



(22) Date de dépôt/Filing Date: 2005/12/09

(41) Mise à la disp. pub./Open to Public Insp.: 2006/06/10

(45) Date de délivrance/Issue Date: 2014/08/19

(30) Priorité/Priority: 2004/12/10 (US11/010,200)

(51) Cl.Int./Int.Cl. *H04L 12/16* (2006.01),
H04L 1/00 (2006.01), *H04L 12/18* (2006.01)


(72) Inventeurs/Inventors:
CLIFFORD, DAVID C., US;
GREEN, DUSTIN L., US;
SMITH, GEOFFREY R., US;
MOHR, GRANT D., US;
BALDWIN, JAMES A., US;
DODD, MICHAEL D., US;
BARRETT, PETER T., US

(73) Propriétaire/Owner:
MICROSOFT CORPORATION, US

(74) Agent: SMART & BIGGAR

(54) Titre : CHANGEMENT DE CANAL ACCELERE DANS DES ENVIRONNEMENTS A DEBIT LIMITE

(54) Title: ACCELERATED CHANNEL CHANGE IN RATE-LIMITED ENVIRONMENTS



(57) Abrégé/Abstract:

Accelerated channel change (ACC) strategies are described for supplying a client module with media information by transitioning between a first delivery of media information provided by first delivery functionality and a second delivery of media information provided by second delivery functionality without exceeding prescribed data rate limitations. The first delivery functionality can represent server-based functionality for delivering a unicast stream of media information at a data rate (or maximum data rate) of some excess amount over the nominal rate of delivery, starting at a selected location within the media information. The second delivery functionality can represent any functionality for delivering the media information at the data rate (or maximum data rate) of the nominal rate of delivery. A join interval separates the first delivery at the excess rate and the second delivery at the nominal rate. In that period, the first delivery functionality can provide media information at a join interval data rate. The join interval data rate, if it is below the nominal data rate, will cause the client module to miss packets of media information. The client module can receive these missing packets by issuing retry requests to the first delivery functionality or to some other entity. The client module can determine the excess amount available by investigating the amount of missing information that occurs at different data rates; or the excess amount can be defined in advance.

ABSTRACT

Accelerated channel change (ACC) strategies are described for supplying a client module with media information by transitioning between a first delivery of media information provided by first delivery functionality and a second delivery of media information provided by second delivery functionality without exceeding prescribed data rate limitations. The first delivery functionality can represent server-based functionality for delivering a unicast stream of media information at a data rate (or maximum data rate) of some excess amount over the nominal rate of delivery, starting at a selected location within the media information. The second delivery functionality can represent any functionality for delivering the media information at the data rate (or maximum data rate) of the nominal rate of delivery. A join interval separates the first delivery at the excess rate and the second delivery at the nominal rate. In that period, the first delivery functionality can provide media information at a join interval data rate. The join interval data rate, if it is below the nominal data rate, will cause the client module to miss packets of media information. The client module can receive these missing packets by issuing retry requests to the first delivery functionality or to some other entity. The client module can determine the excess amount available by investigating the amount of missing information that occurs at different data rates; or the excess amount can be defined in advance.

ACCELERATED CHANNEL CHANGE IN RATE-LIMITED ENVIRONMENTS

TECHNICAL FIELD

This subject matter relates to strategies for streaming media information to client modules in a manner which allows the client modules to quickly begin presenting the media information.

BACKGROUND

Traditional systems for delivering broadcast media information rely on the use of a physical tuner in a client module (e.g., a television set) that tunes to and receives the media information. In this traditional model, a user can quickly change channels, resulting in the virtual instantaneous transition from one program to another. And as such, the user does not typically perceive a delay in the presentation of a new program upon tuning to a new program. However, this simple manner of operation does not apply to the delivery of digital media information in streaming fashion over a network. In this environment, the client module typically must store a prescribed amount of media information in a buffer before it begins to play the media information to the user. It requires a certain amount of time to fill up this buffer when the user first connects to a stream of media information. Further, digital media information is commonly expressed as a series of key frames (e.g., I frames) and difference frames (e.g., B and P frames). A client module must wait for a key frame before it begins to present the media information. As a result of these factors, there will be a noticeable lag prior to the presentation of programs as the user switches from one channel to the next.

Microsoft Corporation of Redmond, Washington, has developed several strategies that reduce the effects of the above-noted problem in media streaming environments. These strategies are generally referred to as Accelerated Channel Change (ACC) paradigms herein, since they permit the client module to display a new program in an

1 expedited manner when a user “tunes” to the new program (or when the user otherwise
2 commences presentation of the new program). In these strategies, the head-end
3 infrastructure delivers an initial burst which feeds a large quantity of media information
4 to the client module via a communication channel, enabling it to begin presentation of the
5 media information on an expedited basis. After the initial burst, the head-end
6 infrastructure delivers the media information to the client module at a regular data rate
7 (e.g., at or below a specified nominal data rate).

8 Generally, the ACC strategies make additional demands on the communication
9 channels beyond that required to transmit the media information at its normal steady-state
10 rate. A system can address these demands by allocating extra bandwidth to implement
11 the ACC strategies. However, many environments allocate a limited amount of
12 bandwidth to the communication channels. The application of ACC strategies in these
13 kinds of environments may present various challenges. For example, there is a potential
14 that the ACC strategies can overburden the communication channels, leading to their
15 over-saturation. Such over-saturation of the channels can result in a poor quality of
16 media presentation (e.g., due to the loss of an unacceptable number of media packets).
17 Such over-saturation can also possibly result in the streams being dropped. For example,
18 consider a household that includes various media presentation units that all “feed off” of
19 a single communication channel. The ACC strategies have the potential of
20 overburdening such a communication channel, particularly when multiple users in the
21 household happen to be consuming media information at the same time (and also
22 potentially invoking the ACC behavior at the time).

23 Other rate-limited environments may allocate a relatively large amount of
24 bandwidth to the channels, thus reducing the potential of over-saturating the channels. It
25 is nevertheless beneficial to adapt the ACC strategies in these environments so as to make

51018-132

1 efficient use of the allotted communication resources. In any environment, it is also
2 beneficial to find ways to streamline the ACC protocol to make it less costly, less error
3 prone, and more efficient.

4 For at least the above-identified reasons, there is a need for an improved paradigm
5 for performing accelerated channel change (ACC) in rate-limited environments of all
6 kinds.

7 8 SUMMARY

9 A method for delivering resource information to a client module is described
10 herein. The method comprises: (a) delivering the resource information, using a first
11 delivery technique, at above-nominal data rate levels during an initial burst period of data
12 transmission to the client module; (b) delivering the resource information at join interval
13 data rate levels during a join interval following the burst period, wherein, following the
14 join interval, the client module switches to a second delivery technique for providing the
15 resource information; and (c) responding to requests from the client module to supply
16 parts of the resource information which the client module potentially missed during the
17 join interval as a result of its receipt of the resource information at the join interval data
18 rate levels.

51018-132

According to one aspect of the present invention, there is provided a method for expediting the presentation of resource information upon a start of the delivery of the resource information within a rate-limited environment, comprising: delivering the resource information to a client module, in a first delivery, using first
5 delivery functionality upon a commencement event; and delivering the resource information to the client module, in a second delivery, using second delivery functionality, wherein the resource information provided by the first delivery and the resource information provided by the second delivery combine together to provide a resultant stream of resource information in which the primary mode of
10 resource information delivery transitions from the first delivery functionality to the second delivery functionality, and wherein the resultant stream has a data rate that remains below a prescribed data rate limit.

According to another aspect of the present invention, there is provided a method for delivering resource information to a client module,
15 comprising: delivering the resource information, using first delivery functionality, at above-nominal data rate levels during an initial burst period of data transmission to the client module; delivering the resource information at join interval data rate levels during a join interval following the burst period, wherein, following the join interval, the client module switches to second delivery functionality for providing
20 the resource information; and responding to requests from the client module to supply parts of the resource information which the client module missed during the join interval as a result of its receipt of the resource information at the join interval data rate levels.

According to still another aspect of the present invention, there is
25 provided a method for receiving resource information at a client module, comprising: receiving, via first delivery functionality, resource information at above-nominal data rate levels during an initial burst period of data transmission; receiving the resource information at join interval data rate levels during a join interval following the burst period; receiving, via second delivery functionality, the
30 resource information following the join interval; requesting parts of the resource

51018-132

information which the client module missed during the join interval as a result of its receipt of the resource information at the join interval data rate levels; and receiving the missing parts of resource information concurrently with the receipt of the resource information via the second delivery functionality.

5 According to yet another aspect of the present invention, there is provided a method for coordinating the delivery of resource information using two delivery functionalities so that an aggregate amount of the resource information delivered on a communication channel does not exceed a prescribed data rate limit, comprising: providing a value which represents an excess amount of
10 bandwidth available to the communication channel; providing a value which represents an amount of time T_{burst} during which first delivery functionality delivers the resource information over the communication channel at a burst data rate levels defined by the excess amount of bandwidth; and providing a value which represents an amount of time T_{join} which separates an end of the delivery of the
15 resource information at the burst data rate levels and a start of a delivery of the resource information using second delivery functionality at data rate levels at or below a nominal data rate, wherein the values of the excess capacity, T_{burst} , and T_{join} are selected such that the ratio of T_{burst} to T_{join} is related in a prescribed manner to a prescribed function of excess capacity.

20 According to a further aspect of the present invention, there is provided one or more computer-readable storage media having stored thereon computer executable instructions, that when executed by one or more processors implements a method as described above or detailed below.

 According to yet a further aspect of the present invention, there is
25 provided a method for expediting presentation of resource information upon a start of delivery of the resource information within a rate-limited environment, comprising: delivering the resource information to a client module, in a first delivery, using first delivery functionality upon a commencement event; and delivering the resource information to the client module, in a second delivery,
30 using second delivery functionality, wherein the resource information provided by

51018-132

the first delivery and the resource information provided by the second delivery combine together to provide a resultant stream of resource information in which a primary mode of resource information delivery transitions from the first delivery functionality to the second delivery functionality, wherein the resultant stream has a data rate that remains below a
5 prescribed data rate limit, and wherein retry functionality receives retry requests from the client module to supply missing parts of the resource information and delivers the missing parts of the resource information to the client module.

According to still a further aspect of the present invention, there is provided a method for delivering resource information to a client module, comprising: delivering the
10 resource information, using first delivery functionality, at above-nominal data rate levels during an initial burst period of data transmission to the client module; delivering the resource information at join interval data rate levels during a join interval following the burst period, wherein, following the join interval, the client module switches to second delivery functionality for providing the resource information; and responding to requests from the
15 client module to supply parts of the resource information which the client module missed during the join interval as a result of its receipt of the resource information at the join interval data rate levels.

According to another aspect of the present invention, there is provided a method for receiving a stream of resource information at a client module, comprising:
20 receiving, via first delivery functionality, resource information at above-nominal data rate levels during an initial burst period of data transmission; receiving the resource information at join interval data rate levels during a join interval following the burst period; receiving, via second delivery functionality, the resource information following the join interval; requesting parts of the resource information which the client module missed during the join interval as a
25 result of its receipt of the resource information at the join interval data rate levels; receiving the parts of resource information which the client module missed concurrently with the receipt of the resource information via the second delivery functionality; and using the received missing parts for filling in holes created in the stream during the join interval.

51018-132

According to yet another aspect of the present invention, there is provided a method for delivering resource information to a client module, comprising: delivering the resource information, using first delivery functionality, at above nominal data rate levels during an initial burst period of data transmission to the client module; delivering the resource
5 information at join interval data rate levels during a join interval following the initial burst period, wherein, following the join interval, the client module switches to second delivery functionality for providing the resource information; and sending, using the first delivery functionality, retry data comprising a part of the resource information which the client module missed.

10 According to another aspect of the present invention, there is provided a method for receiving a stream of resource information at a client module, comprising: receiving, via first delivery functionality, resource information at above-nominal data rate levels during an initial burst period of data transmission; receiving the resource information at
15 join interval data rate levels during a join interval following the burst period; receiving, via second delivery functionality, the resource information following the join interval; requesting parts of the resource information which the client module missed; receiving the missing parts of resource information concurrently with the receipt of the resource information; and using the received missing parts for filling in holes created in the stream.

According to still another aspect of the present invention, there is provided a
20 method for delivering resource information to a client module, comprising: delivering the resource information, using first delivery functionality, at above-nominal data rate levels during an initial burst period of data transmission to the client module; delivering the resource information at join interval data rate levels during a join interval following the initial burst period, wherein, following the join interval, the client module switches to second delivery
25 functionality for providing the resource information; and delivering again, contemporaneously with delivery of the resource information by the second delivery functionality, at least a part of the resource information, delivered using the first delivery functionality, which the client module missed.

51018-132

Additional exemplary implementations are described in the following.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows an exemplary system which can employ an improved rate-limited ACC paradigm.

5 Fig. 2 shows an exemplary server module and client module used in the system of Fig. 1.

1 Fig. 3 shows an exemplary sequence of information flowing between the server
2 module and the client module of Fig. 3 in the course of performing the rate-limited ACC
3 paradigm.

4 Fig. 4 shows a representation of a communication channel that couples the server
5 module and the client module of Fig. 3, which illustrates the amount of bandwidth
6 consumed by the rate-limited ACC strategy as a function of time.

7 Fig. 5 shows a stream of media information presented at three different junctures
8 of the communication channel representation of Fig. 6.

9 Figs. 6-8 present different data-versus-time graphs which illustrate the exemplary
10 operation of the rate-limited ACC paradigm in three different modes of operation
11 governed by different retry assumptions.

12 Fig. 9 shows an exemplary procedure which explains the operation of the server
13 module in performing the rate-limited ACC paradigm.

