Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2868712 A1 2013/10/03

(21) 2 868 712

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2013/03/26
(87) Date publication PCT/PCT Publication Date: 2013/10/03
(85) Entrée phase nationale/National Entry: 2014/09/26
(86) N° demande PCT/PCT Application No.: JP 2013/058756
(87) N° publication PCT/PCT Publication No.: 2013/146770
(30) Priorité/Priority: 2012/03/28 (JP2012-074655)

(51) Cl.Int./Int.Cl. *H04L 12/717* (2013.01)

(71) Demandeur/Applicant:
NEC CORPORATION, JP

(72) Inventeurs/Inventors:
TAKAJO, MAMORU, JP;
TAKASHIMA, MASANORI, JP;
SATO, SHIHOMI, JP

(74) Agent: SMART & BIGGAR

(54) Titre : DISPOSITIF DE COMMUNICATION, DISPOSITIF DE COMMANDE, SYSTEME DE COMMUNICATION, PROCEDE DE COMMUNICATION, PROCEDE DE COMMANDE DE DISPOSITIF DE COMMUNICATION ET PROGRAMME

(54) Title: COMMUNICATION APPARATUS, CONTROL APPARATUS, COMMUNICATION SYSTEM, COMMUNICATION METHOD, METHOD FOR CONTROLLING COMMUNICATION APPARATUS, AND PROGRAM

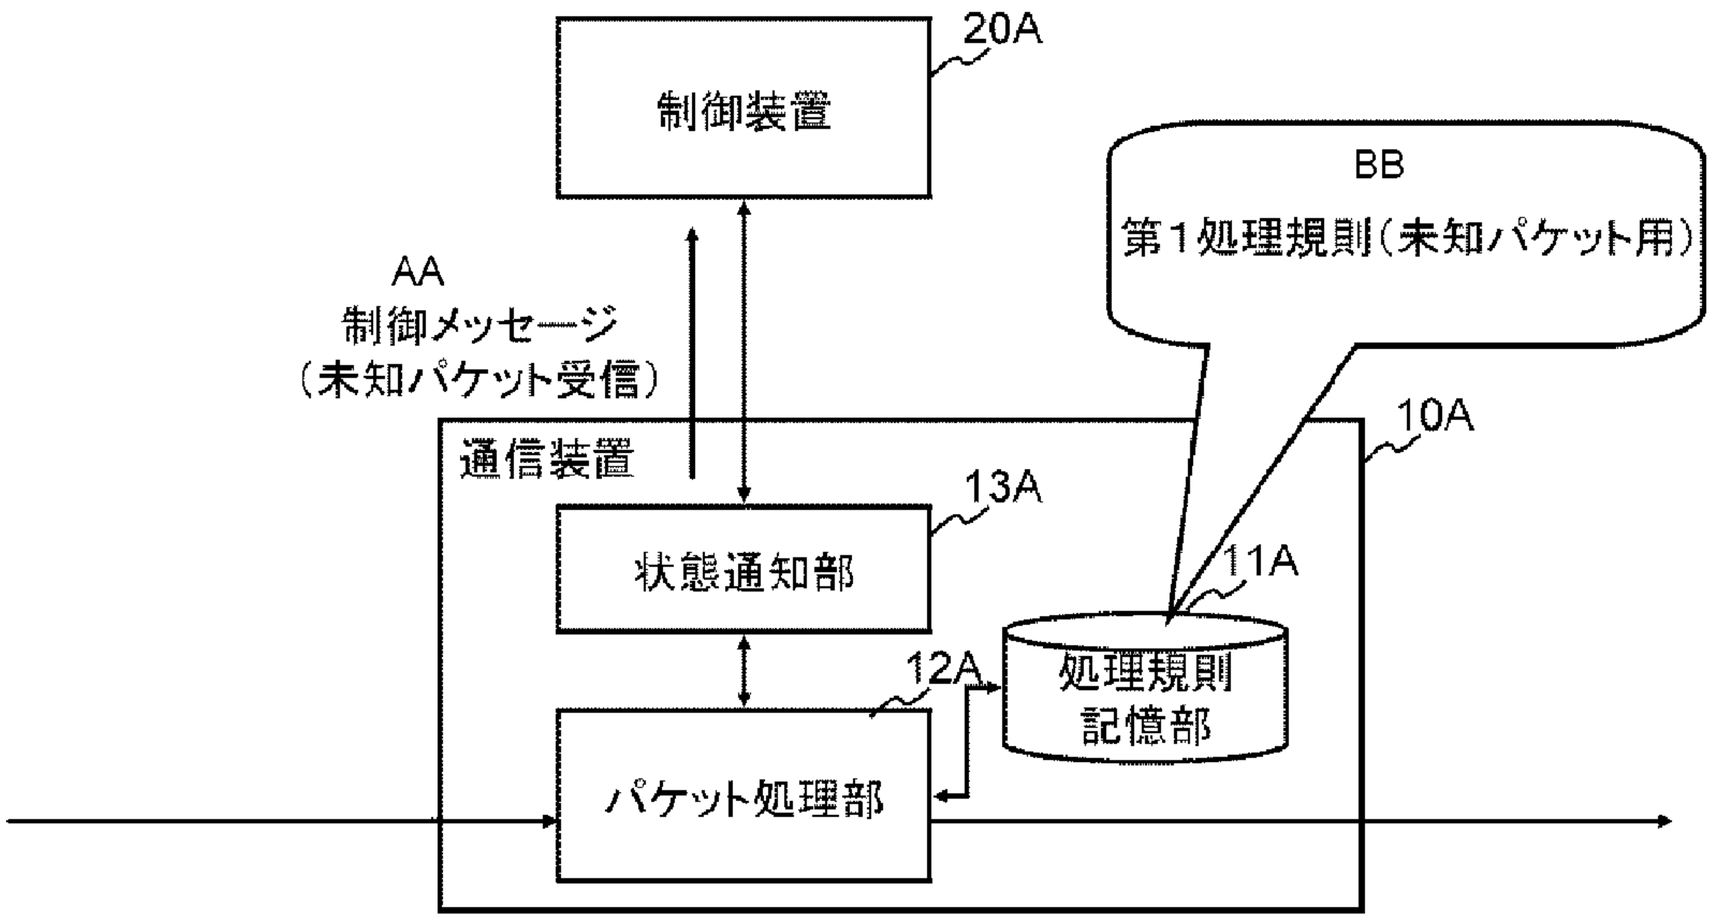

10A Communication device
11A Process-rule storage unit
12A Packet processor
13A State-notification unit
20A Control device
AA Control message (unknown packet reception)
BB First process rule (for unknown packet)

(57) Abrégé/Abstract:
The present invention reduces load on each device in a central control network when a large amount of data has come in, and also inhibits reversal of the order of transmission of the packets. A communication device is provided with: a process-rule storage unit

CA 2868712 A1 2013/10/03

(21) 2 868 712

(13) A1

(57) Abrégé(suite)/Abstract(continued):
capable of storing a first process rule, set in advance, for unknown packets, and a second process rule set by a control device; a packet processor for processing reception packets on the basis of the process rules stored in the process-rule storage unit; and a state-notification unit for transmitting a predetermined control message to the control device when a packet is processed using the first process rule. When the control device receives the predetermined control message, the control device creates the second process rule to be set in the communication device and sets the rule in the communication device, and also discards the packets included in the control message. The communication device continues processing packets using the first process rule until the second process rule is set by the control device.

ABSTRACT

The invention reduces load on respective apparatuses in a centrally controlled network when a large amount of data is inputted, and inhibits reversal of order of transmission of packets. A communication apparatus is provided with: a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus; a packet processor that processes received packets based on processing rules stored in the processing rule storage unit; and a state notification unit that transmits a predetermined control message to the control apparatus when a packet is processed using the first processing rule. On receiving the predetermined control message, the control apparatus creates the second processing rule to be set in the communication apparatus and sets the second processing rule in the communication apparatus, and also discards packets included in the control message. The communication apparatus continues packet processing using the first processing rule until the second processing rule is set by the control apparatus.

WO 2013/146770 A1

集中制御型のネットワークにおいて大量のデータが流入してきた際の各機器の負荷を低減し、また、パケットの送信の順序逆転を抑止する。通信装置は、事前に設定された未知パケット用の第１の処理規則と、制御装置から設定された第２の処理規則とを格納可能な処理規則記憶部と、前記処理規則記憶部に格納された処理規則に基づいて受信パケットを処理するパケット処理部と、前記第１の処理規則を用いてパケットを処理した際に、前記制御装置に対して、所定の制御メッセージを送信する状態通知部とを備える。前記制御装置は、前記所定の制御メッセージを受信すると、前記通信装置に設定する第２の処理規則を作成し、前記通信装置に設定するとともに、前記制御メッセージに含まれるパケットを廃棄する。通信装置は、前記制御装置から前記第２の処理規則が設定されるまで、前記第１の処理規則を用いてパケット処理を継続する。

DESCRIPTION

COMMUNICATION APPARATUS, CONTROL APPARATUS, COMMUNICATION SYSTEM, COMMUNICATION METHOD, METHOD FOR CONTROLLING COMMUNICATION APPARATUS, AND PROGRAM

TECHNICAL FIELD

[0001]

[REFERENCE TO RELATED APPLICATION]

The present application claims priority from Japanese Patent Application No. 2012-074655 (filed on March 28, 2012), the content of which is hereby incorporated in its entirety by reference into this disclosure. The present invention relates to a communication apparatus, a control apparatus, a communication system, a communication method, a method for controlling the communication apparatus, and a program, and in particular relates to a communication apparatus, a control apparatus, a communication system, a method for controlling the communication apparatus, and a program, that perform packet processing in accordance with control information set by the control apparatus.

BACKGROUND ART

[0002]

With the expansion of network bandwidth, services are required that handle large amounts of data such as image distribution and the like. As a component technology thereof, OpenFlow in Non Patent Literature 1 and 2 may be cited. In OpenFlow, switches capable of high-speed forwarding are arranged on a user plane, and by optimally setting path information by an OpenFlow controller, it is possible to adequately respond to high speed forwarding needs.

CITATION LIST

NON-PATENT LITERATURE

[0003]

[NPL 1]

CA 02868712 2014-09-26

Nick McKeown and 7 others, "OpenFlow: Enabling Innovation in Campus Networks", [online], [search conducted February 14, 2012], Internet <URL: http://www.openflow.org/documents/openflow-wp-latest.pdf>

[NPL 2]

"OpenFlow 1.2", [online], [search conducted March 27, 2012], Internet <URL: https://www.opennetworking.org/images/stories/downloads/openflow/openflow-spec-v1.2.pdf>

SUMMARY OF THE INVENTION

TECHNICAL PROBLEM

[0004]

The following analysis is given by the present invention. In a centrally controlled network as represented by OpenFlow in Non Patent Literature 1 and 2, there is a problem in that, until setting of a processing rule (equivalent to a flow entry in OpenFlow) by a control apparatus is completed, it is necessary to perform forwarding (Packet-in, Packet-out messaging in OpenFlow protocol) via a control plane. In this way, in a case of a data burst or a large amount of data being inputted, there are problems of load increase in an entire system including a control apparatus and communication apparatus on an input side, and of packet loss occurrence.

[0005]

As a result of forwarding via a control plane, before and after setting control information by the control apparatus, there is also a problem of order being reversed with regard to packets for which the control apparatus has specified forwarding via the control plane, and packets for which sequential forwarding is started by a newly set processing rule. For example, in FIG. 9, a forwarding instruction (Packet-out; S207) is received via the control plane from a controller, but before that, a Unicast (S206) of a following packet is performed first by a unicast flow entry set in steps S204, S205 (Flow-mod).

[0006]

It is an object of the present invention to provide a

communication apparatus, a control apparatus, a communication system, a communication method, a method for controlling the communication apparatus, and a program, that contribute to reducing load when a relatively large amount of data is inputted, in the abovementioned centrally controlled network, and to inhibiting reversal of order of arrival of the packets described above.

SOLUTION TO PROBLEM

[0007]

According to a first aspect, there is provided a communication apparatus that includes: a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus; a packet processor that processes received packets based on processing rules stored in the processing rule storage unit; and a state notification unit that transmits a predetermined message to the control apparatus when a packet is processed using the first processing rule; wherein the communication apparatus continues packet processing using the first processing rule until the second processing rule is set by the control apparatus.

[0008]

According to a second aspect, there is provided a control apparatus connected to the abovementioned communication apparatus, wherein the control apparatus creates, on receiving the predetermined control message, the second processing rule to be set in the communication apparatus and sets the second processing rule in the communication apparatus, and also discards packets included in the control message.

[0009]

According to a third aspect, there is provided a communication system including the abovementioned communication apparatus and the control apparatus.

[0010]

According to a fourth aspect, there is provided a communication method wherein a communication apparatus including a processing rule storage unit capable of storing a first processing rule for unknown

packets set in advance, and a second processing rule set by a control apparatus, and a packet processor that processes received packets based on processing rules stored in the processing rule storage unit, performs: a step of transmitting a predetermined control message to the control apparatus when a packet is processed using the first processing rule, and a step of continuing packet processing using the first processing rule until the second processing rule is set by the control apparatus. The present method is linked to a specific apparatus known as a communication apparatus that processes packets using processing rules from the control apparatus.

[0011]

According to a fifth aspect, there is provided a control method wherein a control apparatus connected to a communication apparatus that includes a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance and a second processing rule set by the control apparatus, a packet processor that processes received packets based on processing rules stored in the processing rule storage unit, and a state notification unit that transmits a predetermined control message to the control apparatus when a packet is processed using the first processing rule, performs: a step of creating, on receiving the predetermined control message from the communication apparatus, the second processing rule to be set in the communication apparatus, and setting the second processing rule in the communication apparatus, and a step of discarding packets included in the control message. The present method is linked to a specific apparatus known as control apparatus that sets processing rules in response to a request from the communication apparatus.

[0012]

According to a sixth aspect, there is provided a computer program that is executed in each of the communication apparatus and the control apparatus as described above. It is to be noted that the program may be recorded on a computer readable (non-transitory) recording medium. That is, the present invention may be embodied as a computer program product.

CA 02868712 2014-09-26

ADVANTAGEOUS EFFECTS OF INVENTION

[0013]

According to the present invention, in a centrally controlled network it is possible to contribute to reducing load in respective apparatuses when a large amount of data is inputted, and to inhibiting reversal of the order of packet transmission.

BRIEF DESCRIPTION OF DRAWINGS

[0014]

[Fig. 1]

FIG. 1 is a diagram showing a configuration of an exemplary embodiment of the present invention.

[Fig. 2]

FIG. 2 is a diagram for describing operations of the exemplary embodiment of the invention.

[Fig. 3]

FIG. 3 is a diagram showing a configuration of a communication system according to a first exemplary embodiment of the present invention.

[Fig. 4]

FIG. 4 is a diagram showing a configuration of a switch of the first exemplary embodiment of the invention.

[Fig. 5]

FIG. 5 is an example of a flow table held in a communication apparatus of the first exemplary embodiment of the invention.

[Fig. 6]

FIG. 6 is a sequence diagram representing operations of the first exemplary embodiment of the invention.

[Fig. 7]

FIG. 7 is continuation diagram of FIG. 6.

[Fig. 8]

FIG. 8 is a sequence diagram representing operations of a second exemplary embodiment of the invention.

[Fig. 9]

FIG. 9 is a sequence diagram representing operations among switch controllers of Non Patent Literature 1 and 2 showing comparative

examples.

DESCRIPTION OF EMBODIMENTS

[0015]

First, a description is given of an outline of an exemplary embodiment of the present invention, making reference to the drawings. It is to be noted that reference symbols in the drawings attached to this outline are added for convenience to respective elements as examples in order to aid understanding, and are not intended to limit the invention to modes shown in the drawings.

[0016]

In the exemplary embodiment, as shown in FIG. 1, the present invention may be realized by a configuration including at least one communication apparatus 10A and a control apparatus 20A that controls the communication apparatus 10A. More specifically, the communication apparatus 10A is provided with: a processing rule storage unit 11A capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus; a packet processor 12A that processes received packets based on processing rules stored in the processing rule storage unit; and a state notification unit 13A that transmits a predetermined message to the control apparatus 20A when a packet is processed using the first processing rule. The communication apparatus 10A continues packet processing using the first processing rule until the second processing rule is set by the control apparatus 20A.

[0017]

Meanwhile, on receiving the predetermined control message from the communication apparatus 10A, the control apparatus 20A creates the second processing rule to be set in the communication apparatus 10A and sets the second processing rule in the communication apparatus 10A. It is to be noted that the control apparatus 20A discards packets included in the control message without a packet transmission instruction being performed with regard to the communication apparatus 10A as performed in Non Patent Literature 1 and 2 (namely, a packet forwarding operation via a control plane is not performed).

[0018]

By the above configuration, even if a large amount of data is inputted, since processing is performed by the first processing rule, the control apparatus 20A is not overloaded. As shown in FIG. 2, after the second processing rule has been set by the control apparatus 20A, as a known packet, switching to processing by the second processing rule is performed, so that packet transmission order is not reversed.

[0019]

[First Exemplary Embodiment]

Next, a detailed description is given concerning a first exemplary embodiment of the present invention, making reference to the drawings. FIG. 3 is a diagram showing a configuration of a communication system according to the first exemplary embodiment of the present invention. FIG. 3 shows the communication system that including switches 10-1 to 10-3 connected to one another, and a controller 20 that controls the switches 10-1 to 10-3. A server 30 and a client 40 in FIG. 3 can communicate via the switches 10-1 and 10-3, respectively. It is to be noted that solid lines in FIG. 3 indicate network connection relationships, and broken lines represent channels for control between the controller 20 and the switches.

