



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04L	A2	(11) International Publication Number: WO 00/21233 (43) International Publication Date: 13 April 2000 (13.04.00)
---	----	---

(21) International Application Number: PCT/US99/23143

(22) International Filing Date: 6 October 1999 (06.10.99)

(30) Priority Data:
09/167,882 7 October 1998 (07.10.98) US

(71) Applicant (for all designated States except US): NOKIA INTERNET COMMUNICATIONS INC. [US/US]; IPR Dept., 313 Fairchild Drive, Mt. View, CA 94043 (US).

(71)(72) Applicants and Inventors: DIXIT, Sudhir [US/US]; Nokia, 3 Burlington Woods Drive, Burlington, MA 01803 (US). GHANI, Nasir [-/US]; Nokia, 3 Burlington Woods Drive, Burlington, MA 01803 (US).

(74) Agent: ROLNIK, Robert; Nokia Inc., 6000 Connection Drive, Irving, TX 75039 (US).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).


Published

Without international search report and to be republished upon receipt of that report.

(54) Title: ENHANCED ACKNOWLEDGEMENT PACING DEVICE AND METHOD FOR TCP CONNECTIONS

(57) Abstract

An enhanced acknowledgement pacing device (500) and method for TCP connections is disclosed. The invention includes a link layer entity (528) for receiving data packets from a source and forwarding the data packets (540) to a forward data link, the link layer entity (528) storing the received data packets (540) in a data packet buffer (530) until the data packets depart the link layer entity (528) and are forwarded to the forward link and an acknowledgement pacing device (500), coupled to the link layer entity, for pacing acknowledgement packets (532) to be sent to the source in response to receiving the data packets (540) from the source (540). The acknowledgement pacing device (500) further includes an acknowledgement control unit (510) for monitoring congestion at the link layer entity (528) and generating a control signal for controlling the processing of acknowledgement packets (532) based upon whether congestion is occurring at the link layer entity (528), an acknowledgement packet buffer (534), coupled to the acknowledgement control unit (510), for storing acknowledgement packets (532) received from the acknowledgement control unit (510) and a scheduler (520), coupled to the acknowledgement control unit (510) and the acknowledgement buffer (534), the scheduler (520) releasing acknowledgement packets (532) to the source (540) based upon the control signal generated by the acknowledgement control unit (510).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

ENHANCED ACKNOWLEDGMENT PACING DEVICE AND METHOD FOR TCP CONNECTIONS

BACKGROUND OF THE INVENTION

1. Field of the Invention.

5 This invention relates in general to networks, and more particularly to an enhanced acknowledgment pacing device and method for TCP connections.

2. Description of Related Art.

Today, an organization's computer network has become its circulatory system. Organizations have combined desktop work stations, servers, and hosts into
10 Local Area Network (LAN) communities. These Local Area Networks have been connected to other Local Area Networks and to Wide Area Networks (WANs). It has become a necessity of day-to-day operation that pairs of systems must be able to communicate when they need to, without regard to where they may be located in the network.

15 During the early years of network computing, proprietary networking protocols were the standard. However, the development of the Open Systems Interconnection Reference Model introduced by the International Organization for Standardization (ISO) has led to an impressive degree of interworking, which generally allows end-user applications to work very well between systems in a
20 network. Implementations are based on written standards that have been made available by volunteers from dozens of computer vendors, hardware component vendors and independent software companies.

During the last decade, LANs have been proliferating. This has created a recurring problem of how to minimize congestion and optimize throughput that must be solved by network managers. An early solution was to simply divide Local Area Networks into multiple smaller networks serving smaller populations. These
5 segments were connected by bridges to form a single Local Area Network with traffic being segregated locally to each segment.

The evolution of new network types and Wide Area Networks created a need for routers. For example, the Internet is a set of networks connected by gateways, which are sometimes referred to as routers. Routers added filtering and firewalling
10 capability to provide more control over broadcast domains, limit broadcast traffic and enhance security. A router is able to chose the best path through the network due to embedded intelligence. This added intelligence also allowed routers to build redundant paths to destinations when possible. Nevertheless, the added complexity of best path selection capability accorded by the embedded intelligence increased the
15 port cost of routers and caused substantial latency overhead. Shared-media networks comprising distributed client/server data traffic, expanded user populations and more complex applications gave birth to new bandwidth bottlenecks. Such congestion produced unpredictable network response times, the inability to support the delay-sensitive applications and higher network failure rates.

20 Congestion control in modern networks is increasingly becoming an important issue. The explosive growth of Internet applications such as the World Wide Web (WWW) has pushed current technology to its limit, and it clear that faster transport and improved congestion control mechanisms are required. As a result, many equipment vendors and service providers are turning to advanced networking

technology to provide adequate solutions to the complex quality of service (QoS) management issues involved. Examples include asynchronous transfer mode (ATM) networks and emerging IP network services. Nevertheless, there is still the need to support a host of existing legacy IP protocols within these newer paradigms. In particular, the ubiquitous TCP transport-layer protocol has long been the workhorse transport protocol in IP networks, widely used by web-browsers, file/email transfer services, etc.

Transmission Control Protocol (TCP) is a part of the TCP/IP protocol family that has gained the position as one of the world's most important data communication protocols with the success of the Internet. TCP provides a reliable data connection between devices using TCP/IP protocols. TCP operates on top of IP that is used for packing the data to data packets, called datagrams, and for transmitting across the network.

The Internet Protocol (IP) is a network layer protocol that routes data across an Internet. The Internet Protocol was designed to accommodate the use of host and routers built by different vendors, encompass a growing variety of growing network types, enable the network to grow without interrupting servers, and support higher-layer of session and message-oriented services. The IP network layer allows integration of Local Area Network "islands".

However, IP doesn't contain any flow control or retransmission mechanisms. That is why TCP is typically used on top of it. Especially, TCP uses acknowledgments for detecting lost data packets. TCP/IP networks are nowadays probably the most important of all networks, and operate on top of several (physical) networks, such as the ATM networks mentioned above. These underlying networks

may offer some information about the condition of network and traffic, which may be used to provide feedback regarding congestion.

To manage congestion, TCP uses a sliding window mechanism coupled with reactive congestion control to adjust the sender's window size. The protocol adjusts its transmission behavior contingent to returning acknowledgment (ACK) packets sent from the remote receiver's end.

A problem with TCP, however, is that its congestion control mechanism is relatively slow. Most TCP implementations use very coarse timers to measure timeouts, i.e., roughly 200-500 ms. granularity. Further, most TCP implementations rely on ACK delays or packet drops to detect congestion. As a result, excessive source window reductions can result in large amounts of bandwidth being wasted as the TCP source is forced to restart its transmission window. Further, many studies have shown that TCP does not perform very well over ATM networks, especially for larger WAN-type propagation delays.

To combat the above shortcomings with TCP, it is necessary to minimize the chances of network congestion by somehow incorporating faster congestion indication mechanisms in the TCP feedback loop. However, to ensure compatibility with current versions and to expedite market acceptance, any such attempt must preclude changes to the actual TCP protocol or its implementation.

