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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 April 2017 (13.04.2017)

WIPO | PCT

(10) International Publication Number
WO 2017/059932 A1

(51) International Patent Classification:
H04L 12/823 (2013.01) *H04L 12/869* (2013.01)
H04L 12/863 (2013.01) *H04L 29/06* (2006.01)

(21) International Application Number:
PCT/EP2015/079128

(22) International Filing Date:
9 December 2015 (09.12.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/238,209 7 October 2015 (07.10.2015) US

(71) Applicant: **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; S-164 83 Stockholm (SE).

(72) Inventors: **EDWALL, Thomas**; Trevebovägen 35, S-184 32 Åkersberga (SE). **SKOG, Robert**; Gullvivegränd 7, 165 76 Hässelby (SE). **IHLAR, Marcus**; Levertinsgatan 1, 754 30 Uppsala (SE).

(74) Agent: **ERICSSON**; Torshamnsgatan 21-23, 164 80 Stockholm (SE).

(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, 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, ST, 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).

Published:

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

WO 2017/059932 A1

(54) Title: CONTROLLING DOWNSTREAM FLOW OF DATA PACKETS TO ONE OR MORE CLIENTS

Fig. 2

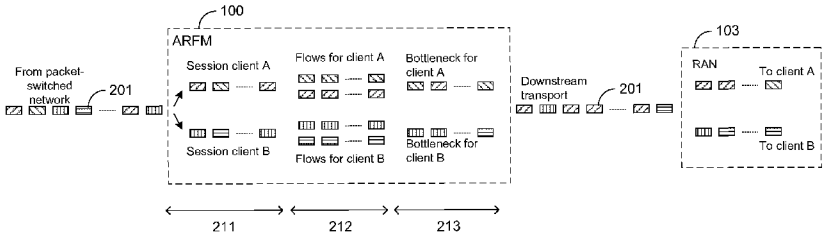

(57) Abstract: A network node (100), referred to as Automatic Rate Flow Manager (ARFM), for controlling a downstream flow of data packets to one or more clients (104) of a Radio Access Network (103) (RAN) is provided. The ARFM comprises a network interface for receiving the data packets from a packet- switched network (101), and processing means being operative to separate the received data packets into different session buffers, one for each client (104), based on a respective destination Internet Protocol (IP) address of each data packet, and, for each client (104), separate the data packets into different queues, one for each flow, de-queue the data packets in accordance with an estimated bitrate of a data connection between the network node (100) and the client (104), and forward the de-queued data packets to the client (104). Thereby, an artificial bottleneck is created at the ARFM (100), reflecting an available downstream (103) path capacity in the network. Advantageously, queue-buildup is thereby limited to the ARFM (100), rather than downstream (103) from the ARFM (100).

CONTROLLING DOWNSTREAM FLOW OF DATA PACKETS TO ONE OR MORE CLIENTS

Technical field

The invention relates to a network node for controlling a downstream flow of data packets to one or more clients, a method of controlling a downstream flow of data packets to one or more clients, a corresponding computer program, and a corresponding computer program product.

Background

Network operators have a need of protecting their networks from congestion, as well as ensuring that the performance of the network is as good as possible.

One way of achieving this is to use Active Queue Management (AQM). In Internet routers, AQM is the practice of intelligently dropping network packets inside a buffer associated with a network interface when the buffer becomes full or gets close to becoming full. The dropping of data packets is often done to reduce overall congestion in the network. The task of dropping packets is performed by a network scheduler, using various algorithms such as Random Early Detection (RED) [S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", IEEE/ACM Transactions on Networking, volume 1, pages 397–413, 1993], Explicit Congestion Notification (ECN) [Internet Engineering Task Force (IETF) RFC 3168], or Controlled Delay (CoDel) [K. Nichols and V. Jacobson, "Controlling queue delay", Queue, volume 10, pages 20–34, ACM, 2012], which are known in the art.

AQM has been an active area of research for decades, and IETF RFC 7567 recommends AQM as a best practice. Deployment of AQM has

been hampered by the need of specific and complex tuning for different types of networks.

Flow-based queue management has been developed as an improvement over traditional AQM. Such algorithms separate flows into different queues and attempt to preserve isolation between different flows, thereby reducing the risk of an aggressive flow taking precedence over, "starving out", a less aggressive flow. As an example, Flow-Queue CoDel (FQ-CoDel) is a state-of-the-art AQM scheme that combines a "zero-tuning AQM" with flow-based scheduling. It is currently being standardized by the IETF [IETF Internet-Draft draft-ietf-aqm-fq-codel-01].

AQM is typically deployed at potential bottlenecks in the network. These bottlenecks include host operating systems, internet routers, Radio Access Network (RAN) nodes, and other types of network equipment. If AQM functionality is deployed at multiple locations in the network, i.e., a distributed deployment, any update of AQM functionality throughout a network is costly and complex.