[0020]

FIG. 4 is a diagram showing a configuration of the switch 10-1 of the first exemplary embodiment of the present invention. FIG. 4 shows the controller 20 that controls the switch 10-1 by setting flow entries in the switch 10-1, and the switch 10-1 that retrieves a flow entry having a match field that matches a received packet, from among flow entries set by the controller 20, to perform packet processing.

[0021]

The switch 10-1 is provided with a flow entry storage unit 11 (equivalent to a processing rule storage unit 11A described above) that stores flow entries, a packet processor 12 (equivalent to a packet processor 12A described above) that processes received packets in accordance with a flow entry stored in the flow entry storage unit 11, and a controller communication unit 13 (equivalent to a state

notification unit 13A described above) that, in a case where a received packet in the packet processing unit 12 matches a flow entry for an Unknown Unicast and forwarding is performed, transmits a control message indicating this to the controller 20.

[0022]

FIG. 5 is a diagram showing an example of a flow entry storage table (flow table) stored in the flow entry storage unit 11. A Priority field in the diagram indicates the priority of each flow entry. The example of FIG. 5 shows an example in which 3 flow entries, Priority = 65535, Priority = 60000, and Priority = "a" (note that "a" < 60000) are registered. Among these, the flow entry with Priority = 65535 is the highest priority flow entry.

[0023]

A retrieval condition field in FIG. 5 stores information of match fields to match received packets. In the flow entry in the diagram where "Exact" is specified, matching is performed according to Exact Matching with a received packet header or the like. For a flow entry where dl_dest = X.X.X.X is specified, a data link layer address of a received packet, that is, where a destination MAC address is X.X.X.X, is matched. In a case of a flow entry where In_port = 0/bb is specified, there is a match with a packet inputted from port 0/bb.

[0024]

An action field in FIG. 5 shows processing content to be applied to a received packet where there is a match with a condition shown in the retrieval condition field. For example, for a packet matching a flow entry with Priority = 65535, processing to forward from port 0/cc is performed (that is, forwarding to switch 10-2). For example, for a packet matching a flow entry with Priority = 60000, processing to forward from port 0/dd is performed (that is, forwarding to switch 10-3). For example, for a packet matching a flow entry with Priority = "a", flooding processing to forward from port 0/cc and port 0/dd (see FLOOD in "4.5 Reserved ports" of Non Patent Literature 2) and transmission of a Packet-in message (forwarding done) giving notification of new packet reception to the controller 20 (that is,

forwarding to switches 10-2, 10-3) are performed. The Packet-in message (forwarding done) differs from a Packet-in message described in Non Patent Literature 1 and 2, and is a control message to which is attached an identifier whereby at least the controller 20 gives notification that forwarding of the received packet has been done, namely, that an output instruction for a received packet in a control plane is unnecessary.

[0025]

In the following description, the flow entry with Priority = 65535 and the flow entry with Priority = 60000 are flow entries (second processing rule) set by the controller 20. The flow entry with Priority = "a" is a flow entry for an Unknown Unicast (a flow entry for an unknown packet) that is set in advance. Therefore, on receiving a packet, the switch 10-1 performs checking of the received packet with flow entries in the priority order of the flow table of FIG. 5, and where there is a match with the flow entry with Priority = 65535 or the flow entry with Priority = 60000, processing determined in the action field of the respective flow entries is performed. On the other hand, for a packet that matches neither the flow entry with Priority = 65535 nor the flow entry with Priority = 60000, flooding processing in accordance with a flow entry for an Unknown Unicast (flow entry for an unknown packet) and transmission of a Packet-in (forwarding done) to the controller 20 are performed.

[0026]

This type of switch 10-1 can be realized by flooding processing when there is a match, with the abovementioned low priority, in an OpenFlow switch of Non Patent Literature 1 or 2, and by setting a flow entry for an Unknown Unicast (flow entry for an unknown packet) that performs transmission of a Packet-in message (forwarding done) to the controller 20. It is to be noted that since the switches 10-2 and 10-3 have a configuration (content (target flow, executed action and the like) of the flow entry storage unit differs as appropriate) similar to the abovementioned switch 10-1, a description is omitted.

[0027]

With regard to the controller 20 of FIG. 3 and FIG. 4, in the OpenFlow controller of Non Patent Literature 1 and 2, on receiving a Packet-in message (forwarding done) a path calculation or flow entry setting operation is performed, but implementation is possible by adding a function to discard packets included in the Packet-in message (forwarding done), without giving a packet output instruction to a switch (in the following, where a switch is not particularly distinguished, "switch 10" is indicated).

[0028]

It is to be noted that functions of the switch 10 and the controller 20 shown in FIG. 3 and FIG. 4 may also be realized in a computer installed in these apparatuses, using hardware thereof, by a computer program that executes the respective processing described above.

[0029]

Next, a detailed description is given concerning operations of the present exemplary embodiment, making reference to the drawings. The following description cites an example in which data is transmitted to the client 40 from the server 30 in FIG. 3. In an initial state, only a flow entry for an Unknown Unicast with Priority = "a" (note that a value is set for "a" that is smaller than the Priority of flow entries set by the controller 20) is set in the flow table of FIG. 5.

[0030]

FIG. 6 is a sequence diagram representing operations of the first exemplary embodiment of the invention. As shown in FIG. 6, first, when an initial packet is transmitted from the server 30 to the client 40 (S001 in FIG. 6), the packet is received at port 0/bb of the switch 10-1.

[0031]

The switch 10-1 retrieves a flow entry having a match condition (retrieve condition) that matches a packet received from the server 30. At this point in time, since a flow entry with Priority = 65535 and a flow entry with Priority = 60000 are not set, a match occurs with a flow entry with Priority = "a" (S002 in FIG. 6). In this case, the switch 10-1 follows the content of the action field of the flow entry (flow entry for an Unknown Unicast) with Priority = "a", to perform flooding

processing (S003 in FIG. 6) and Packet-in (forwarding done) transmission (S004 in FIG. 6).

[0032]

It is to be noted that, while omitted in FIG. 6, processing similar to that of switch 10-2 is performed, and a packet transmitted from the server 30 is forwarded to the client 40 via the switch 10-2.

[0033]

Meanwhile, the controller 20 that receives the control message Packet-in (forwarding done) recognizes that this is a Packet-in (forwarding done), not a normal Packet-in, and starts path calculation and flow entry creation, but a packet output instruction (Packet-out) using a packet included in the Packet-in (forwarding done) or packet information is not performed (S005 in FIG. 6).

[0034]

Below, on receiving a packet from the server 30, the switch 10-1 follows the content of the action field of the flow entry (flow entry for an Unknown Unicast) with Priority = "a", to perform flooding processing (S003 in FIG. 6) and Packet-in (forwarding done) transmission (S004 in FIG. 6).

[0035]

Thereafter, as shown in FIG. 7, when the controller 20 creates a Unicast flow entry (UC flow entry) (S103 in FIG. 7), the controller 20 sets the created Unicast flow entry (UC flow entry) for switch 10-1 and switch 10-2 (S104 in FIG. 7; flow-mod).

[0036]

Setting of the Unicast flow entry (UC flow entry) is completed, and thereafter, on receiving a packet from the server 30, the switches 10-1 and 10-2 perform Unicast forwarding in accordance with the Unicast flow entry (UC flow entry) (S105 in FIG. 7).

[0037]

As described above, according to the present exemplary embodiment, even if packets exceeding the processing capability of the switch 10 and the controller 20 are generated, packet order reversal and packet loss occurrence are effectively inhibited. A reason for this is

that the configuration is such that flooding and Packet-in (forwarding done) transmission to the controller are performed by the flow entry for an Unknown Unicast, a packet transmission instruction is given via the control plane by the controller, and packet output processing accompanying this is inhibited.

[0038]

[Second Exemplary Embodiment]

Next, a description is given concerning a second exemplary embodiment of the present invention in which a function is added to the switch 10 described above, and Packet-in (forwarding done) transmission itself is inhibited. Since the present exemplary embodiment is realized by a configuration similar to the first exemplary embodiment described above, the description below is centered on points of difference in operations thereof.

[0039]

FIG. 8 is a sequence diagram representing operations of the second exemplary embodiment of the invention. Basic operations of the switch 10 and controller 20 are similar to the first exemplary embodiment, but as shown on the left edge of FIG. 8, after the switch 10 transmits a first control message Packet-in (forwarding done), for a predetermined time, second and subsequent Packet-in (forwarding done) transmissions are inhibited. It is to be noted that the predetermined time can be calculated from system scale, assumed flow type, and the like.

[0040]

Specifically, when the first packet is received, the switch 10-1 follows the content of the action field of a flow entry (flow entry for an Unknown Unicast) with Priority = “a”, to perform flooding processing (S003 in FIG. 8) and Packet-in (forwarding done) transmission (S004 in FIG. 8). However, thereafter, until a predetermined time elapses, if a subsequent packet is received, flooding processing (S003 in FIG. 8) is performed but Packet-in (forwarding done) transmission (S004 in FIG. 8) is stopped.

[0041]

When the predetermined time elapses, the switch 10 performs flooding processing (S003 in FIG. 8) and Packet-in (forwarding done) transmission (S004 in FIG. 8). Thereafter, until the predetermined time elapses, even if data forwarding is performed, the Packet-in (forwarding done) transmission is stopped.

[0042]

As described above, according to the present exemplary embodiment it is possible to inhibit load increase on the switch 10 and the controller 20 even more than in the first exemplary embodiment. According to the present exemplary embodiment, since adjustment of Packet-in (forwarding done) transmission interval is enabled, it is possible to economize on resources (memory, CPU and the like).

[0043]

In the abovementioned exemplary embodiment, a description was given where, after transmission of a first control message Packet-in (forwarding done), for a predetermined time, second and subsequent Packet-in (forwarding done) transmissions are inhibited, but instead of a predetermined time, the number of times a target packet for control message Packet-in (forwarding done) is received, or data amount, may be used. In this case, after transmission of the control message Packet-in (forwarding done), the switch 10 inhibits the control message Packet-in (forwarding done), until the number of times a packet is received or the data amount exceed a predetermined threshold. Thereafter when the number of times a packet is received or data amount exceed the predetermined threshold, the switch 10 resets values thereof, and transmits the control message Packet-in (forwarding done). It is to be noted that the number of times a packet is received or the data amount can be implemented by providing a flow statistical information field (Counter), similar to flow entries described in Non Patent Literature 1 and 2.

[0044]

A description of exemplary embodiments of the present invention have been given above, but the invention is not limited to the abovementioned exemplary embodiments, and further modifications,

substitutions and adjustments may be added within a scope that does not depart from fundamental technical concepts of the invention. For example, although the abovementioned exemplary embodiments are not particularly limited, in view of the advantages of the present invention as described above, besides utilization in data centers, the invention may be preferably applied to sequencing, image distribution service or TV telephony services in which packet loss effects are large.

[0045]

In the abovementioned exemplary embodiments, a description was given in which flooding processing is performed, but it is also possible to calculate end-to-end paths where communication is expected to occur giving consideration to network topology, and to set flow entries so as to perform packet forwarding following these paths.

[0046]

Although omitted in the abovementioned exemplary embodiments, it is also possible to restrict whether or not to allow flooding processing in accordance with security level (authentication state or the like) of a connection destination, or flooding target region (limitation of UDP (User Datagram Protocol) packets and the like). These can be realized by rewriting matching conditions (retrieval conditions) of flow entries for an Unknown Unicast (flow entries for unknown packets), by the controller 20. In this way, it is possible to switch control in accordance with connection destination and state thereof.

[0047]

In the first exemplary embodiment and the second exemplary embodiment described above, a description was given in which the controller 20 controls the switch 10, but a control target can also be a communication apparatus such as a mobile telephone terminal, a smart phone, a tablet terminal, a personal computer, a game console, a mobile router or the like, that hold flow entries as described above, and in accordance with content thereof, process received packets or packets from an installed application. In this way, for example, it is possible to inhibit load increase due to packets from a packet processor within a apparatus, or packet order reversal.

[0048]

The first exemplary embodiment and the second exemplary embodiment described above are also preferably operated giving consideration to traffic volume and time slots, which are statistically comprehended, rather than having continuous operation. For example, it is possible to perform control allowing packet forwarding in a control plane as in Non Patent Literature 1 and 2, during nighttime, holidays or the like, when network monitoring level drops, and to perform operations as in the first exemplary embodiment and the second exemplary embodiment, in time slots where large amounts of data communication occur or in states that enable strengthening of monitoring commensurate with large amounts of data communication.

[0049]

Finally, preferred modes of the present invention are summarized.

[Mode 1]

(Refer to the communication apparatus according to the first aspect.)

[Mode 2]

The communication apparatus according to Mode 1, wherein the predetermined message is a control message that requests setting of a second processing rule for packet(s) processed using the first processing rule, with respect to the control apparatus.

[Mode 3]

The communication apparatus according to Mode 1 or Mode 2, wherein the state notification unit transmits, in a case of receiving a packet that matches neither the first nor the second processing rule, a second control message requesting setting of the second processing rule and an instruction to forward the packet, with respect to the control apparatus.

[Mode 4]

The communication apparatus according to any one of Modes 1 to 3, wherein the state notification unit inhibits, after transmission of the predetermined message, transmission of the predetermined control message to the control apparatus, even if a packet is processed using the first processing rule, until a predetermined condition is established.

CA 02868712 2014-09-26

[Mode 5]

The communication apparatus according to Mode 4, wherein the predetermined condition is any of: elapse of a predetermined time, receipt of more than a predetermined number of packets, and receipt of a data amount greater than a predetermined amount.

[Mode 6]

The communication apparatus according to any one of Modes 1 to 5, wherein the first processing rule is a processing rule for performing a broadcast or a multicast, by predetermined path(s).

[Mode 7]

(Refer to the control apparatus according to the second aspect.)

[Mode 8]

(Refer to the communication system according to the third aspect.)

[Mode 9]

(Refer to the communication method according to the fourth aspect.)

[Mode 10]

(Refer to the control method of the communication apparatus according to the fifth aspect.)

[Mode 11]

A program that executes on a computer installed in a communication apparatus provided with a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus, and a packet processor that processes received packets based on processing rules stored in the processing rule storage unit, the program including a process of transmitting a predetermined control message to the control apparatus when a packet is processed using the first processing rule, and a process of continuing packet processing using the first processing rule until the second processing rule is set by the control apparatus.

[Mode 12]

A program that executes on a computer installed in a control apparatus connected to a communication apparatus provided with a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control

apparatus, a packet processor that processes received packets based on processing rules stored in the processing rule storage unit, and a state notification unit that transmits a predetermined control message to the control apparatus when a packet is processed using the first processing rule, a process of creating, on receiving the predetermined control message from the communication apparatus, the second processing rule to be set in the communication apparatus, and setting the second processing rule in the communication apparatus, and a process of discarding packets included in the control message.

It is to be noted that that the abovementioned Modes 7 to 12 may be extended to Modes 2 to 6, similar to Mode 1.

[0050]

It is to be noted that the various disclosures of the abovementioned Non Patent Literature are incorporated herein by reference thereto. Modifications and adjustments of exemplary embodiments and examples may be made within the bounds of the entire disclosure (including the scope of the claims) of the present invention, and also based on fundamental technological concepts thereof. Furthermore, various combinations and selections of various disclosed elements (including respective elements of the respective claims, respective elements of the respective exemplary embodiments and examples, respective elements of the respective drawings, and the like) are possible within the scope of the claims of the present invention. That is, the present invention clearly includes every type of transformation and modification that a person skilled in the art can realize according to the entire disclosure including the scope of the claims and to technological concepts thereof.

REFERENCE SIGNS LIST

[0051]

10-1 to 10-3 switch

10A communication apparatus

11 flow entry storage unit

11A processing rule storage unit

12, 12A packet processor

CA 02868712 2014-09-26

13 controller communication unit

13A state notification unit

20 controller

20A control apparatus

30 server

40 client

CLAIMS:

1. A communication apparatus comprising:

a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus;

a packet processor that processes received packets based on processing rules stored in said processing rule storage unit; and

a state notification unit that transmits a predetermined control message to said control apparatus when a packet is processed using said first processing rule; wherein

said communication apparatus continues packet processing using said first processing rule until said second processing rule is set by said control apparatus.

2. The communication apparatus according to claim 1, wherein said predetermined message is a control message that requests setting of a second processing rule for packet(s) processed using said first processing rule, with respect to said control apparatus.

3. The communication apparatus according to claim 1 or 2, wherein said state notification unit transmits, in a case of receiving a packet that matches neither said first nor said second processing rule, a second control message requesting setting of said second processing rule and an instruction to forward said packet, with respect to said control apparatus.

4. The communication apparatus according to any one of claims 1 to 3, wherein said state notification unit inhibits, after transmission of said predetermined message, transmission of said predetermined control message to said control apparatus, even if a packet is processed using said first processing rule, until a predetermined condition is established.

5. The communication apparatus according to claim 4, wherein said predetermined condition is any of: elapse of a predetermined time, receipt of more than a predetermined number of packets, and receipt of a data amount greater than a predetermined amount.

6. The communication apparatus according to any one of claims 1 to 5, wherein said first processing rule is a processing rule for performing

a broadcast or a multicast, by predetermined path(s).

