

(19)



(11)

EP 3 534 264 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
22.05.2024 Bulletin 2024/21

(21) Application number: **16922519.0**

(22) Date of filing: **23.11.2016**

(51) International Patent Classification (IPC):
G06F 12/02^(2006.01) G06F 12/06^(2006.01)

(52) Cooperative Patent Classification (CPC):
G06F 12/0607; G06F 12/02; G06F 2212/653; G06F 2212/657

(86) International application number:
PCT/CN2016/106987

(87) International publication number:
WO 2018/094620 (31.05.2018 Gazette 2018/22)

(54) MEMORY ALLOCATION METHOD AND APPARATUS

SPEICHERZUWEISUNGSVERFAHREN UND -VORRICHTUNG

PROCÉDÉ ET APPAREIL D'ATTRIBUTION DE MÉMOIRE

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

(43) Date of publication of application:
04.09.2019 Bulletin 2019/36

(73) Proprietor: **Huawei Technologies Co., Ltd.**
Longgang District
Shenzhen, Guangdong 518129 (CN)

(72) Inventors:
• **XIAO, Shihai**
Shenzhen
Guangdong 518129 (CN)
• **HU, Xing**
Shenzhen
Guangdong 518129 (CN)

• **WANG, Kanwen**
Shenzhen
Guangdong 518129 (CN)
• **YANG, Wei**
Shenzhen
Guangdong 518129 (CN)

(74) Representative: **Gill Jennings & Every LLP**
The Broadgate Tower
20 Primrose Street
London EC2A 2ES (GB)

(56) References cited:
CN-A- 1 963 788 CN-A- 102 681 946
US-A1- 2006 248 261 US-A1- 2007 162 715
US-A1- 2012 102 270 US-A1- 2013 268 741
US-A1- 2016 306 567

EP 3 534 264 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description**TECHNICAL FIELD**

[0001] This application relates to the computer field, and in particular, to a memory allocation method and a device.

BACKGROUND

[0002] A computer system includes a processor, a memory controller, and a memory. As shown in FIG. 1, the memory includes a plurality of ranks. A rank may be understood as a matrix formed by a plurality of storage units. A row of the matrix is marked as row (namely, ROW), and a column of the matrix is marked as column (COL for short). There may be a mapping relationship between one rank and a plurality of memory pages. There is a mapping relationship between each memory page and one rank. A memory page includes a plurality of storage units. In FIG. 1, each rank includes four ROWs and four COLs, and a memory page includes two storage units.

[0003] A memory access process of the processor may include the following: After receiving a memory access request that carries a to-be-accessed virtual address and that is sent by an application program, if the processor finds that a memory page has been allocated to the to-be-accessed virtual address, the processor sends, to the memory controller, a memory access request carrying a physical address corresponding to the to-be-accessed virtual address. The memory controller first needs to activate a to-be-accessed storage unit before accessing the storage unit. The memory controller may activate one or more ranks at a time. If finding that no memory page has been allocated to the to-be-accessed virtual address, the processor may randomly allocate an idle memory page to the application program.

[0004] However, according to the foregoing memory allocation method, memory pages allocated by the processor to one application program may be concentrated in one rank. When activating a rank in the memory access process of the processor, the memory controller activates a maximum of four ROWs in the rank at a time. Consequently, a data access rate is reduced.

[0005] US 2016/306567 A1 discloses a method for operating a semiconductor device and a semiconductor system.

[0006] US 2013/268741 A1 discloses a system and method for reducing power consumption in a memory chip.

[0007] US 2012/102270 A1 discloses an apparatus to reduce memory power consumption.

SUMMARY

[0008] The invention is defined in the appended claims.

[0009] The dependent claims set out particular embod-

iments.

[0010] Embodiments of the present invention provide a memory allocation method and a device, to improve a data access rate of a memory controller in a memory access process of a processor, thereby improving overall performance.

[0011] To achieve the foregoing objective, the following technical solutions are used in the embodiments of the present invention.

[0012] According to a first aspect, a memory allocation method is provided according to claim 1.

[0013] According to the memory allocation method provided in this application, because first portions of address information in addresses of ranks in one rank group are the same, there is a mapping relationship between the first portion of address information and information about a first preset location in a physical address of a memory page in the memory, and one memory page may include a plurality of storage units, the processor may distribute storage units in one memory page in ranks of one rank group according to a specific allocation policy, thereby ensuring that a memory page allocated to one application program can span a plurality of ranks. In this way, a memory controller can activate a plurality of ranks in a memory access process of the processor, thereby increasing a data access rate. In addition, the processor determines the target rank group based on the access traffic of the rank groups, and uses the idle memory page corresponding to the target rank group as the to-be-allocated memory page. Compared with a prior-art solution in which any idle memory page is used as a to-be-allocated memory page, in this application, the processor uses an idle memory page corresponding to a rank group with relatively large access traffic as a to-be-allocated memory page, leading to a relatively low probability of rank group switching in a process of accessing the memory page by the processor, thereby reducing rank switching overheads.

[0014] In a possible design, an address of each rank in the memory may further include a second portion of address information, there is a mapping relationship between the second portion of address information and page offset information in a physical address of a memory page.

[0015] In a possible unclaimed design, the selecting, by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups may include: selecting, by the processor from the at least two rank groups, a rank group with largest access traffic as the target rank group. In this way, the processor can have a lower probability of rank group switching in the process of accessing the memory page, thereby reducing rank switching overheads.

[0016] In a possible unclaimed design, the selecting, by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups may include: obtaining, by the processor, access traffic of ranks in the at least two rank groups;

and when determining that access traffic of a first rank in a first rank group of the at least two rank groups exceeds total access traffic of each of the other ranks in the at least two rank groups, selecting a second rank group in the at least two rank groups as the target rank group, where access traffic of the second rank group is less than access traffic of the first rank group. In this way, an idle memory page corresponding to the target rank group may be allocated to equalize the access traffic of the ranks in the memory to some extent, so as to improve overall performance.

[0017] In the present invention, the selecting, by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups may include: obtaining, by the processor, access traffic of ranks in the at least two rank groups; and selecting, from the at least two rank groups, a rank group to which a rank whose access traffic is less than a first preset threshold and is greater than 0 belongs as the target rank group. In this way, an idle memory page corresponding to the target rank group may be allocated to increase access traffic of the target rank group, and equalize the access traffic of the ranks in the memory to some extent, so as to improve overall performance. In addition, a rank group to which a rank whose access traffic is 0 belongs is not selected as the target rank group. This can reduce rank switching overheads.

[0018] In a possible unclaimed design, the selecting, by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups may include: selecting, by the processor from the at least two rank groups, a rank group whose access traffic is less than a second preset threshold and is greater than 0 as the target rank group. In this way, an idle memory page corresponding to the target rank group may be allocated to increase the access traffic of the target rank group, and equalize access traffic of ranks in the memory to some extent, so as to improve overall performance. In addition, a rank group to which a rank whose access traffic is 0 belongs is not selected as the target rank group. This can reduce rank switching overheads.

