agent controller module **702**, a streamlet cache module **704**, and a network controller module **706**. In one embodiment, the agent controller module **702** is configured to interface with a viewer **708**, and transmit streamlets **304** to the viewer **708**. Alternatively, the agent controller module **702** may be configured to simply reassemble streamlets into a single file for transfer to an external device such as a portable video player.

In a further embodiment, the client module **114** may comprise a plurality of agent controller modules **702**. Each agent controller module **702** may be configured to interface with one viewer **708**. Alternatively, the agent controller module **702** may be configured to interface with a plurality of viewers **708**. The viewer **708** may be a media player (not shown) operating on a PC or handheld electronic device.

The agent controller module **702** is configured to select a quality level of streamlets to transmit to the viewer **708**. The agent controller module **702** requests lower or higher quality streams based upon continuous observation of time intervals between successive receive times of each requested streamlet. The method of requesting higher or lower quality streams will be discussed in greater detail below with reference to FIG. **10**.

The agent controller module **702** may be configured to receive user commands from the viewer **708**. Such commands may include play, fast forward, rewind, pause, and stop. In one embodiment, the agent controller module **702** requests streamlets **304** from the streamlet cache module **704** and arranges the received streamlets **304** in a staging module **709**. The staging module **709** may be configured to arrange the streamlets **304** in order of ascending playback time. In the depicted embodiment, the streamlets **304** are numbered 0, 1, 2, 3, 4, etc. However, each streamlet **304** may be identified with a unique filename.

Additionally, the agent controller module **702** may be configured to anticipate streamlet **304** requests and pre-request streamlets **304**. By pre-requesting streamlets **304**, the user may fast-forward, skip randomly, or rewind through the content and experience no buffering delay. In a further embodiment, the agent controller module **702** may request the streamlets **304** that correspond to time index intervals of 30 seconds within the total play time of the content. Alternatively, the agent controller module **702** may request streamlets at any interval less than the length of the time index. This enables a “fast-start” capability with no buffering wait when starting or fast-forwarding through content file **200**. In a fur-

13

ther embodiment, the agent controller module 702 may be configured to pre-request streamlets 304 corresponding to specified index points within the content or within other content in anticipation of the end user 104 selecting new content to view.

In one embodiment, the streamlet cache module 704 is configured to receive streamlet 304 requests from the agent controller module 702. Upon receiving a request, the streamlet cache module 704 first checks a streamlet cache 710 to verify if the streamlet 304 is present. In a further embodiment, the streamlet cache module 704 handles streamlet 304 requests from a plurality of agent controller modules 702. Alternatively, a streamlet cache module 704 may be provided for each agent controller module 702. If the requested streamlet 304 is not present in the streamlet cache 710, the request is passed to the network controller module 706. In order to enable fast forward and rewind capabilities, the streamlet cache module 704 is configured to store the plurality of streamlets 304 in the streamlet cache 710 for a specified time period after the streamlet 304 has been viewed. However, once the streamlets 304 have been deleted, they may be requested again from the web server 116.

The network controller module 706 may be configured to receive streamlet requests from the streamlet cache module 704 and open a connection to the web server 116 or other remote streamlet 304 database (not shown). In one embodiment, the network controller module 706 opens a TCP/IP connection to the web server 116 and generates a standard HTTP GET request for the requested streamlet 304. Upon receiving the requested streamlet 304, the network controller module 706 passes the streamlet 304 to the streamlet cache module 704 where it is stored in the streamlet cache 710. In a further embodiment, the network controller module 706 is configured to process and request a plurality of streamlets 304 simultaneously. The network controller module 706 may also be configured to request a plurality of streamlets, where each streamlet 304 is subsequently requested in multiple parts.

In a further embodiment, streamlet requests may comprise requesting pieces of any streamlet file. Splitting the streamlet 304 into smaller pieces or portions beneficially allows for an increased efficiency potential, and also eliminates problems associated with multiple full-streamlet requests sharing the bandwidth at any given moment. This is achieved by using parallel TCP/IP connections for pieces of the streamlets 304. Consequently, efficiency and network loss problems are overcome, and the streamlets arrive with more useful and predictable timing.

