

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 June 2011 (30.06.2011)

PCT

(10) International Publication Number
WO 2011/078995 A1

- (51) International Patent Classification:
G01R 31/08 (2006.01)
- (21) International Application Number:
PCT/US2010/060517
- (22) International Filing Date:
15 December 2010 (15.12.2010)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
12/647,426 25 December 2009 (25.12.2009) US
- (71) Applicant: CISCO TECHNOLOGY, INC. [US/US];
170 Tasman Drive, San Jose, California 95134 (US).
- (72) Inventors: DUKKIPATI, Nandita; 1003 Peggy Lane,
Menlo Park, California 94025 (US). HA, Sangtae; 465
Meadow Road Apt #5101, Princeton, New Jersey 08540
(US). SUBRAMANIAN, Vijaynarayanan; 1063 More
Ave #9-111, Sunnyvale, California 94089 (US). BONO-
MI, Flavio Giovanni; 526 Lowell Ave, Palo Alto, Cali-
fornia 94301 (US).
- (74) Agents: WILLIAMS, Kirk et al.; P.O. Box 39425, Den-
ver, Colorado 80239-0425 (US).
- (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ,

CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP,
KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,
NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD,
SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG,
ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted
a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of
the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))

(54) Title: INCREASING TRANSMISSION RATE TO A REMOTE DEVICE




FIGURE 1

(57) Abstract: In response to a detected loss of previously transmitted information by an apparatus communicating with a remote device (e.g., using TCP), the rate of transmission of information is increased by the apparatus in response to attributing the detected loss of previously transmitted information as not being caused by congestion. This attribution of the packet loss is typically determined based on roundtrip delays between sent information and received corresponding acknowledgments, which may be used directly or indirectly, such as by estimating network queuing delays based on the measured roundtrip delays.

WO 2011/078995 A1

INCREASING TRANSMISSION RATE TO A REMOTE DEVICE

FIELD OF THE INVENTION

5 The present disclosure relates generally to communicating information with remote device over a network.

10 BACKGROUND OF THE INVENTION

 The communications industry is rapidly changing to adjust to emerging technologies and ever increasing customer demand. This customer demand for new applications and increased performance of existing applications is driving communications network and system providers to employ networks and systems having greater speed and capacity (e.g., greater bandwidth). In trying to achieve these goals, a common approach
15 taken by many communications providers is to use packet switching technology.

 Transmission Control Protocol (TCP) is used extensively in networks, such as the Internet, to reliably communicate information between devices. This reliability is based on the sender receiving acknowledgements from the receiver in response to successfully
20 receiving packets. The performance of TCP also depends upon whether or not acknowledgements are received for sent packets. In response to an acknowledgment, the sender TCP slowly increases the rate at which it sends packets. Conversely, if the sender TCP does not receive such an acknowledgment for a sent packet, it decreases, such as by
25 half, its rate of packet transmission based on the assumption that the packet was lost due to network congestion. Note, the sender TCP will continue slowly increasing its transmission rate of packets in response to receiving an acknowledgement of a sent packet.

SUMMARY OF THE INVENTION

Disclosed are, *inter alia*, methods, apparatus, computer-storage media, mechanisms, and means associated with the rate of transmitting information to a remote device in response to attributing a detected loss of previous information sent to the remote device as not being a result of network congestion. One embodiment includes an apparatus, comprising one or more processing elements and memory configured to communicate information with a remote device over a network, which includes: in response to a detected loss of previously transmitted information by the apparatus, increasing the rate of transmission of information by the apparatus in response to attributing the detected loss of previously transmitted information as not being caused by congestion.

In one embodiment, this communication of information is performed using Transmission Control Protocol (TCP), and wherein said attributing of the detected loss of previously transmitted information as not being caused by congestion includes measuring a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information. In one embodiment, said attributing of the detected loss of previously transmitted information as not being caused by congestion includes estimating that queuing delays within the network between the apparatus and the remote device have not increased therefore not causing the detected loss. In one embodiment, a current queuing delay is determined based on a minimum of measured roundtrip delays, an average of measured roundtrip delays, and/or a median, smooth-filtered maximum of measured roundtrip delays.

BRIEF DESCRIPTION OF THE DRAWINGS

The appended claims set forth the features of the invention with particularity. The invention, together with its advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

- FIG. 1 illustrates a network operating according to one embodiment;
- FIG. 2 illustrates an apparatus or component used in one embodiment;
- FIG. 3 illustrates a process performed in one embodiment;
- FIG. 4 illustrates a process performed in one embodiment;
- FIG. 5 illustrates a process performed in one embodiment; and
- FIGs. 6A-B illustrate a process performed in one embodiment.

DETAILED DESCRIPTION

Disclosed are, *inter alia*, methods, apparatus, computer-storage media, mechanisms, and means associated with the rate of transmitting information to a remote device in response to attributing a detected loss of previous information sent to the remote device as not being a result of network congestion. Embodiments described herein include various elements and limitations, with no one element or limitation contemplated as being a critical element or limitation. Each of the claims individually recites an aspect of the invention in its entirety. Moreover, some embodiments described may include, but are not limited to, *inter alia*, systems, networks, integrated circuit chips, embedded processors, ASICs, methods, and computer-readable media containing instructions. One or multiple systems, devices, components, etc. may comprise one or more embodiments, which may include some elements or limitations of a claim being performed by the same or different systems, devices, components, etc. A processing element may be a general processor, task-specific processor, or other implementation for performing the corresponding processing. The embodiments described hereinafter embody various aspects and configurations within the scope and spirit of the invention, with the figures illustrating exemplary and non-limiting configurations. Note, computer-readable media and means for performing methods and processing block operations (e.g., a processor and memory or other apparatus configured to perform such operations) are disclosed and are in keeping with the extensible scope and spirit of the invention. Note, the term "apparatus" is used consistently herein with its common definition of an appliance or device.

Note, the steps, connections, and processing of signals and information illustrated in the figures, including, but not limited to any block and flow diagrams and message sequence charts, may typically be performed in the same or in a different serial or parallel ordering and/or by different components and/or processes, threads, etc., and/or over different connections and be combined with other functions in other embodiments, unless this disables the embodiment or a sequence is explicitly or implicitly required (e.g., for a sequence of read the value, process said read value - the value must be obtained prior to processing it, although some of the associated processing may be performed prior to, concurrently with, and/or after the read operation). Also note, nothing described or referenced in this document is admitted as prior art to this application unless explicitly so stated.

