

INFORMATION FLOW RECOGNITION METHOD, NETWORK CHIP, AND NETWORK DEVICE

[0001] The disclosure claims the benefits of priority to Chinese Application No. 202110062885.8, filed January 18, 2021, which is incorporated herein by reference in its entirety.

TECHNICAL FIELD

[0002] The present application relates to the field of network communication technologies, and in particular, to an information flow recognition method, a network chip, and a network device.

BACKGROUND

[0003] In a data center network, there are many information flows, and some information flows may need to be recognized in some scenarios. For example, for an Elephant Flow, in order to improve the processing efficiency of the elephant flow, a flow table of the elephant flow is usually placed in an on-chip cache to reduce the number of accesses to an off-chip cache, so it is required to recognize the elephant flow.

[0004] In the prior art, flow table information of each received information flow may be temporarily stored in the off-chip cache, then the flow table information of each information flow may be read from the off-chip cache, and accurate statistics is performed on each information flow according to the read flow table information, thereby recognizing the elephant flow therefrom. However, an access bandwidth of the off-chip cache is often low, resulting in a slow recognition speed and low efficiency.

SUMMARY OF THE DISCLOSURE

[0005] Various aspects of the present application provide an information flow recognition method, a network chip, and a network device, so as to improve the efficiency of information flow recognition.

[0006] An embodiment of the present application provides a network chip, including: a main control chip and an off-chip memory, and the main control chip includes: an on-chip memory, a flow classification module, a flow recognizing module, and a flow screening module; the on-chip memory is configured to cache a message entering the network chip; the flow classification module is configured to map the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow; the flow recognizing module is configured to recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and output the flow table information corresponding to the potential flow to the off-chip memory to be accumulated; and the flow screening module is configured to read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.

[0007] An embodiment of the present application further provides an information flow recognition method, applicable to a main control chip in a network chip, and the method includes: caching a message entering the network chip into an on-chip memory of the network chip, and mapping the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow; recognizing a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information; outputting the flow table information corresponding to the potential flow to the off-chip memory of the network chip to be accumulated; and reading the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is

met, and recognizing a target flow in the potential flow based on the read flow table information.

[0008] An embodiment of the present application provides a network device, including: a memory, a processor, and a network chip; the memory is configured to store a computer program; the processor is coupled to the memory and configured to execute the computer program, for sending a message to the network chip and/or receiving identification information of a target flow reported by the network chip; the network chip includes: a main control chip and an off-chip memory, wherein the main control chip includes: an on-chip memory, a flow classification module, a flow recognizing module, and a flow screening module; the on-chip memory is configured to cache a message entering the network chip; the flow classification module is configured to map the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow; the flow recognizing module is configured to recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and output the flow table information corresponding to the potential flow to the off-chip memory to be accumulated; and

[0009] the flow screening module is configured to read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.

[0010] In the embodiments of the present application, a combination of an on-chip memory and an off-chip memory is used for recognizing an information flow. First, a message entering a network chip is cached in the on-chip memory, the message is mapped to at least one piece of flow table information in the chip, and a potential flow that is a

possible target flow in the information flow is recognized accordingly, then the flow table information corresponding to the potential flow is cached to the off-chip memory to be accumulated, until a set trigger condition is met, the flow table information of the potential flow accumulated in the off-chip memory is read to accurately recognize the target flow in the potential flow accordingly. Firstly, a potential flow is recognized by taking advantage of a large access bandwidth of the on-chip memory, which can increase the speed of recognizing the potential flow. Secondly, only the potential flow is further recognized by taking advantage of a large storage space of the off-chip memory, rather than further recognizing all information flows by the off-chip memory, which can reduce the number of accesses to the off-chip memory and help improve the efficiency and accuracy of recognizing a target flow.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The accompanying drawings described herein are used to provide further understanding of the present application, and construct a part of the present application. Exemplary embodiments and illustrations of the present application are used to explain the present application, and are not intended to form improper limitation on the present application. In the drawings:

[0012] **FIG. 1a** is a schematic structural diagram of a network chip according to an example embodiment of the present application;

[0013] **FIG. 1b** is a schematic structural diagram of another network chip according to an example embodiment of the present application;

[0014] **FIG. 2** is a schematic structural diagram of a data center system according to an example embodiment of the present application;

[0015] **FIG. 3** is a schematic structural diagram of a cloud computing system according to an example embodiment of the present application;

[0016] **FIG. 4** is a schematic diagram of an information flow recognition process according to an example embodiment of the present application;

[0017] **FIG. 5** is a schematic flow chart of an information flow recognition method according to an example embodiment of the present application; and

[0018] **FIG. 6** is a schematic flow chart of a network device according to an example embodiment of the present application.

DETAILED DESCRIPTION

[0019] To make objectives, technical solutions and advantages of the present application more comprehensible, technical solutions of the present application are described clearly and completely through specific embodiments of the present application and corresponding accompanying drawings. Apparently, the described embodiments are merely a part of embodiments of the present application, instead of all embodiments. Based on the embodiments of the present application, all other embodiments derived by a person of ordinary skill in the art without any creative effort shall all fall within the protection scope of the present application.

[0020] For the problem of a low information flow recognition efficiency in the prior art, embodiments of the present application provide a solution, and a basic idea lies in that: a combination of an on-chip memory and an off-chip memory is used for recognizing an information flow. First, a message entering a network chip is cached in the on-chip memory, the message is mapped to at least one piece of flow table information in the on-chip memory, and a potential flow that is a possible target flow in the information flow is recognized accordingly, then the flow table information of the potential flow is cached to the off-chip memory to be accumulated, until a set trigger condition is met, the flow table information of the potential information flow accumulated in the off-chip memory is read to accurately recognize the target flow in the potential flow accordingly. Firstly, a potential

flow is recognized by taking advantage of a large access bandwidth of the on-chip memory, which can increase the speed of recognizing the potential flow. Secondly, only the potential flow is further recognized by taking advantage of a large storage space of the off-chip memory, rather than further recognizing all information flows by the off-chip memory, which can reduce the number of accesses to the off-chip memory and help improve the efficiency and accuracy of recognizing a target flow.

[0021] The technical solutions provided in the embodiments of the present application are described in detail in combination with the accompanying drawings.

[0022] In various embodiments of the present application, data of the network chip that passes through the same network within the same period of time and has certain common features or attributes is referred to as an information flow. For example, data with the same 4-tuple, 5-tuple, or 7-tuple within a certain period of time may be regarded as an information flow, or data from the same IP address within a period of time may also be regarded as an information flow, or data that accesses the same IP address within a period of time may also be regarded as an information flow. In various network systems, there are usually various network devices such as servers, switches, routers, or hubs. These network devices usually include network chips, and the network chip can realize functions such as data transmission and reception, data encapsulation and decapsulation, link management, and data encoding and decoding. In addition, in some application scenarios, the network chip is also required to recognize an information flow.

[0023] For example, the information flow in the network system may be classified into an elephant flow and a Mice Flow. The elephant flow is a process of transferring large amounts of data continuously through a network link, for example, virtual machine migration or data migration and so on belong to the elephant flow. The mice flow is a process of transferring small amounts of short-term data through a network link, for

example, sending an email, browsing a web page, and sending an instant message belong to the mice flow. The study found that once a mice flow appears subsequent to an elephant flow in a queue at an outlet terminal of a network switch, it will cause a long queue delay. In order to optimize or ensure the transmission performance of the mice flow and alleviate the delay, it is very important to recognize the elephant flow.

[0024] For another example, in a cloud computing scenario, a cloud tenant may be interested in or pay more attention to the traffic in some scenarios. For example, the cloud tenant is an online education service provider who has deployed online education services including various subjects such as English, Mathematics, and Chinese on the cloud, and pays more attention to online traffic of the English subject in the online education services. In this case, it is necessary to recognize information flows involved in the English subject specified by the cloud tenant, and perform traffic statistics on the recognized information flows.