In one embodiment, the client module 114 is configured to use multiple TCP connections between the client module 114 and the web server 116 or web cache. The intervention of a cache may be transparent to the client or configured by the client as a forward cache. By requesting more than one streamlet 304 at a time in a manner referred to as "parallel retrieval," or more than one part of a streamlet 304 at a time, efficiency is raised significantly and latency is virtually eliminated. In a further embodiment, the client module allows a maximum of three outstanding streamlet 304 requests. The client module 114 may maintain additional open TCP connections as spares to be available should another connection fail. Streamlet 304 requests are rotated among all open connections to keep the TCP flow logic for any particular connection from falling into a slow-start or close mode. If the network controller module 706 has requested a streamlet 304 in multiple parts, with each part requested on mutually independent TCP/IP connections, the network controller module 706 reassembles the parts to present a complete streamlet 304 for use by all other components of the client module 114.

14

When a TCP connection fails completely, a new request may be sent on a different connection for the same streamlet 304. In a further embodiment, if a request is not being satisfied in a timely manner, a redundant request may be sent on a different connection for the same streamlet 304. If the first streamlet request's response arrives before the redundant request response, the redundant request can be aborted. If the redundant request response arrives before the first request response, the first request may be aborted.

Several streamlet 304 requests may be sent on a single TCP connection, and the responses are caused to flow back in matching order along the same connection. This eliminates all but the first request latency. Because multiple responses are always being transmitted, the processing latency of each new streamlet 304 response after the first is not a factor in performance. This technique is known in the industry as "pipelining." Pipelining offers efficiency in request-response processing by eliminating most of the effects of request latency. However, pipelining has serious vulnerabilities. Transmission delays affect all of the responses. If the single TCP connection fails, all of the outstanding requests and responses are lost. Pipelining causes a serial dependency between the requests.

Multiple TCP connections may be opened between the client module 114 and the web server 116 to achieve the latency-reduction efficiency benefits of pipelining while maintaining the independence of each streamlet 304 request. Several streamlet 304 requests may be sent concurrently, with each request being sent on a mutually distinct TCP connection. This technique is labeled "virtual pipelining" and is an innovation of the present invention. Multiple responses may be in transit concurrently, assuring that communication bandwidth between the client module 114 and the web server 116 is always being utilized. Virtual pipelining eliminates the vulnerabilities of traditional pipelining. A delay in or complete failure of one response does not affect the transmission of other responses because each response occupies an independent TCP connection. Any transmission bandwidth not in use by one of multiple responses (whether due to delays or TCP connection failure) may be utilized by other outstanding responses.

A single streamlet 304 request may be issued for an entire streamlet 304, or multiple requests may be issued, each for a different part or portion of the streamlet. If the streamlet is requested in several parts, the parts may be recombined by the client module 114 streamlet.

In order to maintain a proper balance between maximized bandwidth utilization and response time, the issuance of new streamlet requests must be timed such that the web server 116 does not transmit the response before the client module 114 has fully received a response to one of the previously outstanding streamlet requests. For example, if three streamlet 304 requests are outstanding, the client module 114 should issue the next request slightly before one of the three responses is fully received and "out of the pipe." In other words, request timing is adjusted to keep three responses in transit. Sharing of bandwidth among four responses diminishes the net response time of the other three responses. The timing adjustment may be calculated dynamically by observation, and the request timing adjusted accordingly to maintain the proper balance of efficiency and response times.

The schematic flow chart diagrams that follow are generally set forth as logical flow chart diagrams. As such, the depicted order and labeled steps are indicative of one embodiment of the presented method. Other steps and methods may be conceived that are equivalent in function, logic, or effect to one or more steps, or portions thereof, of the illustrated

15

method. Additionally, the format and symbols employed are provided to explain the logical steps of the method and are understood not to limit the scope of the method. Although various arrow types and line types may be employed in the flow chart diagrams, they are understood not to limit the scope of the corresponding method. Indeed, some arrows or other connectors may be used to indicate only the logical flow of the method. For instance, an arrow may indicate a waiting or monitoring period of unspecified duration between enumerated steps of the depicted method. Additionally, the order in which a particular method occurs may or may not strictly adhere to the order of the corresponding steps shown.