[0019] According to a second aspect, a computing device is provided according to claim 3.

[0020] In a possible design, an address of each rank in the memory further includes a second portion of address information, there is a mapping relationship between the second portion of address information and page offset information in a physical address of a memory page.

[0021] In a possible unclaimed design, the processor may be specifically configured to select, from the at least two rank groups, a rank group with largest access traffic as the target rank group.

[0022] In a possible unclaimed design, the processor may be specifically configured to: when determining that access traffic of a first rank in a first rank group of the at least two rank groups exceeds total access traffic of the

each of other ranks in the at least two rank groups, select a second rank group in the at least two rank groups as the target rank group, where access traffic of the second rank group is less than access traffic of the first rank group.

[0023] In the present invention, the processor may be specifically configured to: obtain access traffic of ranks in the at least two rank groups; and select, from the at least two rank groups, a rank group to which a rank whose access traffic is less than a first preset threshold and is greater than 0 belongs as the target rank group.

[0024] In a possible unclaimed design, the processor may be specifically configured to select, from the at least two rank groups, a rank group whose access traffic is less than a second preset threshold and greater than 0 as the target rank group.

[0025] According to a third unclaimed aspect, a structure of a physical address of a memory is provided and is applied to a computer system. The computer system includes a processor and a memory, the memory includes at least two rank groups, each rank group includes at least two ranks, and first portions of address information in addresses of ranks in one rank group are the same. An address of each rank in the memory further includes a second portion of address information. The physical address includes page number information and page offset information. There is a mapping relationship between information about a first preset location in the page number information and the first portion of address information, and there is a mapping relationship between information about a second preset location in the page offset information and the second portion of address information.

[0026] According to a fourth unclaimed aspect, a method for obtaining an address of a rank is provided. The method may include: obtaining, by a memory controller, information about a first preset location in page number information of a physical address, to obtain a first portion of address information of the rank; obtaining information about a second preset location in page offset information of the physical address, to obtain a second portion of address information of the rank; and obtaining the address of the rank based on the first portion of address information of the rank and the second portion of address information of the rank. A structure of the physical address may be the structure of the physical address described in the third aspect.

[0027] Optionally, based on the fourth aspect, the memory controller may determine ranks whose first portions of address information are the same as one rank group, to calculate access traffic, in a preset time segment, of a rank group to which the rank belongs.

[0028] According to a fifth unclaimed aspect, a computing device is provided. The computing device includes a memory controller, where the memory controller is configured to: obtain information about a first preset location in page number information of a physical address, to obtain a first portion of address information of a rank; obtain

information about a second preset location in page offset information of the physical address, to obtain a second portion of address information of the rank; and obtain an address of the rank based on the first portion of address information of the rank and the second portion of address information of the rank. A structure of the physical address may be the structure of the physical address described in the third aspect.

[0029] According to another aspect, an embodiment of the present invention provides a processor, where the processor includes modules configured to execute the method provided in the first aspect and possible implementations of the first aspect.

[0030] According to another aspect, a computer-readable storage medium is provided. The computer-readable storage medium stores a computer-executable instruction. When at least one processor of a computing device executes the computer-executable instruction, the computing device executes the memory allocation method provided in the first aspect or any possible implementation of the first aspect.

[0031] According to another aspect, an embodiment of the present invention provides a memory controller, where the memory controller includes modules configured to execute the method described in the third aspect.

[0032] According to another aspect, a computer-readable storage medium is provided. The computer-readable storage medium stores a computer-executable instruction. When at least one processor of a computing device executes the computer-executable instruction, the computing device executes the memory allocation method provided in the first aspect.

[0033] It can be understood that any computing device or computer storage medium provided above is configured to execute the memory allocation method provided above. For beneficial effects that can be achieved by the memory allocation method, refer to the beneficial effects of the corresponding memory allocation method provided above. Details are not repeated herein.

BRIEF DESCRIPTION OF DRAWINGS

[0034]

FIG. 1 is a schematic structural diagram of a memory provided in the prior art;

FIG. 2 is a schematic architecture diagram of a system to which technical solutions provided in embodiments of the present invention is applicable;

FIG. 3 is a schematic structural diagram of a memory according to an embodiment of the present invention;

FIG. 4 is a schematic structural diagram of a physical address of a memory page according to an embodiment of the present invention;

FIG. 5 is a flowchart of a memory allocation method according to an embodiment of the present invention;

FIG. 6 is a flowchart of a method for obtaining an address of a rank according to an embodiment of the present invention;

FIG. 7 is a schematic structural diagram of a computing device according to an embodiment of the present invention; and

FIG. 8 is a schematic structural diagram of another computing device according to an embodiment of the present invention.

DESCRIPTION OF EMBODIMENTS

[0035] Technical solutions provided in the embodiments of the present invention may be applicable to a computer system shown in FIG. 2. The computer system shown in FIG. 2 includes a processor 21, a memory controller 22, and a memory 23. The processor 21 is a control center of the computer system. The memory controller 22 is configured to manage data in the memory 23.

[0036] In the embodiments of the present invention, the memory 23 is divided into a plurality of rank groups, and each rank group may include at least two ranks. A quantity of ranks included in each rank group may be preset, or may be changed based on an actual requirement (for example, a requirement on a data access rate of the memory controller and a requirement on rank switching overheads). A rank may be a collection of dynamic random access memory (dynamic random access memory, DRAM) chips. These chips are connected to a same chip select signal.

[0037] For example, one rank group may include two ranks or four ranks. Quantities of ranks included in different rank groups may be the same or may be different. Generally, quantities of ranks included in different rank groups are the same. As shown in FIG. 3, the memory includes two rank groups, each rank group includes two ranks, each rank includes four ROWs and four COLs, each rank group includes three memory pages, and each memory page includes four storage units.

[0038] There may be a mapping relationship between each rank group and one or more memory pages, and the storage units included in each memory page may be distributed in ranks of one rank group. One memory page may include a plurality of storage units. In the embodiments of the present invention, distribution, in one rank group, of storage units in one memory page is not limited. For example, storage units in one memory page may be distributed in ranks of one rank group according to an even allocation policy.

[0039] It should be noted that the distribution, in one rank group, of storage units in one memory page relates to a method for determining an idle memory page. It can be understood that storage units in one memory page provided in the prior art are distributed in one rank, while in the embodiments of the present invention, storage units in one memory page are distributed in ranks of one rank group.

[0040] According to the present invention, An imple-

mentation of distributing one memory page in ranks of one rank group is as follows: Each storage unit has a physical address. The physical address includes page number (PAGE NUM) information and page offset (PAGE OFFSET) information. In other words, the physical address includes a PAGE NUM field and a PAGE OFFSET field. An address of each rank includes a first portion of address information and a second portion of address information. There is a mapping relationship between the first portion of address information and information about a first preset location, and the first preset location is a portion of the page number information in the physical address. In addition, first portions of address information in addresses of ranks in one rank group are the same. In addition, there is a mapping relationship between the second portion of address information and information about a second preset location, and the second preset location is a portion of the page offset information in the physical address. A physical address of a memory page is a physical address of a storage unit included in the memory page.

