



(12) **DEMANDE DE BREVET CANADIEN  
CANADIAN PATENT APPLICATION**  
(13) **A1**

(86) **Date de dépôt PCT/PCT Filing Date:** 2022/11/24  
(87) **Date publication PCT/PCT Publication Date:** 2023/06/01  
(85) **Entrée phase nationale/National Entry:** 2024/05/15  
(86) **N° demande PCT/PCT Application No.:** CN 2022/134035  
(87) **N° publication PCT/PCT Publication No.:** 2023/093805  
(30) **Priorité/Priority:** 2021/11/26 (CN202111421847.3)

(51) **Cl.Int./Int.Cl. G06F 9/54** (2006.01)  
(71) **Demandeur/Applicant:**  
ZTE CORPORATION, CN  
(72) **Inventeurs/Inventors:**  
DAI, SHUZHOU, CN;  
YAN, LIN, CN;  
LIU, QIANGJUN, CN;  
WANG, JUN, CN;  
LIAO, ZHIJIA, CN;  
YU, FENG, CN  
(74) **Agent:** FASKEN MARTINEAU DUMOULIN LLP

(54) **Titre : PROCEDURE DE COMMANDE DE STOCKAGE, CONTROLEUR DE STOCKAGE, PUCE DE STOCKAGE, CARTE RESEAU ET SUPPORT LISIBLE**  
(54) **Title: STORAGE CONTROL METHOD, STORAGE CONTROLLER, STORAGE CHIP, NETWORK CARD, AND READABLE MEDIUM**

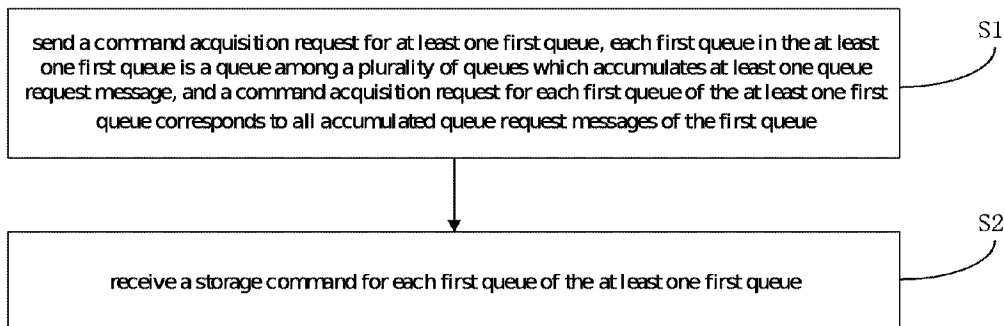


FIG. 1

(57) **Abrégé/Abstract:**

Provided in the present disclosure is a storage control method. The method comprises: sending a command acquisition request for at least one first queue, wherein the at least one first queue is a queue, among a plurality of queues, for which at least one queue request message has accumulated, and a command acquisition request for each of the at least one first queue corresponds to all of the accumulated queue request messages for the first queue; and receiving a storage command for each of the at least one first queue.

## **ABSTRACT**

The present disclosure provides a storage control method, including: sending a command acquisition request for at least one first queue, where the at least one first queue is a queue among a plurality of queues which accumulates at least one queue request message, and a command acquisition request for each first queue of the at least one first queue corresponds to all accumulated queue request messages of the first queue; and receiving a storage command for each first queue in the at least one first queue.



## SUMMARY

[0003] Embodiments of the present disclosure provide a storage control method, a storage controller, a storage chip, a network card, and a computer-readable medium.

5 [0004] In a first aspect, an embodiment of the present disclosure provides a storage control method, including: sending a command acquisition request for at least one first queue, where the at least one first queue is a queue among a plurality of queues which accumulates at least one queue request message, and a command acquisition request for each first queue of the at least one first queue corresponds to all accumulated queue request messages of the first queue; and receiving a storage command for each first  
10 queue in the at least one first queue.

[0005] In some embodiments, before sending the command acquisition request for the at least one first queue, the storage control method further includes: determining a queue corresponding to a flag register which is enabled in a first flag register table of at least one flag register table as the at least one first queue, wherein each of the at least  
15 one flag register table comprises a plurality of flag registers respectively corresponding to the plurality of queues, and the flag register which is enabled represents that the corresponding queue has accumulated the at least one queue request message.

[0006] In some embodiments, the storage control method further includes: in response  
20 to a queue request message for a second queue, enabling a flag register corresponding to the second queue in a second flag register table of the at least one flag register table, wherein the second queue is a queue corresponding to a flag register which is disabled in the first flag register table.

[0007] In some embodiments, the storage control method further includes: in response to  
25 a queue request message for any first queue of the at least one first queue, determining whether a command acquisition request for the any first queue is sent; and when the command acquisition request for the any first queue is sent, enabling a flag register corresponding to the any first queue in the second flag register table.

[0008] In some embodiments, after sending the command acquisition request for the at

least one first queue, the storage control method further includes: for any first queue of the at least one first queue, after a command acquisition request for the any first queue is sent, clearing a flag register corresponding to the any first queue in the first flag register table.

5 [0009] In some embodiments, the storage control method further includes: in a case where all flag registers in the first flag register table are cleared, taking a flag register table used as the second flag register table as the first flag register table.

10 [0010] In some embodiments, sending the command acquisition request for the at least one first queue includes: for any first queue of the at least one first queue, generating a command acquisition request for the any first queue according to an accumulated queue request message of the any first queue, wherein a number of storage commands corresponding to the command acquisition request for the any first queue is equal to a sum of numbers of storage commands corresponding to all accumulated queue request messages of the any first queue; and sending the command acquisition request for the  
15 any first queue.

[0011] In a second aspect, an embodiment of the present disclosure provides a storage controller, including: one or more processing units; and a storage unit having stored thereon one or more programs which, when executed by the one or more processing units, cause the one or more processing units to implement the storage control method  
20 described in the first aspect of the present disclosure.

[0012] In a third aspect, an embodiment of the present disclosure provides a storage chip, including: the storage controller described in the second aspect of the present disclosure; and a memory.

25 [0013] In a fourth aspect, an embodiment of the present disclosure provides a network card, including: a main control chip and a storage chip, wherein the storage chip includes the storage controller described in the second aspect of the present disclosure.

[0014] In a fifth aspect, an embodiment of the present disclosure provides a computer-readable medium storing a computer program which, when executed by a processor, implements the storage control method described in the first aspect of the

present disclosure.

## BRIEF DESCRIPTION OF DRAWINGS

5 [0015] FIG. 1 is a flowchart of a storage control method according to an embodiment of the present disclosure;

[0016] FIG. 2 is a schematic diagram of a queue according to an embodiment of the present disclosure;

[0017] FIG. 3 is a flowchart of a storage control method according to an embodiment of the present disclosure;

10 [0018] FIG. 4 is a schematic diagram of queue management according to an embodiment of the present disclosure;

[0019] FIG. 5 is a flowchart of a storage control method according to an embodiment of the present disclosure;

15 [0020] FIG. 6 is a flowchart of a storage control method according to an embodiment of the present disclosure;

[0021] FIG. 7 is a flowchart of a storage control method according to an embodiment of the present disclosure;

[0022] FIG. 8 is a flowchart of a storage control method according to an embodiment of the present disclosure;

20 [0023] FIG. 9 is a flowchart of some operations in a storage control method according to an embodiment of the present disclosure;

[0024] FIG. 10 is a constituent block diagram of a storage controller according to an embodiment of the present disclosure;

25 [0025] FIG. 11 is a constituent block diagram of a storage chip according to an embodiment of the present disclosure;

[0026] FIG. 12 is a constituent block diagram of a network card according to an embodiment of the present disclosure;

[0027] FIG. 13 is a constituent block diagram of a computer-readable medium according to an embodiment of the present disclosure;

5 [0028] FIG. 14 is a flowchart illustrating management of a plurality of queues according to an example of the present disclosure.

[0029] FIG. 15 is an architecture diagram of management of a plurality of queues according to an example of the present disclosure; and

10 [0030] FIG. 16 is an architecture diagram of management of a plurality of queues according to an example of the present disclosure.