FIG. 8 is a schematic flow chart diagram illustrating one embodiment of a method 800 for processing content in accordance with the present invention. In one embodiment the method 800 starts 802, and the content module 112 receives 804 content from the publisher 110. Receiving content 804 may comprise receiving 804 a digital copy of the content file 200, or digitizing a physical copy of the content file 200. Alternatively, receiving 804 content may comprise capturing a radio, television, cable, or satellite broadcast. Once received 804, the streamlet module 404 generates 808 a plurality of source streamlets 303 each having a fixed duration. Alternatively, the streamlets 303 may be generated with a fixed file size.

In one embodiment, generating 808 streamlets comprises dividing the content file 200 into a plurality of two second streamlets 303. Alternatively, the streamlets may have any length less than or equal to the length of the stream 202. The encoder module 406 then encodes 810 the streamlets 303 into sets 306 of streamlets 304, in a plurality of streams 202 according to an encoding scheme. The quality may be pre-defined, or automatically set according to end user bandwidth, or in response to pre-designated publisher guidelines.

In a further embodiment, the encoding scheme comprises a proprietary codec such as WMV9®. The encoder module 406 then stores 812 the encoded streamlets 304 in the streamlet database 408. Once stored 812, the web server 116 may then serve 814 the streamlets 304. In one embodiment, serving 814 the streamlets 304 comprises receiving streamlet requests from the client module 114, retrieving the requested streamlet 304 from the streamlet database 408, and subsequently transmitting the streamlet 304 to the client module 114. The method 800 then ends 816.

FIG. 9 is a schematic flow chart diagram illustrating one embodiment of a method 900 for viewing a plurality of streamlets in accordance with the present invention. The method 900 starts and an agent control controller module 702 is provided 904 and associated with a viewer 708 and provided with a staging module 709. The agent controller module 702 then requests 906 a streamlet 304 from the streamlet cache module 704. Alternatively, the agent controller module 702 may simultaneously request 906 a plurality of streamlets 304 the streamlet cache module 704. If the streamlet is stored 908 locally in the streamlet cache 710, the streamlet cache module 704 retrieves 910 the streamlet 304 and sends the streamlet to the agent controller module 702. Upon retrieving 910 or receiving a streamlet, the agent controller module 702 makes 911 a determination of whether or not to shift to a higher or lower quality stream 202. This determination will be described below in greater detail with reference to FIG. 10.

In one embodiment, the staging module 709 then arranges 912 the streamlets 304 into the proper order, and the agent controller module 702 delivers 914 the streamlets to the viewer 708. In a further embodiment, delivering 914 streamlets 304 to the end user comprises playing video and or audio streamlets on the viewer 708. If the streamlets 304 are not

16

stored 908 locally, the streamlet request is passed to the network controller module 706. The network controller module 706 then requests 916 the streamlet 304 from the web server 116. Once the streamlet 304 is received, the network controller module 706 passes the streamlet to the streamlet cache module 704. The streamlet cache module 704 archives 918 the streamlet. Alternatively, the streamlet cache module 704 then archives 918 the streamlet and passes the streamlet to the agent controller module 702, and the method 900 then continues from operation 910 as described above.

Referring now to FIG. 10, shown therein is a schematic flow chart diagram illustrating one embodiment of a method 1000 for requesting streamlets 304 within a adaptive-rate shifting content streaming environment in accordance with the present invention. The method 1000 may be used in one embodiment as the operation 911 of FIG. 9. The method 1000 starts and the agent controller module 702 receives 1004 a streamlet 304 as described above with reference to FIG. 9. The agent controller module 702 then monitors 1006 the receive time of the requested streamlet. In one embodiment, the agent controller module 702 monitors the time intervals A between successive receive times for each streamlet response. Ordering of the responses in relation to the order of their corresponding requests is not relevant.

Because network behavioral characteristics fluctuate, sometimes quite suddenly, any given Δ may vary substantially from another. In order to compensate for this fluctuation, the agent controller module 702 calculates 1008 a performance ratio r across a window of n samples for streamlets of playback length S . In one embodiment, the performance ratio r is calculated using the equation

$$r = S \frac{n}{\sum_{i=1}^n \Delta_i}.$$

Due to multiple simultaneous streamlet processing, and in order to better judge the central tendency of the performance ratio r , the agent controller module 702 may calculate a geometric mean, or alternatively an equivalent averaging algorithm, across a window of size m , and obtain a per

