(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
29 January 2015 (29.01.2015)

WIPO | PCT

(10) International Publication Number
WO 2015/012936 A1

(51) International Patent Classification:
H04L 29/06 (2006.01) *H04L 12/801* (2013.01)
H04W 80/06 (2009.01)

(21) International Application Number:
PCT/US2014/036909

(22) International Filing Date:
6 May 2014 (06.05.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/948,279 23 July 2013 (23.07.2013) US

(71) Applicant: **CISCO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, CA 95134-1706 (US).

(72) Inventor: **LAI, Cheng-jia**; 587 Sunnymount Ave., Sunnyvale, CA 94087 (US).

(74) Agents: **FLOAM, D., Andrew** et al.; Edell, Shapiro & Finnan, LLC, 9801 Washingtonian Blvd., Suite 750, Gaithersburg, MD 20878 (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, BN, 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, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, 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, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: NETWORK CONGESTION CONTROL WITH AWARENESS OF RANDOM PACKET LOSSES

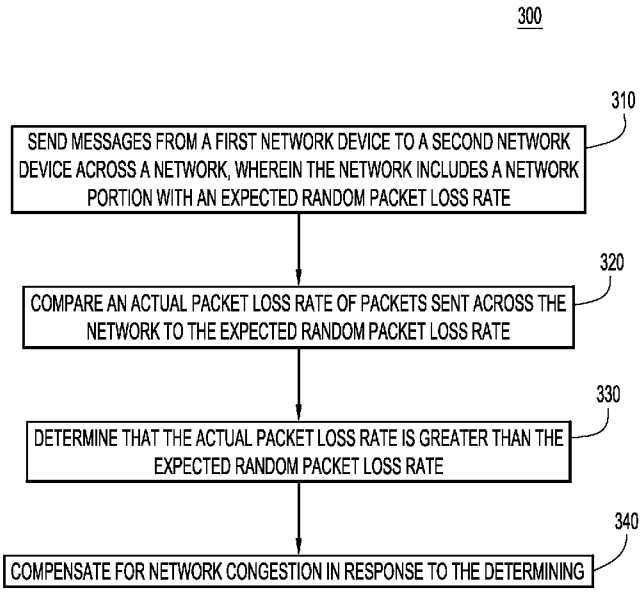


FIG.3

(57) Abstract: Messages are sent from a first network device to a second network device across a network. The network includes a network portion with an expected random packet loss rate. The actual packet loss rate for packets sent across the network is compared to the expected random packet loss rate. A determination is made that the actual packet loss rate is greater than the expected random packet loss rate. Compensation for network congestion is performed in response to the determination that the actual packet loss rate exceeds the expected random packet loss rate.

WO 2015/012936 A1

Published:

— *with international search report (Art. 21(3))*

NETWORK CONGESTION CONTROL WITH AWARENESS OF RANDOM PACKET LOSSES

TECHNICAL FIELD

[0001] The present disclosure relates to computer networks, and in particular, congestion compensation and avoidance in computer networks.

BACKGROUND

[0002] Congestion control mechanisms in Internet transport layers, such as Transmission Control Protocol (TCP) congestion control mechanisms and Stream Control Transmission Protocol (SCTP) congestion control mechanisms, have been designed to assume that packet losses are an indication of network congestion in the path. Additional mechanisms have been developed to provide congestion indications, e.g. explicit congestion notification (ECN), but these mechanisms are not consistently supported. Accordingly, transport layer congestion control mechanism implementations must still treat any packet loss as an indication of congestion.

[0003] Congestion control mechanisms adjust for congestion by resizing the congestion window (cwnd) size for the transmission. Even for layer-4 TCP and SCTP congestion signaling, such as ECN Congestion Experienced (CE) signaling, congestion control mechanisms react to the signaling in the same manner in which they react to actual lost packets.

[0004] As network technology evolves to include wireless networks, networks have begun to see random packet losses. For example, third generation (3G) wireless networks can see 1% packet loss rates even under perfect environmental conditions. Similarly, Long-Term Evolution (LTE) wireless users may see 10^{-4} or 10^{-6} percent packet loss rates even under the best of conditions. As wireless devices move around, they can experience significantly higher random packet loss rates due to the physical location of the device, and irrespective of current network congestion.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 is an illustration of an example network environment configured to employ congestion compensation with random packet loss awareness.

[0006] FIG. 2 is an illustration of an example network environment configured to provide random packet loss awareness through out-of-band signaling between a sending device and network links.

[0007] FIG. 3 is a flowchart illustrating a process for providing congestion compensation with random packet loss awareness.

[0008] FIG. 4 is a flowchart illustrating congestion compensation with random packet loss awareness implemented in a network environment which fast retransmits lost packets in response to a predetermined number of duplicate acknowledgement messages.

[0009] FIG. 5 is a flowchart illustrating congestion compensation with random packet loss awareness implemented in a network environment which retransmits lost packets in response to expiration of a retransmission timer.

[0010] FIG. 6 is an illustration of a plurality of network configurations that may implement congestion compensation with random packet loss awareness.

[0011] FIG. 7 is a block diagram of a network enabled device configured to provide congestion compensation with random packet loss awareness.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Overview

[0012] Techniques are presented herein to provide congestion control with awareness of random packet losses. Messages are sent from a first network device to a second network device across a network. The network includes a network portion with an expected random packet loss rate. The actual packet loss rate for packets sent across the network is compared to the expected random packet loss rate. A determination is made that the actual packet loss rate is greater than the expected random packet loss rate. Compensation for network congestion is performed in response to the determination that the actual packet loss rate exceeds the expected random packet loss rate.

Example Embodiments

[0013] Depicted in FIG. 1 is network environment 100 in which network congestion compensation including random packet loss awareness may be been implemented. Included in network environment 100 is a sending device 110, a receiving device 120, and a network 130. Included within network 130 is portion 140 in which there is an expected random packet loss rate. According to the example of FIG. 1, portion 140 is embodied in an Open System Interconnection (OSI) model Layer 2 network portion, while the remainder of network 130 is embodied in OSI Layer 4 network portions. As a sending device, device 110 sends traffic 150 over network 130 to receiving device 120. The receiving device 120 responds to sending device 110 through return traffic 160, e.g., acknowledgments. Depending on the content of return traffic 160, sending device 110 may determine whether or not all of the packets sent through traffic 150 have been received by receiving device 120. Also included in sending device 110 is random packet loss awareness unit 190, which is configure to allow sending device 110 to apply congestion compensation to traffic 150 which takes into account the random packet loss rate of portion 140.