#### DETAIL DESCRIPTION OF EMBODIMENTS

15 [0031] In order to enable those of ordinary skill in the art to better understand the technical solutions of the present disclosure, a storage control method, a storage controller, a storage chip, a network card, and a computer-readable medium provided by the present disclosure are described in detail below with reference to the accompanying drawings.

20 [0032] Exemplary embodiments will be described more fully below with reference to the accompanying drawings, but the exemplary embodiments illustrated may be implemented in different forms, and should not be interpreted as being limited to the embodiments described below. Rather, the embodiments are provided to make the present disclosure thorough and complete, and are intended to enable those of ordinary skill in the art to fully understand the scope of the present disclosure.

25 [0033] All the embodiments of the present disclosure and the features therein may be combined with each other if no conflict is incurred.

[0034] The term “and/or” used herein includes any and all combinations of one or more associated listed items.

[0035] The terms used herein are merely used to describe particular embodiments, and are not intended to limit the present disclosure. The terms “a” and “the” used herein which indicate a singular form are intended to include a plural form, unless expressly stated in the context. It should be further understood that the terms “include” and/or “be made of” used herein indicate the presence of the described features, integers, steps, operations, elements and/or components, but do not exclude the presence or addition of one or more other features, integers, steps, operations, elements, components and/or combinations thereof.

[0036] Unless otherwise defined, all the terms (including technical and scientific terms) used herein have the same meaning as commonly understood by those of ordinary skill in the art. It should be further understood that terms, such as those defined in commonly used dictionaries, should be interpreted as having a meaning that is consistent with a meaning in the context of the related technology and the background of the present disclosure, and should not be interpreted in an idealized or overly formal sense unless expressly so defined herein.

[0037] In an NVME-based storage system, a process of storage interaction between a host and a storage controller includes: the host sends a management message to the storage controller, and the storage controller creates a plurality of queues according to the management message; after generating a storage command, the host sends a queue request message to the storage controller; the storage controller responds to the queue request message and sends a command acquisition request to the host; the host transmits the storage command to the storage controller according to the command acquisition request; and the storage controller performs subsequent processing according to the storage command.

[0038] With development of the technology, a number of the queues between the host and the storage controller keeps growing. For example, in the field of the virtualization technology, a plurality of virtualized storage devices may be constructed and presented between the host and a user, and each virtualized storage device may be corresponding to one Virtual Function (VF) in the host. Each VF may be bound to a corresponding storage controller, and performs storage interaction with the bound storage controller through a plurality of queues, and the plurality of VFs may correspond to a large number

of queues. Efficiency of acquiring storage commands from the host by the storage controller becomes an important factor which restricts storage performance of the storage system, and management of the large number of queues is an important factor which affects the efficiency of acquiring the storage commands from the host by the storage controller.

5

[0039] In some related technologies, for each queue request message sent by the host, the storage controller separately makes a respond and sends a command acquisition request. For a plurality of queue request messages, the storage controller sequentially and serially respond to queue request messages according to an order in which the queue request messages arrive at the storage controller.

10

[0040] It is found by research that when there are the large number of queues, queue request messages for different queues are frequently sent from the host, and a large number of command acquisition requests are usually generated during a process from responding to the queue request messages to initiating acquisition of storage commands from the host by the storage controller. A Direct Memory Access (DMA) component and a data bus serially process the command acquisition requests sent from the storage controller to the host, which may easily cause blockage and timeout of the command acquisition requests of a storage control system. Moreover, since the storage controller sends one command acquisition request for each queue request message, the DMA component initiates one command data transfer for each command acquisition request, and a number of storage commands corresponding to each queue request message is usually small, a number of storage commands acquired from the host in unit time is relatively small, and efficiency of acquiring the storage commands from the host by the storage controller is very low.

15

20

25

[0041] FIG. 1 is a flowchart of a storage control method according to an embodiment of the present disclosure, and the storage control method includes operations S1 and S2.

[0042] At operation S1, a command acquisition request for at least one first queue is sent, the at least one first queue is a queue among a plurality of queues which accumulates at least one queue request message, and a command acquisition request for each first queue of the at least one first queue corresponds to all accumulated queue

30

request messages of the first queue.

[0043] At operation S2, a storage command for each first queue of the at least one first queue is received.

[0044] In the embodiments of the present disclosure, a storage controller accumulates queue request messages for the plurality of queues through an accumulation process before operation S1 is performed. It should be noted that an accumulated queue request message of a queue refers to a queue request message which is sent for the queue by the host and received by the storage controller during the accumulation process. For example, for any queue, when the storage controller receives one or more queue request messages for the queue during the accumulation process, the one or more queue request messages are accumulated queue request messages of the queue.

[0045] Accordingly, in the embodiments of the present disclosure, when receiving a queue request message sent by the host for any queue, the storage controller does not directly respond to the queue request message for the queue, but performs operation S1 to, for any first queue, uniformly respond to the at least one accumulated queue request message of the first queue with one command acquisition request. Therefore, in the embodiments of the present disclosure, one command acquisition request may correspond to one queue request message for one queue, or may correspond to a plurality of queue request messages for one queue. When operation S1 is performed, a command acquisition request for each of a plurality of first queues may correspond to a plurality of queue request messages, or a command acquisition request for each of part of the first queues may correspond to one queue request message and a command acquisition request for each of the remaining first queues may correspond to a plurality of queue request messages. Compared with the method in which one command acquisition request is sent for each queue request message, the method according to the embodiments of the present disclosure can reduce the number of command acquisition requests, and further reduce a number of times of transferring of the storage commands from the host by the DMA component.

[0046] In the embodiments of the present disclosure, when a command acquisition request for one queue corresponds to a plurality of queue request messages of the queue,

the host may send storage commands corresponding to all the queue request messages of the queue to the storage controller according to the command acquisition request, which can increase the number of the storage commands acquired from the host by the DMA component in one-time transfer.

5 [0047] It should be further noted that, in the embodiments of the present disclosure, the storage controller does not respond to the queue request messages in strict accordance with an order of receiving the queue request messages. By accumulating the queue request messages and then sending the command acquisition request to respond to the queue request messages by performing operation S1, all the received queue request  
10 messages can be responded to, thereby avoiding response omission and response delay.

[0048] In the storage control method provided by the embodiments of the present disclosure, the storage controller first accumulates the received queue request messages and uniformly respond to the accumulated at least one queue request messages; and for any queue, the storage commands corresponding to a plurality of queue request  
15 messages of the queue can be acquired by sending one command acquisition request. Thus, the number of the storage commands acquired from the host by the DMA component in one-time transfer can be increased, the number of times of transferring of the storage commands from the host by the DMA component can be reduced, blockage and timeout caused when the storage controller acquires the storage commands from the  
20 host can be avoided, the efficiency of acquiring the storage commands can be increased, and storage performance of a storage system can be improved.

[0049] How to accumulate the queue request messages is not particularly limited in the embodiments of the present disclosure. In some embodiments, the storage controller maintains a flag register table for storing flag register information of each queue, and  
25 the flag register information of each queue is configured to represent whether the storage controller receives a queue request message for the queue. In the accumulation process, each time the storage controller receives a queue request message for a queue, the storage controller enables a flag register corresponding to the queue, and the flag register which is enabled represents that the queue request message for the queue is  
30 received and also represents that the queue has accumulated at least one queue request message.

[0050] FIG. 2 is a schematic diagram of a queue between the host and the storage controller according to an embodiment of the present disclosure. As shown in FIG. 2, a queue is also called a storage Submission Queue (SQ), and one SQ is equivalent to an area in a memory, and includes a head pointer “SQ head Doorbell” and a tail pointer “SQ tail Doorbell”. The storage controller updates the tail pointer “SQ tail Doorbell” after receiving a queue request message for the SQ, and updates the head pointer “SQ head Doorbell” after acquiring and processing a storage command. For example, assuming that the tail pointer “SQ tail Doorbell” = the head pointer “SQ head Doorbell” = 0, the host generates three storage commands and sends queue request messages to the storage controller, and the storage controller updates the tail pointer “SQ tail Doorbell” of the queue according to the queue request messages, at this time, the head pointer “SQ head Doorbell” = 0 and the tail pointer “SQ tail Doorbell” = 3; and after the storage controller sends a command acquisition message, acquires the three storage commands from the host, and executes the three storage commands, the storage controller updates the head pointer “SQ head Doorbell”, at this time, the head pointer “SQ head Doorbell” = 3 and the tail pointer “SQ tail Doorbell” = 3.