14 Fig. 10 shows an exemplary procedure, which complements the procedure of Fig.
15 9, which explains the operation of the client module in performing the rate-limited ACC
16 paradigm.

17 Fig. 11 shows a variation of the representation of Fig. 4, in which the
18 communication channel allocates additional excess capacity compared to the case of Fig.
19 4.

20 Fig. 12 shows a representation of a communication channel that allocates a
21 variable amount of additional capacity.

22 Fig. 13 shows a system which employs a unicast accelerated channel change
23 (ACC) paradigm to deliver video on demand (VOD) assets in a rate-limited environment.

24 Fig. 14 shows an exemplary computer environment for implementing aspects of
25 the system shown in Fig. 1.

1 The same numbers are used throughout the disclosure and figures to reference like
2 components and features. Series 100 numbers refer to features originally found in Fig. 1,
3 series 200 numbers refer to features originally found in Fig. 2, series 300 numbers refer
4 to features originally found in Fig. 3, and so on.

6 **DETAILED DESCRIPTION**

7 The following description sets forth a strategy for allowing a client module to
8 transition from a first mode of delivering resource information (delivered by first
9 resource delivery functionality) to a second mode of delivering resource information
10 (delivered by second resource delivery functionality) in such a manner that an amount of
11 resource information received by the client module per time interval does not exceed a
12 prescribed maximum bandwidth. In one exemplary implementation, the first mode of
13 delivery uses a server module to provide a unicast stream of media information delivered
14 at burst data rate levels, and the second mode of delivery uses multicast functionality to
15 provide a multicast stream of the media information at nominal data rate levels. The
16 strategy described herein can be invoked when a user wishes to commence the
17 presentation of a stream of media information, such as when the user switches from one
18 first program to another program. In another exemplary implementation, the first mode
19 of delivery uses the server module to provide the unicast stream at the burst-data rate
20 levels, and the second model of delivery also uses the server module to provide the
21 unicast stream, but at the nominal data rate levels.

22 The strategy described herein confers a number of benefits. For instance, the
23 strategy is particularly advantageous in those environments characterized as “rate-limited
24 environments.” These environments allocate the client module (or group of client
25 modules) a prescribed amount of bandwidth to receive one or more streams of media

1 information. By maintaining the amount of consumed bandwidth below the prescribed
2 data rate level, the strategy accommodates the use of an accelerated channel change burst
3 without exceeding the capacity of the communication channel. This, in turn, reduces the
4 chances that one or more streams of resource information will be dropped (or otherwise
5 degraded) as a result of the accelerated channel change behavior.

6 Other features and attendant benefits of the strategy will be set forth in the
7 following disclosure. Generally, this disclosure sets forth the strategy by first explaining
8 the application of the strategy to an environment that delivers resource information at a
9 constant data rate to client modules using a hybrid unicast/multicast ACC technique.
10 Further, some examples assume that the amount of excess capacity provided by the
11 communication channel is less than a nominal data rate of the resource information (so
12 that two full nominal data rate streams cannot “fit” onto the communication channel at
13 the same time). As to the issue of bit rate, the constant data rate implementation provides
14 a convenient pedagogical vehicle for visualizing the functional interdependencies
15 involved in the strategy. For instance, assuming that a constant data rate is available for
16 delivery of resource information, the different stages in the transition from the first mode
17 of delivery to the second mode delivery can be depicted as rectangular blocks on a
18 bandwidth-versus-time graph. However, the strategy can also be applied to other types of
19 environments. For instance, the analysis also applies to worst-case assumptions for the
20 case of capped variable bit rate (VBR) delivery (where the resource information is
21 delivered at the maximum data rate permitted). Moreover, the strategy can be applied to
22 environments that allocate a relatively large amount of bandwidth (such as two times the
23 nominal data rate or greater). Data rate limitations may not present a significant design
24 challenge in these environments, but such environments can benefit from the efficiency at
25

1 which the ACC strategy transitions from one mode of delivery to another, as well as other
2 merits of the strategy.

3 As to the issue of excess capacity, the strategy can also be applied to
4 communication channels that allocate a relatively large amount of excess capacity,
5 including excess capacity equal to or greater than the nominal bit rate of the stream which
6 corresponds to a total capacity equal to or greater than two times the nominal bit rate of
7 the stream. As will be described, the use of large levels of excess capacity to deliver
8 resource information can reduce (or eliminate) the amount of missed packets caused by
9 the transition from the first mode of delivery to the second mode of delivery, and
10 therefore may reduce the need for retries following the transition from the first mode to
11 the second mode.

12 As to the issue of delivery modes, the strategy can be applied to other
13 environments which transition between other kinds of modes of delivery, such as in the
14 case of the pure unicast delivery model described above (where the different modes
15 correspond to two different delivery phases of the same unicast delivery mechanism).

16 As to terminology, the term "resource information" refers to any data represented
17 in electronic form that can be consumed by a user. The resource information may
18 include, or may omit, interactive content. A specific class of resource information may
19 pertain to media information. The media information can include any information that
20 conveys audio and/or video information, such as audio resources (e.g., music, etc.), still
21 picture resources (e.g., digital photographs, etc.), moving picture resources (e.g., audio-
22 visual television programs, movies, etc.), computer programs (e.g., games, etc.), and so
23 on.

24 The term "stream" refers to the representation of the resource information as time-
25 sequenced flow of the resource information. A client module can consume a stream of

1 resource information in piecemeal fashion as it is received, as opposed to requiring the
2 client module to receive the entire body of the resource information (e.g., as a single file)
3 before consuming the resource information.

4 The term “communication channel” broadly refers to the resources that a system
5 uses for handling communication between a particular source module and a particular
6 client module.

7 The term “unicast” refers to the delivery of resource information between a server
8 module and a client module in one-to-one fashion, where specific resources are allocated
9 to the unicast delivery for that particular client module. The term “multicast” refers to
10 the delivery of resource information to potentially multiple client modules in a manner
11 that does not require per-client allocation of communication resources (and particularly,
12 does not require per-client allocation of server resources). The term “multicast” may
13 encompass certain types of Internet Protocol (IP) multicasting technologies, such as the
14 Internet Group Management Protocol (IGMP), but is not limited to such specific
15 multicasting technologies. As very broadly used herein, the term “multicast” is generally
16 synonymous with the term “broadcast.”

17 The term “accelerated channel change (ACC)” refers to any paradigm for
18 allowing a client module to begin presenting a stream of resource information on an
19 expedited basis. The term “rate-limited ACC” refers to the novel strategy summarized
20 above, which permits a client module to switch between first delivery functionality and
21 second delivery functionality in a manner which does not exceed prescribed bandwidth
22 limitations. The term “rate-limited hybrid unicast/multicast ACC” refers to the specific
23 application of the rate-limited ACC strategy to the task of switching from a unicast
24 stream delivered by unicast delivery functionality to a multicast stream delivered by
25 multicast delivery functionality.

1 The term “nominal” refers to the maximum allowable steady-state bandwidth the
 2 stream is designed to fit within. For example, a stream encoder may create the resource
 3 information such that it can be continuously delivered and presented while utilizing no
 4 more than the nominal bit-rate for delivery, and without interruption of presentation of
 5 the resource information, assuming the client module presents resource information no
 6 earlier than the time specified by the stream.

7 The term “data rate levels” (as in, for example, “first data rate levels,” “second
 8 data rate levels,” etc.) can refer to a series of constant or varying data rates that conform
 9 to some pattern or criterion. For instance, the second mode of delivery (multicast)
 10 provides the resource information at “second data rate levels,” meaning generally herein
 11 that it delivers the resource information at levels conforming to a nominal data rate. In
 12 the case of a constant data rate environment, the second data rate levels might describe
 13 the “flat” delivery of the resource information at the nominal data rate itself. In variable
 14 data rate environments, the second data rate levels will be permitted to vary below a
 15 ceiling defined by the nominal data rate.

16 The term “maximum rate” or “*MaxRate*” refers to a total amount of bandwidth
 17 that a channel allocates to handling information exchange.

18 The term “excess data rate” or “excess capacity” refers to available bandwidth in
 19 a communication channel that can be used to handle the special provisions involved in
 20 the ACC strategy. For example, a channel might devote a maximum data rate of
 21 *MaxRate* to handling all information exchange over the communication channel. At a
 22 given time, the steady-state consumption of a multicast stream of resource information
 23 might be *ActualCurrentMulticastRate*. The excess data rate is then *MaxRate* –
 24 *ActualCurrentMulticastRate*. That excess data rate can be used for handling the special
 25 provisions of the ACC strategy, such as handling retries to account for the packets lost

1 when transitioning from the first mode of delivery to the second mode delivery. For
2 constant data rate delivery, the excess capacity can be represented as a constant amount.
3 For variable data rate delivery, the excess capacity will vary (because the amount of
4 bandwidth consumed by the steady-state delivery will vary). In this disclosure, the term
5 “excess data rate” or “excess capacity” refers to excess data rate in units of information
6 quantity per time unit. For steady-state delivery (and for capped VBR delivery when
7 delivering data at the nominal rate), the excess data rate can also be expressed as a
8 fraction E of the nominal data rate (also referred to as the “fractional excess capacity” or
9 “fractional excess data rate” herein).

10 A “primary mode of delivery” of resource information refers to a mode used to
11 deliver the bulk of the resource information being consumed by a client module at any
12 given time. For instance, the client module may be receiving a stream of resource
13 information from a multicast mode of delivery, while still relying on a unicast mode of
14 delivery to receive missing packets in response to retry requests. In this case, the
15 multicast mode of delivery would comprise the primary mode of delivery, since it
16 provides the bulk of the stream of resource information. The ACC strategy described
17 herein involves the transition from one primary mode of delivery to another.

18 This disclosure includes the following sections. Section A describes an
19 exemplary system for implementing the rate-limited ACC strategy. Section B describes
20 the manner of operation of the system of Section A. Section C describes variations to the
21 implementations described in Sections A and B. And Section D describes an exemplary
22 computer environment for implementing aspects of the systems of Section A.

23
24
25

A. Exemplary System

Generally, any of the functions described with reference to the figures can be implemented using software, hardware (e.g., fixed logic circuitry), manual processing, or a combination of these implementations. The term “logic,” “module” or “functionality” as used herein generally represents software, hardware, or a combination of software and hardware. For instance, in the case of a software implementation, the term “logic,” “module,” or “functionality” represents program code that performs specified tasks when executed on a processing device or devices (e.g., CPU or CPUs). The program code can be stored in one or more computer readable memory devices. More generally, the illustrated separation of logic, modules and functionality into distinct units may reflect an actual physical grouping and allocation of such software and/or hardware, or can correspond to a conceptual allocation of different tasks performed by a single software program and/or hardware unit. The illustrated logic, modules and functionality can be located at a single site (e.g., as implemented by a processing device), or can be distributed over plural locations. As used herein, the terms “first functionality” and “second functionality” can correspond to different physical mechanisms for delivering resource information, or the same physical mechanism operating under different conditions.

A.1. Overview of an Exemplary System that Implements Rate-Limited ACC for the Constant and Capped VBR Delivery of Resource Information

The system 100 shown in Fig. 1 represents one exemplary and non-limiting system that can be used to implement the rate-limited ACC strategy, corresponding generally to functionality developed by Microsoft Corporation to deliver media information (such as television programs, movies, pictures, audio resources, etc.) in streaming fashion to users. However, the rate-limited ACC strategy can be implemented

1 on many other types of systems. Namely, the rate-limited ACC strategy described herein
2 can be employed in any system that involves a source module (referred to herein as a
3 server module) which delivers resource information to a target module (referred to as a
4 client module) in streaming fashion. The examples developed in Section A pertain
5 particularly to the delivery of resource information at a constant data rate. The analysis
6 also applies to worst-case assumptions for the case of capped variable bit rate (VBR)
7 delivery (where the resource information is delivered at the maximum data rate
8 permitted). Section C below describes variations of the concepts developed herein.

9 By way of overview, the system 100 includes head-end functionality 102, which
10 comprises acquisition functionality 104 for delivering resource information to a
11 collection of client modules 106 via a coupling mechanism 108. The acquisition
12 functionality 104 can receive information from one or more sources. The sources can
13 represent any kind of entity which produces or provides information, such as
14 conventional cable or satellite television providers, one or more Video-On-Demand
15 (VOD) suppliers of information, one or more publishing houses of information, one or
16 more library sources of information, any kind of Internet-enabled repository of
17 information, and so on. In general, the sources can supply live information or
18 prerecorded information. Live information corresponds to information that captures a
19 current state of an ongoing event (such as a sporting event which is being televised live).
20 Prerecorded information corresponds to information that has already been recorded. (The
21 information may have been recorded in its entirety, before it is delivered to the client
22 modules, or the information may be recorded as it is being delivered, there being a time
23 lag between recordation and delivery.) The acquisition functionality 104 itself can
24 comprise one or more server computers, dedicated hardware, or other mechanism
25

1 dedicated to the task of retrieving the resource information and (optionally) storing the
2 resource information prior to dissemination to the server modules 106.

3 The system 100 can use any kind of coupling mechanism 108 to couple the head-
4 end functionality 102 to the client modules 106. For instance, the coupling mechanism
5 108 can include any kind of network (or combination of networks), such as a wide area
6 network (e.g., the Internet), an intranet, Digital Subscriber Line (DSL) network
7 infrastructure, point-to-point coupling infrastructure, and so on. The coupling
8 mechanism can use or involve any kind of protocol or combination of protocols, such as
9 the Internet Protocol (IP), the Transmission Control Protocol (TCP), the User Datagram
10 Protocol (UDP), the HyperText Transfer Protocol (HTTP), the Simple Object Access
11 Protocol (SOAP), the Real Time Transport Protocol (RTP), Real Time Streaming
12 Protocol (RTSP), and many potential others. In the case where one or more digital
13 networks are used to disseminate information, the coupling mechanism 108 can include
14 various hardwired and/or wireless links, routers, gateways, name servers, and so on. In
15 the case where DSL infrastructure is used to disseminate information, the coupling
16 mechanism 108 can utilize the services, in part, of telephone coupling infrastructure and
17 DSL processing functionality. The rate-limited ACC solutions described herein are not
18 technology-dependent, meaning that these techniques can “sit on top” of different
19 protocols and technologies, or otherwise integrate and interact with such protocols and
20 technologies.