[0014] According to the example of FIG. 1, sending device 110 may send traffic 150 to receiving device 120 in the form of sequentially sent packets 152, 154 and 156. Packets 152, 154 and 156 may include sequence numbers indicative of the order in which they were sent. Accordingly, packet 152 may include a sequence number of "1," packet 154 may include a sequence number of "2" and packet 156 may include a sequence number of "3." These sequence numbers allow the packets 152, 154 and 156 to be reordered if they are received out of sequence. Receiving device 120 will acknowledge the packets it has received through acknowledgements 162 and 166, which will include the sequence number for the next packet it expects to receive. For example, acknowledgment 162 acknowledges receipt of packet 152, and includes the lowest sequence number that it has not yet received, mainly sequence number "2" for packet 154. Acknowledgement 166, which acknowledges receipt of packet 156, will also indicate the sequence number of "2" if packet 154 has not yet been received, as "2" would still be the lowest un-received sequence number.

[0015] If sending device 110 receives a sequence number of "2" in a predetermined number of acknowledgement messages, such as messages 162 and 166, it may assume that packet 154 has been lost, and may proceed to resend packet 154. Sending device 110 may also modify the manner in which it sends packets when it is determined that a packet has been lost. For example, sending device may assume that network congestion is the cause of

the packet loss rate, and sending device 110 may resize a congestion window to attempt to alleviate network congestion to improve the packet loss rate. On the other hand, if the packet loss rate is not due to network congestion, resizing the congestion window will not remedy the lost packet problem, and may further degrade network performance. In order to avoid this problem, sending device 110 is configured to include random packet loss awareness unit 190.

[0016] Random packet loss awareness unit 190 is configured to determine when a packet loss rate exceeds a random packet loss rate caused by a "lossy" portion of network 130. If the packet loss rate is below the random loss rate caused by the lossy portion of the network, random packet loss awareness unit 190 will determine that the packet loss is not due to network congestion, and therefore, packet transmission will not be modified to account for network congestion. On the other hand, if the packet loss rate exceeds the random packet loss rate, random packet loss awareness unit 190 may determine that network congestion is a contributing factor in the packet loss, and may direct sender 110 to take steps necessary to alleviate or compensate for the network congestion. Specifically, random packet loss awareness unit 190 will determine whether the following expression is true or not:

$$\text{Actual Packet Loss Rate} < \text{Expected Random Packet Loss Rate}$$

[0017] If the above-listed expression is true, the random packet loss awareness unit 190 may determine that at least some portion of the actual packet loss rate is due to network congestion. Accordingly, random packet loss awareness unit 190 can take appropriate action to alleviate the congestion. On the other hand, if the above-listed expression is not true, random packet loss awareness unit 190 may determine that the actual packet loss rate is not due to congestion.

[0018] For example, portion 140 of network 130 may be embodied in a wireless network portion, such as a Wi-Fi™ network, a third generation wireless network (e.g. 3G/GPRS/HPSA), or a fourth generation wireless network (e.g. 4G/LTE). These wireless networks are expected to see some packet loss regardless of the amount of traffic currently being handled by network 130. Specifically, wireless networks are subject to signal-to-noise ratio (SNR) optimization to save energy at wireless transceivers, and therefore, the transceivers will generally experience random packet loss. If portion 140 is embodied in a 3G wireless network, users may experience 1% packet loss regardless of network congestion. Similarly, if portion 140 is embodied in an LTE wireless network, users may experience a

random packet loss rate of between 10^{-4} and 10^{-6} percent. Furthermore, mobile devices may experience greater random packet loss rates due to the physical location of the mobile device. Because these loss rates are not due to network congestion, the loss rate will not be improved by altering the rate or number of messages sent over network 130.

[0019] Furthermore, when congestion compensation techniques are applied in response to random packet losses, network performance may be degraded instead of improved. For example, even when the network path has no congestion, if congestion compensation techniques are applied in response to a 1% random packet loss rate, such as those present in 3G networks, the congestion window may be reduced to a value of less than 12.2 packets in equilibrium, and limit throughput over the network to less than 0.73 Mbps. If the random packet loss rate approaches or exceeds 10%, as may happen depending on the geographic location of a wireless device, the congestion window will drop to less than 3.87 packets in equilibrium, and throughput will drop to less than 0.23 Mbps. This drop in throughput due to the congestion compensation may result in significant under-utilization of wireless links. Furthermore, the drop in throughput may result in failures of certain applications, such as video streaming, even though the network should be able to provide sufficient throughput. For example, video streaming applications may experience oscillation of video quality or even complete failure over a network which is more than capable of supporting the video stream.

[0020] Accordingly, random packet loss awareness unit 190 would be aware of the expected random packet loss rate for network portion 140, and utilize this knowledge to control sending device 110 to compensate, or not compensate, for network congestion, as appropriate. Depending on the implementation, random packet loss awareness unit 190 may be made aware of the expected packet loss rate through different means. For example, random packet loss awareness unit 190 may be manually configured with the expected packet loss rate, the sending device 110 may include an application programming interface (API) which locally determines and reports the expected random packet loss rate to the random packet loss awareness unit 190, or the random packet loss awareness unit 190 may have the expected random packet loss rate reported to it through out-of-band signaling messages.

[0021] With reference now made to FIG. 2, depicted therein is an example sender device 110, including random packet loss awareness unit 190, which utilizes a combination of manual configuration and signaling. Specifically, the random packet loss awareness unit 190 may be manually configured with information indicating the type of lossy network portion

140 within network 130. Accordingly, the random packet loss awareness unit 190 will expect at least a minimum packet loss rate, such as the 1% packet loss rate of a 3G network. This manual configuration may be supplemented through an API that receives reports of dynamic packet loss rates from the local links 210a-c in network 130 for network portions 140 and 250-270. According to the example of FIG. 2, network links 210a-c use out-of-band signaling messages 220a-c to provide dynamic packet loss rates for individual portions 140 and 250-270 of network 130. For example, signaling message 220a may provide sender 110 with dynamic packet loss rate information for network portion 140. Similarly, message 220b provides dynamic packet loss rate information for portion 260, while message 220c provides dynamic packet loss rate information for portion 270.

[0022] Upon receipt of signaling messages 220a-c, random packet loss awareness unit 190 may determine the actual packet loss rate, as well as the expected random packet loss rate. By comparing the two values, random packet loss awareness unit 190 is able to appropriately compensate for packet losses over network 130.

[0023] Referring now to FIG. 3, depicted therein is a flowchart 300 depicting a process of providing congestion control with awareness of random packet losses. The process begins in block 310 in which messages are sent from a first network device to a second network device across a network. The network includes a network portion which includes an expected random packet loss rate. For example, the network may include a wireless portion which has an expected random packet loss rate due to SNR optimization.