[0051] In the embodiments of the present disclosure, before operation S1 is performed, the storage controller does not respond to the received queue request messages, and the queue request messages received by the storage controller are accumulated according to a mechanism illustrated by FIG. 2. For example, assuming that the tail pointer “SQ tail Doorbell” = the head pointer “SQ head Doorbell” = 0 in an initial phase, in a first phase, the host generates the three storage commands for the queue and sends the queue request messages to the storage controller, the storage controller updates the tail pointer “SQ tail Doorbell” of the queue according to the queue request messages, at this time, the head pointer “SQ head Doorbell” = 0 and the tail pointer “SQ tail Doorbell” = 3; and in a second phase, the host generates two storage commands for the queue and sends queue request messages to the storage controller, and the storage controller updates the tail pointer “SQ tail Doorbell” of the queue according to the queue request messages, at this time, the head pointer “SQ head Doorbell” = 0 and the tail pointer “SQ tail Doorbell” = 5, ..., thus realizing the accumulation of the queue request messages.

[0052] In some embodiments, the storage controller may determine the queue (i.e., the first queue) which has accumulated the at least one queue request message based on the

flag register table, and then initiate a response.

[0053] FIG. 3 is a flowchart of a storage control method according to an embodiment of the present disclosure.

5 [0054] Referring to FIG. 3, the storage control method further includes operation S3 before operations S1 and S2.

[0055] At operation S3, a queue corresponding to a flag register which is enabled in a first flag register table of at least one flag register table is determined as the at least one first queue, each flag register table of the at least one flag register table includes a plurality of flag registers respectively corresponding to the plurality of queues, and the  
10 flag register which is enabled represents that the corresponding queue has accumulated the at least one queue request message.

[0056] In the embodiments of the present disclosure, the storage controller may maintain one flag register table, or may maintain a plurality of flag register tables, which is not particularly limited in the embodiments of the present disclosure.

15 [0057] In some embodiments, the storage controller maintains two flag register tables, and alternately accumulating and responds based on the two flag register tables, which is also known as PING-PONG rotation switching between the two flag register tables. It should be noted that accumulating based on a flag register table refers to: when the storage controller receives a queue request message for a queue, enabling a flag register  
20 corresponding to the queue in the flag register table; and responding based on a flag register table refers to: performing operations S3, S1 and S2 to respond to an accumulated queue request message of a queue corresponding to a flag register which is enabled in the flag register table, generate a command acquisition request, and acquire a storage command from the host.

25 [0058] FIG. 4 is a schematic diagram of queue management according to an embodiment of the present disclosure. FIG. 4 is a schematic diagram of alternately accumulating and responding for one queue based on two flag register tables. As shown in FIG. 4, at time t1, accumulating is based on a flag register table 1, and responding is based on a flag register table 2; at time t2, responding is based on the flag register table 1, and

accumulating is based on the flag register table 2; and at time t3, accumulating is based on the flag register table 1, and responding is based on the flag register table 2.

[0059] FIG. 5 is a flowchart of a storage control method according to an embodiment of the present disclosure.

5 [0060] Referring to FIG. 5, the storage control method further includes operation S4.

[0061] At operation S4, in response to a queue request message for a second queue, a flag register corresponding to the second queue in a second flag register table of the at least one flag register table is enabled, and the second queue is a queue corresponding to a flag register which is disabled in the first flag register table.

10 [0062] In the embodiments of the present disclosure, for any flag register table, when responding is made based on the flag register table, the flag register table is used as the first flag register table; and when accumulating is based on the flag register table, the flag register table is used as the second flag register table.

[0063] In the embodiments of the present disclosure, the queue request messages can be  
15 totally accumulated based on the second flag register table while responding based on the first flag register table, so that the number of queue request messages corresponding to a single command acquisition request can be further increased, thereby increasing the number of the storage commands acquired in a single storage command transfer, reducing the number of transfer times, and improving the storage performance. In  
20 addition, with the queue request messages totally accumulated based on the second flag register table while responding based on the first flag register table, omission of the queue request messages can be avoided.

[0064] In the embodiments of the present disclosure, when responding based on the first  
25 flag register table, if a queue request message for the first queue is received before the command acquisition request for the first queue is sent, the accumulating of the queue request messages for the first queue is continued based on the first flag register table; and if the queue request message for the first queue is received after the command acquisition request for the first queue is sent, the queue request message for the first queue is accumulated based on the second flag register table.

[0065] FIG. 6 is a flowchart of a storage control method according to an embodiment of the present disclosure.

[0066] Referring to FIG. 6, the storage control method further includes operations S51 and S52.

5 [0067] At operation S51, in response to a queue request message for any first queue, it is determined whether a command acquisition request for the first queue is sent.

[0068] At operation S52, when the command acquisition request for the first queue is sent, a flag register corresponding to the first queue in the second flag register table is enabled.

10 [0069] Through the operations S51 and S52, efficiency of responding to the queue request messages can be increased, and omission of the queue request messages can be avoided.

[0070] It should be noted that, as shown in FIG. 2, when responding based on the first flag register table, if the queue request message for the first queue is received before the  
15 command acquisition request for the first queue is sent, the tail pointer “SQ tail Doorbell” is updated; and when sending the command acquisition request, the command acquisition request is generated based on the updated tail pointer “SQ tail Doorbell”, that is, the queue request message for the first queue which is received before the  
20 command acquisition request for the first queue is sent is accumulated based on the first flag register table.

[0071] In some embodiments, when responding based on the first flag register table, after the command acquisition request for the first queue is sent, the flag register corresponding to the first queue in the first flag register table is cleared, so as to allow the flag register table used as the first flag register table to be alternately used as the  
25 second flag register table.

[0072] FIG. 7 is a flowchart of a storage control method according to an embodiment of the present disclosure.

[0073] Referring to FIG. 7, the storage control method further includes operation S6

after operation S1.

[0074] At operation S6, for any first queue, after a command acquisition request for the first queue is sent, a flag register corresponding to the first queue in the first flag register table is cleared.

5 [0075] In some embodiments, when alternately accumulating and responding based on the two flag register tables, the two flag register tables are switched according to a predetermined cycle. In some embodiments, when alternately accumulating and responding based on the two flag register tables, the switching is carried out when all flag registers in the first flag register table are cleared, thus realizing the PING-PONG  
10 rotation switching.

[0076] FIG. 8 is a flowchart of a storage control method according to an embodiment of the present disclosure.

[0077] Referring to FIG. 8, the storage control method further includes operation S7.

15 [0078] At operation S7, in a case where all the flag registers in the first flag register table are cleared, the flag register table used as the second flag register table is taken as the first flag register table.

[0079] In the embodiments of the present disclosure, the storage controller may respond to the plurality of first queues at the same time, or may respond to the plurality of first queues one by one, which is not particularly limited in the embodiments of the present  
20 disclosure.

[0080] FIG. 9 is a flowchart of some operations in a storage control method according to an embodiment of the present disclosure.

[0081] Referring to FIG. 9, when responding to the plurality of first queues one by one, operation S1 includes operations S11 and S12.

25 [0082] At operation S11, for any first queue, a command acquisition request for the first queue is generated according to an accumulated queue request message of the first queue, and a number of storage commands corresponding to the command acquisition request for the first queue is equal to a sum of the numbers of storage commands

corresponding to all accumulated queue request messages of the first queue.

[0083] At operation S12, the command acquisition request for the first queue is sent.