The use of a Transport Control Protocol (TCP) proxy is an alternative to improve network utilization as well as mitigate network congestion. TCP proxies typically improve network utilization by using specifically tuned Congestion Control Algorithms (CCAs) that perform well in RANs. Fairness between different flows is achieved by utilizing the same type of congestion control for all flows. TCP proxies also have the possibility to accelerate flows by splitting the connection in two connections, a first leg between the upstream end-point of the TCP connection (the TCP server) and the proxy, and a second leg between the proxy and the downstream end-point of the TCP connection (the TCP client). This leads to tighter feedback loops by making TCP ramp up its sending speed faster. A problem with TCP proxies is that the enhanced congestion control only is beneficial for flows with a lifetime which is sufficiently long to leave TCP's slow-start phase. Another problem is that transport flows other than TCP, such as Quick User

Datagram Protocol (UDP) Internet Connections (QUIC), are not handled by the TCP proxy. Adding support for new protocols is complicated by end-to-end encryption of communication sessions which is commonly used between server and client via the proxy.

An operator's radio network is frequently deployed with different radio access technologies, i.e., a is a combination of different types of RANs, such as Global System for Mobile Communications (GSM), Universal Mobile Telecommunications System (UMTS), Long Term Evolution (LTE), and next-generation radio access technologies colloquially termed "5G", each radio-access technology having different characteristics. Operators can configure/tune their RANs in different ways, e.g., using AQM in LTE radio base stations. However, the resulting differences in radio characteristics will have an impact on the characteristics of a transport protocol which is used for transmitting data over the RAN, such as TCP, in particular in the downlink. For instance, a web server located on the Internet, serving mobile users located in different parts of the RAN and accessing the operator network through different radio-access technologies, will have to cater for these differences in order to support as good Quality of Experience (QoE) as possible.

Summary

It is an object of the invention to provide an improved alternative to the above techniques and prior art.

More specifically, it is an object of the invention to provide an improved flow-control of downstream data packets transmitted to one or more clients, in particular clients receiving the data packets over a RAN.

These and other objects of the invention are achieved by means of different aspects of the invention, as defined by the independent claims. Embodiments of the invention are characterized by the dependent claims.

According to a first aspect of the invention, a network node for controlling a downstream flow of data packets to one or more clients is provided. The one or more clients may, e.g., be mobile terminals or User Equipments (UEs) receiving the data packets over a RAN. The network node comprises at least one network interface and processing means. The at least one network interface is arranged for receiving the data packets from a packet-switched network, such as the Internet. The processing means is operative to separate the received data packets into different session buffers, one for each client. The received data packets are separated based on a respective destination Internet Protocol (IP) address of each data packet. The processing means is further operative, for each client, i.e., for each session buffer, to separate the data packets into different queues, one for each flow, de-queue the data packets, and forward the de-queued data packets to the client. Each flow may be defined by a 5-tuple comprising source IP address, destination IP address, source port number, destination port number, and protocol number. The data packets are de-queued in accordance with an estimated bitrate of a data connection between the network node and the client. The de-queued data packets are forwarded by means of the at least one network interface.

According to a second aspect of the invention, a method of controlling a downstream flow of data packets to one or more clients is provided. The method is performed by a network node and comprises receiving the data packets from a packet-switched network, such as the Internet, and separating the received data packets into different session buffers, one for each client. The received data packets are separated based on a respective destination IP address of each data packet. The method further comprises, for each client, i.e., for each session buffer, separating the data packets into different queues, one for each flow, de-queueing the data packets, and forwarding the de-queued data packets to the client. Each flow may be defined by a 5-tuple comprising source IP address, destination IP address,

source port number, destination port number, and protocol number. The data packets are de-queued in accordance with an estimated bitrate of a data connection between the network node and the client.

According to a third aspect of the invention, a computer program is provided. The computer program comprises computer-executable instructions for causing a device to perform the method according to an embodiment of the second aspect of the invention, when the computer-executable instructions are executed on a processing unit comprised in the device.

According to a fourth aspect of the invention, a computer program product is provided. The computer program product comprises a computer-readable storage medium which has the computer program according to the third aspect of the invention embodied therein.

The invention makes use of an understanding that an improved flow-control of downstream data packets which are transmitted to one or more clients can be achieved by means of a network node operating as Automatic Rate Flow Manager (ARFM), which is described herein. The ARFM is a single ingress (entry) and egress (exit) point for a certain client and may, e.g., be deployed between an external packet-switched network, such as the Internet, and an entry point of an operator network or a RAN.

The ARFM attempts to reflect an available path capacity in the network to ensure that queue-buildup is preferably limited to the point in the network where the ARFM is deployed, rather than further down the path, i.e., downstream from the ARFM. In the prior art, bottlenecks may, e.g., build up in radio access nodes of the RAN due to limitations in air-interface capacity. As a consequence of creating an artificial bottleneck for each client at an upstream location, the build-up of queues at potential bottlenecks downstream from the ARFM, e.g., in the operator network or the RAN, is avoided.

The ARFM simplifies deployment of AQM in an operator network

since a single node is sufficient to manage the entire network downstream from the point of deployment. The ARFM may advantageously be deployed in networks where path characteristics may differ considerably between connections to different clients, such as mobile terminals or UEs receiving data over a RAN, or different RANs having different radio characteristics and/or radio access technologies. A typical example of such a network is 3GPP-type of network in which different clients have separate scheduling queues.

According to an embodiment of the invention, the data packets are de-queued at a bitrate which is substantially equal to the estimated bitrate for each client. Optionally, the bitrate of the data connection between the network node and the client may be estimated, for each client, by the network node. For instance, the bitrate may be estimated by measuring the rate of acknowledgements pertaining to a transport protocol used by the client, correlated with the amount of data packets sent to the client. Alternatively, the bitrate may be estimated based on feedback information received from a downstream network node, e.g., a Radio Resource Management (RRM) node of a RAN, such as a Base Station Controller (BSC) in a GSM network, a Radio Network Controller (RNC) in a UMTS network, an eNodeB in an LTE network, or a corresponding RRM node in a 5G network.