[0024] In block 320, the actual packet loss rate is compared to the expected random packet loss rate. The expected random packet loss rate may be predetermined based on the type of network link in which the network portion having the expected random packet loss rate is embodied. In other examples, the expected random packet loss rate may be dynamically calculated by the sending device or through signaling messages sent to the sending device from one or more network links. The actual packet loss rate may be determined according to the content of acknowledgment messages received by the sending device, or in response to retransmission timeouts (RTOs), as will be described below with reference to FIGs. 4 and 5.

[0025] In block 330, a determination is made as to whether or not the actual packet loss rate exceeds the expected random packet loss rate. Specifically, it is determined that the

actual packet loss rate is greater than the expected random packet loss rate. In block 340, network congestion is compensated for in response to the determination.

[0026] With reference now made to FIG. 4, depicted therein is a flowchart 400 illustrating a process of providing congestion control with awareness of random packet losses in a Layer 4 network which includes a Layer 2 portion having an expected random packet loss rate. Furthermore, the Layer 4 network is configured to retransmit packets in response to receiving a predetermined number of duplicate acknowledgement messages.

[0027] The process begins in block 410 where a predetermined number of duplicate acknowledgement messages are received, in this case, three duplicate acknowledgement messages. As previously discussed above with reference to FIG. 1, the duplicate acknowledgement messages may serve as an indication that packets are being lost between the sending and receiving devices. According to the example of FIG. 4, if three acknowledgement messages are received which all indicate the same packet as the next expected packet, it may be determined that the indicated packet has been lost.

[0028] In block 420, the packet indicated by the duplicate acknowledgment messages is retransmitted to the receiving device. This retransmission is considered a "fast" retransmission because the sending device does not wait for a retransmission timer to time-out before retransmitting the packet. Instead, the packet is retransmitted as soon as the three duplicate acknowledgment messages are received. Because the three duplicate acknowledgement messages will often be received prior to expiration of the transmission timer, the retransmission in response to the duplicate acknowledgement messages is considered "fast" relative to other retransmission techniques.

[0029] Next, if the sending device was not operating with awareness of expected random packet loss rates in its congestion compensation methods, the sending device may immediately compensate for network congestion. According to the present example, the actual packet loss rate is first calculated in step 430. Specifically, the Layer 4 packet loss rate (L4-PLR) is calculated. According to the example of FIG. 4, Layer 4 packet loss rate is calculated in response to the fast retransmission of block 420. According to other examples, the Layer 4 packet loss rate may be calculated in response to fast retransmitting a predetermined number of packets, or the Layer 4 packet loss rate may be calculated at predetermined intervals.

[0030] In step 440, the Layer 4 packet loss rate is compared to the expected random packet loss rate, which according to the example of FIG. 4 is the Layer 2 packet loss rate (L2-PLR). Specifically, it is determined whether or not the Layer 4 packet loss rate is greater than the Layer 2 packet loss rate. If the Layer 4 packet loss rate is less than or equal to the Layer 2 packet loss rate, it may be determined that the Layer 4 packet loss rate is simply due to the random losses in the Layer 2 portion of the network, and therefore, performing congestion compensation would be ineffective. Accordingly, the process would continue to block 460, where the process ends until duplicate acknowledgment messages indicate that another packet has been lost.

[0031] On the other hand, if the Layer 4 packet loss rate is greater than the Layer 2 packet loss rate, it may be determined that network congestion contributes, at least partially, to the Layer 4 packet loss rate. Accordingly, the process proceeds to block 450 where a congestion window (cwnd) and the slow start threshold (ssthresh) are reduced according to a slow-start congestion compensation method. Specifically, slow-start congestion compensation avoids sending more data than a network is capable of handling by decreasing the number of packets that can be sent without receiving an acknowledgement (decreasing the congestion window) and decreasing the threshold for when more congestion compensation or avoidance should be performed (decreasing the slow start threshold). Upon completion of the congestion compensation or avoidance of block 450, the process ends in block 460, and the sending device continues sending packets until another packet is lost.

[0032] With reference now made to FIG. 5, depicted therein is flowchart 500 illustrating a process similar to that illustrated in flowchart 400 of FIG. 4. The process of flowchart 500 differs from that of FIG. 4 in that the sending device waits for retransmission timeouts (RTOs) before determining a packet has been lost. Accordingly, the process begins in block 510 where a packet RTO timer expires. In response to expiration of the timer, the sending device retransmits the packet. Next, if the sending device was not operating with awareness of expected random packet loss rates in its congestion compensation methods, the sending device may immediately compensate for network congestion. Instead, because the sending device does implement or operate with awareness of expected random packet loss rates, the runtime Layer 4 packet loss rate is calculated in step 530. In step 540, it is determined whether or not the Layer 4 packet loss rate exceeds the expected random packet loss rate, in this case, the Layer 2 packet loss rate. If the Layer 4 packet loss rate is less than or equal to

the expected random Layer 2 packet loss rate, the process ends at step 560 without any congestion compensation or avoidance being performed.

[0033] If, on the other hand, the Layer 4 packet loss rate exceeds the expected Layer 2 packet loss rate, network congestion compensation and/or avoidance is performed in step 550. Specifically, because the Layer 4 packet loss rate exceeds the expected Layer 2 packet loss rate, it may be determined that network congestion contributes, at least partially, to the Layer 4 packet loss rate. In response, the sending device will reduce the congestion window and slow start threshold in order to compensate for the network congestion. Once the congestion window and slow start threshold are reduced, the process completes in block 560 until another lost packet is detected.

[0034] With reference now made to FIG. 6, depicted therein is a plurality of example networks 610-670 over which packets are sent from sending device 110 to receiving device 120. As illustrated, the entirety of network 610 between sending device 110 and receiving device 120 may be comprised of wired or optical networks, without any intervening wireless network portions. Nevertheless, so long as some or all of network 610 has an expected random packet loss rate, the techniques described herein will apply. For networks 620 and 630, the wireless or lossy portion of the network does not need to be located at an intermediate network link. Instead, the wireless or lossy portion of the network may be local to either of sending device 110 or receiving device 120. Finally, as illustrated, networks 640-670 are not limited to a single wireless or lossy portion. Instead, each of networks 640-670 is constructed from two or more wireless or lossy portions.

[0035] Turning to FIG. 7, depicted therein is a network enabled device configured to implement network congestion compensation including random packet loss awareness, such as sending device 110. The sending device 110 comprises one or more network interface units 710 (e.g., network interface cards) to enable communication over a network, processor(s) 720, bus 730, and memory 740. The memory 740 contains software instructions for operating system 745 and random packet loss awareness unit 790.

