



(11) **EP 1 535 130 B9**

(12) **CORRECTED EUROPEAN PATENT SPECIFICATION**

(15) Correction information:  
**Corrected version no 1 (W1 B1)**  
**Corrections, see**  
**Claims DE**  
**Claims EN**  
**Claims FR**

(51) Int Cl.:  
**H04L 12/56 (2006.01)**

(86) International application number:  
**PCT/US2003/023513**

(87) International publication number:  
**WO 2004/015520 (19.02.2004 Gazette 2004/08)**

(48) Corrigendum issued on:  
**25.02.2009 Bulletin 2009/09**

(45) Date of publication and mention  
of the grant of the patent:  
**10.09.2008 Bulletin 2008/37**

(21) Application number: **03784828.0**

(22) Date of filing: **29.07.2003**

(54) **QUALITY OF SERVICE MANAGEMENT IN NETWORK GATEWAYS**

**DIENSTQUALITÄTSVERWALTUNG IN NETZWERK-GATEWAYS**

**GESTION DE LA QUALITE DE SERVICE DANS DES PASSERELLES DE RESEAU**

(84) Designated Contracting States:  
**DE FR GB IT**

• **KAMEL, Ibrahim**  
**Lawrenceville, NL 08643 (US)**

(30) Priority: **12.08.2002 US 217229**

(74) Representative: **Chapman, Paul Nicholas**  
**Atkinson Burrington Limited**  
**28 President Buildings**  
**President Way**  
**Sheffield**  
**S4 7UR (GB)**

(43) Date of publication of application:  
**01.06.2005 Bulletin 2005/22**

(73) Proprietor: **Panasonic Corporation**  
**Kadoma-shi**  
**Osaka 571-8501 (JP)**

(56) References cited:  
**EP-A- 0 948 168** **WO-A-99/62259**  
**US-A- 6 009 106** **US-A1- 2003 084 108**  
**US-A1- 2003 152 028** **US-A1- 2004 022 191**  
**US-A1- 2004 136 379** **US-B1- 6 307 839**  
**US-B1- 6 765 931**

(72) Inventors:  
• **AKINLAR, Cuneyt**  
**Dayton, NJ 08810 (US)**  
• **ELBASSIONI, Khaled**  
**New Brunswick, NJ 08901 (US)**

**EP 1 535 130 B9**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

**Description**

## FIELD OF THE INVENTION

5 **[0001]** The present invention relates to ensuring quality-of-service in networks, and in particular to ensuring quality-of-service for networks transmitting and receiving realtime and non-realtime data-streams.

## BACKGROUND OF THE INVENTION

10 **[0002]** Network users are able to access various types of information from the Internet and other sources. The type of information that the network users can access can be broadly divided into two categories: realtime streams and non-realtime streams. For example, a typical user will be receiving realtime data streams of video or audio and non-realtime data streams like e-mail, web pages, or File-Transfer Protocol (FTP) downloads. Realtime data-streams are generally required to be transmitted or processed within some small upper limit of time. Non-realtime data-streams are broadly understood to be not requiring processing or transmission within the time constraints such as those required for the realtime data-streams. Realtime and non-realtime data-streams have differing characteristics as described next.

15 **[0003]** The chief characteristics of realtime and non-realtime data-streams of relevance here are their respective bandwidth requirements for providing different levels of Quality-of-service (QoS). QoS is broadly understood as the set of performance properties of given a network service, generally including throughput, transit, delay and priority. In the present context of realtime streams, the additional QoS parameters include bandwidth availability, delay and jitter among other parameters. Those skilled in the art will appreciate that relevancy and importance of any given QoS parameters will depend upon the nature of the realtime data stream used in a particular application. The invention covers and supports any set of QoS parameters for a given realtime data stream. Realtime streams need a guaranteed QoS for providing relatively fast and time constrained information transmission. Non-realtime streams are generally transmitted using the transmission control protocol (TCP)/Internet Protocol (IP). Contrastingly, non-realtime streams do not generally require the QoS similar to that required for the realtime streams. A typical example of a network handling realtime and non-realtime data-streams is described next.

20 **[0004]** A network can be configured to receive both realtime and non-realtime data-streams from an external source. A single transmission channel generally links the user's network to the Internet service provider (ISP). The same transmission channel concurrently carries both the realtime and non-realtime streams. The bandwidth capacity of such a transmission channel generally remains fixed. Therefore, it becomes necessary to balance the allocation of available bandwidth between the conflicting demands made by the realtime and non-realtime streams. The problem of bandwidth allocation is illustrated next in the context of a typical user.

25 **[0005]** A network user is usually connected to a network like Internet through a service provider who may provide Internet-access and possibly other services like video-on-demand, IP telephony, streaming audio and video. The service provider is linked to the network user by a transmission channel like a dial-up telephone line, xDSL, ISDN, etc. The connecting device at the service provider's end may be an edge router, and at the network user end it would generally be a gateway.

30 **[0006]** Realtime data-streams require an almost fixed allocation of bandwidth. Realtime data-streams offer little flexibility in adjusting bandwidth requirements without compromising the QoS parameters. In contrast, the non-realtime data-streams are relatively flexible about their bandwidth demands, because they do not usually require a relatively strict QoS. Bandwidth availability may change over a given period of time. Therefore, the non-realtime stream traffic from the service provider to the network user needs to be controlled in order to ensure that the realtime streams get the required bandwidth for maintaining its QoS. Possible methods for controlling the sharing of bandwidths are considered next.

35 **[0007]** A conventional approach using a packet pacing method is discussed next. Non-realtime traffic transmitted from the router located at the service provider to the gateway will generally be the Internet communication traffic transmitted using the TCP protocol. The TCP sender at the Internet site controls the non-realtime traffic by pacing the non-realtime packets to ensure that the realtime traffic gets the required bandwidth. The packet pacing method and its associated problems are described next.

40 **[0008]** Packet pacing is generally performed by controlling the rate of packet transmission. Implementing such packet pacing method requires significant changes in the operations of a TCP sender. In a typical network user scenario the TCP sender, i.e., a HTTP server, is under control of an external agency like a university, hospital, or company. The ISP may not be expected to employ any particular bandwidth management techniques. An ISP typically will be servicing a large number of users in a situation where each one of the users has several active TCP connections operating at the same time. Such packet pacing approach is not feasible to implement at an ISP site due to scalability problems associated with supporting a large number of users. Thus, there is a need for an improved bandwidth management technique that is implemented at the gateway side of the network.

45 **[0009]** Another approach involves controlling the TCP traffic for the non-realtime streams from a conventional user

gateway. The difficulty with this approach is that the TCP-receiver at the user gateway has almost no operatively effective control over the TCP-sender, which is typically a Hypertext Transfer Protocol (HTTP) server or a FTP server. Hence, there is a need for an apparatus and method that allows controlling the non-real time traffic at the gateway end, and which is feasible in a TCP environment without using any special apparatus at the user end.

5 [0010] Above described known methods for bandwidth management in networks where realtime and non-realtime traffic share the available bandwidth of a channel have several drawbacks as described above. Thus, there is a need for a bandwidth management solution that allows controlling the non-realtime streams bandwidth demands so that the realtime streams can provide a desired QoS. Further, there is a need for implementing such a solution on the gateway located at the user's end of the network.

10 [0011] United States Patent 6,307,839 discloses a dynamic bandwidth allocation system used to optimise transmission over a twisted pair between an intelligent services director (ISD) at a customer premises and a facilities management platform (FMP) at a local office. Both the ISD and FMP have the capability to sense and seize available bandwidth and decide the optimal bandwidth allocation scheme for managing requested services. In an example, all available bandwidth is used for data transmission until a voice call is received. A necessary portion of the bandwidth is then reallocated from data usage to transmission of the voice call.

15 [0012] EP 0 948 168 A1 discloses a method of controlling the flow of data from a sender to a receiver over a packet exchange connection. One or both of the partners in the connection monitor bandwidth values of links in the connection, and the flow of data from the sender is controlled by employing the bandwidth values.