The term "one embodiment" is used herein to reference a particular embodiment, wherein each reference to "one embodiment" may refer to a different embodiment, and the use of the term repeatedly herein in describing associated features, elements and/or limitations does not establish a cumulative set of associated features, elements and/or limitations that each and every embodiment must include, although an embodiment typically may include all these features, elements and/or limitations. In addition, the terms "first," "second," etc. are typically used herein to denote different units (e.g., a first element, a second element). The use of these terms herein does not necessarily connote an ordering such as one unit or event occurring or coming before another, but rather provides a mechanism to distinguish between particular units. Moreover, the phrases "based on x" and "in response to x" are used to indicate a minimum set of items "x" from which something is derived or caused, wherein "x" is extensible and does not necessarily describe a complete list of items on which the operation is performed, etc. Additionally, the phrase "coupled to" is used to indicate some level of direct or indirect connection between two elements or devices, with the coupling device or devices modifying or not modifying the coupled signal or communicated information. Moreover, the term "or" is used herein to identify a selection of one or more, including all, of the conjunctive items. Additionally, the transitional term "comprising," which is synonymous with "including," "containing," or "characterized by," is inclusive or open-ended and does not exclude additional, unrecited elements or method steps. Finally, the term "particular machine," when recited in a method claim for performing steps, refers to a particular machine within the 35 USC § 101 machine statutory class.

Disclosed are, *inter alia*, methods, apparatus, computer-storage media, mechanisms, and means associated with the rate of transmitting information to a remote device in response to attributing a detected loss of previous information sent to the remote device as not being a result of network congestion. In one embodiment, in response to a detected loss of previously transmitted information by an apparatus communicating with a remote device (e.g., using TCP), the rate of transmission of information is increased by the apparatus in response to attributing the detected loss of previously transmitted information as not being caused by congestion. This attribution of the packet loss is typically determined based on roundtrip delays between sent information and received

corresponding acknowledgments, which may be used directly or indirectly, such as by estimating network queuing delays based on the measured roundtrip delays.

One embodiment includes an apparatus, comprising: one or more processing elements and memory configured to communicate information with a remote device over a network; wherein said communication with the remote device includes, in response to a detected loss of previously transmitted information by the apparatus: increasing the rate of transmission of information by the apparatus in response to attributing the detected loss of previously transmitted information as not being caused by congestion.

In one embodiment, said communication of information is performed using Transmission Control Protocol (TCP), and wherein said attributing of the detected loss of previously transmitted information as not being caused by congestion includes measuring a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information.

In one embodiment, said communication using TCP with the remote device includes a loss recovery phase of said TCP communications triggered by the detected loss of previously transmitted information, with the loss recovery phase including performing active congestion control, with said active congestion control including said increasing the rate of transmission of information by the apparatus in response to attributing said detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information; and wherein the loss recovery phase includes reducing the rate of transmission of information by the apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion. In one embodiment, communication using TCP with the remote device includes using a congestion window (cwnd) for reacting to network conditions; and wherein said increasing the rate of transmission of information includes increasing the congestion window (cwnd); and wherein said reducing the rate of transmission includes decreasing the congestion window (cwnd). In one embodiment, said attributing the detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an

acknowledgement received by the apparatus corresponding to said sent information includes determining a plurality of queuing delays based on the plurality of measured roundtrip delays and said attributing said detected loss of previously transmitted information as not being caused by congestion based on the plurality of queuing delays.

5 In one embodiment, said attributing the detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information includes determining a plurality of queuing delays based on the plurality of measured
10 roundtrip delays and said attributing said detected loss of previously transmitted information as not being caused by congestion based on the plurality of queuing delays. In one embodiment, a current queuing delay is determined based on a minimum of a plurality of said measured roundtrip delays and one or more of said measured roundtrip delays. In one embodiment, the current queuing delay is determined based on a difference of said one
15 or more of said measured roundtrip delays with said minimum of a plurality of said measured roundtrip delays. In one embodiment, the current queuing delay is determined based on an average of a plurality of said measured roundtrip delays, a minimum of a plurality of said measured roundtrip delays, and a median, smooth-filtered maximum of a plurality of said measured roundtrip delays.

20 In one embodiment, said attributing of the detected loss of previously transmitted information as not being caused by congestion includes estimating that queuing delays within the network between the apparatus and the remote device have not increased therefore not causing the detected loss. In one embodiment, said communication of information is performed using Transmission Control Protocol (TCP). In one embodiment,
25 a current queuing delay of said queuing delays is determined based on a difference of one or more measured roundtrip delays with a minimum of a plurality of measured roundtrip delays. In one embodiment, a current queuing delay of said queuing delays is determined based on an average of a plurality of a plurality of measured roundtrip delays, a minimum of a plurality of the plurality of measured roundtrip delays, and a median, smooth-filtered
30 maximum of a plurality of the plurality of measured roundtrip delays.

On embodiment includes a method, comprising: communicating, by a particular apparatus, with a remote device over a network; wherein said communication with the

remote device includes, in response to a detected loss of previously transmitted information by the particular apparatus: increasing the rate of transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as not being caused by congestion based on a current queuing delay determined based on measured roundtrip delays between information sent from the particular apparatus to the remote device and an acknowledgement received by the particular apparatus corresponding to said sent information.

In one embodiment, in response to the detected loss of previously transmitted information by the particular apparatus: decreasing the rate of transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion based on the current queuing delay.

In one embodiment, said communicating with the remote device over the network includes using Transmission Control Protocol (TCP). In one embodiment, communicating using TCP with the remote device includes using a congestion window (cwnd) for reacting to

network conditions; and wherein said increasing the rate of transmission of information includes increasing the congestion window (cwnd); and wherein said decreasing the rate of transmission of information includes decreasing the congestion window (cwnd). One

embodiment includes determining the current queuing delay based on a minimum of a plurality of said measured roundtrip delays and one or more of said measured roundtrip

delays. One embodiment includes determining the current queuing delay based on a difference of said one or more of said measured roundtrip delays with said minimum of a plurality of said measured roundtrip delays. One embodiment includes determining the

current queuing delay based on an average of a plurality of said measured roundtrip delays, a minimum of a plurality of said measured roundtrip delays, and a median, smooth-filtered maximum of a plurality of said measured roundtrip delays. In one embodiment, in response

to the detected loss of previously transmitted information by the particular apparatus: decreasing the rate of transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as being caused by

congestion based on the current queuing delay. In one embodiment, said attributing said detected loss of previously transmitted information as not being caused by congestion

based on the current queuing delay includes determining that an average queuing delay is not greater than a previously determined average queuing delay not including the current

queuing delay. One embodiment includes: in response to the detected loss of previously transmitted information by the particular apparatus: decreasing the rate of transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion based on the current
5 queuing delay.

One embodiment includes an apparatus, comprising: one or more processing elements and memory configured to communicate using Transmission Control Protocol (TCP) with a remote device over a network; wherein said communication using TCP with the remote device includes using a congestion window (cwnd) for reacting to network
10 conditions; wherein said communication with the remote device using TCP includes characterizing the roundtrip delay between the apparatus and the remote device based on measured roundtrip delay for a plurality of sent information and corresponding received acknowledgements; and wherein said communication with the remote device using TCP includes: in response to an identified loss of particular information sent to the remote
15 device and a determination based on said maintained characterization of the roundtrip delay, increasing the congestion window (cwnd). In one embodiment, the determination based on said maintained characterization of the roundtrip delay distinguishes between information loss due to network congestion and information loss due to another cause.