[0025] In the present embodiment, in order to realize information flow recognition under various scenarios or requirements, and improve the efficiency and accuracy of recognizing an information flow, a network chip is provided, and network chip may use a combination of an on-chip memory and an off-chip memory to recognize the information flow, so as to meet the purpose of information flow recognition under various scenarios or requirements. The network chip has a hardware structure and may be implemented based on an integrated circuit chip such as a Field Programmable Gate Array (FPGA), a Complex Programmable Logic Device (CPLD), or a single-chip microcomputer. In the present embodiment, an internal implementation structure of the network chip is not limited. **FIG. 1a** is a schematic structural diagram of a network chip according to an exemplary embodiment of the present application. As shown in **FIG. 1a**, a network chip 100 at least includes: a main control chip 101 and an off-chip memory 102 disposed outside the main

control chip 101. The main control chip 101 includes: an on-chip memory 101a, a flow classification module 101b, a flow recognizing module 101c, and a flow screening module 101d. Optionally, the on-chip memory 101a may be an on-chip cache, and the off-chip memory 102 may be an off-chip cache.

[0026] The on-chip memory 101a is located on the main control chip 101, and the off-chip memory 102 is located outside the main control chip 101. Relatively speaking, the access bandwidth of the on-chip memory 101a is higher than that of the off-chip memory 102. However, the storage capacity of the off-chip memory 102 may be made larger, while the capacity of the on-chip memory 101a is relatively small. In view of this, the on-chip memory 101a is mainly configured to temporarily cache some information to be processed by the main control chip 101. The off-chip memory 102 is mainly configured to cache information that does not need to be processed immediately. For the information stored in the off-chip memory 102, the main control chip 101 may first load the information from the off-chip memory 102 to the on-chip memory 101a when need to be processed.

[0027] In a practical application, the main control chip 101 needs to perform various processing on a message entering the network chip 100. The processing includes, but is not limited to, recognizing an information flow where the message is located. The embodiment of the present application focuses on the description of a process of the main control chip 101 performing the information flow recognition. For example, in an application scenario of traffic control, the elephant flow may be recognized, so that speed limit processing such as speed limit or packet loss may be performed on the elephant flow in time to solve the congestion problem. For another example, in a cloud computing scenario, an information flow specified by a cloud tenant may be recognized, so that the cloud tenant may perform subsequent processing or operations on the specified information flow. In a scenario of information flow recognition, after receiving a message,

the network chip 100 first caches the message in the on-chip memory 101a, and with the advantage of a large access bandwidth of the on-chip memory 101a, the flow classification module 101b and the flow recognizing module 101c cooperate with each other to complete preliminary flow recognition.

[0028] Specifically, the flow classification module 101b may read the cached message from the on-chip memory 101a and map the message to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow, and different pieces of flow table information correspond to different information flows. The implementation in which the flow classification module 101b reads a message from the on-chip memory 101a is not limited. For example, when the on-chip memory 101a caches a new message, the flow classification module 101b reads the new message from the on-chip memory 101a, and maps the new message to a piece of flow table information. Alternatively, the flow classification module 101b periodically reads messages cached by the on-chip memory 101a, and maps the read messages to corresponding flow table information, respectively, for example, a reading cycle may be 1s, 2s, 5s, or the like, which is not limited here.

[0029] In the present embodiment, the flow classification module 101b can not only map the message cached in the on-chip memory 101a to at least one piece of flow table information, but also provide at least one piece of flow table information to the flow recognizing module 101c. Optionally, the flow classification module 101b may output at least one piece of flow table information to the flow recognizing module 101c, or the flow recognizing module 101c may read at least one piece of flow table information from the flow classification module 101b. The flow table information corresponding to each information flow mainly includes some information related to the information flow, such as statistical information related to the information flow, some header information (such as

5-tuple) of the information flow, and the next operation (referred to as action information of the information flow) corresponding to the information flow. In the present embodiment, the statistical information related to the information flow is focused on, and the statistical information may vary due to different flow recognition requirements. For example, to recognize an elephant flow, the statistical information in the flow table information may include but not limited to: the quantity of messages that have appeared in the information flow, a time stamp of the message, identification information of the information flow, or the like. For another example, to recognize the information flow (referred to as a specified flow) specified by the user, the flow table information may include but not limited to: whether the message in the information flow carries label information, the time stamp of the message, the identification information of the information flow, or the like. In the present embodiment, the label information is used for distinguishing whether the message is a message in the specified flow, if it is a message in the specified flow, the label information may be carried, if it is not a message in the specified flow, the label information may not be carried. In addition, if the message carries the label information, the label information may also be used as a specific flow table information, which will be reflected in the flow table information corresponding to the information flow to which the message belongs. For ease of description, in the embodiment of the present application, the information flow that needs to be recognized is referred to as a target flow, and the target flow may be an elephant flow, a specified flow, or include both an elephant flow and a specified flow.

[0030] In the present embodiment, the flow recognizing module 101c may recognize a potential flow that is a possible target flow from an information flow according to at least one piece of flow table information provided by the flow classification module 101b. The potential flow refers to an information flow that may be a target flow. In the

present embodiment, in order to improve the recognition efficiency, save the storage space of the on-chip memory 101a, and ensure that each information flow entering the network chip can be recognized, in a preliminary recognition process, the information flow may be initially recognized based on flow table information obtained from a small quantity of messages. For example, for each information flow, flow table information (mainly referring to statistical information) corresponding to messages arriving within 1s or a specified quantity (such as 20) of messages may be acquired, the flow recognizing module 101c roughly recognizes the information flow according to the flow table information.

[0031] For the potential flow recognized by the flow recognizing module 101c that may be the target flow, the potential flow may be further recognized in combination with the flow table information (mainly statistical information) obtained from a larger quantity of messages, so as to improve the accuracy of the recognition result. In order to facilitate the accumulation of the flow table information corresponding to the potential flow, after the potential flow is recognized, the flow table information corresponding to the potential flow cached in the on-chip memory 101a may be output to the off-chip memory 102 to be accumulated. The advantage of large storage space of the off-chip memory 102 is used to continue to accumulate flow table information corresponding to potential flows, so that the flow table information is more abundant. Optionally, for other information flows that are not recognized as potential flows, flow table information corresponding to these information flows may be deleted from the on-chip memory 101a, so that the storage space of the on-chip memory 101a may be released, and the space utilization of the on-chip memory 101a may be improved.

[0032] In the present embodiment, the flow screening module 101d is responsible for further precisely recognizing, according to the flow table information obtained from more message statistics, the potential flow recognized by the flow recognizing module

101c, so as to determine whether the potential flow is a target flow. A trigger condition may be set for the flow screening module 101d. When the set trigger condition is met, the flow screening module 101d reads the flow table information corresponding to the accumulated potential flows in the off-chip memory 102, and recognize the target flow in the potential flow based on the read flow table information. It should be noted that, for the potential flow, the flow table information read by the flow screening module 101d from the off-chip memory 102 is usually different from the flow table information output from the on-chip memory 101a, and the flow table information in the off-chip memory 102 is obtained based on a larger quantity of messages, while the flow table information in the on-chip memory 101a is obtained based on a relatively small quantity of messages. Compared with the flow table information output by the on-chip memory 101a, the flow table information read by the flow screening module 101d from the off-chip memory 102 is more abundant and more accurate. Of course, in some special cases, there may be no messages in the information flow for a long period of time, and therefore, the two types of flow table information may also be the same.

[0033] The set trigger condition may be that the calculated quantity of messages in the flow table information cached in the off-chip memory 102 exceeds a set quantity threshold, so that when the calculated quantity of messages in the flow table information cached in the off-chip memory 102 exceeds the set data threshold, the flow screening module 101d may read the flow table information corresponding to the accumulated potential flows in the off-chip memory 102. Alternatively, the flow recognizing module 101c maintains the above-mentioned quantity threshold, calculate an accumulative quantity of messages in the flow table information output each time, and sends a trigger signal to the flow screening module 101d when the accumulative quantity of messages exceeds the set quantity threshold. When receiving the trigger signal sent by the flow recognizing

module 101c, the flow screening module 101d reads the flow table information corresponding to the accumulated potential flows in the off-chip memory 102.

Alternatively, a timing cycle may also be set, and expiration of the timing cycle may be used as a set trigger condition, so that the flow screening module 101d may read the flow table information corresponding to the accumulated potential flows in the off-chip memory 102 when each timing cycle expires.