7. A control apparatus, connected to a communication apparatus that includes:

a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus,

a packet processor that processes received packets based on processing rules stored in said processing rule storage unit, and

a state notification unit that transmits a predetermined control message to said control apparatus when a packet is processed using said first processing rule; wherein

said control apparatus creates, on receiving said predetermined control message, said second processing rule to be set in said communication apparatus and sets said second processing rule in said communication apparatus, and also discards packets included in said control message.

8. A communication system comprising: a communication apparatus including

a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus,

a packet processor that processes received packets based on processing rules stored in said processing rule storage unit, and

a state notification unit that transmits a predetermined control message to said control apparatus, when a packet is processed using said first processing rule, wherein

said communication apparatus continues packet processing using said first processing rule until said second processing rule is set by said control apparatus; and

a control apparatus, connected to said communication apparatus, that creates, on receiving said predetermined control message, said second processing rule to be set in said communication apparatus and sets said rule in said communication apparatus, and also discards packets included in said control message.

9. A communication method wherein a communication apparatus that includes a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance, and a second processing rule set by a control apparatus, and a packet processor that processes received packets based on processing rules stored in said processing rule storage unit, performs:

a step of transmitting a predetermined control message to said control apparatus when a packet is processed using said first processing rule, and

a step of continuing packet processing using said first processing rule until said second processing rule is set by said control apparatus.

10. A control method wherein a control apparatus, connected to a communication apparatus that includes a processing rule storage unit capable of storing a first processing rule for unknown packets set in advance and a second processing rule set by a control apparatus, a packet processor that processes received packets based on processing rules stored in said processing rule storage unit, and a state notification unit that transmits a predetermined control message to said control apparatus when a packet is processed using said first processing rule, performs:

a step of creating, on receiving said predetermined control message from said communication apparatus, said second processing rule to be set in said communication apparatus, and setting said second processing rule in said communication apparatus, and

a step of discarding packets included in said control message.

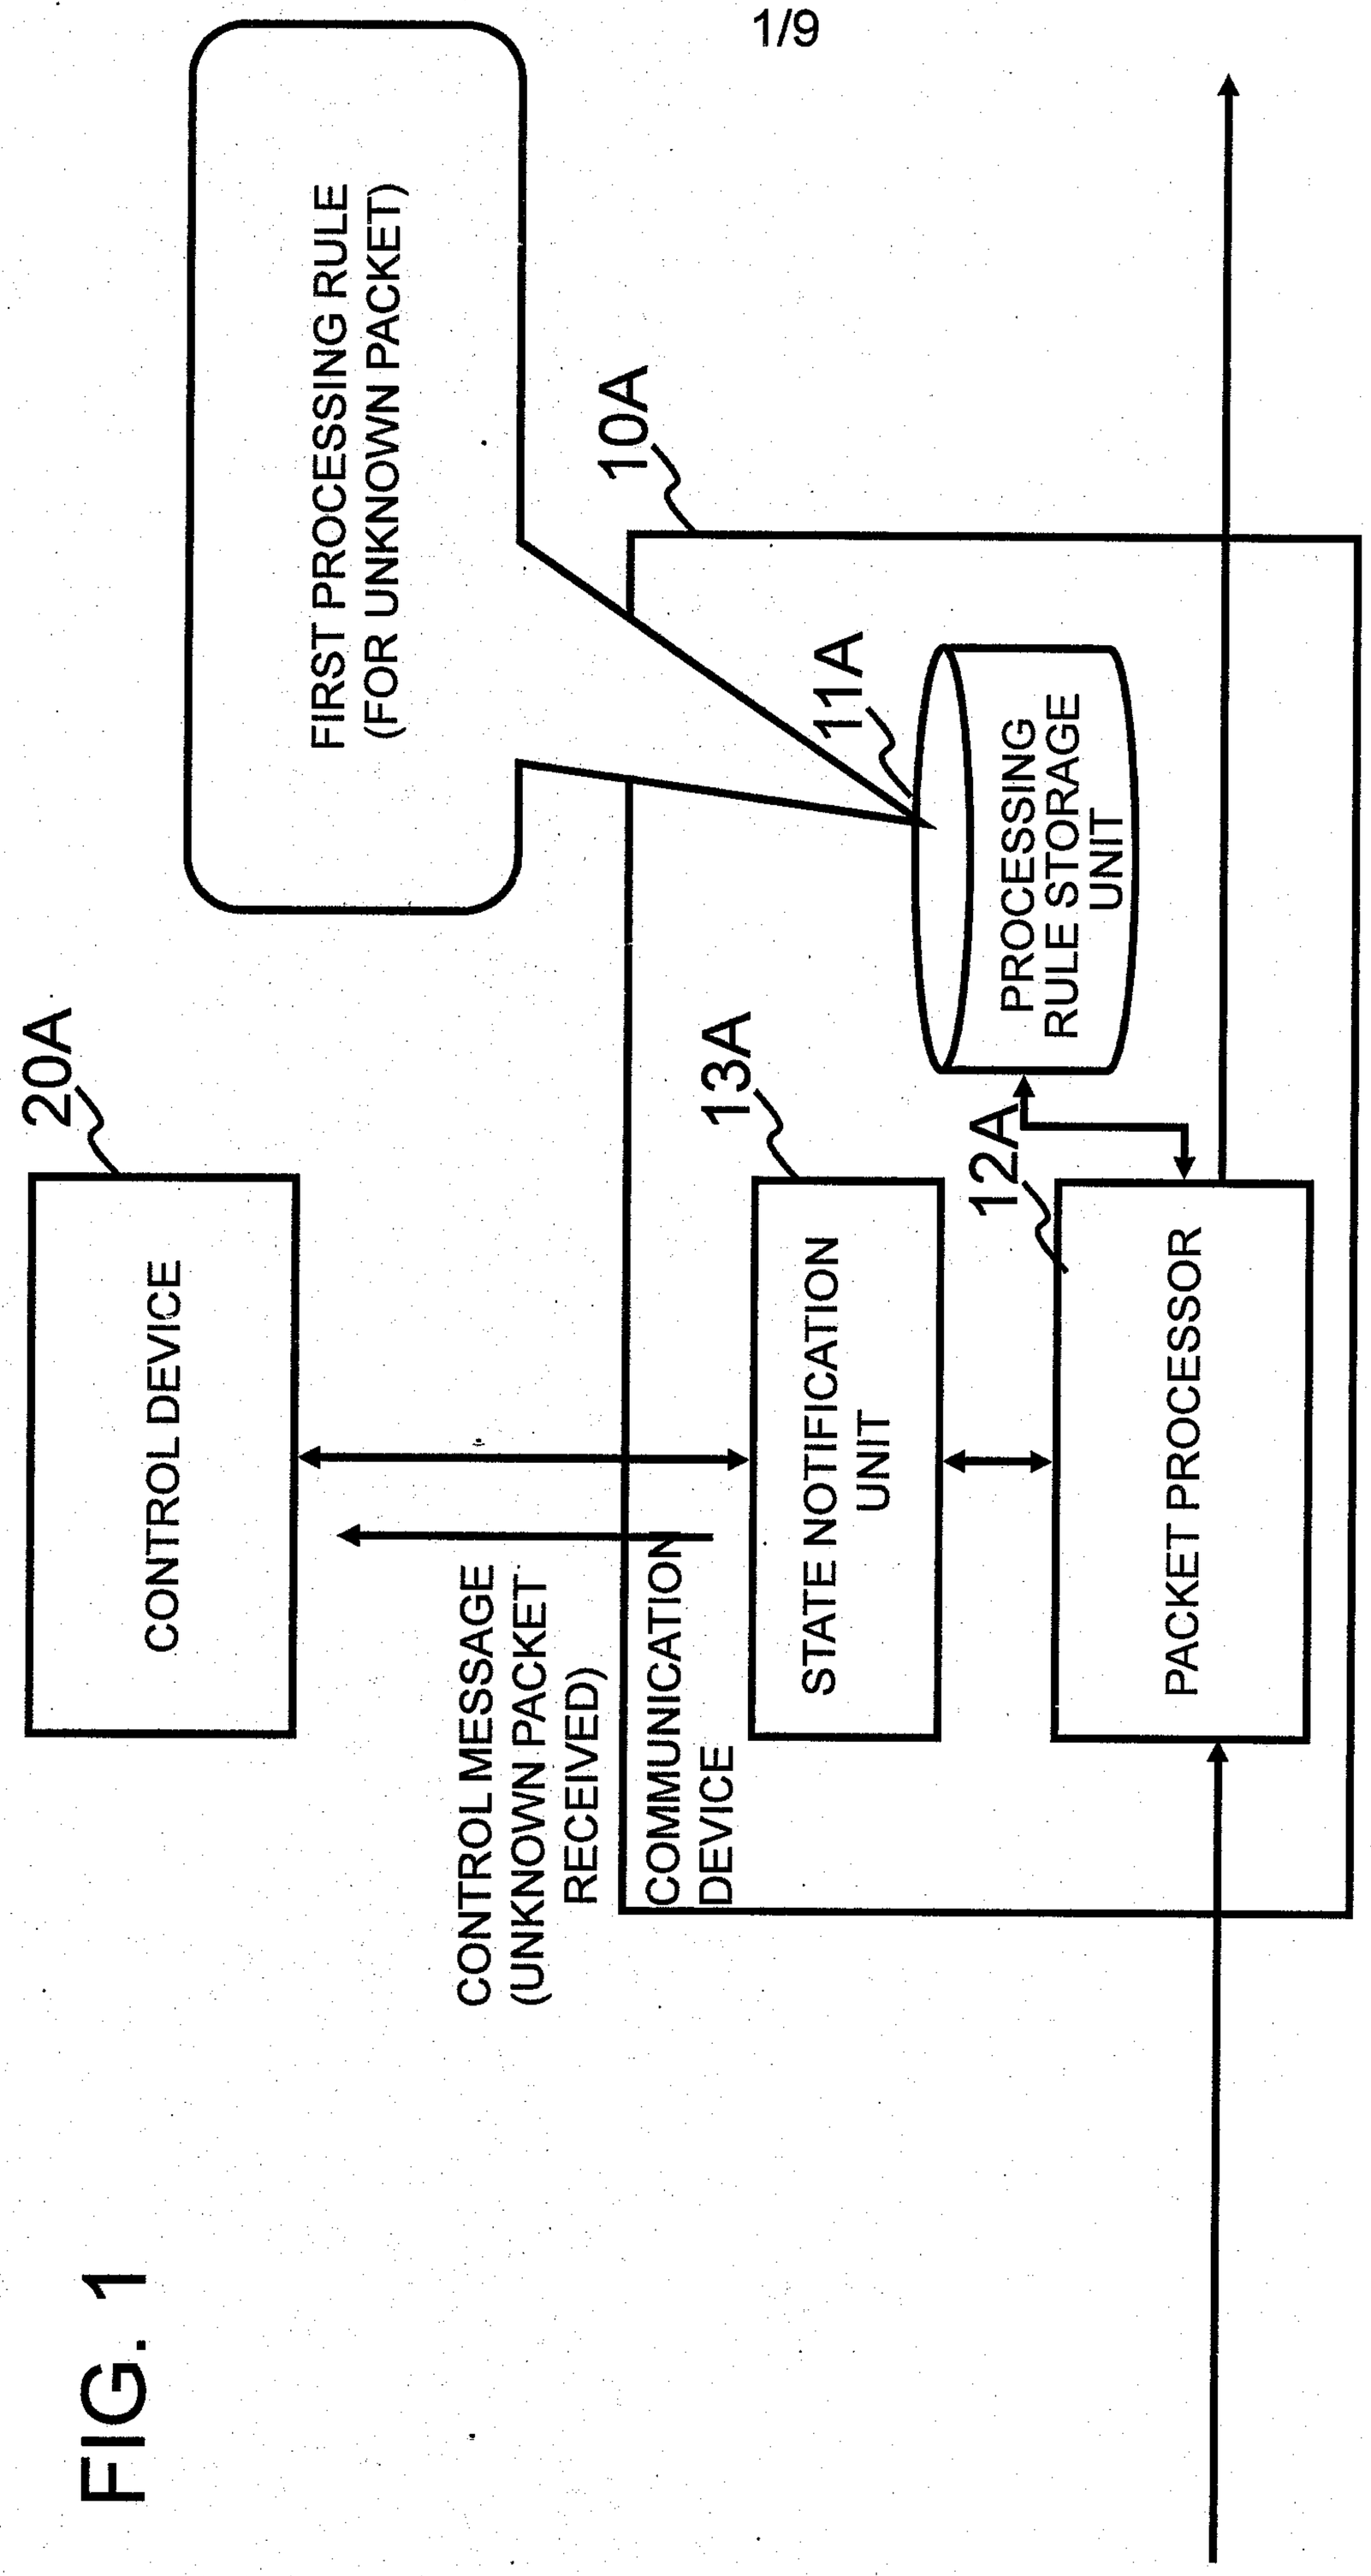
FIRST PACKET PROCESSING RULE
(FOR UNKNOWN PACKET)
10A
11A
PROCESSING RULE STORAGE UNIT
20A
13A
12A
CONTROL DEVICE
CONTROL MESSAGE (UNKNOWN PACKET RECEIVED)
COMMUNICATION DEVICE
STATE NOTIFICATION UNIT
PACKET PROCESSOR