Expressly turning to the figures, FIG. 1 illustrates a network operating according to
20 one embodiment. As shown, network device 101 (e.g., apparatus, router, switch, computer, personal device, etc.) is configured to communicate information, such as, but not limited to, using Transmission Control Protocol, with network device 103 across network 102. Network device 101 is configured to increase its rate of transmission of information despite sent information be lost, in response to attributing the loss of information as not
25 being congestion related. Typically, prior TCP implementation will decrease its rate of transmission of information for any cause, or at least not increasing its rate of transmission in response to a packet loss. Note, in one embodiment, network device 103 is also configured to increase its rate of transmission of information despite sent information be lost, in response to attributing the loss of information as not being congestion related.

FIG. 2 is block diagram of an apparatus or component 200 used in one embodiment
30 associated with the rate of transmitting information to a remote device in response to attributing a detected loss of previous information sent to the remote device as not being a

result of network congestion. In one embodiment, system or component 200 performs one or more processes corresponding to one of the flow diagrams illustrated or otherwise described herein.

In one embodiment, apparatus or component 200 includes one or more processing
5 elements 201, memory 202, storage device(s) 203, specialized component(s) 205
(e.g., optimized hardware such as for performing operations, etc.), and interface(s) 207 for
communicating information (e.g., sending and receiving packets, user-interfaces,
displaying information, etc.), which are typically communicatively coupled via one or
more communications mechanisms 209, with the communications paths typically tailored
10 to meet the needs of the application. In one embodiment apparatus or component 200
corresponds to, or is part of, network device 101 of FIG. 1.

Various embodiments of apparatus or component 200 may include more or less
elements. The operation of apparatus or component 200 is typically controlled by
processing element(s) 201 using memory 202 and storage device(s) 203 to perform one or
15 more tasks or processes. Memory 202 is one type of computer-readable/computer-storage
medium, and typically comprises random access memory (RAM), read only memory
(ROM), flash memory, integrated circuits, and/or other memory components. Memory 202
typically stores computer-executable instructions to be executed by processing element(s)
201 and/or data which is manipulated by processing element(s) 201 for implementing
20 functionality in accordance with an embodiment. Storage device(s) 203 are another type of
computer-readable medium, and typically comprise solid state storage media, disk drives,
diskettes, networked services, tape drives, and other storage devices. Storage device(s) 203
typically store computer-executable instructions to be executed by processing element(s)
201 and/or data which is manipulated by processing element(s) 201 for implementing
25 functionality in accordance with an embodiment.

FIG. 3 illustrates a process performed in one embodiment. Processing begins with
process block 300. In process block 302, the apparatus starts communicating with a remote
device, such as, but not limited, using TCP. As determined in process block 303, in
response to identifying a loss of sent information (e.g., based on received
30 acknowledgements), then processing proceeds to process block 311 (described hereinafter).
Otherwise, as determined in process block 305, in light of no detected loss of information,
if the rate of sending information should be increased, then in process block 306, it is

increased; otherwise, in process block 308, the current rate of transmitting information is maintained. Continuing with process block 311, as determined therein such as, but not limited to, based on estimates of queuing delays in the network, and/or measured roundtrip delays between sending of information and receiving corresponding acknowledgements: if the loss of information is not attributed to congestion in the network, then in process block 314, the rate of transmission of the information is increased; otherwise, in process block 312, the rate of transmission of the information is typically decreased or sometimes maintained. Processing of the flow diagram of FIG. 3 returns to process block 303.

FIG. 4 illustrates a process performed in one embodiment. Processing begins with process block 400. In process block 402, the apparatus communicates with a remote device, such as, but not limited to using Transmission Control Protocol, (TCP), including measuring roundtrip delays between information sent and corresponding received acknowledgements. In process block 404, in response to a detected loss of previously transmitted information by the apparatus: increasing the rate of transmission of information by the apparatus (e.g., increasing a congestion window) in response to attributing the detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and corresponding acknowledgements received by the apparatus (e.g., directly, determining one or more queuing delays determined therefrom such based on minimum, average and/or median, smooth-filtered maximum of roundtrip delays). Processing of the flow diagram of FIG. 4 is complete as illustrated by process block 406.

FIG. 5 illustrates a process performed in one embodiment. Processing begins with process block 500, and in process block 502, the apparatus receives an acknowledgement for information previously sent. As determined in process block 505, in response to identifying a loss of transmitted data (e.g., based on some number of received acknowledgments allowing the identification of lost data), the apparatus enters a loss recovery phase of TCP communications (510), which includes process blocks 512-518.

In process block 512, the apparatus determines whether or not to attribute the packet loss to congestion based on measured roundtrip delays of information sent and corresponding acknowledgement received. One embodiment uses these roundtrip delays directly, possibly in combination with other information. In one embodiment, these

roundtrip delays are used to estimate queuing delays, which are used to identify whether or not to attribute the packet loss to congestion in the network. In one embodiment, a queuing delay is determined based on an average of measured roundtrip delays, a minimum of measured roundtrip delays, and/or a median, smooth-filtered maximum of measured
5 roundtrip delays. As determined in process block 515, if the loss of information was attributed to congestion in the network, then in process block 518, the apparatus implements fast recovery, which typically includes decreasing the rate of transmission of information (e.g., decreasing a congestion window (cwnd), or setting cwnd to the TCP slow-start threshold value (ssthresh)). Otherwise, in process block 516, in response to
10 attributing the loss of transmitted information as not being caused by congestion in the network, the rate of transmission of information is increased (e.g., increasing a congestion window (cwnd)).

Continuing, as determined in process block 520, if a congestion window is greater than a threshold value, then in process block 522, the rate of transmission is increased
15 (e.g., a congestion window (cwnd) is increased). Otherwise, in process block 524, the apparatus implements TCP slow start. Processing of the flow diagram of FIG. 5 is complete as indicated by process block 529.

FIGs. 6A and 6B illustrate a process performed in one embodiment. As shown, this process is described using pseudo code using BIC-TCP (e.g., typically implemented in
20 Linux). Note, known functions, such as `bictcp_recalc_ssthresh()` and `bictcp_update()`, and `bictcp_reset()`, and/or known variables or other items in known implementation of TCP may be referenced.

Code 601 illustrates Initialization, wherein variables are initialized and `blast_reset()`
613 is performed. This code includes initializing a low threshold variable γ_L and high
25 threshold variable γ_H for providing a band for identifying whether or not to attribute a packet loss to network congestion (as described hereinafter).

Code 602 illustrates `On_each_ACK`, which is performed in response to each received TCP acknowledgement to maintain some statistics about measured roundtrip and/or queuing delays. Variable `dMin` is used to maintain a lowest measured roundtrip
30 time; `sRTT` is set a median, smooth-filtered maximum of the measured roundtrip delays; `dMax` maintains the maximum median, smooth-filtered maximum of the measured

roundtrip delays; and sumRTT and cntRTT are maintained for calculating an average of the measured roundtrip delays.