Along these lines, a variety of ACK pacing schemes have been proposed. These ACK pacing schemes basically modulate the spacing of TCP ACK packets to limit source emissions during periods of congestion. ACK pacing is well-suited at the boundary of high speed (sub)networks, such as ATM, gigabit IP (i.e., optical WDM), or satellite. In essence this technique performs TCP traffic shaping at the

access nodes. Such methodologies are specifically beneficial for advanced ATM data services, i.e., underlying ABR flow control or per-connection queuing, where congestion tends to buildup at the periphery of the ATM network, i.e., in the access nodes. If the forward link is congested, as indicated via some congestion metric,

5 ACK packets are appropriately delayed before being sent to the source.

Other authors have proposed modifying fields in the ACK packets themselves, i.e., receiver-window size, to improve performance. However, such schemes either require accurate round-trip delay measurements or cannot maintain tight buffer control. Furthermore, rewriting ACK packet fields will require

10 expensive checksum recomputations.

Although ACK pacing is an effective way of controlling TCP source behaviors, many of the proposed schemes are either too complex and/or overly sensitive to network parameter settings. Since studies have shown that TCP's throughput and fairness levels can be low in many high-speed network scenarios, it

15 is necessary to devise efficient, practical schemes to enhance its performance.

Although amending the protocol's functionality itself is also an option, this may not be a feasible alternative in the short-to-medium time frame. It is along these lines that the ACK pacing methods can provide significant benefits.

It can be seen that there is a need for a more robust, comprehensive scheme

20 for ACK pacing.

It can also be seen that there is a need for ACK pacing that provides high throughput and precise levels of bandwidth fairness.

It can also be seen that there is a need for ACK pacing that significantly reduces TCP buffering delays and is applicable to a wide range of network scenarios.

It can also be seen that there is a need for ACK pacing that provides faster
5 congestion indication without modifying the TCP protocol.

SUMMARY OF THE INVENTION

To overcome the limitations in the prior art described above, and to overcome other limitations that will become apparent upon reading and understanding the present specification, the present invention discloses an enhanced acknowledgment pacing device and method for TCP connections.

The present invention solves the above-described problems by providing a more robust, comprehensive scheme for ACK pacing. The ACK pacing according to the present invention provides high throughput and precise levels of bandwidth fairness. Further, the ACK pacing significantly reduces TCP buffering delays and is applicable to a wide range of network scenarios. Thus, the ACK pacing provides faster congestion indication without modifying the TCP protocol.

A system in accordance with the principles of the present invention includes a link layer entity for receiving data packets from a source and forwarding the data packets to a forward data link, the link layer entity storing the received data packets in a data packet buffer until the data packets depart the link layer entity and are forwarded to the forward data link and an acknowledgment pacing device, coupled to the link layer entity, for pacing ACK packets to be sent to the source in response to receiving the data packets from the source. The acknowledgment pacing device further includes an acknowledgment control unit for monitoring congestion at the link layer entity and generating a control signal for controlling the processing of acknowledgment packets based upon whether congestion is occurring at the link layer entity, an acknowledgment packet buffer, coupled to the acknowledgment control unit, for storing acknowledgment packets received from the acknowledgment control unit and a scheduler, coupled to the acknowledgment control unit and the acknowledgment buffer, the scheduler releasing acknowledgment packets to the source based upon the control signal generated by the acknowledgment control unit.

Other embodiments of a system in accordance with the principles of the invention may include alternative or optional additional aspects. One such aspect of the present invention is that the scheduler chooses ACK packets to release based upon a queuing strategy.

5 Another aspect of the present invention is that the queuing strategy includes sending a head-of-line ACK packet when aggregate ACK packet buffering is used.

Another aspect of the present invention is that the queuing strategy includes a weighted-round-robin (WRR) process for selecting an ACK packet in the ACK packet buffer for release when per-class or per-flow ACK packet buffering is used.

10 Another aspect of the present invention is that the weighted-round-robin (WRR) process uses weights for weighting the selection of the ACK packet for release inversely proportionally to a TCP maximum segment size (MSS) to the lessen a bias against smaller MSS flows.

Another aspect of the present invention is that the queuing strategy includes a fair-queuing (FQ) process for selecting an ACK packet in the ACK packet buffer for release when per-class or per-flow ACK packet buffering is used.

Another aspect of the present invention is that the acknowledgment control unit further includes an ACK packet pacing processor, the ACK packet pacing processor generating the control signal for controlling the processing of acknowledgment packets using a ACK packet arrival processor and a data packet departure processor.

20 Another aspect of the present invention is that the ACK packet arrival processor controls the processing of ACK packets to the ACK packet buffer by checking the congestion at the link layer entity and deciding whether to hold ACK packets in the ACK packet buffer or to send the ACK packet directly to the source without buffering the ACK packets in the ACK packet buffer.

25

Another aspect of the present invention is that the ACK packet arrival processor decides whether to hold ACK packets in the ACK packet buffer or to send the ACK packet directly to the source without buffering the ACK packets in the ACK packet buffer by determining if the link layer entity is congested, determining if the ACK packet buffer is empty, storing an ACK packet in the buffer if the link layer entity is congested or the ACK packet buffer is not empty, and forwarding the ACK packet to the source if the ACK packet buffer is empty and the link layer entity is not congested.

Another aspect of the present invention is that the ACK packet is stored in the ACK packet buffer and gated out by the scheduler if the ACK packet is a first ACK packet to be buffered in the ACK packet buffer during congestion.

Another aspect of the present invention is that the ACK control unit increases a spacing between ACK packets gated from the ACK packet buffer if the ACK packet is the first ACK packet to be buffered in the ACK packet buffer during congestion.

Another aspect of the present invention is that the data packet departure processor decreases a spacing between the release of ACK packets from the ACK packet buffer if congestion in the link layer entity has abated.

These and various other advantages and features of novelty which characterize the invention are pointed out with particularity in the claims annexed hereto and form a part hereof. However, for a better understanding of the invention, its advantages, and the objects obtained by its use, reference should be made to the drawings which form a further part hereof, and to accompanying descriptive matter, in which there are illustrated and described specific examples of an apparatus in accordance with the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

Fig. 1 illustrates the OSI model which includes seven layers;

5 Fig. 2 illustrates a comparison of the Internet Protocol Network Layer and the OSI seven layer model;

Fig. 3 illustrates a packet stream and a TCP sliding window;

Fig. 4 illustrates a network system wherein a receiver provides acknowledgments to the source as well as receives data from the source;

10 Fig. 5 illustrates the enhanced ACK pacing device according to the present invention;

Fig. 6 illustrates the pseudocode for the TCP ACK arrival and data departure 660 method according to the present invention;

Fig. 7 illustrates pseudocode for a congestion status method using two
15 hysteresis queue thresholds, QL and QH; and

Fig. 8 illustrates the pseudocode for the enhanced TCP ACK arrival and ATM cell departure methods according to the present invention

DETAILED DESCRIPTION OF THE INVENTION

In the following description of the exemplary embodiment, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration the specific embodiment in which the invention may be practiced.

5 It is to be understood that other embodiments may be utilized as structural changes may be made without departing from the scope of the present invention.

The present invention provides an enhanced acknowledgment pacing device and method for TCP connections. A more robust, comprehensive scheme for ACK pacing is provided which allows high throughput and precise levels of bandwidth
10 fairness. Further, the ACK pacing significantly reduces TCP buffering delays and is applicable to a wide range of network scenarios. Thus, the ACK pacing according to the present invention provides faster congestion indication without modifying the TCP protocol.