[0034] For example, it is assumed that the flow recognizing module 101c recognizes a potential flow based on the flow table information obtained from messages arriving within 1 second (that is, the first time-length), while the flow screening module 101d further precisely recognizes a potential flow based on the flow table information obtained from messages arriving within 3 seconds (that is, the second time-length). In this way, if the flow table information output by the flow recognizing module 101c from the on-chip memory 101a to the off-chip memory 102 each time is obtained from the messages arriving within 1 second, as the flow recognizing module 101c continuously outputs the flow table information to the off-chip memory 102 to be accumulated, the flow table information corresponding to the potential flows accumulated in the off-chip memory 102 may become increasingly more abundant, so that when the timing cycle of 3 seconds expires, the flow screening module 101d may read the flow table information obtained according to statistics of messages arriving within 3 seconds from the off-chip memory. The flow screening module 101d can recognize potential flows based on the flow table information obtained from more message statistics, and therefore, the information on which flow recognition is based is more abundant and more representative, which is conducive to improving the accuracy of the recognition result. The target flow may be accurately recognized from the potential flows. The first-time length is shorter than the

second-time length, the first time-length being 1 second and the second time-length being 3 seconds is just an example of values, which is not limited thereto.

[0035] In the embodiments of the present application, a combination of an on-chip memory and an off-chip memory is used for recognizing an information flow. First, a message entering a network chip is cached in the on-chip memory, the message is mapped to at least one piece of flow table information on the chip (that is, the main control chip 101), and a potential flow that is a possible target flow in the information flow is recognized accordingly, then the flow table information corresponding to the potential flow is output to the off-chip memory to be accumulated, until a set trigger condition is met, the flow table information corresponding to the potential flow cached in the off-chip memory is read to recognize the target flow in the potential flows accordingly. Firstly, a potential flow is recognized by taking advantage of a large access bandwidth of the on-chip memory, which can increase the speed of recognizing the potential flow. Secondly, only the flow table information corresponding to the potential flow is further accumulated by taking advantage of a large storage space of the off-chip memory, and the potential flow is further recognized, rather than further recognizing all information flows by the off-chip memory, which can reduce the number of accesses to the off-chip memory and help improve the efficiency and accuracy of recognizing a target flow.

[0036] In this implementation, the target flow may be an elephant flow or a specified information flow. According to different target flows, a detailed process of recognizing the target flow by the network chip 100 is also different. Illustration is made below for different cases.

[0037] **Case A1: The target flow is an elephant flow.**

[0038] In the present embodiment, a message entering the network chip 100 is first sent to the on-chip memory 101a on the main control chip 101 for caching. On the main

control chip 101, the flow classification module 101b maps the message cached in the on-chip memory 101a to at least one piece of flow table information. In the present embodiment, the implementation in which the flow classification module 101b maps the message cached in the on-chip memory 101a to at least one piece of flow table information is not limited. Illustration is made below by using an example.

[0039] In an optional embodiment, the flow classification module 101b may map the message cached in the on-chip memory 101a to at least one piece of flow table information according to feature information of the message. The feature information of the message may be a 4-tuple, a 5-tuple, or a 7-tuple of the message. The 4-tuple includes: a source IP address, a destination IP address, a source port, and a destination port. The 5-tuple includes: a source IP address, a destination IP address, a protocol number, a source port, and a destination port. The 7-tuple includes: a source IP address, a destination IP address, a protocol number, a source port, a destination port, a service type, and an interface index. For example, messages with the same 4-tuple may be mapped to the same piece of flow table information, or messages with the same 5-tuple may be mapped to the same piece of flow table information, which is not limited.

[0040] In yet another optional embodiment, the message cached in the on-chip memory 101a may be mapped to at least one piece of flow table information by using a Hash function. Specifically, a hash table is maintained in the on-chip memory 101a, the hash table includes a certain quantity of entries, each entry corresponding to an information flow, which on the one hand is responsible for recording the flow table information corresponding to the information flow, and on the other hand recording a hash result of the information flow, different hash results representing different information flows. In other words, an entry records the hash result and flow table information of the information flow corresponding to the entry. For the message cached in the on-chip memory 101a, the flow

classification module 101b maps, by using the hash function, the message cached in the on-chip memory 101a to the hash table in the on-chip memory 101a, specifically, to an entry in the hash table.

[0041] Further optionally, each time when the on-chip memory 101a caches a new message, the flow classification module 101b may read the new message from the on-chip memory 101a, extract feature information of the new message, and hash the feature information of the new message by using the hash function to obtain a hash result of the new message. For example, if the feature information of the new message is a source IP address and a source port in the message, the source IP address and the source port in the new message may be arranged sequentially to form a string of numbers, and the string of numbers may be hashed to obtain the hash result of the new message. Messages having the same source IP address and source port may correspond to the same hash result, that is, they may be mapped to the same entry. After the hash result of the new message is obtained, the hash result may be matched in the hash table. If there is a target entry corresponding to the hash result in the hash table, it indicates that an information flow to which the message belongs has been recorded in the hash table, the new message is mapped to the target entry, and flow table information in the target entry is updated, for example, adding 1 to the quantity of messages and recording a time stamp of the new message. If there is no target entry corresponding to the hash result in the hash table, it indicates that an information flow to which the new message belongs is the latest information flow, and the information flow has not been recorded in the hash table, an idle entry is acquired from the hash table, and the new message is mapped to the idle entry, which specifically refers to recording flow table information of the information flow to which the new message belongs in the idle entry, for example, recording the quantity of

messages as 1, and recording a time stamp of the new message and an information flow identifier.

[0042] Further optionally, the quantity of entries in the hash table is limited, for example, there are 10, 20, or 25 entries. Acquiring an idle entry from the hash table may be classified into two situations: one is that there are idle entries in the hash table, and one idle entry may be directly selected from the idle entries, and the new message is mapped to the selected idle entry; the other is that there is no idle entry in the hash table, then flow table information in an entry that meets an output condition in the hash table may be output to the flow recognizing module 101c, and the flow table information in the entry may be deleted after being output to the flow recognizing module 101c, thus obtaining a blank entry. The output condition may be that if the quantity of messages exceeds a set threshold, an entry is randomly selected from the entries whose quantities of messages exceed the set threshold, and flow table information in the entry is output to the flow recognizing module 101c, thereby obtaining a blank entry. Alternatively, the output condition may be outputting flow table information in an entry corresponding to the largest quantity of messages in the hash table to the flow recognizing module 101c; the output condition may also be that a maximum difference between message time stamps in entries exceeds a set threshold, an entry is randomly selected from the entries whose maximum difference between message time stamps exceeds the set threshold and flow table information in the entry is output to the flow recognizing module 101c, which is not limited.

[0043] In the present embodiment, on the basis of the above hash table, the flow recognizing module 101c and the flow classification module 101b may adopt a pipeline operation method, that is, the flow recognizing module 101c waits for the flow classification module 101b to output flow table information to flow recognizing module 101c, and each time when receiving the flow table information output by the flow

classification module 101b, makes statistics, according to the received flow table information, on whether an information flow corresponding to the received flow table information is a potential flow. When the target flow is an elephant flow, the flow recognizing module 101c calculates the bandwidth of the information flow corresponding to the received flow table information, and recognizes the information flow as a potential flow when the bandwidth is greater than a set bandwidth threshold. Specifically, a difference between time stamps included in the flow table information may be calculated, and a ratio of the quantity of messages in the flow table information to the time stamp difference may be calculated, which is the bandwidth of the information flow. For example, a bandwidth threshold of an elephant flow may be set, and if the bandwidth of the information flow is greater than the set bandwidth threshold of the elephant flow, the information flow is considered to be a potential elephant flow.

[0044] In the present embodiment, considering that the flow recognizing module 101c recognizes whether the information flow belongs to an elephant flow based on the flow table information obtained from the statistics on a small quantity of messages, some information flows with burst traffic in a short period of time may be misrecognized as elephant flows. In other words, there may actually be some mice flows with short-term burst traffic in the elephant flows recognized by the flow recognizing module 101c. Therefore, the elephant flows recognized by the flow recognizing module 101c are referred to as potential elephant flows, which means that the potential elephant flows need to be further screened to obtain a true elephant flow. Based on this, after the potential elephant flow is recognized, flow table information corresponding to the potential elephant flow cached in the on-chip memory 101a may be output to the off-chip memory 102 to be accumulated, so as to accumulate more abundant flow table information by using the space advantage of the off-chip memory 102. In conjunction with this, a flow screening module

101d is arranged on the main control chip 101, and the flow screening module 101d is responsible for reading the flow table information corresponding to the potential elephant flows that has been accumulated in the off-chip memory 202 when a set trigger condition is met, and performing further recognition on the potential elephant flows based on the read flow table information, so as to recognize a true elephant flow in the potential elephant flows and eliminate a elephant flow not true. There may be one potential elephant flow or a plurality of potential elephant flows.