20 SUMMARY OF THE INVENTION

[0013] According to an aspect of the present invention there is provided an apparatus for ensuring quality-of-service in a network as claimed in claim 1.

25 [0014] A system for ensuring quality of service in a network is disclosed. The network uses a single communication channel sharing realtime and non-realtime transmissions, e.g. TCP traffic, that is connected to a gateway. The non-realtime streams are transmitted using non-realtime senders that have flow control parameters or windows. The gateway is further connected to a network including various network elements. The gateway includes a bandwidth control unit that controls the bandwidth demands of the non-realtime transmissions by adjusting the flow control parameter on the non-realtime senders. The realtime streams require consistent bandwidth to support quality of service parameters like delay and jitter. The bandwidth control regulates the non-realtime connections bandwidth requirement, and hence ensures the bandwidth required by realtime streams. The bandwidth control can also dynamically allocate bandwidth between multiple non-realtime TCP connections, so that the unused bandwidth available during a TCP slow-start of a given TCP connection can be allocated to other steady state TCP connection.

30 [0015] Further areas of applicability of the present invention will become apparent from the detailed description provided hereinafter. It should be understood that the detailed description and specific examples, while indicating the preferred embodiment of the invention, are intended for purposes of illustration only and are not intended to limit the scope of the invention.

40 BRIEF DESCRIPTION OF DRAWINGS

[0016] The present invention will become more fully understood from the detailed description and the accompanying drawings, wherein:

[0017] Figure 1 shows a network configuration for illustrating the invention's implementation of bandwidth management;

45 [0018] Figure 2 shows an exemplary network configuration having a single TOP sender and implementing the invention's bandwidth management;

[0019] Figure 3 is a graph showing the average bandwidth for the realtime and non-realtime streams in the absence of any bandwidth management;

[0020] Figure 4 shows the average inter-packet time for the VoD stream;

50 [0021] Figure 5 shows network configuration having multiple TCP connections and implementing bandwidth management;

[0022] Figure 6 is a graph showing the average bandwidth of each stream and the aggregate bandwidth of all streams; and

[0023] Figure 7 is a graph showing performance characteristics of dynamic bandwidth management.

55 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0024] The following description of the preferred embodiment(s) is merely exemplary in nature and is in no way intended to limit the invention, its application, or uses.

**[0025]** Figure 1 shows a network configuration 10 for illustrating bandwidth management. Bandwidth management mechanism employing the principle of the invention will be illustrated using the exemplary network configuration 10. Hence, the network configuration 10 is described next in detail. The constituent elements of the network configuration 10 are described next. A VoD server 12 and an Internet Protocol (IP) phone 14 are connected to a private network 16. The VoD server 12 provides video-on-demand transmissions and the IP phone 14 provides phone-like communication service using the IP protocol to the network user. A FTP server 18 and a HTTP server 20 are connected to an Internet server 22. Typically, the VoD server 12 and the IP phone 14 transmit in a realtime manner with strict time constraints. In contrast, the FTP server 18 and the HTTP server 20 transmit information in a non-realtime manner with relatively stringent time constraints. Those skilled in the art will appreciate that the following description of the network is only an illustration and that invention covers any suitable type of network configuration.

**[0026]** Internet server 22 and the private network 16 are both connected to an ISP edge router 24. An access segment 26 connects the edge router 24 to a gateway 28. Access segment 26 is a communication channel for transmitting and receiving data to and from said ISP edge router 24 and the gateway 28. Access segment 26 is generally a broadband connection like xDSL, ISDN or coaxial cable, but it can also be a dial-up telephone connection. Access segment 26 simultaneously carries both realtime and non-realtime streams transmitted via the edge router 24. Streams are logical channels of data-flow. Realtime streams carry realtime data for applications like video-on demand. Non realtime streams in the present context are generally TCP or similar logical communication exchange using an appropriate protocol. Those skilled in the art will appreciate that the term "streams" is used in a generally broad manner to indicate a sequence of data or information.

**[0027]** Realtime streams share the bandwidth of the same access segment 26 with the non-realtime TCP traffic, and hence bandwidth management methods or algorithms are required to apportion the available bandwidth between realtime and non-realtime streams. Such a bandwidth management method or algorithm should limit the incoming TCP traffic in such a manner that sufficient bandwidth out of the aggregate bandwidth is left for the realtime streams that have strict QoS requirements. Preferable characteristics of the edge router 24 are described next.

**[0028]** Edge router 24 is a connection point for the network user communicating with the service provider. Edge router 24 can be any general purpose device generally having the capability to forward packets from the Internet hosts to the gateway 28. Edge router 24 must be able to transmit and receive IP packets between the gateway 28 and the Internet server 22. Hence, any router providing such service can be used here as the edge router 24. Edge router 28 may have other capabilities, e.g. ability to multiplex Internet and other traffic, but additional capabilities are not relevant for the present invention. The only capability that is relevant here is the ability to transmit and receive IP packets from the Internet hosts to the gateway 28. Next, the features of the data-streams carried over the access segment 26 are described.

**[0029]** The realtime media streams may be transmitted as IP or non-IP traffic. One of the characteristics of the realtime streams that is considered relevant here is that they are packetized, i.e., sent in packets, and have stringent time constraints, and any packet delays are detrimental to their performance characteristics. Realtime streams carried over the access segment 26 have strict QoS requirements such as sufficient bandwidth, minimal delay and minimal or no jitter. The media streams from the VoD server 12 and the IP phone 14 are merely examples of any realtime streams, and the invention is not limited by the number or types of realtime streams of information. The modalities of transmitting realtime streams are not considered here as they may vary across configurations. The principle of the invention encompasses any type of realtime transmission having certain definite bandwidth requirements necessary to ensure a given set of QoS parameters. The preferable network location for implementing the bandwidth management method is the gateway 28. Preferable characteristics of the gateway 28 are described next.

**[0030]** Gateway 28 connects the access segment 26 and the network 30. Network 30 may be constructed by using multiple networking technologies, for example, IEEE 1394, Ethernet, 802.11 wireless LAN, and powerline networks. Gateway 28 can be a conventional gateway or a router installed as a separate unit. The gateway 28 is an interconnection point in the network for connecting the network 30 to the edge router 24. The gateway 28 can also be integrated in an appliance such as digital television or a set-top-box. The function of the gateway 28 is to provide a point of connection between the Internet connection on its one interface and the network 30 connected components on its other interface. Gateway 28 serves as a connection point for a network 30 to the edge router 24. The network 30 and associated network elements are described next.

**[0031]** A range of network elements can be connected to the gateway 28 through the network 30. Various devices like television 32, computers 34 and IP phones 14 can be connected to the network 30. Those skilled in the art will appreciate that the network elements shown here are used only to illustrate the type of devices that can be connected to the network 30. Above listed devices or appliances are merely illustrations and many other devices can also be connected to the network 30.

**[0032]** Figure 2 shows an exemplary network configuration having a single TCP sender and implementing bandwidth management. The invention can be better understood by those in skilled in the art from a comparison between a network with bandwidth management and a network without bandwidth management. Bandwidth management for a network having a single TCP sender is described in two steps. In the first step, the performance of the network is simulated and

analyzed while assuming that no bandwidth management is performed. Such a simulation provides a background for making a comparison between the network without and with bandwidth management. In the second step, the same network is simulated and analyzed, but with a bandwidth control being used to implement the invention's principle for a single TCP sender network. Therefore, first the network shown in Figure 2 is simulated and analyzed assuming that

there is no bandwidth control as described below.

**[0033]** The following description establishes the need for bandwidth control. An assumption is made that the bandwidth control shown in the figure 2 does not exist in order to provide a comparison further-on in the description below. The network described here is used to simulate the performance characteristics of a typical network that does not use any bandwidth management. VoD server 12, HTTP server 20, Internet server 22, edge router 24 and the gateway 28 are interconnected as described in the context of figure 1.

**[0034]** In the present network the access segment 26 is a dedicated asymmetric link, for example, an ADSL link, with a downstream bandwidth of 2.0 Mbps and an upstream bandwidth of 0.7 Mbps. The delay in context of the access segment 26 is a negligible 1 ms, since typically the edge router 24 and the gateway 28 would be relatively close to each other.