[0041] As shown in FIG. 4, FIG. 4 is a schematic structural diagram of a physical address of a memory page according to an embodiment of the present invention. A PAGE NUM field may include a first portion of address information of a rank, a ROW in which the physical address is located, and a COL in which the physical address is located. A PAGE OFFSET field includes a second portion of address information of the rank, a BA, a BG, a COL, and a block offset (block offset). BA is short for bank address, indicating an address of a bank. BG is short for bank group address, indicating an address of a bank group. One rank may include at least two bank groups, and a bank group may include at least two banks. For a relationship between the rank, the bank group, and the bank, refer to the prior art. Details are not provided herein.

[0042] For example, if one rank group includes two ranks: rank1 and rank2, a 2-bit binary number may be used to mark an address of each rank. For example, an address of the rank1 may be marked as 00, and an address of the rank2 may be marked as 01. In this case, in the address of each rank, a high-order bit of the binary number may be located in a PAGE NUM field, and a low-order bit of the binary number may be located in a PAGE OFFSET field. Because one memory page includes a plurality of storage units, the manner may be used to configure some storage units included in the memory page to be in the rank1, and configure the other storage units to be in the rank2, thereby ensuring that one memory page is distributed in ranks of one rank group.

[0043] For another example, if one rank group includes four ranks: rank1, rank2, rank3, and rank4, a three-bit binary number may be used to mark an address of each rank. For example, an address of the rank1 may be marked as 000, an address of the rank2 may be marked as 001, an address of the rank3 may be marked as 010, and an address of the rank4 may be marked as 011. In

this case, in the address of each rank, a highest-order bit of the binary number (namely, "0") may be located in a PAGE NUM field, and the other bits of the binary number (for example, "00", "01", "10", and "11") may be located in a PAGE OFFSET field. Because one memory page includes a plurality of storage units, the manner may be used to configure some storage units included in the memory page to be in the rank1, configure some storage units to be in the rank2, configure some storage units to be in the rank3, and configure the other storage units to be in the rank4, thereby ensuring that one memory page is distributed in ranks of one rank group.

[0044] In this specification, the terms "first", "second", and so on are intended to distinguish between different objects but do not indicate a particular order of the objects. The term "a plurality of" in this specification means two or more than two.

[0045] As shown in FIG. 5, FIG. 5 is a flowchart of a memory allocation method according to an embodiment of the present invention. The method may be specifically executed by an operating system in a processor. The method may include the following steps S101 to S104.

[0046] S101. The processor receives a memory access request, where the memory access request carries a to-be-accessed virtual address.

[0047] The memory access request may be a read request, used to read data stored in a storage unit marked by a physical address corresponding to the to-be-accessed virtual address. The memory access request may alternatively be a write request, used to write data into the storage unit marked by the physical address corresponding to the to-be-accessed virtual address. If the memory access request is a write request, the memory access request may further carry to-be-written data.

[0048] The operating system in the processor provides a virtual address for an application program, allocates a memory space to the application program based on a requirement of the application program, and records a mapping relationship between the virtual address and a physical address. When receiving the memory access request, the processor may first translate the to-be-accessed virtual address carried in the memory access request to the physical address, and then queries whether a mapping relationship between the to-be-accessed virtual address and the physical address has been recorded. If it is found that the mapping relationship between the to-be-accessed virtual address and the physical address has not been recorded, no memory page has been allocated to the to-be-accessed virtual address. In this case, the following steps S102 to S104 may be performed to allocate a memory page to the virtual address. Then, a memory access request carrying the physical address may be sent to a memory controller. After receiving the memory access request carrying the physical address, the memory controller may access the storage unit marked by the physical address. If the mapping relationship between the to-be-accessed virtual address and the physical address has been recorded, a memory page

has been allocated to the to-be-accessed virtual address. In this case, a memory access request carrying the physical address may be sent to a memory controller. After receiving the memory access request carrying the physical address, the memory controller may access the storage unit marked by the physical address.

[0049] S102. When determining that no memory page has been allocated to the to-be-accessed virtual address, the processor obtains access traffic of rank groups in a memory.

[0050] Access traffic of a rank group may be a quantity of times for which the processor accesses the rank group. The memory controller may calculate access traffic of rank groups in at least two rank groups periodically or upon triggering, and then proactively send the access traffic of the rank groups in the at least two rank groups to the processor. Alternatively, the processor proactively obtains the access traffic of the rank groups in the at least two rank groups from the memory controller.

[0051] After receiving the memory access request that carries the physical address and that is sent by the processor, the memory controller may parse out the physical address to obtain an address of a rank, so as to calculate access traffic, in a preset time segment, of the rank, and further calculate access traffic, in the preset time segment, of a rank group to which the rank belongs. As shown in FIG. 6, a method, provided in this embodiment of the present invention, used by the memory controller to parse out the physical address to obtain the address of the rank may include the following steps S1 and S2.

[0052] S1. The memory controller obtains information about a first preset location in page number information of the physical address, to obtain a first portion of address information of the rank; and obtains information about a second preset location in page offset information of the physical address, to obtain a second portion of address information of the rank.

[0053] In the physical address shown in FIG. 4, the memory controller may obtain the information about the first preset location in the PAGE NUM field, to obtain the first portion of address information of the rank; and obtain the information about the second preset location in the PAGE OFFSET field, to obtain the second portion of address information of the rank.

[0054] S2. The memory controller obtains the address of the rank based on the first portion of address information of the rank and the second portion of address information of the rank.

[0055] Based on the example in S1, assuming that the first portion of address information is "0" and the second portion of address information is "01", the address of the rank may be "001".

[0056] Based on S1 and S2, the memory controller may determine ranks whose first portions of address information are the same as one rank group, to calculate the access traffic, in the preset time segment, of the rank group to which the rank belongs.

[0057] S103. The processor selects a target rank group

from at least two rank groups based on the access traffic of the rank groups.

[0058] The processor may obtain, in a preset time, the access traffic of the rank groups. Alternatively, when determining that no memory page has been allocated to the to-be-accessed virtual address, the processor may obtain the access traffic of the rank groups. The processor may obtain the access traffic of the rank groups from the memory controller, or obtain the access traffic of the rank groups from a cache. The access traffic of the rank groups in the cache may be access traffic of the rank groups that is obtained by the processor before the processor receives the memory access request.

[0059] Optionally, S103 may, in an unclaimed embodiment, implemented in the following manner.

[0060] Manner 1: The processor selects, from the at least two rank groups, a rank group with largest access traffic as the target rank group.

