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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
23 July 2009 (23.07.2009)

PCT

(10) International Publication Number
WO 2009/092003 A1

(51) International Patent Classification:
G06F 11/00 (2006.01)

(21) International Application Number:
PCT/US2009/031286

(22) International Filing Date: 16 January 2009 (16.01.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/021,890 17 January 2008 (17.01.2008) US

(71) Applicant *(for all designated States except US)*: **BIGFOOT NETWORKS, INC.** [US/US]; 3925 W. Braker Lane, 3rd Floor, Austin, Texas 78759 (US).

(72) Inventors: **BEVERLY, Harlan, T.**; 586 Brown Road, McDade, Texas 78650 (US). **MUSTA, Charles, A.**; 6629 Whitemarsh Valley Walk, Austin, Texas 78746 (US).

(74) Agent: **LARSON NEWMAN ABEL & POLANSKY, LLP**; 5914 West Courtyard Drive, Suite 200, Austin, Texas 78730 (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, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

(54) Title: NETWORK MESSAGE MANAGEMENT DEVICE AND METHODS THEREOF

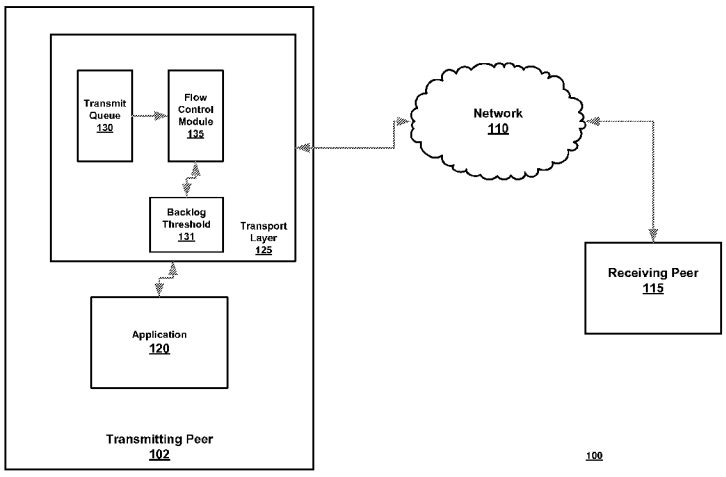


FIG. 1

(57) Abstract: A method of managing communication of messages via a network includes storing messages received from an application in a queue at a transport layer. The transport layer monitors the amount of information stored at the queue. When the amount of information exceeds a threshold amount, the transport layer provides a warning message to the application. In an embodiment, the transport layer provides the warning message before the queue assigned to store messages from the application is full, so that the application can continue to provide messages to the transport layer after the warning message. The application can take appropriate action in response to the message, such as reducing the amount or frequency of information provided at the transport layer, thereby reducing the likelihood of a communication bottleneck at the transport layer.

WO 2009/092003 A1

NETWORK MESSAGE MANAGEMENT DEVICE AND METHODS THEREOF

BACKGROUND

Field of the Disclosure

The present disclosure relates to network communication, and more particularly to management of communication flow between electronic devices.

Description of the Related Art

In network communication, such as communication between peers in a peer-to-peer network, an application at a source communication node sends messages to a destination node via a network layer. In one configuration, the application provides each message to a transport layer, which parses each message into a set of packets, and provides the packets to the network layer for communication to the destination node. The transport layer thus provides a layer of abstraction for the application, so that the application does not have to manage packet communications with the destination node.

However, the transport layer typically has a limited number of resources, such as memory resources, to prepare and communicate packets. This can become problematic when the messages are provided by the application much more quickly than packets can be reliably communicated by the transport layer. For example, if the transport layer is implementing an ordered communication protocol, such as TCP/IP, the transport layer typically stores a transmitted packet in a queue until an acknowledgement is received from the destination node indicating that the packet has been properly received at the destination node, thereby allowing the transport layer to retire the packet from the queue. The queue is typically limited in size, so that if the application provides messages such that packets cannot be retired quickly enough, a bottleneck can result, causing loss of packets or other communication problems. The bottleneck problem can be addressed by having the application set the rate of communications of messages at a fixed rate that is known to avoid bottleneck of the queue. However, this can result in underutilization of resources at the transport layer.

BRIEF DESCRIPTION OF THE DRAWINGS

The present disclosure may be better understood, and its numerous features and advantages made apparent to those skilled in the art by referencing the accompanying drawings.

FIG. 1 is a block diagram of a communications network in accordance with one embodiment of the present disclosure.

FIG. 2 is a flow diagram of a method of providing messages from an application to a transport layer in accordance with one embodiment of the present disclosure.

FIG. 3 is a method of managing the flow of messages provided from an application to a transport layer in accordance with one embodiment of the present disclosure.

FIG. 4 is a block diagram of a communication network in accordance with another embodiment of the present disclosure.

FIG. 5 is a block diagram of a communication network in accordance with another embodiment of the present disclosure.

FIG. 6 is a block diagram of a communication network in accordance with another embodiment of the present disclosure.

DETAILED DESCRIPTION

A method of managing communication of messages via a network includes storing messages received from an application in a queue at a transport layer. The transport layer monitors the amount of information stored at the queue. When the amount of information exceeds a threshold amount, the transport layer provides a warning message to the application. In an embodiment, the transport layer provides the warning message before the queue assigned to store messages from the application is full, so that the application can continue to provide messages to the transport layer after the warning message. The application can take appropriate action in response to the message, such as reducing the amount or frequency of information provided at the transport layer, thereby reducing the likelihood of a communication bottleneck at the transport layer.

Referring to FIG. 1, a block diagram of a particular embodiment of a communications network 100 is illustrated. The communications network 100 includes a transmitting peer 102, a network 110, and a receiving peer 115. The transmitting peer 102 and receiving peer are each connected to the network 110. It will be appreciated that the transmitting peer 102 and receiving peer 115 are referred to as transmitting and receiving for illustrative purposes, and each peer can both transmit and receive information via the network 110. Thus, receiving peer 115 can also transmit messages to transmitting peer 102, which can receive the transmitted messages. It will further be appreciated that, although for purposes of discussion transmitting peer 102 and receiving peer 115 are illustrated as communication nodes in a peer to peer network, in other embodiments the peers can be arranged in a client server fashion. In addition, the transmitting peer 102 and receiving peer 115 are sometimes referred to as "nodes" or "communication nodes." As used herein a node refers to an electronic device in a network that is able to receive, transmit, or route a communication message. The node that receives a message is referred to as the target node of the message, while the node that creates and transmits the message is referred to as the source node. For purposes of discussion herein, it is assumed that network 110 is a wide-area packet-switched network, such as the Internet.