Fig. 1 illustrates the OSI model 100 which includes seven layers, including
15 an Application Layer 110, Presentation Layer 120, Session Layer 130, Transport Layer 140, Network Layer 150, Data Link Layer 160, and Physical Layer 170. The OSI model 100 was developed by the International Organization for Standardization (ISO) and is described in ISO 7498, entitled "The OSI Reference Model", and which is incorporated by reference herein.

20 Each layer of the OSI model performs a specific data communications task, a service to and for the layer that precedes it (e.g., the Network Layer provides a service for the transport layer). The process can be likened to placing a letter in a series of envelopes before it is sent through the postal system. Each succeeding envelope adds another layer of processing or overhead information necessary to
25 process the transaction. Together, all the envelopes help make sure the letter gets to the right address and that the message received is identical to the message sent.

Once the entire package is received at its destination, the envelopes are opened one by one until the letter itself emerges exactly as written.

In a data communication transaction, however, each end user is unaware of the envelopes, which perform their functions transparently. For example, an automatic bank teller transaction can be tracked through the multi-layer OSI system. One multiple layer system (Open System A) provides an application layer that is an interface to a person attempting a transaction, while the other multiple layer system (Open System B) provides an application layer that interfaces with applications software in a bank's host computer. The corresponding layers in Open Systems A and B are called peer layers and communicate through peer protocols. These peer protocols provide communication support for a user's application, performing transaction related tasks such as debiting an account, dispensing currency, or crediting an account.

Actual data flow between the two open systems (Open System A and Open System B), however, is from top 180 to bottom 182 in one open system (Open System A, the source), across the communications line, and then from bottom 182 to top 180 in the other open system (Open System B, the destination). Each time that user application data passes downward from one layer to the next layer in the same system more processing information is added. When that information is removed and processed by the peer layer in the other system, it causes various tasks (error correction, flow control, etc.) to be performed.

The ISO has specifically defined all seven layers, which are summarized below in the order in which the data actually flows as they leave the source:

Layer 7, the Application Layer 110, provides for a user application (such as getting money from an automatic bank teller machine) to interface with the OSI

application layer. That OSI application layer 110 has a corresponding peer layer in the other open system, the bank's host computer.

Layer 6, the Presentation Layer 120, makes sure the user information (a request for \$50 in cash to be debited from your checking account) is in a format (i.e., syntax or sequence of ones and zeros) the destination open system can understand.

Layer 5, the Session Layer 130, provides synchronization control of data between the open systems (i.e., makes sure the bit configurations that pass through layer 5 at the source are the same as those that pass through layer 5 at the destination).

Layer 4, the Transport Layer 140, ensures that an end-to-end connection has been established between the two open systems and is often reliable (i.e., layer 4 at the destination confirms the request for a connection, so to speak, that it has received from layer 4 at the source).

Layer 3, the Network Layer 150, provides routing and relaying of data through the network (among other things, at layer 3 on the outbound side an address gets placed on the envelope which is then read by layer 3 at the destination).

Layer 2, the Data Link Layer 160, includes flow control of data as messages pass down through this layer in one open system and up through the peer layer in the other open system.

Layer 1, the Physical Interface Layer 170, includes the ways in which data communications equipment is connected mechanically and electrically, and the means by which the data moves across those physical connections from layer 1 at the source to layer 1 at the destination.

Fig. 2 is a comparison 200 illustrating where the Internet Protocol Network Layer 202 fits in the OSI seven layer model 204. In Fig. 2, the transport layers 210 provides data connection services to applications and may contains mechanisms that

guarantee that data is delivered error-free, without omissions and in sequence. The transport layer 210 in the TCP/IP model 212 sends segments by passing them to the IP layer 202, which routes them to the destination. The transport layer 210 accepts incoming segments from IP 202, determines which application is the recipient, and
5 passes the data to that application in the order in which it was sent.

Thus, the Internet Protocol 202 performs Network Layer functions and routes data between systems. Data may traverse a single link or may be relayed across several links in an Internet. Data is carried in units called datagrams, which include an IP header that contains layer 3 220 addressing information. Routers examine the
10 destination address in the IP header in order to direct datagrams to their destinations. The IP layer 202 is called connectionless because every datagram is routed independently and the IP layer 202 does not guarantee reliable or in-sequence delivery of datagrams. The IP layer 202 routes its traffic without caring which application-to-application interaction a particular datagram belongs to.

15 The TCP layer 210 provides a reliable data connection between devices using TCP/IP protocols. The TCP layer 210 operates on top of the IP layer 202 that is used for packing the data to data packets, called datagrams, and for transmitting the across the underlying network via physical layer 230.

However, the IP protocol doesn't contain any flow control or retransmission
20 mechanisms. That is why the TCP layer 210 is typically used on top of the IP layer 202. In contrast, TCP protocols provide acknowledgments for detecting lost data packets.

Fig. 3 illustrates a packet stream 300 and a TCP sliding window 310. One of the main features of a TCP source is that it uses a sliding window 310 that
25 determines the bytes and, consequently, the IP packets that can be sent before an

acknowledgment is received from the receiver. This makes it possible to adjust the effective transmission rate of the source.

When the TCP source increases the size of the sliding window 310, its average transmission rate increases, too. The sliding window 310 is on top of octets 5 12-19. Octets up to 11 have been transmitted and the sliding window 310 has moved past them. Inside the sliding window 310, there are two octet groups 320, 322. The first octet group 320 is the octets from 12 to 16, which have been transmitted 330. The second group of octets 322 in the sliding window 310 are octets 17-19, which have not yet been transmitted. The second group of octets 322 10 can be sent immediately 340. Finally, octets 20 and upwards 350 cannot be transmitted 360. Octet 12 has to be acknowledged and the sliding window slid forward before octet 20 may be transmitted. Thus, TCP provides retransmission of lost data packets and flow control using this TCP sliding window 310. The sliding window 310 is actually the minimum of the congestion window of the window 15 advertisement which is sent by the receiver.

Fig. 4 illustrates a TCP network system 400 wherein a receiver 410 provides acknowledgments 420 to the source 430 as well as receives data 440 from the source 430. The receiver 410 sends acknowledgment packets 420 that also include window advertisement data 450 for informing the source 430 of the capacity of the receiver 20 410 to handle incoming data 440. Thus, the receiver 410 can advertise a suitable window size 450 for flow control purposes. In practice, the window advertisement 450 specifies how many additional octets of data the receiver 410 is prepared to accept. The source 430 is supposed to adjust its sliding window according to this advertisement, unless the congestion window 460 maintained by the source 430 is 25 smaller.

The second window, the congestion window 460, is used internally at the TCP source 430 for dropping the size of the sliding window. This occurs if a timer expires telling that a data packet has been sent, but no acknowledgment has arrived within a certain time period. This means that the data packet has been lost which is most probably caused by network congestion. In order not to make the congestion worse, the TCP source 430 drops its transmission rate by reducing the size of the sliding window. The relation of these windows can be expressed as:

$$T_w = \text{MIN}(\text{window advertisement, congestion window}),$$

where T_w refers to the transmission window, i.e., the sliding window.

In principle, the congestion window 460 and feedback information included in the advertisement window 450 provided by the underlying network can be used for the same purpose, namely to adjust the transmission rate of the TCP source 430 according to the load and congestion of the network. However, one important difference between the congestion window 460 and feedback information included in the advertisement window 450 is that the congestion window 460 works on the end-to-end basis and is typically quite slow to react to changes due to relatively long timeouts. Thus, the congestion window 460 can not also give any detailed information. The TCP source 410 simply knows that a packet has been discarded which may not give the exact picture about the network condition. Feedback information included in the advertisement window 450, on the other hand, may be more accurate and may react faster to the changing conditions.