$$\varphi_{current} = \left(\prod_{j=1}^m r_j \right)^{\frac{1}{m}}$$

The policy determination about whether or not to upshift 1010 playback quality begins by comparing $\varphi_{current}$ with a trigger threshold Θ_{up} . If $\varphi_{current} \geq \Theta_{up}$, then an up shift to the next higher quality stream may be considered 1016. In one embodiment, the trigger threshold Θ_{up} is determined by a combination of factors relating to the current read ahead margin (i.e. the amount of contiguously available streamlets that have been sequentially arranged by the staging module 709 for presentation at the current playback time index), and a minimum safety margin. In one embodiment, the minimum safety margin may be 24 seconds. The smaller the read ahead margin, the larger Θ_{up} is to discourage upshifting until a larger read ahead margin may be established to withstand network disruptions. If the agent controller module 702 is able to sustain 1016 upshift quality, then the agent controller module 702 will upshift 1017 the quality and subsequently request higher quality streams. The determination of whether use of the higher quality stream is sustainable 1016 is made by

17

comparing an estimate of the higher quality stream's performance factor, ϕ_{higher} , with Θ_{up} . If $\phi_{higher} \geq \Theta_{up}$, then use of the higher quality stream is considered sustainable. If the decision of whether or not the higher stream rate is sustainable **1016** is "no," the agent controller module **702** will not attempt to upshift **1017** stream quality. If the end of the stream has been reached **1014**, the method **1000** ends **1016**.

If the decision on whether or not to attempt upshift **1010** is "no", a decision about whether or not to downshift **1012** is made. In one embodiment, a trigger threshold Θ_{down} is defined in a manner analogous to Θ_{up} . If $\phi_{current} > \Theta_{down}$, then the stream quality may be adequate, and the agent controller module **702** does not downshift **1018** stream quality. However, if $\phi_{current} \leq \Theta_{down}$, the agent controller module **702** does downshift **1018** the stream quality. If the end of the stream has not been reached **1014**, the agent controller module **702** begins to request and receive **1004** lower quality streamlets and the method **1000** starts again. Of course, the above described equations and algorithms are illustrative only, and may be replaced by alternative streamlet monitoring solutions.

The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. An apparatus for multi-bitrate video and/or audio content streaming, the apparatus comprising:

one or more computing devices including,

a receiving module configured to receive videos;

a streamlet module configured to segment the media content received videos to generate for each a plurality of sequential raw streamlets that collectively store data to playback the entire video and that individually store data to playback only a portion that starts at a unique time index and whose duration is less than the entire duration of the corresponding video;

an encoding module configured to encode each raw streamlet to generate, for each of said raw streamlets, a set including an encoded streamlet for each bitrate supported by the multi-bitrate content streaming, wherein each of the encoded streamlets in each of the sets is stored as a separate content file, wherein the encoded streamlets within each of the sets have the same time index as their corresponding raw streamlet such that the encoded streamlets of the same set independently yield on playback the same portion of the video, wherein the separate content files within each of the sets are independently requestable by end user stations, and wherein shifts between the different bit rates are made at the time indexes during streaming of each of the videos; and

a web server to be coupled over the Internet to the end user stations to receive requests from the end user stations for different ones of the separate content files from different ones of the sets and to transmit the requested ones of the separate content files to the requesting one of the end user stations, wherein each of the end user stations initiate each of the shifts between the different bit rates during streaming of each of the videos through a request for the separate

18

content file storing a different bit rate one of the encoded streamlets for a subsequent one of the time indexes.

2. The apparatus of claim 1, wherein the encoding module is configured to encode multiple of the plurality of sequential raw streamlets in parallel.

3. The apparatus of claim 2, wherein the encoding module is further configured to encode the multiple streamlets in a time-selective manner so that the plurality of sequential raw streamlets are processed successively at a regular time interval.

4. The apparatus of claim 3, wherein the one or more computing devices further comprise a streamlet database, wherein the encoding module is further configured to transmit the separate content files each storing one of the encoded streamlets to the streamlet database at the regular time interval to be stored, and wherein the web server is coupled to the streamlet database to retrieve the requested ones of the separate content files from the streamlet database to transmit to the requesting one of the end user stations.

5. The apparatus of claim 1, wherein the portions of each of the plurality of sequential raw streamlets has a predetermined length of time in the range of between about 0.1 and 5 seconds.