[0036] Memory 740 may comprise ROM, RAM, magnetic disk storage media devices, optical storage media devices, flash memory devices, electrical, optical, or other physical/tangible (e.g., non-transitory) memory storage devices. The processor 720 is, for example, a microprocessor or microcontroller that executes instructions, including instructions contained random packet loss awareness unit 790. Thus, in general, the memory

740 may comprise one or more tangible (non-transitory) computer readable storage media (e.g., a memory device) encoded with software comprising computer executable instructions and when the software is executed (by the processor 720), and in particular random packet loss awareness unit 790, it is operable to perform the operations described herein in connection with FIGs. 1-5. Specifically, random packet loss awareness unit 790 includes instructions that allow processor 720 to implement congestion compensation including expected random packet loss awareness.

[0037] By implementing the techniques and devices described above, networks may achieve improved Transmission Control Protocol/Stream Control Transmission Protocol (TCP/SCTP) throughput and network utilization for better user experiences, particularly when the Layer 4 end-to-end path includes a wireless link or other portion which experiences random packet losses, particularly in Layer 2 of the OSI model. Furthermore the techniques and devices described herein may be applied as an enhancement to any existing TCP and SCTP congestion control mechanism (or other Layer 4 protocols). Specifically, these techniques may be particularly advantageous in Layer 4 networks supporting Selective Acknowledge (SACK), (e.g. TCP-SACK and SCTP) which allows for more accurate runtime estimation of the Layer 4 packet loss rate. Furthermore, the when signaling is implemented to dynamically determine Layer 2 random packet loss, the Layer 2 packet loss rate can be advantageously collected instead of relying solely on acknowledgement messages generated at the last hop of the transmission path.

[0038] The above description is intended by way of example only.

What is claimed is:

1. A method comprising:

sending messages from a first network device to a second network device across a network, wherein the network includes a network portion with an expected random packet loss rate;

comparing an actual packet loss rate of packets sent across the network to the expected random packet loss rate;

determining that the actual packet loss rate is greater than the expected random packet loss rate; and

compensating for network congestion in response to the determining.

2. The method of claim 1, wherein comparing the actual packet loss rate to the expected random packet loss rate comprises comparing an actual Layer 4 packet loss rate to an expected Layer 2 packet loss rate.

3. The method of claim 1, wherein compensating for network congestion comprises reducing a congestion window of the first network device.

4. The method of claim 1, wherein compensating for network congestion comprises reducing a slow-start threshold of the first network device.

5. The method of claim 1, further comprising fast retransmitting a lost packet from the first network device to the second network device.

6. The method of claim 1, further comprising receiving at the first device a message reporting the expected random packet loss rate of the network portion.

7. The method of claim 6, wherein receiving the message comprising receiving a locally generated message.

8. The method of claim 6, wherein receiving the message comprises receiving an out-of-band signaling message.

9. The method of claim 8, wherein receiving the out-of-band signaling message comprises receiving a plurality of messages from devices associated with a corresponding one of a plurality of network links, wherein each of the plurality of messages indicates an expected random packet loss rate for a respective network link.

10. The method of claim 1, further comprising manually configuring the first network device with the expected random packet loss rate.

11. The method of claim 1, further comprising dynamically calculating the actual packet loss rate at the first network device.

12. The method of claim 1, further comprising receiving a plurality of duplicate acknowledgement messages;

wherein comparing the actual packet loss rate to the expected random packet loss rate comprises comparing in response to receiving the plurality of duplicate acknowledgment messages.

13. The method of claim 1, wherein comparing the actual packet loss rate to the expected random packet loss rate comprises comparing the actual packet loss rate to the expected packet loss rate in response to expiration of a retransmission time-out timer.

14. The method of claim 1, wherein sending the messages from the first network device to the second network device comprises sending the messages across a wireless portion of the network.

15. An apparatus comprising:

a network interface unit configured to enable network communications;

a memory; and

a processor coupled to the network interface unit and the memory, wherein the processor is configured to:

send messages to a receiving network device across a network, wherein the network includes a network portion with an expected random packet loss rate;

compare an actual packet loss rate of packets sent across the network to the expected random packet loss rate;

determine that the actual packet loss rate is greater than the expected random packet loss rate; and

compensate for network congestion in response to determining that the actual packet loss rate is greater than the expected random packet loss rate.

16. The apparatus of claim 15, wherein the processor is further configured to compare an actual Layer 4 packet loss rate to an expected Layer 2 packet loss rate.

17. The apparatus of claim 15, wherein the processor is further configured to receive a message reporting the expected random packet loss rate of the network portion.

18. A tangible computer readable medium comprising instructions that when executed by a processor cause the processor to:

send messages to a receiving network device across a network, wherein the network includes a network portion with an expected random packet loss rate;

compare an actual packet loss rate of packets sent across the network to the expected random packet loss rate;

determine that the actual packet loss rate is greater than the expected random packet loss rate; and

compensate for network congestion in response to determining that the actual packet loss rate is greater than the expected random packet loss rate.

19. The computer readable medium of claim 18, wherein the instructions further cause the processor to compare an actual Layer 4 packet loss rate to an expected Layer 2 packet loss rate.

20. The computer readable medium of claim 18, wherein the instructions further cause the processor to receive a message reporting the expected random packet loss rate of the network portion.