FIG. 1

FIG. 2

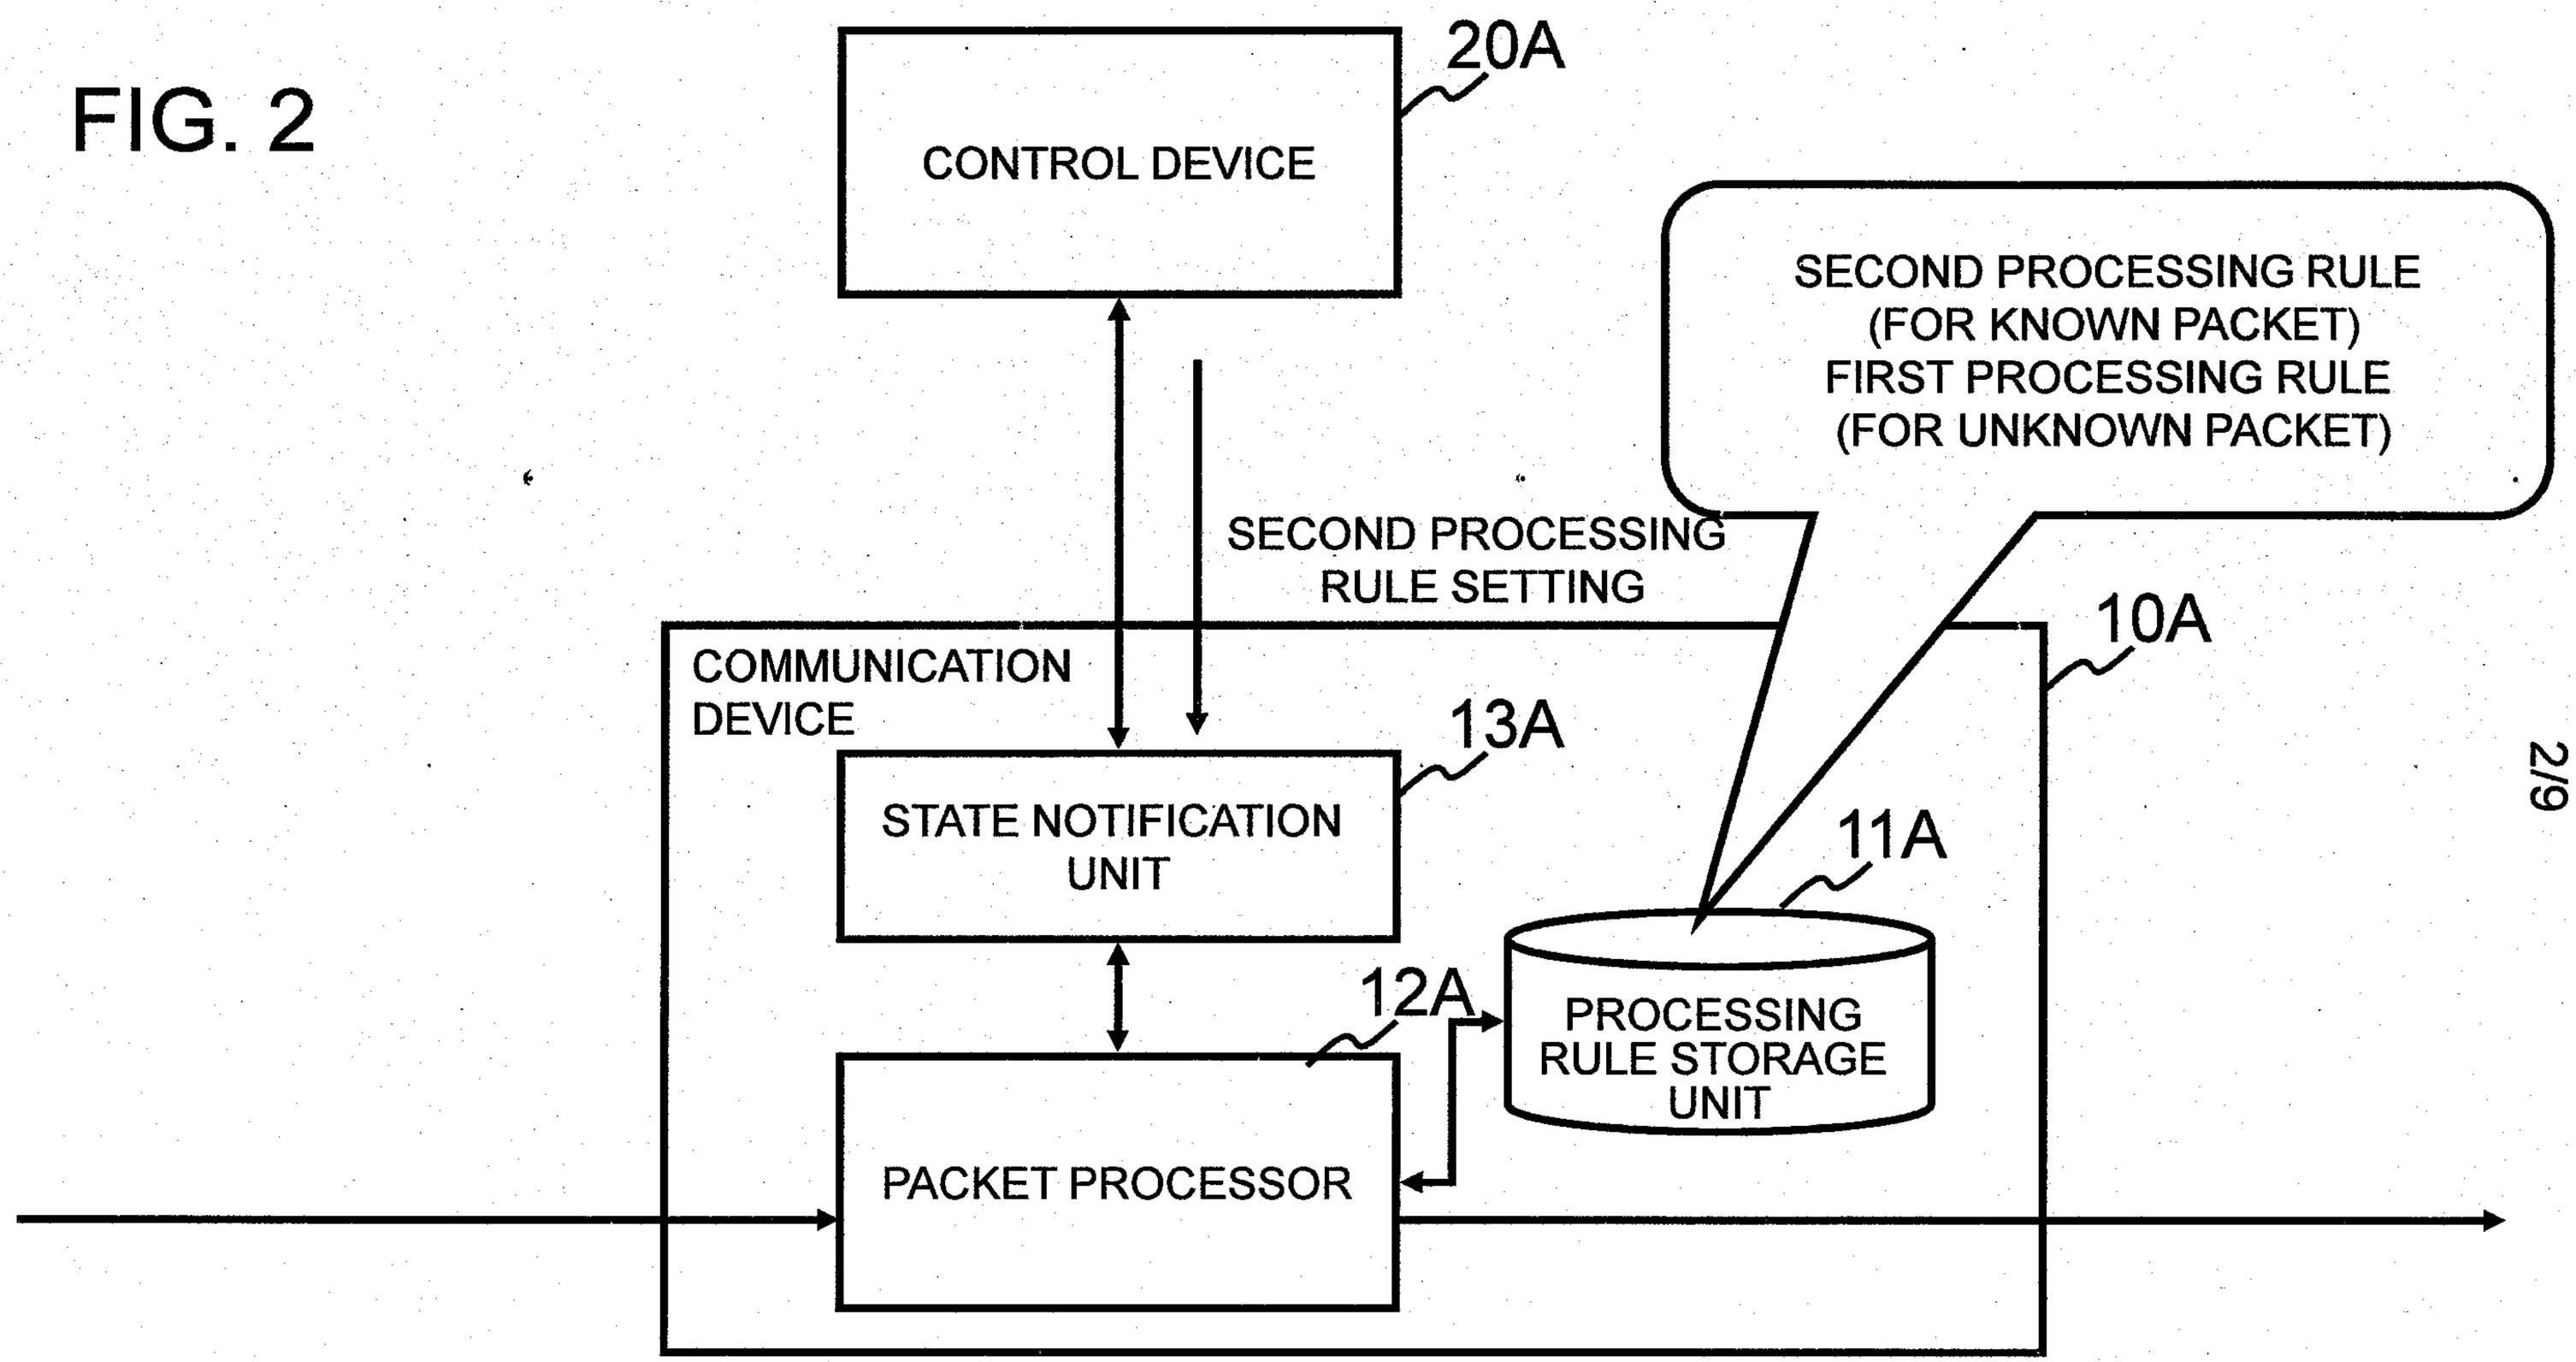

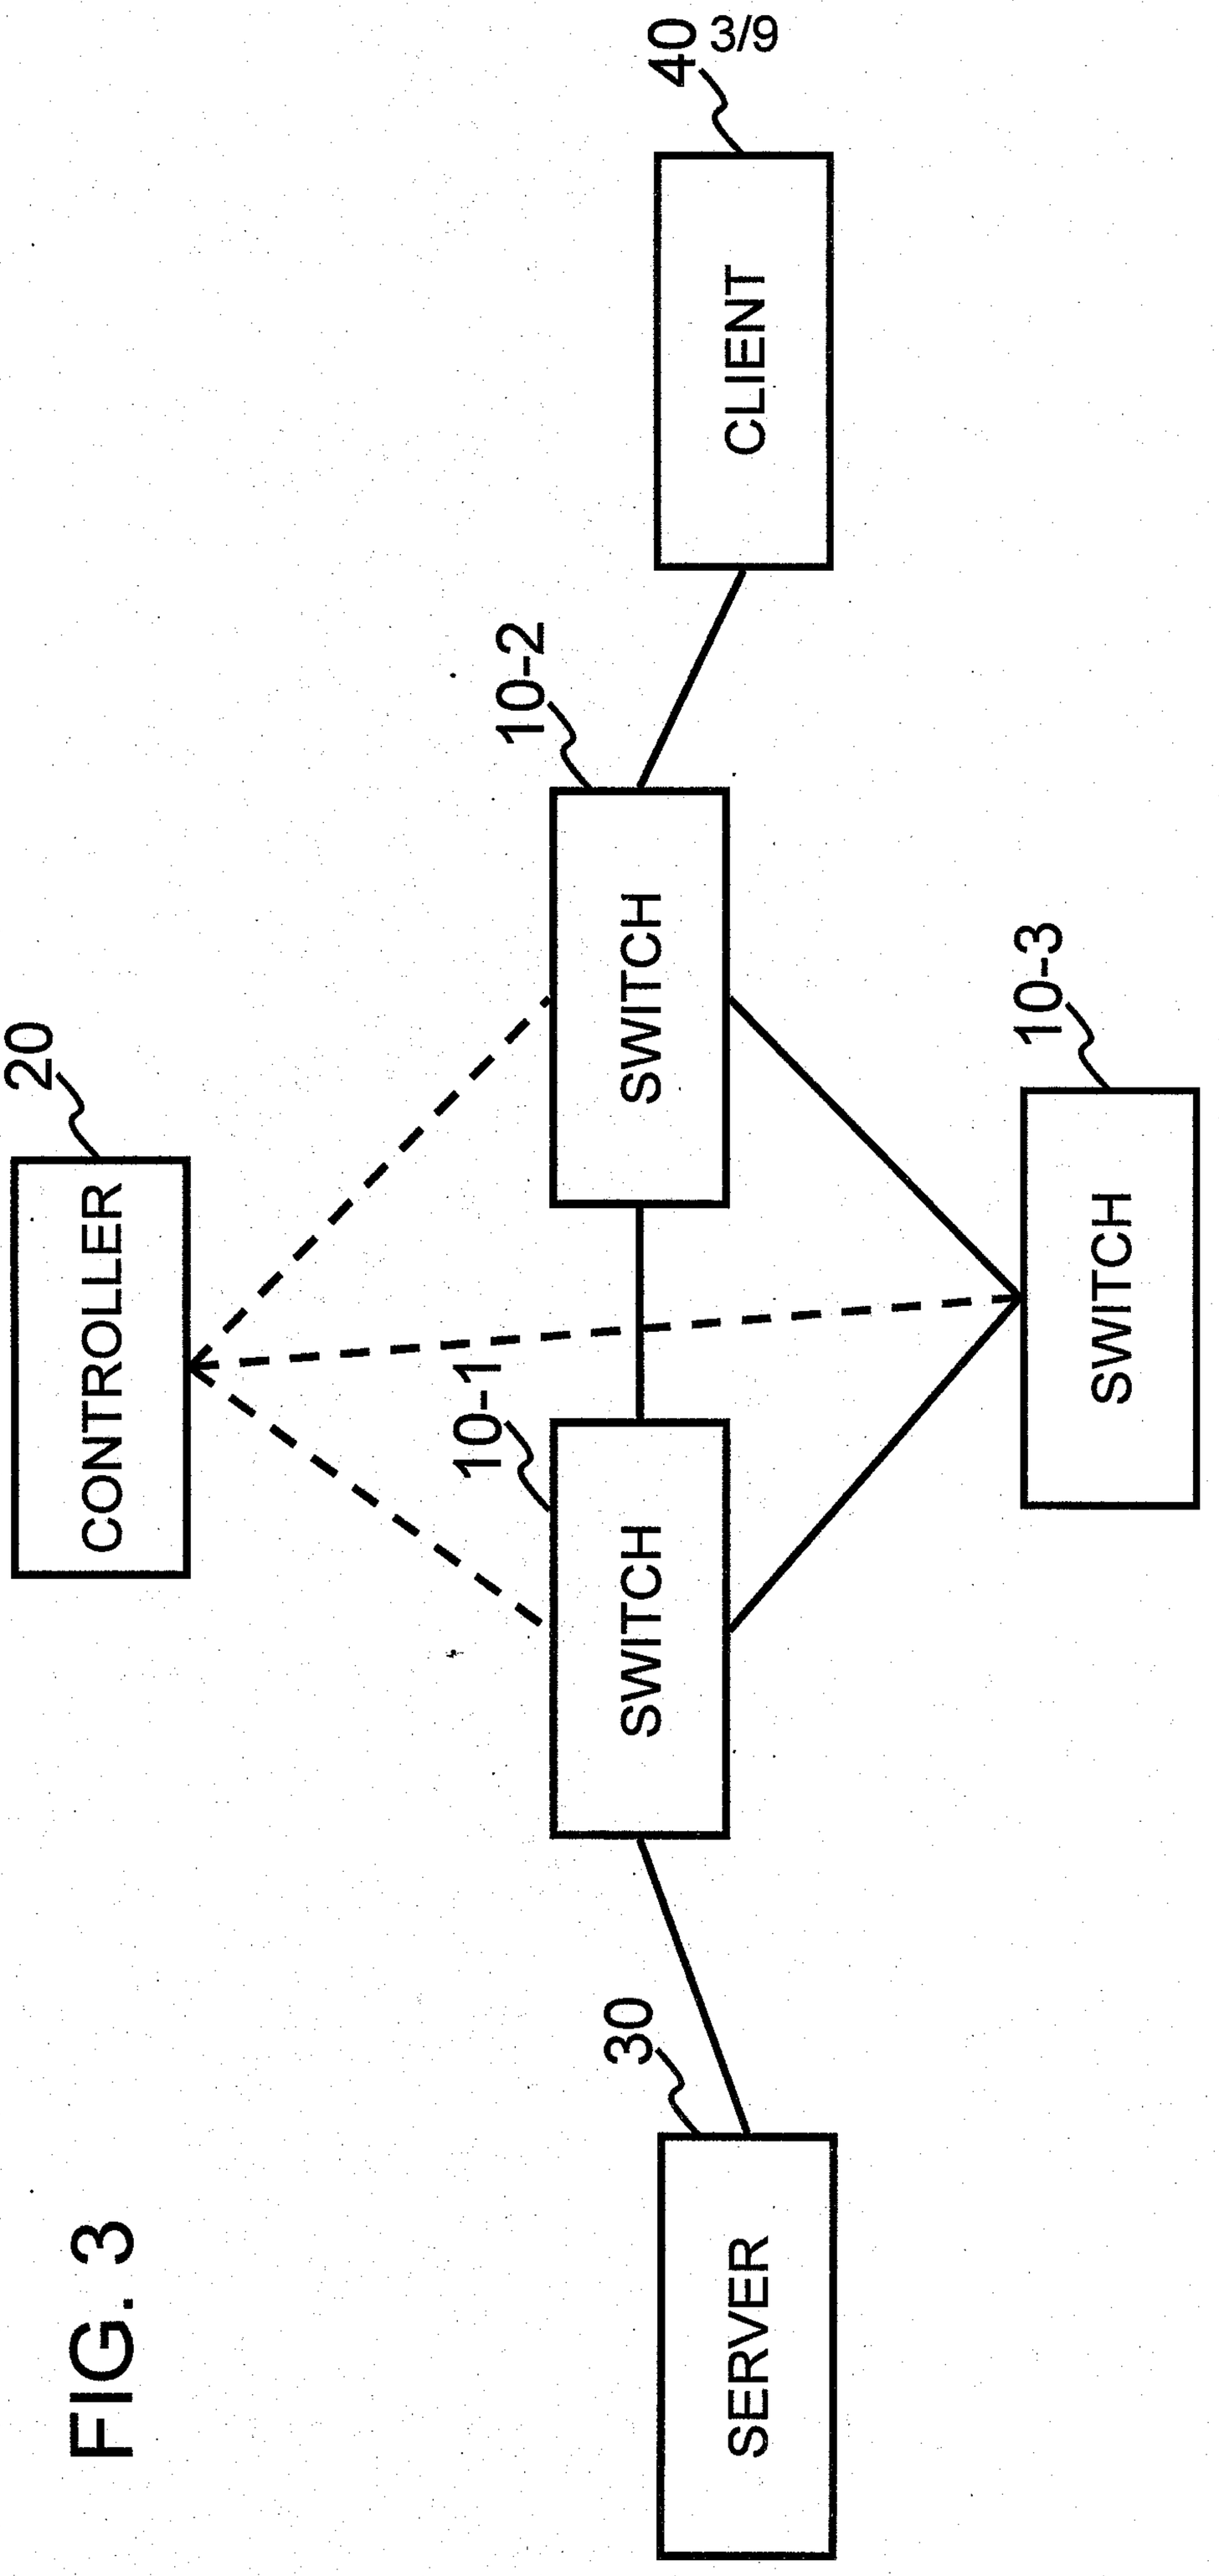
40
CLIENT
10-2
SWITCH
20
CONTROLLER
10-3
SWITCH
10-1
SWITCH
30
SERVER

FIG. 3

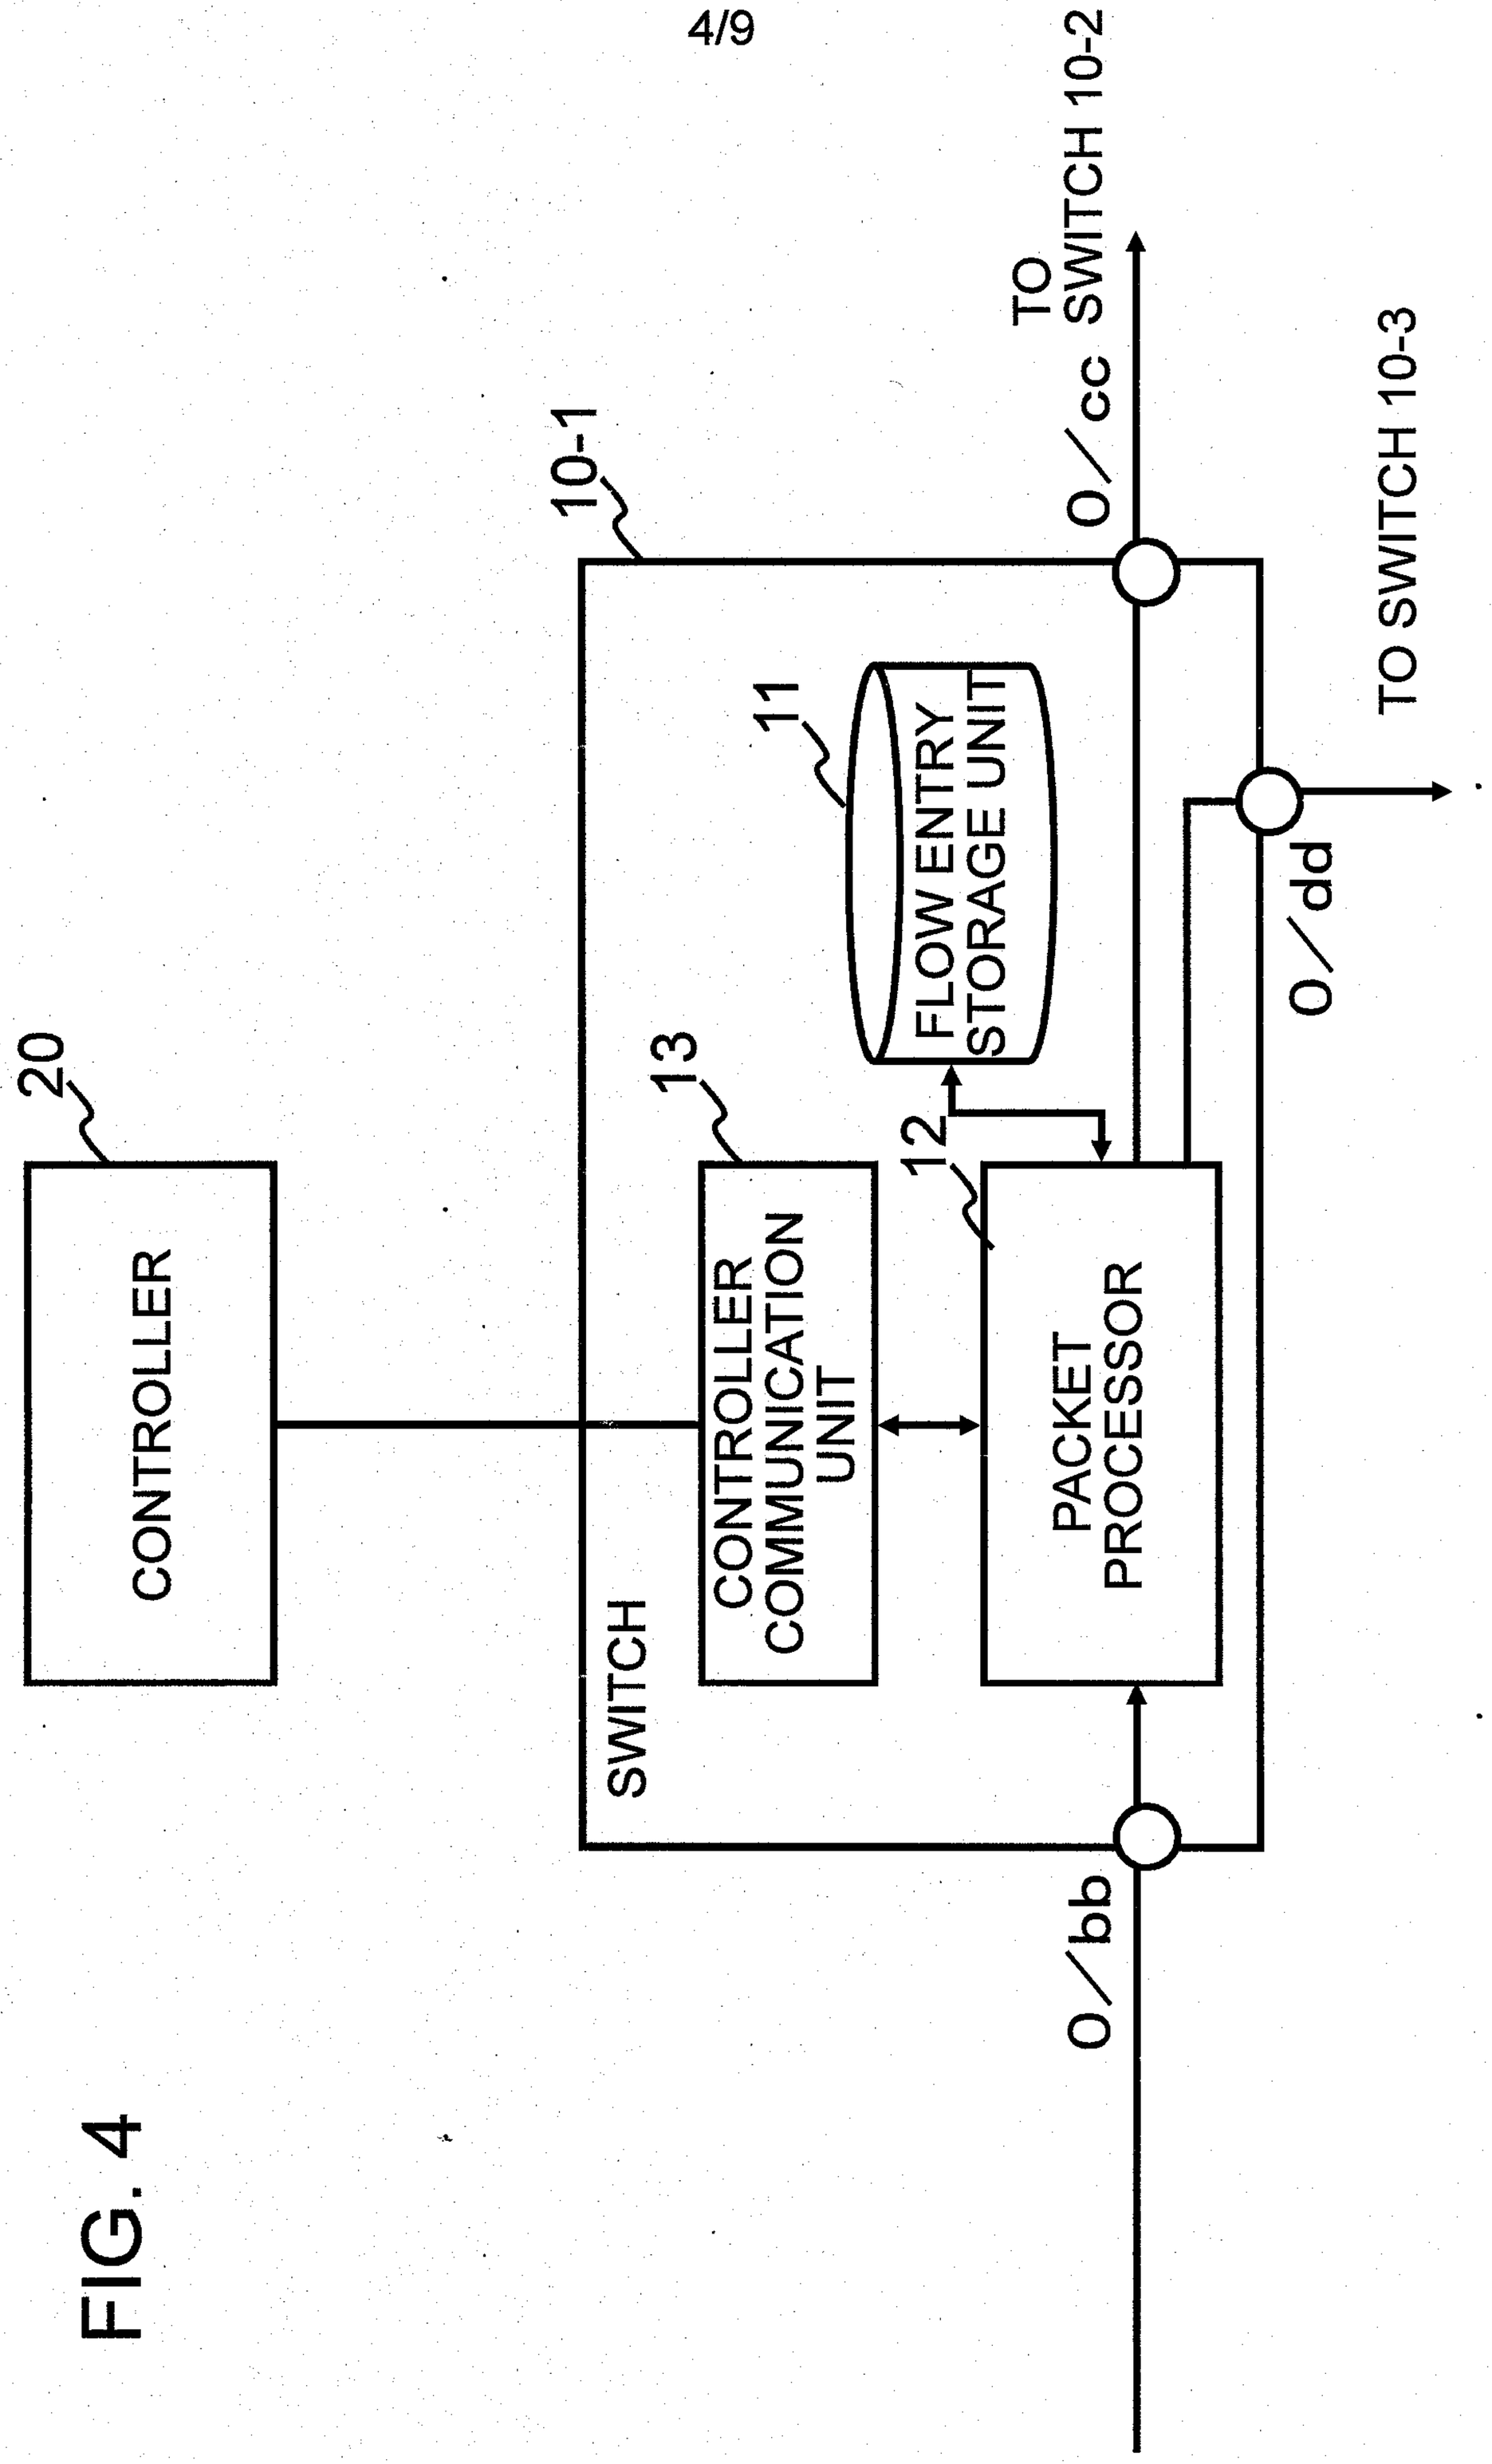
10-1
TO
O/cc SWITCH 10-2
TO SWITCH 10-3
11
FLOW ENTRY STORAGE UNIT
13
12
O/dd
20
CONTROLLER
SWITCH
CONTROLLER COMMUNICATION UNIT
PACKET PROCESSOR
O/bb

FIG. 4

FIG. 5

Priority	RETRIEVAL CONDITION	ACTION	DESCRIPTION
65535	Exact	Output=0/cc	HIGH PRIORITY FLOW ENTRY => SECOND PROCESSING RULE
⋮			
60000	dl_dest=x.x.x.x	Output=0/dd	UNICAST FLOW ENTRY => SECOND PROCESSING RULE
⋮			
a	In_port=0/bb	FLOODING PACKET-IN (FORWARDING DONE)	UNKNOWN UNICAST FLOW ENTRY => FIRST PROCESSING RULE