6. The apparatus of claim 1, wherein the portions of each of the plurality of sequential raw streamlets has a predetermined length of time of the video in the range of between about 1 and 3 seconds.

7. The apparatus of claim 1, wherein the portions of each of the plurality of sequential raw streamlets has a predetermined length of time in the range of between about 1.5 and 2.5 seconds.

8. The apparatus of claim 1, wherein the receiving module is configured to receive the videos including encoded video and encoded audio and convert the encoded video to raw video and the encoded audio to raw audio.

9. The apparatus of claim 1, wherein the encoding module comprises a master module configured to assign an encoding job to one of a plurality of hosts to generate at least one of the separate content files from at least one of the plurality of sequential raw streamlets in response to an encoding job completion bid received from the one host.

10. The apparatus of claim 9, wherein each said host is configured to submit an encoding job completion bid based on a plurality of computing variables.

11. The apparatus of claim 10, wherein the plurality of computing variables is selected from a group consisting of current encoding job completion percentage, average encoding job completion time, processor speed, and physical memory capacity.

12. The apparatus of claim 1, wherein the encoding module is configured to encode the plurality of sequential raw streamlets with an offset parallel processing scheme such that the sets are published at regular time intervals to match the real-time submission rate at which each of the plurality of sequential raw streamlets are received at the encoding module, the time intervals being less than the time of encoding one of the sets and sufficient to play the video according to the real-time submission rate on at least one of the end user stations.

13. The apparatus of claim 1, wherein the encoding module is configured to encode the plurality of sequential raw streamlets with an offset parallel processing scheme such that the sets are published at regular time intervals to match the real-time submission rate at which each of the plurality of sequential raw streamlets are received at the encoding module, the time intervals being greater than the time of encoding one of

19

sets and sufficient to play the video according to the real-time submission rate on at least one of the end user stations.

14. A system for multi-bitrate video and/or audio content streaming, the system comprising:

end user stations coupled to the Internet; and

one or more computing devices including,

a receiving module configured to receive videos;

a streamlet module configured to segment the received videos to generate for each a plurality of sequential raw streamlets that collectively store data to playback the entire video and that individually store data to playback only a portion that starts at a unique time index and whose duration is less than the entire duration of the corresponding video; content having a predetermined length of time;

an encoding module configured to encode each raw streamlet to generate, for each of said raw streamlets, a set including an encoded streamlet for each bitrate supported by the multi-bitrate content streaming, wherein each of the encoded streamlets in each of the sets is stored as a separate content file that individually stores data to playback only the same video portion as the raw streamlet from which it was generated; wherein the separate content files within each of the sets are independently requestable by the end user stations; and

a web server to be coupled over the Internet to the end user stations to receive requests from the end user stations for different ones of the separate content files from different ones of the sets and to transmit the requested ones of the separate content files to the requesting one of the end user stations, wherein shifts between the different bit rates are made at the time indexes during streaming of each of the videos wherein each of the end user stations initiate each of the shifts between the different bit rates during streaming of each of the videos through a request for the separate content file storing a different bit rate one of the encoded streamlets for a subsequent one of the time indexes.

15. The system of claim **14**, wherein the encoding module is configured to encode multiple of the plurality of sequential raw streamlets in parallel.

16. The system of claim **15**, wherein the encoding module is further configured to encode the multiple streamlets in a time-selective manner so that the plurality of sequential raw streamlets are processed successively at a regular time interval.

17. The system of claim **16**, wherein:

the encoding module comprises a master module configured to assign an encoding job to one of a plurality of host computing modules to generate at least one of the separate content files from at least one of the plurality of sequential raw streamlets in response to an encoding job completion bid received from the one host, and

the one or more computing devices further comprise a streamlet database, and wherein the encoding module is further configured to transmit the separate content files each storing one of the encoded streamlets to the streamlet database at the regular time interval to be stored, and wherein the web server is coupled to the streamlet database to retrieve the requested ones of the separate content files from the streamlet database to transmit to the requesting one of the end user stations.

18. The system of claim **14**, wherein the predetermined length of time is in the range of between about 0.1 and 5 seconds.

20

19. The system of claim **14**, wherein the predetermined length of time is in the range of between about 1 and 3 seconds.