**[0035]** VoD server 12 is connected to the edge router 24 by a full-duplex VoD link 36 having a 20 ms delay. VoD server 12 transmits video signal having a constant bit rate (CBR) in a realtime manner at the rate of 1.0 Mbps. The HTTP server 20 is configured as a constituent part of the Internet server 22 (as shown) or it may be externally connected (as shown in figure 1) to the Internet server 22. The HTTP server 20 transmits non-realtime data packets over a full duplex HTTP link 38 having 1.5 Mbps bandwidth and a 20 ms delay.

**[0036]** Edge router 24 includes a first-in-first-out (FIFO) queue 40 having a capacity to hold 20 packets. Realtime stream from the VoD server 12 requires 1 Mbps bandwidth from the aggregate 2 Mbps downward capacity 42 of the access segment 26. As a result, the HTTP server 20 can transmit packet traffic that uses up to 1.0 Mbps maximum capacity for non-realtime traffic directed toward the gateway 28.

**[0037]** Figures 3 and 4 deal with a network simulation for a network using no bandwidth control for management bandwidth requirements of non-realtime streams. First, some foundational information about the simulation technique is described below.

**[0038]** A network simulator is used to perform simulations for comparing performance with and without bandwidth management. Any conventional network simulator capable of performing simulation as described next may be used. For example, the 'ns UCB/LBNL/VINT' network simulator can be used. The VoD Stream is modeled by a CBR source sending data using the user datagram protocol (UDP) with a packet size of 576 bytes. A TCP connection for non-realtime data is modeled using a TCP/Reno source at the HTTP server 20 and a TCP sink at the gateway 28. The maximum segment size is set to 576 bytes, and the size of an initial flow control window 46 is set to 32 packets. The TCP flow control window 46 sizes are 16 KB or 32 KB for most operating systems. Hence, the HTTP server 20 always sends TCP packets of size 576 bytes, and does not have more than 64 unacknowledged TCP segments in transit. In the present context the TCP sender in the description below would mean the HTTP server 20, which transmits the HTTP download 52 to the gateway 28. Following the foundational information for network simulation, the specific simulations are described below in context of the appropriate figures.

**[0039]** Figure 3 is a graph showing the average bandwidth for the realtime and non-realtime streams in the absence of any bandwidth management. Time measured in seconds is plotted on the X-axis of the graph, and average bandwidth measured in Mbps is plotted on the Y-axis. The average bandwidth is calculated at the gateway 28 over a period of 0.1 seconds. The graph clearly shows that the VoD stream 50 is not able to receive a consistent bandwidth of 1 Mbps, which is bandwidth required for the realtime VoD stream 50 to satisfy the QoS parameters like delay and jitter.

**[0040]** Further, the HTTP download 52 also shows chaotic behavior due to packet drops at the edge router 24 (see figure 2). Whenever the HTTP server 20 (see figure 2) starts pumping the HTTP download 52 requiring more than 1 Mbps bandwidth, the edge router 24 starts dropping packets from both the realtime and non-realtime streams. This problem occurs at 1.40 seconds, at 3.5 seconds and then repeats itself periodically.

**[0041]** Just as the HTTP download 52 starts losing packets at 1.40 seconds, the TCP congestion window 48 (see figure 2) gets smaller to adjust for the congestion. This causes the HTTP server 20 to reduce its rate of packet transmission. Such a rate reduction for the HTTP download 52 is a positive factor for the VoD stream 52, but the access segment 26 bandwidth remains under-utilized since the reduced HTTP download 52 leaves some bandwidth unused. The under-utilization of bandwidth due to reduced HTTP download 52 continues till the TCP sender recovers its lost packets and increases its transmission rate. Once the TCP sender fully recovers and starts transmitting above 1 Mbps limit at around 3.5 seconds the edge router 24 again drops packets from both streams causing the same behavior that occurred at 1.4 seconds as described above and this cycle continues till the end of simulation.

**[0042]** Figure 4 shows the average inter-packet time for the VoD stream 50. Time measured in seconds is plotted on the X-axis of the graph, and average inter-packet time measured in Mbps is plotted on the Y-axis.

**[0043]** Jitter is an undesirable characteristic in a realtime transmission. The average inter-packet time is a good indicator of jitter. Variable inter-packet time leads to more jitter. As soon as the HTTP server 20 starts pumping HTTP

download 52 (see figure 3) above the 1 Mbps level, the VoD stream 52 (see figure 3) packets are delayed in the FIFO queue 40 (see figure 2), and thus causing the inter-packet time to increase. Later when the HTTP server 20 detects the packet drops and reduces its rate of packet transmission, the packets that have been queued in the FIFO queue 40 are transmitted in quick succession leading to decrease in the inter-packet time. Both of these cases of increase or decrease in the inter-packet time are undesirable and either cause underflow or overflow at the other end on the gateway 28. Ideally, the inter-packet time should remain constant at the level shown by the no-jitter line 54. The sections where jitter occurs during the transmission and causes problems is shown by jitter-regions 56. In the second step as referred to above, the full network as shown in figure 2 is now discussed below including bandwidth management for a single TCP sender network using the principle of the invention.

**[0044]** A network configuration that uses realtime and non-realtime transmission over a single transmission channel faces bandwidth allocation problems as described above. The description below is in the context of a single transmission channel, but those skilled in the art will appreciate that the invention can operate over multiple transmission channels also.

**[0045]** Bandwidth management is required to adhere to the QoS requirements of realtime streams. Bandwidth management also improves the overall channel utilization of the access segment 26 and the throughput of the non-realtime network traffic. Bandwidth management ensures the QoS criteria for realtime streams. Bandwidth management requires a choice to be made of a location in the network for implementing the bandwidth control methods. Gateway 28 is the present invention's preferred location in the network for implementing bandwidth control.

**[0046]** The bandwidth management technique of the present invention for a single TCP sender will be described next while referring back to the figure 2. In a typical network setting, a TCP sender like the HTTP sever 20 would not know in advance the available bandwidth in the path of its transmission to a TCP receiver. The TCP receiver like the gateway 28 of the present embodiment would have that knowledge of available path bandwidth. Here, the gateway 28 knows in advance that the TCP traffic should not exceed 1 Mbps, because the realtime traffic needs an assured bandwidth of 1 Mbps from the overall access segment 26's downward capacity of 2.0 Mbps.

**[0047]** The gateway 28 uses its knowledge of bandwidth requirements of the realtime and non-realtime streams to control the non-realtime, i.e., TCP traffic, coming from the HTTP server 20 so that the TCP traffic does not exceed the bandwidth available for non-realtime streams. Hence, the realtime streams are able to satisfy the required QoS criteria.

**[0048]** The bandwidth control 60 makes it possible to ensure the QoS requirements for the realtime streams are satisfied by controlling the data flow from the gateway 28 end. The bandwidth control 60 can be implemented in hardware, as a software module or as a combination of hardware and software. Controlling the flow of non-realtime traffic from the gateway 28 end eliminates the possible scalability problems associated with the solutions that control traffic from the edge router 24 at the Internet service provider side.

**[0049]** If a bandwidth management solution is employed at the edge router 24 end then a separate protocol is required to coordinate the bandwidth negotiation process between the gateway 28 and the edge router 24 for each realtime and non-realtime traffic stream. Implementing the bandwidth control 60 not the gateway eliminates this coordination problem. The details of how the bandwidth is managed from the gateway 28 are described next.

**[0050]** The description next refers to a single TCP connection, and then later-on multiple TCP connections are considered. A TCP sender, i.e. here the HTTP server 20, sends 'wnd' number of packets to the TCP receiver, i.e., here the gateway 28, within each round trip time ("*rtt*") segment. The 'wnd' number of packets to be sent in each *rtt* segment is calculated as  $wnd = \min \{cwnd, fwnd\}$ , which is the active window size. The TCP sender always maintains two variables or parameters called "*cwnd*" and "*fwnd*". The *cwnd* parameter represents a congestion control window and is computed by the TCP sender based upon the received acknowledgement and packet drops. The *cwnd* parameter is strictly controlled by the TCP sender. The *fwnd* parameter represents the flow control window and is set by the TCP receiver.