Transmitting peer 102 includes an application 120 and a transport layer 125. The application 120 is a software application embodied on a computer readable medium including instructions to

manipulate a processor in order to execute specified tasks. One such task includes creating messages for communication to receiving peer 115. In particular, in response to specified conditions, such as a user input, request from another application, or other condition, the application 120 can create one or more messages including one or more data payloads for communication to receiving peer 115. The application 120 provides the messages to the transport layer 125 for communication via the network 110.

The transport layer 125 is configured to receive messages from the application 120 and form packets based on the messages. The transport layer 125 then provides the packets to the network 110 for communication to the target node for each packet. In an embodiment, the transport layer 125 can provide each packet to another communication layer, such as a logical or physical communication layer, for provision to the network 110.

The transport layer 125 includes a transmit queue 130, which is configured to store each packet that has been provided to the network 110 for communication. The transport layer 125 awaits an acknowledgement message for each packet from the target node of each packet and, in response to receiving the acknowledgement message, retires the packet from the transmit queue 130. If the transport layer receives an warning message associated with a particular packet from the target node of that packet, the transport layer 125 can retransmit the packet. The transmit queue 130 thereby improves the reliability of communications with target nodes.

This can be better understood with reference to an example. The transport layer 125 can receive messages from application 120 targeted to receiving peer 115. In response, the transport layer 125 forms packets including the data payloads of the received messages. The transport layer provides each packet to the network 110 for communication to the receiving peer 115. In addition, the transport layer stores each packet at the transmit queue 130.

The receiving peer 115 monitors the network 110 for packets targeted to the peer. In response to receiving such a packet, the receiving peer stores the packet at a local queue (not shown). For each packet received, the receiving peer 115 sends an acknowledgement message to the source node via the network 110. The receiving peer can also monitor received packets to determine if an expected packet has not been received. For example, if packets are being communicated to the receiving peer 115 according to an ordered protocol, the transport layer at the receiving peer 115 can determine that an packet expected in the order has not been received. In response, the receiving peer 115 can communicate an warning message to indicate the expected packet has not been received.

The transport layer 125 receives the acknowledgment messages and warning messages provided by the receiving peer 115. In response to receiving an acknowledgment message, the transport layer 125 retires the packet associated with the acknowledgement message from the transmit queue 130, thereby freeing space in the queue. In response to receiving an warning message, the transport layer 125 retransmits the packet associated with the warning message.

The transport layer 125 includes a flow control module 135 that monitors the amount of unretired data that has been communicated to the receiving node 115. That is, the flow control module 135 monitors the amount of data associated with packets at the transmit queue 130. The amount of unretired data communicated to a target node is referred to for purposes of discussion as the "backlog" for the target node. The flow control module 130 compares the amount of backlog to a backlog threshold 131. If the amount of backlog exceeds the backlog threshold, the flow control module 130 can communicate an warning message to the application 120. In response to receiving the warning message, the application 120 can take appropriate action, such as throttling (e.g. reducing) the rate at which messages targeted to the receiving peer 115 are communicated to the transport layer. The flow control module 135 thereby reduces the likelihood of a communication bottleneck at the transport layer 125 or the receiving peer 115. As used herein, a communication bottleneck refers to a situation at the transport layer 125 where the layer is unable to communicate messages quickly enough so that layer resources, such as memory, reach a defined or programmable threshold.

The backlog threshold 131 indicates an amount of available resources, such as memory space, at the transport layer 125 or the receiving peer 115. In an embodiment, the backlog threshold 131 is set so that the flow control module communicates an warning message before the amount of available resources indicated by the threshold is exhausted. For example, the backlog threshold 131 can be set so that a warning message is communicated before the transmit queue 130 is full. This reduces the likelihood of a bottleneck at the queue. In particular, the transport layer can continue to receive messages from the application 120 after the warning message is communicated, as resources are still available to store the messages.

In one embodiment, the backlog threshold 131 can be a programmable value that is set based on a request from the application 120. In another embodiment, the backlog threshold 131 is set based on a request from the receiving peer 115. For example, during negotiation of a communication session, the receiving peer 115 can determine an amount of resources, such as memory, available for communications from the transmitting peer 102. Based on the amount of resources, the receiving peer 115 requests the backlog threshold be set to a value that reduces the likelihood of a communication bottleneck at the receiving peer 115.

Accordingly, as discussed above and further herein, the communication flow, such as the rate of messages being provided by the application 120, is managed or influenced at the transport layer 125. This allows for a reduced likelihood of communication bottleneck without the application 120 having to monitor communication overhead. This can provide for simplification of the application 120, as well as increased communication bandwidth. Further, this allows for a reduced use of memory resources and computer cycles for the application 120. For example, in response to an indication from the flow control module 135 that the threshold has been reached, the application 120, or a processor executing the application 120, can enter a low power mode for a defined amount of time, rather than continue to attempt to send messages via transport layer 125. The power consumption of peer 402 can thereby be reduced.

The operation of the communication network 100 can be better understood with reference to FIGs. 2and 3. FIG 2 illustrates a flow diagram of a particular embodiment of a method of communicating messages by the application 120. At block 202, the application 120 determines messages to be sent to the receiving peer 115. At block 204, the application 120 provides the next message to be sent to the transport layer 125. At block 206, the application 120 determines if a backlog warning message has been received from the flow control module 135. If not, method flow returns to block 204 and the application 120 provides the next message to the transport layer 125 at the present communication rate. If a backlog warning message has been received, the method flow moves to block 208 and the application 120 adjusts (throttles) the rate at which messages are provided to the transport layer 125. For example, the application 120 can delay providing messages to the transport layer 125 for a programmable or defined amount of time. The method flow moves to block 204 and the application 120 provides the next message at the new communication rate. Thus, the application 120 adjusts the rate at which it provides messages to the transport layer 125 based on warning messages received from the layer, thereby reducing the chance of communication bottleneck.

FIG. 3 illustrates a flow diagram of a particular embodiment of a method of communicating packets at the transport layer 125. At block 302, the transport layer 125 receives a message from the application 120. In response, at block 304, the transport layer 125 parses the received message into one or more packets for communication via the network 110. At block 306, the transport layer 125 communicates the packets to the network 110, and stores the communicated packets at the transmit queue 130. At block 308, the flow control module 135 determines if the backlog exceeds the backlog threshold. If not, the method flow returns to block 306. If so, the method flow moves to block 310 and the transport layer 125 communicates a backlog warning message to the application 120. The transport layer 125 thus provides backpressure for messages from the application 120 via the backlog warning messages.