20. The system of claim **14**, wherein the predetermined length of time is in the range of between about 1.5 and 2.5 seconds.

21. The system of claim **14**, wherein the receiving module is configured to receive the videos including encoded video and encoded audio and convert the encoded video to raw video and the encoded audio to raw audio.

22. The system of claim **17**, wherein each said host computing module is configured to submit an encoding job completion bid based on a plurality of computing variables.

23. The system of claim **22**, wherein the plurality of computing variables is selected from a group consisting of current encoding job completion percentage, average encoding job completion time, processor speed, and physical memory capacity.

24. A method for multi-bitrate video and/or audio content streaming, the method comprising:

capturing a single video;

segmenting the single video and generating a plurality of sequential raw streamlets that collectively store data to playback the entire video and that individually store data to playback only a portion that starts at a unique time index and whose duration is less than the entire duration of the corresponding video;

encoding each raw streamlet to generate, for each of said raw streamlets, a set including an encoded streamlet for each bitrate supported by the multi-bitrate content streaming, wherein each of the encoded streamlets in each of the sets is stored as a separate content file, wherein the encoded streamlets within each of the sets have the same time index as their corresponding raw streamlet such that the encoded streamlets of the same set independently yield on playback the same portion of the single video, wherein the separate content files within each of the sets are independently requestable by end user stations, and wherein shifts between the different bit rates are made at the time indexes during streaming of the single video;

receiving requests from the end user stations over the Internet for different ones of the separate content files from different ones of the sets; and

transmitting the requested ones of the separate content files to the requesting one of the end user stations, wherein each of the end user stations initiate each of the shifts between the different bit rates during streaming of the single video through a request for the separate content file storing a different bit rate one of the encoded streamlets for a subsequent one of the time indexes.

25. The method of claim **24**, wherein said encoding comprises encoding multiple of the plurality of sequential raw streamlets in parallel.

26. The method of claim **25**, wherein said encoding the multiple streamlets comprises encoding the multiple streamlets in a time-selective manner so that the plurality of sequential raw streamlets are processed successively at a regular time interval.

27. The method of claim **25**, further comprising:

transmitting the separate content files each storing one of the encoded streamlets to a streamlet database at the regular time interval; and

storing the separate content files each storing one of the encoded streamlets in the streamlet database, wherein

21

said transmitting comprises retrieving the requested ones of the separate content files from the streamlet database.

28. The method of claim 24, wherein said receiving comprises receiving the single video including encoded video and encoded audio, and the method further comprises converting the encoded video to raw video and the encoded audio to raw audio.

29. The method of claim 24, wherein said encoding comprises assigning an encoding job to one of a plurality of host computing modules to generate at least one of the separate content files from at least one of the plurality of sequential raw streamlets in response to an encoding job completion bid received from the one host.

30. The method of claim 29, wherein said encoding comprises receiving the encoding job completion bid based on a plurality of computing variables.

31. A non-transitory computer readable storage medium storing a computer readable program of computer instructions, wherein the computer readable program when executed on a computer causes the computer to carry out operations for multi-bitrate video and/or audio content streaming, the operations comprising:

capturing a single video;

segmenting the single video and generating a plurality of sequential raw streamlets that collectively store data to playback the entire video and that individually store data to playback only a portion that starts at a unique time index and whose duration is less than the entire duration of the corresponding video;

encoding each raw streamlet to generate, for each of said raw streamlets, a set including an encoded streamlet for each bitrate supported by the multi-bitrate content streaming, wherein each of the encoded streamlets in each of the sets is stored as a separate content file, wherein the encoded streamlets within each of the sets have the same time index as their corresponding raw streamlet such that the encoded streamlets of the same set independently yield on playback the same portion of the single video, wherein the separate content files within each of the sets are independently requestable by end user stations, and wherein shifts between the different bit rates are made at the time indexes during streaming of the single video;

receiving requests from the end user stations over the Internet for different ones of the separate content files from different ones of the sets; and

transmitting the requested ones of the separate content files to the requesting one of the end user stations, wherein each of the end user stations initiate each of the shifts between the different bit rates during streaming of the single video through a request for the separate content file storing a different bit rate one of the encoded streamlets for a subsequent one of the time indexes.

32. The computer readable storage medium of claim 31, wherein said encoding comprises encoding multiple of the plurality of sequential raw streamlets in parallel.