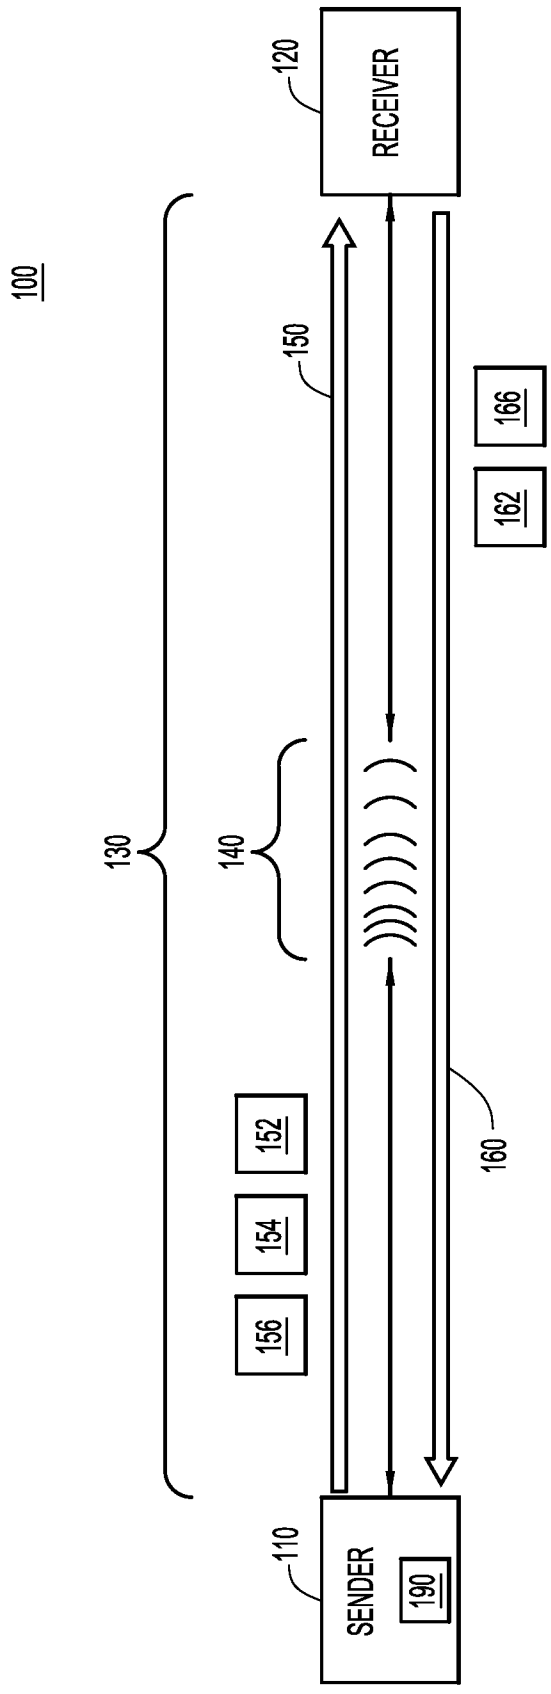

100
120
RECEIVER
150
166
162
130
140
152
154
156
160
110
SENDER
190

FIG.1

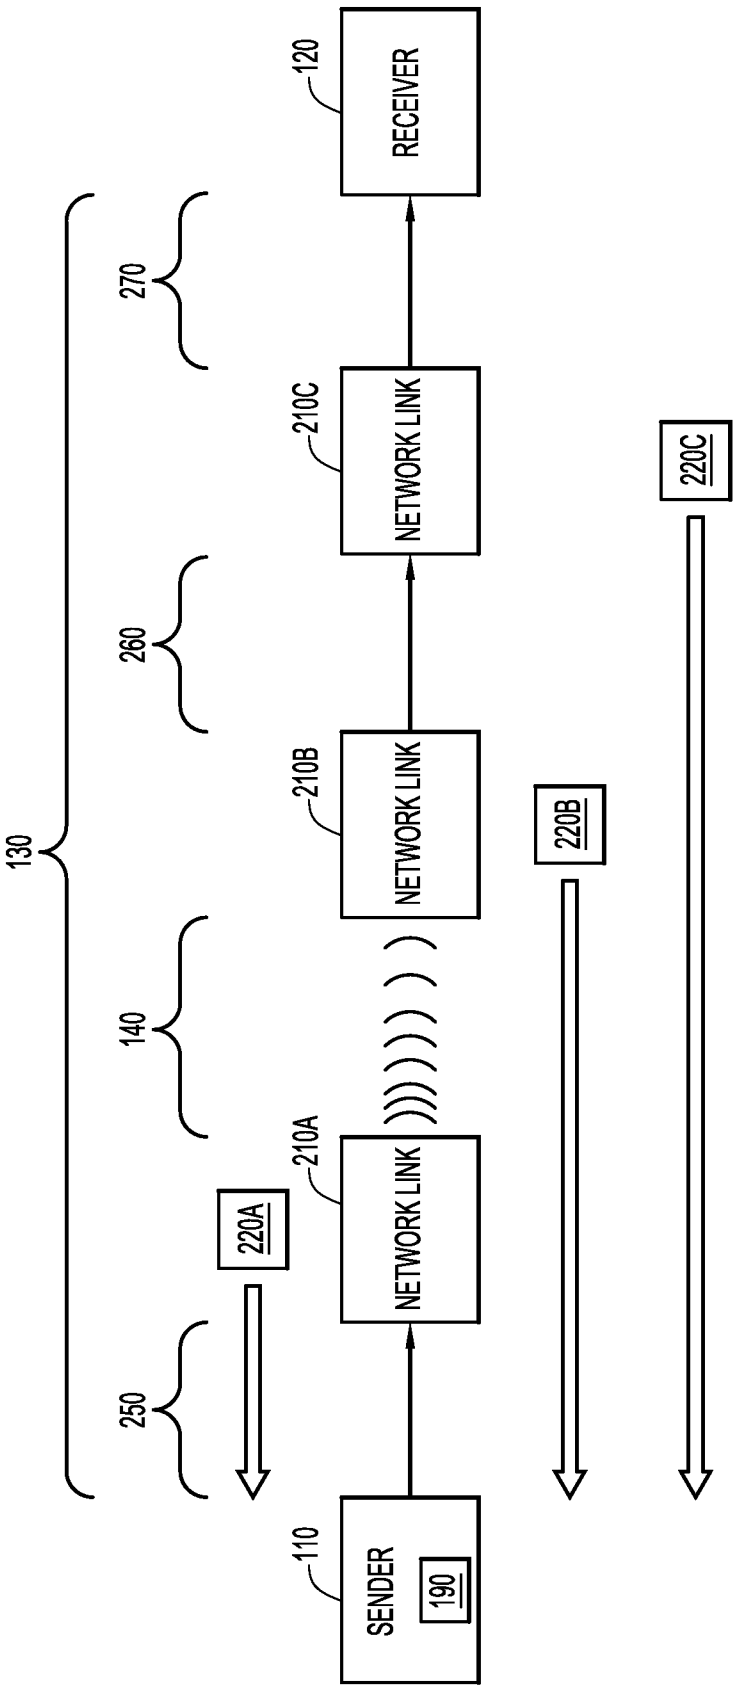

130
250
140
260
270
110
SENDER
190
210A
NETWORK LINK
210B
NETWORK LINK
210C
NETWORK LINK
120
RECEIVER
220A
220B
220C

FIG.2

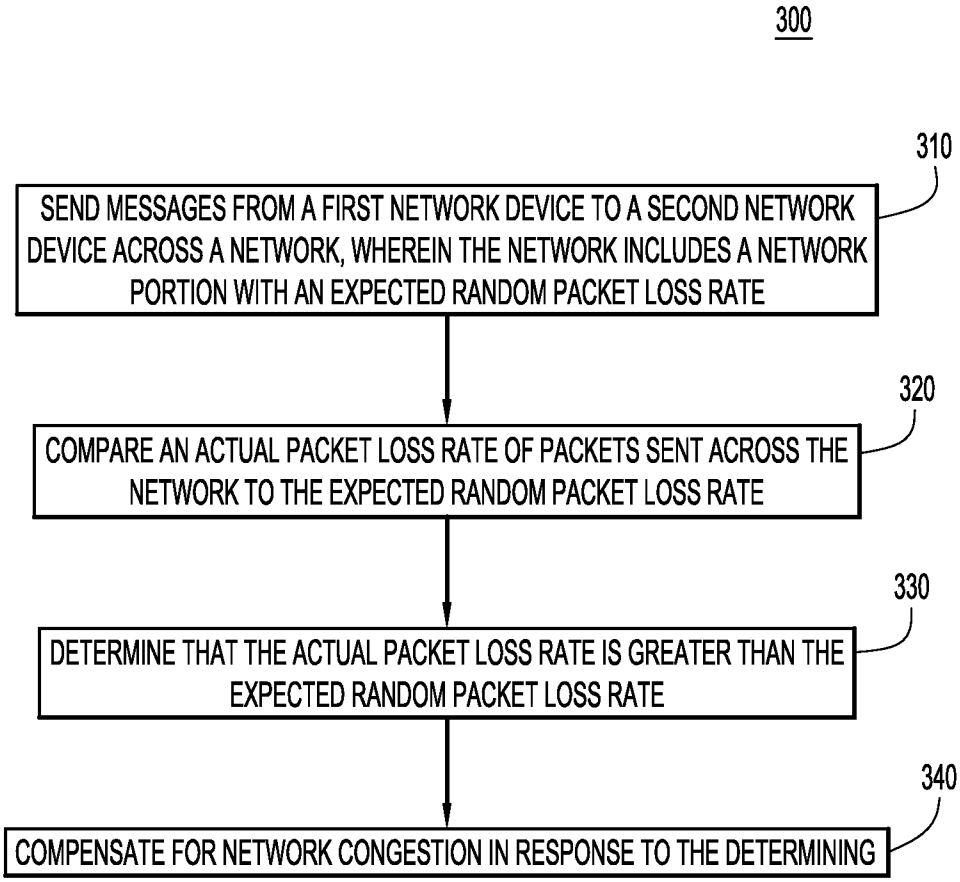

300
310
SEND MESSAGES FROM A FIRST NETWORK DEVICE TO A SECOND NETWORK DEVICE ACROSS A NETWORK, WHEREIN THE NETWORK INCLUDES A NETWORK PORTION WITH AN EXPECTED RANDOM PACKET LOSS RATE