CA 02868712 2014-09-26

FIG. 6

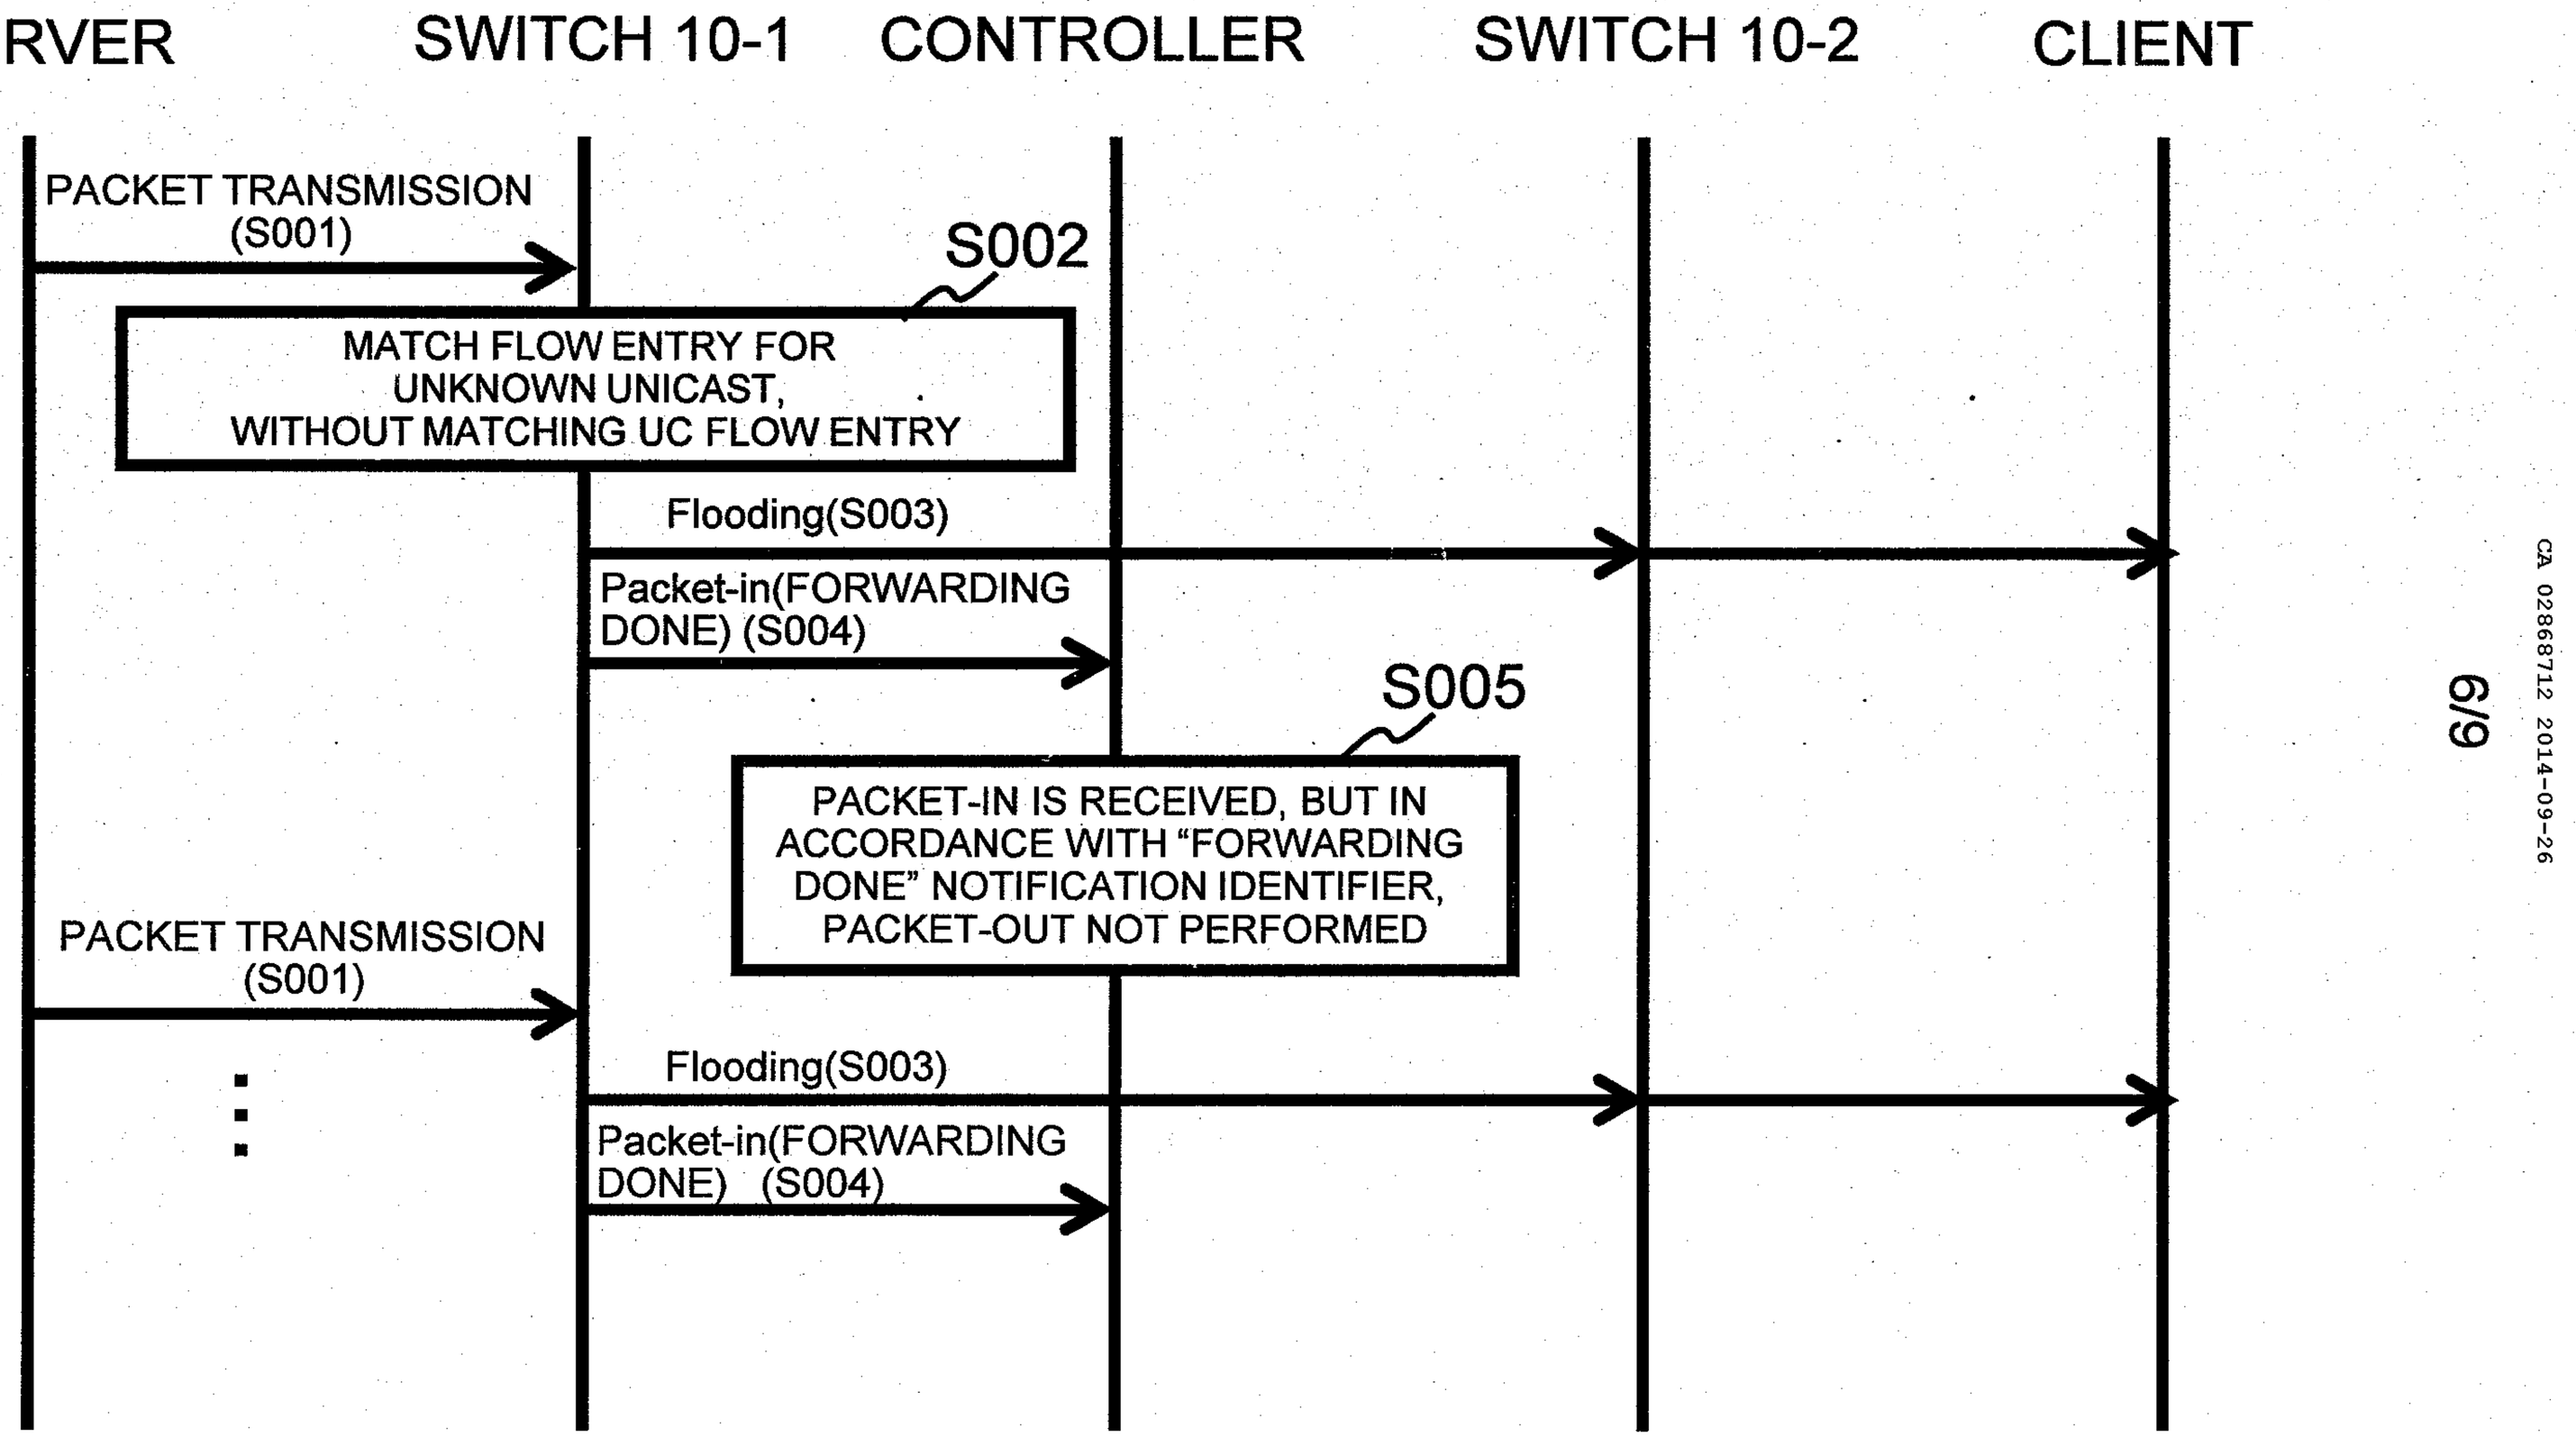

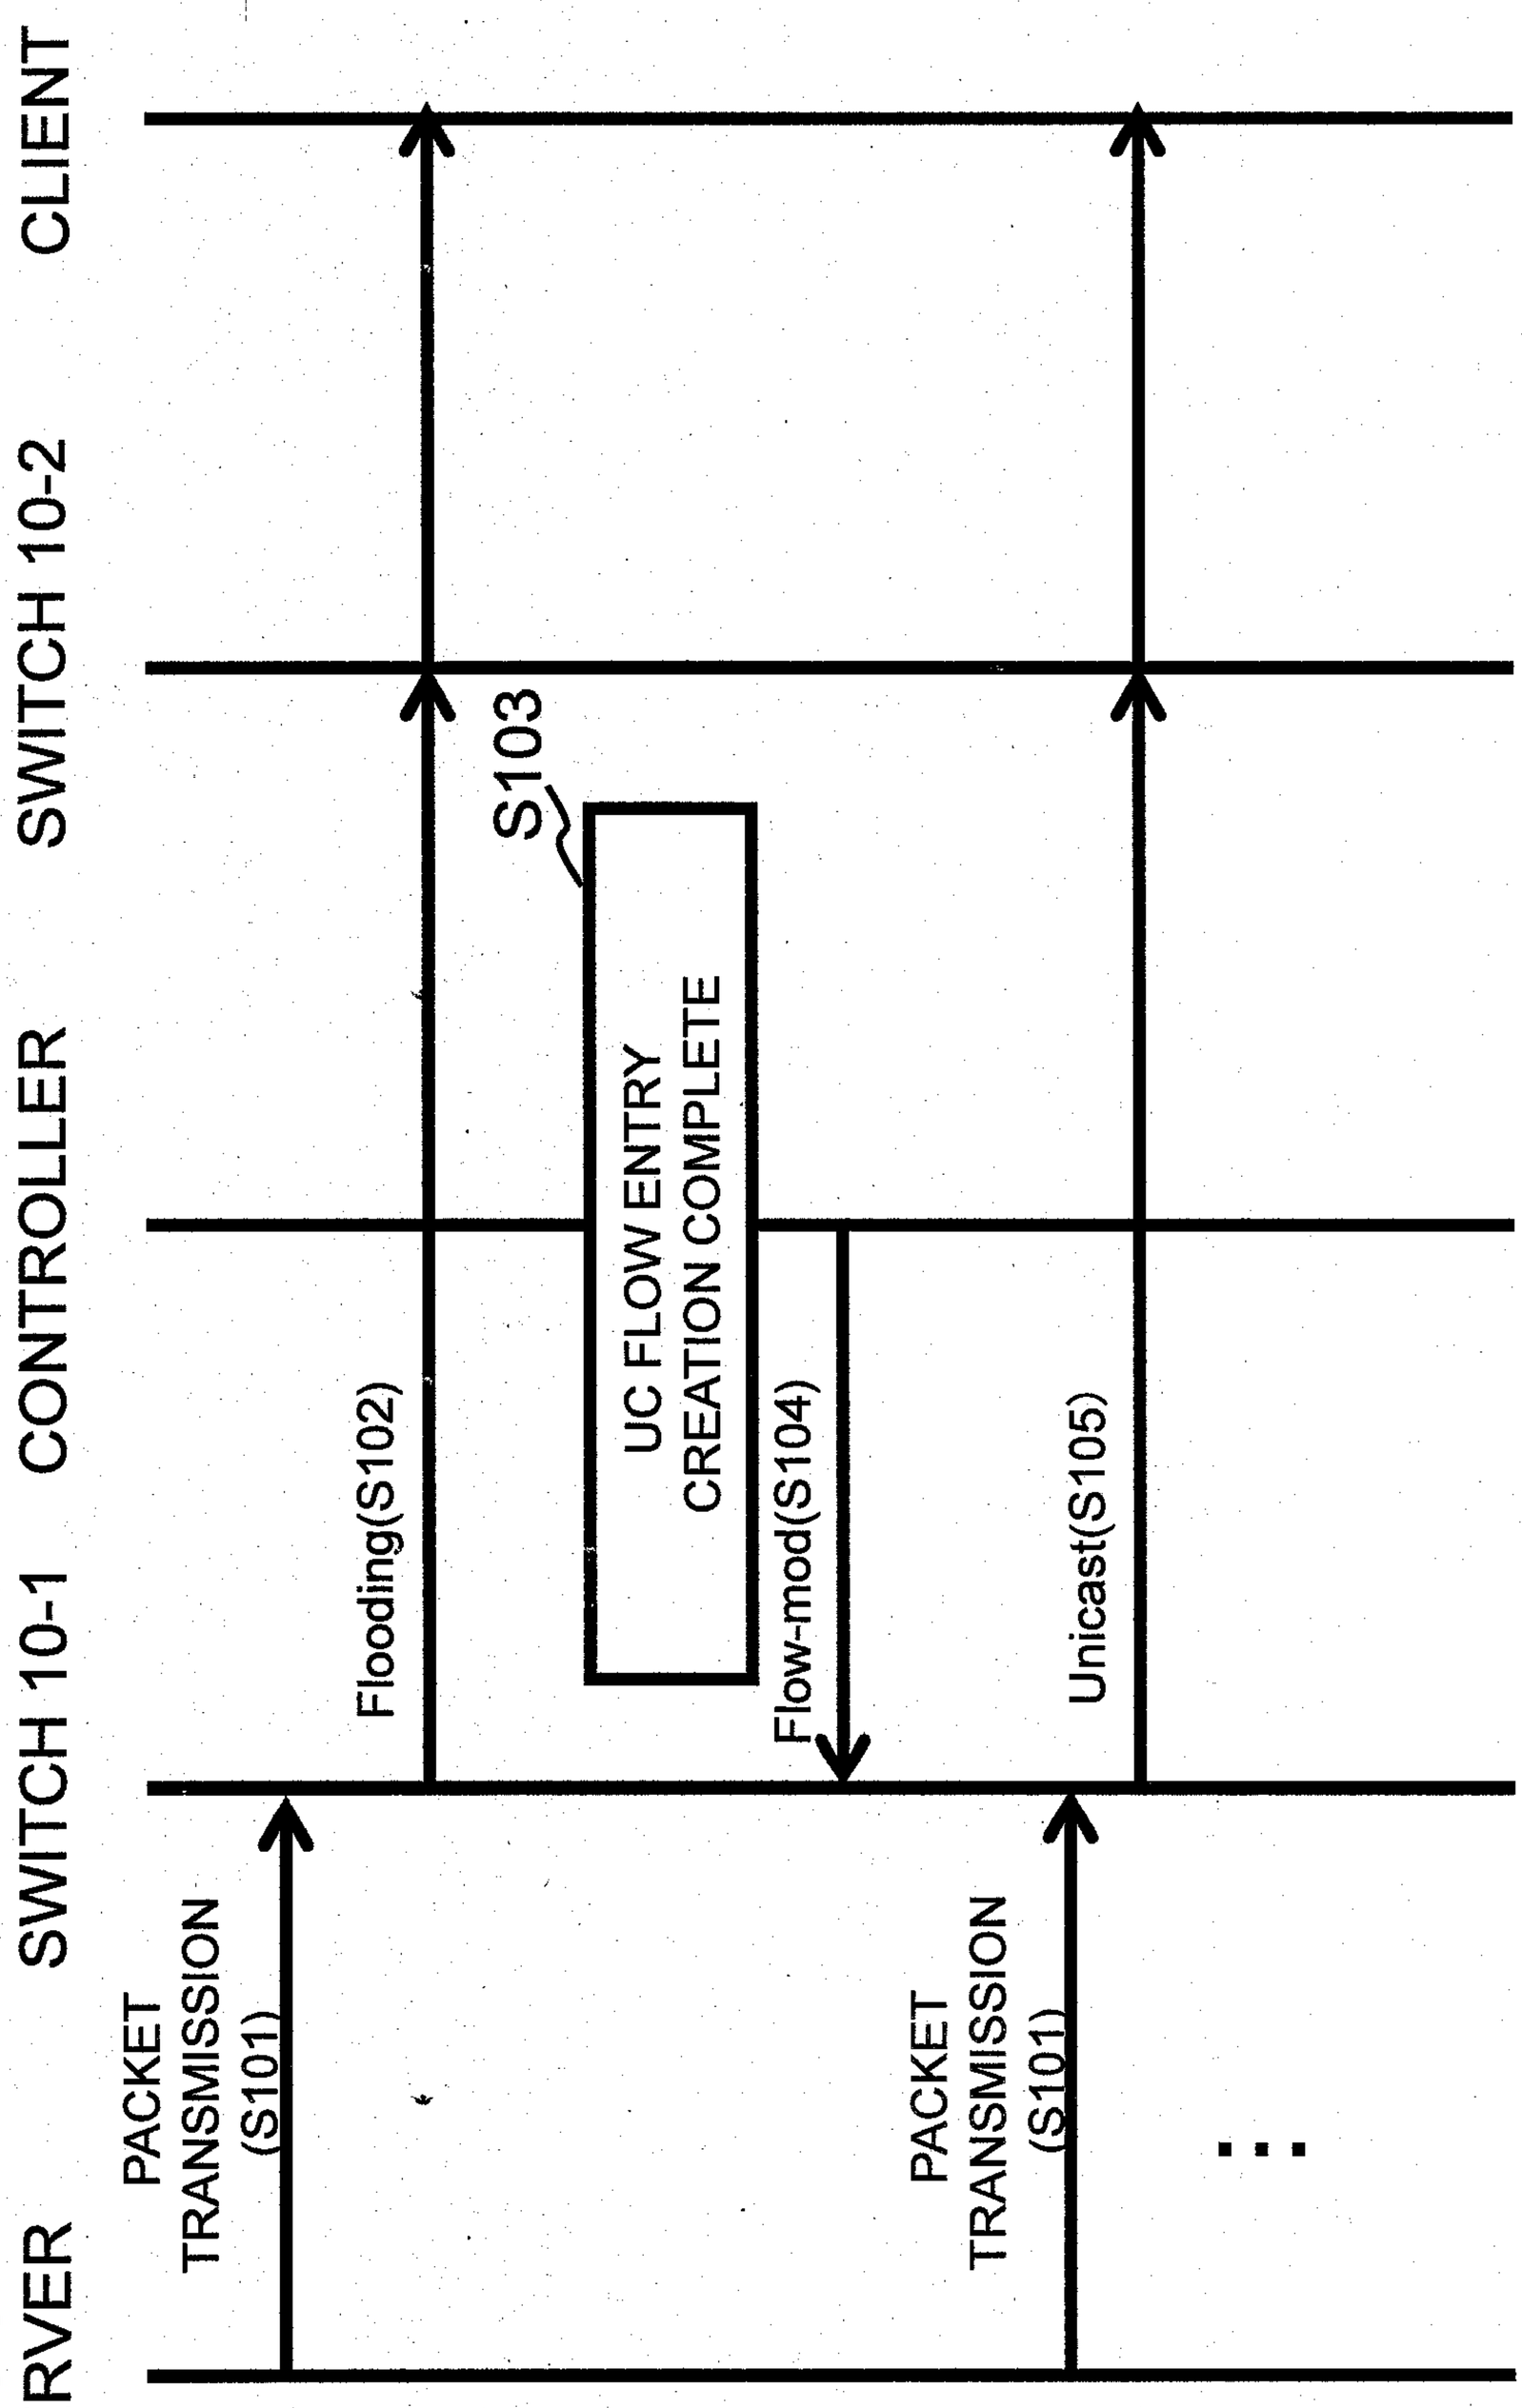
CLIENT
SWITCH 10-2
CONTROLLER
SWITCH 10-1
SERVER
PACKET TRANSMISSION (S101)
Flooding(S102)
S103
UC FLOW ENTRY CREATION COMPLETE
Flow-mod(S104)
PACKET TRANSMISSION (S101)
Unicast(S105)
...