According to an embodiment of the invention, the data packets may be de-queued using a fair scheduler or a weighted scheduler, in accordance with a policy decided by an operator of the network node.

According to an embodiment of the invention, the data packets are de-queued in accordance with a maximum queueing delay. Thereby, the sending rate per flow can be regulated. The maximum queueing delay may be configured by an operator of the network node.

Even though advantages of the invention have in some cases been described with reference to embodiments of the first aspect of the invention, corresponding reasoning applies to embodiments of other aspects of the

invention.

Further objectives of, features of, and advantages with, the invention will become apparent when studying the following detailed disclosure, the drawings and the appended claims. Those skilled in the art realize that different features of the invention can be combined to create embodiments other than those described in the following.

Brief description of the drawings

The above, as well as additional objects, features and advantages of the invention, will be better understood through the following illustrative and non-limiting detailed description of embodiments of the invention, with reference to the appended drawings, in which:

Fig. 1 shows possible deployment points for the network node for controlling a downstream flow of data packets to one or more clients, in accordance with embodiments of the invention.

Fig. 2 illustrates flow-control of downstream data packets, in accordance with embodiments of the invention.

Fig. 3 shows a flow chart illustrating a method of controlling a downstream flow of data packets to one or more clients, in accordance with embodiments of the invention.

Fig. 4 shows a network node for controlling a downstream flow of data packets to one or more clients, in accordance with an embodiment of the invention.

Fig. 5 shows a network node for controlling a downstream flow of data packets to one or more clients, in accordance with another embodiment of the invention.

All the figures are schematic, not necessarily to scale, and generally only show parts which are necessary in order to elucidate the invention, wherein other parts may be omitted or merely suggested.

Detailed description

The invention will now be described more fully herein after with reference to the accompanying drawings, in which certain embodiments of the invention are shown. This invention may, however, be embodied in many different forms and should not be construed as limited to the embodiments set forth herein. Rather, these embodiments are provided by way of example so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art.

The solutions disclosed herein utilize a network node 100, herein referred to as Automatic Rate Flow Manager (ARFM), as a single ingress (entry) and egress (exit) point for a certain client 104, such as a mobile terminal or a UE. The ingress/egress point may, e.g., be the default routing between a packet router of the RAN 103 or an operator network 102, such as the Gateway General Packet Radio Service (GPRS) Support Node (GGSN) or the Packet Data Network Gateway (PGW or PDN-GW), and an external packet-switched network 101, such as the Internet. An example deployment of the ARFM is illustrated in Fig. 1, in which the ARFM 100 is located between an external packet-switched network 101 and an entry point of an operator network 102 (illustrated in diagram 110). Alternatively, the ARFM 100 may be located at an entry point of a RAN 103 (illustrated in diagram 120). As a further alternative, ARFM 100 may be deployed at the Gi interface of a 3GPP network.

ARFM 100 attempts to reflect an available path capacity in the downstream network to ensure that queue buildup is mainly limited to the point in the network where ARFM 100 is deployed, rather than further down the path, i.e., downstream from ARFM 100. In the prior art, bottlenecks may, e.g., build up in radio access nodes of the RAN 103 due to limitations in air-interface capacity.

In the following, operation of ARFM 100 is described in more detail with reference to Fig. 2, which illustrates the flow of downstream data packets 201 from left to right. Data packets 201 reaching ARFM 100 from an upstream source, such as the Internet or any other external packet-switched network 101, are separated 211 into different session buffers based on their respective destination IP address, i.e., based on the client 104 which they are addressed to. In Fig. 2, two clients 104, "client A" and "client B" are illustrated, but embodiments of the invention controlling the flow of downstream data packets 201 to any number of clients 104 maybe envisaged. Data packets 201 received from the upstream network are addressed to either client A 104 (downward and upward diagonal pattern in Fig. 2) or client B 104 (horizontal and vertical pattern in Fig. 2). Each client 104 may, e.g., be a mobile terminal or a UE.

The respective path capacity (herein also referred to as bandwidth), more specifically the bitrate, from ARFM 100 to each of the clients 104 is estimated, preferably regularly or continuously. This may, e.g., be achieved by correlating the amount of sent data segments with corresponding acknowledgements received from clients 104. If TCP is used as a transport protocol, the acknowledgement received from clients 104 are signaled by means of the ACK flag.

The flow of data packets 201 within ARFM 100 is regulated on a per-client basis so as to match the path capacity estimated for each client 104. This is achieved by creating artificial bottlenecks, one for each client 104. As a consequence, a queue will build up at an artificial bottleneck if the downstream path between ARFM 100 and client 104 has a lower capacity than the upstream path, between ARFM 100 and the source of the data packets, i.e., the external packet-switched network 101.

For each client 104, the data packets are separated into different queues 212 based on which flow they belong to. Each flow is defined by its 5-tuple, i.e., source IP address and port number, destination IP address and

port number, and the protocol used for the flow. In Fig. 2, each client 104 is illustrated as receiving two flows (client A: downward diagonal pattern and upward diagonal pattern, respectively; client B: horizontal pattern and vertical pattern, respectively). Data packets are de-queued 213, using either a fair or weighted scheduler, and forwarded downstream to the clients 104, with a bitrate which is substantially equal to the estimated bitrate for each client 104. As a consequence of creating an artificial bottleneck 213 for each client 104 at an upstream location, the build-up of queues at potential bottlenecks downstream from ARFM 100, e.g., in the operator network 102 or the RAN 103, is avoided.

