



- (51) **International Patent Classification:**
G06F 12/02 (2006.01)
- (21) **International Application Number:**
PCT/US2015/062801
- (22) **International Filing Date:**
25 November 2015 (25.11.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
201410686872.8 25 November 2014 (25.11.2014) CN
- (71) **Applicant:** ALIBABA GROUP HOLDING LIMITED
[—/US]; Fourth Floor, One Capital Place, P.O. Box 847,
George Town, Grand Cayman (KY).
- (72) **Inventors:** NIU, Gongbiao; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen YI Road, Yu Hang District, Hangzhou, 311121 (CN). HUANG, Zheng; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen YI Road, Yu Hang District, Hangzhou, 311121 (CN).
- (74) **Agent:** JIA, Alexandra; Murabito, Hao & Barnes, LLP, 2 N. Market St., 3rd Floor, San Jose, CA 95113 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

[Continued on next page]

(54) **Title:** METHOD AND APPARATUS FOR MEMORY MANAGEMENT

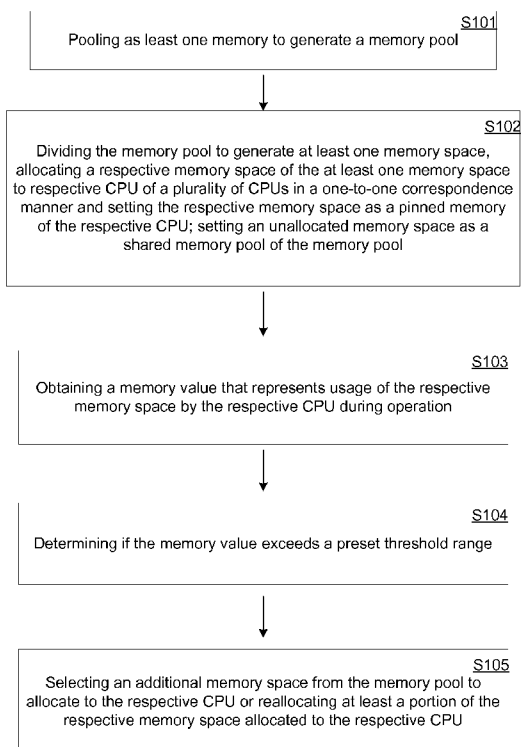


FIG. 1A

(57) **Abstract:** A method and apparatus of memory management are disclosed. Pooling of at least one memory to generate a memory pool, dividing the memory pool to generate at least one memory space, and allocating a respective memory space to a respective CPU in a one-to-one correspondence manner are performed. Further, the respective memory space allocated to the respective CPU is set as a pinned memory of the respective CPU. Additionally, setting unallocated memory space as a shared memory pool, obtaining a memory value that represents usage of the respective memory space by the respective CPU during operation, and determining if the memory value exceeds a preset threshold range are performed. Selecting, if the memory value exceeds the preset threshold range, additional memory space from the memory pool to allocate to the respective CPU or reallocating at least a portion of the respective memory space allocated to the CPU are performed.



(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

METHOD AND APPARATUS FOR MEMORY MANAGEMENT

CROSS-REFERENCE TO RELATED APPLICATION

[001] This application claims priority to and the benefit of Chinese Patent Application No. 201410686872.8, filed on November 25th, 2014, which is incorporated herein by reference in its entirety.

TECHNICAL FIELD

[002] The disclosure relates generally to the field of memory management and, more particularly, to a method and apparatus for managing computer memory.

BACKGROUND

[003] Generally, there is a fixed ratio between memory size and CPU processor capability in a single server. When there is a need to increase the power of a CPU or the memory size, the common manner in which this process is performed is to add one or more servers. As a result, there is an increase in the overall capability of the CPU or the size of the memory. By simply adding additional servers in order to increase memory space or CPU processing capability, there can be a resulting over-capacity that results in a decrease in the utilization of the overall available CPU processing capability or memory usage that can be considered a waste of resources.