**[0051]** The data rate ("*b*") for a TCP connection within a *rtt* segment is given by  $b = \frac{wnd}{rtt}$ . Considering the slow start phase in a given TCP connection, if the connection starts at time  $t^0$  as measured by the sender's clock, the sender will transmit packets, at any time  $t > t^0$  assuming no packet drops and given by:

$$f(t) = \min \{fwnd(t), g(t)\} ,$$

where

$$g(t) \stackrel{\text{def}}{=} 2^{\left\lfloor \frac{t-t_0}{m} \right\rfloor} \quad \text{[Equation no. 1]}$$

If packet drops are taken into consideration then more complex throughput formulas can be derived by known methods.

**[0052]** A "steady state" connection is one for which the number of packets transmitted is completely determined by the *fwnd* parameter. For a given connection to be in the steady state, value of the *cwnd* must be greater than the value of the *fwnd* parameter. In the steady state the TCP sender's output can be completely controlled by the TCP receiver and is shown by the above equation no. 1.

**[0053]** Gateway 28 controls the non-realtime traffic by manipulating the size of the flow control window *fwnd<sub>i</sub>*, which is located within the TCP sender at the other end. In the present illustration the non-realtime traffic, i.e., TCP traffic, should not exceed the 1 Mbps limit. If the maximum TCP segment size is 576 bytes, and the round trip time between the gateway 28 and the HTTP server 20 is 47 ms, which is obtained through simulation, then the gateway 28 sets the

flow control window size, i.e., the value of the *fwnd<sub>i</sub>* to  $\frac{1\text{Mbps} * 47\text{ms}}{576 * 8\text{bits}} = 10 \text{ packets}$ . Setting the value of *fwnd<sub>i</sub>*

to 10 packets ensures that regardless of the congestion window size, the maximum data rate of the TCP connection can never exceed 1 Mbps.

**[0054]** Figure 5 shows a network configuration having multiple TCP connections and implementing bandwidth management. The configuration of the network shown is conceptually and structurally similar to that described in context of the figure 2 above. The network in figure 5 includes an additional TCP sender in the form of the FTP server 18 that transmits information requiring 1.5 Mbps bandwidth and has a delay of 20 ms. The bandwidth control 60 is used to manipulate the flow control from the gateway 28.

**[0055]** Figure 6 shows the average bandwidth of each stream along with the aggregate bandwidth of all streams, i.e., realtime and non-realtime combined. At time 0, the VoD server 12 starts pumping realtime CBR video at the rate of 1 Mbps. At time 1 second, a network user starts a FTP download from the FTP server 18. At around 3 seconds the network user starts a webpage download from the HTTP server 20. It is assumed that the webpage to be downloaded has multiple items like images and audio-clips to be downloaded as parts of the webpage. The web-browser (not shown) starts four simultaneous TCP connections, i.e., HTTP downloads 52<sub>a</sub>, 52<sub>b</sub>, 52<sub>c</sub>, 52<sub>d</sub>, to download the webpage and the associated multiple items. The four HTTP downloads 52<sub>a</sub>, 52<sub>b</sub>, 52<sub>c</sub>, 52<sub>d</sub> finish around 4.5 seconds, and the simulation is terminated at 5 seconds.

**[0056]** VoD stream 50 clearly achieves a sustained 1 Mbps regardless of the number of active TCP connections. The bandwidth control 60 reduces the data rate for the FTP download 58 when the HTTP downloads 52 start around 2 seconds. Bandwidth control 60's reduction of the data rate ensures that the aggregate non-realtime traffic never exceeds the available 1 Mbps bandwidth for the non-real time traffic. Bandwidth control 60 adjusts the individual non-realtime data connections so that the realtime streams receive the guaranteed bandwidth sufficient to service its QoS requirements. Thus, the bandwidth control 60 adjusts the aggregate non-realtime bandwidth by manipulating the individual flow control windows on the several TCP senders.

**[0057]** Bandwidth management technique using the principle of the present invention for multiple TCP connections is described next. To illustrate a set N of *n* non-realtime connections is considered. Each non-realtime connection is typically a HTTP or FTP connection. Let *rtt<sub>i</sub>* be the estimate of the round-trip time of a given connection *i*. The *rtt<sub>i</sub>* is calculated as described next.

**[0058]** Gateway 28 makes an initial estimate of the time required to get an acknowledgement at the time of setting up the connection. Let R be the set of realtime streams and *B<sub>i</sub>* be the bandwidth required by a given realtime stream *i* ∈ R, where the streams in R require constant bit rate. The above described parameters *n*, *i*, *rtt<sub>i</sub>*, *B<sub>i</sub>* and the sets N and R are assumed to be functions of time and will be denoted as such, if necessary.

**[0059]** The goal is to maximize the throughput  $b_i = \frac{wnd_i}{rtt_i}$  for each connection, since the TCP sender *i* sends *wnd<sub>i</sub>* = *min* {*cwnd<sub>i</sub>*, *fwnd<sub>i</sub>*}. The throughput maximization is subject to the inequality given below:

$$\sum_{i \in N} \frac{wnd_i}{rtt_i} \leq B \quad \text{[Equation no. 2]}$$

5 where,

$$10 \quad B = B_c - \sum_{i \in R} B_i$$

and where  $B_c$  is the total capacity of the access segment 26.

15 **[0060]** If the connections are all identically important, then the steady state flow control window size for each  $i$ , subject to the equation no. 2 is given by the conservative bound as given by the equation below:

$$20 \quad fwnd_i = \frac{B \times rtt_i}{n}$$

**[0061]** A static scheduling point is defined as a point in time at which either a new connection is established or an existing connection is terminated. The static bandwidth allocation procedure or algorithm is as shown below: for each static scheduling point  $t$  do the following

25

if the operation is a non-realtime connection-add then

set  $n \leftarrow n + 1$

30

else if the operation is a non-realtime connection-drop then

set  $n \leftarrow n - 1$

35

else // realtime add or drop

re-calculate  $B$  as given by  $B = B_c - \sum_{i \in R} B_i$

40

endif

for  $i = 1$  to  $n$

45

set  $fwnd_i \leftarrow \frac{B \times rtt_i}{n}$

set the flow-control window to  $fwnd_i$  in the first acknowledgement

sent to the sender of the connection  $i$

50

endfor

endfor

55

// END

**[0062]** Figure 7 is a graph showing performance characteristics of dynamic bandwidth management. The algorithm described in paragraph [0046] (hereafter called "the algorithm") can be further improved as described next. The algorithm

works in a static manner and limits the aggregate non-realtime TCP traffic bandwidth for ensuring QoS guarantees for the realtime traffic. The algorithm is invariant with respect to the number of non-realtime connections. Further improvements to performance of the non-realtime connections and to the total channel utilization are possible by using dynamic rather than static bandwidth allocation. Dynamic bandwidth allocation techniques of the present invention are described next with an illustration.

**[0063]** To illustrate the improvement, the following table is used as an example:

Period	1	2	3	4	5	...
TCP-1 <i>fwnd</i> with static bw allocation	16	8	8	8	8	...
TCP-1 <i>fwnd</i> with static bw allocation	-	1	2	4	8	...
Extra (unused) BW	0	7	6	4	0	...
TCP-1 <i>fwnd</i> with dynamic bw allocation	16	15	14	12	8	...

First we consider the algorithm operation. Initially there is only one TCP connection with a round-trip-time of 1 second. If the available capacity of the access segment 26 (see figure 5) is 16 packets/second then it is fully used by the first TCP connection since it is the only one. At the beginning of the second period another TCP connection arrives that has a round-trip-time of 1 second. According to the algorithm the available bandwidth is split among the first and second TCP connections with each connection getting 8 packets/second. However, the second TCP connection does not immediately start its share of 8 packets/second, because of the TCP slow start. The second TCP connection sends only 1, 2 and 4 packets in periods 2, 3, and 4 respectively before reaching the steady state rate of 8 packets/second. The static bandwidth allocation does not compensate for the TCP slow start mechanism. Thus, with static bandwidth allocation implemented using the algorithm there remains the unused bandwidths of 7, 6 and 4 packets in the periods 2,3 and 4 respectively.