Embodiments of the invention simplify deployment of AQM in an operator network since a single node, ARFM 100, is sufficient to manage the entire network downstream from the point of deployment. In contrast to a TCP proxy, ARFM 100 can ensure flow-fairness and flow-isolation for all types of traffic entering the network, in particular both TCP- and UDP-based traffic. Advantageously, short-lived flows may experience an increased responsiveness due to the scheduling on flow level. For example, a Domain Name System (DNS) response does not have to share a queue with a TCP download of bulk data.

The purpose of ARFM 100 is to reduce buffering of data packets 201 in nodes without AQM or with sub-optimal AQM. This is achieved by creating an artificial bottleneck at a desired location in the network in a controlled manner, thereby preventing bottlenecks from arising at locations downstream from the ARFM, such as in RAN 103, or shifting existing bottlenecks from a downstream location in the network to ARFM 100.

ARFM 100 may be deployed in networks where path characteristics may differ considerably between connections to different clients 104, such as mobile terminals or UEs receiving data over a RAN 103 or different RANs having different radio characteristics and/or radio access technologies. A

typical example of such a network is 3GPP-type of network in which different clients have separate scheduling queues.

In Fig. 3, an embodiment 300 of the method of controlling a downstream flow of data packets 201 to one or more clients 104 is illustrated. Method 300 is performed by a network node, such as ARFM 100, and comprises receiving 301 data packets 201 from a packet-switched network 101, and separating 302–304 the received data packets 201 into different session buffers 211, one for each client 104, based on a respective destination IP address of each data packet 201. Method 300 further comprises, for each client 104, separating 305–307 data packets 201 into different queues 212, one for each flow, de-queueing 308–309 data packets 201 in accordance with an estimated bitrate of a data connection between ARFM 100 and client 104, and forwarding 310 de-queued data packets 201 to client 104. Preferably, ARFM 100 is deployed as an ingress node at an entry point to a network, e.g., an entry point of a RAN 103 or a network 102 of an operator, as is illustrated in Fig. 1.

Downstream data packets 201 received from an external packet-switched network 101 are routed via ARFM 100 where they are separated 302–304 into different session buffers 211, one for each client 104, based on their destination IP address. If 302 a corresponding session buffer does not yet exists, e.g., because a new client 104 has accessed RAN 103, a new buffer 211 is created 303 for the session of the new client.

Subsequently, within each session 211, data packets 201 are separated 305–307 into different queues 212 based on the packets' 201 5-tuple (source IP address, destination IP address, source port number, destination port number, protocol number). Data packets 201 with the same 5-tuple are referred to as a flow. If 305 a corresponding flow queue does not yet exists, e.g., because a client 104 has started using a new protocol, a new queue 212 is created 306.

Then, for each session 212 (i.e., for each client 104) an artificial bottleneck 213 is created by de-queueing 308–309 data packets 201 in accordance with an estimated capacity of the data connection between ARFM 100 and the respective client 104. To this end, the bandwidth of each artificial bottleneck 213, i.e., its bitrate, is adjusted so as to create a bottleneck of substantially the same capacity/bandwidth as that of the connection between ARFM 100 and the respective client 104.

The capacity or bandwidth, i.e., the available bitrate between ARFM 100 and a client 104, may be estimated based on measurements of the capacity of the downstream path towards the respective client 104. For instance, the path capacity may be estimated by measuring the rate of TCP (or other transport protocol) acknowledgements correlated to the amount of sent downstream data packets 201.

Additionally, or alternatively, explicit feedback information from a bottleneck in the network, such as a radio base station or any other network node, may also serve as input for path capacity estimation. Such information may, e.g., be received from a BSC, an RNC, an eNodeB, a corresponding RRM node of a 5G network, or any other network node handling RRM tasks. The estimates serve as input to a Proportional-Integral-Derivative (PID) controller (or similar) which generates the artificial bottleneck capacity by adjusting parameters of the queue scheduler accordingly. More specifically, the artificial bottleneck 213 is realized by regulating the de-queuing rate of the individual flows 212. As a result, queues will build up in ARFM 100 if the path between an upstream source 101 and ARFM 100 has a higher capacity than the path from ARFM 100 to the destination, i.e., client 104. Data packets 201 within a session are de-queued using a flow scheduler at a rate matching the measured downstream path capacity. The scheduling can either be fair or weighted, based on policy decided by the operator (e.g., to give precedence to a certain flow).

In order to regulate the sending rate per flow (from the source 101), each flow 212 can individually be managed by AQM. This may be achieved by utilizing a maximum queueing delay. When the maximum queuing delay is reached, the ARFM can either drop the data packet packet to make the sender slow down, or de-queue the data packet earlier than in accordance with the estimated bitrate. The maximum queuing delay (per flow) may be decided based on policy and may be configured by the operator of ARFM 100.

The solution described herein may be implemented in a network node suitable for routing data packets between different networks or network entities, such as a proxy. Embodiments of such a network node are described in the following, with reference to Figs. 4 and 5.