FIG. 4 illustrates a communication network 400 in accordance with one embodiment of the present disclosure. The communication network 400 includes a transmitting peer 402, a network 410, and a receiving peer 415, each of which is connected and configured similarly to the corresponding items of FIG. 1. The transport layer 425 of the transmitting peer 402 includes multiple backlog thresholds, including backlog threshold 431, backlog threshold 432, and backlog threshold 433. Each of these backlog thresholds stores a different threshold value. The flow control module 435 can compare the amount of backlog associated with receiving peer 415 to each of the backlog thresholds 431-433 and, based on the comparison, determine whether a backlog warning message should be communicated to the application 420. In an embodiment, the transport layer 425 can communicate a different warning message in response to each of the backlog thresholds 431-433 being exceeded. For example, in one embodiment the backlog thresholds 431-433 include progressively larger values, with backlog threshold 431 having the smallest value. In response to determining that the amount of backlog exceeds backlog threshold 431, the transport layer 425 can provide a backlog warning message of a first type to the application 420. In response to the backlog exceeding the backlog threshold 432, the transport layer 425 can provide an warning message of different type to the application 420. Based

on the type of received warning message, the application 420 can take different actions. For example, the application 420 can throttle the rate at which messages are provided for communication at a different rate based on the warning message type. Thus, based on the backlog thresholds 431-433, the transport layer 425 can provide different warning messages to indicate that a bottleneck is becoming progressively more likely, allowing the application 420 to take appropriate action.

FIG. 5 illustrates a communication network 500 in accordance with one embodiment of the present disclosure. The communication network 500 includes a transmitting peer 502, a network 510, and a receiving peer 515, each of which is connected and configured similarly to the corresponding items of FIG. 1. The transmitting peer 502 includes a processor 503 configured to execute the application 520 and also includes a network interface 523 configured to execute the transport layer 525. In an embodiment, the network interface 523 is a network interface device, such as a network interface card or processor, configured to provide a physical layer interface between the network 510 and the processor 503. The network interface 523 can be located in the same physical housing as the processor 503, but remote from it, such as on another card or integrated circuit die. The transport layer 525 can perform the flow control functions described herein. Thus, in the illustrated embodiment of FIG. 5 flow control functions can be implemented at the network interface 523, thereby reducing load at the processor 503 and improving communication bandwidth. In one embodiment, in response to a warning message from the transport layer 525, the processor 502 can be placed in a sleep mode or other reduced power state for a defined amount of time, thereby reducing resources while the network interface 523 continues to transmit packets, thereby clearing the potential bottleneck at transport layer 525.

FIG.6 illustrates a communication network 600 in accordance with one embodiment of the present disclosure. The communication network 600 includes a transmitting peer 602, a network 610, and a receiving peer 615, each of which is connected and configured similarly to the corresponding items of FIG. 1. Communication network 600 also includes a receiving peer 616, which is configured similarly to receiving peer 615. The transmitting peer 602 includes applications 620 and 621, and transport layer 625. Transport layer 625 includes a flow control module 635, transmit queues 630 and 638, and backlog thresholds 631 and 632.

In operation, applications 620 and 621 can establish independent communication sessions with receiving peers 615 and 616 respectively. In response, transport layer 625 assigns memory space for each connection to establish transmit queues 630 and 638, respectively, for the corresponding communication session. Thus, transmit queue 630 is associated with the communication session between application 620 and receiving peer 615, while transmit queue 638 is associated with the communication session between application 621 and receiving peer 616. The transport layer 625 also establishes backlog thresholds for each communication session, illustrated as backlog thresholds 631 and 632, respectively. Specifically, transport layer 625 establishes the backlog threshold 631 for the communication session between application 620 and receiving peer 615, while backlog threshold 632 is established for the communication session between application 621 and receiving peer 616.

The transport layer 625 receives messages from each of the applications 620 and 621 and stores each received message in the associated transmit queue. As the messages are retired (in response to an acknowledgement from the associated receiving peer), they are removed from the respective queue. When the number of stored messages at a transmit queue reaches the associated backlog threshold, the transport layer 625 provides a warning message to the associated application. Thus, in the illustrated embodiment of FIG. 6, the transport layer 625 provides bottleneck warnings for each established communication session. Further, the backlog threshold for each session can be different, allowing for flexible management of transport layer resources.

It will be appreciated that the methods disclosed herein can be performed by a computer program stored on a computer readable medium that tangibly embodies the program. The computer program can include instructions to manipulate a processor to perform one or more of the methods described herein.

The above disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments that fall within the true spirit and scope of the present invention. Thus, to the maximum extent allowed by law, the scope of the present invention is to be determined by the broadest permissible interpretation of the following claims and their equivalents, and shall not be restricted or limited by the foregoing detailed description.

WHAT IS CLAIMED IS:

1. A method, comprising:
obtaining a first message at a transport layer from a first application;
determining at the transport layer a first potential communication bottleneck associated with a target node of the first message;
in response to determining the first potential communication bottleneck, providing a first warning message from the transport layer to the first application, the first warning message indicative of a reduced amount of transmission resources available at the transport layer, the transport layer able to accept obtain additional messages from the first application with the reduced amount of transmission resources.

2. The method of claim 1, wherein determining the first potential communication bottleneck comprises:
determining the first potential communication bottleneck based on a comparison of an amount of data stored in a first queue at the transport layer with a first threshold value.

3. The method of claim 3, further comprising:
determining the first threshold value based on a request from the first application.

4. The method of claim 2, further comprising:
determining the first threshold value based on a request from the target node.

5. The method of claim 2, further comprising:
determining a second potential communication bottleneck associated with the target node based on a comparison of the amount of data stored at the first queue with a second threshold value, the second threshold value different from the first; and
in response to determining the second potential communication bottleneck, providing a second warning message from the transport layer to the first application.

6. The method of claim 5, wherein a type of the first warning message is different than a type of the second warning message.

7. The method of claim 6, wherein the type of the first warning message requests the first application to change a rate of message communication by a first amount and the type of the second warning message requests the first application to change the rate of message communication by a second amount different from the first.

8. The method of claim 2, further comprising:

obtaining a second message at the transport layer from a second application;
determining at the transport layer a first potential communication bottleneck associated with a target node of the second message, the target node of the second message different than the target node of the first message; and
in response to determining the first potential communication bottleneck associated with the target node of the second message, providing a second warning message from the transport layer to the second application.

9. The method of claim 8, wherein determining the second potential communication bottleneck comprises:
determining the second potential communication bottleneck based on a comparison of an amount of data stored in a second queue at the transport layer with a second threshold value.