Code 603 illustrates `At_the_start_of_each_RTT_round`, which resets variables maintained for each round of roundtrip times maintained.

5 Code 604 illustrates `Packet_loss`, which is performed in response to a packet loss to recalculate the `ssthresh` variable (e.g., perform TCP slow start, such as in process block 524 of FIG. 5).

Code 605 illustrates `Fast_Recovery_for_each_ACK`, which is performed in response to an identified packet loss (e.g., process blocks 512-518 of FIG. 5). Note, 10 `loss_predictor()` 608 (FIG. 6B) determines whether or not to attribute the packet loss to congestion, and in response to attributing the packet loss as not due to congestion, the congestion window (`cwnd`) continues to be evolved (e.g., increased to increase the rate of transmission of information).

Code 606 illustrates `Timeout`, which is performed in response to a TCP timeout in 15 order to reset the TCP code accordingly.

Code 607 illustrates `recalc_ssthresh()`, which is performed to recalculate the slow start threshold value, with this calculation being different if the packet loss is attributed to network congestion or not to network congestion.

Code 608 illustrates `loss_predictor()`, which determines whether to attribute a 20 packet loss to network congestion or not to network congestion. As shown, if the average queuing delay (e.g., average roundtrip delay minus minimum observed roundtrip delay) is lower than a predetermined low threshold value amount of the longest queuing delay (e.g., longest roundtrip delay minus minimum roundtrip delay), then attribute the packet loss as not being related to network congestion. Also, if the average queuing delay is greater than 25 a predetermined high threshold value amount of the longest queuing delay, then attribute the packet loss as being related to network congestion. Further, inside this band defined by the low and high threshold variables times the longest queuing delay, `loss_predictor()` determines to attribute the packet loss to congestion if the average queuing delay is currently increasing; else it attributes the packet loss as not being due to congestion.

30 Code 609 illustrates `median_filter(RTT)` used to median filter a measured roundtrip time based on multiple history values. One embodiment uses more and/or different history values, which may change the accuracy (possibly increasing) of `loss_predictor()`.

Code 610 illustrates `avg_delay()` used to determine the average queuing delay (e.g., the average roundtrip delay minus the minimum measured roundtrip delay).

Code 611 illustrates `low_threshold()` used to determine the predetermined low threshold value amount of the longest queuing delay.

5 Code 612 illustrates `high_threshold()` used to determine the predetermined high threshold value amount of the longest queuing delay.

Code 613 illustrates `blast_reset()` used to reset the BIC TCP process and initialize some variables.

10 In view of the many possible embodiments to which the principles of our invention may be applied, it will be appreciated that the embodiments and aspects thereof described herein with respect to the drawings/figures are only illustrative and should not be taken as limiting the scope of the invention. For example, and as would be apparent to one skilled in the art, many of the process block operations can be re-ordered to be performed before, after, or substantially concurrent with other operations. Also, many different forms of data
15 structures could be used in various embodiments. The invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

CLAIMS

What is claimed is:

1. An apparatus, comprising:

one or more processing elements and memory configured to communicate

5 information with a remote device over a network;

wherein said communication with the remote device includes, in response to a detected loss of previously transmitted information by the apparatus: increasing the rate of transmission of information by the apparatus in response to attributing the detected loss of previously transmitted information as not being caused by congestion.

10 2. The apparatus of claim 1, wherein said communication of information is performed using Transmission Control Protocol (TCP), and wherein said attributing of the detected loss of previously transmitted information as not being caused by congestion includes measuring a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus
15 corresponding to said sent information.

3. The apparatus of claim 2, wherein said communication using TCP with the remote device includes a loss recovery phase of said TCP communications triggered by the detected loss of previously transmitted information, with the loss recovery phase including performing active congestion control, with said active congestion control including said
20 increasing the rate of transmission of information by the apparatus in response to attributing said detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information; and wherein the loss recovery phase
25 includes reducing the rate of transmission of information by the apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion.

4. The apparatus of claim 3, wherein communication using TCP with the remote device includes using a congestion window (cwnd) for reacting to network conditions; and wherein said increasing the rate of transmission of information includes increasing the congestion window (cwnd); and wherein said reducing the rate of transmission includes
5 decreasing the congestion window (cwnd).

5. The apparatus of claim 3, wherein said attributing the detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information
10 includes determining a plurality of queuing delays based on the plurality of measured roundtrip delays and said attributing said detected loss of previously transmitted information as not being caused by congestion based on the plurality of queuing delays.

6. The apparatus of claim 2, wherein said attributing the detected loss of previously transmitted information as not being caused by congestion based on a plurality of measured
15 roundtrip delays between information sent from the apparatus to the remote device and an acknowledgement received by the apparatus corresponding to said sent information includes determining a plurality of queuing delays based on the plurality of measured roundtrip delays and said attributing said detected loss of previously transmitted information as not being caused by congestion based on the plurality of queuing delays.

20 7. The apparatus of claim 6, wherein a current queuing delay is determined based on a minimum of a plurality of said measured roundtrip delays and one or more of said measured roundtrip delays.

8. The apparatus of claim 7, wherein the current queuing delay is determined based on a difference of said one or more of said measured roundtrip delays with said minimum
25 of a plurality of said measured roundtrip delays.

9. The apparatus of claim 6, wherein the current queuing delay is determined based on an average of a plurality of said measured roundtrip delays, a minimum of a plurality of said measured roundtrip delays, and a median, smooth-filtered maximum of a plurality of said measured roundtrip delays.

10. The apparatus of claim 1, wherein said attributing of the detected loss of previously transmitted information as not being caused by congestion includes estimating that queuing delays within the network between the apparatus and the remote device have not increased therefore not causing the detected loss.

5 11. The apparatus of claim 10, wherein said communication of information is performed using Transmission Control Protocol (TCP).

12. The apparatus of claim 10, wherein a current queuing delay of said queuing delays is determined based on a difference of one or more measured roundtrip delays with a minimum of a plurality of measured roundtrip delays.

10 13. The apparatus of claim 12, wherein a current queuing delay of said queuing delays is determined based on an average of a plurality of a plurality of measured roundtrip delays, a minimum of a plurality of the plurality of measured roundtrip delays, and a median, smooth-filtered maximum of a plurality of the plurality of measured roundtrip delays.

15 14. A method, comprising:
communicating, by a particular apparatus, with a remote device over a network;
wherein said communication with the remote device includes, in response to a detected loss of previously transmitted information by the particular apparatus: increasing the rate of transmission of information by the particular apparatus in response to attributing
20 said detected loss of previously transmitted information as not being caused by congestion based on a current queuing delay determined based on measured roundtrip delays between information sent from the particular apparatus to the remote device and an acknowledgement received by the particular apparatus corresponding to said sent information.

25 15. The method of claim 14, including: in response to the detected loss of previously transmitted information by the particular apparatus: decreasing the rate of transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion based on the current queuing delay.