[0045] In the present embodiment, an implementation in which the flow screening module 101d recognizes the target flow from the potential elephant flow based on the read flow table information corresponding to the potential elephant flow is not limited. For example, the flow screening module 101d may perform statistics on the bandwidth of the potential elephant flow according to the read flow table information corresponding to the potential elephant flow, and recognize a potential elephant flow whose bandwidth exceeds a set elephant flow bandwidth threshold as a true elephant flow; and recognize a potential elephant flow whose bandwidth does not exceed the set elephant flow bandwidth threshold as a mice flow. For another example, the bandwidth of each potential elephant flow may be calculated, and the potential elephant flows may be sorted according to the bandwidth from large to small, and first N potential elephant flows that occupy 80% to 100% of the network bandwidth may be recognized. The first N potential elephant flows are recognized as true elephant flows, and remaining potential elephant flows are recognized as mice flows.

[0046] **Case A2: The target flow is a specified flow.**

[0047] In the present embodiment, as shown in **FIG. 1b**, the main control chip is further provided with a flow labeling module 101e. The flow labeling module 101e is responsible for determining, before a message entering the network chip 100 is cached to

the on-chip memory 101a, whether the message entering the network chip 100 is a message in a specified flow according to feature information of the message, and if so, adding label information in the message, wherein the label information is used for identifying whether the message belongs to the specified flow, and the label information may be used as a type of specific flow table information, and may be recorded in flow table information of the specified flow. The label information may be a number, a character string, or a combination of a number and a character string, which is not limited here.

[0048] In the present embodiment, an upper-layer application may deliver feature information of the message in the specified flow to the flow labeling module 101e in advance, or a relevant operator may configure the feature information of the message in the specified flow to the flow labeling module 101e. In this way, when receiving the message, the flow labeling module 101e may extract feature information from the message, and compare the feature information extracted from the message with the feature information of the message in the specified flow stored locally; if the two are the same, determine that the message belongs to the specified flow, add label information in the message, and send the message after adding the label information into the on-chip memory 101a for caching; and if the two are different, directly send the message into the on-chip memory 101a for caching. For example, if a source IP address of the message in the specified flow is IP1 and a source port is M, it may be determined whether a source IP address in the feature information of the message is IP1 and whether a source port is M, if so, the message is considered to be the message in the specified flow, and label information corresponding to the specified flow is added to the message; and if not, the message is considered not to be a message in the specified flow.

[0049] In the present embodiment, the flow classification module 101b can map the message cached in the on-chip memory to at least one piece of flow table information, and

output at least one piece of flow table information to the flow recognizing module 101c. When the flow classification module 101b recognizes that the message carries label information, the label information is recorded as a specific flow table information into the flow table information corresponding to the information flow. In other words, the flow table information includes the quantity of messages, the time stamp, and whether there is label information added by the flow labeling module 101e. For a specified flow, its flow table information includes label information; and for a non-specified flow, its flow table information does not include label information. For the implementation in which the flow classification module 101b maps the message cached in the on-chip memory to at least one piece of flow table information, reference may be made to the foregoing embodiment, and details are not repeated here.

[0050] When the target flow is the specified flow, after receiving the flow table information output by the flow recognizing module 101c, the flow recognizing module 101d may determine whether the flow table information includes label information, recognize an information flow corresponding to the flow table information including the label information as a potential flow, and output the flow table information corresponding to the potential flow to the off-chip memory 102 to be accumulated, so as to cache a larger amount of and more abundant flow table information.

[0051] In the present embodiment, the flow screening module 101d may read the flow table information corresponding to the accumulated potential flow in the off-chip memory 102 when the set trigger condition is met, and performs further recognition on the potential flow based on the read flow table information, so as to recognize the specified flow in the potential flow. For related descriptions about the set trigger condition, reference may be made to the foregoing embodiments, and details are not repeated here. For example, the flow screening module 101d may determine whether the calculated quantity

of messages in the flow table information carrying the label information is greater than a set threshold; if yes, determine that the potential flow belongs to the specified flow; and if no, determine that the potential flow is not the specified flow, but a misjudged information flow.

[0052] Here, on the basis that the main control chip 101 includes the flow labeling module 101e, the main control chip 101 can not only recognize the specified flow, but also recognize the elephant flow, and even recognize the specified flow at the same time in the process of recognizing the elephant flow. The process of the main control chip 101 simultaneously recognizing the elephant flow and the specified flow may be obtained with reference to the following case A3.

[0053] Case A3: The target flow includes both the elephant flow and the specified flow.

[0054] In the present embodiment, an information flow entering the network chip 100 first enters the flow labeling module 101e, and the flow labeling module 101e determines whether the message is a message in the specified flow according to feature information of the message, if so, adds label information to the message and send the message with the added label information into the on-chip memory 101a for caching; and if not, directly send the message into the on-chip memory 101a for caching. For details about adding the label information to the message by the flow labeling module 101e, reference may be made to the description in case A2, which may not be repeated here.

[0055] In the present implementation, the message entering the network chip 100 may be cached in the on-chip memory 101a no matter whether the label information is added thereto by the flow labeling module 101e or not. On the main control chip 101, the flow classification module 101b may map the message cached in the on-chip memory 101a to at least one piece of flow table information, and may output at least one piece of flow

table information to the flow recognizing module 101c when an output condition is met, wherein the flow table information includes: the quantity of messages, the time stamp, and whether it includes the label information added by the flow labeling module 101e. For the implementation in which the flow classification module 101b maps the message cached in the on-chip memory 101a to at least one piece of flow table information, reference may be made to the case A1, and details are not repeated here.

[0056] In the present embodiment, the flow recognizing module 101c recognizes a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and outputs the flow table information corresponding to the potential flow cached in the on-chip memory to the off-chip memory to be accumulated. In the present embodiment, when the target flow includes both the elephant flow and the specified flow, the potential flow recognized by the flow recognizing module 101c may include a potential elephant flow and/or a potential specified flow. For a piece of flow table information output by the flow classification module 101b, the flow recognizing module 101c may first recognize whether an information flow corresponding to the flow table information is a potential elephant flow, and then recognize whether the information flow is a potential specified flow; or first recognize whether the information flow corresponding to the flow table information is a potential specified flow, and then recognize whether the information flow is a potential elephant flow; or recognize whether the information flow corresponding to the flow table information is a potential elephant flow and a potential specified flow at the same time, which is not limited here.

[0057] The flow recognizing module 101c first recognizing whether the information flow corresponding to the flow table information is a potential elephant flow, and then recognizing whether the information flow is a potential specified flow is taken as an example. Specifically, each time when receiving flow table information output by the

flow classification module 101b, the flow recognizing module 101c calculates the bandwidth of an information flow corresponding to the received flow table information according to the received flow table information, and recognizes the information flow as a potential elephant flow when the bandwidth is greater than a set threshold, and outputs the flow table information corresponding to the potential elephant flow to the off-chip memory 102 to be accumulated. When the bandwidth is less than or equal to the set threshold, it may be further determined whether the received flow table information includes label information, if the label information is included, the information flow corresponding to the information flow corresponding to the flow table information is recognized as a potential specified flow and the flow table information corresponding to the potential specified flow is output to the off-chip memory 102. For another information flow that is not recognized as a potential flow, corresponding flow table information thereof may be deleted from the on-chip memory 101a, so as to save the bandwidth of the on-chip memory 101a.

[0058] The flow screening module 101d reads the flow table information corresponding to the accumulated potential flow in the off-chip memory 102 when a set trigger condition is met. The flow screening module 101d may distinguish whether flow table information corresponding to a potential elephant flow or flow table information corresponding to a potential specified flow is read according to whether the read flow table information includes the label information. For the potential specified flow, the flow screening module 101d may further determine whether the information flow is a specified flow according to the flow table information read from the off-chip memory 102 and including the label information. For the potential element flow, the flow screening module 101d may further determine whether the information flow is an elephant flow according to the flow table information read from the off-chip memory 102 and not including the label information. Regarding the implementation of further recognizing the elephant flow and

the specified flow, reference may be made to the descriptions in cases A1 and A2, and details are not repeated here.