**[0064]** Considering the previous example, the first TCP connection will be allocated the unused bandwidth till the second TCP connection achieves a steady state. Therefore, the first TCP connection will send 15, 14 and 12 packets during the periods 2, 3 and 4 respectively. The second TCP connection reaches steady state in period 5, and then uses all of its allocated 8 packets and hence the first TCP connection also uses its allocated 8 packets.

**[0065]** Simulation of dynamic allocation of bandwidth using the bandwidth control 60 is shown. The network used to simulate the system is the same as shown in Figure 5. Only the bandwidths for the non-realtime streams HTTP download 52 and FTP download 58 are shown, because the realtime stream shows the same performance as in case of static bandwidth allocation. Both static and dynamic bandwidth control ensure that the QoS requirements of the realtime streams are met. Dynamic bandwidth allocation improves the throughput, and hence performs better than the static bandwidth allocation.

**[0066]** In the simulation shown there is only one opportunity around 3 seconds for the dynamic allocation to take effect. The FTP download with dynamic allocation 58<sub>a</sub> is already in steady state when the four connections, i.e., HTTP downloads 52<sub>a</sub>, 52<sub>b</sub>, 52<sub>c</sub>, 52<sub>d</sub>, for the HTTP download 52 are started. In case of the static bandwidth allocation as shown by the algorithm, the bandwidth control will immediately distribute the available bandwidth among all the active connections. But in the case of dynamic allocation, the fact that the recently initiated four HTTP connections would be in a slow start mode is used to allocate the unused bandwidth available during the slow start of the HTTP download to the FTP download, which is already in a steady state. The FTP download performance is improved as seen by the shifting of the FTP download with dynamic allocation 58<sub>a</sub> to the right around 3 seconds. This can be compared in the graph to the plot for FTP download without dynamic allocation 58<sub>b</sub>. Therefore the dynamic allocation improves the utilization of the aggregate available bandwidth. The data rate for the FTP connections is gradually reduced to the steady state rate as the HTTP download 52 reaches the steady state. Preceding is the description of the dynamic bandwidth management. Below is the further description of the above referred bandwidth control 60.

**[0067]** The bandwidth control 60 (see figure 5) can be designed to work with a dynamic bandwidth allocation algorithm instead of the algorithm described above. The dynamic bandwidth allocation method achieves improved performance by allocating the unused bandwidth to the TCP connection that is already in steady state.

**[0068]** A particular application of the present invention is described in the context of a home network user. All above figures are used to provide context the description of the invention in context of the home user. The home network user is typically connected to the Internet through a home gateway, which is a specific type of gateway 28. The user is connected to the other services like video-on-demand and IP telephony through the same home gateway. The home network 30 (see figure 1) that connects to the home or residential gateway can be connected to a wide variety of devices like computers, televisions, telephones, radios etc.

**[0069]** The above described problems bandwidth management are present in the home user scenario, because it

would be difficult to implement bandwidth management techniques at the Internet service provider end. The home user would normally not have any control over the Internet service provider's mechanism of implementing TCP connections. Hence, it becomes necessary to implement bandwidth management at the residential or home gateway.

**[0070]** The principle of the present invention is used to incorporate a bandwidth control 60 into the home gateway. The operation of the bandwidth control is described above in detail. Above description applies equally to a home network user. In particular, the home user will be typically sharing the communication channel of access segment 60 for both realtime and non-realtime TCP traffic as the home user may find it expensive to use dedicated channels for realtime datastreams. Hence, the invention is beneficial to the home user using a shared channel for accessing realtime and non-realtime data.

## Claims

1. An apparatus for ensuring quality-of-service in a network, said apparatus comprising:

at least one first stream (18, 20) sender having a flow control parameter and operable to transmit a first stream therefrom in accordance with the flow control parameter; and  
a network interconnection (28) adapted to receive said first stream (38) over the network from said first stream sender and at least one second stream (36) transmitted over the network from a second stream sender (12),

### characterised in that:

said flow control parameter specifies an amount of data which can be sent before a specified event occurs; and  
said apparatus further comprises a bandwidth control (60) associated with said network interconnection and operable to adjust the flow control parameter of the first stream sender to meet a performance parameter associated with the second stream (36).

2. The apparatus of claim 1 further comprising:

a first network connection including said first stream sender; and  
a second network including said network interconnection.

3. The apparatus of claim 2 wherein said network interconnection being a home gateway and said second network being a home network.

4. The apparatus of claim 2 further comprising:

at least one channel connecting said first stream sender and said second stream sender to said network interconnection, said channel having a bandwidth capacity shared by said first stream and said second stream.

5. The apparatus of claim 2 wherein said first stream being a non-realtime stream and said second stream being a real-time stream, said second stream consistently requiring a part of said bandwidth capacity.

6. The apparatus of claim 5 wherein said bandwidth control adjusting said flow control parameter so that said first stream using a share of said bandwidth capacity that is less than or equal to the difference between said bandwidth capacity and the bandwidth requirements of said second stream.

7. The apparatus of claim 1 wherein said performance parameter being selected from a set of predetermined quality-of-service parameters associated with said second stream.

8. The apparatus of claim 7 wherein said bandwidth control maintaining quality-of-service parameters for said second stream by adjusting said flow control parameter and controlling bandwidth requirements of said first stream.

9. The apparatus of claim 1 wherein said flow control parameter regulates the flow of said first stream.

10. The apparatus of claim 1 wherein said flow control parameter regulating the bandwidth usage of said first stream by regulating the flow of said first stream.

11. The apparatus of claim 1 wherein said first stream sender operates in accordance with Transmission Control Protocol (TCP) and said flow control parameter comprises a flow control window as defined by TCP.
- 5 12. The apparatus of claim 1 wherein said network interconnection is a device chosen from a group consisting of routers, protocol converters and gateways.
- 10 13. The apparatus of claim 11 wherein the network interconnection dynamically allocates bandwidth amongst realtime streams and non-realtime streams so that bandwidth allocated to non-realtime streams does not exceed the bandwidth available to the non-realtime streams.
14. The apparatus of claim 13 wherein the network interconnection dynamically allocates bandwidth by increasing flow control window of a non-realtime stream that is in a steady state during a slow start period of another non-realtime stream.
- 15 15. The apparatus of claim 14 wherein the network interconnection dynamically allocates bandwidth by decreasing flow control parameter on the non-realtime stream when the another non-realtime stream achieves a steady state.

### Patentansprüche

- 20 1. Vorrichtung zur Sicherstellung einer Dienstqualität in einem Netzwerk, wobei die Vorrichtung aufweist:

mindestens einen ersten Datenstromsender (18, 20), der einen Strömungssteuerungsparameter aufweist und betrieben werden

25 kann, um einen ersten Datenstrom davon entsprechend dem Strömungssteuerungsparameter zu senden, und eine Netzwerkverbindung (28), die angepasst ist, um den ersten Datenstrom (38) über das Netzwerk von dem ersten Datenstromsender und mindestens einen zweiten Datenstrom (36) zu empfangen, der über das Netzwerk von einem zweiten Datenstromsender (12) übertragen wird,

30 **dadurch gekennzeichnet, dass**

der Strömungssteuerungsparameter eine Datenmenge spezifiziert, die gesendet werden kann, bevor ein spezifiziertes Ereignis auftritt, und

35 die Vorrichtung weiterhin eine Bandbreitensteuerung (60) aufweist, die der Netzwerkverbindung zugeordnet ist und betrieben werden kann, um den Strömungssteuerungsparameter des ersten Datenstromsenders einzustellen, um einen dem zweiten Datenstrom (36) zugeordneten Leistungsparameter zu erfüllen.