16. The method of claim 15, wherein said communicating with the remote device over the network includes using Transmission Control Protocol (TCP).

17. The method of claim 16, wherein communicating using TCP with the remote device includes using a congestion window (cwnd) for reacting to network conditions; and
5 wherein said increasing the rate of transmission of information includes increasing the congestion window (cwnd); and wherein said decreasing the rate of transmission of information includes decreasing the congestion window (cwnd).

18. The method of claim 14, including determining the current queuing delay based on a minimum of a plurality of said measured roundtrip delays and one or more of said
10 measured roundtrip delays.

19. The method of claim 18, including determining the current queuing delay based on a difference of said one or more of said measured roundtrip delays with said minimum of a plurality of said measured roundtrip delays.

20. The method of claim 14, including determining the current queuing delay based
15 on an average of a plurality of said measured roundtrip delays, a minimum of a plurality of said measured roundtrip delays, and a median, smooth-filtered maximum of a plurality of said measured roundtrip delays.

21. The method of claim 20, including: in response to the detected loss of previously transmitted information by the particular apparatus: decreasing the rate of
20 transmission of information by the particular apparatus in response to attributing said detected loss of previously transmitted information as being caused by congestion based on the current queuing delay.

22. The method of claim 20, wherein said attributing said detected loss of previously transmitted information as not being caused by congestion based on the current
25 queuing delay includes determining that an average queuing delay is not greater than a previously determined average queuing delay not including the current queuing delay.

23. The method of claim 22, including: in response to the detected loss of
previously transmitted information by the particular apparatus: decreasing the rate of
transmission of information by the particular apparatus in response to attributing said
detected loss of previously transmitted information as being caused by congestion based on
5 the current queuing delay.

24. An apparatus, comprising:
one or more processing elements and memory configured to communicate using
Transmission Control Protocol (TCP) with a remote device over a network;
wherein said communication using TCP with the remote device includes using a
10 congestion window (cwnd) for reacting to network conditions;

wherein said communication with the remote device using TCP includes
characterizing the roundtrip delay between the apparatus and the remote device based on
measured roundtrip delay for a plurality of sent information and corresponding received
acknowledgements; and

15 wherein said communication with the remote device using TCP includes: in
response to an identified loss of particular information sent to the remote device and a
determination based on said maintained characterization of the roundtrip delay, increasing
the congestion window (cwnd).

25. The apparatus of claim 24, wherein the determination based on said maintained
20 characterization of the roundtrip delay distinguishes between information loss due to
network congestion and information loss due to another cause.