[0059] In an optional embodiment, as shown in **FIG. 1b**, the main control chip 101 is further provided with a processor 101f, and the flow screening module 201d may report identification information of the target flow to the processor when the target flow is recognized. The identification information refers to information that can uniquely identify an information flow, such as a source IP address, a destination IP address, or a port number of the information flow, which is not limited.

[0060] In the present embodiment, after receiving the identification information of the target flow, the processor 101f may perform various operations based on the identification information. In an optional embodiment, the processor 101f may perform flow control on the target flow according to the identification information of the target flow. For example, if the target flow is an elephant flow, and the processor detects that its own utilization rate exceeds a set threshold, for example, whether it exceeds 90%, and if the utilization rate of the processor 101f exceeds the set threshold, in order to ensure the stability of the processor 101f, the processor 101f may perform flow-limiting processing on the elephant flow. In yet another optional embodiment, no matter the elephant flow or the mice flow may occupy a piece of flow table information, the flow table information is the basis for information flow forwarding, and entries of the flow table information integrate network configuration information of various level in the network, so that more abundant rules may be used during data forwarding. Normally, the flow table information is stored in the off-chip memory 102. When an information flow arrives, the processor accesses the off-chip memory 102, reads flow table information of the information flow, and performs corresponding operations. However, for the elephant flow, the processor 101f needs to frequently access the off-chip memory 102 to read the flow table information, and

the off-chip memory 102 has a low access bandwidth and low efficiency. Based on this, the processor 101f may load the flow table information of the target flow from the off-chip memory 102 to the on-chip memory 101a according to the identification information of the target flow. The access bandwidth of the on-chip memory 101a is higher than that of the off-chip memory 102, and therefore, the efficiency of the processor 101f may be improved. In yet another optional embodiment, after receiving the identification information of the target flow, the processor 101f may perform accurate statistics on the target flow based on the identification information of the target flow, for example, calculate the bandwidth of the target flow, or perform interception, risk control, or special processing on the target flow.

[0061] In some embodiments of the present application, the network chip 100 may be implemented as a network card, and may be applied to a switch, a router, or a server in any network system to implement information flow sending and receiving, recognition, and other processing. A network card-based information flow recognition process in a data center system and a cloud computing system is introduced below.

[0062] **FIG. 2** is a schematic structural diagram of a data center system according to an example embodiment of the present application. As shown in **FIG. 2**, the data center system 200 includes: a plurality of servers 201 and a plurality of network switching devices 202. The server 201 is mainly responsible for performing various computing tasks, and may be regarded as a terminal side device. The server 201 is only an example of a terminal side device, which is not limited thereto. The network switching device 202 mainly functions to realize interconnection between the servers 201, and may be regarded as a network side device. The plurality of servers 201 are interconnected through the plurality of network switching devices 202, and network data (such as various messages) between servers 201 may be forwarded through the network switching devices 202.

[0063] As shown in **FIG. 2**, a server 201 may be directly in communication connection with one, two, or more network switching devices 202, or be directly in communication connection with another server 201, and use the other server 201 as a relay to be indirectly in communication connection with one, two, or more network switching devices 202. The communication connection here may be a wired connection or a wireless connection.

[0064] It should be noted that, in addition to the servers 201 and the network switching devices 202, the data center system further includes some photoelectric lines for implementing interconnection between the servers 201 and the network switching devices 202. In the present embodiment, the quantity of servers 201 and the quantity of network switching devices 202 are not limited, and may be determined by the scale of the data center system. For example, in some large-scale data center systems, a single cluster may include thousands of network switching devices, tens of thousands of servers, and hundreds of thousands of photoelectric circuits.

[0065] In the present embodiment, the implementation form of the network switching device 202 is not limited, may be any device with functions such as device interconnection, data exchange, and forwarding, and may include, for example, a router, a switch, a hub, or the like. For example, in the data center system shown in **FIG. 2**, the network switching device 202 includes, but is not limited to, a switch and a router.

[0066] In the present embodiment, as shown in **FIG. 2**, an application program is installed on a server F1, the application program may generate data, and the data generated by the application program may reach a network card W1 on the server F1 after being processed by various layers. The network chip encapsulates the data into a message, and sends the message to a switch F2 through a photoelectric link between it and the switch F2. The switch F2 also has a network card W2, and after receiving the message, the network

card W2 recognizes that it needs to be forwarded to a server F4. The data packet may be forwarded to the server F4 via a router F3 and its network card W3. After receiving the message, the network card W4 on the server F4 provides the message to a processor of the server F4, and the message is further processed by the processor from the perspective of software.

[0067] In the present embodiment, in order to realize functions such as load balancing, congestion control, and network anomaly detection, the server needs to recognize an elephant flow when transmitting an information flow. After the elephant flow is recognized, flow control of the elephant flow may be performed, or flow table information of the elephant flow may be loaded from the off-chip memory to the on-chip memory, for improving the efficiency of the data center system. Both the server 201 and the network switching device 202 include a network chip, and the internal structure of the network chip may be obtained with reference to the aforementioned embodiments, which may not be described in detail here. A process of recognizing the elephant flow is described in detail may be described in detail below by using the network card W4 on the server F4.

[0068] After receiving the message on the network link, the network card W4 temporarily caches the message in the on-chip memory. Each time when there is a new message cached in the on-chip memory, the flow classification module in the network card W4 hashes feature information of the message by using a hash function to obtain a hash result V of the new message. If there is a target entry corresponding to the hash result V in the hash table, the new message is mapped to the target entry, the quantity of messages in the flow table information in the target entry is increased by 1, and a time stamp of the message is recorded. If there is no target entry corresponding to the hash result V in the hash table, an idle entry is acquired from the hash table, the new message is mapped to the

idle entry, the quantity of messages in the idle entry is increased by 1, and a time stamp of the message is recorded. If there is no idle entry in the hash table, the quantity of messages and the time stamp of the message recorded in an entry having the largest quantity of messages in the hash table are output to the flow recognizing module in the network card W4, so as to obtain a blank entry. The quantity of messages and the time stamp of the message recorded in the entry belong to the flow table information corresponding to the information flow, which may be all or part of the flow table information, which is not limited.

[0069] Each time when receiving the quantity of messages and the message time stamp that are output by the flow classification module, the flow recognizing module calculates the bandwidth of the corresponding information flow according to a ratio of the received quantity of messages to a message time stamp difference, and when the bandwidth is greater than a set elephant flow threshold, the information flow is recognized as a potential elephant flow, and flow table information such as the quantity of messages and time stamp corresponding to the potential information flow cached in the on-chip memory is output to the off-chip memory to be accumulated.

[0070] As time goes by, there are increasingly more flow table information corresponding to the potential flow in the off-chip memory, and the flow screening module may read the flow table information corresponding to the potential elephant flow stored in the off-chip memory at intervals, and calculate the bandwidth of the potential flow. If the bandwidth of the potential flow exceeds the set elephant flow bandwidth threshold, the potential flow is considered as an elephant flow.

[0071] After recognizing the elephant flow, the flow screening module may provide identification information of the elephant flow to the processor, and the processor may load the flow table information of the elephant flow corresponding to the

identification information from the off-chip memory to the on-chip memory based on the identification information, and monitor a utilization rate of the processor. When the utilization rate of the processor exceeds 90%, the elephant flow is limited to prevent network congestion.

[0072] **FIG. 3** is a schematic structural diagram of a cloud computing system according to an example embodiment of the present application. As shown in **FIG. 3**, the cloud computing system 300 includes: a central management and control device 301, at least one edge cloud node 302, and a terminal 303. Each edge cloud node 302 includes a series of edge infrastructures, including but not limited to: a distributed data center (DC), a wireless computer room or cluster, a communication network of an operator, a core network device, a base station, an edge gateway, a home gateway, a computing device, and/or storage device, and other edge devices and corresponding network environments, and the like. It should be noted here that locations, capabilities, and included infrastructures of different edge cloud nodes 302 may be the same or different. For the edge cloud node 302, it may provide various resources externally, for example, computing resources such as CPU and GPU, storage resources such as memory and hard disk, network resources such as bandwidth, and the like.