10. The method of claim 9, wherein the first threshold value is different than the second threshold value.

11. The method of claim 2, wherein the first threshold value is indicative of an amount of data less than the size of the first queue.

12. The method of claim 1, wherein obtaining the first message comprises obtaining the first message from a first processor executing the first application, and wherein the transport layer is executed at a network interface device different from the first processor, the network interface device configured to provide an interface between the first processor and a network.

13. A device, comprising:
a network interface device, the network interface device configured to:
obtain a first message at a transport layer from a first application executing at a processor remote from the network interface device;
determine at the transport layer a first potential communication bottleneck associated with a target node of the first message;
in response to determining the first potential communication bottleneck, provide a first warning message from the transport layer to the first application; and
communicate the first message to the target node via a network.

14. The device of claim 13, wherein the network interface device is configured to:
determine the first potential communication bottleneck based on a comparison of an amount of data stored in a first queue at the transport layer with a first threshold value.

15. The device of claim 14, wherein the network interface device is configured to:
determine the first threshold value based on a request from the first application.

16. The device of claim 14, wherein the network interface device is configured to:
determining the first threshold value based on a request from the target node.

17. The device of claim 14, wherein the network interface device is configured to:
determine a second potential communication bottleneck associated with the target node based on a comparison of the amount of data stored at the first queue with a second threshold value, the second threshold value different from the first; and
in response to determining the second potential communication bottleneck, provide a second warning message from the transport layer to the first application.

18. The device of claim 17, wherein a type of the first warning message is different than a type of the second warning message.

19. The device of claim 18, wherein the type of the first warning message requests the first application to change a rate of message communication by a first amount and the type of the second warning message requests the first application to change the rate of message communication by a second amount different from the first.

20. The device of claim 14, wherein the network interface device is configured to:
obtain a second message at the transport layer from a second application;
determine at the transport layer a second potential communication bottleneck associated with a target node of the second message, target node of the second message different than the target node of the first message; and
in response to determining the second potential communication bottleneck, provide a second warning message from the transport layer to the second application.

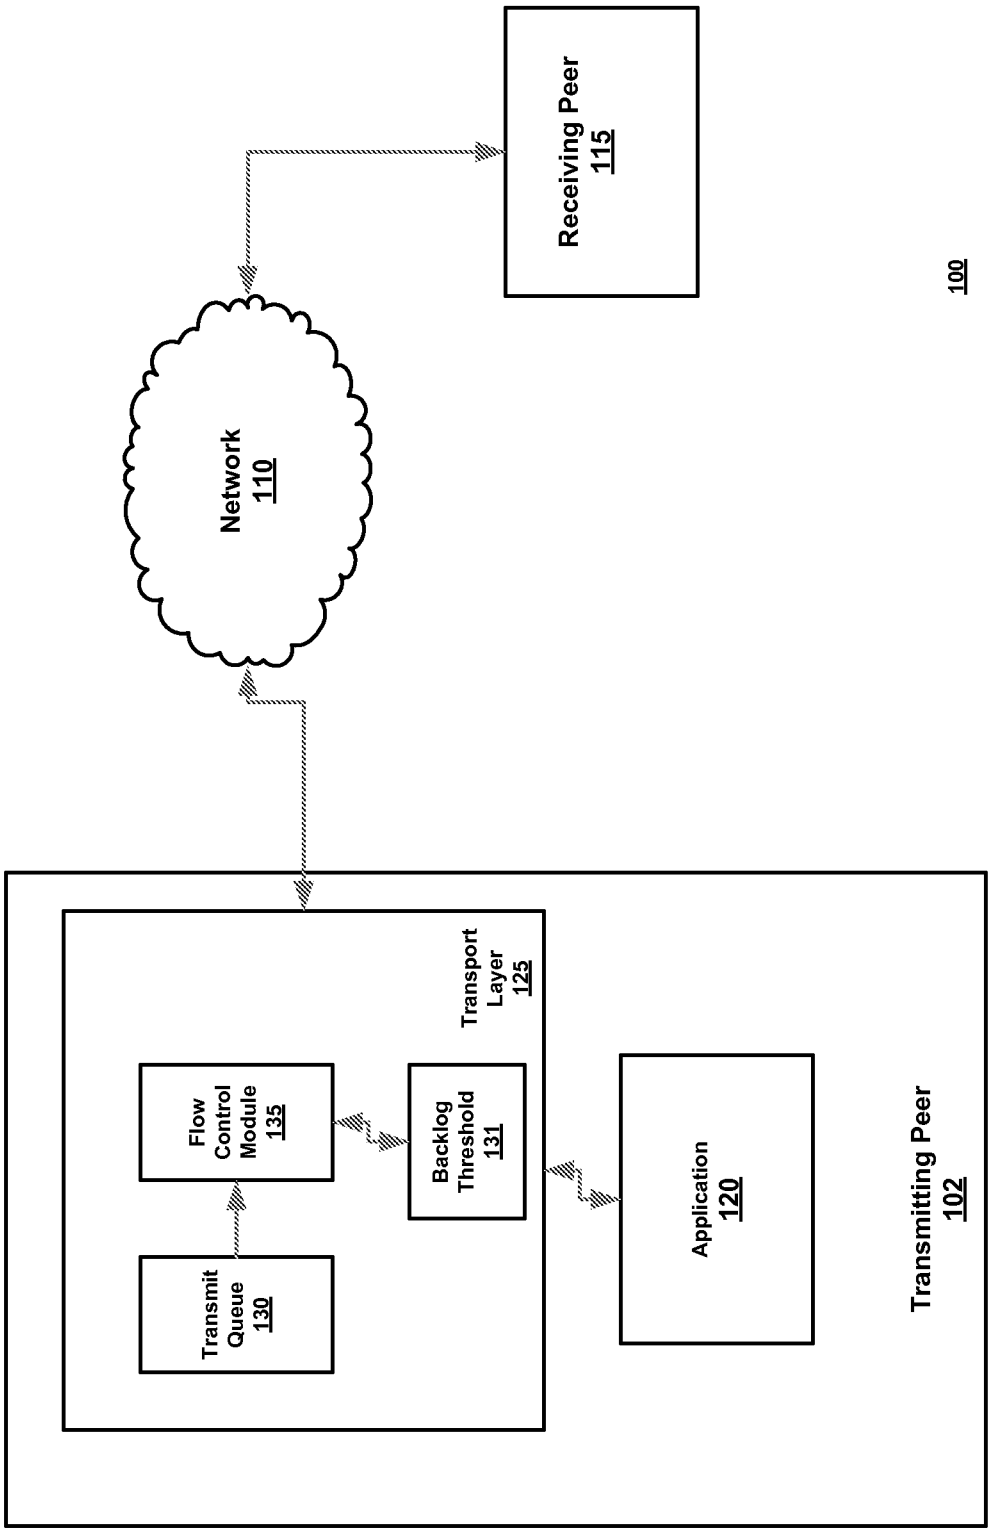

Receiving Peer
115
Network
110
100
Transport
Layer
125
Flow
Control
Module
135
Backlog
Threshold
131
Application
120
Transmit
Queue
130
Transmitting Peer
102