[0061] It can be understood that the access traffic provided above is historical access traffic. The historical access traffic provides reference for future access traffic. Specifically, it can be considered that a rank group with larger historical access traffic has a higher probability of being accessed in the future. Therefore, if an idle memory page corresponding to the rank group with the largest access traffic is used as a to-be-allocated memory page, the processor can have a lower probability of rank group switching in a process of accessing the memory page, thereby reducing rank switching overheads. Optionally, in actual implementation, to expand a selection range, the processor may use a rank group with relatively large access traffic as the target rank group. In other words, the processor may select, from the at least two rank groups, any group with access traffic greater than a preset threshold as the target rank group.

[0062] S104. The processor selects, from idle memory pages, a to-be-allocated memory page for the memory access request, where information about a first preset location in a physical address of the to-be-allocated memory page is the same as first portions of address information in addresses of ranks in the target rank group.

[0063] A memory page may be an idle memory page or a non-idle memory page. No data is stored in any storage unit in the idle memory page, and data is stored in one or more storage units in the non-idle memory page. In a memory allocation process, the processor uses an idle memory page as the to-be-allocated memory page. If there is a mapping relationship between the target rank group and a plurality of idle memory pages, the processor may use any one of the plurality of idle memory pages as the to-be-allocated memory page. If there is no mapping relationship between the target rank group and an idle memory page, the processor may re-determine a target rank group based on access traffic of another rank group.

[0064] According to the memory allocation method provided in this embodiment of the present invention, because first portions of address information in addresses

of ranks in one rank group are the same, there is a mapping relationship between the first portion of address information and information about a first preset location in a physical address of a memory page in the memory, and one memory page may include a plurality of storage units, the processor may distribute storage units in one memory page in ranks of one rank group according to a specific allocation policy, thereby ensuring that a memory page allocated to one application program can span a plurality of ranks. In this way, the memory controller can activate a plurality of ranks in a memory access process of the processor, thereby increasing a data access rate. In addition, the processor determines the target rank group based on the access traffic of the rank groups, and uses the idle memory page corresponding to the target rank group as the to-be-allocated memory page. Compared with a prior-art solution in which any idle memory page is used as a to-be-allocated memory page, in this application, the processor uses an idle memory page corresponding to a rank group with relatively large access traffic as a to-be-allocated memory page, leading to a relatively low probability of rank group switching in a process of accessing the memory page by the processor, thereby reducing rank switching overheads.

[0065] Optionally, S103 may, in an unclaimed embodiment, be implemented in any of the following manners.

[0066] Manner 2: The processor obtains access traffic of ranks in the at least two rank groups, and when determining that access traffic of a first rank in a first rank group of the at least two rank groups exceeds total access traffic of the other ranks in the at least two rank groups, selects a second rank group in the at least two rank groups as the target rank group, where access traffic of the second rank group is less than access traffic of the first rank group.

[0067] Access traffic of a rank is a quantity of times for which the processor accesses the rank. That the access traffic of the first rank in the first rank group of the at least two rank groups exceeds the total access traffic of the other ranks in the at least two rank groups indicates that access traffic of the memory is concentrated in the rank. This case may be considered as an abnormal case. In this case, if the second rank group whose access traffic is less than that of the first rank group is used as the target rank group, an idle memory page corresponding to the target rank group may be allocated to equalize the access traffic of the ranks in the memory to some extent, so as to improve overall performance. Optionally, the processor may use a rank group whose access traffic is less than that of the first rank group as the second rank group.

[0068] Manner 3: The processor obtains access traffic of ranks in the at least two rank groups, and then selects, from the at least two rank groups, a rank group to which a rank whose access traffic is less than a first preset threshold and greater than 0 belongs as the target rank group.

[0069] In "a rank group to which a rank whose access

traffic is less than a first preset threshold and greater than 0 belongs", the "access traffic" is access traffic specific to a rank. If the rank group to which the rank whose access traffic is less than the first preset threshold belongs is used as the target rank group, an idle memory page corresponding to the target rank group may be allocated to increase access traffic of the target rank group, and equalize the access traffic of the ranks in the memory to some extent, so as to improve overall performance. A rank group to which a rank whose access traffic is 0 belongs is not selected as the target rank group, because if the rank group to which the rank whose access traffic is 0 belongs is selected, the rank group not expected to be accessed previously starts to be accessed, or in other words, rank switching inevitably occurs in the process, increasing rank switching overheads.

[0070] Manner 4: The processor selects, from the at least two rank groups, a rank group whose access traffic is less than a second preset threshold and is greater than 0 as the target rank group.

[0071] In "a rank group whose access traffic is less than a second preset threshold and is greater than 0", the "access traffic" is access traffic specific to a rank group. If the rank group whose access traffic is less than the second preset threshold is used as the target rank group, an idle memory page corresponding to the target rank group may be allocated to increase the access traffic of the target rank group, and equalize access traffic of ranks in the memory to some extent, so as to improve overall performance. A rank group whose access traffic is 0 is not selected as the target rank group, because if the rank group whose access traffic is 0 is selected, the rank group not expected to be accessed previously starts to be accessed, or in other words, rank switching inevitably occurs in the process, increasing rank switching overheads.

[0072] The solution provided in the embodiments of the present invention is described above mainly from a perspective of the processor. It can be understood that, to implement the foregoing functions, the processor includes corresponding hardware structures and/or software modules for executing the functions. A person skilled in the art should be easily aware that, the modules and algorithm steps in the examples described with reference to the embodiments disclosed in this specification can be implemented by hardware or a combination of hardware and computer software in the present invention. Whether the functions are executed by hardware or computer software driving hardware depends on particular applications and design constraint conditions of the technical solutions.

[0073] In the embodiments of the present invention, the processor may be divided into function modules according to the foregoing method example. For example, function modules corresponding to various functions may be obtained through division, or two or more functions may be integrated into one processing module. The integrated module may be implemented in a form of hard-

ware, or may be implemented in a form of a software function module. It should be noted that the module division in the embodiments of the present invention is used as an example, is merely logical function division, and may be other division in actual implementation.

[0074] FIG. 7 is a schematic structural diagram of a processor 7. The processor 7 may include a receiving module 701 and a selection module 703. Optionally, the processor 7 may further include an obtaining module 702. The receiving module 701 is configured to support the processor 7 in performing S101 in FIG. 4, and/or used in another process of the technology described in this specification. The obtaining module 702 is configured to support the processor 7 in performing S102 in FIG. 4, and/or used in another process of the technology described in this specification. The selection module 703 is configured to support the processor 7 in performing S103 and S104 in FIG. 4, and/or used in another process of the technology described in this specification. In addition, the processor 7 may further include a storage module. The storage module is configured to store corresponding program code and data that are used by the processor 7 to perform any memory allocation method provided above.