320
COMPARE AN ACTUAL PACKET LOSS RATE OF PACKETS SENT ACROSS THE NETWORK TO THE EXPECTED RANDOM PACKET LOSS RATE
330
DETERMINE THAT THE ACTUAL PACKET LOSS RATE IS GREATER THAN THE EXPECTED RANDOM PACKET LOSS RATE
340
COMPENSATE FOR NETWORK CONGESTION IN RESPONSE TO THE DETERMINING

FIG.3

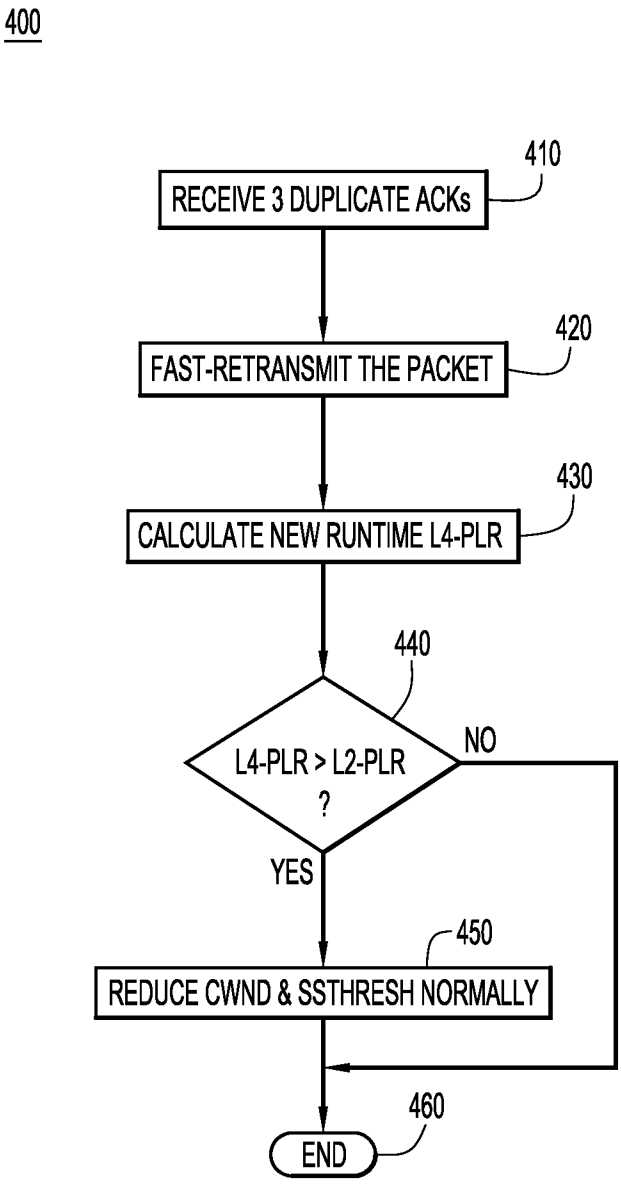

400
410
RECEIVE 3 DUPLICATE ACKs
420
FAST-RETRANSMIT THE PACKET
430
CALCULATE NEW RUNTIME L4-PLR
440
L4-PLR > L2-PLR ?
NO
YES
450
REDUCE CWND & SSTHRESH NORMALLY
460
END