FIGURE 1

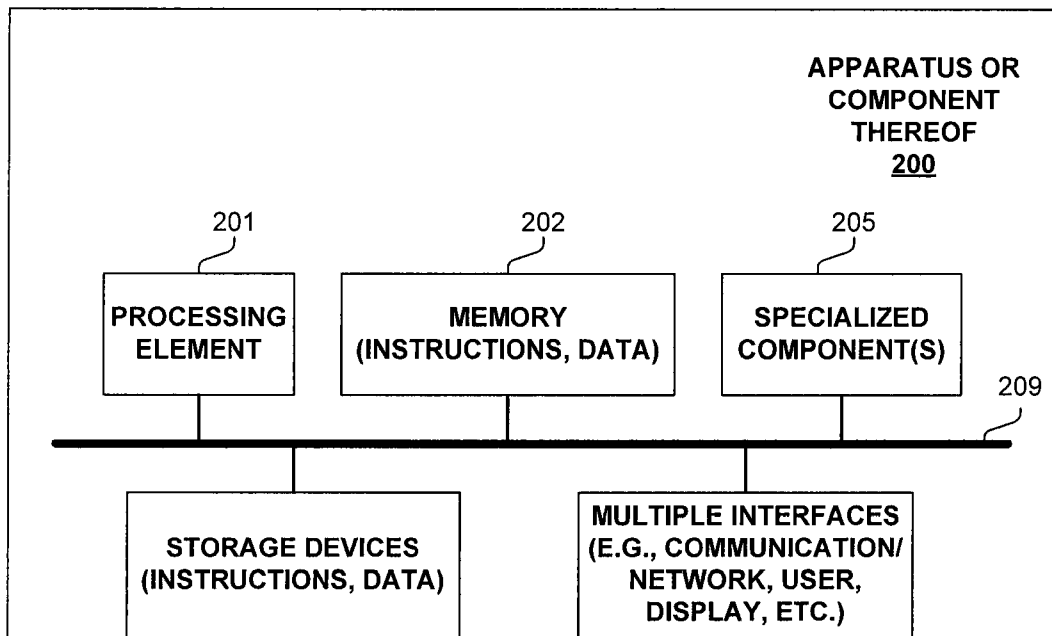


FIGURE 2

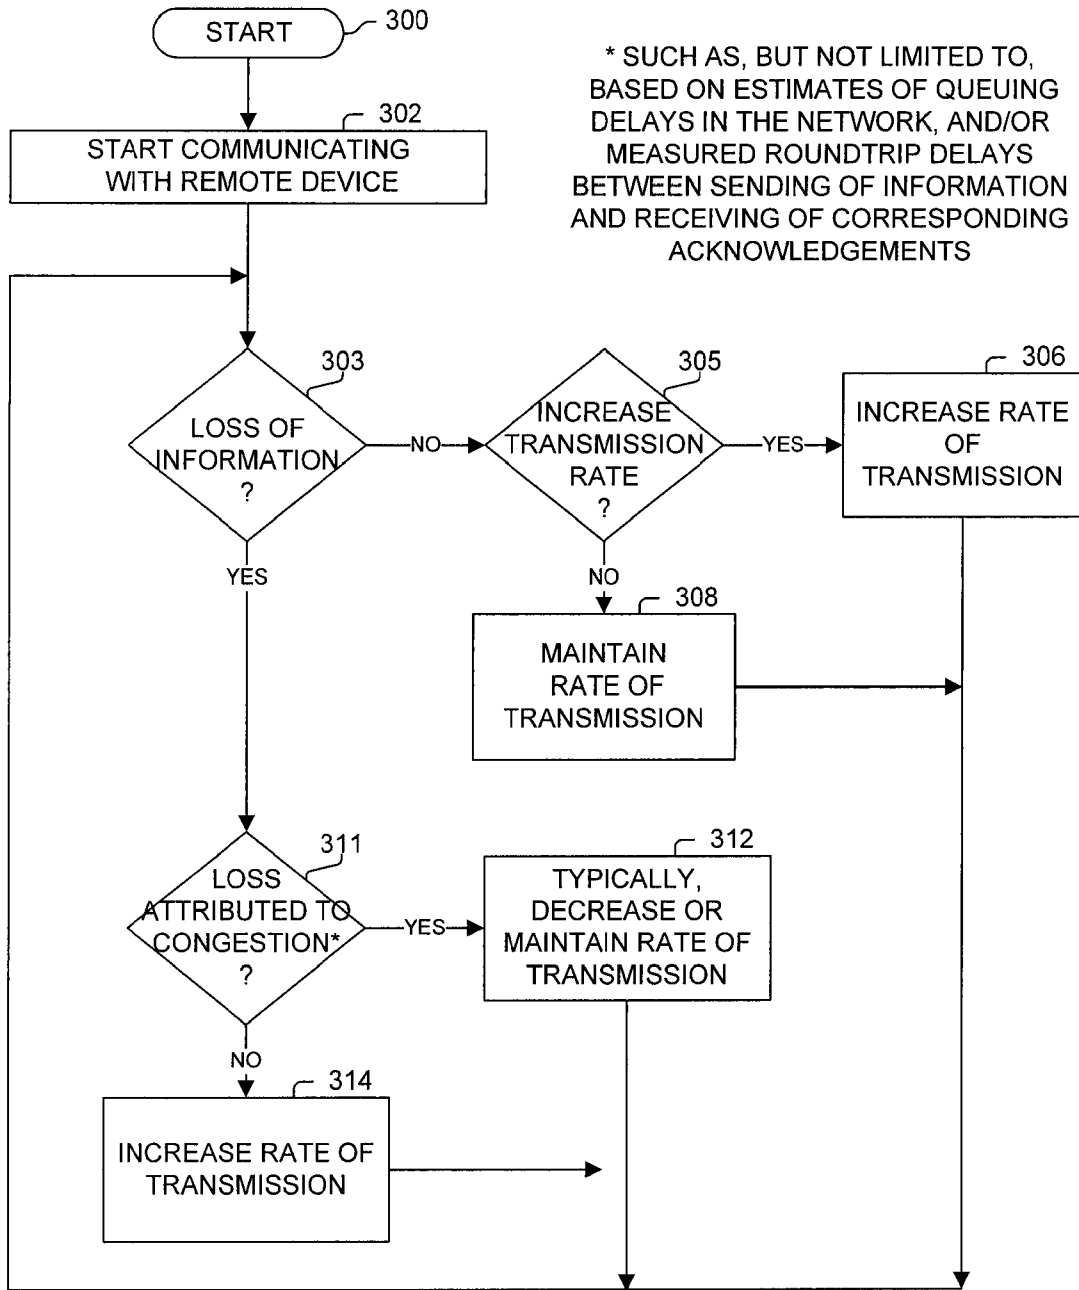


FIGURE 3

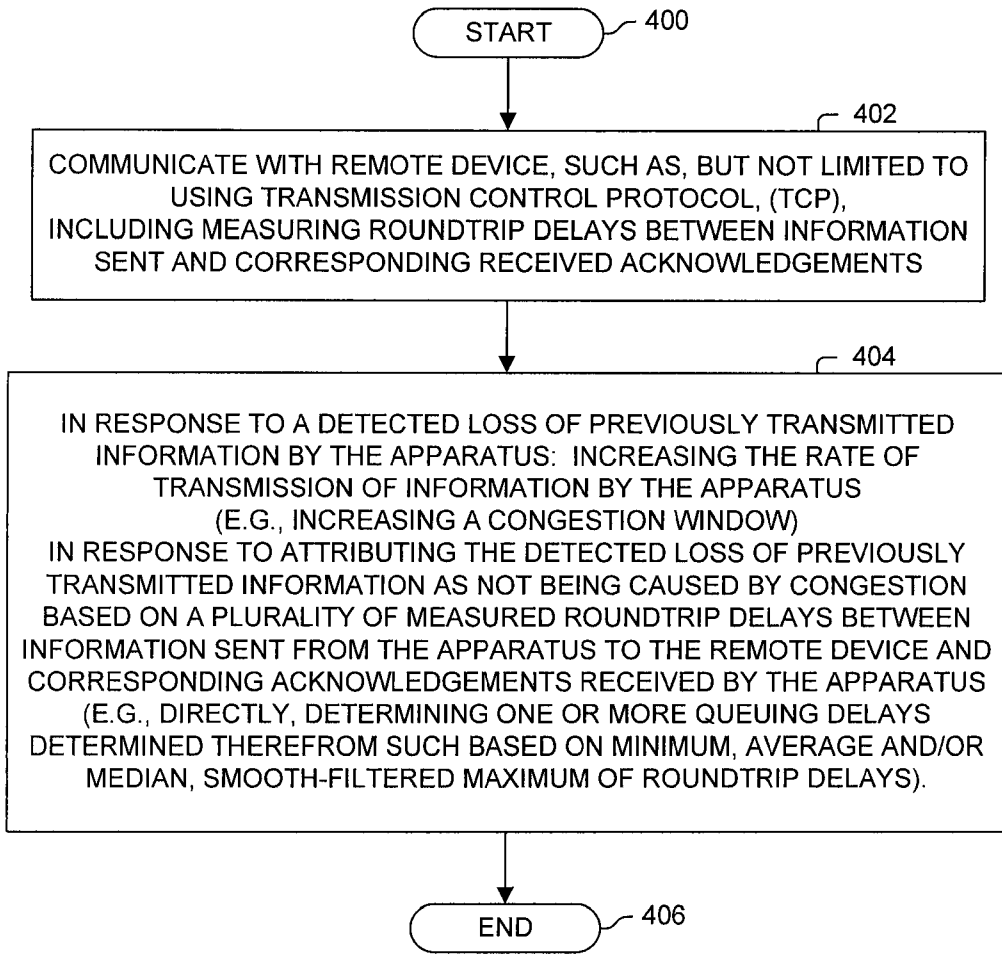


FIGURE 4

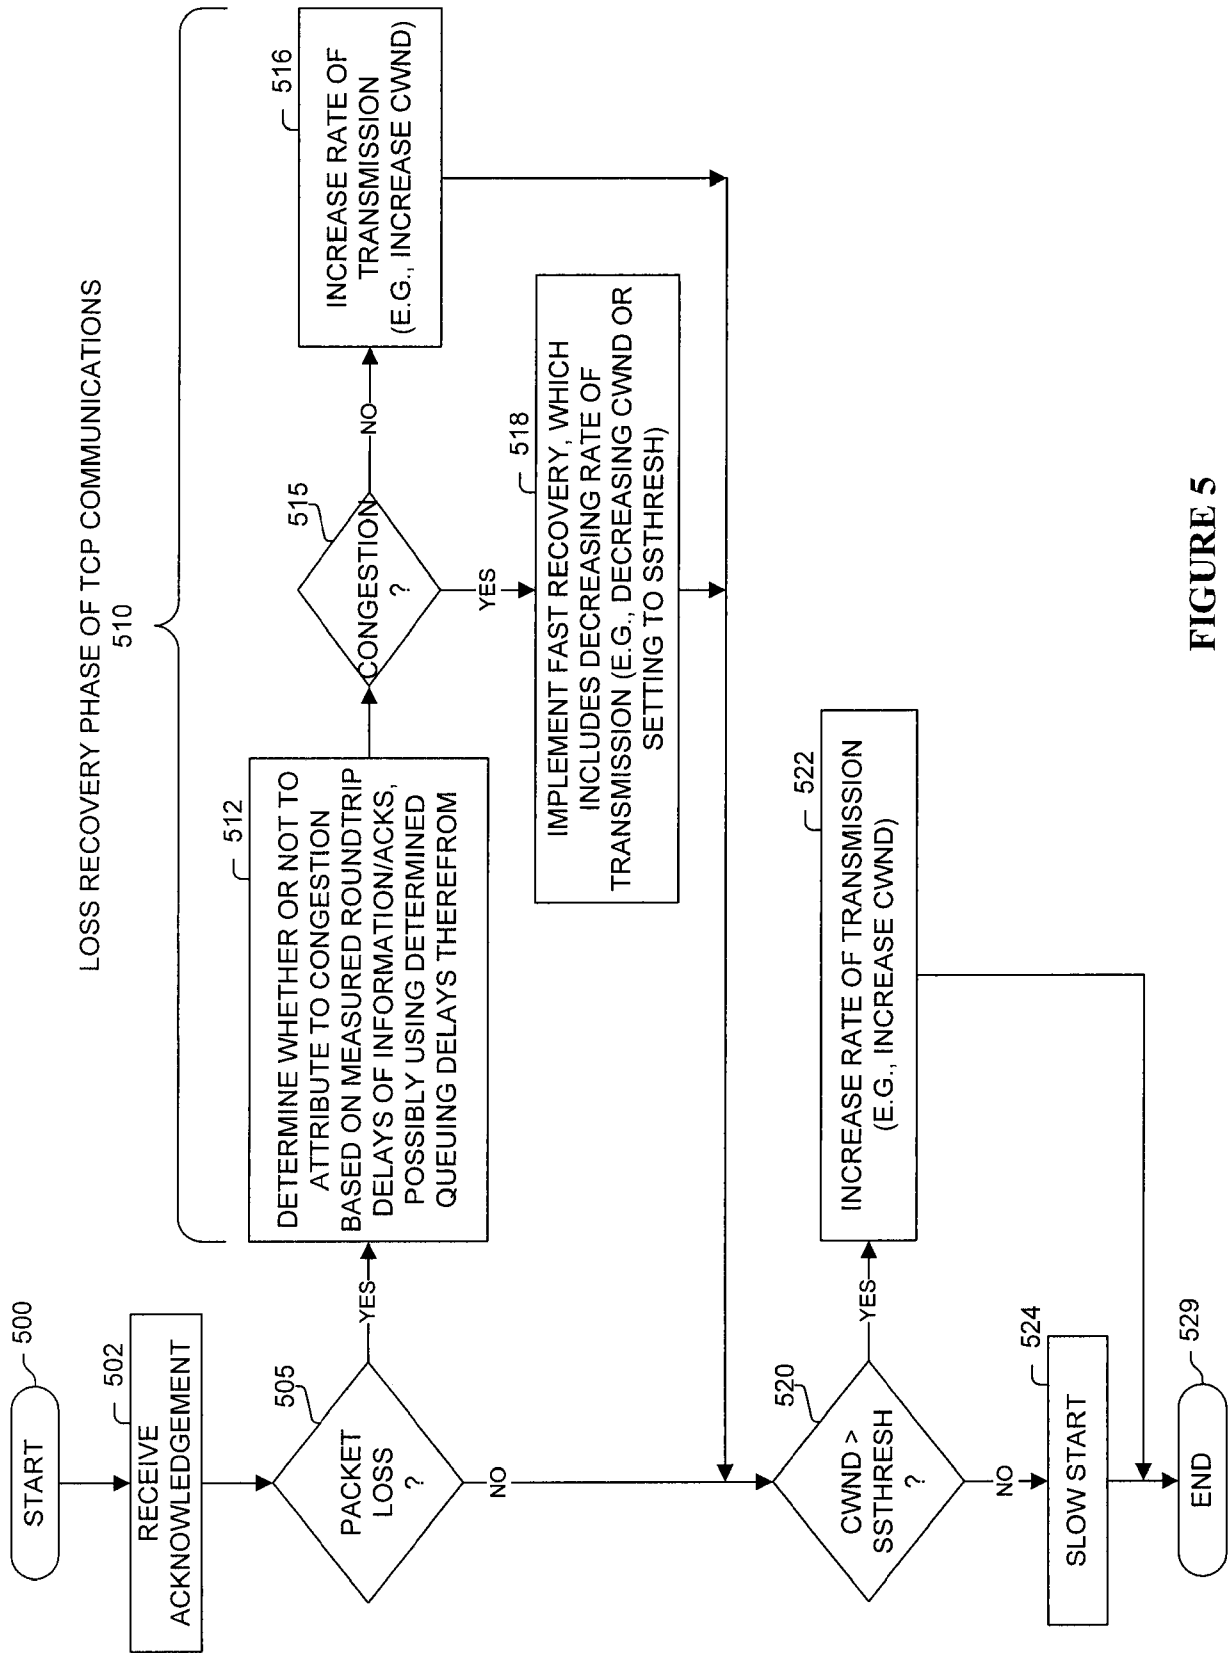


FIGURE 5

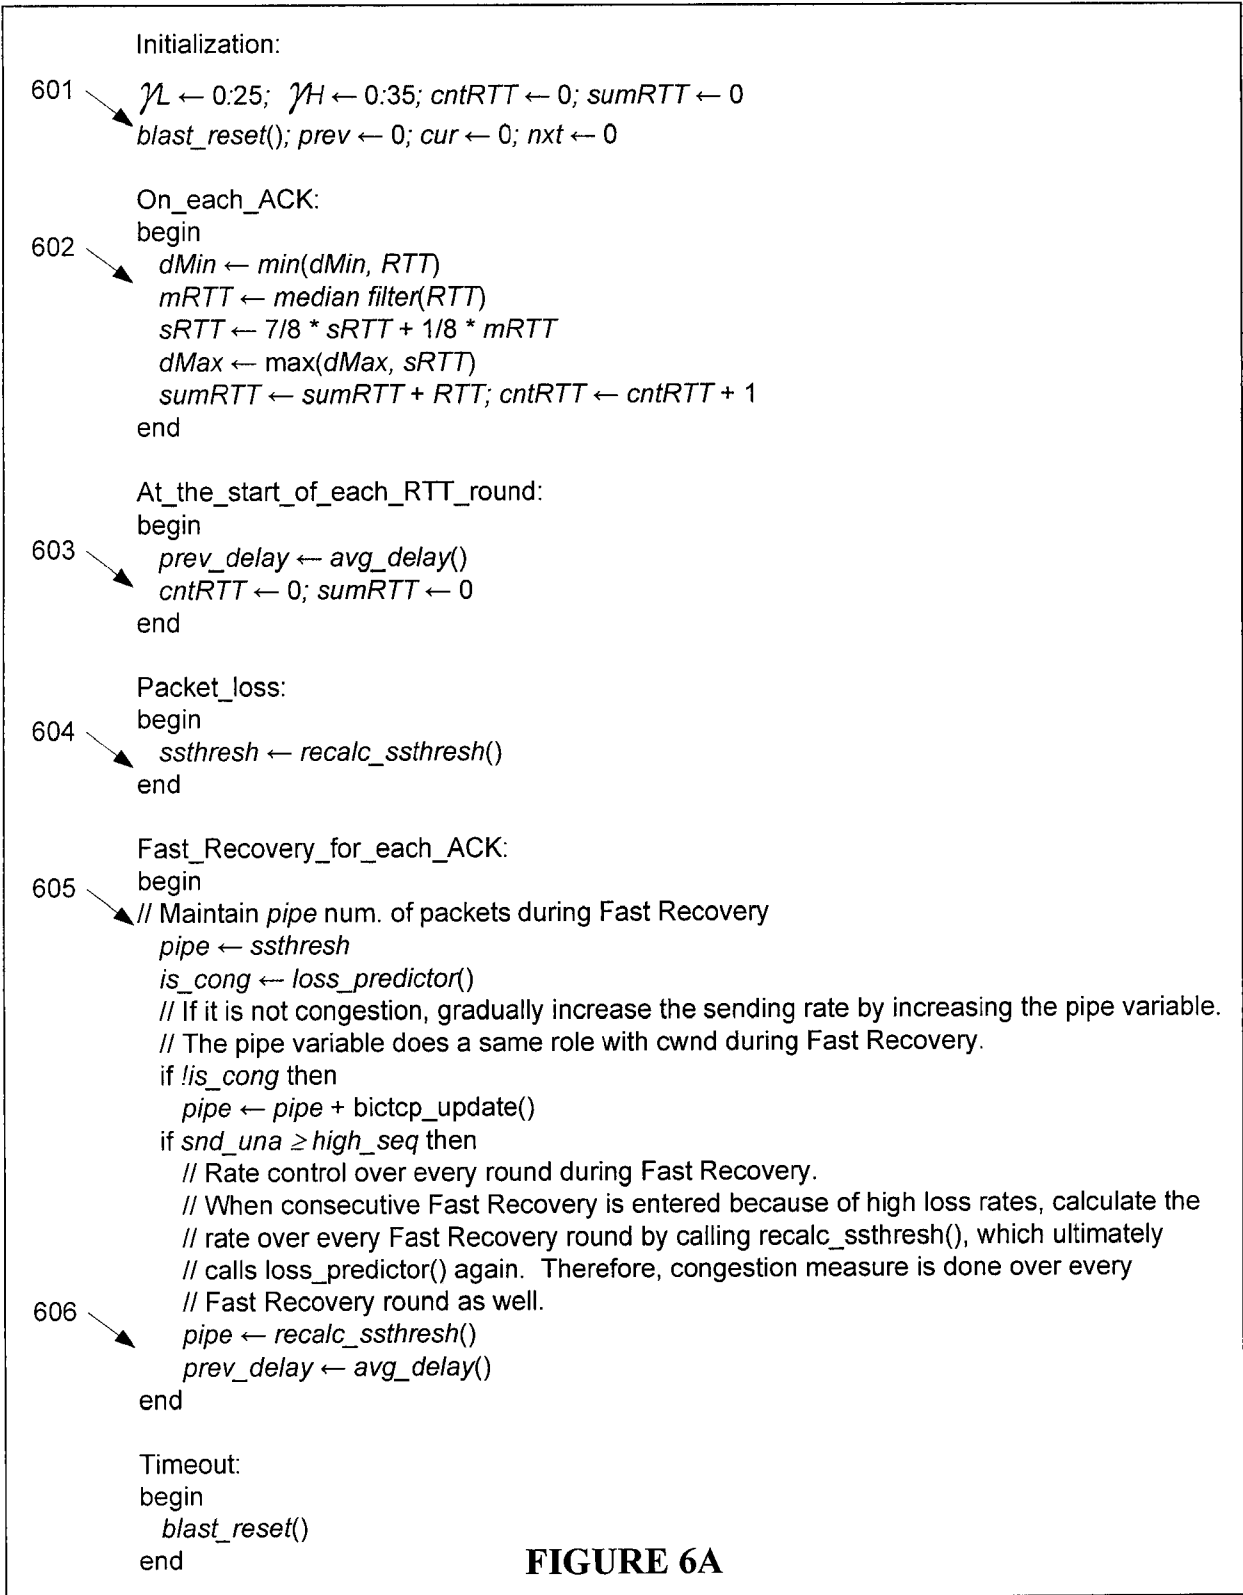


FIGURE 6A

```

607 → recalc_ssthresh():
begin
  is_cong ← loss_predictor()
  if is_cong then ssthresh ← bictcp_recalc_ssthresh()
  else ssthresh ← cwnd
  return ssthresh
end

608 → loss_predictor():
begin
  // Identify the congestion-related losses using D
  if avg_delay() ≤ low_threshold() then is_cong ← 0
  else if avg_delay() > high_threshold() then is_cong ← 1
  else
  // Is the delay increasing in the grey zone ?
  if avg_delay() > prev_delay then is_cong ← 1
  else is_cong ← 0
  return is_cong
end