- 40 2. Vorrichtung gemäß Anspruch 1, weiterhin umfassend:

eine erste Netzwerkverbindung, die den ersten Datenstromsender enthält, und ein zweites Netzwerk, das die Netzwerkverbindung enthält.

- 45 3. Vorrichtung gemäß Anspruch 2, wobei die Netzwerkverbindung eine Heimschnittstelle ist und das zweite Netzwerk ein Heimnetzwerk ist.

- 50 4. Vorrichtung gemäß Anspruch 2, weiterhin umfassend:

mindestens einen Kanal, der den ersten Datenstromsender und den zweiten Datenstromsender mit der Netzwerkverbindung verbindet, wobei der Kanal eine Bandbreitenkapazität aufweist, die von dem ersten Datenstrom und dem zweiten Datenstrom geteilt wird.

- 55 5. Vorrichtung gemäß Anspruch 2, wobei der erste Datenstrom ein Nicht-Echtzeit-Strom und der zweite Datenstrom ein Echt-Zeit-Strom ist, wobei der zweite Datenstrom ständig einen Teil der Bandbreitenkapazität benötigt.

6. Vorrichtung gemäß Anspruch 5, wobei die Bandbreitensteuerung den Strömungssteuerungsparameter so einstellt, dass der erste Datenstrom einen Teil der Bandbreitenkapazität verwendet, der kleiner oder gleich der Differenz zwischen der Bandbreitenkapazität und dem Bandbreitenbedarf des zweiten Datenstroms ist.

7. Vorrichtung gemäß Anspruch 1, wobei der Leistungsparameter aus einem Satz von vorbestimmten Dienstquali-

tätsparametern ausgewählt wird, die dem zweiten Datenstrom zugeordnet sind.

- 5 8. Vorrichtung gemäß Anspruch 7, wobei die Bandbreitensteuerung Dienstqualitätsparameter für den zweiten Datenstrom beibehält durch eine Einstellung des Strömungssteuerungsparameters und eine Steuerung des Bandbreitenbedarfs des ersten Datenstroms.
9. Vorrichtung gemäß Anspruch 1, wobei der Strömungssteuerungsparameter den Fluss des ersten Datenstroms reguliert.
- 10 10. Vorrichtung gemäß Anspruch 1, wobei der Strömungssteuerungsparameter die Bandbreitensteuerung des ersten Datenstroms durch eine Regulierung des Flusses des ersten Datenstroms reguliert.
11. Vorrichtung gemäß Anspruch 1, wobei der erste Datenstromsender entsprechend einem Transmission-Control-Protocol (TCP) arbeitet und der Strömungssteuerungsparameter ein Strömungssteuerungsfenster aufweist, wie es durch das TCP definiert wird.
- 15 12. Vorrichtung gemäß Anspruch 1, wobei die Netzwerkverbindung ein Gerät ist, das aus einer Gruppe ausgewählt ist, die aus Routern, Protokollwandlern und Schnittstellen besteht.
- 20 13. Vorrichtung gemäß Anspruch 11, wobei die Netzwerkverbindung Bandbreite dynamisch zwischen Echt-Zeit-Strömen und Nicht-Echtzeit-Strömen zuteilt, so dass die den Nicht-Echtzeit-Strömen zuteilte Bandbreite nicht die Bandbreite übersteigt, die für die Nicht-Echtzeit-Ströme verfügbar ist.
- 25 14. Vorrichtung gemäß Anspruch 13, wobei die Netzwerkverbindung Bandbreite dynamisch zuteilt, indem das Strömungssteuerungsfenster eines Nicht-Echtzeit-Stroms vergrößert wird, der sich in einem stationären Zustand befindet, während einer langsamen Startperiode eines anderen Nicht-Echtzeit-Stroms.
- 30 15. Vorrichtung gemäß Anspruch 14, wobei die Netzwerkverbindung Bandbreite dynamisch zuteilt durch Verringern des Strömungssteuerungsparameters auf dem Nicht-Echtzeit-Strom, wenn der andere Nicht-Echtzeit-Strom einen stationären Zustand erreicht.

## Revendications

- 35 1. Dispositif pour assurer une qualité de service dans un réseau, le dispositif comprenant :
- au moins un premier émetteur (18, 20) de flux ayant un paramètre de réglage de débit et pouvant fonctionner pour transmettre un premier flux en fonction du paramètre de réglage de débit; et  
une interconnexion (28) de réseau conçue pour recevoir le premier flux (38) sur le réseau à partir du premier  
40 émetteur de flux et au moins un deuxième flux (36) transmis sur le réseau à partir d'un deuxième émetteur (12) de flux,
- caractérisé en ce que :**
- 45 le paramètre de réglage de débit spécifie une quantité de données qui peut être envoyée avant qu'un événement spécifique se produise ; et  
le dispositif comprend, en outre, un réglage (60) de largeur de bande associé à l'interconnexion de réseau et pouvant fonctionner pour régler le paramètre de réglage de débit du premier émetteur de flux pour satisfaire un paramètre de performance associé au deuxième flux (36).
- 50 2. Dispositif suivant la revendication 1, comprenant en outre :
- une première connexion de réseau comprenant le premier émetteur de flux ; et  
un deuxième réseau comprenant l'interconnexion de réseau.
- 55 3. Dispositif suivant la revendication 2, dans lequel l'interconnexion de réseau est une passerelle locale et le deuxième réseau est un réseau local.

4. Dispositif suivant la revendication 2 comprenant, en outre :

au moins un canal reliant le premier émetteur de flux et le deuxième émetteur de flux à l'interconnexion de réseau, le canal ayant une capacité de largeur de bande partagée par le premier flux et par le deuxième flux.

5  
5. Dispositif suivant la revendication 2, dans lequel le premier flux est un flux qui n'est pas en temps réel et le deuxième flux est un flux en temps réel, le deuxième flux exigeant d'une manière continue une partie de la capacité de largeur de bande.

10  
6. Dispositif suivant la revendication 5, dans lequel le réglage de la largeur de bande règle le paramètre de réglage de débit de façon à ce que le premier flux utilise une partie de la capacité de largeur de bande qui est inférieure ou égale à la différence entre la capacité de largeur de bande et les exigences en largeur de bande du deuxième flux.

15  
7. Dispositif suivant la revendication 1, dans lequel le paramètre de performance est choisi dans un jeu de paramètres de qualité de service déterminé à l'avance et associé au deuxième flux.

20  
8. Dispositif suivant la revendication 7, dans lequel le réglage de largeur de bande maintient les paramètres de qualité de service du deuxième flux en réglant le paramètre de réglage de débit et en réglant les exigences de largeur de bande du premier débit.

9. Dispositif suivant la revendication 1, dans lequel le paramètre de réglage de débit régule le débit du premier flux.

25  
10. Dispositif suivant la revendication 1, dans lequel le paramètre de réglage de débit régule l'usage de largeur de bande du premier flux en régulant le débit du premier flux.

30  
11. Dispositif suivant la revendication 1, dans lequel le premier émetteur de flux fonctionne suivant le Transmission Control Protocol (TCP) et le paramètre de réglage de débit comprend une fenêtre de réglage de débit telle que définie par TCP.

35  
12. Dispositif suivant la revendication 1, dans lequel l'interconnexion de réseau est un dispositif choisi dans un groupe consistant en des routeurs, des transformateurs de protocole et des passerelles.

40  
13. Dispositif suivant la revendication 11, dans lequel l'interconnexion de réseau alloue dynamiquement de la largeur de bande entre des flux en temps réel et des flux qui ne sont pas en temps réel, de sorte que de la largeur de bande allouée à des flux qui ne sont pas en temps réel ne dépasse pas la largeur de bande disponible pour les flux qui ne sont pas en temps réel.

45  
14. Dispositif suivant la revendication 13, dans lequel l'interconnexion de réseau alloue dynamiquement de la largeur de bande, en augmentant la fenêtre de réglage de débit d'un flux qui n'est pas en temps réel qui est dans un état stationnaire, pendant un temps de démarrage lent d'un autre flux qui n'est pas en temps réel.