An underlying network can use the receiver's window advertisements 450 carried in acknowledgment packets 420 for controlling the transmission speed of a TCP source 410. This may be accomplished by adding device or network functionality, herein referred to as Feedback Information Converter (FIC).

Thus, TCP uses a sliding window protocol where the source 430 adjusts its window size based upon returning ACK packets 420 from the receiver 410. Hence the window's 460 growth rate will be related to the rate of these returning packets 420. Therefore, it is evident that by modifying the timing of the returning ACK stream 420, the growth of the source window 460 can be controlled. It is this fundamental principle upon which ACK pacing methods are based. Specifically, these methods appropriately delay returning ACK packets 420 in congested network elements, e.g., access nodes and IP routers, to limit excessive emissions by the source 430. When properly done, ACK pacing can reduce TCP timeouts, limit queue buildups, and thereby improve overall connection goodputs.

Due to the largely asymmetric nature of TCP traffic profiles, ACK pacing, is really only required at the TCP source side 430. This is a noteworthy point, since it implies that the required ACK pacing functionality need only be limited to large web-servers/file-hosts. Hence no expensive upgrades are required for a much larger, diverse user access base. It should be mentioned, however, that ACK pacing assumes good rate control inside the network. This essentially abstracts the network to a fairly constant bandwidth, causing congestion to occur primarily at the access nodes.

Advanced ATM bearer capabilities, e.g., VBR-nrt, ABR and GFR, can realistically achieve these conditions. Furthermore, it is expected that emerging rate guarantees in high-speed IP routers will also yield conditions favorable towards ACK pacing.

However, as stated above, many of the current ACK pacing methods are not particularly amenable to implementation. For example, fast-TCP (F-TCP) required

knowledge of the underlying data "clearing" rate in the forward direction. This can either be the link capacity or for the case of the ATM available bit rate (ABR) service category, the connection's allowed cell rate (ACR), etc. The computed delays for the ACK packets 420 are based upon this rate.

5 Clearly, such schemes require more advanced information processing methods and can be problematic if the underlying rate varies widely. Furthermore, delayed emission of ACK packets 420 by the remote TCP client 410 can compound the sensitivity issues and significantly degrade the performance of such schemes. Also, no explicit fairness provision can be provided by these schemes since they
10 simply buffer returning ACK packets 420 in an aggregate manner, i.e., first-in-first-out (FIFO). Another ACK pacing method, the ACK bucket scheme, requires too much per-flow state, essentially "tracking" the windowing behaviors of each TCP flow.

 Fig. 5 illustrates the enhanced ACK pacing device 500 according to the
15 present invention. The ACK pacing device 500 relies on queue length information to infer congestion levels and does not require any additional (expensive) timer mechanisms. The ACK pacing device 500 illustrated in Fig. 5 is very generic and can be tailored to fit a wide range of networks.

 In Fig. 5, an ACK control unit 510 is provided. The ACK control unit 510
20 controls the processing of ACK packets during both overload (i.e., congestion) and underload periods along with the operation of the ACK scheduler unit 520. The ACK control unit 510 relies on traffic measurements and data transmit notifications 522 from the underlying link-layer entity 528 and the data packet buffer 530. During congestion periods, returning ACK packets 532 are stored in the ACK

buffers 534 using appropriate classification granularities (aggregate, per-class, per-flow) and gated out at an appropriately chosen rate. Specifically, the emission of ACK packets 532 during congested periods are performed so as to allow the buffers 534 to empty in reasonable time. When congestion subsides, the ACK emission rate is then increased to allow for improved bandwidth utilization. Note the ACK control unit 510 activates the ACK scheduler unit 520 to emit ACK packets 532 in the buffers 534 in all cases.

Since TCP is an expansive protocol, it always attempts to increase its transmission quota barring any receiver window limitations. This means that for large (bulk) file transfers, the regular TCP protocol will repeatedly increase its window size, loose packets, and then slow down. As the volume of data in the system increases, so does the number of ACK packets 532.

This point has a very subtle implication for ACK pacing schemes. Namely, the data packet 540 growth in the regular TCP protocol will be "replaced" by ACK packet 532 growth. This is referred to as the ACK buffer "drift" phenomenon. The rate of this drift will be linear (i.e., fast) for the case of ACK pacing with TCP connections in slow start phases, and will be sub-linear for ACK pacing with TCP connections in congestion avoidance phases.

There are two possible methods to address this problem. The simpler approach is to provide ACK buffers 534 with sufficient capacity for the ACK packets 532 and employ drop-from-front strategies in rare event of ACK buffer 534 exhaustion. Typically, ACK numbers in the front will most likely pertain to lower sequence numbers than those for arriving packets 540. This buffering approach is very reasonable, since ACK packets 532 are small (40 bytes), and most file transfers

are not infinite. For example, an ACK buffer 534 of 64kB of RAM can hold approximately 1,700 ACK packets 532, which is more than adequate for 155 Mb/s WAN links.

Another approach would be to track TCP sequence numbers using two variables, i.e., per-flow accounting for the last-in and last-out values. This approach can yield smaller memory requirements, but requires ACK number re-writing (i.e., check sum recomputations). Also, if ACK packets 532 arrive out of sequence, special considerations are necessary. Furthermore, it is likely that the other fields present in ACK packets 532, such as receiver window sizes and URG/RST flags, may also contain non-redundant information which can complicate matters further.

Thus, the ACK buffering approach is more feasible from an implementation perspective. The ACK buffering approach posts minimal additional constraints and does not tamper with any fields in the TCP packet 532.

As shown in Fig. 5, in the forward direction of data flow, the link layer entity 528 can be representative of a wide range of underlying technologies. Examples include dedicated links or ATM VC's, or IP flow classes. Furthermore, the link layer entity 528 can be either dedicated to a single TCP flow, e.g., ATM VC, etc., or be shared among a group of TCP flows (traffic aggregation). Similarly, the ACK pacing in the reverse direction can be done on different levels. For example, if per-flow queuing is done in the forward direction., then per-flow ACK pacing is also necessary in the reverse direction, i.e., per flow data/per-flow ACK.

If, however, aggregate or class-based ACK pacing is done in forward direction, then it may be desirable to do likewise in the reverse direction (aggregate data/aggregate ACK, per-class data/per-class ACK). Others may decide to do

simple, aggregate queuing in the forward direction, yet more advanced per-flow ACK buffering in the reverse direction. This approach improves fairness amongst flows aggregated onto the same link-layer entity 530, without requiring high speed per-flow buffering and scheduling techniques in the forward direction. Although
5 per-flow ACK accounting is still required for incoming ACK packets 532, it is restricted to the network edge where the processing rate requirements are also significantly reduced since ACK packets 532 pertain to larger IP packet sizes. By choosing scheduler allocations, i.e., weights, inversely proportional to a flow's TCP maximum segment size (MSS), the bias against smaller MSS flows can be lessened
10 (to an extent).

With the ongoing standardization efforts for a differentiated services architecture, the latter philosophy fits in quite nicely. Namely, per-flow accounting/overhead is limited to the access parts of the network, i.e., where ACK
pacing is done, reducing complexity within the backbone. Since most access nodes
15 will carry much fewer connections than backbone devices, this approach is very feasible in emerging networks.