33. The computer readable storage medium of claim 32, wherein said encoding the multiple streamlets comprises encoding the multiple streamlets in a time-selective manner so that the plurality of sequential raw streamlets are processed successively at a regular time interval.

34. The computer readable storage medium of claim 31, wherein the operations further comprise:

transmitting the separate content files each storing one of the encoded streamlets to a streamlet database at the regular time interval; and

22

storing the separate content files each storing one of the encoded streamlets in the streamlet database, wherein said transmitting comprises retrieving the requested ones of the separate content files from the streamlet database.

35. An apparatus for multi-bitrate video and/or audio content streaming, the apparatus comprising one or more computing devices including,

means for capturing videos;

means for segmenting the single video and generating a plurality of sequential raw streamlets that collectively store data to playback the entire video and that individually store data to playback only a portion that starts at a unique time index and whose duration is less than the entire duration of the corresponding video; and

means for encoding each raw streamlet to generate, for each of said raw streamlets, a set including an encoded streamlet for each bitrate supported by the multi-bitrate content streaming, wherein each of the encoded streamlets in each of the sets is stored as a separate content file, wherein the encoded streamlets within each of the sets have the same time index as their corresponding raw streamlet such that the encoded streamlets of the same set independently yield on playback the same portion of the single video, wherein the separate content files within each of the sets are independently requestable by end user stations, and wherein shifts between the different bit rates are made at the time indexes during streaming of the single video;

means for receiving requests from the end user stations over the Internet for different ones of the separate content files from different ones of the sets; and

means for transmitting the requested ones of the separate content files to the requesting one of the end user stations, wherein each of the end user stations initiate each of the shifts between the different bit rates during streaming of each of the videos through a request for the separate content file storing a different bit rate one of the encoded streamlets for a subsequent one of the time indexes.

36. The apparatus of claim 1, wherein the encoding module is further configured to start the generation of a second one of the sets from a second of the plurality of sequential raw streamlets prior to completing the generation of a first of the sets from a first of the plurality of sequential raw streamlets to generate the separate content files of the first and second sets in parallel, and wherein the second sequential raw streamlet sequentially follows the first sequential raw streamlet.

37. The apparatus of claim 1, wherein the encoding module is further configured to multi-pass encode each of the plurality of sequential raw streamlets of at least one of the videos, and wherein the at least one of the videos is of a live event.

38. The apparatus of claim 1, wherein least one of the videos is of a live event, wherein the encoding module is configured to receive the plurality of sequential raw streamlets of the at least one video according to a real-time submission rate, and wherein the encoding module is configured to encode the plurality of sequential raw streamlets at regular intervals using an offset parallel processing scheme, wherein the regular intervals match the intervals at which each of the plurality of sequential raw streamlets is received at the encoding module so the output timing of the encoding module does not fall behind the real-time submission rate.

39. The apparatus of claim 1, wherein the one or more computing devices further comprising a streamlet database,

23

and wherein the encoding module is configured to transmit the separate content files each storing one of the encoded streamlets to the streamlet database to be stored, and wherein the web server is coupled to the streamlet database to retrieve the requested ones of the separate content files from the streamlet database to transmit to the requesting one of the end user stations.

40. The apparatus of claim 9, wherein at least one of the plurality of hosts is configured to generate the separate content files within one of the sets.

41. The apparatus of claim 9, wherein one of the plurality of host is configured to generate a first of the separate content files according to a first one of the supported bitrates for one of the plurality of sequential raw streamlets, and wherein another one of the plurality of hosts is configured to generate a second of the separate content files according to a second one of the supported bitrates for the same one of the plurality of sequential raw streamlets.

42. The system of claim 14, wherein the encoding module is further configured to start the generation of a second one of the sets from a second of the plurality of sequential raw streamlets prior to completing the generation of a first of the sets from a first of the plurality of sequential raw streamlets to generate the separate content files of the first and second sets in parallel, and wherein the second sequential raw streamlet sequentially follows the first sequential raw streamlet.

43. The system of claim 14, wherein the encoding module is further configured to multi-pass encode each of the plurality of sequential raw streamlets of at least one of the videos, and wherein the at least one of the videos is of a live event.