[0084] How to respond to the plurality of first queues one by one is not particularly limited in the embodiments of the present disclosure. In some embodiments, the plurality of first queues between the storage controller and the host are numbered, and the plurality of first queues is responded to one by one in an order of the numbered plurality of first queues. In some embodiments, the plurality of queues between the storage controller and the host have different priority levels, and the plurality of queues is responded to one by one in an order of the priority levels of the plurality of queues. In some embodiments, responding may be made in a randomly selected and determined order of the plurality of first queues.

[0085] FIG. 10 is a constituent block diagram of a storage controller according to an embodiment of the present disclosure.

[0086] Referring to FIG. 10, the storage controller includes one or more processing units 101 and a storage unit 102.

[0087] The storage unit 102 has stored thereon one or more programs which, when executed by the one or more processing units 101, cause the one or more processing units 101 to implement the storage control methods described above.

[0088] The processing unit 101 is a device having data processing capability, and includes, but is not limited to, a Central Processing Unit (CPU); and the storage unit 102 is a device having data storage capability, and includes, but is not limited to, a Random Access Memory (RAM, more specifically, a Synchronous Dynamic RAM (SDRAM), a Double Data Rate SDRAM (DDR SDRAM), etc.), a Read-Only Memory (ROM), an Electrically Erasable Programmable Read-Only Memory (EEPROM), and a flash memory (FLASH).

[0089] FIG. 11 is a constituent block diagram of a storage chip according to an embodiment of the present disclosure.

[0090] Referring to FIG. 11, the storage chip includes: a storage controller 201 and a

memory 202.

[0091] The storage controller 201 may be the storage controller described with reference to FIG. 10.

[0092] In some embodiments, the storage chip is a storage acceleration chip.

5 [0093] FIG. 12 is a constituent block diagram of a network card according to an embodiment of the present disclosure.

[0094] Referring to FIG. 12, the network card includes a main control chip 301 and a storage chip 302.

10 [0095] The storage chip 302 includes the storage controller described with reference to FIG. 10.

[0096] In some embodiments, the network card is a storage-acceleration smart network card.

[0097] FIG. 13 is a constituent block diagram of a computer-readable medium according to an embodiment of the present disclosure.

15 [0098] Referring to FIG. 13, the computer-readable medium has stored thereon a computer program which, when executed by a processor, implements the storage control methods described above.

20 [0099] In order to enable those of ordinary skill in the art to understand the technical solutions provided by the embodiments of the present disclosure more clearly, the technical solutions provided by the embodiments of the present disclosure are illustrated in detail below by specific examples.

#### Example One

25 [0100] In this example, the storage control system includes a plurality of storage controllers. FIG. 14 is a flowchart illustrating management of a plurality of queues according to an example of the present disclosure. Referring to FIG. 14, an interaction process for managing a plurality of storage SQs between the storage controllers and the host includes operations E1 to E5.

[0101] At E1, each storage controller creates a plurality of storage SQs according to a request of the host and numbers all of the storage SQs.

5 [0102] After the host initializes the storage controllers in the storage control system, the host sends a management message to each storage controller for creating the plurality of storage SQs. Each storage controller acquires a corresponding management command from the host according to the management message, creates the plurality of storage SQs according to the management command, and sequentially numbers the plurality of storage SQs.

10 [0103] At operation E2, the host generates storage commands and sends queue request messages to each storage controller.

[0104] After each storage controller creates the storage SQs, the host generates the storage commands and stores the storage commands in address spaces corresponding to the storage SQs in the host, and at the same time, the host sends the queue request messages to each storage controller.

15 [0105] At operation E3, a message flag management mechanism is established.

[0106] Each storage controller maintains the flag register tables, creates message flag registers corresponding to all the storage SQs created by the storage controller, and enables flag registers of the storage SQs corresponding to the received queue request messages according to the received queue request messages, that is, each time the  
20 storage controller receives one queue request message, the storage controller enables a flag register of a corresponding storage SQ (for example, set the flag register to 1).

[0107] At operation E4, the flag registers of all the numbered storage SQs are subjected to hand-raising monitoring.

[0108] Each storage controller takes a storage SQ, which is found to have a hand-raising  
25 action through the monitoring (a corresponding flag register is 1), as a first queue, responds to the plurality of storage SQs according to an order of the numbered storage SQs, transfers the storage commands from the host to the storage controller, and clear the corresponding flag registers after the storage commands are transferred.

[0109] At operation E5, while performing operations E1 to E4 based on the first flag register table, a next round of accumulating and aggregating is performed based on the second flag register table, the PING-PONG rotation switching is performed, and so on.

#### Example Two

5 [0110] FIG. 15 is an architecture diagram of management of a plurality of queues according to an example of the present disclosure. FIG. 15 is a schematic architecture diagram of a storage controller and a host. As shown in FIG. 15, a process of performing management interaction on a plurality of storage SQs includes the following operations.

10 [0111] The host sends a queue request message, i.e., a Doorbell message, and the queue request message is transmitted to a message processing and response module in the storage controller through a Peripheral Component Interconnect express (PCIe) module and a DMA control module.

15 [0112] The message processing and response module sends a command acquisition request for a corresponding storage SQ to the host through the DMA control module and the PCIe module.

[0113] The host may generate, according to the command acquisition request, a PCIe packet including a storage command for the corresponding storage SQ, and the PCIe packet is transmitted to a command processing module in the storage controller through the PCIe module and the DMA control module for subsequent processing.

#### 20 Example Three

[0114] FIG. 16 is an architecture diagram of management of a plurality of queues according to an example of the present disclosure. FIG. 16 is a schematic architecture diagram of performing management interaction on a plurality of storage SQs. In this embodiment, the storage control system includes a plurality of storage controllers.

25 [0115] The host may initialize each storage controller in the storage control system through the DMA component, and each storage controller creates a plurality of storage SQs. After the storage SQs are created, the host sends queue request messages (Doorbell messages) corresponding to the storage SQs to each storage controller in the storage control system, and the Doorbell messages are transmitted to message monitoring units

of the corresponding storage controllers.

[0116] Each storage controller assigns a corresponding number to each created storage SQ in a flag register table component. For example, if 2048 storage SQs are created, the corresponding numbers are 0 to 2047, each number is corresponding to a unique SQ, and each SQ has one flag register, a flag register being 1 indicates that an update of a Doorbell message is performed, and a flag bit being 0 indicates that there is no update of a Doorbell message. Each time a message monitoring unit monitors that an update of a Doorbell message is performed for a corresponding storage SQ, the flag register corresponding to the storage SQs is set to 1.

[0117] The message monitoring unit monitors and judges that the 2048 flag registers concurrently. If there is a plurality of flag registers being 1, the storage controller responds to the storage SQs whose flag registers are 1. The 2048 flag registers form the flag register table component.

[0118] For each SQ whose flag register is found to be 1 through the monitoring, a queue response execution unit compares a value of “SQ tail Doorbell” corresponding to the SQ with a value of “SQ head Doorbell” corresponding to the SQ to generate a command acquisition request, and sends the command acquisition request to the host, then a storage command is transferred from the host to the storage controller through the DMA component, and at the same time, the flag register table component is informed to clear the flag register of the SQ.

[0119] In this example, a PING-PONG rotation switching design is adopted, and two flag register tables are designed and are respectively denoted as Table A and Table B. In a first round, the message monitoring unit monitors the Table A in the flag register table component; for a first queue whose flag register is 1 in the Table A, if an update of a new Doorbell message needs to be performed for the first queue before a command acquisition request for the first queue is sent, accumulation and aggregation of the doorbell messages of the first queue may be continued based on the Table A and are managed by a message aggregation control unit. While responding based on the Table A, if an update of a Doorbell message needs to be performed for a second queue whose flag register in the Table A is not 1, the flag register of the second queue in the Table B is set

to 1; and if a plurality of Doorbell messages are continuously issued for individual second queues, accumulation and aggregation of the Doorbell messages of the individual second queues are performed based on the Table B. When all the flag registers in the Table A are cleared, a flag PING-PONG switching unit controls the switching from the Table A to the Table B and informs the message monitoring unit to monitor the Table B in a next round, thus realizing switching from a PING round to a PONG round. After the flag registers in the Table B are processed, the Table A is processed in a next round, thus realizing switching from the PONG round to the PING round. In this way, the Table A and the Table B are subjected to the PING-PONG switching and the switching is repeated cyclically. With all the above operations, no new Doorbell messages of the corresponding SQs are missed, and the Doorbell messages can be accumulated and the aggregated to the greatest extent.