FIG. 1

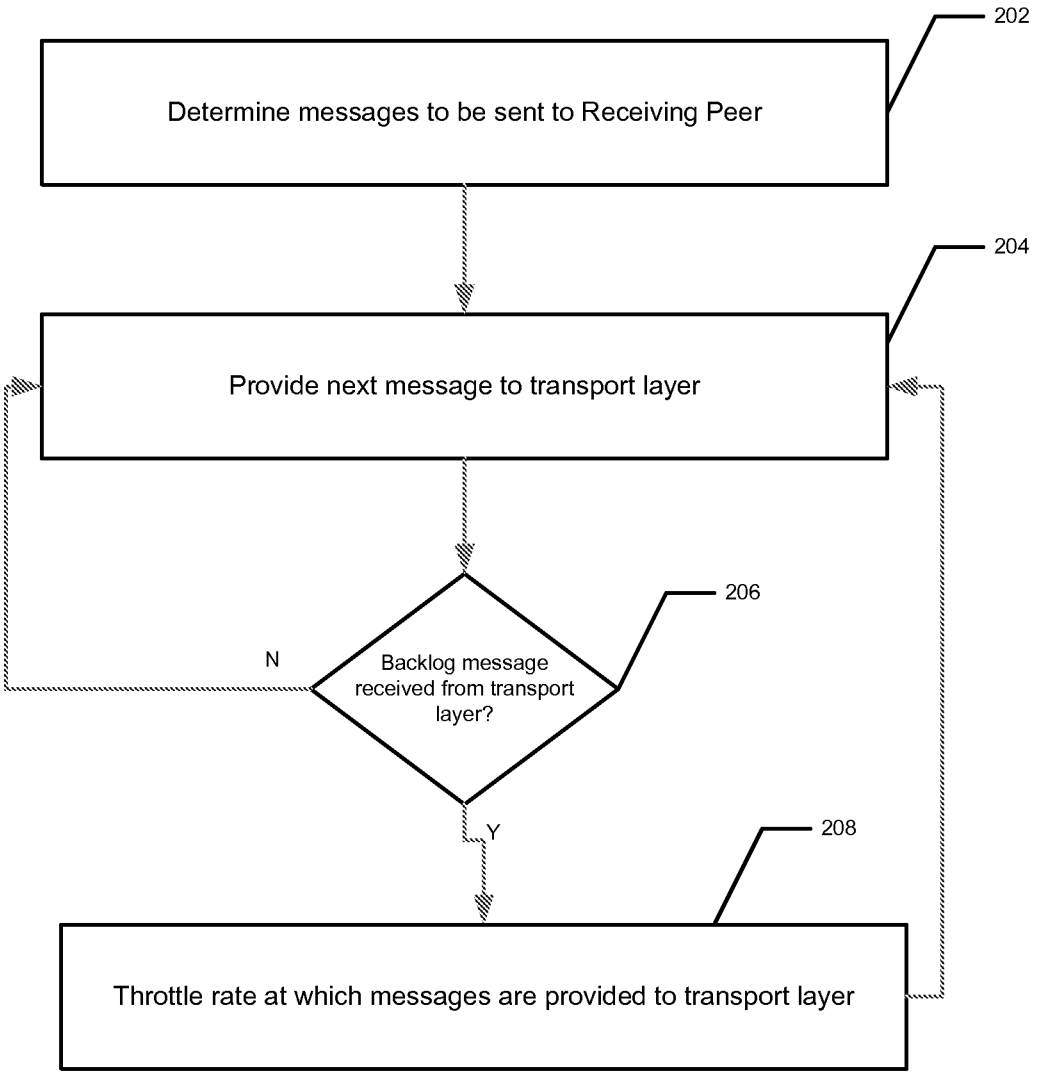

202
Determine messages to be sent to Receiving Peer
204
Provide next message to transport layer
206
Backlog message received from transport layer?
N
Y
208
Throttle rate at which messages are provided to transport layer