In Fig. 4, an embodiment 400 of network node 100 for controlling a downstream flow of data packets 201 to one or more clients 104 is illustrated. Network node 400 comprises at least one network interface 401, and processing means comprising at least one processor 402 and memory 403, coupled by interconnects. Network interface 401 enables network node 400 to communicate with other network nodes or networks, such as external packet-switched network 101, operator network 102, or RAN 103. In particular, network interface 401 is arranged for receiving data packets 201 from packet-switched network 101. Processor 402 may, e.g., be a general purpose processor. Memory 403 is a computer-readable storage medium, such as a Random Access Memory (RAM), a Flash memory, or any other type of non-transitory computer-readable storage medium. Memory 403 contains computer-executable instructions 404, i.e., a computer program, which, when executed on processor 402 implement the functionality described herein. The computer-executable 402 instructions may also be embodied in a transitory computer-readable medium such as a signal or carrier wave, and may be downloaded over a network. To this end, network node 400 is operative to separate received data packets 201 into different

session buffers 211, one for each client 104, based on a respective destination IP address of each data packet 201, and, for each client 104, separate data packets 201 into different queues 212, one for each flow. Network node 400 is further operative to de-queue data packets 201 in accordance with an estimated bitrate of a data connection between network node 400 and client 104, and forward the de-queued data packets 201, by means of network interface 401, to client 104. Network node 400 may further be operative to perform alternative or additional steps in accordance with embodiments of the invention described herein, and may additionally include further components for providing additional functionality, including any functionality described herein and/or any functionality necessary to support the solution described herein.

In Fig. 5, an alternative embodiment 500 of network node 100 for controlling a downstream flow of data packets 201 to one or more clients 104 is illustrated. Network node 500 comprises at least one network interface 501, a queue module 502, and a scheduler 503. Network interface 501 enables network node 500 to communicate with other network nodes or networks, such as external packet-switched network 101, operator network 102, or RAN 103. In particular, network interface 501 is arranged for receiving data packets 201 from a packet-switched network 101. Queue module 502 is arranged for separating the received data packets 201 into different session buffers 211, one for each client 104, based on a respective destination IP address of each data packet 201. Queue module 502 is further arranged for, for each client 104, separating data packets 201 into different queues 212, one for each flow. Scheduler 503 is arranged for de-queueing data packets 201 in accordance with an estimated bitrate of a data connection between network node 500 and client 104, and forwarding the de-queued data packets 210, by means of network interface 501, to client 104. Network interface 501, queue module 502, and scheduler 503, may further be arranged to perform alternative or additional steps in accordance with

embodiments of the invention described herein, and may additionally include further components for providing additional functionality, including any functionality described herein and/or any functionality necessary to support the solution described herein.

Modules 501–503, as well as any additional modules comprised in network node 500, may be implemented by any kind of electronic circuitry, e.g., any one, or a combination of, analogue electronic circuitry, digital electronic circuitry, and processing means executing a suitable computer program.

The person skilled in the art realizes that the invention by no means is limited to the embodiments described above. On the contrary, many modifications and variations are possible within the scope of the appended claims.

CLAIMS

1. A network node (100; 400; 500) for controlling a downstream flow of data packets (201) to one or more clients (104), the network node comprising:

at least one network interface (401; 501) arranged for receiving the data packets from a packet-switched network (101), and

processing means (402–404; 502, 503) being operative to:

separate the received data packets into different session buffers (211), one for each client, based on a respective destination Internet Protocol, IP, address of each data packet, and

for each client:

separate the data packets into different queues (212), one for each flow,

de-queue (213) the data packets in accordance with an estimated bitrate of a data connection between the network node and the client, and

forward the de-queued data packets, by means of the at least one network interface, to the client.

2. The network node according to claim 1, the processing means being operative to de-queue the data packets at a bitrate which is substantially equal to the estimated bitrate for each client.

3. The network node according to claim 1 or 2, the processing means being operative to de-queue the data packets using one of a fair scheduler and a weighted scheduler.

4. The network node according to any one of claims 1 to 3, the processing means being further operative to, for each client, estimate the bitrate of the data connection between the network node and the client.

5. The network node according to claim 4, the processing means being operative to estimate the bitrate by measuring the rate of acknowledgements pertaining to a transport protocol used by the client, correlated with the amount of data packets sent to the client.

6. The network node according to claim 4, the processing means being operative to estimate the bitrate based on feedback information received from a downstream network node.

7. The network node according to claim 6, wherein the feedback information is received from a Radio Resource Management, RRM, node of a Radio Access Network (103), RAN.

8. The network node according to any one of claims 1 to 7, the processing means being operative to de-queue the data packets in accordance with a maximum queueing delay.

9. The network node according to any one of claims 1 to 8, wherein each of the one or more clients is a mobile terminal receiving the data packets over a Radio Access Network (103), RAN.

10. The network node according to any one of claims 1 to 9, wherein the network node is deployed as an ingress node at an entry point of one of a Radio Access Network (103), RAN, and an operator network (102).

11. A method (300) of controlling a downstream flow of data packets (201) to one or more clients (104), the method being performed by a network node (100; 400) and comprising:

receiving (301) the data packets from a packet-switched network (101),

separating (302–304) the received data packets into different session buffers (211), one for each client, based on a respective destination Internet Protocol, IP, address of each data packet, and

for each client:

separating (305–307) the data packets into different queues (212), one for each flow,

de-queueing (308, 309; 213) the data packets in accordance with an estimated bitrate of a data connection between the network node and the client, and

forwarding (310) the de-queued data packets to the client.