50  
15. Dispositif suivant la revendication 14, dans lequel l'interconnexion de réseau alloue dynamiquement de la largeur de bande, en diminuant le paramètre de réglage de débit du flux qui n'est pas en temps réel, lorsque l'autre flux qui n'est pas en temps réel atteint un état stationnaire.

55

50

55

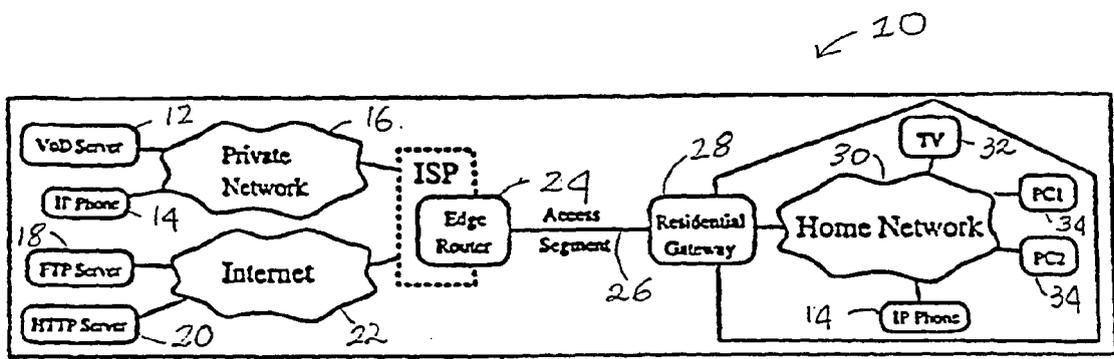


FIGURE 1

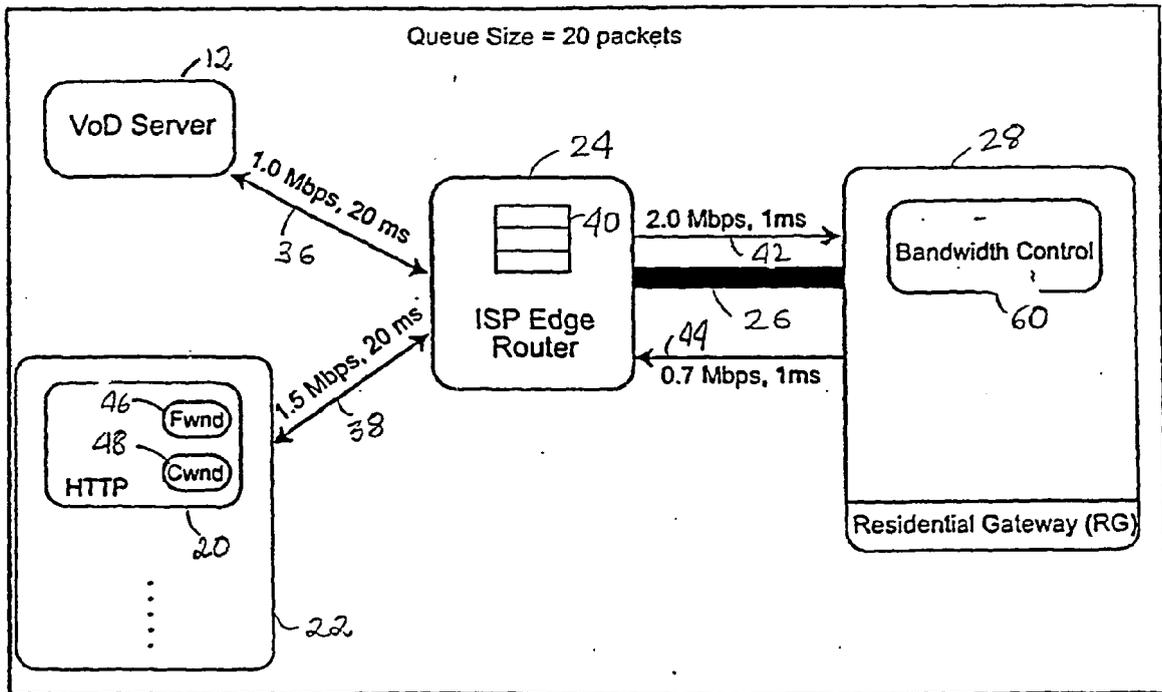


FIGURE 2

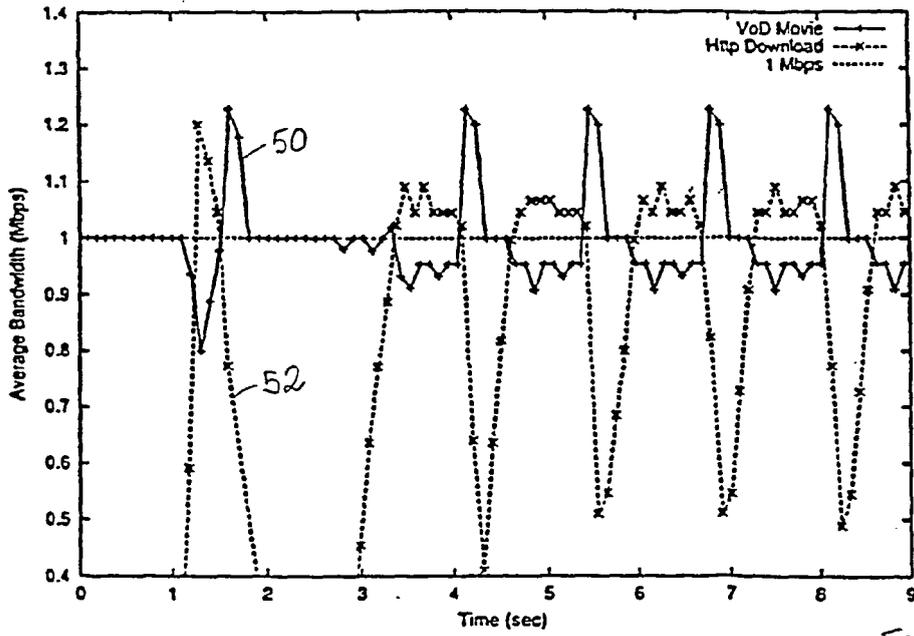