FIG. 2

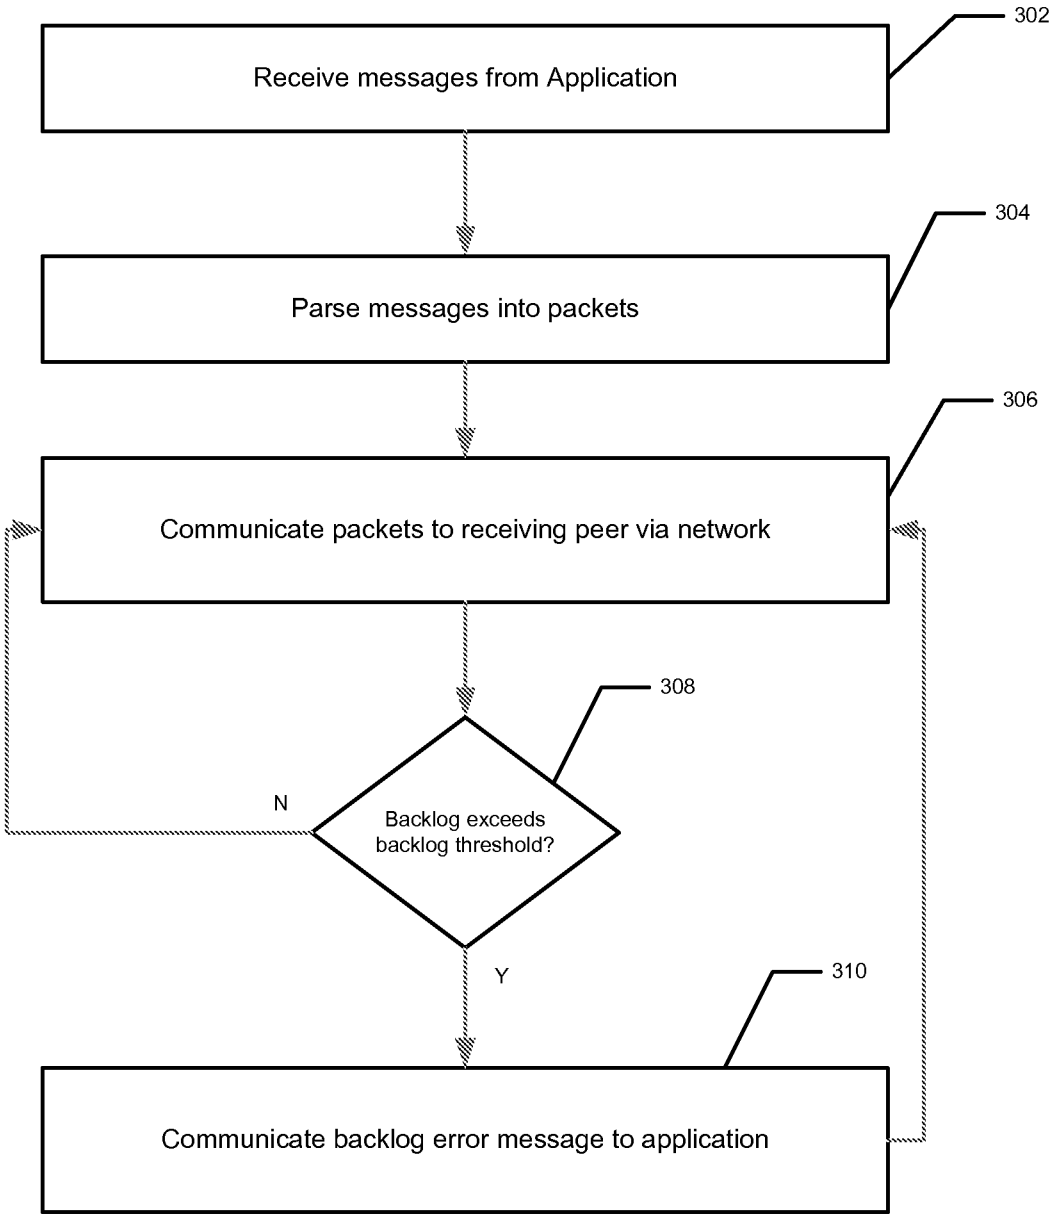

302
Receive messages from Application
304
Parse messages into packets
306
Communicate packets to receiving peer via network
308
Backlog exceeds backlog threshold?
N
Y
310
Communicate backlog error message to application