[0073] In the present embodiment, the cloud computing system is used as an example to implement load balancing, flow control, and the like, and there may be migration of virtual machines, cloud computing services, or data. In the entire process, elephant flows may be generated. It is found by research that in a queue at an export terminal of a network switch, a mice flow often follows an elephant flow, which causes a long term queue delay. In order to optimize or ensure the performance and improve the delay, it is very important to recognize an elephant flow.

[0074] In addition, a tenant on the cloud may submit a request for deploying a cloud computing service to the central management and control device 301 through a human-computer interaction interface provided by the central management and control device 301. The cloud computing service may be a short video service, an online education service, an online live streaming service, a game service, or the like, and the central management and control device 301 selects a suitable edge cloud node 302 for the tenant cloud, deploys a cloud computing service for the cloud tenant on the edge cloud node 302. The terminal 303 may use the cloud computing service deployed on the nearest edge cloud node 302. In **FIG. 3**, for example, the cloud tenant is an online education service provider, and an online education service for a plurality of subjects such as Chinese, mathematics, and English is deployed on the edge cloud node 302, the online education service provider hopes to focus on the traffic bandwidth of the English subject.

[0075] Specifically, if the cloud tenant hopes to acquire the traffic bandwidth of the English subject, the cloud tenant may send a request of acquiring the traffic bandwidth of the English subject to the central management and control device 301, and the central management and control device 301 may deliver, according to the deployment of the English subject on the edge cloud node 302, feature information of an information flow of the English subject to a server deployed with the online education service of the English subject in the edge cloud node 302. The server includes a network chip, and the feature information of the information flow of the English subject is specifically delivered to the network chip. According to the request of the terminal user, the online education service of the English subject provides the terminal user with teaching content of the English subject through the network chip. Specifically, the network chip receives a message provided by the online education service of the English subject, and recognizes, based on the feature information of the information flow corresponding to the English subject, whether the

message belongs to the information flow corresponding to the English subject, and if so, adds label information Y to the message, as shown in **FIG. 4**. The label information Y may uniquely label the information flow corresponding to the English subject. Specifically, when receiving the message, the flow labeling module may extract feature information from the message, and compare the feature information extracted from the message with the feature information of the message in the specified flow stored locally; if the two are the same, determine that the message belongs to the specified flow, add label information in the message, and send the message after adding the label information into the on-chip memory for caching; and if the two are different, directly send the message to into the on-chip memory for caching.

[0076] As shown in **FIG. 4**, each time when a new message is cached in the on-chip memory, the flow classification module hashes feature information of the new message by using a hash function to obtain a hash result X of the new message; and determines whether there is a target entry corresponding to the hash result X in a hash table. If there is a target entry corresponding to the hash result X in the hash table, the new message is mapped to the target entry, and the quantity of messages in the flow table information in the target entry is increased by 1. If there is no target entry corresponding to the hash result X in the hash table, it is determined whether there is a spare entry in the hash table. If there is a spare entry in the hash table, an idle entry is acquired from the hash table, the new message is mapped to the idle entry, flow table information is inserted into the idle entry, and the quantity of messages is set to 1. If there is no idle entry in the hash table, flow table information recorded in an entry with the largest quantity of messages in the hash table is output to the flow recognizing module to obtain a blank entry. The flow table information includes: the quantity of messages, a time stamp of the message, and whether the message carries label information. The flow table information of the

information flow corresponding to the English subject includes the label information Y, and flow table information of an information flow corresponding to a non-English subject does not include the label information Y.

[0077] As shown in **FIG. 4**, each time when receiving the flow table information output by the flow classification module, the flow recognizing module calculates the bandwidth of the corresponding information flow according to the ratio of the quantity of messages in the received flow table information to the message time stamp difference, and determines whether the calculated bandwidth is greater than the set elephant flow bandwidth threshold. If the bandwidth is greater than the set elephant flow bandwidth threshold, the information flow is recognized as a potential elephant flow, and flow table information corresponding to the potential elephant flow cached in the on-chip memory is output to the off-chip memory to be accumulated. If the bandwidth is less than or equal to the set elephant flow bandwidth threshold, it is further determined whether the flow table information includes label information, and if the label information is included, an information flow corresponding to the flow table information is considered as a specified flow of the English subject. If the flow table information does not include the label information, the flow table information corresponding to the flow table information in the on-chip memory is deleted.

[0078] As time goes by, there are increasingly more flow table information corresponding to the potential flow in the off-chip memory, and the flow screening module may read the flow table information corresponding to the potential flow stored in the off-chip memory at intervals, and recognize a true elephant flow and a specified flow therefrom. After recognizing the true elephant flow and the specified flow, the flow screening module may provide identification information of the elephant flow and the specified flow to the processor. On the one hand, based on the received identification

information, the processor may load the flow table information of the elephant flow corresponding to the identification information from the off-chip memory to the on-chip memory, and monitor a utilization rate of the processor. When the utilization rate of the processor exceeds 90%, the elephant flow is limited to prevent network congestion. On the other hand, the processor may also perform accurate statistics on the traffic bandwidth of the specified English subject. Further, the processor provides the traffic bandwidth of the English subject to the central management and control device through the network connection between the edge cloud node and the central management and control device. The central management and control device summarizes the traffic bandwidth of a plurality of edge cloud nodes, and returns a summary result to the online education service provider.

[0079] In the present embodiment, in addition to providing a network chip, an information flow recognition method is further provided, the method is applicable to a main control chip in a network chip, as shown in **FIG. 5**, and the method includes:

[0080] S501: Cache a message entering the network chip into an on-chip memory of the network chip, and map the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow;

[0081] S502: Recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information;

[0082] S503: Output the flow table information corresponding to the potential flow to an off-chip memory of the network chip to be accumulated; and

[0083] S504: Read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.

[0084] In an optional embodiment, the method provided in the present embodiment further includes: performing accurate statistics on the target flow according to the identification information of the target flow when the target flow is recognized, or performing flow control on the target flow, or loading the flow table information of the target flow from the off-chip memory to the on-chip memory.

[0085] In an optional embodiment, mapping the message in the on-chip memory to at least one piece of flow table information includes: mapping the message in the on-chip memory to a hash table in the on-chip memory by using a hash function, wherein one entry in the hash table corresponds to one information flow, and the entry records corresponding flow table information.

[0086] In an optional embodiment, mapping the message in the on-chip memory to a hash table in the on-chip memory by using a hash function includes: hashing, each time when a new message is cached in the on-chip memory, feature information of the new message by using the hash function to obtain a hash result of the new message; if there is a target entry corresponding to the hash result in the hash table, mapping the new message to the target entry and updating flow table information in the target entry; and if there is no target entry corresponding to the hash result in the hash table, acquiring an idle entry from the hash table, mapping the new message to the idle entry, and recording flow table information of an information flow to which the new message belongs in the idle entry.

[0087] In an optional embodiment, the method provided in the present embodiment further includes: if there is no idle entry in the hash table, acquiring flow table information in an entry meeting an output condition in the hash table, and deleting the flow table information after outputting the entry information to obtain a blank entry.

[0088] In an optional embodiment, acquiring the flow table information in the entry meeting the output condition in the hash table includes: acquiring flow table information in an entry with the largest quantity of messages in the hash table.

[0089] In an optional embodiment, recognizing the potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information includes: calculating the bandwidth of an information flow corresponding to the acquired flow table information according to the acquired flow table information bandwidth, and recognizing the information flow as a potential flow if the bandwidth is greater than a set threshold.

[0090] In an optional embodiment, before the message entering the network chip is cached into the on-chip memory, it is determined, according to the feature information of the message, whether the message is a message in the specified flow, and if so, label information is added to the message. Correspondingly, recognizing the potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information includes: determining whether at least one piece of flow table information includes label information, and recognizing the information flow corresponding to the flow table information including the label information as a potential flow.

[0091] In the embodiments of the present application, a combination of an on-chip memory and an off-chip memory is used for recognizing an information flow. First, a message entering a network chip is cached in the on-chip memory, the cached message is mapped to at least one piece of flow table information in the chip, and a potential flow that is a possible target flow in the information flow is recognized, then the flow table information corresponding to the potential flow is cached to the off-chip memory to be accumulated, until a set trigger condition is met, the flow table information corresponding

to the potential flow cached in the off-chip memory is read to accurately recognize the target flow in the potential flow accordingly. Firstly, a potential flow is recognized by taking advantage of a large access bandwidth of the on-chip memory, which can increase the speed of recognizing the potential flow. Secondly, only the flow table information of the potential flow is accumulated by taking advantage of a large storage space of the off-chip memory, and the potential flow is further recognized, rather than further recognizing all information flows by the off-chip memory, which can reduce the number of accesses to the off-chip memory and help improve the efficiency and accuracy of recognizing a target flow.