[0075] As shown in FIG. 8, FIG. 8 is a schematic structural diagram of a computing device 8 according to an embodiment of the present invention. The computing device 8 may include a processor 801, a memory controller 802, a memory 803, a transceiver 804, and a bus 805. The processor 801, the memory controller 802, the memory 803, and the transceiver 804 are connected to each other by using the bus 805. The processor 801 may be a CPU, a general purpose processor, a digital signal processor (digital signal processor, DSP), an application-specific integrated circuit (application-specific integrated circuit, ASIC), a field programmable gate array (field programmable gate array, FPGA) or another programmable logic device, a transistor logic device, a hardware component, or any combination thereof. The processor 801 may implement or execute various example logic blocks, modules, and circuits described with reference to content disclosed in the present invention. Alternatively, the processor may be a combination implementing a computing function, for example, a combination of one or more microprocessors, or a combination of the DSP and a microprocessor. The bus 805 may be a peripheral component interconnect (peripheral component interconnect, PCI) bus, an extended industry standard architecture (extended industry standard architecture, EISA) bus, or the like. The bus may be divided into an address bus, a data bus, a control bus, and the like. For ease of representation, the bus is represented by using only one bold line in FIG. 8. However, it does not mean that there is only one bus or only one type of bus.

[0076] The methods or algorithm steps described with reference to the content disclosed in the present invention may be implemented in a hardware manner, or may be implemented in a manner of executing a software instruction by a processing module. The software instruc-

tion may include a corresponding software module. The software module may be stored in a random access memory (random access memory, RAM), a flash memory, a read-only memory (read-only memory, ROM), an erasable programmable read-only memory (erasable programmable ROM, EPROM), an electrically erasable programmable read-only memory (electrically EPROM, EEPROM), a register, a hard disk, a removable hard disk, a compact disc read-only memory (CD-ROM), or a storage medium in any other forms well known in the art. A storage medium used as an example is coupled to the processor, so that the processor can read information from the storage medium, and can write information into the storage medium. Certainly, the storage medium may alternatively be a part of the processor. The processor and the storage medium may be located in an ASIC.

[0077] A person skilled in the art should be aware that in one or more of the foregoing examples, the functions described in the present invention may be implemented by using hardware, software, firmware, or any combination thereof. When the functions are implemented by using software, these functions may be stored in a computer-readable medium or transmitted as one or more instructions or code in the computer-readable medium. The computer-readable medium includes a computer storage medium and a communications medium, where the communications medium includes any medium that enables a computer program to be transmitted from one place to another. The storage medium may be any available medium accessible to a general or dedicated computer.

[0078] In the foregoing specific implementations, the objectives, technical solutions, and beneficial effects of the present invention are further described in detail. It should be understood that the foregoing descriptions are merely specific implementations of the present invention, and are not intended to limit the protection scope of the present invention.

40 Claims

1. A memory allocation method, wherein the method is applied to a computer system, the computer system comprises a processor and a memory, the memory comprises at least two rank groups, each rank group comprises at least two ranks, first portions of address information in addresses of ranks in one rank group are the same, there is a mapping relationship between the first portion of address information and a portion of page number information in a physical address of a memory page in the memory; and the method comprises:

receiving (S101), by the processor, a memory access request, wherein the memory access request carries a to-be-accessed virtual address; when determining that no memory page has been allocated to the to-be-accessed virtual ad-

dress, selecting (S103), by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups; and

selecting (S104), by the processor from idle memory pages, a to-be-allocated memory page for the memory access request, wherein the portion of page number information in a physical address of the to-be-allocated memory page is the same as the first portions of address information in addresses of ranks in the target rank group; wherein the selecting (S103), by the processor, a target rank group from the at least two rank groups based on access traffic of the at least two rank groups comprises:

obtaining (S102), by the processor, access traffic of ranks in the at least two rank groups; and

selecting, by the processor from the at least two rank groups, a rank group to which a rank whose access traffic is less than a first preset threshold and is greater than 0 belongs as the target rank group.

2. The memory allocation method according to claim 1, wherein an address of each rank in the memory further comprises a second portion of address information, there is a mapping relationship between the second portion of address information and a portion of page offset information in a physical address of a memory page.

3. A computing device (8), wherein the computing device comprises a processor (801) and a memory (803), the memory comprises at least two rank groups, each rank group comprises at least two ranks, first portions of address information in addresses of ranks in one rank group are the same, there is a mapping relationship between the first portion of address information and a portion of page number information in a physical address of a memory page in the memory; and the processor is configured to:

receive a memory access request, wherein the memory access request carries a to-be-accessed virtual address;

when determining that no memory page has been allocated to the to-be-accessed virtual address, select a target rank group from the at least two rank groups based on access traffic of the at least two rank groups; and

select, from idle memory pages, a to-be-allocated memory page for the memory access request, wherein the portion of page number information in a physical address of the to-be-al-

located memory page is the same as the first portions of address information in addresses of ranks in the target rank group; wherein the processor (801) is specifically configured to:

obtain access traffic of ranks in the at least two rank groups; and

select, from the at least two rank groups, a rank group to which a rank whose access traffic is less than a first preset threshold and is greater than 0 belongs as the target rank group.

4. The computing device according to claim 3, wherein an address of each rank in the memory further comprises a second portion of address information, there is a mapping relationship between the second portion of address information and a portion of page offset information in a physical address of a memory page.

5. A non-transitory computer readable storage medium comprising one or more computer-executable instructions, wherein said one or more computer-executable instructions, when executed on one or more processors of a computer system, cause the computer system to perform the memory allocation method according to any one of claims 1-2.

Patentansprüche

1. Speicherzuweisungsverfahren, wobei das Verfahren auf ein Computersystem angewendet wird, das Computersystem einen Prozessor und einen Speicher umfasst, der Speicher mindestens zwei Ranggruppen umfasst, jede Ranggruppe mindestens zwei Ränge umfasst, erste Abschnitte von Adressinformationen in Adressen von Ränge in einer Ranggruppe gleich sind, es eine Zuordnungsbeziehung zwischen dem ersten Abschnitt der Adressinformationen und einem Teil der Seitenzahlinformationen in einer physischen Adresse einer Speicherseite im Speicher gibt; und das Verfahren ferner Folgendes umfasst:

Empfangen (S101) einer Speicherzugriffsanforderung durch den Prozessor, wobei die Speicherzugriffsanforderung eine virtuelle Adresse enthält, auf die zugegriffen werden soll;

wenn festgestellt wird, dass der virtuellen Adresse, auf die zugegriffen werden soll, keine Speicherseite zugewiesen wurde, Auswählen (S103) einer Zielranggruppe aus den mindestens zwei Ranggruppen durch den Prozessor auf der Grundlage des Zugriffsverkehrs der mindestens zwei Ranggruppen; und