In light of the above discussion, the ACK scheduler 520 can be specified as being fairly generic, borrowing from a variety of packet scheduling methods to improve fairness. For example, in the simplest form of aggregate (FIFO) ACK
20 buffering, the scheduler 520 merely has to send the head-of-line (HOL) ACK. For more advanced per-class or per-flow ACK buffering strategies, a weighted-round-robin (WRR) or fair-queuing (FQ) scheduler can be implemented to "choose" the next suitable ACK for transmission.

The ACK control unit 510 includes a ACK pacing processor 512 for controlling the pacing of ACK packets to the source. The ACK pacing processor includes two main components: the data packet departure processor 514 and the ACK arrival processor. The ACK arrival processor 516 checks congestion levels and decides whether or not to hold incoming ACK packets. The data packet departure processor 514 monitors congestion levels via the link layer entity 528 and data packet buffer 530 and decides when to "clock" out ACK packets 532 to the source.

Fig. 6 illustrates the pseudocode 600 for the TCP ACK arrival 610 and data departure 660 method according to the present invention. The TCP ACK arrival 610 and data departure 660 method are executed for all incoming TCP ACK packets.

In Fig. 6, it is assumed that queue objects exist for enqueueing/dequeueing ACK packets, and that a running count of the number of buffered ACK packets is kept, e.g., num_ACK 604. In case of link-layer congestion and/or a non-empty ACK buffer 612, an incoming ACK packet is stored in the buffer 614. The buffered ACK packets are kept in the buffer and can only be sent out appropriately by the data packet departure method 660.

As discussed previously, the ACK buffering can be done on an aggregate basis or more selective per-class/per-flow basis as discussed above with reference to Fig. 5. If the ACK packet arrives at an empty buffer and there is no congestion 640, it is simply forwarded onwards to the TCP source (i.e., transparent pass-through) 642. However, if this is the first ACK to be buffered 630, then in order to "jump-start" the ACK emission process, this ACK packet must be gated out after an appropriate interval 632.

To avoid any dependencies on expensive timer mechanisms, the emission of ACK packets should be associated with the underlying data packet departure process 660 in the link-layer entity. Namely, during congestion, ACK packets are sent after every α_1 data packets have been emitted, where α_1 is termed an (integral) slow-down factor.

From an implementation perspective, the above functionality can be achieved elegantly by using a simple counter variable, e.g., `pkt_counter` 644. For the first ACK packet, the counter variable value is set to α_1 and then decremented per data packet departure according to the data departure process 660. When the counter variable value reaches zero, a buffered ACK packet is released and the counter is reset.

Contrary to some expectations, a given value of α_1 does not imply a TCP source slow-down of equivalent magnitude. Here, the issue is complicated by the many features of the TCP protocol, such as slow-start/congestion-avoidance phases, the "ACK-every-other property", delayed ACK timers, etc. For example, in the idealized case of infinite sources sending full-sized segments, with the ubiquitous "ACK-every-other" feature enabled, it can be shown that a value of $\alpha_1 > 3$ is required to throttle a TCP source. Alternatively, if the TCP source's end-system behaviors are unknown, then very large values of α_1 can be used to "guarantee" queue length control. In other words, such values essentially inhibit all ACK emissions until congestion subsides (i.e., on/off type control), but usually give increased queue oscillations.

Note that in order to present a generic, more flexible specification, the pseudocode in Fig. 6 does not explicitly specify the congestion detection method.

Specifically, the congestion_status() routine 620 simply returns a boolean value indicating whether or not the link layer entity is congested. Clearly a whole variety of congestion indication mechanisms can be used here. Some examples include queue lengths, averaged queue lengths, input rate overload measurements, and data loss rates. However, preferably the queue length should be used to simplify implementation complexities.

Pseudocode for a sample method 700 using two hysteresis queue thresholds, QL 710, QH 714, is illustrated in Fig. 7. In Fig. 7, the congestion status is checked and a binary flag is returned 702. Hysteresis queue thresholds, QL 710, QH 712 are used 720. If congestion exists and the queue length is less than QL 722, then congestion abatement status change is stored by setting the flag to a first state 724, i.e., congested_flag=OFF. Alternatively, if congestion does not exist and the queue length is greater than QH 730, the congestion onset status change is stored by setting the flag to a second state 732, i.e., the congested_flag=ON. The state of the binary flag is then returned 740.

Results show that if these thresholds 710, 712 are appropriately sized based upon the round-trip delays between the sources and access nodes (i.e., access network delays), near loss-less performance can be achieved. Since such delays are usually many times smaller than the end-to-end delays observed in WAN networks, sizeable reductions in the buffering requirements are possible with ACK pacing schemes according to the present invention.

Referring again to Fig. 6, the data packet departure method 660 is executed whenever a packet departs the link layer. The goal is to release stored ACK packets in a timely fashion, thereby properly controlling the congestion (queue) levels at the

access node's link-layer buffer, i.e., minimizing packet losses. The method first checks to see if there are any buffered ACK packets awaiting transmission 662 and whether the ACK emission counter, i.e., `pkt_counter`, has reached zero 664. If this is the case, a buffered ACK packet is released to the source 670.

5 After this, if congestion still exists 680, the inter-ACK packet spacing is maintained at one per α_1 data packets by resetting `pk_counter` to α_1 682. This allows the data buffer in the data link layer to drain further. If congestion has abated 684, however, then the inter-ACK spacing is reduced to α_2 data packets 686, allowing sources to send faster. The α_2 parameter is termed an integral speedup factor, and
10 necessarily $\alpha_1 > \alpha_2$. If the counter is non-zero 688, then it is simply decremented 690. However, to prevent bandwidth under-utilization after congestion periods, if the counter is larger than α_2 , it is simply reset to α_2 (i.e., especially for larger α_2 values).

 Again, due to complications arising from TCP specifics, an $\alpha_2 = 1$ value does
15 not imply that TCP source rates will (approximately) equal the underlying link entity's rate. More specifically, for idealized conditions with the "ACK-every-other" feature, a value of $\alpha_2 = 2$ performs better.

 Referring again to Fig. 5, note that the eligible ACK packets 532 are chosen based upon the queuing strategy used by the ACK scheduler 520. This overall
20 mechanism does not require any expensive timer mechanisms to release stored ACK packets 532 as required in prior methods.

 Fig. 8 illustrates the pseudocode 800 for the enhanced TCP ACK arrival 810 and ATM cell departure 860 methods according to the present invention. In the TCP ACK arrival process 810, a determination is made as to whether the link-layer entity

is congested or whether the ACK buffer is non-empty 812. If the link layer entity is congested and the ACK buffer is not empty 814, incoming ACK are stored in the queue 816 (FIFO, per-class, or per-flow). ACK packets are stored at the tail of the respective ACK queues and the ACK count is incremented 818. Next, a check is
5 made to determine if this ACK packet is the first ACK packet buffered 820. If this is the first ACK packet to be buffered 822, the cell counter is set to α_1 * packet cells so that the counter is set to a larger spacing 824. Otherwise 826, the ACK packet is sent to the TCP source 828.