[0120] After the queue response execution unit is notified by the message monitoring unit, the queue response execution unit may send a command acquisition request corresponding to each storage SQ that is subjected to full accumulation and aggregation of the Doorbell messages to the DMA component, and the DMA component then initiates a command acquisition operation to the host and acquires a storage command.

[0121] In this example, the number of transfer times of the DMA component can be reduced, and the number of the storage commands transferred by the DMA component in a single transfer can be increased, so that the efficiency of acquiring the storage commands can be increased, and the storage performance can be improved.

[0122] It should be understood by those of ordinary skill in the art that the functional modules/units in all or some of the operations, the systems and the devices in the method disclosed above may be implemented as software, firmware, hardware, or suitable combinations thereof. If implemented as hardware, the division between the functional modules/units stated above is not necessarily corresponding to the division of physical components; for example, one physical component may have a plurality of functions, or one function or operation may be performed through cooperation of several physical components. Some or all of the physical components may be implemented as software executed by a processor, such as a central processing unit, a digital signal processor or a microprocessor, or may be implemented as hardware, or may be

implemented as an integrated circuit, such as an application specific integrated circuit. Such software may be distributed on a computer-readable medium, which may include a computer storage medium (or a non-transitory medium) and a communication medium (or a transitory medium). As well known by those of ordinary skill in the art, the term

5 “computer storage medium” includes volatile/nonvolatile and removable/non-removable media used in any method or technology for storing information (such as computer-readable instructions, data structures, program modules and other data). The computer storage medium includes, but is not limited to, an RAM, an ROM, an EEPROM, a flash memory or other memory techniques, a Compact Disc Read Only

10 Memory (CD-ROM), a Digital Versatile Disc (DVD) or other optical discs, a magnetic cassette, a magnetic tape, a magnetic disk or other magnetic storage devices, or any other medium which can be configured to store desired information and can be accessed by a computer. In addition, it is well known by those of ordinary skill in the art that the communication media generally include computer-readable instructions, data structures,

15 program modules, or other data in modulated data signals such as carrier wave or other transmission mechanism, and may include any information delivery medium.

[0123] The present disclosure has disclosed the exemplary embodiments using specific terms, but the terms are merely used and should be merely interpreted as having general illustrative meanings, rather than for the purpose of limitation. Unless expressly stated,

20 it is apparent to those of ordinary skill in the art that features, characteristics and/or elements described in connection with a particular embodiment may be used alone or in combination with features, characteristics and/or elements described in connection with other embodiments. Therefore, it should be understood by those of ordinary skill in the art that various changes in the forms and the details may be made without departing

25 from the scope of the present disclosure of the appended claims.

**What is claimed is:**

1. A storage control method, comprising:

sending a command acquisition request for at least one first queue, wherein the at least one first queue is a queue among a plurality of queues which accumulates at least one queue request message, and a command acquisition request for each first queue of the at least one first queue corresponds to all accumulated queue request messages of the first queue; and

receiving a storage command for each first queue of the at least one first queue.

2. The storage control method of claim 1, wherein before sending the command acquisition request for the at least one first queue, the storage control method further comprises:

determining a queue corresponding to a flag register which is enabled in a first flag register table of at least one flag register table as the at least one first queue, wherein each of the at least one flag register table comprises a plurality of flag registers respectively corresponding to the plurality of queues, and the flag register which is enabled represents that the corresponding queue has accumulated the at least one queue request message.

3. The storage control method of claim 2, further comprising:

in response to a queue request message for a second queue, enabling a flag register corresponding to the second queue in a second flag register table of the at least one flag register table, wherein the second queue is a queue corresponding to a flag register which is disabled in the first flag register table.

4. The storage control method of claim 3, further comprising:

in response to a queue request message for any first queue of the at least one first queue, determining whether a command acquisition request for the any first queue is sent; and

when the command acquisition request for the any first queue is sent, enabling a flag register corresponding to the any first queue in the second flag register table.

5. The storage control method of claim 3 or 4, wherein after sending the command acquisition request for the at least one first queue, the storage control method further comprises:

for any first queue of the at least one first queue, after a command acquisition request for the any first queue is sent, clearing a flag register corresponding to the any first queue in the first flag register table.

6. The storage control method of claim 5, further comprising:

in a case where all flag registers in the first flag register table are cleared, taking a flag register table used as the second flag register table as the first flag register table.

7. The storage control method of any one of claims 1 to 4, wherein sending the command acquisition request for the at least one first queue comprises:

for any first queue of the at least one first queue, generating a command acquisition request for the any first queue according to an accumulated queue request message of the any first queue, wherein a number of storage commands corresponding to the command acquisition request for the any first queue is equal to a sum of numbers of storage commands corresponding to all accumulated queue request messages of the any first queue; and

sending the command acquisition request for the any first queue.

8. A storage controller, comprising:

one or more processing units; and

a storage unit having stored thereon one or more programs which, when executed by the one or more processing units, cause the one or more processing units to implement the storage control method of any one of claims 1 to 7.

9. A storage chip, comprising:

the storage controller of claim 8; and

a memory.

10. A network card, comprising:  
a main control chip and a storage chip,  
wherein the storage chip comprises the storage controller of claim 8.

11. A computer-readable medium storing a computer program which, when executed by a processor, implements the storage control method of any one of claims 1 to 7.