609 → median_filter(RTT):
begin
  prev ← cur; cur ← nxt; nxt ← RTT
  cur ← max(min(prev, nxt), cur)
  cur ← min(max(prev, nxt), cur)
  return cur
end

610 → avg_delay():
begin
  if cntRTT > 0 and sumRTT > 0 then
  return max( sumRTT / cntRTT - dMin, 0)
  else return 0
end

611 → low_threshold():
begin
  return  $\gamma_L * (dMax - dMin)$ 
end

612 → high_threshold():
begin
  return  $\gamma_H * (dMax - dMin)$ 
end

613 → blast_reset():
begin
  bictcp_reset(); prev_delay ← 0; dMin ←  $\infty$ ; dMax ← 0
end

```

FIGURE 6B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2010/060517

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G01R 31/08 (2011.01)

USPC - 370/232

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G01R 31/08; H04L 12/26; G06F 15/173 (2011.01)

USPC - 370/232, 230,235,374; 709/224, 233

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

MicroPatent, Google Patents

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2005/0201284 A1 (CHERITON) 15 September 2005 (15.09.2005) entire document	1, 10-11, and 14-16
--		-----
Y		2-9, 12-13, and 17-25
Y	US 2006/0034286 A1 (KONING et al) 16 February 2006 (16.02.2006) entire document	2-9, 12, 13, and 17-25
A	US 5,974,028 A (RAMAKRISHNAN) 26 October 1999 (26.10.1999) entire document	1-25
A	US 2004/0218617 A1 (SAGFORS) 04 November 2004 (04.11.2004) entire document	1-25

Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

08 February 2011

Date of mailing of the international search report

17 FEB 2011

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774