FIG. 7

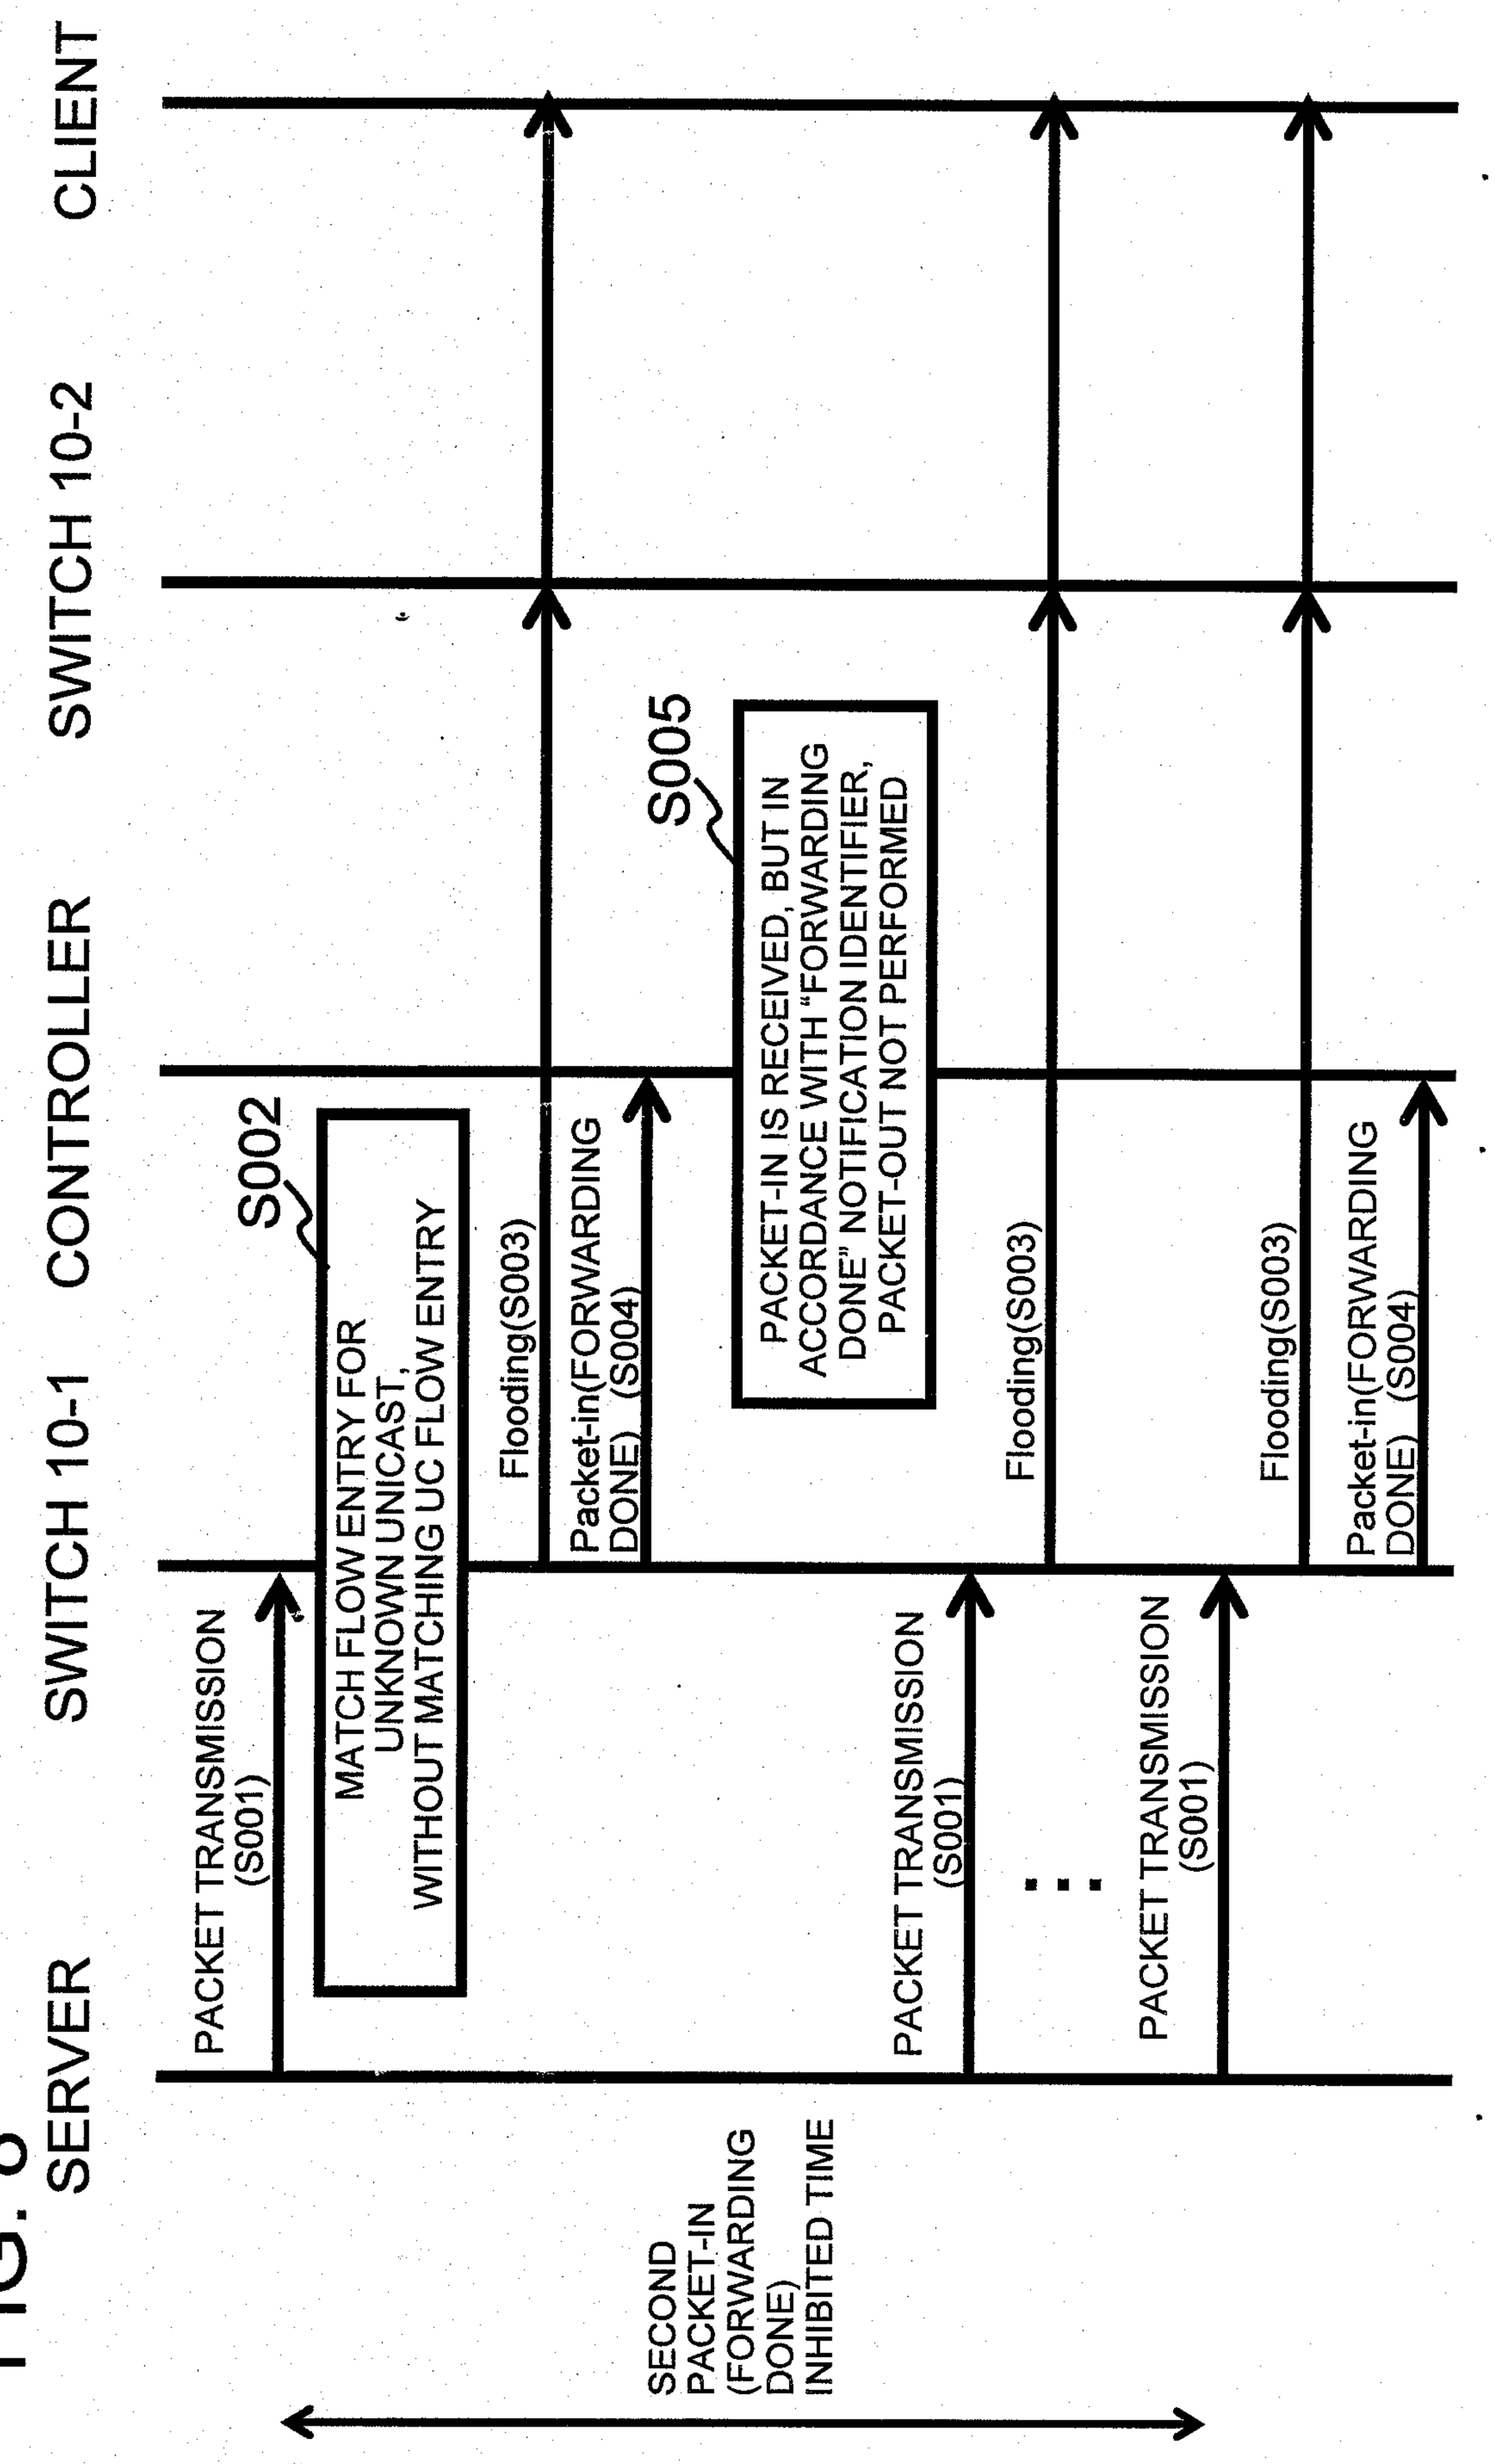
CLIENT
SWITCH 10-2
CONTROLLER
SWITCH 10-1
SERVER
PACKET TRANSMISSION (S001)
MATCH FLOW ENTRY FOR UNKNOWN UNICAST, WITHOUT MATCHING UC FLOW ENTRY
S002
Flooding(S003)
Packet-in(FORWARDING DONE) (S004)
S005
PACKET-IN IS RECEIVED, BUT IN ACCORDANCE WITH "FORWARDING DONE" NOTIFICATION IDENTIFIER, PACKET-OUT NOT PERFORMED
PACKET TRANSMISSION (S001)
Flooding(S003)
PACKET TRANSMISSION (S001)
Flooding(S003)
Packet-in(FORWARDING DONE) (S004)
SECOND PACKET-IN (FORWARDING DONE) INHIBITED TIME