21 The head-end functionality 102 relies on a combination at least two different
22 delivery systems to deliver the resource information supplied by the acquisition
23 functionality 104 (or supplied by some other source). (Although, as will be described in
24 Section C, the ACC strategy can also be performed using a single delivery mechanism.)
25 First delivery functionality comprises a collection of server modules 110 for delivering a

1 first stream of resource information to the client modules 106. Second delivery
2 functionality comprises multicast functionality 112 for delivering a second stream of the
3 resource information to the client modules 106. The construction and cooperative
4 operation of the server modules 110 and multicast functionality 112 are explained in
5 greater detail below.

6 Each client module can interact with the head-end functionality via a
7 communication channel. As broadly used herein, the term “communication channel”
8 represents the resources that the system 100 uses for handling the communication
9 between a particular client module and the head-end functionality 102. The resources
10 that comprise the communication channel may include functionality supplied by the
11 head-end functionality 102, resources provided by the coupling mechanism 108, and/or
12 functionality provided by the client modules 106 themselves.

13 The system 100 is a rate-limited environment. This means that there are threshold
14 limits that restrict the amount of resources that the system 100 allocates to each
15 communication channel, which place restrictions on the bandwidth of each
16 communication channel. For example, assume that client module A is coupled to the
17 head-end functionality 102 by a collection of system resources that define a
18 communication channel as broadly used herein. The system 100 imposes a limit on the
19 amount of resource information in a single stream that can be fed to module A. Or
20 assume that client module A has the capacity to display plural streams of resource
21 information at the same time (which might be the case, for example, where client module
22 A has the capacity to provide a picture-in-picture presentation). In this case, the system
23 100 imposes a limit on the combined amount of information supplied by plural streams of
24 resource information that can be consumed at any time.

1 In another example, assume that plural communication modules (e.g., client
 2 module A and client module B) are allocated to a single communication channel to
 3 interact with the head-end functionality 102. This scenario can correspond to the case
 4 where the system 100 allocates a prescribed amount of communication resources that a
 5 household containing plural media presentation units can use to interact with the head-
 6 end infrastructure 102. In this case, the system 100 imposes a limit on the amount of
 7 resource information that can be supplied to the plural client modules at the same time.

8 In one exemplary implementation, the system 100 can implement its rate
 9 constraints by allocating a certain amount of bandwidth over an amount of bandwidth
 10 required to deliver the stream in a steady-state condition. Stated generally, assume that a
 11 communication channel allocates a total amount *MaxRate* of bandwidth to a
 12 communication channel. The excess capacity as a function of time can therefore be
 13 expressed as $Excess(t) = MaxRate - ActualCurrentRate(t)$, where *Excess(t)* refers the
 14 amount of excess bandwidth as a function of time, and *ActualCurrentRate(t)* refers to the
 15 amount of bandwidth being consumed to deliver the resource information in its steady
 16 state condition as a function of time. (*MaxRate* can potentially change as a function of
 17 time too.) For the case of constant data rate delivery, the *ActualCurrentRate* can be
 18 assumed to be constant, so the excess amount can also be assumed to be constant. (Such
 19 simplifications can also be made in the case of capped VBR to provide analysis results
 20 for worst-case assumptions.) In these scenarios, one way to express the excess capacity
 21 is by forming a fractional excess capacity value *E*, defined with respect to the nominal
 22 data rate of the stream. For instance, if the system 100 allocates 30 percent beyond that
 23 which is required to deliver the stream at the nominal data rate, then *E* would be 0.3. In
 24 this way of representing the excess capacity, the system 100 allocates a total amount of
 25 bandwidth of $(1 + E) * Nominal$ to the communication channel, where *E* represents the

1 fractional amount of excess capacity of the communication channel above and beyond
2 what is required to deliver the resource information in the steady-state condition
3 (*Nominal*). (Again, the nominal data rate (*Nominal*) refers to the peak data rate at which
4 the stream is expected to be delivered in the steady-state condition.)

5 The system 100 can select the parameter E in different ways. In one case, the
6 system 100 can be pre-provisioned to operate using a prescribed value for E . Different
7 values of E may be appropriate for different streaming environments, depending on
8 various technical and business considerations. As noted above, in one exemplary and
9 non-limiting environment, E can be selected such that each communication channel is
10 configured to operate at approximately 30 percent over the nominal rate which is required
11 to present a stream of resource information in normal conditions. Many other values can
12 be selected for E . If E is set at 100 percent, then the system can accommodate two
13 complete versions of the stream at the same time. The fact that E is less than 1.0 has
14 certain consequences which will be set forth in detail below. As will be described in
15 Section C, E can also be greater than 1.0, with the consequences set forth in that section.

16 In another case, the system 100 can automatically determine the level E . For
17 example, a client module can progressively increase the data rate at which it interacts
18 with the head-end functionality 102. The communication channel has certain limitations
19 regarding how much data it can stream to the client module. This means that the system
20 100 will perform progressively worse near the limits of the channel's capacity. Namely,
21 the communication channel will begin to "drop" packets. To remedy the problem of lost
22 packets, the client module will send retry requests to the head-end functionality 102 to
23 ask for the missing packets. However, as the data rate increases, the system 100 may be
24 unable to fill in the missing packets using this retry mechanism before the packets are
25 consumed by the client module's decoder. The missing packets are referred to as "holes"

1 herein. The consumption of resource information with “holes” produces a presentation
2 having poor quality.

3 Further, the operation of sending retry requests to a server module and supplying
4 retry packets to the client module itself obviously consumes the resources of the
5 communication channel, potentially overtaxing the communication channel yet further.
6 If the data rate increases to a certain level beyond the capacity of the communication
7 channel, the loss of packets and the consequent flood of retry information can cause a
8 deleterious feedback loop, leading to a state referred to as a packet storm. Such a
9 condition can completely disable the presentation of the stream.

10 In any event, the client module can determine the level E by iteratively increasing
11 the data rate, determining the level of packet loss at each data rate, and selecting the value
12 E at which the level of packet loss is acceptable, but beyond which becomes
13 unacceptable. The floor for the value E is effectively 0 percent, which yields the nominal
14 rate of the stream.

15 Once the value of E is established, the main objective of the system 100 is to
16 supply resource information to the client modules 106 in a manner which does not exceed
17 the prescribed rate limits. To carry out this objective, the system 100 must ensure that the
18 streams of resource information supplied by the server modules 110 and the multicast
19 functionality 112 are combined in such a manner that the aggregate demand on
20 communication resources never exceeds the prescribed data rate limit. This provision is
21 advantageous because it prevents the system 100 from inundating the communication
22 channels with too much resource information. It therefore prevents too many packets
23 from being dropped, and the consequent degradation in resource presentation when the
24 decoder consumes such resource information having holes. It further helps prevent the
25

1 extreme case where one or more streams are simply dropped altogether because the
2 system 100 becomes too overburdened.

3 The precise manner of combining the streaming output of the server modules 106
4 and the multicast functionality 112 will be set forth in greater detail below. As a
5 preliminary matter, the description will provide more explanation pertaining to the
6 configuration and operation of the server modules 110 and the multicast functionality
7 112.

8 The server modules 110 can include one or more server computers (e.g., a so-
9 called "farm" of computer devices) or other mechanisms. The server modules 110 can
10 each basically perform the same functions. Routing functionality and/or load balancing
11 functionality (not shown) can assign any of one of the server modules 110 to provide
12 services to any one of the client modules 106. Due to use of redundant components at the
13 head-end functionality 102, the following discussion will often revert to the singular
14 when discussing the server modules 110. Indeed, it should be noted that the system 100
15 can function with even one server module.

16 From a high level perspective, the server modules 110 perform two basic roles:
17 (1) sending a burst of resource information upon commencement of the stream; and (2)
18 handling retry requests from the client modules 106. As to the first role, these server
19 modules 110 are responsible for supplying the resource information to the client modules
20 106 when the client modules first "tune to" resource information. Such initial events
21 occur when the user turns on a client module to consume whatever program happens to
22 be presented on an initial channel; such initial events also occur when the user switches
23 channels, thereby tuning to a new program. The goal of the server modules is to feed a
24 client module resource information on an expedited basis when the client module first
25 tunes to a stream of the resource information. This, in turn, enables the client module to

1 present the resource information more quickly, e.g., hopefully fast enough so that the user
 2 will not experience a delay when changing channels.

3 In one exemplary and non-limiting implementation, a server module performs the
 4 expedited buffer-fill operation by streaming the resource information to a client module
 5 at the full capacity of the communication channel that couples the server module to the
 6 client module. In other words, the server module sends an initial burst of resource
 7 information at a level $(1+E)*Nominal$, meaning that the burst is performed at the normal
 8 nominal rate of the stream plus some fractional rate E times the nominal rate of the
 9 stream. According to another feature, the server module supplies this burst to the client
 10 module in a unicast fashion, meaning that the server module generally devotes resources
 11 to communicate with the client module on a one-to-one basis. The unicast can be
 12 performed using “datagram” connectionless style communication, such as UDP. The
 13 unicast can alternatively be performed using TCP, or some other protocol or mechanism.
 14 As will be explained in greater detail in connection with Fig. 2, the server module also
 15 starts the unicast burst at a key frame of the resource information. This allows the client
 16 module to immediately begin presenting the resource information (because, without this
 17 provision, the client module would otherwise have to wait until it received a key frame if
 18 it initially received a series of difference frames).

19 Fig. 1 generally illustrates the unicast coupling between a server module and a
 20 client module by the unicast stream 114. This stream 114 has two components. The
 21 larger of the two arrows represents the transmission of a stream of resource information
 22 to the client module in unicast fashion. The stream can comprise a plurality of packets,
 23 each containing pieces of the resource information (e.g., one or more frames of audio-
 24 video information). The resource information can be expressed using any standard, such
 25 as MPEG-2, among many other standards. The client module that receives the stream of

1 resource information can reconstruct the resource information based on the received
2 packets. The smaller of the two arrows represents any commands and other information
3 that are transmitted by the client module to the server module. For example, the client
4 module can send retry requests to the server module. The client module will generate
5 retry requests when it determine that it has missed one or more packets in the stream of
6 information supplied by the server module based on holes in the received information.
7 (Alternatively, as will be described in Section C, the client module can rely on another
8 entity to receive missing packets.)

9 In connection with these retry requests, the second role performed by the server
10 module is to honor such retry requests (if possible) by supplying any packets of missing
11 resource information to the requesting client module. In one implementation, the server
12 module can place the retry packets into the stream of resource information that it supplies
13 to the client module, although, in one exemplary implementation, the retry packets are
14 not otherwise distinct from new (first-transmitted) packets in the stream. (In an
15 alternative exemplary implementation, the retry packets are distinct from the first-
16 transmitted packets.) Upon receiving the retry packets, the client module can insert them
17 into their proper location in the client module's buffer. The client module can order the
18 packets based on sequence numbers associated with the packets.

19 On the other hand, the multicast functionality 112 performs the basic role of
20 delivering resource information to the client modules 106 in a multicast mode. As
21 broadly used herein, the term multicast refers to the transmission of resource information
22 to a client module in any manner that does not require one-to-one allocation of head-end
23 resources to the client module. One kind of multicast functionality 112 can be
24 implemented using the Internet Group Management Protocol (IGMP). IGMP
25 functionality can supply resource information to a group of recipients by a tree of

1 distribution nodes which receive the resource information from an ultimate source. A
2 new client module can join the ongoing multicast by tapping into the multicast
3 distribution; it can perform this task by locating an appropriate distribution node within
4 the tree of such nodes. Indeed, in the system 100 of Fig. 1, both the server modules 110
5 themselves and the client modules 106 can couple to a multicast distribution framework
6 which ultimately receives resource information from the acquisition functionality 104.
7 However, the IGMP protocol is merely one non-limiting strategy for sending resource
8 information in a multicast fashion to the client modules. The term multicast transmission
9 should be construed to encompass any kind of broadcast (non-unicast) transmission of
10 resource information, including, for instance, broadcast satellite transmission of resource
11 information.

12 In one case, a single head-end infrastructure can supply logic which implements
13 both the server modules 110 and the multicast functionality 112. In another case, two
14 different infrastructures can be devoted to implementing the server modules 110 and the
15 multicast functionality 112. These server modules 110 and multicast functionality 112
16 can be maintained and operated by the same commercial entity, or different respective
17 commercial entities.

18 However implemented, Fig. 1 generally illustrates the multicast coupling between
19 the multicast functionality 112 and one of the client modules 106 by the multicast stream
20 116. This stream 116 has two components. The larger of the two arrows represents the
21 transmission of a stream of resource information to the client module in multicast
22 fashion. Again, the stream can comprise a plurality of packets, each containing pieces of
23 the resource information (e.g., one or more frames of audio-video information). The
24 client module that receives the stream of resource information can reconstruct the
25 resource information based on the received packets. The smaller of the two arrows

1 represents any commands or information that is transmitted by the client module to the
 2 multicast functionality 112. For example, the client module can issue a join request via
 3 this route to ask the multicast functionality 112 if the client module may join an ongoing
 4 multicast stream.