12. The method according to claim 11, wherein the data packets are de-queued at a bitrate which is substantially equal to the estimated bitrate for each client.

13. The method according to claim 11 or 12, wherein the data packets are de-queued using one of a fair scheduler and a weighted scheduler.

14. The method according to any one of claims 11 to 13, further comprising, for each client, estimating the bitrate of the data connection between the network node and the client.

15. The method according to claim 14, wherein the bitrate is estimated by measuring the rate of acknowledgements pertaining to a transport

protocol used by the client, correlated with the amount of data packets sent to the client.

16. The method according to claim 14, wherein the bitrate is estimated based on feedback information received from a downstream network node.

17. The method according to claim 16, wherein the feedback information is received from a Radio Resource Management, RRM, node of a Radio Access Network (103), RAN.

18. The method according to any one of claims 11 to 17, wherein the data packets are de-queued in accordance with a maximum queueing delay.

19. The method according to any one of claims 11 to 18, wherein each of the one or more clients is a mobile terminal receiving the data packets over a Radio Access Network (103), RAN.

20. The method according to any one of claims 11 to 19, wherein the network node is deployed as an ingress node at an entry point of one of a Radio Access Network (103), RAN, and an operator network (102).

21. A computer program (404) comprising computer-executable instructions for causing a device to perform the method according to any one of claims 12 to 20, when the computer-executable instructions are executed on a processing unit (402) comprised in the device.

22. A computer program product comprising a computer-readable storage medium (403), the computer-readable storage medium having the computer program (404) according to claim 21 embodied therein.

Fig. 1

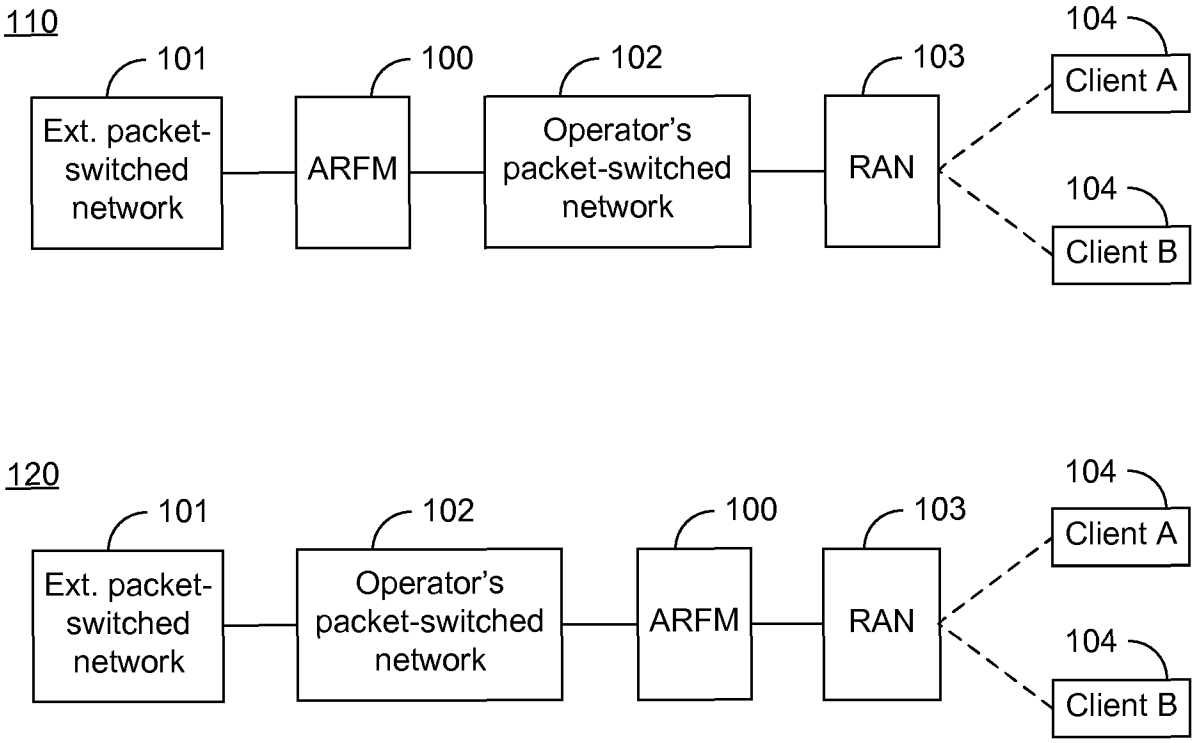
110
101
Ext. packet-switched network
100
ARFM
102
Operator's packet-switched network
103
RAN
104
Client A
104
Client B
120
101
Ext. packet-switched network
102
Operator's packet-switched network
100
ARFM
103
RAN
104
Client A
104
Client B