[0092] It should be noted that, the steps of the method provided in the above embodiment may all be performed by the same device, or the method may also be performed by different devices. For example, step S501 to step S503 may be performed by a device A, for another example, steps S501 and S502 may be performed by a device A, and step S503 may be performed by a device B; or the like.

[0093] In addition, in some of the processes described in the above embodiments and accompanying drawings, a plurality of operations appearing in a specific order are included, but it should be clearly understood that these operations may not be performed in the order herein or performed in parallel, serial numbers of the operations, such as S501 and S502, are only used for distinguishing different operations, and the serial numbers themselves do not represent any performing order. Additionally, these processes may include more or fewer operations, and these operations may be performed sequentially or in parallel. It should be noted that the descriptions of "first," "second," and the like in herein are used for distinguishing different messages, devices, modules, and the like, which do not represent the order of priority, nor limit that "first" and "second" are different types.

[0094] **FIG. 6** is a schematic structural diagram of a network device according to an example embodiment of the present application. As shown in **FIG. 6**, the network device includes: a network chip 601, a memory 602, and a first processor 603. The network chip 601 includes a main control chip 601a and an off-chip memory 601b, and the main control chip 601a includes an on-chip memory 6001, a flow classification module 6002, a flow recognizing module 6003, and a flow screening module 6004. Further optionally, as shown in **FIG. 6**, the main control chip 601a further includes: a second processor 6005 and a flow labeling module 6006. The first processor 603 refers to a processor of the network device, which may be regarded as a device processor. The second processor 6005 refers to a processor on the main control chip 601a in the network chip 601, which may be regarded as an on-chip memory.

[0095] The memory 602 is configured to store a computer program, and may be configured to store other various types of data to support operations on network devices. Examples of such data include instructions for any application or method operated on a network device.

[0096] The memory 602 may be realized by any type of volatile or non-volatile storage device or a combination thereof, such as a static random access memory (SRAM), an electrically erasable programmable read-only memory (EEPROM), an erasable programmable read-only memory (EPROM), a programmable read-only memory (PROM), a read-only memory (ROM), a magnetic memory, a flash memory, a magnetic disk, or an optical disc.

[0097] The first processor 603 is coupled to the memory 602 and is configured to execute a computer program in the memory 602, so as to: send a message to the network chip 601 and/or receive information of a target flow reported by the network chip 601. The on-chip memory 6001 is configured to cache a message entering the network chip; the flow

classification module 6002 is configured to map the message in the on-chip memory to at least one piece of flow table information, and provide at least one piece of flow table information to the flow recognizing module; the flow recognizing module 6003 is configured to recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and output the flow table information corresponding to the potential flow cached in the on-chip memory to the off-chip memory to be accumulated; and the flow screening module 6004 is configured to read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.

[0098] In an optional embodiment, the flow screening module 6004 is further configured to: report identification information of the target flow to the second processor 6005 when the target flow is recognized; and the second processor 6005 is configured to perform accurate statistics on the target flow according to the identification information of the target flow, or perform flow control on the target flow, or load flow table information of the target flow from the off-chip memory to the on-chip memory.

[0099] In an optional embodiment, the flow classification module 6002, when mapping the message in the on-chip memory to at least one piece of flow table information, is specifically configured to map the message in the on-chip memory to a hash table cached in the on-chip memory by using a hash function, wherein one entry in the hash table corresponds to one information flow, and the entry records flow table information of the corresponding information flow.

[00100] In an optional embodiment, the flow classification module 6002, when mapping the message in the on-chip memory to a hash table in the on-chip memory by using a hash function, is specifically configured to hash, each time when a new message is

cached in the on-chip memory, feature information of the new message by using the hash function to obtain a hash result of the new message; if there is a target entry corresponding to the hash result in the hash table, map the new message to the target entry and update flow table information in the target entry; and if there is no target entry corresponding to the hash result in the hash table, acquire an idle entry from the hash table, map the new message to the idle entry, and record flow table information of an information flow to which the new message belongs in the idle entry.

[00101] In an optional embodiment, the flow classification module 6002, when outputting at least one piece of flow table information to the flow recognizing module, is specifically configured to output, if there is no idle entry in the hash table, flow table information in an entry meeting an output condition in the hash table to the flow recognizing module 6003 to obtain a blank entry.

[00102] In an optional embodiment, the flow classification module 6002, when outputting the flow table information in the entry meeting the output condition in the hash table to the flow recognizing module, is specifically configured to output the flow table information in an entry with the largest quantity of messages in the hash table to the flow recognizing module.

[00103] In an optional embodiment, the flow recognizing module 6003, when recognizing a potential flow that is a possible target flow in the information flow according to at least one piece of flow table information, is specifically configured to calculate, each time when the flow table information output by the flow classification module is received, a bandwidth of an information flow corresponding to the received flow table information according to the received flow table information, and recognize the information flow as a potential flow when the bandwidth is greater than a set threshold.

[00104] In an optional embodiment, the flow labeling module 6006 is configured to determine, before caching the message entering the network chip into the on-chip memory, whether the message is a message in a specified flow according to the feature information of the message, and if so, add label information to the message, wherein the label information is a type of information in the flow table information. The flow recognizing module 6003, when recognizing a potential flow that is a possible target flow in the information flow according to at least one piece of flow table information, is specifically configured to determine whether at least one piece of flow table information includes label information, and recognize the information flow including the label information in the flow table information as a potential flow.

[00105] In an optional embodiment, the network chip 601 is a network card.

[00106] In the network device provided in the embodiments of the present application, a combination of an on-chip memory and an off-chip memory is used for recognizing an information flow. First, a message entering a network chip is cached in the on-chip memory, the message is mapped to at least one piece of flow table information in the chip, and a potential flow that is a possible target flow in the information flow is recognized accordingly, then the flow table information corresponding to the potential flow is cached to the off-chip memory to be accumulated, until a set trigger condition is met, the flow table information of the potential flow accumulated in the off-chip memory is read to accurately recognize the target flow in the potential flow accordingly. Firstly, a potential flow is recognized by taking advantage of a large access bandwidth of the on-chip memory, which can increase the speed of recognizing the potential flow. Secondly, only the potential flow is further recognized by taking advantage of a large storage space of the off-chip memory, rather than further recognizing all information flows by the off-chip

memory, which can reduce the number of accesses to the off-chip memory and help improve the efficiency and accuracy of recognizing a target flow.

[00107] Further, as shown in **FIG. 6**, the network device further includes: a communication component 606, a display 607, a power supply component 608, an audio component 609, and other components. **FIG. 6** only schematically shows some components, which does not mean that the network device only includes the components shown in **FIG. 6**. It should be noted that a component in a dotted line box in **FIG. 6** is an optional component rather than a mandatory component, specifically depending on a product form of the network device.

[00108] The above communication component in **FIG. 6** is configured to facilitate wired or wireless communication between a device where the communication component is located and other devices. The device where the communication component is located may access a wireless network based on a communication standard, such as WiFi, 2G, 3G, 4G/LTE, 5G, and other mobile communication networks, or a combination thereof. In an example embodiment, the communication component receives a broadcast signal or broadcast related information from an external broadcast management system via a broadcast channel. In an example embodiment, the communication component further includes a near field communication (NFC) module to facilitate short-range communication. For example, the NFC module may be implemented based on the radio frequency identification (RFID) technology, infrared data association (IrDA) technology, ultra-wideband (UWB) technology, Bluetooth (BT) technology, and other technologies.

[00109] The above display in **FIG. 6** includes a screen, and the screen may include a liquid crystal display (LCD) and a touch panel (TP). If the screen includes a touch panel, the screen may be implemented as a touch screen to receive an input signal from a user. The touch panel includes one or more touch sensors to sense touches, swipes, and gestures

on the touch panel. The touch sensor may not only sense a boundary of a touch or swipe action, but also detect a duration and pressure associated with the touch or swipe operation.