5 With the above preliminary explanation, more can be said on how the head-end
 6 functionality 102 presents the unicast stream 114 and the multicast stream 116 in such a
 7 manner so as not to exceed the rate limit, e.g., $(1+E)*Nominal$. To review, upon a change
 8 event (e.g., when a client module tunes to a new program), a server module comes into
 9 action by sending a burst of unicast resource information at the exemplary rate of
 10 $(1+E)*Nominal$. This burst rapidly fills up the buffer of the client module and then
 11 terminates. Near the end of the burst, the server module will instruct the client module to
 12 switch to the multicast stream 116. The client module will comply by sending a request
 13 to the multicast functionality 112, asking to connect to the multicast stream 116. When
 14 this connection is established, the client module then receives the multicast stream from
 15 the multicast functionality 112 at less than or equal to the nominal data rate (*Nominal*).
 16 (That is, this section assumes that the multicast rate remains steady at *Nominal*, but it
 17 need not. In other cases, the multicast rate will vary but will be capped at the *Nominal*
 18 rate. Section C sets forth an implementation which takes advantage of this fact by
 19 dynamically computing the excess capacity in the channel due to the varying multicast
 20 data rate.)

21 In summary, the ACC strategy therefore has the effect of transitioning from a first
 22 mode of delivery of the resource information (e.g., unicast via the server module) to a
 23 second mode of delivery of the resource information (e.g., multicast via the multicast
 24 functionality 112). The use of the terms “unicast stream” 114 and “multicast stream”
 25

1 116 should not obscure the fact that these streams contain the same pieces of data
2 delivered using different delivery techniques.

3 There is a join interval between the end of the unicast burst 114 and the start of
4 the multicast stream 116. In this join interval, the client module is attempting to establish
5 a connection to the multicast functionality 112. For example, in an IGMP multicast
6 solution, the join interval may depend, in part, on a delay in router(s) handling the IGMP
7 join request. In a satellite broadcast solution, the join interval may depend, in part, on a
8 time required to switch the client module's tuner (not shown) to a new transponder (to
9 receive the broadcast). Still other solutions are possible, each having its own respective
10 "join behavior."

11 There are different events that can be used to demarcate this join interval. In one
12 case, the server module continues to transmit at the burst rate up until the time it issues a
13 request for the client module to switch over to the multicast stream 116. In another case,
14 the server module can continue to transmit at the burst rate beyond this point up until the
15 minimum time period in which the client module can successfully satisfy the server
16 module's request that it join the multicast stream 116. The second scenario will shorten
17 the length of the join interval. In any event, the general goal in the rate-limited
18 environment is to prevent two full-rate streams from being received simultaneously,
19 because this may exceed the $(1+E)*Nominal$ capacity threshold (where E is understood as
20 being less than 1 in this exemplary and non-limiting example).

21 Another issue that this disclosure addresses at length (below) is what activity can
22 occur in the join interval. In a first scenario, the server module that is servicing the client
23 module sends no resource information to the client module in this join interval. This
24 would mean that the below-nominal rate is 0. In another case, the server module
25 transmits resource information taken from unicast stream 114 at a below-nominal rate. In

one particular case, the server module sends the client module resource information at the rate of $E \cdot \text{Nominal}$. The value of E can be fixed in advance. Or, in the event that the server module (and/or client module) is continually monitoring the available excess capacity of the communication channel, then the excess capacity can fluctuate depending on the time-varying excess capacity (as will be described in Section C). Generally, it is useful to send at least some data from the unicast stream 114 during the join interval because this allows the client module's buffer to receive at least some resource information in what would otherwise be a "dead" period.

However, regardless of what the server module does during the join interval, the client module will generally miss some of the packets that would have been transmitted during this interval (that is, if there was no break in the full-rate transmission). Stated another way, there will be holes in the received information caused by the join interval. Without corrective measures, the client module's output will suffer degradation if the client module consumes the resource information containing the holes. Such degradation could manifest itself in various artifacts, or possibly the complete failure of the presentation (e.g., where the system 100 effectively drops the stream).

According to the novel technique employed herein, the client module can overcome the problems that would be caused by the holes by sending retry requests to the server module. These retry requests ask the server module to supply any packets that the client module has identified as missing. The client module can identify missing packets by ordering the received packets according to their sequence numbers. A discontinuity in the sequence numbers suggests that a packet has been lost (or, if corrupted, effectively lost). However, with some transports (such as UDP) because the client module can expect to receive packets out of order (upon which it reassembles them into the proper order), it may assess a loss of packets only when the packets are not received within a

prescribed amount of time. In any event, the client module can make its retry requests in one or more submissions to the server module. The client module can make these requests after the join interval, when the multicast stream 116 is already being received. The server module then responds by supplying the missing packets requested by the client module. Again, the server module supplies the missing packets simultaneously with the presentation of the multicast stream. So as not to exceed the overall limitations on data rate, the server module will supply the retry packets at a rate that, when combined with the actual or nominal bandwidth requirements of the multicast stream, does not exceed the $(1+E)*Nominal$ threshold requirements. The effect of receiving the retry packets after the join interval is to supply the client module with any missing information before this missing information negatively impacts the performance of the client module. In other words, the retry packets fill in the holes in the received stream before these holes are consumed by the client module's decoder.

The above features describe the run-time behavior of the rate-limited ACC strategy. Another aspect of the strategy pertains to the manner in which the system 100 is configured in order to successfully achieve the above-described rate-limited ACC streaming behavior. Namely, a number of different parameters govern the behavior of the rate-limited streaming. Exemplary such parameters include:

- The amount of fractional excess capacity E that is to be employed;
- The amount of excess resource information that the server module needs to forward to the client module in order to accommodate the rate-limited ACC strategy, and the amount of time required to transmit this excess resource information (e.g., the length of the unicast burst);
- The amount of buffer that the server module needs to allocate for storing the excess resource information;

- The amount of buffer than the client module needs to allocate to store the excess resource information (and a determination of whether it has such capacity);
- The type of unicast transmission that is permitted during the join interval (including receiving no unicast transmission or receiving unicast transmission at the fractional excess rate of E);
- The length of time associated with the join interval (which generally corresponds to the technology-specific time required to make a multicast connection);
- Various features regarding the characteristics of the retry behavior that takes place after the join interval;

The selection of one parameter generally impacts other parameters. The general goal in configuring the system 100 is to select a combination of such parameters that yields the desired rate-limited ACC behavior.

One exemplary configuration tool allows a designer to define various combinations of parameters that will provide desired rate-limited ACC behavior by ensuring that the ratio (T_{burst}/T_{join}) is related to some function of the fractional excess capacity E in a defined way, where T_{burst} describes the length of time of the unicast burst, and T_{join} describes the amount of time required to join the multicast stream 116. Calculation of this nature can be performed by any of the “actors” within the system 100, such as any of the server modules 110 and/or any of the client modules 106.

Alternatively, the system 100 may employ a separate configuration tool 118 to make any of the computations described herein. The configuration tool 118 can comprise a computer device, which can be either a stationary or portable computer device. Further, the computer device 118 can be either a stand-alone or networked computer device. The configuration tool 118 can rely on a human operator to manually input relevant parameters, or can automatically extract relevant parameters from the system 100.

1 Finally, the system 100 can rely on a human operator to implement the recommendations
 2 of the configuration tool. Or the configuration tool 118 can automatically supply its
 3 recommendations to the relevant actors in the system 100, and these actors can
 4 automatically reconfigure themselves based on the recommendations.

5 As a conclusion to this subsection, it can be seen that the rate-limited ACC
 6 strategy described above achieves the benefits of rapidly supplying the client modules
 7 106 with resource information upon channel change events, to thereby improve user
 8 experience (by not subjecting the user to annoying delays when he or she changes
 9 channels). At the same time, the rate-limited ACC strategy does not demand that the
 10 system 100 devote large bandwidth capacity to accommodate this feature. Stated in
 11 another way, the system eliminates the need to provide a bandwidth that can
 12 accommodate the concurrent presentation of two full-rate streams (e.g., a full-rate unicast
 13 stream at the same time as a full-rate multicast stream). Instead, the strategy staggers the
 14 burst-rate unicast stream 114 and the multicast stream 116 in series, with a join interval in
 15 between; any packets that are missed due to this join interval can be filled in by retry
 16 requests, where the retry activity takes place concurrently with the presentation of the
 17 stream and reception of the multicast stream (because the multicast stream 116 takes
 18 place at or below the nominal rate, allowing some excess capacity for retry activity). The
 19 system is therefore suitable for those environments where there are significant constraints
 20 on bandwidth due to any number of technical and/or business-related reasons.

21 The following sections provide more details on the topics introduced above.

22 *A.2. Exemplary Server Module and Client Module for Rate-Limited ACC*

23 Fig. 2 describes the rate-limited ACC functionality on a next level of detail by
 24 showing more information regarding any of the exemplary server modules 106 and client
 25 modules 110 that appear in Fig. 1. Namely, Fig. 2 shows an exemplary server module

1 202 which interacts with an exemplary client module 204 via a communication channel
2 206 to deliver resource information. Again, the case of a constant bit rate or worst-case
3 capped VBR is assumed, although, as will be described in Section C, the principles
4 described herein can be extended to other environments.

5 The server module 202 can be implemented as any kind of computer device with
6 hardware and/or software that enables it to function in the role of a supplier of services to
7 client modules. As a computer device, the server module 202 can include any array of
8 conventional hardware (not shown), such as one or more processor devices (CPUs),
9 RAM memory, ROM memory, various media drives and associated media (e.g., hard
10 drives and associated hard disc), various buses to couple the components together,
11 various communication interfaces, various input/output functionality, and so forth.

12 The client module 204 can include any kind of equipment for interacting with the
13 server module 202 (and for interacting with the multicast functionality 112). In one
14 exemplary case, the client module 204 can correspond to set-top boxes, game-type
15 consoles (such as Microsoft Corporation's XboxTM game consoles), software and/or
16 hardware functionality integrated into the associated presentation devices, general
17 purpose or special purpose computer devices, and so forth. In any case, the client module
18 204 can include at least the basic hardware identified above with respect to server module
19 202, as well as any additional hardware specifically targeted to its role as a consumer
20 (rather than supplier) of resource information.

21 Fig. 2 shows that the client module 204 comprises processing functionality 208
22 for processing resource information received from the server module 202, as well as
23 presentation functionality 210 that outputs the processed resource information to a
24 consumer. For instance, the client processing functionality 208 can correspond to a set-
25 top box, and the presentation functionality 210 can correspond to a television unit which

1 couples to the set-top box. Or the client processing functionality 208 can correspond to a
 2 general purpose computer and the presentation functionality 210 can correspond to a
 3 computer monitor which couples to the computer.

4 Fig. 2 illustrates the rate-limited channel 206 by showing a bandwidth 212 that is
 5 allocated to steady-state "normal" streaming of the resource information at a nominal
 6 peak rate, and an excess bandwidth 214 which is added to the nominal bandwidth 212.
 7 The extra bandwidth 214 can be expressed as a bit rate or as a fractional value E . The
 8 channel 206 is rate-limited in the sense that the instantaneous data rate should not exceed
 9 the ceiling threshold defined by the aggregate of the nominal bandwidth 212 and the extra
 10 bandwidth 214, that is, in aggregate, $(1 + E)$ of the nominal bandwidth.

11 An example will make these concepts more concrete. Consider the illustrative
 12 case where a client module is allowed to use two one megabit streams and two 100
 13 kilobit streams simultaneously. Assume further that the system that supplies these
 14 streams is provisioned with three megabits to implement the communication channel,
 15 which leaves 0.8 megabits free on the channel. This yields an excess percentage amount
 16 of $0.8/2.2$ or $.36$. This fractional rate limitation can be applied to each component stream
 17 supported by the communication channel. That is, each one megabit stream can use $.36$
 18 megabits and each 100 kilobit stream can use 36 kilobits extra without the risk of any of
 19 these streams interfering with each other. It is therefore seen that the rate restrictions can
 20 be parsed out in fractional fashion to apply individually to component streams supplied to
 21 a single client module (or supplied to multiple client modules which feed off of the same
 22 communication channel). However, to simplify the discussion, it will generally be
 23 assumed that the rate restrictions apply to the case where a single client module is
 24 presenting a single stream of resource information.

Additional information will next be provided regarding the composition of the server module 202. Namely, the server module 202 includes server functionality 216 which enables server module 202 to function in the manner described herein. In part, the server functionality 216 includes: logic (hardware and/or software) for receiving resource information from the acquisition functionality 104 (or elsewhere); logic for supplying a burst of unicast resource information to a client module 204 in response to a channel change event; logic for optionally transmitting a below-nominal rate unicast stream to the client module 204 during the join interval (when the client module is switching from unicast delivery to multicast delivery); logic for receiving requests from the client module 204 to supply missing packets, and so forth. To function in this manner, the server module 202 maintains a buffer 218. For the case of the delivery of live resource information streamed from the acquisition functionality 104, the buffer 218 can comprise a sliding circular buffer. The buffer 218 stores recently received resource information (from the acquisition functionality 104), including the resource information that will populate: the initial unicast burst; the below-nominal unicast transmission during the join interval; and the information that will be used to satisfy the client module 204's retry requests (if possible). Generally speaking, the buffer 218 comprises a FIFO-type buffer; it stores new resource information that is supplied by the acquisition functionality 104 by generally deleting the oldest resource information within the buffer 218. The server module 202 can also function as a forwarder of information that it stores in advance. Namely, instead of streaming in live media information from the acquisition functionality 104, the server module 202 can stream video on demand (VOD) information that it stores in its memory. Fig. 1 represents this scenario by the bubble labeled "VOD Source" associated with the topmost exemplary server module. In this scenario, the server module 202 can employ a static buffer 218' to delivery the VOD assets, rather than the

1 sliding circular buffer 218 shown in Fig. 2. Section C provides additional information
2 regarding the delivery of VOD assets.