[004] Currently, software-based solutions are available to address the above-described problems of the decreased usage rate of CPU or memory as well as the associated waste of resources. Such software-based solutions manage and dispatch memory in a service system by protocol of encapsulation.

[005] However, the current solutions have drawbacks. For example, a memory size is not capable to perform dynamic propagation when memory usage needs to be replicated. Moreover, due to the use of protocols of encapsulation at the software level, a CPU in a server system tends to utilize a relatively large amount of memory to

process these protocols, causing a certain amount of delay in time and decrease in the processing efficiency of the server system.

SUMMARY OF THE DISCLOSURE

[006] Embodiments of the disclosure provide methods and apparatus of memory management. Also, embodiments of the disclosure provide techniques to address the shortcomings of conventional solutions by providing better memory management and CPU processing allocation.

[007] Accordingly, one embodiment of the disclosure provides a method of memory management. The method includes generating a memory pool, dividing the memory pool to generate at least one memory space, and allocating a respective memory space to a respective CPU in a one-to-one correspondence manner. The respective memory space is set as a pinned memory of the respective CPU. The method further includes setting an unallocated memory space as a shared memory pool of the memory pool. That is, the unallocated memory space includes memory space that remains unallocated in the memory pool. The method further includes obtaining a memory value that represents usage of the respective memory space by the respective CPU during operation and determining if the memory value exceeds a preset threshold range. The method yet further includes selecting, if the memory value exceeds the preset threshold range, additional memory space from the memory pool to allocate to the respective CPU and/or reallocating a portion of the respective memory space allocated to the respective CPU.

[008] According to another embodiment of the disclosure, a device for memory management is provided. The device includes a memory pooling unit to generate a memory pool. The device further includes a memory divider to divide to memory pool to generate at least one memory space. Also, the memory divider is operable to allocate a respective memory space to a respective CPU in a one-to-one correspondence manner. The respective memory space is set as a pinned memory of the respective CPU. An unallocated memory space is set as a shared memory pool of the memory pool. That is,

the unallocated memory space includes memory space that remains unallocated or is not pinned memory in the memory pool. The device yet further includes a memory value obtaining unit to obtain a memory value that represents usage of the respective memory space by the respective CPU during operation. Also, the device includes a memory value determination unit to determine if the memory value exceeds a preset threshold range. The device further includes a manager to select additional memory space from the memory pool to allocate to the respective CPU and/or to reallocate at least a portion of the respective memory space allocated to the respective CPU. The memory value determination unit is operable to commence the manager if the memory value exceeds the preset threshold range.

[009] Further, another embodiment of the disclosure provides a system for memory management. The system includes instructions to generate at least one memory pool, instructions to divide the memory pool to generate at least one memory space, and instructions to allocate a respective memory space to a respective CPU in a one-to-one correspondence manner. The respective memory space is set as a pinned memory of the respective CPU. The system further includes instructions to set an unallocated memory space as a shared memory pool of the memory pool. That is, the unallocated memory space includes memory space that remains unallocated in the memory pool. The system further includes instructions to obtain a memory value that represents usage of the respective memory space by the respective CPU during operation and instructions to determine if the memory value exceeds a preset threshold range. The system yet further includes instructions to select, if the memory value exceeds the preset threshold range, additional memory space from the memory pool to allocate to the respective CPU and/or instructions to reallocate at least a portion of the respective memory space allocated to the respective CPU.

[010] This summary includes, by necessity, simplifications, generalizations, and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the disclosure, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

[011] Embodiments of the disclosure will be better understood from a reading of the following detailed description, taken in conjunction with the accompanying drawings, in which like reference characters designate like elements.

[012] Figure 1A is a flow diagram of a method of memory management according to an embodiment.

[013] Figure 1B is a flow diagram of a method of memory management according to an embodiment.

[014] Figure 2 is a block diagram of an apparatus of memory management according to an embodiment.