FIG. 8
SE

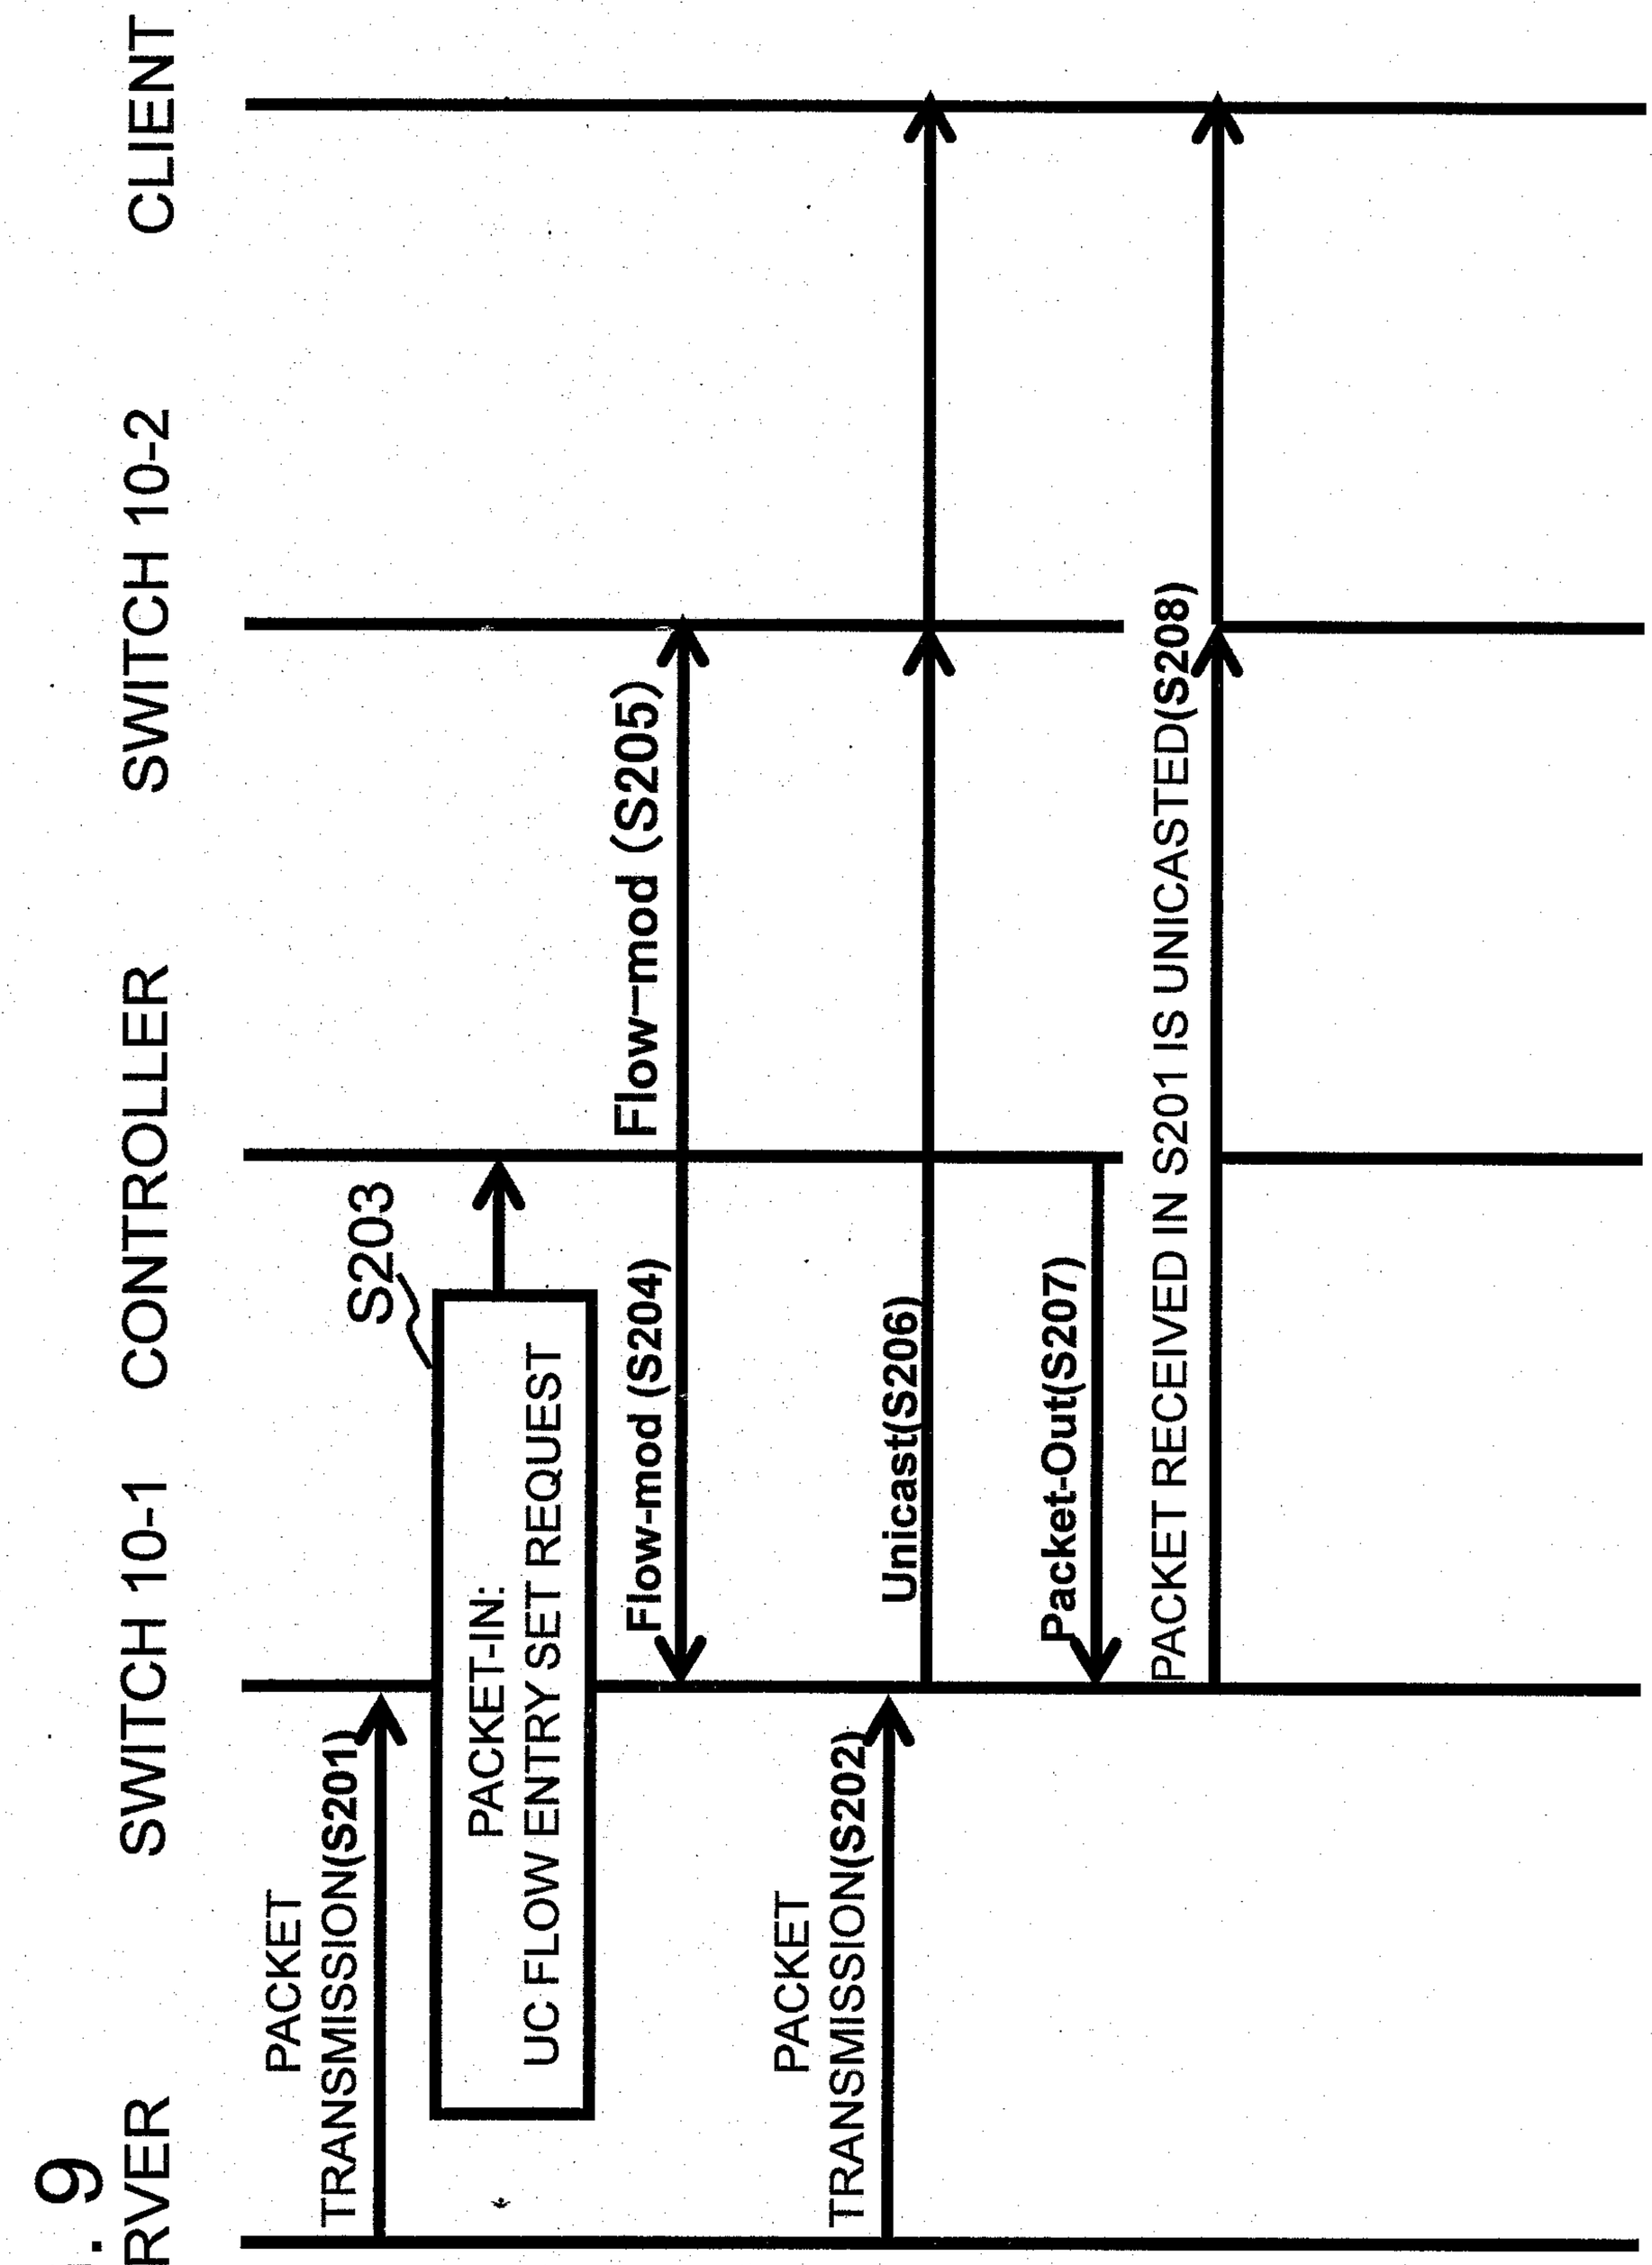
CLIENT
SWITCH 10-2
CONTROLLER
SWITCH 10-1
CONTROLLER
PACKET TRANSMISSION(S201)
PACKET-IN: UC FLOW ENTRY SET REQUEST
S203
Flow-mod (S205)
Flow-mod (S204)
PACKET TRANSMISSION(S202)
Unicast(S206)
Packet-Out(S207)
PACKET RECEIVED IN S201 IS UNICASTED(S208)

FIG. 9

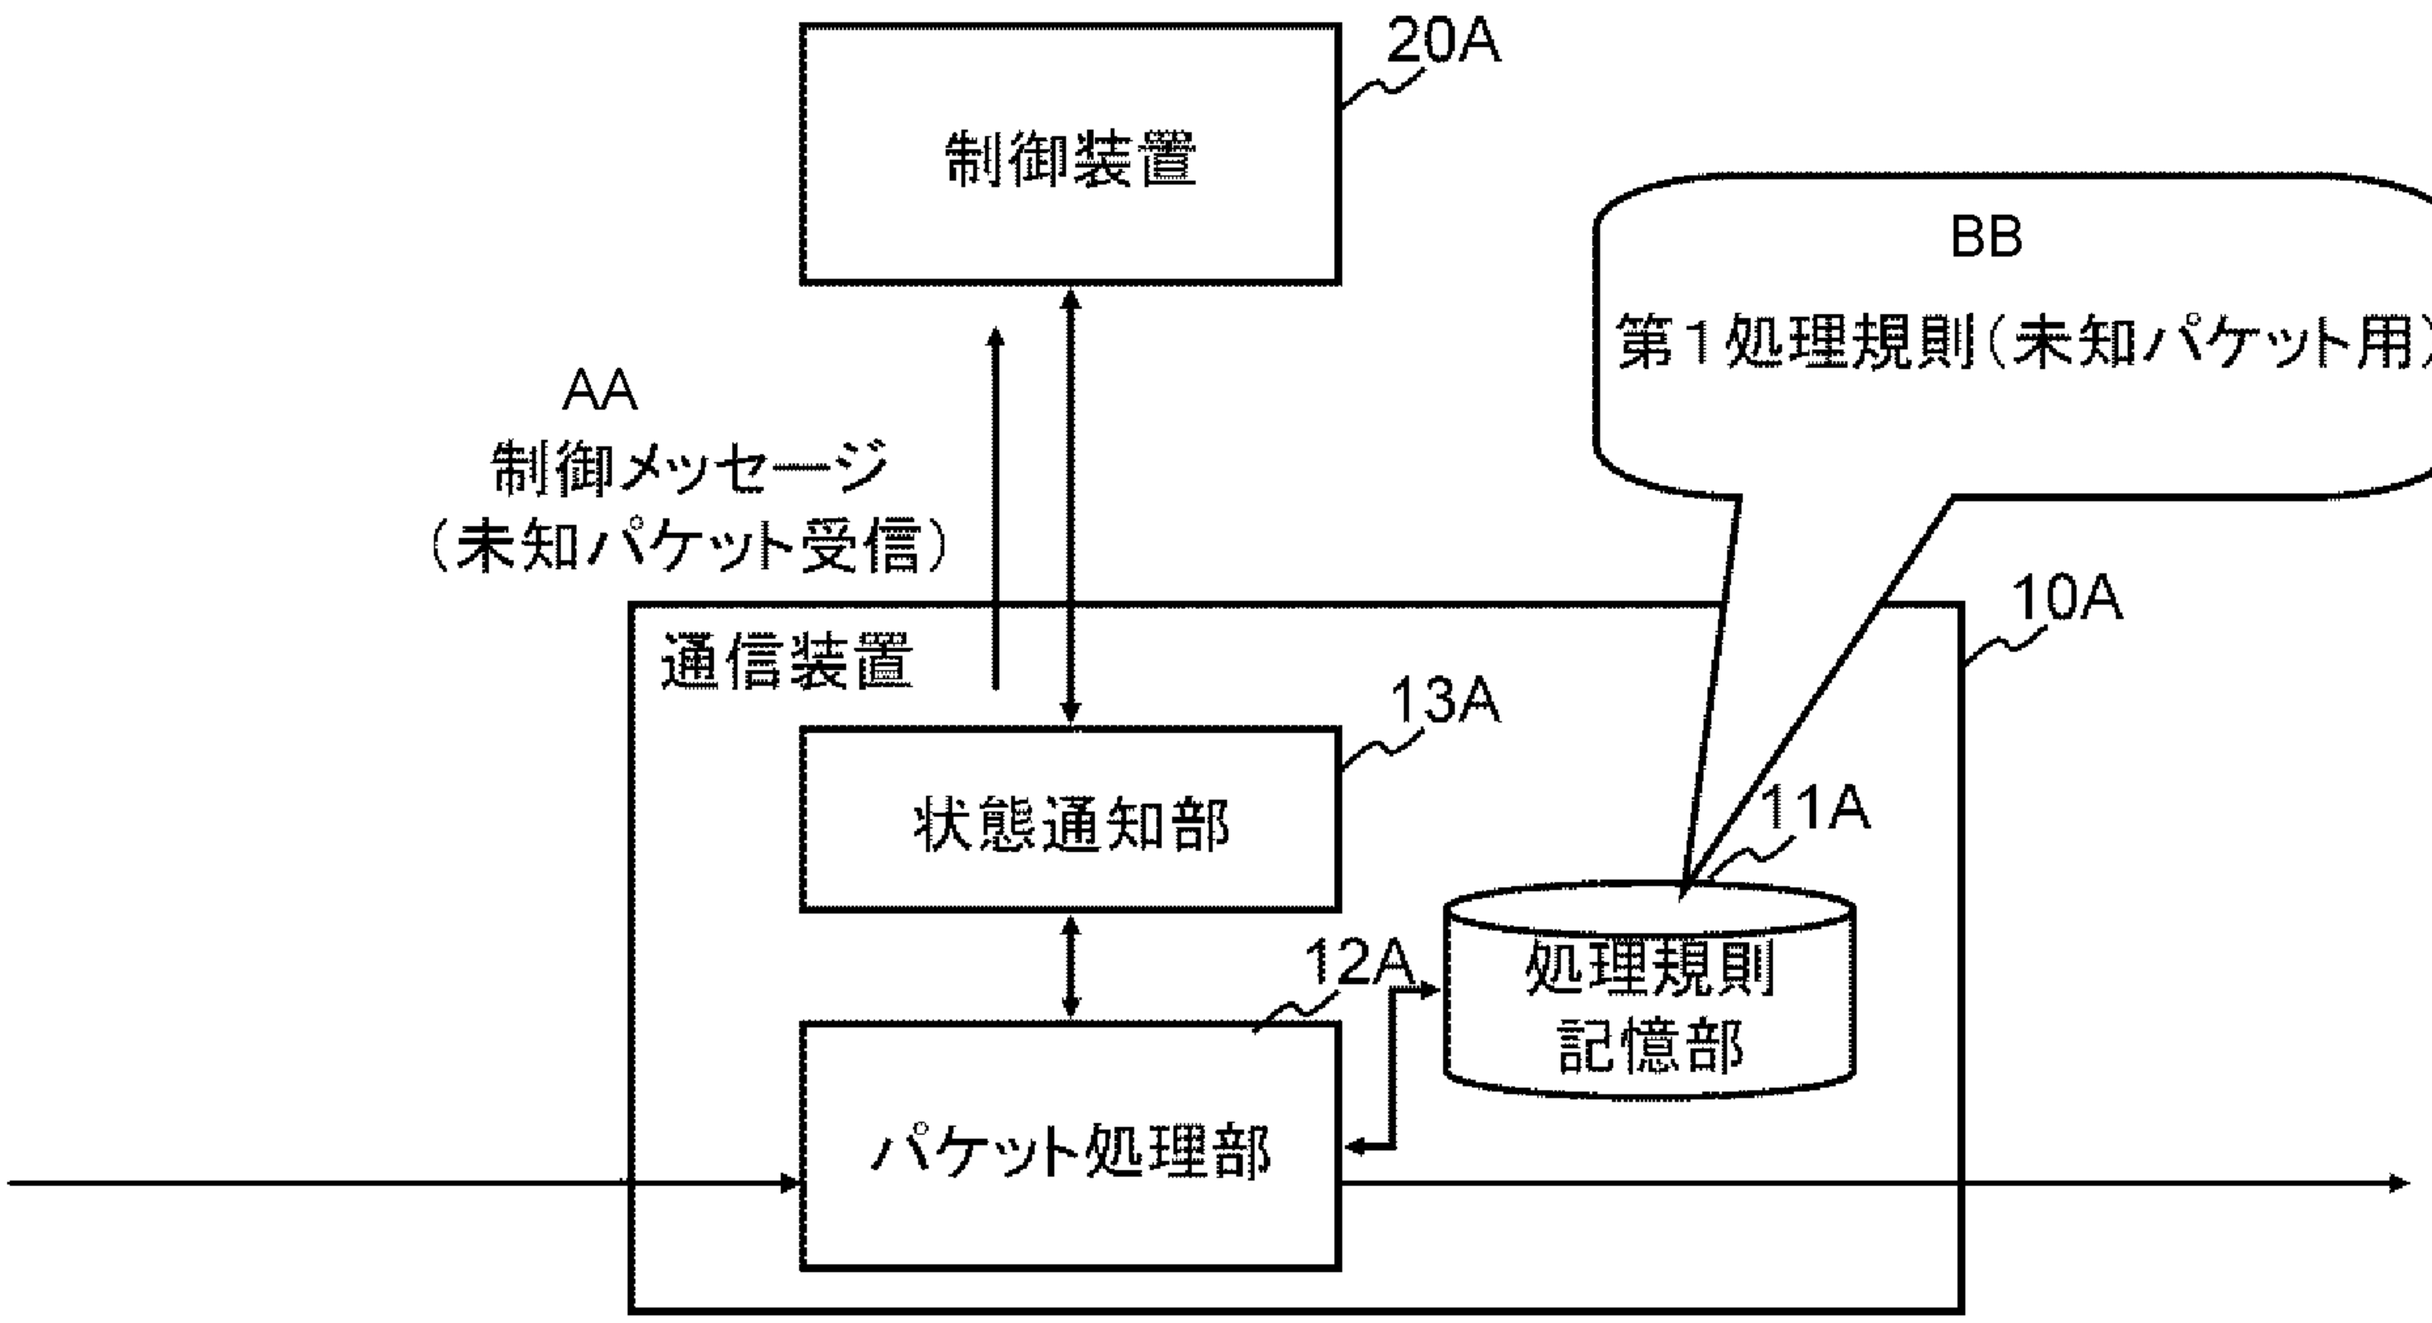

10A Communication device
11A Process-rule storage unit
12A Packet processor
13A State-notification unit
20A Control device
AA Control message (unknown packet reception)
BB First process rule (for unknown packet)