3 During the initial burst at the $(1+E)*Nominal$ data rate, the server module 202
4 will send the burst for a prescribed amount of time, referred to as T_{burst} . This means that
5 the server module 202 will reach back to pull an appropriate amount of resource
6 information from the buffer 218 to compose the burst having the duration T_{burst} . Fig. 2
7 illustrates a segment of resource information 220 that is presented in the exemplary
8 unicast burst. This segments may generally correspond to resource information that was
9 received by the acquisition functionality 104 about 10 seconds prior to current time. The
10 minimum amount of excess buffer that the server module must maintain is equal to
11 $T_{delay}*Nominal$ (which is the same as $E*T_{burst}*Nominal$).

12 The resource information sent by the server module 202 can be expressed in
13 various formats. In one implementation, the resource information comprises video
14 information. The video information is composed of a series of key frames (also referred
15 to as I frames, random access points, etc.) which contain complete video frames that are
16 independently decodable and don't depend on previous resource information in order to
17 decode to a complete video frame that can be displayed to the user. The video
18 information can also include a variety of difference frames (Δ) which represent a video
19 frame for different times by expressing how the video information differs from one or
20 more other frames. MPEG-2, for example, uses this approach by presenting a
21 combination of I-frames, B frames and P frames (where the I frames constitute key
22 frames and the B and P frames constitute difference frames). In this technology, as
23 described above, the server module 202 will select the segment of resource information to
24 be sent to the client module 204 in the burst so that it begins at a key frame. This is
25 beneficial because the client module 204 cannot begin rendering resource information

1 until it receives a key frame. By forcing the server module 202 to start its transmission
 2 with a key frame, the client module 204 will not have to receive a series of difference
 3 frames before receiving a key frame. (Note that, while Fig. 2 and the other figures
 4 represent the resource information as a series of frames to facilitate discussion, it is
 5 important to note that the server module 202 will break this information up into a series
 6 of packets and transmit the information to the client module 204 in units of packets, not
 7 frames.)

8 Now turning to the composition of the client module 204, this module 204
 9 includes client functionality 222 which enables the client module 204 to function in the
 10 manner described herein. In part, the client functionality 222 includes: logic (hardware
 11 and/or software) for establishing a connection with the server module 202 as well as the
 12 multicast functionality 112; logic for receiving and rendering the unicast stream 114 and
 13 the multicast stream 116; and logic for sending retry requests to the server module 202.
 14 To function in this manner, the client module 204 maintains a client buffer 224 (if so
 15 equipped). The client buffer 224 stores a prescribed amount of resource information
 16 received from the server module 202. The client functionality 222 draws from the client
 17 buffer 224 when it presents the resource information. The maximum amount of extra
 18 buffer required is $T_{delay} * Nominal$. In other words, the client module 204 should be pre-
 19 provisioned with a buffer that can store $T_{delay} * Nominal$ amount of information (which is
 20 the same as $E * T_{burst} * Nominal$).

21 The client module 204 can communicate its capabilities to the server module 202
 22 on its own initiative, or when queried by the server module 202. If the server module 202
 23 determines that the client module 204 has insufficient buffer capacity, it can refrain from
 24 using the rate-limited ACC technique when communicating with the client module 204,
 25

1 or can use another kind of accelerated start technique that the client module 204 can
2 accommodate.

3 In operation, in one mode, the client module 208 can immediately start playing
4 resource information upon receipt of the resource information (e.g., without requiring any
5 preliminary buffering). This solution may result in some unsteady performance when
6 first tuning to a channel. In another mode, the client module 208 can require that a
7 prescribed amount of information be buffered prior to the presentation of the video
8 information. This will result in a short presentation delay, but such a delay is reduced by
9 virtue of the use of the ACC initial burst. In still another mode, the client module 208
10 can display the initial key frame as soon as it is received, then buffer a prescribed amount
11 before beginning full-motion playback.

12 Finally, Fig. 2 shows a series of timelines 226 that illustrate the manner in which
13 the client module 204 consumes resource information. Consider the case of video
14 information. Many protocols tag video information using two separate time stamps: a
15 delivery timestamp and a presentation timestamp. A delivery timestamp associated with
16 a packet describes when that packet is intended to be delivered to a receiver. The
17 presentation timestamp describes when the decoded version of that packet should actually
18 end up on the presentation functionality 210 (e.g., a television screen). The maximum
19 gap permitted between the delivery and presentation timestamps is a configurable
20 parameter of an encoder, referred to as the “max PCR/PTS delay.”

21 A series of delivery timestamps defines a delivery timeline. The first timeline in
22 the group of timelines 226 shows a delivery timeline. A series of presentation
23 timestamps defines a presentation timeline. The second timeline in the group of timelines
24 226 shows a presentation timeline. The arrows connecting the delivery timeline and the
25 presentation timeline illustrate how packets in the delivery timeline correspond to the

1 same packets in the presentation timeline. The arrows are skewed because the packets
2 are intended to be presented some time after their respective delivery times. Fig. 2 also
3 shows that the packets are evenly spaced on the delivery timeline, but are bunched up on
4 the presentation timeline. This is because video information is delivered at a constant
5 rate (the nominal rate of the stream) in this exemplary and non-limiting case (and
6 therefore satisfies the worst-case for the steady-state stream bandwidth), but the video
7 information is not consumed at a constant rate (e.g., the key frames are much larger than
8 the difference frames, but are presented on a display for the same amount of time).

9 The group of timelines 226 also shows a third timeline that describes when the
10 packets are actually presented (as opposed to when they were intended to be presented).
11 The arrows connecting the presentation timeline and the actual timeline illustrate how
12 packets in the presentation timeline correspond to the same packets in the actual timeline.
13 The arrows are skewed because the packets are actually presented some time after they
14 were intended to be presented. Namely, the large buffer 224 used by the client module
15 204 imposes a delay in the presentation of the video information (because there will
16 typically be some amount of time between when the packet is placed in the buffer 224
17 and when it is actually output to the presentation functionality 210). Different client
18 modules may have different mechanisms (e.g., buffer capacities) for delaying the
19 presentation of the video information; hence, different client modules may follow
20 different respective actual presentation timelines. Other reasons that the client modules
21 may have different actual presentation timelines include: (a) client modules tuning to the
22 stream at slightly different times; and (b) client modules requiring different amounts of
23 delay relative to a live state for smooth playback, depending on the loss rate of the
24 channel; and (c) client modules requiring different amounts of delay relative to a live
25 state for smooth playback, depending on the delivery time jitter of the channel.

1 Consider the following example to illustrate the presentation of video information
 2 for two client modules that request the same program from the server module 202. If the
 3 client modules make the request about one second apart from each other, then the server
 4 module 202 might start bursting data to both client modules starting from the same key
 5 frame. If network conditions are similar for the two client modules, the client module
 6 which connected to the server module 202 earlier in time would have an actual
 7 presentation timeline exactly one second before the second client module to connect. As
 8 long as those two client modules remain connected to the program, they will remain one
 9 second apart in actual display time.

10 In some broadcast systems the delivery timestamp serves as an authoritative clock
 11 reference. One exemplary implementation of the present system, however, uses a shared
 12 clock to control playback of the video information. For example, the system can use a
 13 Network Time Protocol (NTP) clock reference as the shared clock. In operation, the
 14 client module 204 determines the offset between the presentation timestamp information
 15 in the video stream and the shared clock, and then attempts to maintain this offset
 16 throughout the reception of the video information.

17 In the present system, the delivery timestamp also serves as floor threshold when
 18 performing the rate-limited ACC technique. Namely, as will be explored in detail below,
 19 the rate-limited ACC technique will attempt to feed video information to the client
 20 module 204 in such a way that the client module 204 never falls behind on the delivery
 21 timeline (as shown in the first timeline of the group 226 of Fig. 2).

22 *A.3. Exemplary Interaction between the Server Module and the Client Module*

23 Fig. 5 is an exemplary series of operations that explain the interaction between the
 24 exemplary client module 204 (introduced in Fig. 2) and, on one hand, the exemplary
 25 server module 202 (introduced in Fig. 2), and, on the other hand, the multicast

1 functionality 112 (introduced in Fig. 1). The order and nature of the operations are
 2 exemplary. Other implementations of the rate-limited ACC strategy can apply the
 3 operations in somewhat different order, and/or can include additional types of operations,
 4 and/or can omit one or more operations shown in Fig. 5.

5 To begin with, in operation (1), when the client module 204 initiates the
 6 presentation of a stream of resource information, it sends a request to the server module
 7 202. The head-end functionality 102 may assign the client module 204 to one of its
 8 server modules 110 using any kind of assignment algorithm (such as a fixed assignment,
 9 a load-balancing assignment, and so forth). Operation (1) may generally correspond to
 10 an event in which a user changes channels to receive a new program, or otherwise
 11 performs any act which prompts the client module 240 to request connection to a new
 12 program (or a new part of a same program, such as in response to a fast-forward
 13 command, etc.).

14 In operation (2), the server module 202 responds by setting up a unicast to the
 15 client module 204. A unicast involves a one-to-one allocation of prescribed server
 16 resources to the client module 204. The server module 202 then initiates the stream by
 17 first sending an initial burst of the resource information at a rate of $(1+E)*Nominal$,
 18 where *Nominal* defines the nominal rate of the stream, and *E* defines some prescribed
 19 fractional amount over the nominal rate (such as approximately 20 percent). This
 20 transmission defines the start of the unicast stream 114 introduced in Fig. 1.

21 Moreover, in operation (1), the client module 204 can optionally inform the server
 22 module 202 of its buffer capacity. The client module 204 might not have any extra
 23 buffer, or may have only a limited amount of extra buffer. If the server module 202
 24 determines that the client module 204 cannot accommodate the rate-limited ACC
 25 technique described herein, it will not send the burst in operation (2). It may simply

1 forward the resource information at the nominal rate (which will incur a delay before the
 2 resource information can be presented), or it will use some other ACC technique that
 3 makes fewer demands on the buffer capacity of the client module 204.

4 Presuming that the client module 204 does have the requisite buffer capacity,
 5 then, for the case of live delivery of resource information, the server module 202 sends a
 6 unicast burst that is offset by a delay time of T_{delay} in presentation timeline terms. This
 7 burst is also configured to start at a key frame.

8 Similar provisions apply to the delivery of pre-stored VOD assets. More
 9 specifically, in the VOD case, it is generally possible to perform the burst by starting
 10 from a point that is as “far back” as desired in the stream, since all the resource
 11 information is available for distribution at the server module 202. But the client module
 12 204’s buffer has a limited capacity, and there is limited need for buffering based on the
 13 characteristics of the network. In view of these factors, one exemplary strategy (for
 14 VOD) is to provide a burst having a duration (T_{delay}/E) that is sufficient to allow enough
 15 time for retries to be requested and sent well before the resource information is needed,
 16 even if the same packet is missed, for example, one or two times. In extremely lossy
 17 networks, it may be desirable to allow for enough time to attempt retry more than twice.
 18 In order to do this, T_{delay} can be made larger, which makes the burst last longer.

19 As indicated by operation 2(a), the client module 204 begins presenting the
 20 resource information a short time after receiving the start of the initial burst. This time is
 21 brief enough so that the user will not perceive an unacceptable delay in presentation upon
 22 changing channels.

23 In operation (3), at the tail end of the unicast burst, the server module sends a
 24 command to the client module 204 to join the multicast stream 116.
 25

1 In operation (4), the client module 204 responds to the request of the server
 2 module 202 by making its own request to the multicast functionality 112 to set up the
 3 multicast stream 116. As described above, this may constitute issuing an IGMP join
 4 command. This command will prompt the multicast functionality 112 to connect the
 5 client module 204 to an appropriate node in which it can receive an ongoing multicast
 6 transmission. As described above, in one case, the server modules 110 themselves may
 7 receive the resource information using a multicast distribution scheme from the ultimate
 8 source defined by the acquisition functionality 104. The join operation in this case would
 9 entail connecting the client module 204 to this same ongoing multicast stream 116.

10 Assume that the multicast join is successful and the client module 204
 11 successfully connects to the multicast stream 116. In operation (5), the multicast
 12 functionality 112 forwards its first multicast packet to the client module 204. The
 13 ensuing multicast stream is received at a data rate that is less than or equal to the nominal
 14 data rate.

15 As described in a preceding subsection, there will be a time lag between the last
 16 unicast packet received and the first multicast packet that is received, referred to herein as
 17 the join interval. In one case, the server module 202 sends no unicast information in the
 18 join interval. In another case (which is the case of Fig. 5), at the start of operation (3), the
 19 server module 202 starts sending the unicast stream at a rate of $E \cdot \text{Nominal}$. In any event,
 20 the failure to send streaming information at the full (nominal) data rate will have the
 21 effect of creating "holes" in the resource information received in the client module 204's
 22 buffer 224. These holes define packets that have been lost, because they never were
 23 received or have been received but are corrupted. These holes will lead to degradation of
 24 the output presentation if they are not filled in by the time the client module 204 decodes
 25 this resource information.

As such, in operation (5a), the client module 204 determines what holes are present in the received resource information. It can do this by ordering the received packets according to their sequence numbers and then noting any discontinuities in the sequence numbers; the discontinuities define holes.

In operation (5b), the client module 204 sends a command to the server module 202 that requests the server module 202 to send the missing packets. The client module 204 can send multiple commands for each respective missing packet or can send commands that each identify several of the missing packets.