[00110] The above power supply component in **FIG. 6** provides power for various components of a device where the power supply component is located. The power supply component may include a power management system, one or more power supplies, and other components associated with generating, managing, and distributing power for the device where the power supply component is located.

[00111] The above audio component in **FIG. 6** may be configured to output and/or input an audio signal. For example, the audio component includes a microphone (MIC), and the microphone is configured to receive an external audio signal when the device where the audio component is located is in an operation mode, such as a call mode, a recording mode, and a speech recognition mode. The received audio signal may be further stored in the memory or sent via the communication component. In some embodiments, the audio component further includes a speaker for outputting the audio signal.

[00112] Persons skilled in the art should understand that, the embodiments of the present application may be provided as a method, an apparatus, or a computer program product. Therefore, the present application may be implemented as a completely hardware embodiment, a completely software embodiment, or an embodiment combining software and hardware. Moreover, the present application may be in the form of a computer program product implemented on one or more computer usable storage media (including, but not limited to, a magnetic disk memory, a CD-ROM, an optical memory and the like) including computer usable program codes.

[00113] The present application is described with reference to flow charts and/or block diagrams according to the method, device (system), and computer program product according to the embodiments of the present application. It should be understood that a

computer program instruction may be used to implement each process and/or block in the flow charts and/or block diagrams and combinations of processes and/or blocks in the flow charts and/or block diagrams. These computer program instructions can be provided to a processor of a general-purpose computer, a special-purpose computer, an embedded processor, or another programmable data processing device to produce a machine, so that the instructions executed by the processor of the computer or another programmable data processing device produce an apparatus for realizing the functions specified in one or more flows of the flowcharts and/or one or more blocks of the block diagrams.

[00114] These computer program instructions may also be stored in a computer-readable memory that can direct a computer or another programmable data processing device to work in a specific manner, so that the instructions stored in this computer-readable memory produce an article of manufacture including an instruction apparatus which implements the functions specified in one or more flows of the flowcharts and/or one or more blocks of the block diagrams.

[00115] These computer program instructions can also be loaded onto a computer or another programmable data processing device, so that a series of operation steps are performed on the computer or another programmable device to produce computer-implemented processing, so that the instructions executed on the computer or another programmable device provide steps for implementing the functions specified in one or more flows of the flowcharts and/or one or more blocks of the block diagrams.

[00116] In a typical configuration, the computing device includes one or more central processing units (CPUs), an input/output interface, a network interface, and a memory.

[00117] The memory may include a volatile memory, a random access memory (RAM), and/or a non-volatile memory in computer-readable media, e.g., a read-only

memory (ROM) or a flash RAM. The memory is an example of the computer readable medium.

[00118] The computer-readable medium includes permanent and non-permanent, removable and non-removable media, which may implement storage of information by using any method or technology. Information may be a computer readable instruction, a data structure, and a module of a program or other data. Examples of computer storage media include, but are not limited to, a phase change memory (PRAM), a static random access memory (SRAM), a dynamic random access memory (DRAM), other type of random access memory (RAM), a read-only memory (ROM), an electrically erasable programmable read-only memory (EEPROM), a flash memory or other memory technology, a compact disc read-only memory (CD-ROM), a digital versatile disc (DVD) or other optical storage, a cassette magnetic tape, tape and disk storage or other magnetic storage device, or any other non-transmission media, which may be configured to store information accessible by a computing device. As defined herein, the computer-readable medium does not include transitory computer readable media (transitory media), such as a modulated data signal and a carrier.

[00119] It should be further noted that the terms "include," "comprise," or any other variations thereof are intended to cover non-exclusive inclusion, so that a process, method, commodity, or device including a series of elements not only includes the elements, but also includes other elements not explicitly listed, or further includes elements inherent to the process, method, commodity, or device. In a case without any more limitations, an element defined by "including a..." does not exclude that the process, method, commodity or device including the element further has other identical elements.

[00120] The above descriptions are merely embodiments of the present application, and are not intended to limit the present application. For those skilled in the art, the present

application can have various alterations and changes. Any modification, equivalent replacement, improvement or the like made without departing from the spirit and principle of the present application should all fall within the scope of claims of the present application.

WHAT IS CLAIMED IS:

[Claim 1] A network chip, comprising: a main control chip and an off-chip memory, wherein the main control chip comprises: an on-chip memory, a flow classification module, a flow recognizing module, and a flow screening module;

the on-chip memory is configured to cache a message entering the network chip;

the flow classification module is configured to map the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow;

the flow recognizing module is configured to recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and output the flow table information corresponding to the potential flow to the off-chip memory to be accumulated; and

the flow screening module is configured to read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.

[Claim 2] The network chip according to claim 1, wherein the main control chip is further provided with a processor;

the flow screening module is further configured to report identification information of the target flow to the processor when the target flow is recognized; and

the processor is configured to perform accurate statistics on the target flow according to the identification information of the target flow, or perform flow control on the target flow, or load the flow table information of the target flow from the off-chip memory to the on-chip memory.

[Claim 3] The network chip according to claim 1, wherein the flow classification

module is specifically configured to:

map the message in the on-chip memory to a hash table in the on-chip memory by using a hash function, wherein an entry in the hash table records the flow table information of the corresponding information flow.

[Claim 4] The network chip according to claim 3, wherein the flow classification module is specifically configured to:

perform, each time when a new message is cached in the on-chip memory, hash processing on feature information of the new message by using a hash function to obtain a hash result of the new message;

map, if there is a target entry corresponding to the hash result in the hash table, the new message to the target entry, and update flow table information in the target entry;

acquire, if there is no target entry corresponding to the hash result in the hash table, an idle entry from the hash table, map the new message to the idle entry, and record the flow table information of the information flow to which the new message belongs in the idle entry.

[Claim 5] The network chip according to claim 4, wherein the flow classification module is specifically configured to:

output, if there is no idle entry in the hash table, flow table information in an entry meeting an output condition in the hash table to the flow recognizing module to obtain a blank entry.

[Claim 6] The network chip according to claim 5, wherein the flow classification module is specifically configured to:

output flow table information in an entry corresponding to the largest quantity of messages in the hash table to the flow recognizing module.

[Claim 7] The network chip according to claim 5, wherein the flow recognizing module is specifically configured to:

calculate, each time when receiving the flow table information output by the flow classification module, a bandwidth of an information flow corresponding to the received flow table information according to the received flow table information, and recognize the information flow as a potential flow when the bandwidth is greater than a set threshold.

[Claim 8] The network chip according to any one of claims 1 to 7, wherein the main control chip is further provided with a flow labeling module;

the flow labeling module is configured to determine, before caching the message entering the network chip into the on-chip memory, whether the message is a message in a specified flow according to feature information of the message, and if so, add label information to the message, wherein the label information is a type of information in the flow table information; and

the flow recognizing module is specifically configured to determine whether the at least one piece of flow table information comprises the label information, and recognize an information flow corresponding to the flow table information comprising the label information as a potential flow.

[Claim 9] The network chip according to any one of claims 1 to 7, wherein the network chip is a network card.

[Claim 10] An information flow recognition method, applicable to a main control chip in a network chip, the method comprising:

caching a message entering the network chip into an on-chip memory of the network chip, and mapping the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow;

recognizing a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information;

outputting the flow table information corresponding to the potential flow to an off-chip memory of the network chip to be accumulated; and

reading the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognizing a target flow in the potential flow based on the read flow table information.

[Claim 11] A network device, comprising: a memory, a processor, and a network chip;

the memory is configured to store a computer program; the processor is coupled to the memory and configured to execute the computer program, for sending a message to the network chip and/or receiving identification information of a target flow reported by the network chip;

the network chip comprises: a main control chip and an off-chip memory, wherein the main control chip comprises: an on-chip memory, a flow classification module, a flow recognizing module, and a flow screening module;

the on-chip memory is configured to cache a message entering the network chip;

the flow classification module is configured to map the message in the on-chip memory to at least one piece of flow table information, wherein one piece of flow table information corresponds to one information flow;

the flow recognizing module is configured to recognize a potential flow that is a possible target flow in the information flow according to the at least one piece of flow table information, and output the flow table information corresponding to the potential flow to the off-chip memory to be accumulated; and

the flow screening module is configured to read the flow table information corresponding to the potential flow accumulated in the off-chip memory when a set trigger condition is met, and recognize a target flow in the potential flow based on the read flow table information.