FIG. 3

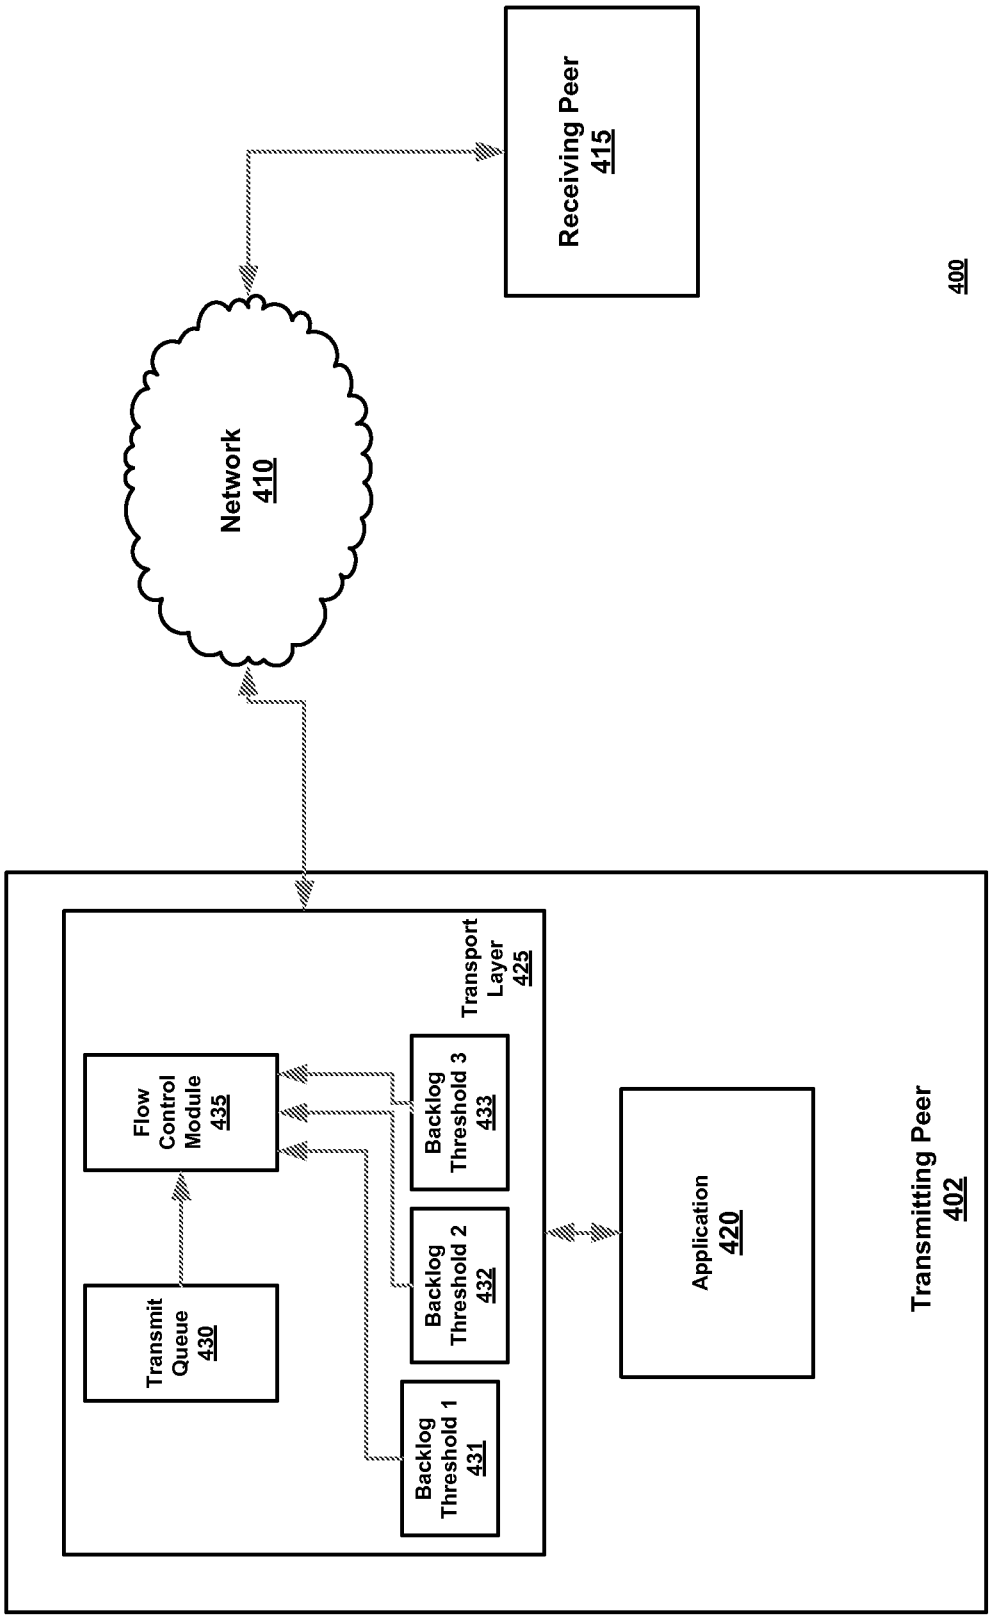

Receiving Peer
415
Network
410
400
Transport
Layer
425
Flow
Control
Module
435
Backlog
Threshold 3
433
Backlog
Threshold 2
432
Backlog
Threshold 1
431
Transmit
Queue
430
Application
420
Transmitting Peer
402

FIG. 4

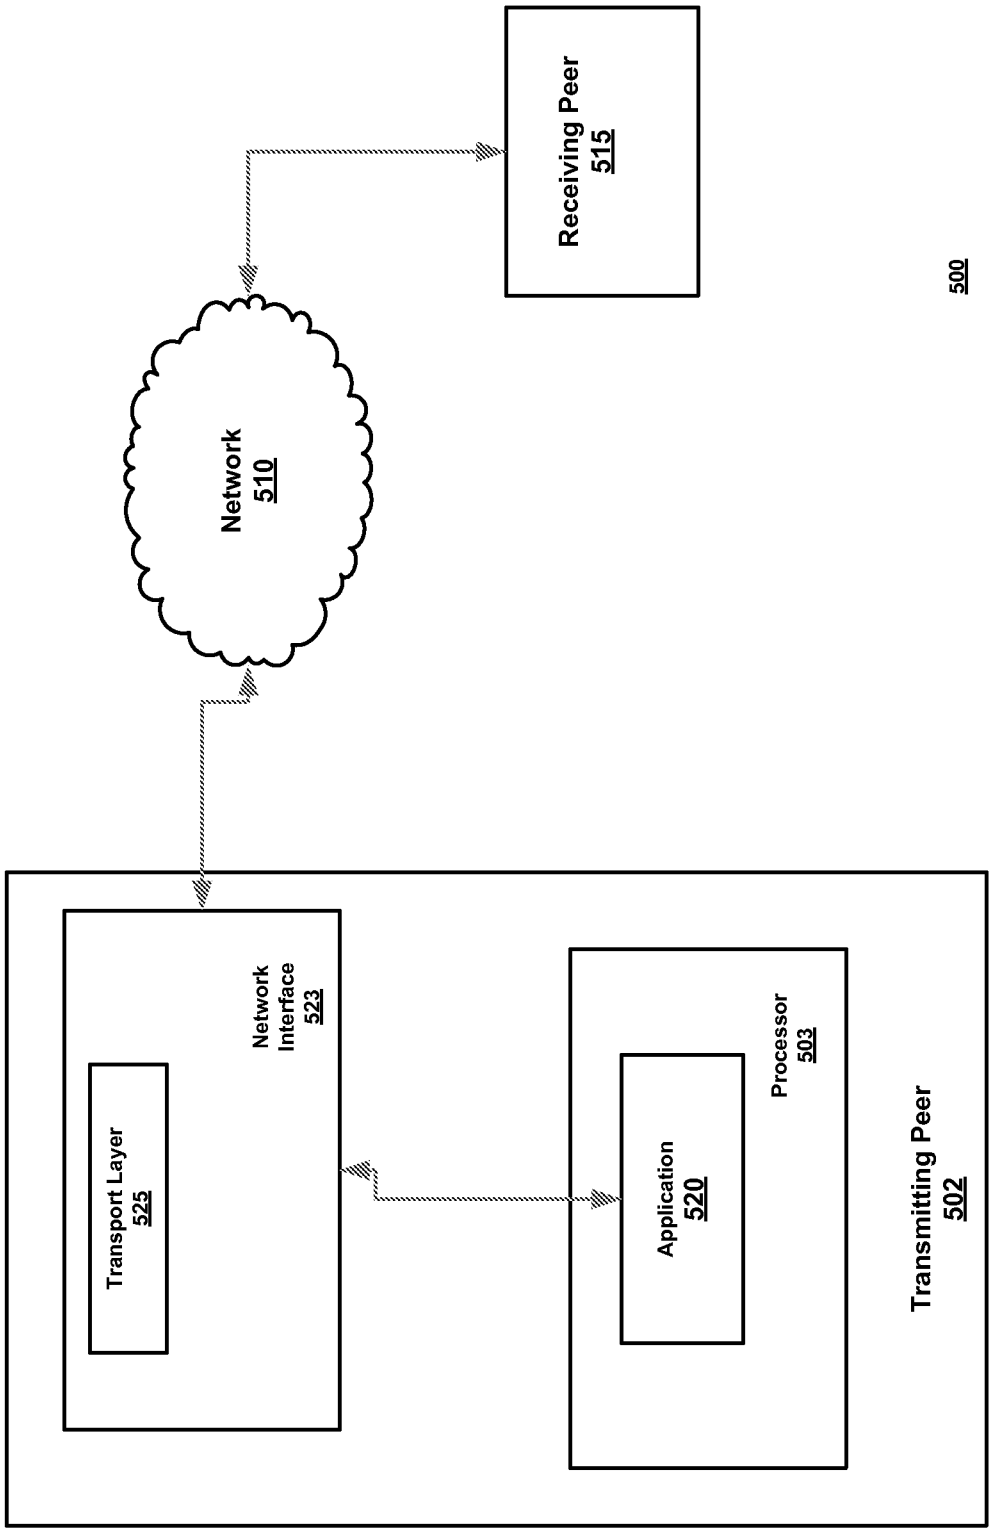

Receiving Peer
515
Network
510
500
Transport Layer
525
Network
Interface
523
Application
520
Processor
503
Transmitting Peer
502