DETAILED DESCRIPTION

[015] Reference will now be made in detail to embodiments, examples of which are illustrated in the accompanying drawings. While the disclosure will be described in conjunction with the embodiments, it will be understood that they are not intended to limit the disclosure to these embodiments. On the contrary, the disclosure is intended to cover alternatives, modifications, and equivalents which may be included within the spirit and scope of the disclosure as defined by the appended claims. Furthermore, in the following detailed description of embodiments, numerous specific details are set forth in order to provide a thorough understanding of the disclosure. However, it will be recognized by one of ordinary skill in the art that embodiments may be practiced without these specific details. In other examples, well-known methods, procedures, and components have not been described in detail so as not to unnecessarily obscure aspects of embodiments. Further, embodiments of the disclosure provide a method and an apparatus of memory management in a server system. There can be embodiments based on other types of systems which is not limited here.

[016] FIG. 1A shows a method of memory management. In step S101, at least one memory is pooled to generate a memory pool.

[017] The memory mentioned in step S101 may be a memory managed by a memory controller configured by a server in a server system.

[018] Pooling includes, but not limited to, the setting of each memory managed by a memory controller configured by the server in the server system as a node, and the connecting of the nodes in the server system through a logic chip or a logic device to create a memory pool in units of the nodes.

[019] In an embodiment of the disclosure, at least one memory is pooled. There are additional steps after the pooling operation to create the memory pool.

[020] Each CPU configured for each server in the server system is set as a CPU node. The CPU nodes are interconnected through a logic chip or a logic device to create a CPU pool that includes the CPU nodes.

[021] The pooling of the memories in the server system and the pooling of the CPUs realize a physical decoupling of the memory and the CPU in the server system. Therefore, as a result of this physical decoupling, unconstrained configurations between the memories and CPUs in the server system can be accomplished. That is, the number of the memories in the server system can be increased or decreased individually due to these unconstrained configurations. Additionally, the number of the CPUs in the server system may also be increased or decreased individually by the unconstrained configurations to avoid a needless lack of usage of memory and CPU resources.

[022] Additionally, there may be other ways or examples to implement or configure the pooling of memories and the CPUs that are not being addressed in detail herein.

[023] In step S102, the memory pool is divided to generate at least one memory space, each of the memory spaces being allocated respectively to a plurality of CPUs in a one-to-one correspondence manner. The memory space allocated to a CPU is set as a pinned memory of the CPU, and the unallocated memory space is set as a shared memory pool of the memory pool. That is, the unallocated memory space includes memory space that remains unallocated in the memory pool.

[024] Assuming the number of the CPUs in the server system is “n” and the number of the memory controllers is “m,” where ($n < m$), the steps to divide the memory

pool to generate at least one memory space to allocate to a respective CPU and set the allocated memory space as the pinned memory of the respective CPU are described next.

[025] (1) Initially, selecting a number of memory controllers equal to the number of the CPUs in the plurality of CPUs. Here, from the above-described “m” memory controllers, “n” memory controllers are randomly selected as the “n” memory controller corresponding to the “n” CPUs to establish a one-to-one mapping relationship.

[026] According to an embodiment of the disclosure, the relationships between the memory controllers and the CPUs are one-to-one correspondence relationships. There are other types of correspondence relationships between the memory controllers and the CPUs such as in a one-to-many or many-to-one manner, which can be easily configured but will not be addressed in detail herein.

[027] (2) Next, generating a set of relationships between “n” Quick Path Interconnect “QPI” port addresses of the “n” CPUs with “n” port addresses of the “n” memories in a one-to-one correspondence manner; allocating a corresponding memory space of at least one dual in-line memory module “DIMM” slot in at least one channel under each memory controller of the “n” memory controllers to the CPU that has a one-to-one corresponding relationship with the memory controller, the corresponding memory space to be set as the pinned memory of the CPU. The allocating is repeated until all the CPUs in “n” CPUs have obtained pinned memories.

[028] The Quick Path Interconnect “QPI” is a packet transmission-based high-speed serial point-to-point connecting agreement that is used by CPU to access data. The DIMM or dual in-line memory module is a module that includes one or more random-access memory chips on a small circuit board with pins that connect to the computer motherboard.