In the ATM Cell departure process 860, a determination is made as to
10 whether the ACK buffer is non-empty 862, i.e., are there ACK packets to send? If there are ACK packets to send 864, scheduler determines the next eligible ACK packet 870. The next eligible ACK packet is dequeued from the head of the eligible ACK queue 872 and is sent to the TCP source 874. The ACK count is decremented 876 and the cell counter is reset appropriately 880. If congestion exists 882, the cell
15 counter is set to equal α_1 * packet cells to increase the spacing 884. Otherwise 886, the cell counter is set to equal α_2 * packet cells to reduce the spacing 888.

If the cell counter is not zero 890, a determination is made as to whether congestion has abated 892. If congestion has abated and the cell counter value is greater than α_2 * packet cells 894, then the cell counter is set to equal α_2 * packet
20 cells 895. Otherwise 896, the cell counter is decremented 898.

Granted that the above packet handling methods are very generic, more flexibility exists for the case of ATM networks which use smaller packet (cell) sizes. Specifically, it is possible to perform ACK emission per data (fractional) packet emission and at the same time circumvent the use of any expensive timer

mechanisms, i.e., the counter is now in terms of cells not packets (cell counter).

Since a cell size is typically much smaller than a TCP MSS-sized packet, packets can now be emitted with more fined-grained time granularities. Namely, the α_1 and α_2 factors do not have to be integers anymore, as is the case in packet-based

5 schemes. Consider a constant value, *packet_cells*, namely the number of cells in (IP) TCP MSS:

$$packet_cells = \left(\left\lceil \frac{(TCP_MSS + 40)}{48} \right\rceil + 1 \right). \quad (1)$$

10 With reference again to Fig. 5, during congestion ACK packets 532 are emitted after every $\alpha_1 \cdot packet_cells$ and during underload, $\alpha_2 \cdot packet_cells$. For the most part, the ACK arrival and cell departure methods are identical to their packet-based counterparts. For example, after every cell emission, the cell counter is decremented and when it reaches zero, a ACK packet 532 in buffer 534 is released
15 by the ACK scheduler 520. During period when the ACK buffer 534 is empty, the counter value is reset appropriately.

In summary, the performance of the TCP protocol over ATM networks is an important area. Recently, various ACK pacing schemes have been proposed to improve TCP's interaction with the more advanced underlying ATM transport
20 categories (i.e., ABR flow control, per-connection queuing). However, these schemes suffer from parameter sensitivity issues and may be difficult to realize in practice. Accordingly, an enhanced ACK pacing device has been disclosed that is capable of performing in a wide range of network scenarios. The scheme uses (more direct) queue-length congestion information to delay TCP ACK packets and can

implement a wide range of fairness criterion. The method provides a robust means of improving end-to-end TCP throughput and bandwidth fairness. The buffering requirements in the access nodes are also significantly for a wide range of subnetworks.

- 5 The foregoing description of the exemplary embodiment of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not with this detailed description,
- 10 but rather by the claims appended hereto.

WHAT IS CLAIMED IS:

- 1 1. An acknowledgment pacing device for pacing acknowledgment
2 packets to be sent to a source in response to receiving data packets from the source,
3 comprising:
4 an acknowledgment control unit for monitoring loading of a network and
5 generating a control signal for controlling the processing of acknowledgment packets
6 based upon the loading of the network;
7 an acknowledgment packet buffer, coupled to the acknowledgment control
8 unit, for storing acknowledgment packets received from the acknowledgment control
9 unit; and
10 a scheduler, coupled to the acknowledgment control unit and the
11 acknowledgment buffer, the scheduler releasing acknowledgment packets based
12 upon the control signal generated by the acknowledgment control unit.
- 1 2. The acknowledgment pacing device of claim 1 where the scheduler
2 chooses acknowledgment packets to release based upon a queuing strategy.
- 1 3. The acknowledgment pacing device of claim 2 wherein the queuing
2 strategy comprises sending a head-of-line acknowledgment packet when aggregate
3 acknowledgment packet buffering is used.
- 1 4. The acknowledgment pacing device of claim 2 wherein the queuing
2 strategy comprises a weighted-round-robin (WRR) process for selecting an
3 acknowledgment packet in the acknowledgment packet buffer for release when per-
4 class or per-flow acknowledgment packet buffering is used.

1 5. The acknowledgment pacing device of claim 4 wherein the weighted-
2 round-robin (WRR) process uses weights for weighting the selection of the
3 acknowledgment packet for release inversely proportionally to a TCP maximum
4 segment size (MSS) to the lessen a bias against smaller MSS flows.

1 6. The acknowledgment pacing device of claim 2 wherein the queuing
2 strategy comprises a fair-queuing (FQ) process for selecting an acknowledgment
3 packet in the acknowledgment packet buffer for release when per-class or per-flow
4 acknowledgment packet buffering is used.

1 7. The acknowledgment pacing device of claim 1 wherein the
2 acknowledgment control unit further comprises an acknowledgment packet pacing
3 processor, the acknowledgment packet pacing processor generating the control
4 signal for controlling the processing of acknowledgment packets using a
5 acknowledgment packet arrival processor and a data packet departure processor.

1 8. The acknowledgment pacing device of claim 7 wherein the
2 acknowledgment packet arrival processor controls the processing of
3 acknowledgment packets to the acknowledgment packet buffer by checking a
4 congestion level of the network and deciding whether to hold acknowledgment
5 packets in the acknowledgment packet buffer or to send the acknowledgment packet
6 directly to the source without buffering the acknowledgment packets in the
7 acknowledgment packet buffer.

1 9. The acknowledgment pacing device of claim 7 wherein the
2 acknowledgment packet arrival processor decides whether to hold acknowledgment
3 packets in the acknowledgment packet buffer or to send the acknowledgment packet
4 directly to the source without buffering the acknowledgment packets in the
5 acknowledgment packet buffer by determining if the network is congested,
6 determining if the acknowledgment packet buffer is empty, storing an
7 acknowledgment packet in the buffer if the network is congested or the
8 acknowledgment packet buffer is not empty, and forwarding the acknowledgment
9 packet to the source if the acknowledgment packet buffer is empty and the network
10 is not congested.

1 10. The acknowledgment pacing device of claim 9 wherein the
2 acknowledgment packet is stored in the acknowledgment packet buffer and gated out
3 by the scheduler if the acknowledgment packet is a first acknowledgment packet to
4 be buffered in the acknowledgment packet buffer during congestion.

1 11. The acknowledgment pacing device of claim 10 wherein the
2 acknowledgment control unit increases a spacing between acknowledgment packets
3 gated from the acknowledgment packet buffer if the acknowledgment packet is the
4 first acknowledgment packet to be buffered in the acknowledgment packet buffer
5 during congestion.

1 12. The acknowledgment pacing device of claim 11 wherein the
2 acknowledgment control unit increases the spacing between acknowledgment
3 packets by setting a packet counter variable to a first predetermined value.

1 13. The acknowledgment pacing device of claim 12 wherein the packet
2 counter variable is decremented as when a data packet departs from the network.

1 14. The acknowledgment pacing device of claim 13 wherein the
2 scheduler releases a buffered acknowledgment packet when the packet counter
3 variable is decremented to zero and the acknowledgment control unit resets the
4 packet counter variable.

1 15. The acknowledgment pacing device of claim 8 wherein the
2 congestion level of the network is determined by analyzing a queue length
3 representing a capacity for a data packet buffer.

1 16. The acknowledgment pacing device of claim 15 wherein the network
2 is indicated as being non-congested when the queue length is less than a low
3 threshold.

1 17. The acknowledgment pacing device of claim 15 wherein the network
2 is indicated as being congested when the queue length is greater than a high
3 threshold.