Auswählen (S104) einer zuzuordnenden Speicherseite für die Speicherzugriffsanforderung durch den Prozessor aus freien Speicherseiten, wobei der Abschnitt der Seitennummerinformationen in einer physischen Adresse der zuzuordnenden Speicherseite derselbe ist als erste Abschnitte der Adressinformationen in Adressen von Rängen in der Zielranggruppe; wobei das Auswählen (S103) einer Zielranggruppe aus den mindestens zwei Ranggruppen durch den Prozessor basierend auf dem Zugriffsverkehr der mindestens zwei Ranggruppen Folgendes umfasst:

- Erhalten (S102) des Zugriffsverkehrs von Rängen in den mindestens zwei Ranggruppen durch den Prozessor; und Auswählen einer Ranggruppe, zu der ein Rang gehört, dessen Zugriffsverkehr kleiner als ein erster voreingestellter Schwellenwert und größer als 0 ist, durch den Prozessor aus den mindestens zwei Ranggruppen als Zielranggruppe.
2. Speicherzuweisungsverfahren nach Anspruch 1, wobei eine Adresse jedes Rangs im Speicher außerdem eines zweiten Abschnitts der Adressinformationen umfasst und in a eine Zuordnungsbeziehung zwischen dem zweiten Abschnitt der Adressinformationen und einem Abschnitt der Seitenversatzinformationen in einer physischen Adresse einer Speicherseite besteht.
3. Computergerät (8), wobei das Computergerät einen Prozessor (801) und einen Speicher (803) umfasst, der Speicher mindestens zwei Ranggruppen umfasst, jede Ranggruppe mindestens zwei Ränge umfasst, erste Abschnitte von Adressinformationen in Adressen von Rängen in einer Ranggruppe gleich sind, eine Zuordnungsbeziehung zwischen dem ersten Abschnitt der Adressinformationen und einem Abschnitt der Seitenzahlinformationen in einer physischen Adresse einer Speicherseite im Speicher besteht; und der Prozessor konfiguriert ist zum:

Empfangen einer Speicherzugriffsanforderung, wobei die Speicherzugriffsanforderung eine virtuelle Adresse enthält, auf die zugegriffen werden soll; wenn festgestellt wird, dass der virtuellen Adresse, auf die zugegriffen werden soll, keine Speicherseite zugewiesen wurde, Auswählen einer Zielranggruppe aus den mindestens zwei Ranggruppen basierend auf dem Zugriffsverkehr der mindestens zwei Ranggruppen; und Auswählen einer zuzuordnenden Speicherseite für die Speicherzugriffsanforderung aus freien Speicherseiten, wobei der Abschnitt der Seiten-

nummerinformationen in einer physischen Adresse der zuzuordnenden Speicherseite derselbe ist als erste Abschnitte der Adressinformationen in Adressen von Rängen in der Zielranggruppe; wobei der Prozessor (801) insbesondere konfiguriert ist zum:

- Erhalten von Zugriffsverkehr von Rängen in den mindestens zwei Rängen; und Auswählen einer Ranggruppe, zu der ein Rang gehört, dessen Zugriffsverkehr kleiner als ein erster voreingestellter Schwellenwert und größer als 0 ist, aus den mindestens zwei Ranggruppen als Zielranggruppe.
4. Computergerät nach Anspruch 3, wobei eine Adresse jedes Rangs im Speicher außerdem eines zweiten Abschnitts der Adressinformationen umfasst und in a eine Zuordnungsbeziehung zwischen dem zweiten Abschnitt der Adressinformationen und einem Abschnitt der Seitenversatzinformationen in einer physischen Adresse einer Speicherseite besteht.
5. Nichtflüchtiges computerlesbares Speichermedium, das eine oder mehrere computerausführbare Anweisungen umfasst, wobei die eine oder mehreren computerausführbaren Anweisungen, wenn sie auf einem oder mehreren Prozessoren eines Computersystems ausgeführt werden, das Computersystem dazu veranlassen, das Speicherzuweisungsverfahren nach einem der Ansprüche 1-2 auszuführen.

Revendications

1. Procédé d'attribution de mémoire, dans lequel le procédé est appliqué à un système informatique, le système informatique comprend un processeur et une mémoire, la mémoire comprend au moins deux groupes de rangs, chaque groupe de rangs comprend au moins deux rangs, des premières parties des informations d'adresse dans des adresses de rangs dans un groupe de rangs sont identiques, il existe une relation de mappage entre la première partie des informations d'adresse et une partie des informations de numéro de page dans une adresse physique d'une page de mémoire dans la mémoire ; et le procédé comprend :

la réception (S101), par le processeur, d'une demande d'accès à la mémoire, dans lequel la demande d'accès à la mémoire transporte une adresse virtuelle à laquelle accéder ; lorsqu'il est déterminé qu'aucune page de mémoire n'a été attribuée à l'adresse virtuelle à laquelle accéder, la sélection (S103), par le pro-

cesseur, d'un groupe de rangs cible parmi les au moins deux groupes de rangs sur la base du trafic d'accès des au moins deux groupes de rangs ; et

la sélection (S104), par le processeur parmi des pages de mémoire inactives, d'une page de mémoire à attribuer pour la demande d'accès à la mémoire, dans lequel la partie des informations de numéro de page dans une adresse physique de la page de mémoire à attribuer est identique aux premières parties des informations d'adresse dans des adresses de rangs dans le groupe de rangs cible ;

dans lequel la sélection (S103), par le processeur, d'un groupe de rangs cible parmi les au moins deux groupes de rangs sur la base du trafic d'accès des au moins deux groupes de rangs comprend :

l'obtention (S102), par le processeur, du trafic d'accès de rangs dans les au moins deux groupes de rangs ; et

la sélection, par le processeur parmi les au moins deux groupes de rangs, d'un groupe de rangs auquel appartient un rang dont le trafic d'accès est inférieur à un premier seuil prédéfini et supérieur à 0 en tant que groupe de rangs cible.

2. Procédé d'attribution de mémoire selon la revendication 1, dans lequel une adresse de chaque rang dans la mémoire comprend en outre une deuxième partie des informations d'adresse, il existe une relation de mappage entre la deuxième partie des informations d'adresse et une partie des informations de décalage de page dans une adresse physique d'une page de mémoire.

3. Dispositif informatique (8), dans lequel le dispositif informatique comprend un processeur (801) et une mémoire (803), la mémoire comprend au moins deux groupes de rangs, chaque groupe de rangs comprend au moins deux rangs, de premières parties des informations d'adresse dans des adresses des rangs dans un groupe de rangs sont identiques, il existe une relation de mappage entre la première partie des informations d'adresse et une partie des informations de numéro de page dans une adresse physique d'une page de mémoire dans la mémoire ; et le processeur est configuré pour :

recevoir une demande d'accès à la mémoire, dans lequel la demande d'accès à la mémoire transporte une adresse virtuelle à laquelle accéder ;

lorsqu'il est déterminé qu'aucune page de mémoire n'a été attribuée à l'adresse virtuelle à laquelle accéder, sélectionner un groupe de rangs

cible parmi les au moins deux groupes de rangs sur la base du trafic d'accès des au moins deux groupes de rangs ; et

sélectionner, parmi des pages de mémoire inactives, une page de mémoire à attribuer pour la demande d'accès à la mémoire, dans lequel la partie des informations de numéro de page dans une adresse physique de la page de mémoire à attribuer est identique aux premières parties des informations d'adresse dans des adresses de rangs dans le groupe de rangs cible ; dans lequel le processeur (801) est spécifiquement configuré pour :

obtenir le trafic d'accès de rangs dans les au moins deux groupes de rangs ; et sélectionner, parmi les au moins deux groupes de rangs, un groupe de rangs auquel appartient un rang dont le trafic d'accès est inférieur à un premier seuil prédéfini et supérieur à 0 en tant que groupe de rangs cible.