Fig. 2

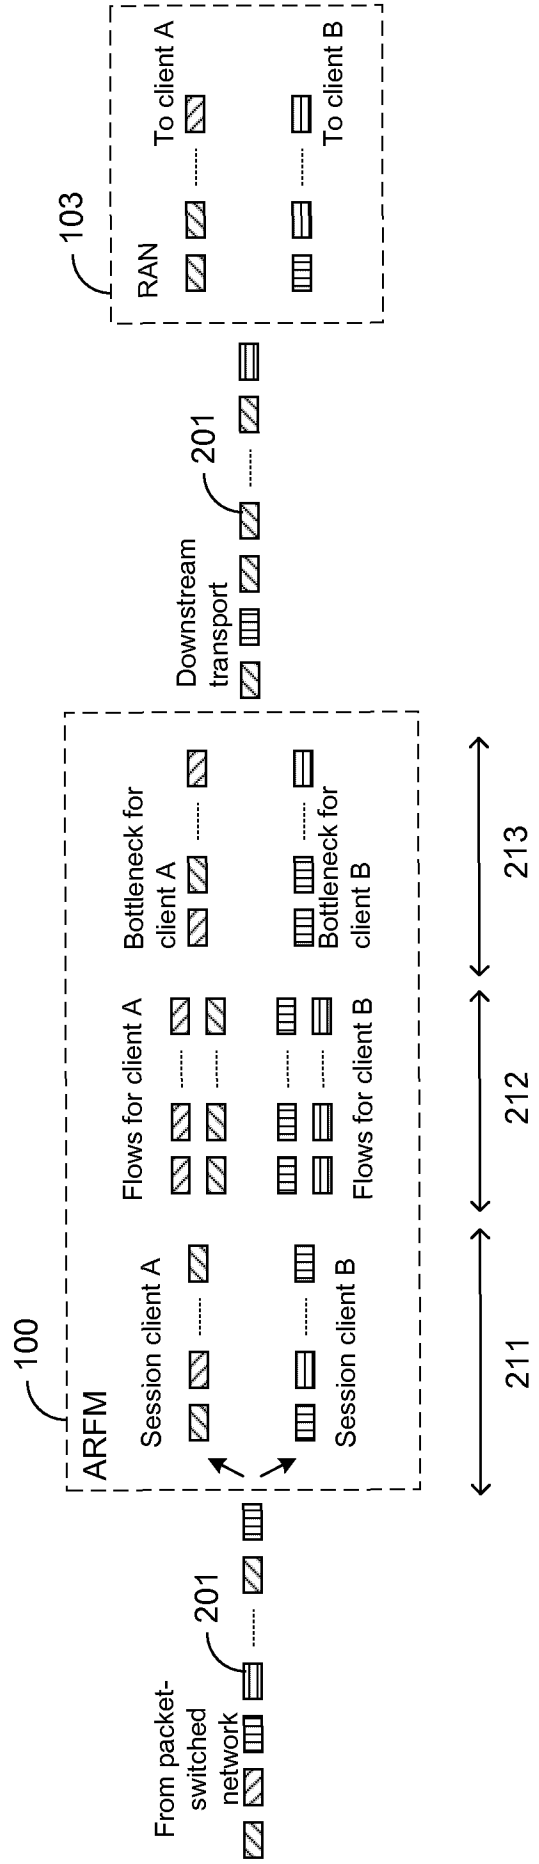
103
RAN
To client A
To client B
201
Downstream transport
100
ARFM
Session client A
Flows for client A
Bottleneck for client A
Session client B
Flows for client B
Bottleneck for client B
211
212
213
201
From packet-switched network

Fig. 3

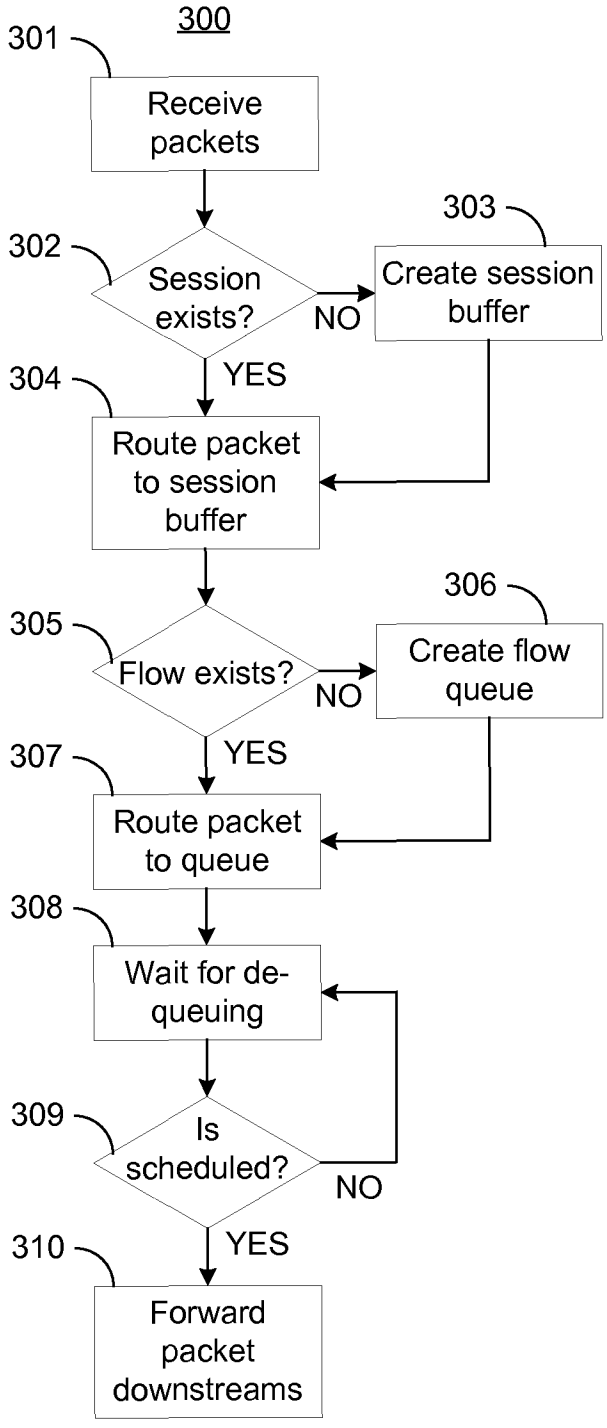
300
301
Receive packets
302
Session exists?
NO
303
Create session buffer
YES
304
Route packet to session buffer
305
Flow exists?
NO
306
Create flow queue
YES
307
Route packet to queue
308
Wait for de-queuing
309
Is scheduled?
NO
YES
310
Forward packet downstreams