1 18. The acknowledgment pacing device of claim 7 wherein the data
2 packet departure processor controls the release of acknowledgment packets from the
3 acknowledgment packet buffer by monitoring congestion levels of the network and
4 deciding when to gate acknowledgment packets from the acknowledgment buffer to
5 the source.

1 19. The acknowledgment pacing device of claim 18 wherein the data
2 packet departure processor decides when to gate acknowledgment packets from the
3 acknowledgment buffer to the source by checking if acknowledgment packets are in
4 the acknowledgment packet buffer awaiting transmission and if a packet counter
5 variable set by the acknowledgment control unit has a value of zero, and releasing a
6 buffered acknowledgment packet in the acknowledgment packet buffer to the source
7 when the packet counter variable has a value of zero.

1 20. The acknowledgment pacing device of claim 19 wherein the data
2 packet departure processor increases the spacing between the release of
3 acknowledgment packets from the acknowledgment packet buffer if congestion still
4 exists in the network.

1 21. The acknowledgment pacing device of claim 20 wherein the spacing
2 between the release of acknowledgment packets is increased by resetting the packet
3 counter variable to a first predetermined value.

1 22. The acknowledgment pacing device of claim 21 wherein the data
2 packet departure processor decrementing the packet counter variable if the value of
3 the packet counter variable is non-zero.

1 23. The acknowledgment pacing device of claim 22 wherein the data
2 packet departure processor resets the packet counter variable to the second
3 predetermined value to prevent bandwidth under-utilization after congestion periods
4 if the packet counter variable is larger than the second predetermined value.

1 24. The acknowledgment pacing device of claim 19 wherein the data
2 packet departure processor decreases a spacing between the release of
3 acknowledgment packets from the acknowledgment packet buffer if congestion in
4 the network has abated.

1 25. The acknowledgment pacing device of claim 24 wherein the data
2 packet departure processor decreases a spacing between the release of
3 acknowledgment packets by resetting the packet counter variable to a second
4 predetermined value, the second predetermined value being less than the first
5 predetermined value.

1 26. The acknowledgment pacing device of claim 25 wherein the data
2 packet departure processor decrementing the packet counter variable if the value of
3 the packet counter variable is non-zero.

1 27. The acknowledgment pacing device of claim 26 wherein the data
2 packet departure processor resets the packet counter variable to the second
3 predetermined value to prevent bandwidth under-utilization after congestion periods
4 if the packet counter variable is larger than the second predetermined value.

1 28. The acknowledgment pacing device of claim 19 wherein the
2 congestion level of the network is determined by analyzing a queue length
3 representing a capacity for a data packet buffer.

1 29. The acknowledgment pacing device of claim 28 wherein the network
2 is indicated as being non-congested when the queue length is less than a low
3 threshold.

1 30. The acknowledgment pacing device of claim 28 wherein the network
2 is indicated as being congested when the queue length is greater than a high
3 threshold.

1 31. The acknowledgment pacing device of claim 1 wherein the
2 acknowledgment packet buffer buffers acknowledgment packets on an aggregate
3 basis.

1 32. The acknowledgment pacing device of claim 1 wherein the
2 acknowledgment packet buffer buffers acknowledgment packets by flow type, and
3 wherein the scheduler releases the acknowledgment packets in the acknowledgment
4 buffer taking into account the type of flows for the buffered acknowledgment
5 packets.

1 33. An access node device, comprising;
2 a link layer entity for receiving data packets from a source and forwarding
3 the data packets to a forward data link, the link layer entity storing the received data
4 packets in a data packet buffer until the data packets depart the link layer entity and
5 are forwarded to the forward data link; and
6 an acknowledgment pacing device, coupled to the link layer entity, for pacing
7 acknowledgment packets to be sent to the source in response to receiving the data
8 packets from the source, the acknowledgment pacing device further comprising:
9 an acknowledgment control unit for monitoring congestion at the link
10 layer entity and generating a control signal for controlling the processing of
11 acknowledgment packets based upon whether congestion is occurring at the link
12 layer entity;
13 an acknowledgment packet buffer, coupled to the acknowledgment
14 control unit, for storing acknowledgment packets received from the acknowledgment
15 control unit; and
16 a scheduler, coupled to the acknowledgment control unit and the
17 acknowledgment buffer, the scheduler releasing acknowledgment packets to the
18 source based upon the control signal generated by the acknowledgment control unit.

1 34. The access node device of claim 33 where the scheduler chooses
2 acknowledgment packets to be released based upon a queuing strategy.

1 35. The access node device of claim 34 wherein the queuing strategy
2 comprises sending a head-of-line acknowledgment packet when aggregate
3 acknowledgment packet buffering is used.

1 36. The access node device of claim 34 wherein the queuing strategy
2 comprises a weighted-round-robin (WRR) process for selecting an acknowledgment
3 packet in the acknowledgment packet buffer for release when per-class or per-flow
4 acknowledgment packet buffering is used.

1 37. The access node device of claim 36 wherein the weighted-round-
2 robin (WRR) process uses weights for weighting the selection of the
3 acknowledgment packet for release inversely proportionally to a TCP maximum
4 segment size (MSS) to the lessen a bias against smaller MSS flows.

1 38. The access node device of claim 34 wherein the queuing strategy
2 comprises a fair-queuing (FQ) process for selecting an acknowledgment packet in
3 the acknowledgment packet buffer for release when per-class or per-flow
4 acknowledgment packet buffering is used.

1 39. The access node device of claim 33 wherein the acknowledgment
2 control unit further comprises an acknowledgment packet pacing processor, the
3 acknowledgment packet pacing processor generating the control signal for
4 controlling the processing of acknowledgment packets using a acknowledgment
5 packet arrival processor and a data packet departure processor.

1 40. The access node device of claim 39 wherein the acknowledgment
2 packet arrival processor controls the processing of acknowledgment packets to the
3 acknowledgment packet buffer by checking the congestion at the link layer entity
4 and deciding whether to hold acknowledgment packets in the acknowledgment
5 packet buffer or to send the acknowledgment packet directly to the source without
6 buffering the acknowledgment packets in the acknowledgment packet buffer.

1 41. The access node device of claim 39 wherein the acknowledgment
2 packet arrival processor decides whether to hold acknowledgment packets in the
3 acknowledgment packet buffer or to send the acknowledgment packet directly to the
4 source without buffering the acknowledgment packets in the acknowledgment
5 packet buffer by determining if the link layer entity is congested, determining if the
6 acknowledgment packet buffer is empty, storing an acknowledgment packet in the
7 buffer if the link layer entity is congested or the acknowledgment packet buffer is
8 not empty, and forwarding the acknowledgment packet to the source if the
9 acknowledgment packet buffer is empty and the link layer entity is not congested.

1 42. The access node device of claim 41 wherein the acknowledgment
2 packet is stored in the acknowledgment packet buffer and gated out by the scheduler
3 if the acknowledgment packet is a first acknowledgment packet to be buffered in the
4 acknowledgment packet buffer during congestion.

1 43. The access node device of claim 42 wherein the acknowledgment
2 control unit increases a spacing between acknowledgment packets gated from the
3 acknowledgment packet buffer if the acknowledgment packet is the first
4 acknowledgment packet to be buffered in the acknowledgment packet buffer during
5 congestion.