After receipt of the retry request(s), the server module 202 prepares the missing packets for transmission to the client module 204 (if it can). The server module 202 performs this task by pulling the missing packets from its buffer 216 (if they are still present). In operation (6), the server module 202 sends the missing packets to the client module 204. As will be apparent in the discussion of the next figure, the server module 202 sends the retry information as a normal stream of packets at a data rate of $E \cdot \text{Nominal}$, concurrently with the receipt of the multicast stream at or below the nominal data rate, such that the total bandwidth is again $(1+E) \cdot \text{Nominal}$. The retry packets are “normal packets” in the sense that they are not specifically earmarked as retry packets. The client module 204 can identify retry packets because it has requested them; they are identified upon receipt by their sequence numbers. Section C describes variations on this protocol, in which the server module 202 can take advantage of dips in the data rate of the multicast stream by sending additional retry information (such that the server module 202 sends information at a rate of $\text{MaxRate} - \text{ActualCurrentMulticastRate}$, which can be expressed also as $E \cdot \text{Nominal} + (\text{Nominal} - \text{ActualCurrentMulticastRate})$).

Fig. 4 shows a bandwidth-vs.-time depiction 400 that illustrates how the operations described in Fig. 5 correspond to different utilizations of bandwidth at

different times. Namely, the horizontal axis of this depiction 400 corresponds to time. The vertical dimension any of the rectangles shown in the depiction 400 corresponds to the amount of bandwidth that is being consumed. The largest encompassing rectangle has a vertical height corresponding to the maximum bit rate permitted by the rate-limited environment, namely $(1+E)*Nominal$. As a preliminary matter, it should be noted that the dimensions of the other rectangles shown within the larger rectangle may not be drawn to scale; namely, the dimensions of some of the rectangles may be expanded or compressed to facilitate illustration and discussion. Further, again note that the simplifying case of constant bit rate is assumed (which also applies to the worst-case scenario of capped VBR). Section C describes how this analysis can be extended for other scenarios.

The enumerated operations shown in Fig. 4 correspond to the same numbered operations shown in Fig. 5, thus there is no need to repeat their description in the context of Fig. 4. The focus in Fig. 5 will center on the relative timing of various events, and the amount of information transmitted at different junctures in the streaming operation. Note first that the unicast burst 402 fills up all of the available bandwidth, e.g., $(1+E)*Nominal$. In contrast, the multicast 404 occurs at or below the nominal rate ($Nominal$). The rectangle 406 denotes a period of time between the unicast burst 402 and the multicast 404 having a horizontal dimension equal to the join interval T_{join} . This defines the period of time that the system 100 is attempting to establish a multicast connection to the multicast functionality 112.

The remaining rectangles shown in Fig. 4 illustrate activity that occurs in the extra bandwidth region (E). A first optional rectangle 408 corresponds to a unicast stream transmitted by the server module 202 at the rate $E*Nominal$. This stream is optional because the server module 202 can alternatively transmit 0 data in this region. Also, the

1 server may choose to additionally use any portion of Nominal left unused by the
 2 multicast. In any event, the “empty” join rectangle 406 will lead to holes in the received
 3 stream of resource information. To address these holes, the client module 204 will send
 4 retry requests to the server module 202, asking the server module 202 for packets that
 5 have been assessed as lost due to the empty join box 406. The small rectangle 410
 6 represents the small amount of bandwidth consumed by the request operation itself. The
 7 following retry rectangle 412 corresponds to the resultant the flow of unicast retry
 8 information from the server module 202 to the client module 204. That is, the retry
 9 information contains the missing packets requested by the client module 204, transmitted
 10 at the rate of $E \cdot \text{Nominal}$ (for the case of constant data rate). The area of the rectangle
 11 412 is equal to or approximated by the area of rectangle 406 minus the area of rectangle
 12 408. More intuitively stated, this means that the length of the join interval and the
 13 amount of data that is sent in the join interval determine the area of rectangle 412. If no
 14 unicast information is transmitted during the join interval, then rectangle 412 will have a
 15 correspondingly greater area.

16 Again, the notable benefit of this technique is that it permits ACC to be
 17 performed, thereby quickly supplying a program to user upon switching to it, within an
 18 environment that has significant limitations on bandwidth.

19 Fig. 5 provides yet further information regarding the operation of the system 100.
 20 This figure particularly shows the contents of the client module 204’s buffer 224 at three
 21 junctures, labeled X, Y and Z, which are also identified in Fig. 4. The contents are shown
 22 arranged on the presentation timeline, not according to when the contents are received.
 23 This figure particularly illustrates the buffer contents in terms of frames because this
 24 form of representation is instructive. It should be noted, however, that the buffer 224 will
 25 contain packets of information, and that the packets will themselves include parts of the

stream of frames. As another simplification, the buffer 224 will typically include more information than is shown in Fig. 5. Moreover, the horizontal time scale of the presentation timeline 224 is appropriately compressed and expanded in Fig. 5 to facilitate illustration and discussion.

At the start of the unicast burst, the client module's buffer 224 is empty. At location X shown in Fig. 4, the buffer 224 has a state shown by the topmost depiction in Fig. 5. Location X corresponds to the juncture at which the entire unicast burst 402 has been sent. Location X also marks the start of the join interval 406. At this juncture, the entire span of the buffer (A+B) corresponds to the amount of resource information that has been received by the burst 402. The vertical line 502 demarcates region A from region B. Region A corresponds to the amount of information that would have been received if the server module 202 had transmitted at the nominal rate, instead of at the excess rate of $(1+E)*Nominal$. Region B defines the amount of extra information that has been transmitted due to the rate of $E*Nominal$. Thus, $(A + B)/A = E$. The line 502 also marks the location where the client module 204 is consuming packets from the perspective of the delivery timeline (again refer to the timelines shown in grouping 226 of Fig. 2).

At location Y shown in Fig. 4, the buffer has a state shown in the middle part of Fig. 5. Location Y corresponds to a juncture at which the unicast transmission at the rate of $E*Nominal$ has just come to an end.

The span of presentation timeline formed by the aggregate of segments C and D corresponds to the join interval (T_{join}), namely, the amount of time required to set up the multicast stream 116. The C segment of the presentation timeline represents information the server module 202 sent at the rate of $E*Nominal$. Thus, the amount of information transmitted in segment C is $T_{join}*E*Nominal$. This segment takes a time of T_{join} to be

delivered. Fig. 5 shows that the client module 204 has received a complete key frame during the unicast transmission at the rate of $E \cdot \text{Nominal}$, although this is merely representative.

The segment D marks the amount of information that has been lost due to the failure to transmit at the nominal bit rate during the join interval. The quantity of information lost during the join interval is denoted by $T_{\text{join}}(1-E) \cdot \text{Nominal}$. The hole defined by segment D has been shown as a continuous block of missing data; but the hole can be comprised of a general span having a mixture of received packets and lost packets. Again, the horizontal span of presentation timeline illustrated in Fig. 5 may not be drawn to proportion.

The vertical line 504 again represents the current point of consumption from the perspective of the delivery timeline. The span of time 506 stretches from the current consumption time 504 to the beginning of the hole (represented by segment D). This segment 506 represents the amount of time that the client module 204 has to fill in the hole. In other words, the client module 204 must fill in the hole before the current consumption time 504 reaches the hole. After the hole, the information 508 denotes the receipt of the first multicast packets from the multicast stream 116.

At location Z shown in Fig. 4, the buffer 224 has a state shown at the bottom part of Fig. 5. Location Z corresponds to a juncture at which the retry packets have all been received in response to a request (or plural requests) by the client module 204. At this juncture, note that the retry information has successfully filled in the hole denoted by span D. Further note that the retry was also successful in the sense that the hole has been filled in before the current point of consumption 510 (from the perspective of the delivery timeline) has reached the start of the hole. Then, after what was once the hole, the client

1 module 204 receives the remainder of the resource information via the multicast stream
2 116.

3 *A.4. Exemplary Mathematical Analysis and Configuration of the Rate-Limited* 4 *ACC Streaming Operation*

5 As can be appreciated from the previous sections, the behavior of the system 100
6 shown in Fig. 1 is governed by a combination of several interrelated parameters. Thus, in
7 order for the system 100 to achieve the benefits described above, an appropriate
8 combination of parameters must be selected. The above subsection A.1 enumerated
9 several exemplary parameters that can be taken into account when configuring the
10 system, various combinations of which can achieve the desired retry behavior that has
11 been described.

12 This section particularly describes mathematical analyses for deriving
13 relationships between the amount of time in the initial burst (T_{burst}), the amount of time
14 required by the join interval (T_{join}), and the amount of excess capacity E . The
15 relationships define combinations of these variables that will yield the desired retry
16 behavior. More specifically, Figs. 6-10 describe the computation of such mathematical
17 relationships for different retry-related assumptions. Any functionality can be used to
18 perform the computations described herein, such as the configuration tool 118 introduced
19 in Fig. 1. Again note that the simplifying case of constant bit rate is assumed (which also
20 applies to the worst-case scenario of capped VBR). Section C describes how this
21 analysis can be extended for other scenarios.

22 To begin with, Fig. 6 corresponds to the case where the client module 204
23 demands that it completely fill in the hole of missing packets with retries before the
24 consumption point (from the perspective of the delivery timeline) reaches the end of the
25 information provided by the initial burst alone. Stated in other terms, this condition

1 specifies that the client module 204 must completely fill in region D shown in Fig. 5
 2 before the consumption arrow (e.g. 504) reaches the end of region B. The top graph in
 3 Fig. 6 (scenario 1-A) shows a scenario in which the above-defined condition has been
 4 reached. The bottom graph in Fig. 6 (scenario 1-B) shows a scenario in which the above-
 5 defined condition has not been reached.

6 Starting with the scenario 1-A, the vertical axis defines an amount of information
 7 received by the client module 204, and the horizontal axis defines progressing time. The
 8 thick diagonal line 602 describes the client module 204's demand for the resource
 9 information. The area beneath the demand curve 602 defines an underflow region 604.
 10 This underflow region 604 corresponds to an operational domain in which the client
 11 module 204 will experience the insufficient flow of resource information. If the
 12 operational state enters this region, the buffer 224 will develop holes that will be
 13 consumed by the client module 204's decoder, resulting in degraded performance. (The
 14 demand curve is actually assuming the maximum data usage. The actual demand curve,
 15 for any particular piece of content, may be lower than this "worst-case" demand curve.
 16 But because the described scheme works for the worst-case scenario, it works for all
 17 variable bit rate cases of lower demand.)

18 The first segment 606 of scenario 1-A describes the receipt of resource
 19 information during the initial unicast burst at rate $(1+E)*Nominal$. Note that the slope of
 20 segment 606 extends at a steeper angle than the demand curve 602, and thereby
 21 progressively extends away from the demand curve 602. This separation equates to
 22 desirable extra resource information that client module 204's buffer 224 accumulates
 23 during the initial burst, which reaches a maximum amount at the end of the first segment
 24 606.

1 The second segment 608 graphically describes what happen when the server
 2 module 202 transmits at the rate of $E \cdot \text{Nominal}$ following the initial burst. This
 3 transmission lasts for the length of the join interval T_{join} . The server module 202 can
 4 optionally transmit 0 amount of information in this period, whereupon the second
 5 segment 608 would have a slope of zero.

6 The third segment 610 graphically describes what happens during the
 7 combination of events corresponding to: (a) the receipt of the multicast stream 116 from
 8 the multicast functionality 112; and (b) the request for and the receipt of retry information
 9 (where the server module 202 sends the retry information in unicast fashion). The
 10 aggregate bit rate of this combined activity is $(1+E) \cdot \text{Nominal}$, and therefore has the same
 11 slope as the first segment 606.

12 The last segment 612 graphically describes what happens when the server module
 13 202 finishes sending the retry information, whereupon the client module 204 continues to
 14 receive only the multicast stream 116 at the nominal rate.

15 To repeat, the constraint on the scenarios shown in Fig. 6 is that all of the holes
 16 caused by the T_{join} period be filled in before the excess buffer accumulated during the
 17 burst period is exhausted. This can be graphically visualized in Fig. 6 by drawing a
 18 dotted line 612 with slope=0 from the end of the first line segment 606 to the end of third
 19 segment. The end of this dotted line does not enter into the underflow region 604, which
 20 means that the stated condition of Fig. 6 has been reached.

21 Scenario 1-A corresponds to the exemplary and hypothetical case in which $E =$
 22 0.5 , $T_{burst} = 6$, and $T_{join} = 1$. By contrast, scenario 1-B, describes the same conditions,
 23 except that the server module 202 transmits at a reduced extra capacity, $E = 0.25$. This
 24 decrease in E (in combination with the specified burst time) makes it impossible to
 25 establish enough reserve in the client module 204's buffer 224 to meet the stated

conditions (that the accumulated extra buffer alone provide enough information so that the client module 204 can completely fill in the gap prior to depletion of the extra buffer).

The graph for scenario 1-B confirms the above conclusion. All of the line segments 606-612 have the same meaning as described above in connection with scenario 1-A, although these segments may trace a different path than the segments shown in scenario 1-A. Of principal note in scenario 1-B is that the dotted line 612 now extends into the underflow region 604, signaling that the stated conditions have not been met.

The constraint identified in Fig. 6 is a particularly conservative one, as information continues to be received during the join interval. This additional information can be used to extend the deadline at which all holes must be filled in.

Fig. 7, for instance, presents two scenarios (scenario 2-A and scenario 2-B) that adopt the more liberal constraint that all holes created by the join period be filled in before running out of resource information collected during both: (a) the initial burst period; and (b) the join interval itself. Recall that the server module 202 optionally transmits unicast information at rate $E \cdot \text{Nominal}$ during the join interval. In the context of Fig. 5, this constraint means that the client module 204 must fill the gap D before the consumption timeline 504 reaches the end of segment C (because segment C represents the amount of information, $T_{\text{join}} \cdot E \cdot \text{Nominal}$) that was acquired during the join interval.

The segments shown in scenarios 2-A and 2-B have the same interpretation as the segments described in Fig. 6, and therefore share the same numeric labels. However, now note that the horizontal dotted line, labeled as line 702, starts from the end of the second segment 608. In scenario 2-A ($E = 0.5$), this horizontal dotted line does not cross the demand curve 602 into the underflow region 604. But in scenario 2-B ($E=0.25$), the dotted horizontal line 702 does extend into the underflow region 604.