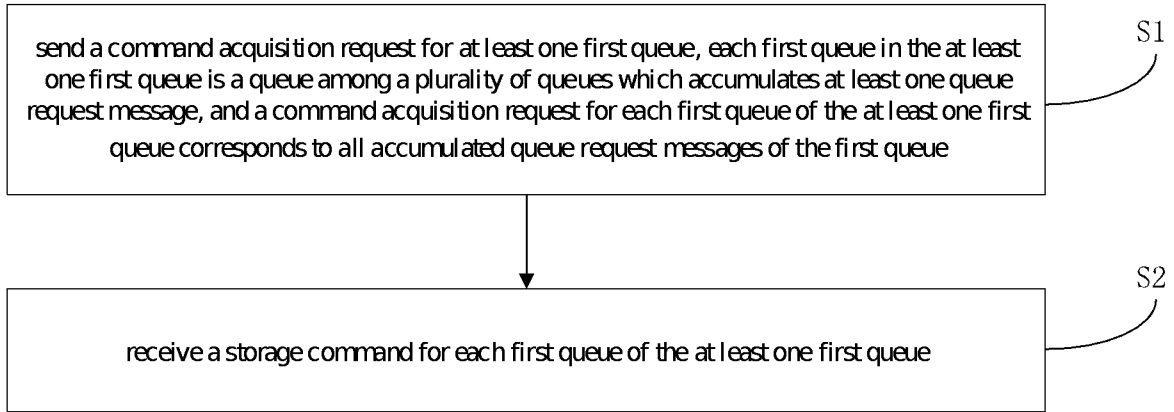


FIG. 1




FIG. 2

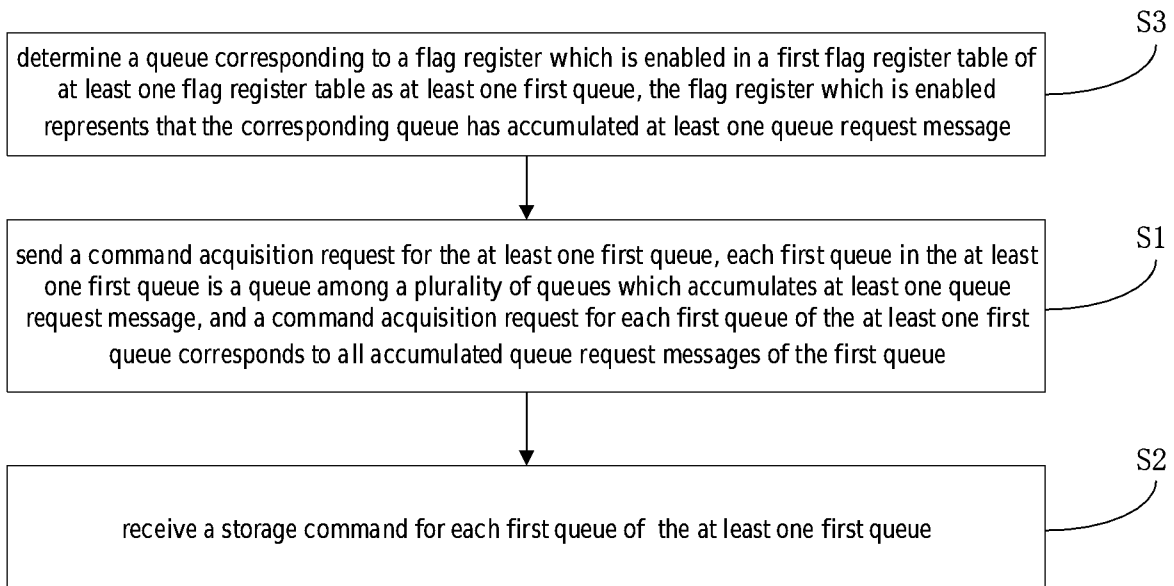


FIG. 3

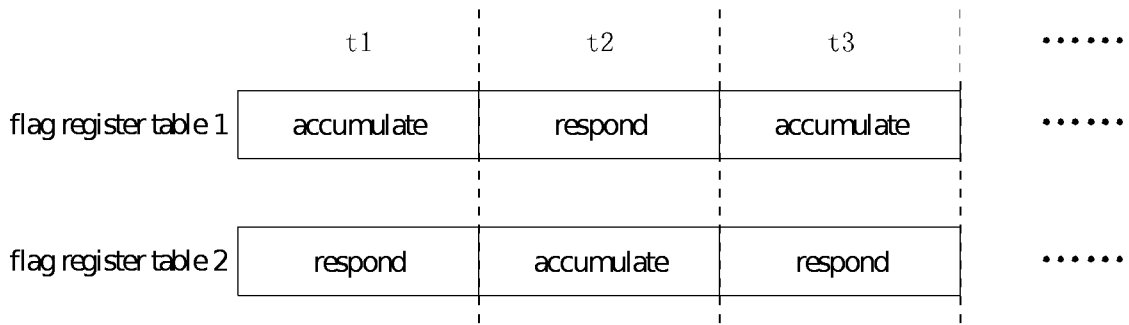


FIG. 4

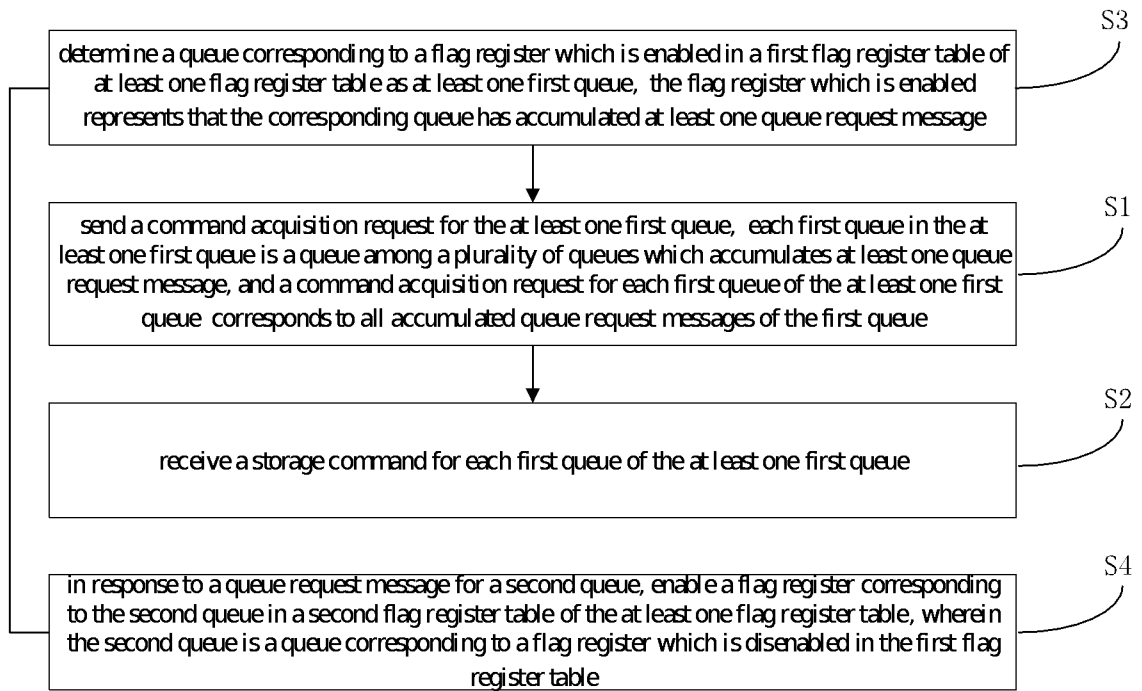


FIG. 5

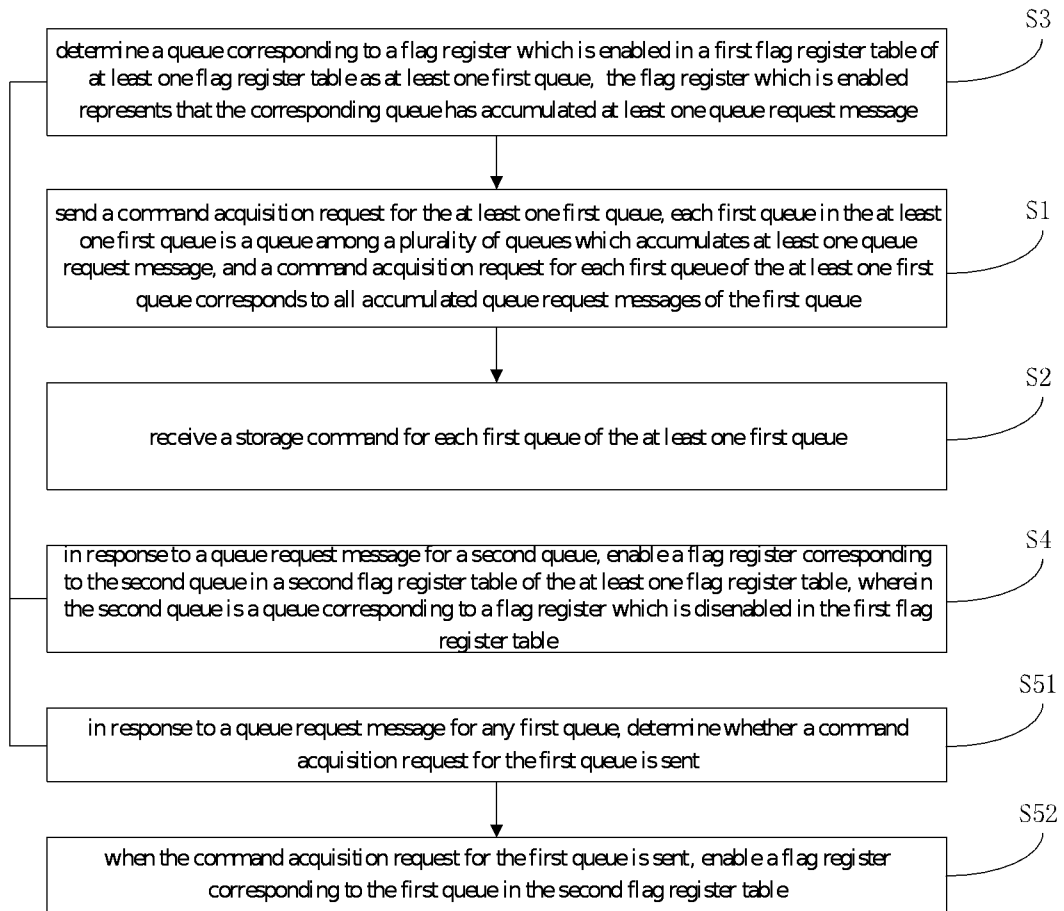


FIG. 6

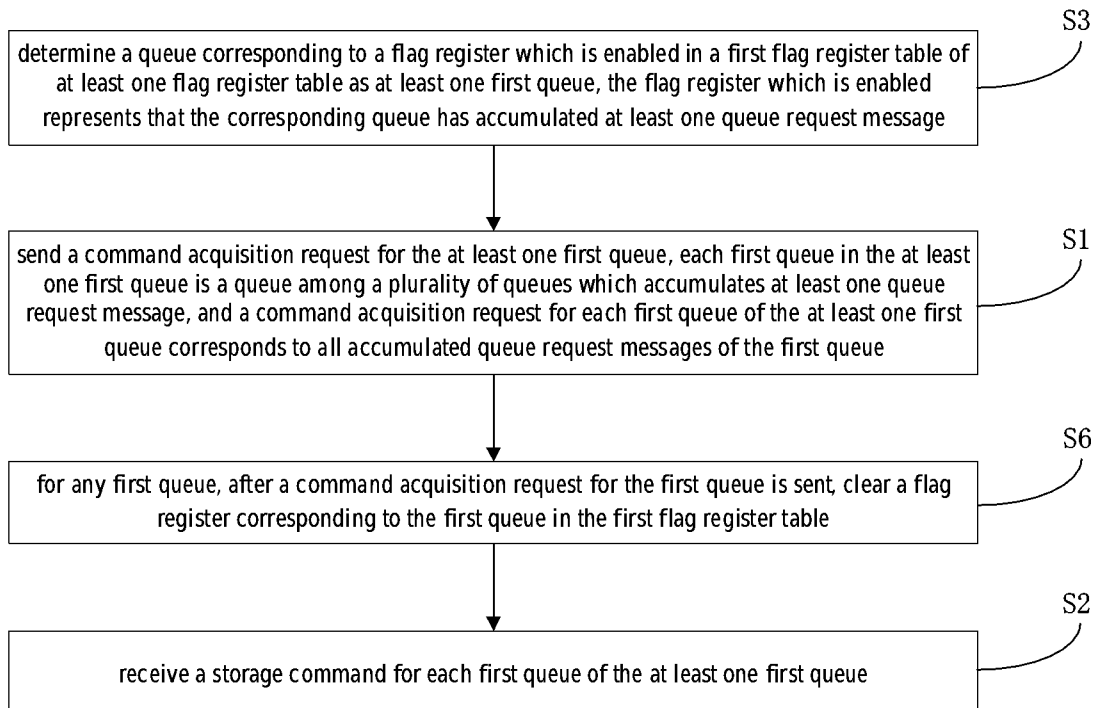


FIG. 7

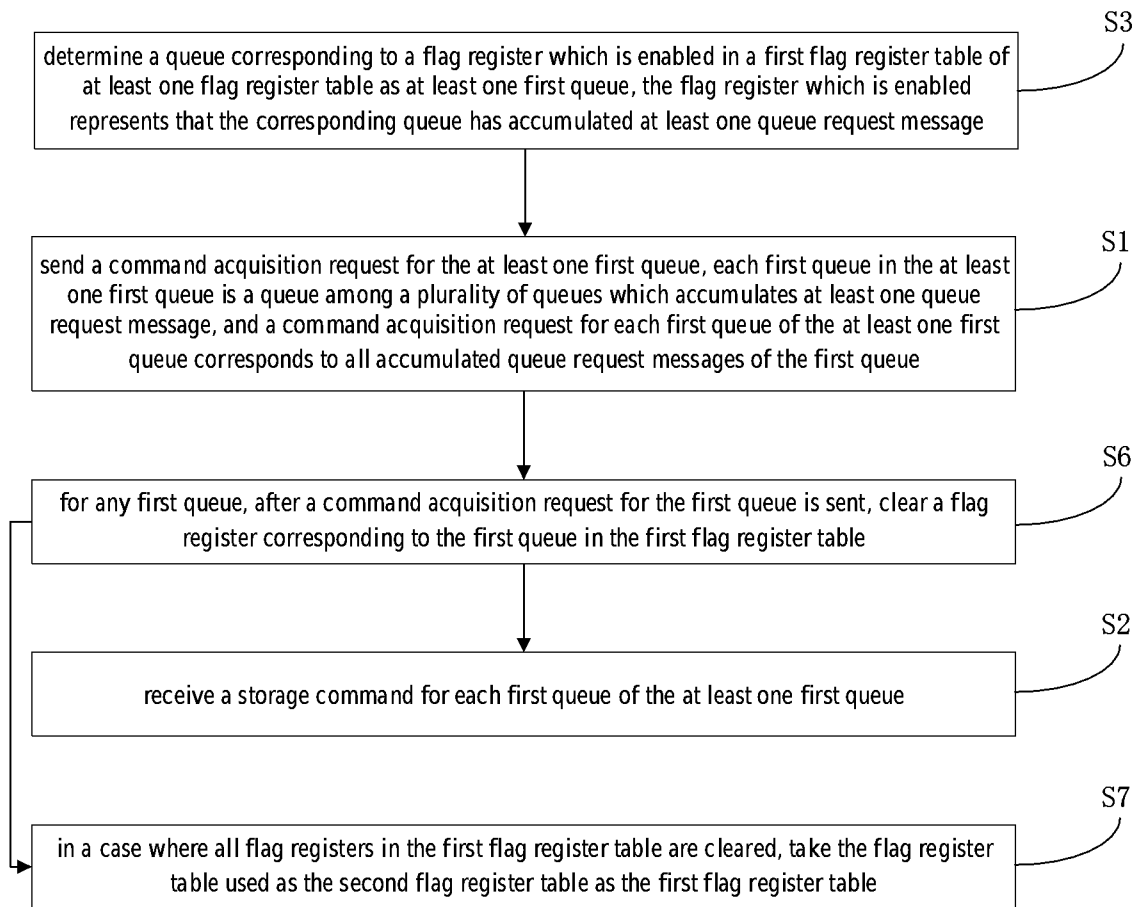


FIG. 8

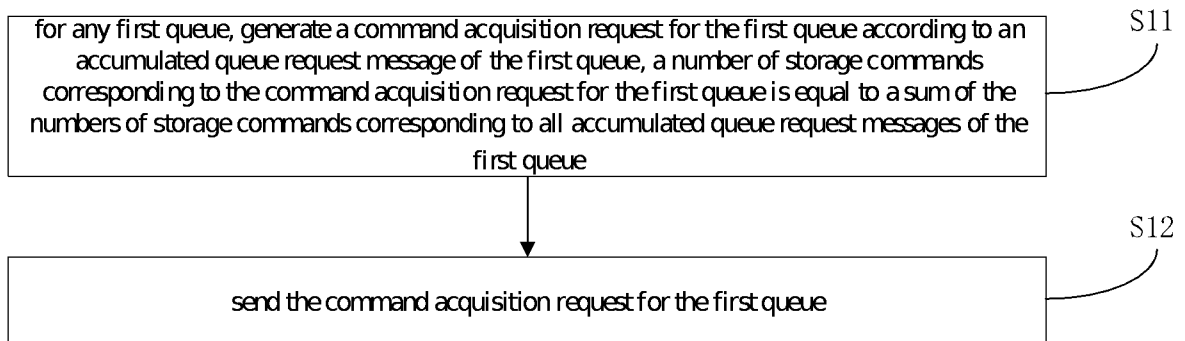


FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13

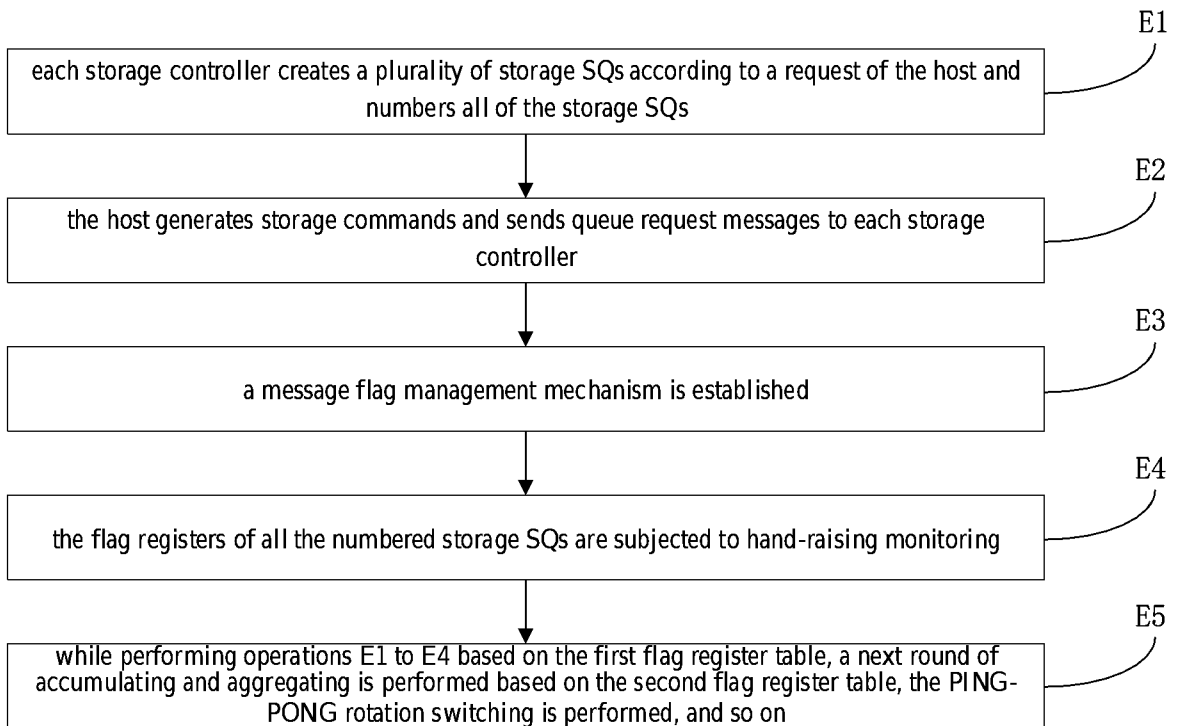


FIG. 14

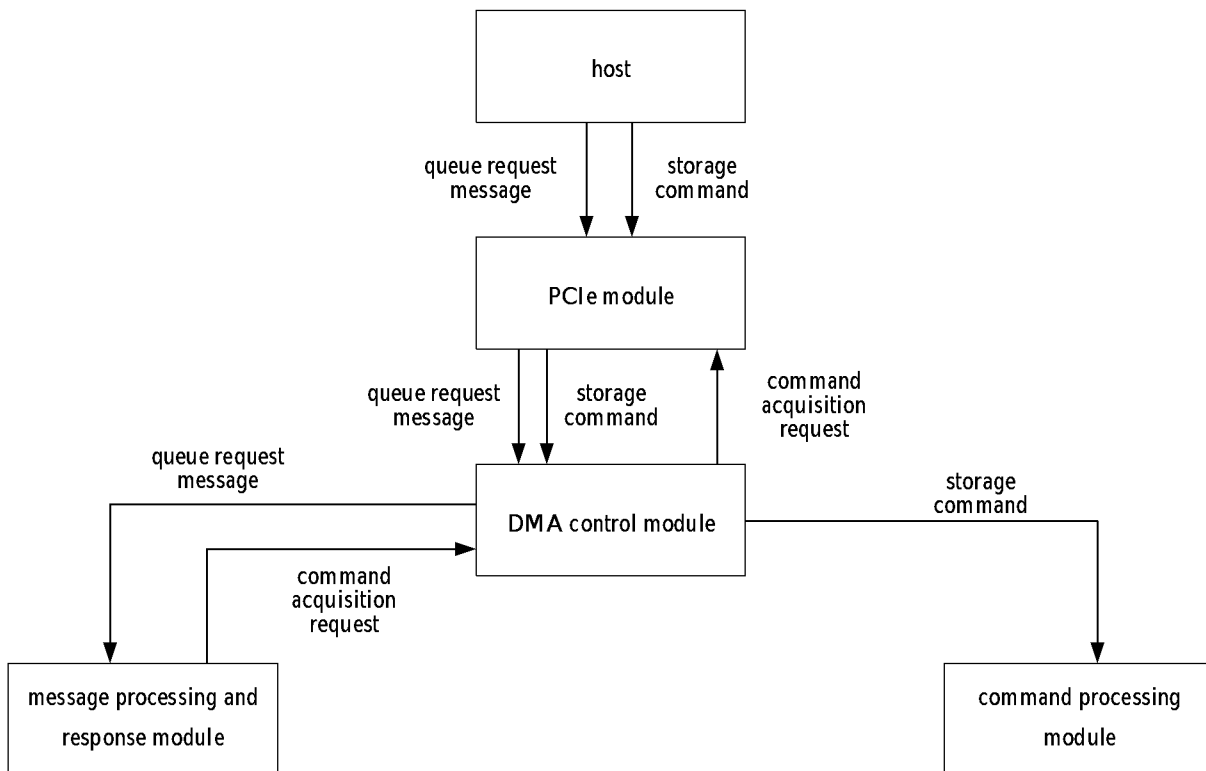


FIG. 15

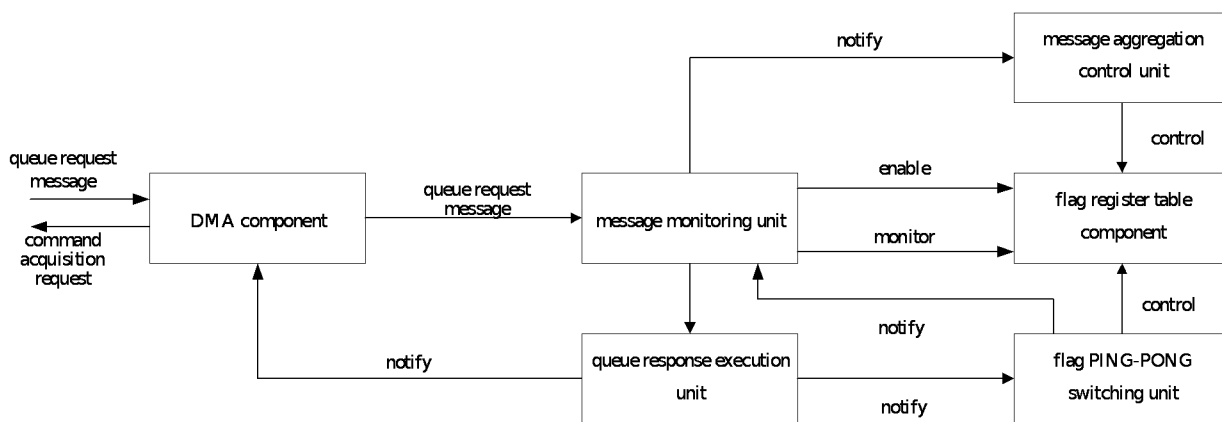


FIG. 16

发送至少一个第一队列的命令获取请求，至少一个第一队列中的每个第一队列为多个队列中的已累积至少一个队列请求消息的队列，至少一个第一队列中的每个第一队列的命令获取请求对应该第一队列的已累积的全部队列请求消息

S1

S1 Send a command acquisition request for at least one first queue, wherein each of the at least one first queue is a queue, among a plurality of queues, for which at least one queue request message has accumulated, and a command acquisition request for each of the at least one first queue corresponds to all of the accumulated queue request messages for the first queue

S2 Receive a storage command for each of the at least one first queue

接收至少一个第一队列中的每一个第一队列的存储命令

S2

图1