Fig. 4

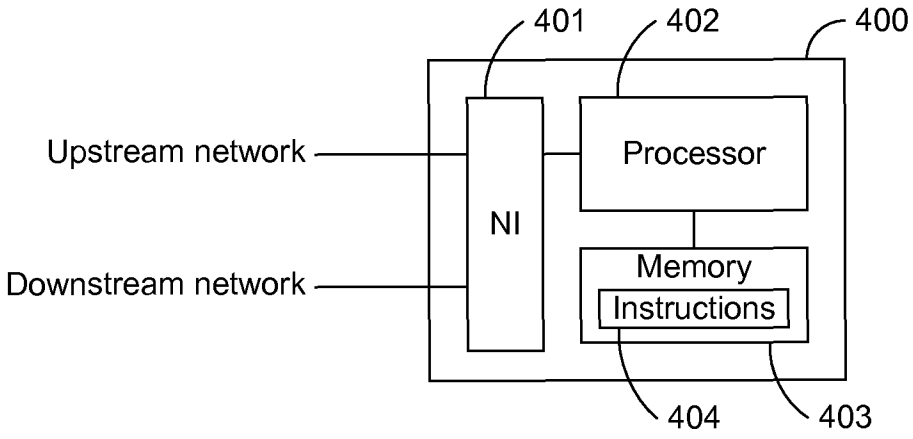
401
402
400
Upstream network
Downstream network
NI
Processor
Memory
Instructions
404
403

Fig. 5

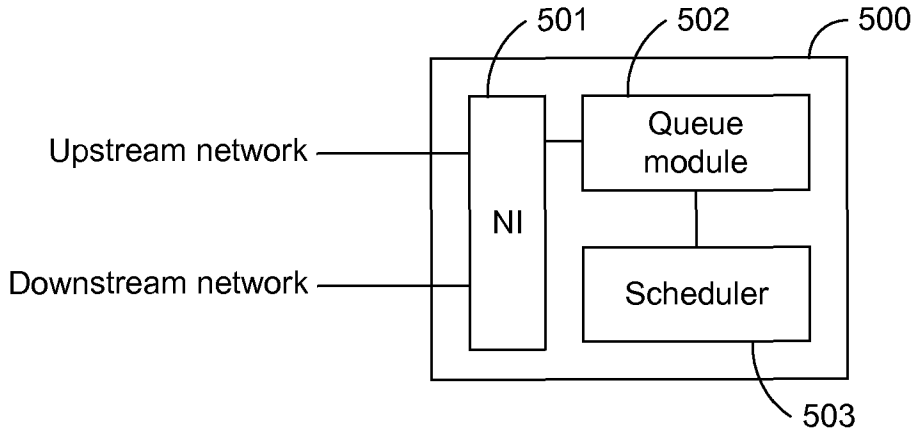
501
502
500
Upstream network
Downstream network
NI
Queue module
Scheduler
503

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2015/079128

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/823 H04L12/863 H04L12/869 H04L29/06
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)
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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/098668 A1 (KAKADIA DEEPAK [US] ET AL) 10 April 2014 (2014-04-10) figures 1A, 1B, 5, 6 paragraph [0011] - paragraph [0024] paragraph [0039] - paragraph [0049] -----	1-22
X	EP 2 764 747 A1 (ERICSSON TELEFON AB L M [SE]) 13 August 2014 (2014-08-13) figures 1, 2, 3, 4, 6 paragraph [0008] - paragraph [0010] paragraph [0014] - paragraph [0036] paragraph [0041] - paragraph [0047] ----- -/--	1-22

[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 June 2016	29/06/2016

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	Gutiérrez García, J

1

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2015/079128

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2012/114328 A1 (CELTRO LTD [IL]; LIOR AYAL [IL]; GORODEISKY SHAHAR [IL]; BRENER-SHEM-T) 30 August 2012 (2012-08-30) figures 1, 2, 3 paragraph [0004] - paragraph [0004] paragraph [0014] - paragraph [0026] paragraph [0042] - paragraph [0103] -----	1-22

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2015/079128

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014098668 A1	10-04-2014	NONE	
EP 2764747 A1	13-08-2014	CN 103959882 A	30-07-2014
		EP 2764747 A1	13-08-2014
		US 2014254381 A1	11-09-2014
		WO 2013050062 A1	11-04-2013
WO 2012114328 A1	30-08-2012	NONE	