[029] A memory controller selects a memory for use to manage the exchange of data between the memory and the CPU. Each memory controller can support 2-4 DDR (Double data rate synchronous dynamic random-access memory) channels, each channel can support 1-3 DIMM slots, for 1-3 memory chips. The memory space of each of memory chips corresponds to the memory space of the DIMM slots.

[030] In another embodiment according to the disclosure, the memory space which corresponds to at least one DIMM slot in at least one channel that is controlled by

each memory controller is allocated to the CPU. The allocation of the memory space is performed in a one-to-one corresponding relationship as the pinned memory of the CPU.

[031] For example, the corresponding memory spaces of 4 DIMM slots in 2 channels controlled by each memory controller are allocated to the CPU that is in a one-to-one corresponding relationship with the memory controller, and the corresponding memory spaces being set as the pinned memory of the CPU, which is shown in Table 1 including:

[032] The corresponding memory space of 3 DIMM slots (DIMM 1, DIMM 2 and DIMM 3) in the first channel (channel 01) under each memory controller and the corresponding memory space of 1 DIMM slot (DIMM 1) in the second channel (channel 02) are allocated respectively to each CPU of “n” CPUs as the pinned memory of each CPU in “n” CPUs.

[033] The unallocated memory space, which remains in the memory pool, is set as the shared memory pool. The shady parts in the Table 1 including the corresponding memory spaces of DIMM 2 and DIMM 3 in channel 02, and the corresponding memory spaces of DIMM 1, DIMM 2, and DIMM 3 in channel 03 and channel 04 controlled by each memory controller are set as the shared memory pool.

Table 1

CPU-QPI port address	Memory controller port address	Channel	DIMM slot
1	1	01	DIMM 1
			DIMM 2
			DIMM 3
		02	DIMM 1
			DIMM 2
			DIMM 3
		03	DIMM 1
			DIMM 2
			DIMM 3

		04	DIMM 1
			DIMM 2
			DIMM 3
2	2	01	DIMM 1
			DIMM 2
			DIMM 3
		02	DIMM 1
			DIMM 2
			DIMM 3
		03	DIMM 1
			DIMM 2
			DIMM 3
		04	DIMM 1
			DIMM 2
			DIMM 3
3	3
...
n	n
	n+1
	n+2
	n+3

	m

[034] In the above embodiment according to the disclosure, the address of the memory space of the memory chip is a fixed address. That is to say the address of the corresponding memory space of the DIMM slot is a fixed address.

[035] There are other possible embodiments to implement the addresses of the corresponding memory space of a DIMM slot. An example is allocating a real time

address by the server system. Other address implementation techniques will not be explained into detail here.

[036] In step S103, a memory value that represents usage of the respective memory space by the respective CPU during operation is obtained.

[037] In this embodiment, the memory values that represent usage of the respective memory space by each CPU during operation is obtained for each of the above-described “n” CPUs.

[038] It should be noted that the memory value used by each CPU during operation of all “n” CPUs obtained by the operating system in the server system includes the memory value used by the operating system and the memory value used by application programs.

[039] Other than the above-mentioned ways to obtain the used memory values, there are additional ways that are not going to be described in detail here.

[040] In this embodiment, the memory values that represent usage of the respective memory space by the respective CPU during operation in “n” CPUs are obtained. These memory values are obtained by setting each CPU of “n” CPUs as a determining unit to determine if the memory value used by the CPU exceeds a preset threshold range. According to the determination, further operation is taken such as whether to allocate additional memory space from the shared memory pool to the CPU, or to free part of original memory space that has been allocated to the CPU.

[041] Besides the method mentioned above, there are other ways for determining the memory space needed. For example, to obtain memory values used by each CPU during operation of “n” CPUs, all the memory values used by at least one CPU during operation of “n” CPUs can be summed up to obtain an overall memory value used by at least one CPU during operation. Then, it is determined if the overall memory value used by at least one CPU during operation exceeds a preset threshold range. According to the determination, either additional memory space is selected from the shared memory pool, which is allocated to each of the CPUs, or the memory space is freed or released.