FIG.4

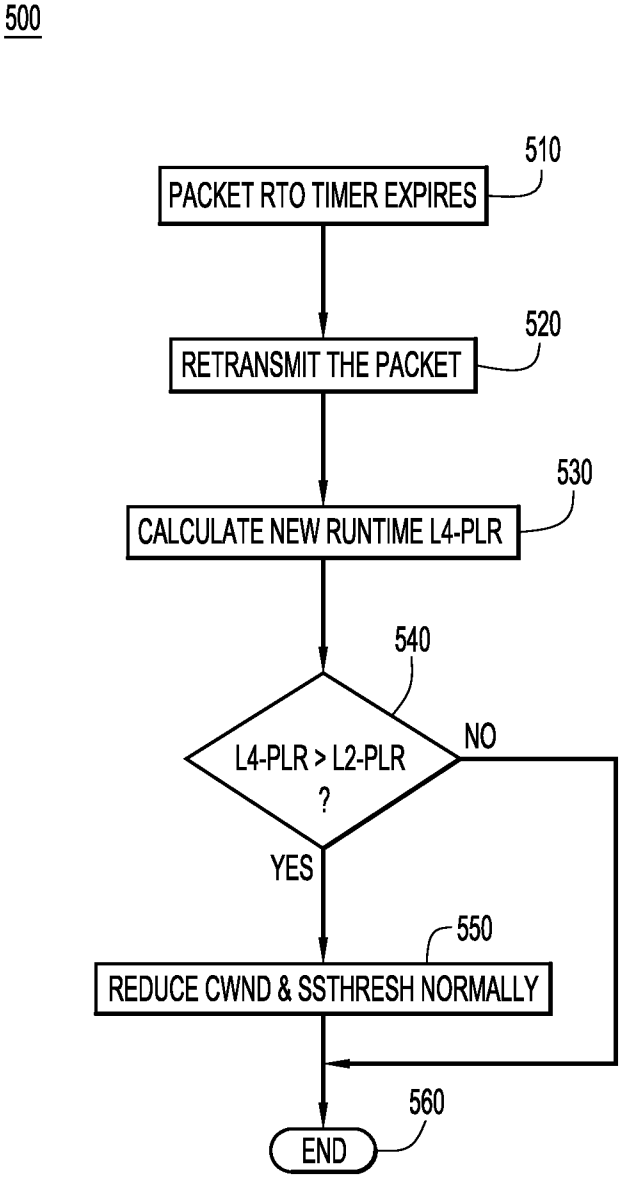

500
510
PACKET RTO TIMER EXPIRES
520
RETRANSMIT THE PACKET
530
CALCULATE NEW RUNTIME L4-PLR
540
L4-PLR > L2-PLR ?
NO
YES
550
REDUCE CWND & SSTHRESH NORMALLY
560
END