Fig. 8 presents two other scenarios (scenario 3-A and scenario 3-B) that adopt a yet more liberal constraint that stipulates that the current point of consumption must simply not overtake any part of the buffer 224 with holes in it. This more liberal assumption can be satisfied based on resource information received as a result of any of: (a) the initial burst period; (b) the join interval; and (c) retry activity following the join interval.

The segments shown in scenarios 3-A and 3-B of Fig. 10 have the same interpretation as the segments described in Fig. 6, and therefore share the same numeric labels. However, now note that the dotted line, labeled as line 1002, again starts from the end of the second segment 608, but now has a non-zero slope E . The non-zero slope represents the fact that resource information continues to be received at fractional rate E by virtue of the retry activity, so the holes continue to be filled in at this rate. In scenario 2-A ($E=0.5$), this dotted line 802 does not cross the demand curve 602 into the underflow region 604. But in scenario 2-B ($E=0.25$), the dotted line 802 does extend into the underflow region 604.

The conditions that demarcate successful operation from unsuccessful operation for the above-described scenarios can be derived as a function of T_{burst} (the amount of time in the initial unicast burst), T_{join} (the amount of time required by the multicast join), and E (the amount of extra capacity allocated to the communication channel).

Starting with scenarios 1-A and 1-B of Fig. 6, the constraint imposed is that the amount of excess buffer attributed to the initial burst itself be sufficient to fill in all of the holes. The following analysis derives the mathematical relationship between T_{burst} , T_{join} and E for this constraint.

$$T_{burst}/T_{join} \geq 1/E^2 \Rightarrow \text{satisfiable}$$

One premise of this derivation is that there is a certain amount of time buffered in advance following the burst. Because the entire time to perform retries to account for the join interval is just the join interval divided by the extra bandwidth ratio, the amount of time buffered before the join interval must be greater than that retry time for this scenario to succeed.

In scenarios 2-A and 2-B of Fig. 9, the constraint imposed is that the amount of excess buffer attributed to the initial burst itself, combined with the $E \cdot \text{Nominal}$ unicast stream transmitted during the join interval, be sufficient to fill in all of the holes. The following analysis derives the mathematical relationship between T_{burst} , T_{join} and E for this constraint.

$$T_{burst}/T_{join} \geq (1 - E^2)/E^2 \Rightarrow \text{satisfiable}$$

One premise of this derivation is that, once the client module 204 has been instructed to switch to the multicast stream 116, advance buffer is being lost at the rate of $(1 - E)$ per-time unit. Thus, the amount of time it will take to retry (after the join interval has completed) is the total amount that would have been lost had none been sent during the join interval, minus the quantity that was received during the join interval, divided by the retry bandwidth. There must be at least that much advance buffer available to the client module 204 at the time that the join succeeds. Note that for all $E < 1$, this yields a lower value than the above first scenario (of Fig. 6).

In scenarios 3-A and 3-B of Fig. 10, the constraint imposed is that the amount of excess buffer attributed to the initial burst itself, combined with the $E \cdot \text{Nominal}$ unicast stream, and also combined with any retry information received thereafter, be sufficient to fill in all of the holes. In other words, the constraint is simply that the current point of consumption (from the perspective of the delivery timeline) keep abreast of any holes in the client module 204's buffer 224. The following analysis derives the mathematical relationship between T_{burst} , T_{join} and E for this constraint.

$$T_{burst} / T_{join} \geq (1 - E) / E^2 \Rightarrow \text{satisfiable}$$

1 This derivation starts in a similar way to the derivation for Fig. 9. The significant
2 difference is that, instead of using the straight "buffered after join" as the allowable retry
3 time, the present derivation takes account of the fact that extra data is still arriving.
4 Because the extra data is arriving from the beginning of the hole, and is arriving at the
5 rate of E , the amount of time the buffer has to survive is the original time, divided by $(1 -$
6 $E)$.

8 **B. Exemplary Method of Operation**

9 Figs. 9 and 10 together describe the operation of the functionality of the preceding
10 figures in flow chart form. To facilitate discussion, certain operations are described as
11 constituting distinct steps performed in a certain order. Such implementations are
12 exemplary and non-limiting. Certain steps described herein can be grouped together and
13 performed in a single operation, and certain steps can be performed in an order that
14 differs from the order employed in the examples set forth in this disclosure.

15 More specifically, Figs. 9 and 10 restate the operations already described in the
16 context of Fig. 5 in flowchart form. As this material has already been carefully set forth,
17 the discussion of Figs. 9 and 10 will serve as a summary; the reader is referred back to
18 the discussion of Fig. 5 for a more complete explanation of the subject matter. Again
19 note that the simplifying case of constant bit rate is assumed (which also applies to the
20 worst-case scenario of capped VBR). Section C describes how the procedures can be
21 extended for other scenarios.

22 Fig. 9 is a procedure 900 which defines operations performed by the server
23 module 202 when a channel change event occurs (or when some other event associated
24 with the start up of a stream occurs). In step 902, the server module 202 receives a
25 request from the client module 204 to receive the initial burst to facilitate start up.

1 Although not shown, this initial period can also involve an investigation of what kind of
 2 excess buffer is required by the client module 204, and a determination of whether the
 3 client module 204 can satisfy such a requirement.

4 In step 904, the server module 202 starts the burst at the rate of $(1+E)*Nominal$.

5 In step 906, the server module 202 sends a join mcast instruction to the client
 6 module 204, which instructs the client module 204 to join the multicast stream 96. In this
 7 step 906, the server module 202 also begins transmitting at the rate of $E*Nominal$ (and
 8 continues at this rate for the ensuring join interval).

9 In step 908, the server module 202 receives a request from the client module 204
 10 to stop streaming altogether (because the client module 204 has successfully connected to
 11 the multicast stream 116). In this step 908, the server module 202 also receives retry
 12 requests from the client module 204 corresponding to holes that occurred due to the join
 13 interval.

14 In step 910, the server module 110 sends the retry packets to the client module
 15 204 in response to the retry requests.

16 Fig. 10 is a procedure 1000 that defines activities performed by the client module
 17 204, and is the counterpart of the server-based activities of Fig. 9.

18 In step 1002, the client module 204 sends a join request to the server module 202,
 19 requesting that the server module 202 commence the transmission of the stream.

20 In step 1004, the client module 204 starts receiving the server module 202's burst
 21 at rate $(1+E)*Nominal$, and shortly thereafter, begins playing the resource information in
 22 the stream.

23 In step 1006, the client module 204 receives a join mcast instruction from the
 24 server module 202, which instructs the client module 204 to join the multicast stream 116
 25

1 enabled by the multicast functionality 112. The client module 204 also receives the last
 2 of the unicast burst, followed by a unicast stream at rate $E \cdot \text{Nominal}$.

3 In step 1008, the client module 204 responds to the server module 202's
 4 instruction by sending its own join request to the multicast functionality 112.

5 In step 1010, in response to the client module 204's request, the client module 204
 6 begins receiving multicast information from the multicast functionality 112. This event
 7 marks the end of the join interval.

8 In step 1012, the client module 204 determines the identity of any holes that have
 9 occurred in the stream due to the below-nominal stream transmission during the join
 10 interval. In this step 1012, the client module 204 sends a retry request to the server
 11 module 202 that asks the server module 202 to supply the missing packets.

12 In step 1014, the client module 204 receives the missing packets from the server
 13 module 202. This has the effect of filling in the holes created by the join interval. If the
 14 holes can be filled in prior to the consumption of these holes, then the procedure is
 15 deemed successful.

17 **C. Exemplary Variations**

18 The following section describes a number of variations to the systems and
 19 procedures set forth above.

20 *C.1. ACC Strategies in Environments with $E \geq 1$.*

21 Several of the examples developed above assumed that E was less than 1.0.
 22 However, in other cases, E can be set as equal to or greater than 1.0. This means that the
 23 communication channel can accommodate at least two nominal data rate streams at the
 24 same time.
 25

Fig. 11 shows a representation 1100 of a hybrid unicast/multicast ACC strategy in which $E \geq 1.0$. More specifically, as before, a unicast burst 1102 is followed by a multicast stream 1104. The burst 1102 occurs at a *MaxRate* level that is at least two times the nominal rate of the stream. A join interval 1106 separates the unicast burst 1102 from the multicast stream 1104. In the join interval, the server module 202 (of Fig. 2) can deliver unicast resource information at some rate, such as 0, or $E \cdot \text{Nominal}$ (for the constant data rate or capped VBR case). In the former case, the join interval will cause packets to be missed. These packets can be provided concurrently with the multicast stream 1104 at a rate of $E \cdot \text{Nominal}$. In the latter case (which applies to the depiction shown in Fig. 11), a unicast stream 1108 of $E \cdot \text{Nominal}$ can supply data throughout the join interval. Since the unicast stream now remains at a nominal rate or above in the join interval, there is potentially no loss of packets due to the join interval.

C.2. ACC Strategies in Environments with Variable Bit Rates

Several of the examples developed above assumed that the data rate was constant. However, in other cases, a variable bit rate (VBR) can be used, such as a capped variable bit rate. Delivery that is governed by a capped variable bit rate maintains the data rate so that it remains below the maximum rate level (i.e., the cap level) of the stream.

Fig. 12 shows a representation of a channel 1200 that supports a variable bit rate. Namely, at any given time, the steady-state transmission of resource information consumes an actual amount 1202 of a total allotted data rate *MaxRate*, leaving, at any given time, a variable extra capacity 1204. The ACC strategy can utilize the extra capacity 1204 to accomplish the transition between the first mode of delivery to the second mode of delivery. For example, in a first phase, the ACC strategy can pack as much additional information as it can into the unicast burst. In a second phase, the ACC strategy can pack as much unicast data as it can during the join interval. In a third phase,

1 the ACC strategy can pack as much retry data (either supplied from a unicast source or a
 2 multicast source) as it can, concurrently with its supply of the multicast stream.
 3 Consider, for instance, the last-mentioned case in more detail. The ACC strategy can
 4 actively monitor the amount of bandwidth consumed by the multicast stream,
 5 *ActualCurrentMulticastRate*. The extra capacity at any given time available for handling
 6 retries is $MaxRate - ActualCurrentMulticastRate$.

7 The ACC strategy can measure the available capacity in different ways. In one
 8 technique, any actor within the system 100 (such as a client module) can report the data
 9 rate that is currently being used to delivery resource information over a particular
 10 communication channel. In addition, any actor (such as a client module) can actively
 11 probe the available capacity by increasing the data rate at which it receives information,
 12 and noting when the communication channel becomes saturated.

13 The calculations described above in preceding sections can be modified
 14 appropriately to reflect a time-varying amount of excess capacity. The existing
 15 calculations nevertheless provide valuable insight, as they yield worst-case results (in
 16 which a stream always takes up its nominal bit-rate). If a system satisfies these
 17 constraints, it can satisfy any scenario that always provides at least $(1+E)*Nominal$ bit-
 18 rate. Additional bit-rate available for retries can only improve the performance of the
 19 described methods.

20 In still other scenarios, it is possible to vary the ceiling amount *MaxRate* of the
 21 communication channel.

22 In still other scenarios, the ACC strategy can be applied in environments that do
 23 not impose any *MaxRate* limitations.

1 *C.3. ACC Strategies for the Unicast Transmission of VOD Assets*

2 Several of the examples developed above assumed that the ACC strategy involved
 3 splicing together unicast delivery of resource information with multicast delivery of
 4 resource information. However, in other cases, other kinds of delivery modes can be
 5 combined together to yield the ACC behavior. Generally, the different modes will still
 6 deliver the same pieces of data, but may differ in various respects, such as by providing
 7 different timing profiles that govern the delivery of the packets, using different delivery
 8 mechanisms, and so forth.

9 Fig. 13 shows a system 1300 in which the server module 202 delivers VOD assets
 10 1302 to the client module 204 using a pure unicast mode of delivery. In the first delivery
 11 mode, the server module 202 bursts the VOD assets at an above-nominal data rate 1304.
 12 In the second delivery mode, the server module 202 sends the VOD assets at or below a
 13 nominal data rate 1306. (In an alternative environment, two different sources can provide
 14 the first and second delivery modes, respectively.) Note that T_{join} may be zero in the
 15 above-described case of pure unicast. (In an alternative pure unicast environment, T_{join}
 16 may not be zero.) The delivery is constrained at all times by a maximum data rate
 17 *MaxRate*. Again, Fig. 13 shows the case of constant data rates to facilitate discussion,
 18 but the system 1300 can also deliver VOD assets using a variable data rate scheme.

19 *C.4. Retry Implementation Variations*

20 Several of the examples developed above assumed that the ACC strategy involved
 21 using the server module 202 to handle retry requests from the client module 204. But
 22 other entities can perform this task. For example, the multicast functionality 112 can
 23 perform this task, some other multitask functionality (not shown) can perform this task
 24 (providing the retry packets in a separate multicast stream), or some other actor (not
 25 shown) can perform this task.

D. Exemplary Computer Environment

In one exemplary implementation, certain aspects of the system 100 shown in Fig. 1 can be implemented as computer code executed by computer devices. In this case, Fig. 14 provides information regarding an exemplary computer environment 1400 that can be used to implement any such computer devices. Namely, the acquisition functionality 104 can be implemented by a server type computer device, each of the server modules 110 can be implemented by a server type computer device, the configuration tool 118 can be implemented by any kind of computer device, and each of the computer modules 106 can be implemented as any kind of computer device, etc.