[042] In step S104, a determination is made as to whether the memory value exceeds a preset threshold range.

[043] In this step, the memory values of each CPU during operation of “n” CPUs are used to determine if each memory value exceeds a preset threshold range of each CPU.

[044] As an example, one CPU of “n” CPUs is used in the process to determine if the memory value exceeds a preset threshold range of the CPU, including:

[045] (1) Determining if the memory value used by the CPU during operation exceeds a preset first threshold.

[046] If the memory value used by the CPU during operation exceeds a preset first threshold, then step S105 is initiated to apply for additional memory space from the shared memory pool according to a certain proportion of the memory value, and to allocate the additional memory space to the corresponding CPU.

[047] If the memory value used by the CPU during processing does not exceed a preset first threshold, a next step is initiated to determine if the memory value used by the CPU during processing is less than a preset second threshold.

[048] (2) Determining if the memory value used by the CPU during operation is less than the preset second threshold.

[049] If the memory value used by the CPU during operation is less than the preset second threshold, in step S105, the memory space allocated to the CPU is partially released according to a certain proportion of the memory value. The process of releasing memory continues until the memory space allocated to the CPU reaches the pinned memory.

[050] If the memory value used by the CPU during operation is not less than a preset second threshold, the process returns to step S103 to obtain the memory values used by each of the CPUs of “n” CPUs.

[051] In addition to these solutions, there are other ways to implement the steps. For example, it can be determined in turn for “n” CPUs if the memory values used by each CPU during operation of “n” CPUs exceed a preset first threshold, and/or it can be determined in turn for “n” CPUs if the memory values used by each CPU during operation of “n” CPUs is less than a preset second threshold. Then the mentioned determining steps are repeated “n” times until all “n” CPUs have been processed by the determining process. This repetition process will not be discussed in further detail here.

[052] For the embodiment according to the disclosure, corresponding to step S103, to obtain the memory values used by each CPU during operation of “n” CPUs, all the memory values used by at least one CPU during operation of “n” CPUs are summed up to obtain an overall memory value used by at least one CPU during processing. Next, it is determined if the overall memory value used by at least one CPU during operation exceeds a preset threshold range. According to the determination, either additional memory space is selected from the shared memory pool, which is allocated to each of the CPUs, or the memory space is freed or reallocated.

[053] In step S104, the determining of whether the overall memory value used by at least one CPU during operation exceeds a preset threshold range can be implemented as described next.

[054] (1) Determining if the overall memory value used by at least one CPU of the “n” CPUs during operation exceeds a preset first threshold.

[055] If the overall memory value used by at least one CPU during operation exceeds a preset first threshold, at least one additional memory space is applied from the memory pool according to a certain proportion of the memory value such that to be allocated to each CPU in the at least one CPU in a one-to-one corresponding manner.

[056] If the overall memory value used by at least one CPU during operation does not exceed a preset first threshold, a next step is initiated to determine if the overall memory value used by at least one CPU during operation is less than a preset second threshold.

[057] (2) Determining if the overall memory value used by at least one CPU during operation is less than the preset second threshold.

[058] If the overall memory value used by at least one CPU during operation is less than the preset second threshold, the memory space allocated to each CPU in the at least one CPU is partially released according to a certain proportion of the memory value. The process of releasing memory continues until the memory spaces allocated to each CPU in the at least one CPU reaches the size of the pinned memory.

[059] If the overall memory value used by at least one CPU during operation is not less than the preset first threshold, the process returns to step S103 to obtain the memory values used by each CPU of “n” CPUs.

[060] In step S105, additional memory space is applied from the memory pool to allocate to the CPU, or the memory space allocated to the CPU is reallocated or freed.

[061] A precondition for implementing step S105 is that, in step S104, it is determined the memory value used by a CPU during processing exceeds a preset first threshold, or the memory value used by a CPU during processing is less than a preset second threshold.