44. The system of claim 14, wherein at least one of the videos is of a live event, wherein the encoding module is configured to receive the plurality of sequential raw streamlets of the at least one video according to a real-time submission rate, and wherein the encoding module is configured to encode the plurality of sequential raw streamlets at regular intervals using an offset parallel processing scheme, wherein the regular intervals match the intervals at which each of the plurality of sequential raw streamlets is received at the encoding module so the output timing of the encoding module does not fall behind the real-time submission rate.

45. The system of claim 14, comprising wherein:

the encoding module comprises a master module configured to assign an encoding job to one of a plurality of host computing modules to generate at least one of the separate content files from at least one of the plurality of sequential raw streamlets in response to an encoding job completion bid received from the one host, and

the one or more computing devices further comprise a streamlet database, and wherein the encoding module is further configured to transmit the separate content files each storing one of the encoded streamlets to the streamlet database to be stored, and wherein the web server is coupled to the streamlet database to retrieve the requested ones of the separate content files from the streamlet database to transmit to the requesting one of the end user stations.

46. The system of claim 45, wherein at least one of the plurality of hosts is configured to generate the separate content files within one of the sets.

47. The system of claim 45, wherein one of the plurality of host is configured to generate a first of the separate content files according to a first one of the supported bitrates for one of the plurality of sequential raw streamlets, and wherein another one of the plurality of hosts is configured to generate

24

a second of the separate content files according to a second one of the supported bitrates for the same one of the plurality of sequential raw streamlets.

48. The method of claim 24, wherein said encoding further comprises starting the generation of a second one of the sets from a second of the plurality of sequential raw streamlets prior to completing the generation of a first of the sets from a first of the plurality of sequential raw streamlets to generate the separate content files of the first and second sets in parallel, and wherein the second sequential raw streamlet sequentially follows the first sequential raw streamlet.

49. The method of claim 24, wherein said encoding further comprises multi-pass encoding each of the plurality of sequential raw streamlets of the video, and wherein the video is of a live event.

50. The method of claim 24, wherein the video is of a live event, wherein said capturing comprises capturing the video of the live event according to a real-time submission rate, and wherein said encoding comprises encoding each of the plurality of sequential raw streamlets of the video at regular intervals using an offset parallel processing scheme, wherein the regular intervals match the intervals at which each of the plurality of sequential raw streamlets is generated so the output timing of said encoding does not fall behind the real-time submission rate.

51. The method of claim 24, further comprising:

transmitting the separate content files each storing one of the encoded streamlets to a streamlet database; and storing the separate content files each storing one of the encoded streamlets in the streamlet database, wherein said transmitting comprises retrieving the requested ones of the separate content files from the streamlet database.

52. The method of claim 29, wherein said assigning the encoding job comprises assigning at least one of the plurality of hosts to generate the separate content files within one of the sets.

53. The method of claim 29, wherein said assigning the encoding job comprises assigning one of the plurality of host to generate a first of the separate content files according to a first one of the supported bitrates for one of the plurality of sequential raw streamlets, and wherein the method further comprises assigning a second encoding job to another one of the plurality of hosts to generate a second of the separate content files according to a second one of the supported bitrates for the same one of the plurality of sequential raw streamlets.

54. The apparatus of claim 35, wherein at least one of the videos is of a live event, wherein said means for encoding receives the at least one video of the live event according to a real-time submission rate, and wherein said means for encoding encodes each of the plurality of sequential raw streamlets of the at least one video at regular intervals using an offset parallel processing scheme, wherein the regular intervals match the intervals at which each of the plurality of sequential raw streamlets is received at said encoding means so the output timing of said encoding means does not fall behind the real-time submission rate.

55. The apparatus of claim 35, wherein the means for encoding further comprises means for multi-pass encoding each of the plurality of sequential raw streamlets of at least one of the videos, and wherein the at least one video is of a live event.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION


PATENT NO. : 7,818,444 B2
APPLICATION NO. : 11/673483
DATED : October 19, 2010
INVENTOR(S) : Brueck et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 22, line 55 claim 38, add the word --at-- after the word wherein and before the word least.

Signed and Sealed this
Twenty-sixth Day of April, 2011

A handwritten signature in black ink, reading "David J. Kappos". The signature is written in a cursive, flowing style with a large initial "D" and a stylized "K".

David J. Kappos
Director of the United States Patent and Trademark Office