1 44. The access node device of claim 39 wherein the data packet departure
2 processor decreases a spacing between the release of acknowledgment packets from
3 the acknowledgment packet buffer if congestion in the link layer entity has abated.

1 45. A method for providing acknowledgment pacing for acknowledgment
2 packets to be sent to a source in response to receiving data packets from the source,
3 comprising:
4 monitoring loading of a network;
5 generating a control signal for controlling the processing of acknowledgment
6 packets based upon the loading of the network;
7 storing acknowledgment packets received from the acknowledgment control
8 unit in an acknowledgment packet buffer ; and
9 releasing acknowledgment packets based upon the control signal.

1 46. The method of claim 45 wherein the releasing further comprises
2 choosing acknowledgment packets to release based upon a queuing strategy.

1 47. The method of claim 46 wherein the queuing strategy comprises
2 sending a head-of-line acknowledgment packet when aggregate acknowledgment
3 packet buffering is used.

1 48. The method of claim 46 wherein the queuing strategy comprises a
2 weighted-round-robin (WRR) process for selecting an acknowledgment packet in the
3 acknowledgment packet buffer for release when per-class or per-flow
4 acknowledgment packet buffering is used.

1 49. The method of claim 48 wherein the weighted-round-robin (WRR)
2 process uses weights for weighting the selection of the acknowledgment packet for
3 release inversely proportionally to a TCP maximum segment size (MSS) to the
4 lessen a bias against smaller MSS flows.

1 50. The method of claim 46 wherein the queuing strategy comprises a
2 fair-queuing (FQ) process for selecting an acknowledgment packet in the
3 acknowledgment packet buffer for release when per-class or per-flow
4 acknowledgment packet buffering is used.

1 51. The method of claim 45 wherein the generating the control signal for
2 controlling the processing of acknowledgment packets comprises a acknowledgment
3 packet arrival processes and a data packet departure process.

1 52. The method of claim 51 wherein the acknowledgment packet arrival
2 process controls the processing of acknowledgment packets to the acknowledgment
3 packet buffer by checking a congestion level of the network and deciding whether to
4 hold acknowledgment packets in the acknowledgment packet buffer or to send the
5 acknowledgment packet directly to the source without buffering the
6 acknowledgment packets in the acknowledgment packet buffer.

1 53. The method of claim 51 wherein the deciding whether to hold
2 acknowledgment packets in the acknowledgment packet buffer or to send the
3 acknowledgment packet directly to the source without buffering the
4 acknowledgment packets in the acknowledgment packet buffer further comprises
5 determining if the network is congested, determining if the acknowledgment packet
6 buffer is empty, storing an acknowledgment packet in the buffer if the network is
7 congested or the acknowledgment packet is not empty, and forwarding the
8 acknowledgment packet to the source if the acknowledgment packet buffer is empty
9 and the network is not congested.

1 54. The method of claim 53 wherein the acknowledgment packet is
2 stored in the acknowledgment packet buffer and gated out if the acknowledgment
3 packet is a first acknowledgment packet to be buffered in the acknowledgment
4 packet buffer.

1 55. The method of claim 54 further comprising increasing a spacing
2 between acknowledgment packets gated from the acknowledgment packet buffer if
3 the acknowledgment packet is the first acknowledgment packet to be buffered in the
4 acknowledgment packet buffer.

1 56. The method of claim 55 wherein increasing the spacing further
2 comprises setting a packet counter variable to a first predetermined value.

1 57. The method of claim 56 further comprising decrementing the packet
2 counter variable when a data packet departs from the network.

1 58. The method of claim 57 further comprising releasing a buffered
2 acknowledgment packet when the packet counter variable is decremented to zero
3 and resetting the packet counter variable.

1 59. The method of claim 52 wherein the congestion level of the network
2 is determined by analyzing a queue length representing a capacity for a data packet
3 buffer.

1 60. The method of claim 59 further comprising indicating the network is
2 not congested when the queue length is less than a low threshold.

1 61. The method of claim 59 further comprising indicating the network is
2 congested when the queue length is greater than a high threshold.

1 62. The method of claim 51 wherein the releasing further comprises
2 monitoring congestion levels of the network and deciding when to gate
3 acknowledgment packets from the acknowledgment buffer to the source.

1 63. The method of claim 62 wherein the deciding further comprises
2 checking if acknowledgment packets are in the acknowledgment packet buffer
3 awaiting transmission and if a packet counter variable set by the acknowledgment
4 control unit has a value of zero, and releasing a buffered acknowledgment packet in
5 the acknowledgment packet buffer to the source when the packet counter variable
6 has a value of zero.

1 64. The method of claim 63 wherein the data packet departure process
2 further comprises increasing the spacing between the release of acknowledgment
3 packets from the acknowledgment packet buffer if congestion still exists in the
4 network.

1 65. The method of claim 64 wherein the increasing further comprises
2 resetting the packet counter variable to a first predetermined value.

1 66. The method of claim 65 wherein the data packet departure process
2 further comprises decrementing the packet counter variable if the value of the packet
3 counter variable is non-zero.

1 67. The method of claim 66 wherein the data packet departure process
2 further comprises resetting the packet counter variable to the second predetermined
3 value to prevent bandwidth under-utilization after congestion periods if the packet
4 counter variable is larger than the second predetermined value.

1 68. The method of claim 63 wherein the data packet departure process
2 further comprises decreasing a spacing between the release of acknowledgment
3 packets from the acknowledgment packet buffer if congestion in the network has
4 abated.

1 69. The method of claim 68 wherein the data packet departure process
2 further comprises decreasing a spacing between the release of acknowledgment
3 packets by resetting the packet counter variable to a second predetermined value, the
4 second predetermined value being less than the first predetermined value.

1 70. The method of claim 69 wherein the data packet departure process
2 further comprises decrementing the packet counter variable if the value of the packet
3 counter variable is non-zero.

1 71. The method of claim 70 wherein the data packet departure process
2 further comprises resetting the packet counter variable to the second predetermined
3 value to prevent bandwidth under-utilization after congestion periods if the packet
4 counter variable is larger than the second predetermined value.

1 72. The method of claim 63 wherein the determining of the congestion
2 level of the network further comprises analyzing a queue length representing a
3 capacity for a data packet buffer.

1 73. The method of claim 72 further comprising indicating the network is
2 not congested when the queue length is less than a low threshold.

1 74. The method of claim 72 further comprising indicating the network is
2 congested when the queue length is greater than a high threshold.

1 75. The method of claim 45 wherein the storing further comprises
2 buffering acknowledgment packets on an aggregate basis.

1 76. The method of claim 45 wherein the storing further comprises
2 buffering acknowledgment packets by flow type, and wherein the releasing further
3 comprises scheduling the release of acknowledgment packets in the
4 acknowledgment buffer by taking into account the type of flows for the buffered
5 acknowledgment packets.




Fig. 1




Fig. 2




Fig. 3




Fig. 4





Fig. 5



700)

```

702 /* Check congestion status and return binary flag */
boolean congestion_status()
710 {
712     /* Use hysteresis queue thresholds (QL, QH) */
720 if ((congested_flag==ON) && (queue_length < QL))
722 {
724     /* Store congestion abatement status change in flag */
congested_flag=OFF
730 }
732 else if ((congested_flag==OFF) && (queue_length > QH))
740 {
/* Store congestion onset status change in flag */
congested_flag=ON
}
/* Return congestion status */
return (congested_flag)
}

```

Fig. 7

800




Fig. 8