FIG. 5

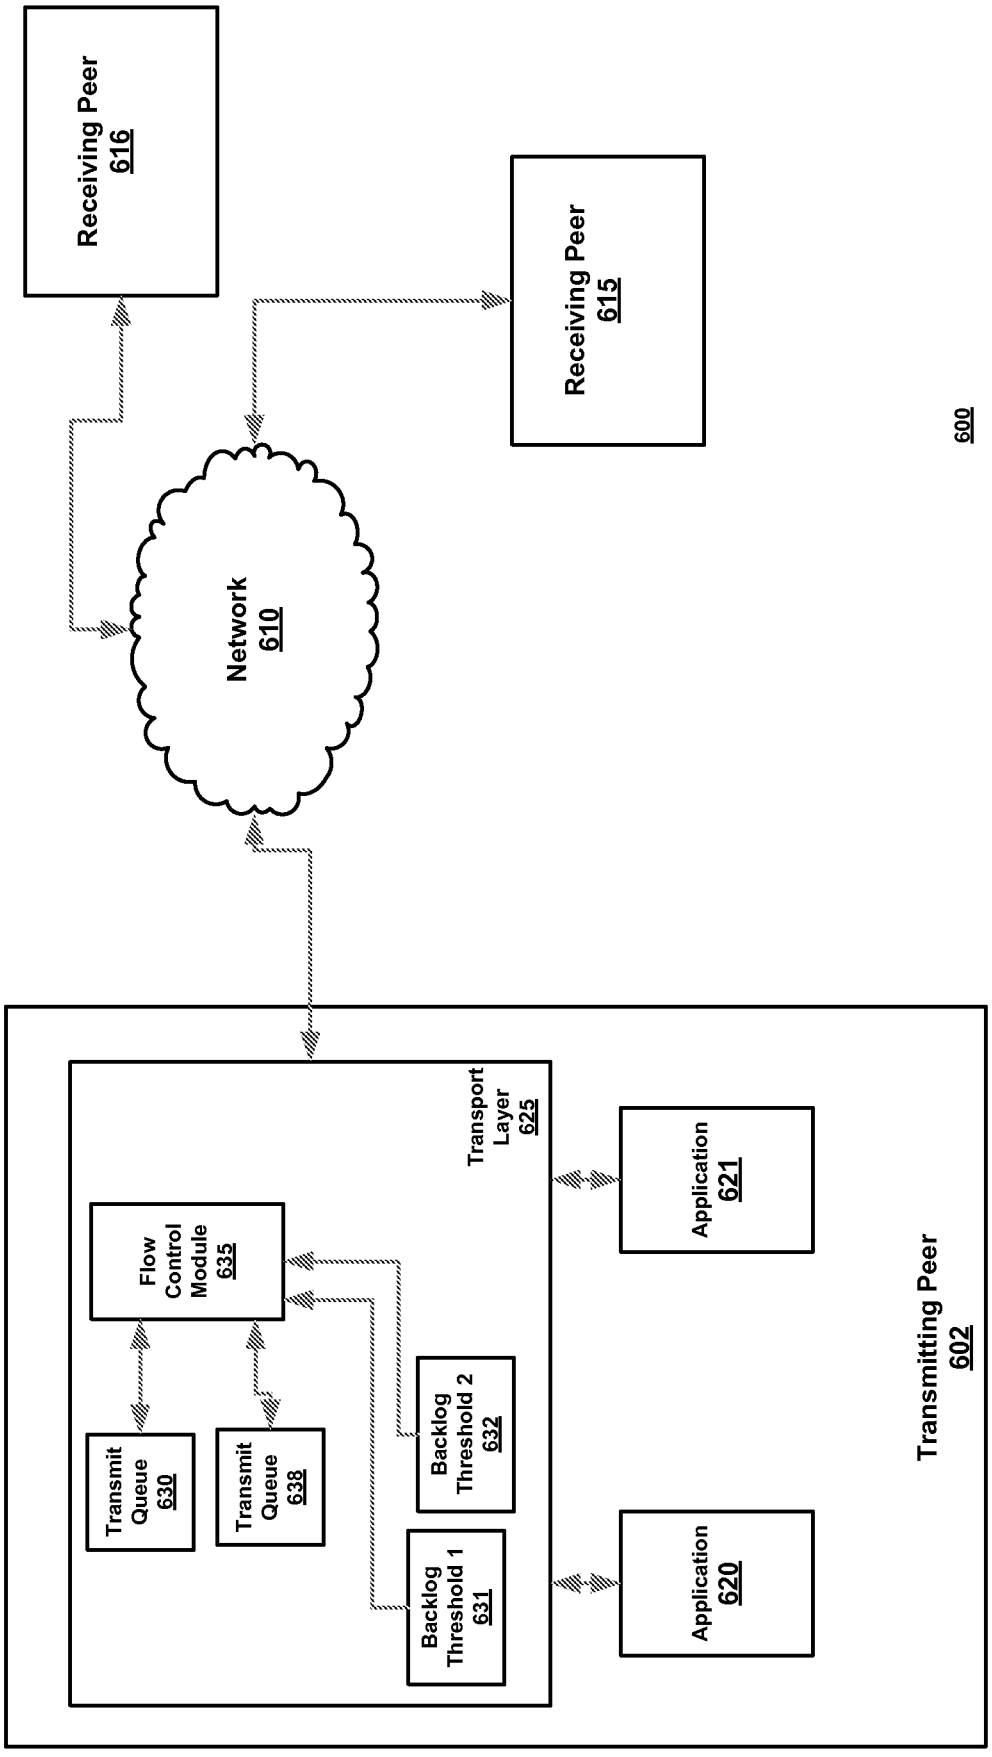

Receiving Peer
616
Receiving Peer
615
Network
610
600
Transmitting Peer
602
Transport
Layer
625
Application
621
Application
620
Flow
Control
Module
635
Transmit
Queue
630
Transmit
Queue
638
Backlog
Threshold 2
632
Backlog
Threshold 1
631

FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 09/31286

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 11/00 (2009.01)
USPC - 370/235
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)
USPC - 370/235

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC - 370/229-237 (keyword limited, see terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PubWEST(PGPB,USPT,EPAB,JPAB); Google. Search Terms Used: transport layer, TCP, network, network device, warning, alarm, queue, threshold, message, identify, notify, feedback, peer, determine, bottleneck, congestion, another, further, subsequent, additional, second, repeat, communication, reduce, decrease, limit, user

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007/0223379 A1 (Sivakumar et al.) 27 September 2007 (27.09.2007) entire document, especially: Fig. 1, Fig. 9, para. [0020], [0021], [0061]-[0069]	1-20
A	US 2004/0218617 A1 (Sagfors) 04 November 2004 (04.11.2004) *entire document*	1-20
A	US 6,975,876 B1 (Cast et al.) 13 December 2005 (13.12.2005) entire document	1-20

☐ 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	Date of mailing of the international search report
24 February 2009 (24.02.2009)	10 MAR 2009
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: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (April 2007)