4. Dispositif informatique selon la revendication 3, dans lequel une adresse de chaque rang dans la mémoire comprend en outre une deuxième partie des informations d'adresse, il existe une relation de mappage entre la deuxième partie des informations d'adresse et une partie des informations de décalage de page dans une adresse physique d'une page de mémoire.

5. Support de stockage non transitoire lisible par ordinateur comprenant une ou plusieurs instructions exécutables par ordinateur, dans lequel lesdites une ou plusieurs instructions exécutables par ordinateur, lorsqu'elles sont exécutées sur un ou plusieurs processeurs d'un système informatique, amènent le système informatique à réaliser le procédé d'attribution de mémoire selon l'une quelconque des revendications 1 à 2.




FIG. 1




FIG. 2

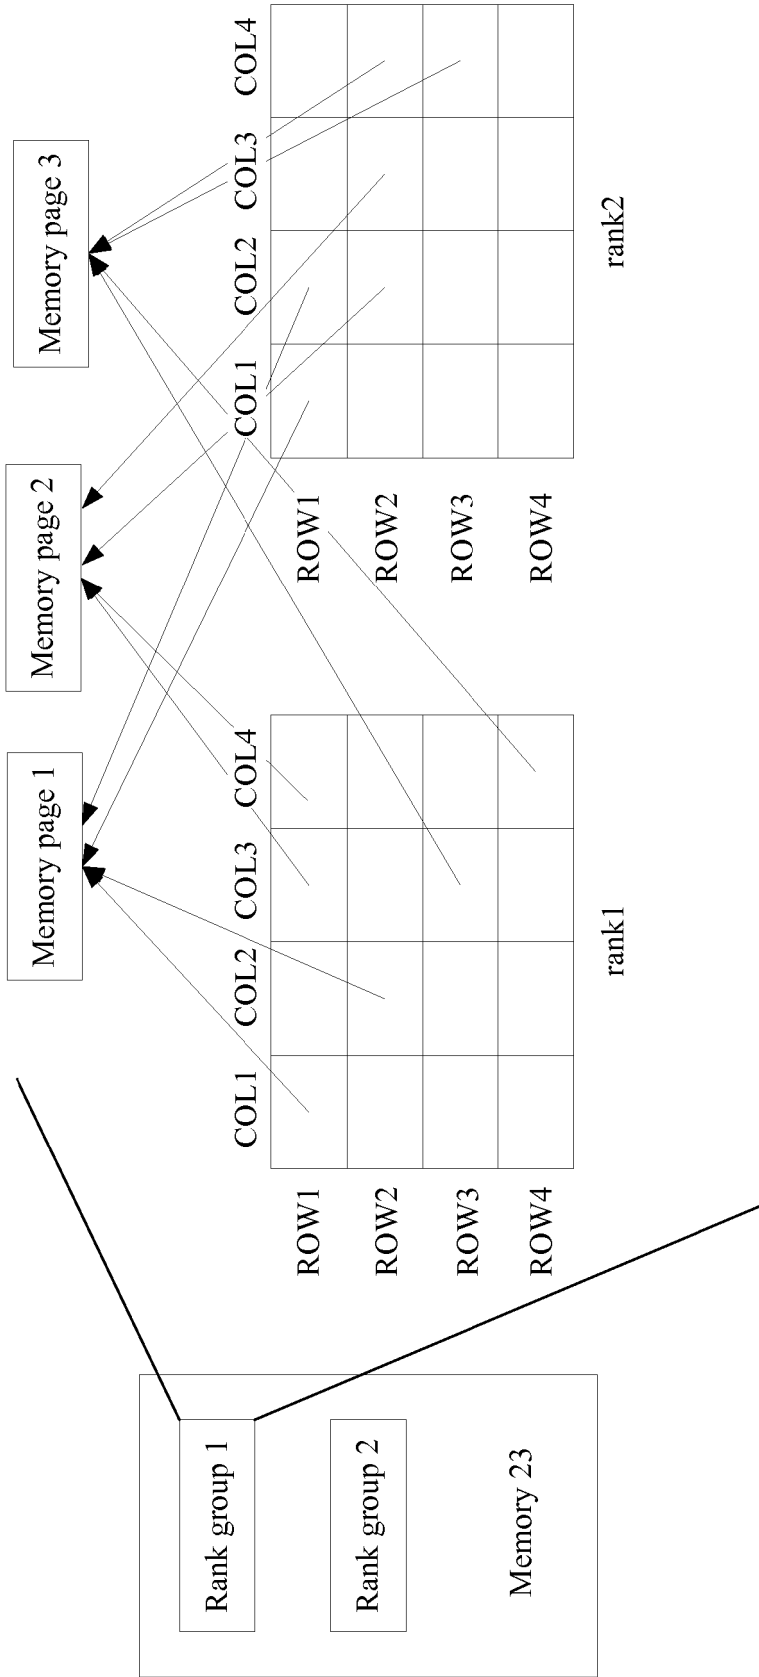


FIG. 3




FIG. 4

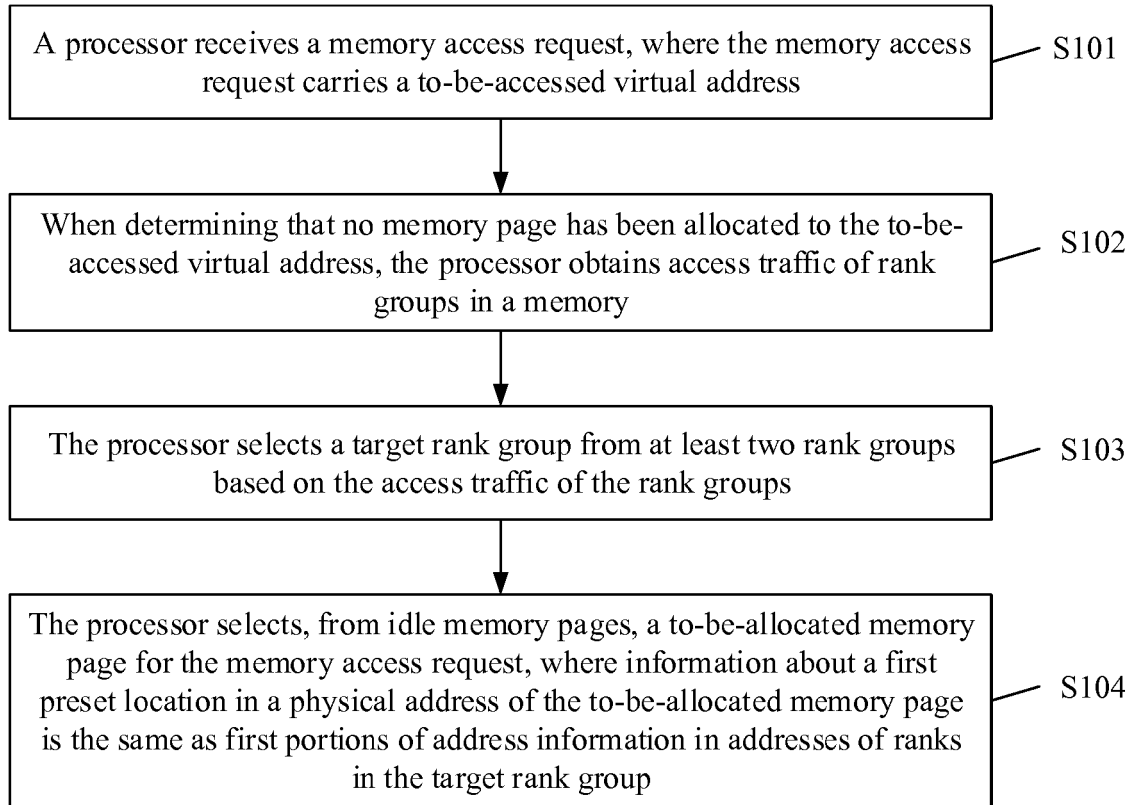


FIG. 5

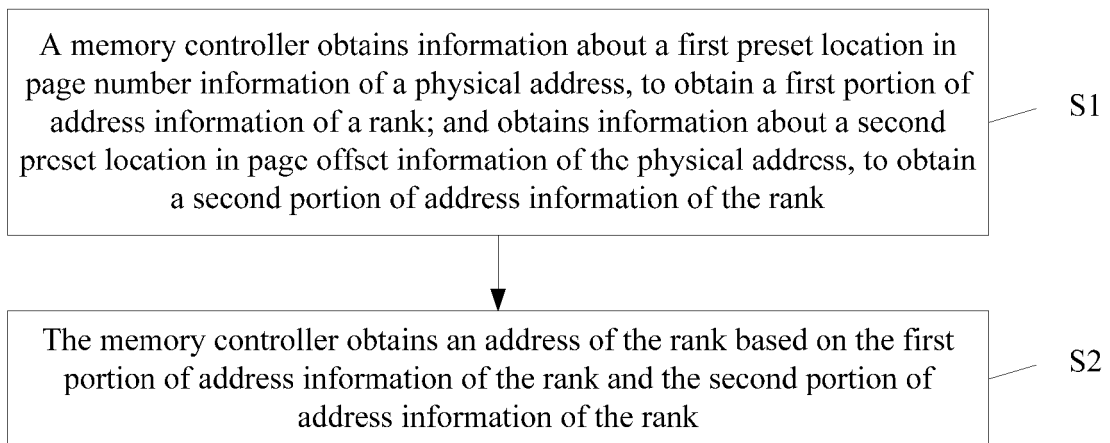


FIG. 6

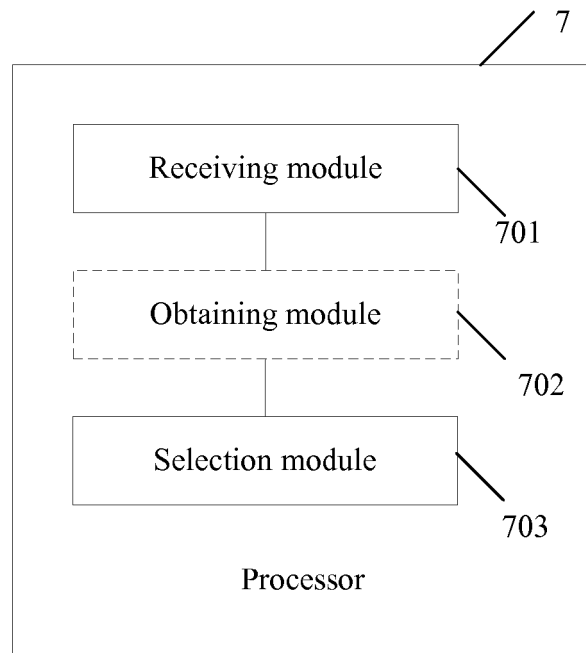


FIG. 7

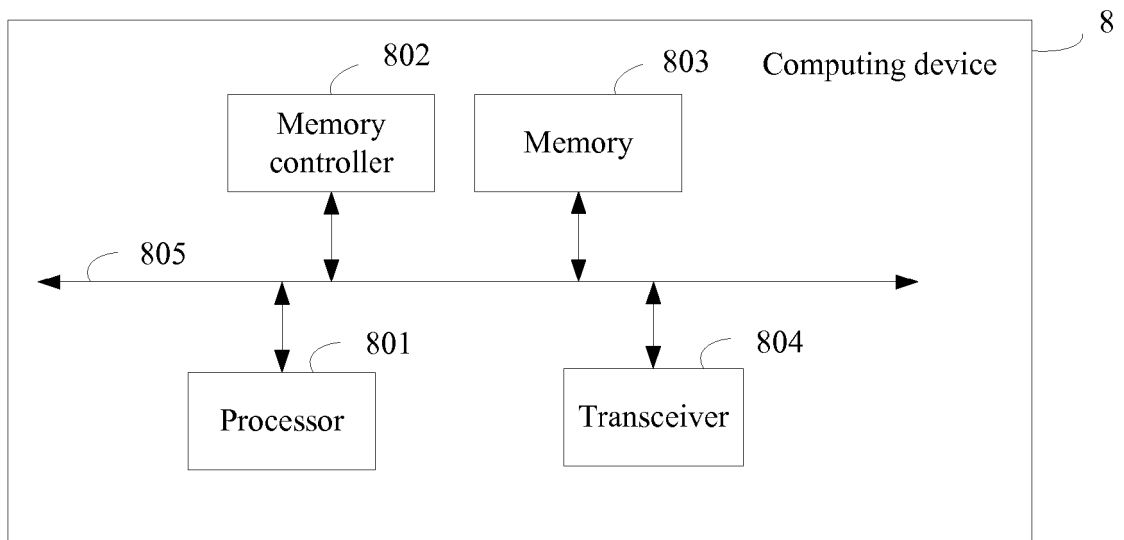


FIG. 8

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2016306567 A1 [0005]
- US 2013268741 A1 [0006]
- US 2012102270 A1 [0007]