FIG.5

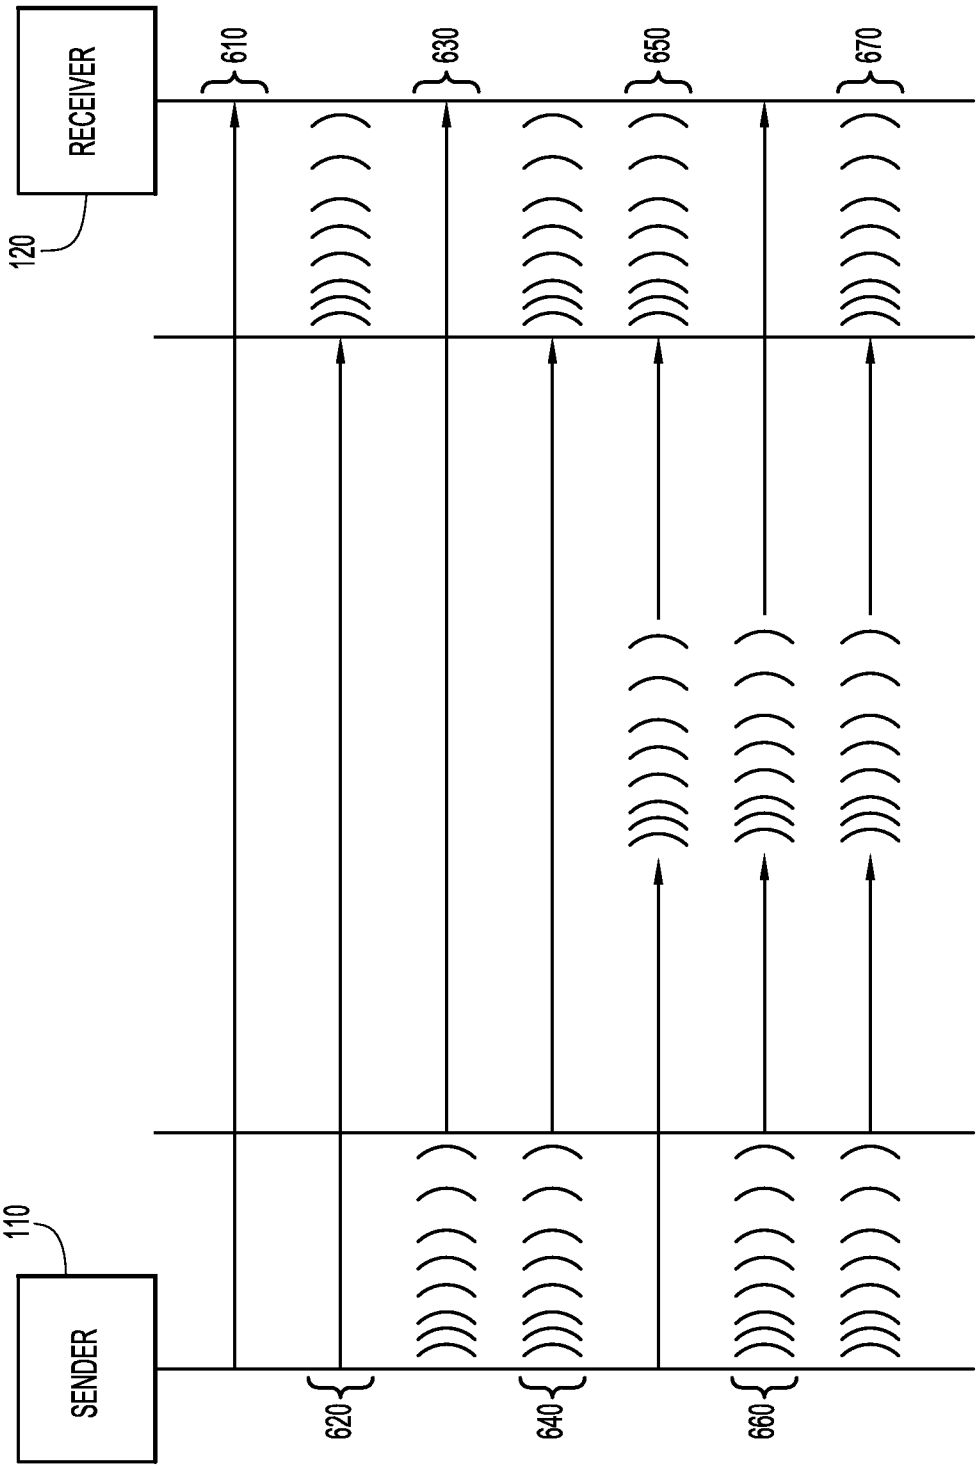


FIG.6

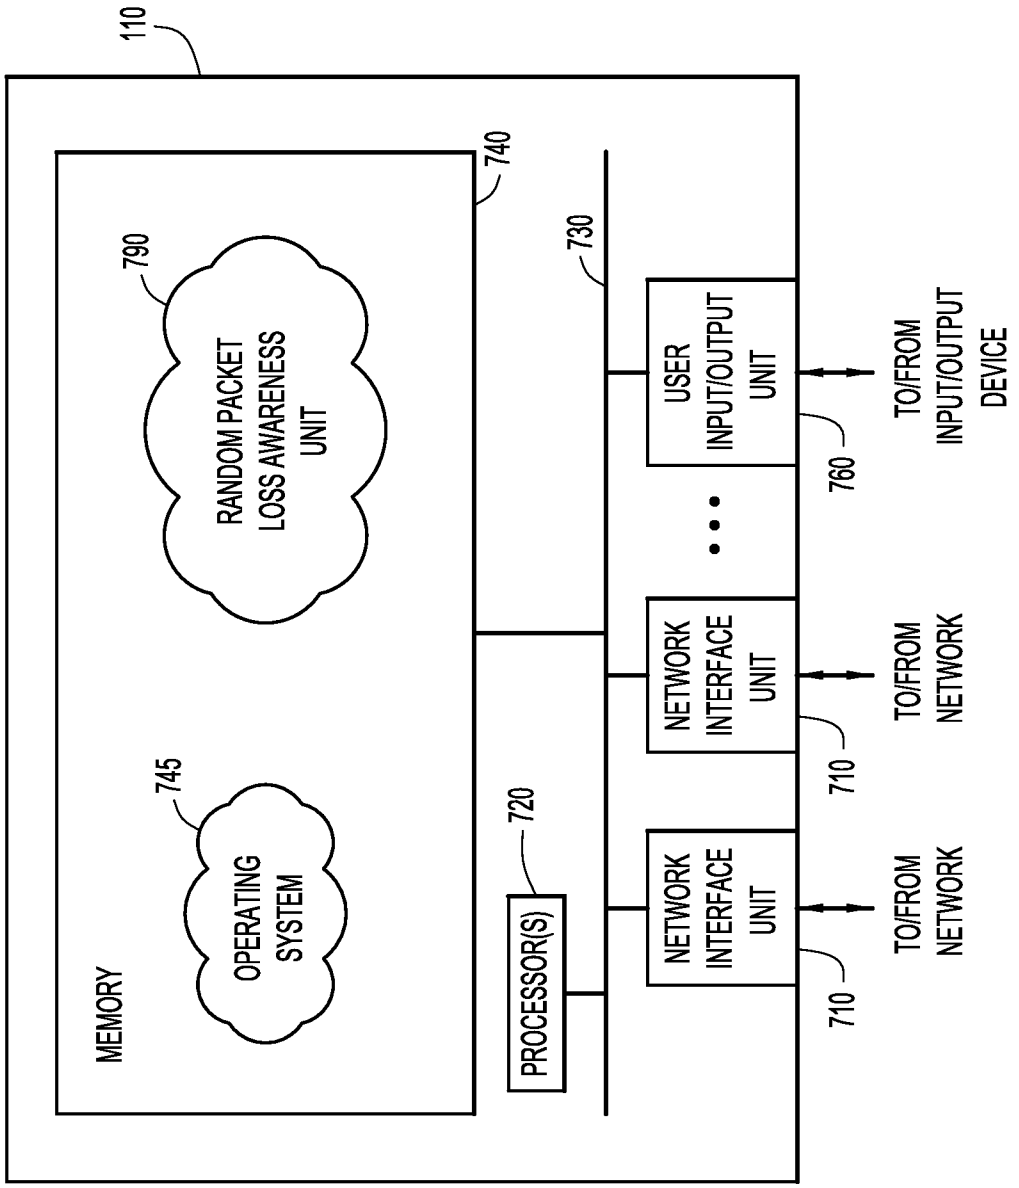

110
740
790
RANDOM PACKET LOSS AWARENESS UNIT
730
745
OPERATING SYSTEM
MEMORY
720
PROCESSOR(S)
NETWORK INTERFACE UNIT
NETWORK INTERFACE UNIT
USER INPUT/OUTPUT UNIT
710
710
760
TO/FROM NETWORK
TO/FROM NETWORK
TO/FROM INPUT/OUTPUT DEVICE

FIG.7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/036909

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06 H04W80/06 H04L12/801
ADD.

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)
H04W H04L

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)

EPO-Internal, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 02/19654 A2 (UNIV CALIFORNIA [US]) 7 March 2002 (2002-03-07) page 5, line 18 - line 34 ----- -/--	1-20

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* 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	Date of mailing of the international search report
22 September 2014	30/09/2014

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	de la Peña Álvarez

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/036909

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	DONG PROTOCOL H ZHAO NANJING UNIV OF POSTS AND TELECOM Y: "A Wireless Channel Model Based Rate Control (WMRC) Scheme in RTP/UDP for Real Time Multimedia Transmissions over Wired-Wireless Networks(WMRC- RTP/UDP); draft-dong-wmrc-rtpcontrol-00.txt", A WIRELESS CHANNEL MODEL BASED RATE CONTROL (WMRC) SCHEME IN RTP/UDP FOR REAL TIME MULTIMEDIA TRANSMISSIONS OVER WIRED-WIRELESS NETWORKS(WMRC- RTP/UDP); DRAFT-DONG-WMRC-RTPCONTROL-00.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, I, 24 December 2010 (2010-12-24), pages 1-14, XP015073213, [retrieved on 2010-12-24] section 4.2 -----	1-20
X	CHENG PENG FU ET AL: "TCP Veno: TCP Enhancement for Transmission Over Wireless Access Networks", IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, IEEE SERVICE CENTER, PISCATAWAY, US, vol. 21, no. 2, February 2003 (2003-02), XP011065582, ISSN: 0733-8716, DOI: 10.1109/JSAC.2002.807336 section II -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/036909

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 0219654 A2	07-03-2002	AU 8858901 A	13-03-2002
		US 2002089930 A1	11-07-2002
		US 2002154602 A1	24-10-2002
		WO 0219654 A2	07-03-2002