The computing environment 1400 includes a general purpose or sever-type computer 1402 and a display device 1404. However, the computing environment 1400 can include other kinds of computing equipment. For example, although not shown, the computer environment 1400 can include hand-held or laptop devices, set top boxes, game consoles, mainframe computers, etc. Further, Fig. 14 shows elements of the computer environment 1400 grouped together to facilitate discussion. However, the computing environment 1400 can employ a distributed processing configuration. In a distributed computing environment, computing resources can be physically dispersed throughout the environment.

Exemplary computer 1402 includes one or more processors or processing units 1406, a system memory 1408, and a bus 1410. The bus 1410 connects various system components together. For instance, the bus 1410 connects the processor 1406 to the system memory 1408. The bus 1410 can be implemented using any kind of bus structure or combination of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures.

1 Computer 1402 can also include a variety of computer readable media, including
2 a variety of types of volatile and non-volatile media, each of which can be removable or
3 non-removable. For example, system memory 1408 includes computer readable media in
4 the form of volatile memory, such as random access memory (RAM) 1412, and non-
5 volatile memory, such as read only memory (ROM) 1414. ROM 1414 includes an
6 input/output system (BIOS) 1416 that contains the basic routines that help to transfer
7 information between elements within computer 1402, such as during start-up. RAM
8 1412 typically contains data and/or program modules in a form that can be quickly
9 accessed by processing unit 1406.

10 Other kinds of computer storage media include a hard disk drive 1418 for reading
11 from and writing to a non-removable, non-volatile magnetic media, a magnetic disk drive
12 1420 for reading from and writing to a removable, non-volatile magnetic disk 1422 (e.g.,
13 a "floppy disk"), and an optical disk drive 1424 for reading from and/or writing to a
14 removable, non-volatile optical disk 1426 such as a CD-ROM, DVD-ROM, or other
15 optical media. The hard disk drive 1418, magnetic disk drive 1420, and optical disk drive
16 1424 are each connected to the system bus 1410 by one or more data media interfaces
17 1428. Alternatively, the hard disk drive 1418, magnetic disk drive 1420, and optical disk
18 drive 1424 can be connected to the system bus 1410 by a SCSI interface (not shown), or
19 other coupling mechanism. Although not shown, the computer 1402 can include other
20 types of computer readable media, such as magnetic cassettes or other magnetic storage
21 devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical
22 storage, electrically erasable programmable read-only memory (EEPROM), etc.

23 Generally, the above-identified computer readable media provide non-volatile
24 storage of computer readable instructions, data structures, program modules, and other
25 data for use by computer 1402. For instance, the readable media can store the operating

1 system 1430, application-specific functionality 1432, other program modules 1434, and
2 program data 1436.

3 The computer environment 1400 can include a variety of input devices. For
4 instance, the computer environment 1400 includes the keyboard 1438 and a pointing
5 device 1440 (e.g., a “mouse”) for entering commands and information into computer
6 1402. The computer environment 1400 can include other input devices (not illustrated),
7 such as a microphone, joystick, game pad, satellite dish, serial port, scanner, card reading
8 devices, digital or video camera, etc. Input/output interfaces 1442 couple the input
9 devices to the processing unit 1406. More generally, input devices can be coupled to the
10 computer 1402 through any kind of interface and bus structures, such as a parallel port,
11 serial port, game port, universal serial bus (USB) port, etc.

12 The computer environment 1400 also includes the display device 1404. A video
13 adapter 1444 couples the display device 1404 to the bus 1410. In addition to the display
14 device 1404, the computer environment 1400 can include other output peripheral devices,
15 such as speakers (not shown), a printer (not shown), etc.

16 Computer 1402 operates in a networked environment using logical connections to
17 one or more remote computers, such as a remote computing device 1446. The remote
18 computing device 1446 can comprise any kind of computer equipment, including a
19 general purpose personal computer, portable computer, a server, etc. Remote computing
20 device 1446 can include all of the features discussed above with respect to computer
21 1402, or some subset thereof.

22 Any type of network 1448 can be used to couple the computer 1402 with remote
23 computing device 1446, a LAN, etc. The computer 1402 couples to the network 1448 via
24 network interface 1450, which can utilize broadband connectivity, modem connectivity,
25 DSL connectivity, or other connection strategy. Although not illustrated, the computing

1 environment 1400 can provide wireless communication functionality for connecting
2 computer 1402 with remote computing device 1446 (e.g., via modulated radio signals,
3 modulated infrared signals, etc.).

4 In closing, a number of examples were presented in this disclosure in the
5 alternative (e.g., case A or case B). In addition, this disclosure encompasses those cases
6 which combine alternatives in a single implementation (e.g., case A and case B), even
7 though this disclosure may not have expressly mention these conjunctive cases in every
8 instance.

9 Moreover, a number of features were described herein by first identifying
10 exemplary problems that these features can address. This manner of explication does not
11 constitute an admission that others have appreciated and/or articulated the problems in
12 the manner specified herein. Appreciation and articulation of the problems present in the
13 video processing, compression and delivery arts are to be understood as part of the
14 present invention. More specifically, there is no admission herein that the features
15 described in the Background section of this disclosure constitute prior art.

16 Although the invention has been described in language specific to structural
17 features and/or methodological acts, it is to be understood that the invention defined in
18 the appended claims is not necessarily limited to the specific features or acts described.
19 Rather, the specific features and acts are disclosed as exemplary forms of implementing
20 the claimed invention.

51018-132

CLAIMS:

1. A method for expediting presentation of resource information upon a start of delivery of the resource information within a rate-limited environment, comprising:

5 delivering the resource information to a client module, in a first delivery, using first delivery functionality upon a commencement event; and

delivering the resource information to the client module, in a second delivery, using second delivery functionality,

10 wherein the resource information provided by the first delivery and the resource information provided by the second delivery combine together to provide a resultant stream of resource information in which a primary mode of resource information delivery transitions from the first delivery functionality to the second delivery functionality, wherein the resultant stream has a data rate that remains below a prescribed data rate limit, and

15 wherein retry functionality receives retry requests from the client module to supply missing parts of the resource information and delivers the missing parts of the resource information to the client module.

2. The method of claim 1, wherein the prescribed data rate limit provides an amount of excess capacity beyond an amount required to deliver the resource information in a steady-state condition.

3. The method of claim 1, wherein:

20 the first delivery functionality comprises unicast server-based functionality, and the first delivery of resource information provides a unicast stream of the resource information; and

25 the second delivery functionality comprises multicast functionality, and the second delivery of resource information provides a multicast stream of the resource information.

51018-132

4. The method of claim 1, wherein:

the first delivery functionality comprises unicast server-based functionality configured to deliver the resource information in a unicast stream at first data rate levels for an initial period of time; and

5 the second delivery functionality comprises the unicast server-based functionality configured to deliver the resource information in a unicast stream at second data rate levels following the initial period of time.

5. The method of claim 1, wherein:

10 the first delivery comprises delivering the resource information at first data rate levels for an initial period of time; and

the second delivery comprises delivering the resource information at second data rate levels.

6. The method of claim 5, wherein the first data rate levels are data rates which exceed a nominal data rate used to deliver the resource information by an excess amount, and
15 the second data rate levels are data rates which remains at or below the nominal data rate.

7. The method of claim 5, wherein there is a join interval that separates the delivery of resource information at the first data rate levels and the delivery of resource information at the second data rate levels, wherein, in the join interval, the first delivery of resource information is provided at join interval data rate levels.

20 8. The method of claim 7, wherein a length of the join interval is based on an amount of time required by the client module to couple to the second delivery functionality to receive the second delivery of the resource information.

9. The method of claim 7, wherein the join interval data rate levels are data rates below a nominal data rate used to deliver file resource information.

51018-132

10. The method of claim 9, wherein the delivery of the resource information at below-nominal join interval rates during the join interval causes parts of the first stream of resource information to be missed by the client module.

11. The method of claim 7, wherein a plurality of parameters govern the behavior
5 of the delivery in the rate-limited environment, and further comprising configuring one or more of the parameters to ensure that a combined amount of the first delivery of resource information and the second delivery of resource information is constrained so that it remains below the prescribed data rate limit.

12. The method of claim 1, wherein the resource information is media information.

10 13. The method of claim 1, wherein the commencement event is associated with a change by the client module from a previous state to a new state, wherein the new state involves a selection of a new program that provides the resource information to be consumed by the client module.

14. The method of claim 13, wherein the commencement event is associated with a
15 change by the client module from an old program, associated with the previous state, to the new program, associated with the new state.

15. The method of claim 2, further comprising:

measuring the amount of excess capacity; and

selecting data rate levels at which the resource information is delivered using
20 the first delivery functionality and/or the second delivery functionality based on the measured amount of excess capacity.

16. The method of claim 15, wherein the amount of excess capacity is measured by investigating the amount of degradation produced by transmission of resource information at different data rates.

25 17. The method of claim 4, wherein the resource information is video on demand (VOD) resource information.

51018-132

18. The method of claim 5, wherein the first data rate levels and the second data rate levels define constant respective data rates.

19. The method of claim 5, wherein the first data rate levels and the second data rate levels define varying respective data rates.

5 20. The method of claim 6, wherein the excess amount is less than the nominal data rate.

21. The method of claim 6, wherein the excess amount is equal to or greater than the nominal data rate.

22. The method of claim 6, wherein the excess amount varies depending on a time-
10 varying amount of excess capacity available.

23. The method of claim 6, wherein the excess amount is specified as a fraction E of the nominal data rate.

24. The method of claim 8, wherein the join interval is shortened to a minimum amount of time that is required by the client module to couple to the second delivery
15 functionality, wherein the first delivery functionality continues to deliver the resource information in a time period defined by that minimum amount of time.

25. The method of claim 7, wherein the join interval data rates dynamically vary based on a time-varying amount of excess capacity available to deliver resource information.

26. The method of claim 7, wherein the join interval data rate levels are data rates
20 at or above a nominal data rate used to deliver the resource information.

27. The method of claim 9, wherein the join interval data rate levels are data rates of 0.

28. The method of claim 9, wherein the join interval data rates are specified as a fraction E of the nominal data rate.

51018-132

29. The method of claim 11, wherein the parameters include at least the following parameters:

an excess amount of capacity that is available;

5 a length of time T_{burst} during which the first delivery functionality supplies the resource information at the first data rate levels; and

a length of time T_{join} of the join interval.

30. The method of claim 1, further comprising determining whether the client module has sufficient buffer resources to receive the resource information prior to delivery of the resource information using the first delivery functionality.

10 31. The method of claim 30, further comprising using a combination of the first delivery and the second delivery of the resource information in the event that the client module has sufficient buffer resources, and omitting use of the first delivery of resource information in the event that the client module does not have sufficient buffer resources.

15 32. One or more computer-readable storage medium having stored thereon computer executable instructions, that when executed by one or more processors implements the method of any one of claims 1 to 31.

33. A method for delivering resource information to a client module, comprising:

20 delivering the resource information, using first delivery functionality, at above-nominal data rate levels during an initial burst period of data transmission to the client module;

delivering the resource information at join interval data rate levels during a join interval following the burst period, wherein, following the join interval, the client module switches to second delivery functionality for providing the resource information; and

51018-132

responding to requests from the client module to supply parts of the resource information which the client module missed during the join interval as a result of its receipt of the resource information at the join interval data rate levels.

34. The method of claim 33, wherein the first delivery functionality and the second
5 delivery functionality comprise unicast delivery functionality and multicast delivery functionality, respectively.

35. The method of claim 33, wherein the first delivery functionality and the second delivery functionality both comprise unicast delivery functionality.

36. One or more computer-readable storage medium having stored thereon
10 computer executable instructions, that when executed by one or more processors implements the method of any one of claims 33 to 35.

37. A method for receiving a stream of resource information at a client module, comprising:

receiving, via first delivery functionality, resource information at above-
15 nominal data rate levels during an initial burst period of data transmission;

receiving the resource information at join interval data rate levels during a join interval following the burst period;

receiving, via second delivery functionality, the resource information following the join interval;

20 requesting parts of the resource information which the client module missed during the join interval as a result of its receipt of the resource information at the join interval data rate levels;

receiving the parts of resource information which the client module missed concurrently with the receipt of the resource information via the second delivery functionality;

25 and

51018-132

using the received missing parts for filling in holes created in the stream during the join interval.

38. The method of claim 37, wherein the first delivery functionality and the second delivery functionality comprise unicast delivery functionality and multicast delivery

5 functionality, respectively.

39. The method of claim 37, wherein the first delivery functionality and the second delivery functionality both comprise unicast delivery functionality.

40. One or more computer-readable storage medium having stored thereon computer executable instructions, that when executed by one or more processors implements

10 the method of any one of claims 37 to 39.




Fig. 1




Fig. 2




Fig. 3







Fig. 4

At Location X**At Location Y****At Location Z****Fig. 5**

SCENARIO 1-A: Sample data consumption graph for rate-limited ACC ($E=.5$, $T_{burst} = 6$, $T_{join} = 1$)





Fig. 6

SCENARIO 1-B: Sample data consumption graph for rate-limited ACC ($E=.25$, $T_{burst} = 6$, $T_{join} = 1$)



SCENARIO 2-A: Sample data consumption graph
for rate-limited ACC ($E=.5$, $T_{burst} = 6$, $T_{join} = 1$)





Fig. 7

SCENARIO 2-B: Sample data consumption graph
for rate-limited ACC ($E=.25$, $T_{burst} = 6$, $T_{join} = 1$)



SCENARIO 3-A: Sample data consumption graph for rate-limited ACC ($E=.5$, $T_{burst} = 6$, $T_{join} = 1$)





Fig. 8

SCENARIO 3-B: Sample data consumption graph for Multicast ICC ($E=.25$, $T_{burst} = 6$, $T_{join} = 1$)






Fig. 9




Fig. 10




Fig. 11




Fig. 12




Fig. 13




Fig. 14