Figure 3

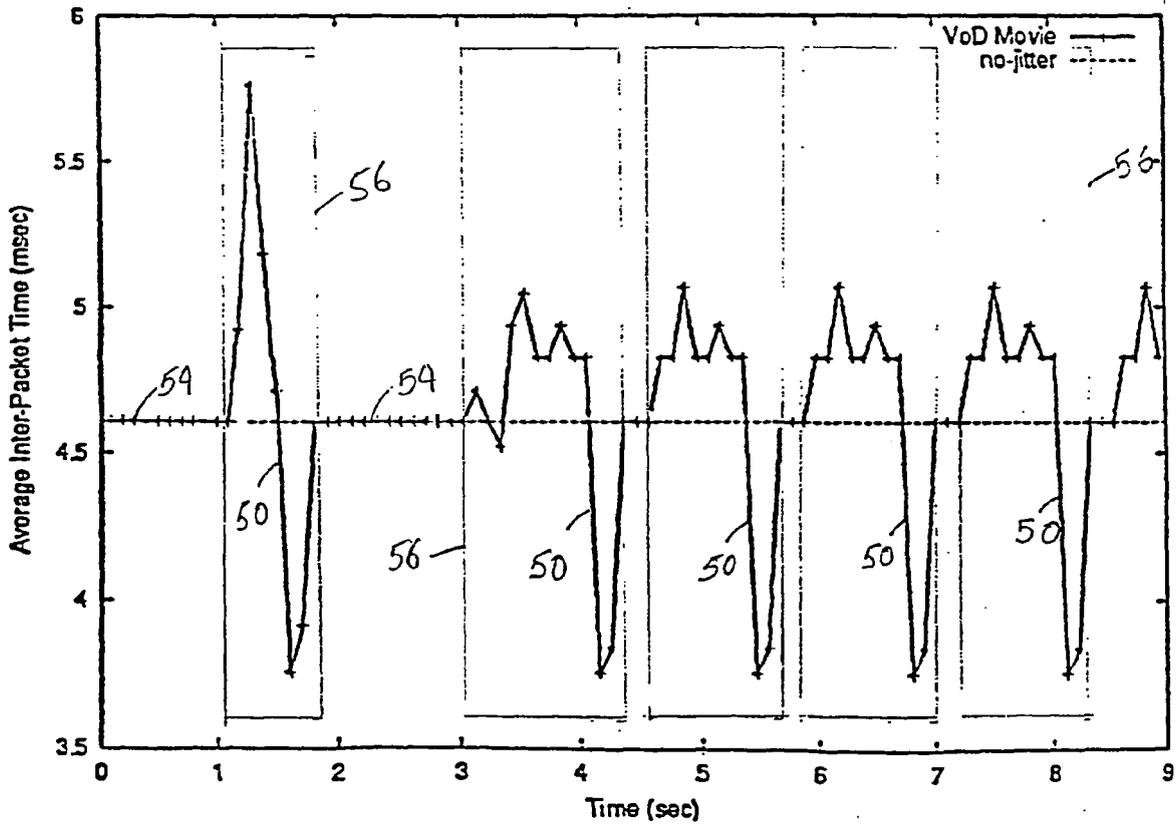


Figure 4

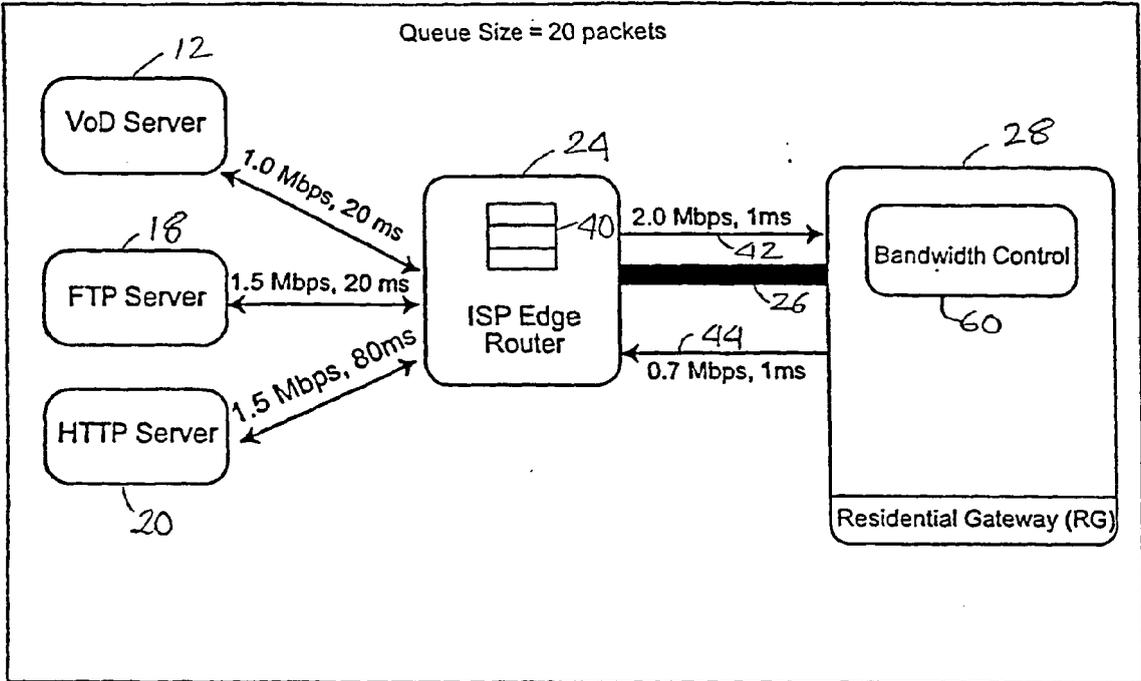


Figure 5

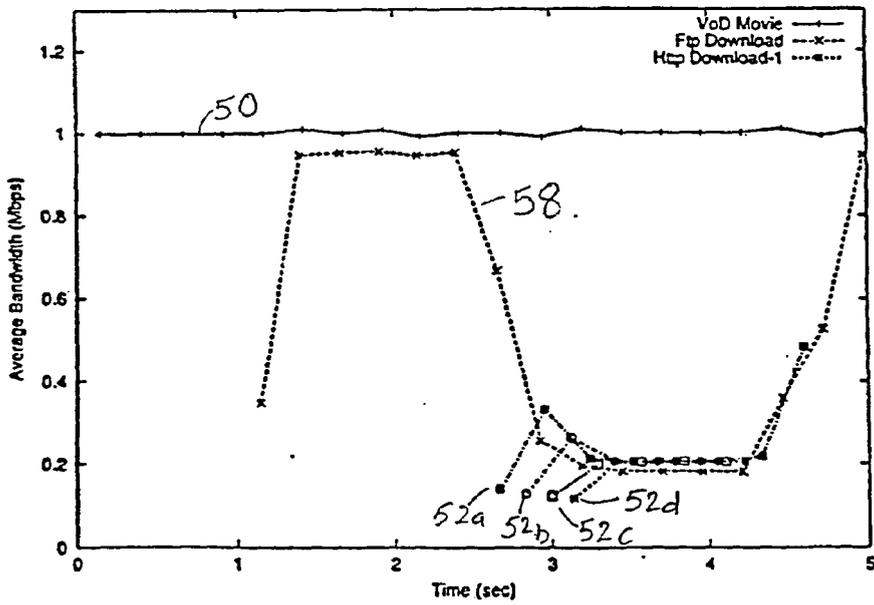


Figure 6

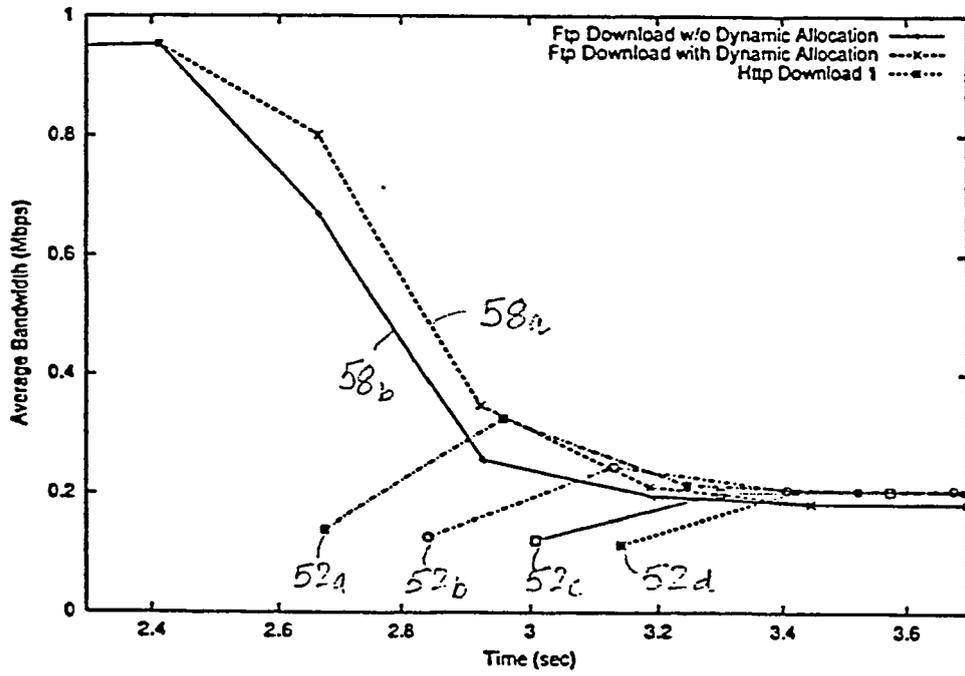


Figure 7

**REFERENCES CITED IN THE DESCRIPTION**

*This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.*

**Patent documents cited in the description**

- US 6307839 B [0011]
- EP 0948168 A1 [0012]