[062] If the memory value used by a CPU during processing exceeds the preset first threshold, one additional memory space is applied from the memory pool according to a certain proportion of the memory value to be allocated to the CPU.

[063] If the memory value used by a CPU during operation is less than the preset second threshold, the memory space allocated to the CPU is partially released until the memory spaces allocated to the CPU reaches the size of the pinned memory.

[064] It should be noted that, in this embodiment, when applying additional memory spaces from the memory pool according to a certain proportion of the memory value to be allocated to the CPU, the corresponding memory space of the DIMM slot in the same channel or the corresponding memory space of DIMM slot in the same memory controller will be allocated to the CPU, as described next.

[065] (1) Initially, determining if there exists a DIMM slot in the memory pool located in the same channel controlled by the same memory controller as the DIMM slot that corresponds to the pinned memory of the CPU.

[066] If there is such a DIMM slot in the memory pool located in the same channel under the same memory controller as the DIMM slot that corresponds to the pinned memory, then the corresponding memory space of the DIMM slot is allocated to the CPU according to a certain proportion of the memory value.

[067] For example, in a memory pool for the corresponding memory space of the DIMM slot in the same channel controlled by the same memory controller with the DIMM slot corresponding to the pinned memory of the CPU, the size of the memory space applied from the memory space of the corresponding DIMM slot will be equivalent to 20% of the memory value used by the CPU during operation. This is the amount to be allocated to the CPU.

[068] If there does not exist a DIMM slot in the memory pool located in the same channel under the same memory controller as the DIMM slot that corresponds to the pinned memory, a next step is initiated to determine if there exists a DIMM slot in the memory pool located in the same memory controller as the DIMM slot that corresponds to the pinned memory, which is described in the following.

[069] (2) It is determined if there is a DIMM slot in the memory pool locating in the same memory controller as the DIMM slot that corresponds to the pinned memory. If there is a DIMM slot in the memory pool located in the same memory controller as the DIMM slot that corresponds to the pinned memory, the corresponding memory space of the DIMM slot is allocated to the CPU according to a certain proportion of the memory value.

[070] For example, in the memory pool, for the corresponding memory space of the DIMM slot in the same memory controller with the DIMM slot corresponding to the pinned memory of the CPU, the size of the memory space applied from the memory space of the corresponding DIMM slot will be equivalent to 20% of the memory value used by the CPU during operation, which is allocated to the CPU.

[071] If there does not exist a DIMM slot in the memory pool located in the same memory controller as the DIMM slot that corresponds to the pinned memory, the memory space in the shared memory pool is allocated to the CPU according to a certain proportion of the memory value.

[072] For example, the memory space applied from the shared memory pool will be equivalent to the size of 20% of the memory value used by the CPU during operation, which is allocated to the CPU.

[073] There are other possible embodiments to implement the above-described application of memory space from the shared memory pool for allocation to a CPU, which will not be explained in detail herein.

[074] Similarly, the step of releasing or reallocating memory spaces from the memory pool according to a certain proportion of the memory value may be implemented in other ways corresponding to the above mentioned method.

[075] For example, for the corresponding memory space of the DIMM slot in the same memory controller with the DIMM slot corresponding to the pinned memory

of the CPU, the size of the memory space released will be equivalent to 20% of the memory value used by the CPU during operation.

[076] There are other possible methods to implement the releasing or reallocating of the memory space, which will not be explained in detail herein.

[077] It should be noted that steps S103, S104, and S105 are in a cyclic process of determination. The memory values used by each CPU during operation of “n” CPUs are obtained in step S103. Whether the memory values used by each CPU during operation of “n” CPUs exceeds a preset threshold range is determined in step S104. Based on the result of step S104, in step S105, additional memory space is selected from the shared memory pool to be allocated to the CPU or the memory space allocated to the CPU is reallocated or freed.

[078] The cyclic determination process of steps S103, S104, and S105 implement a real time monitoring of the memory values used by the CPU during processing. According to the result of the monitoring, certain memory space is allocated to the CPU or certain memory space of the CPU is released such that the memory is utilized in a proper way and the functionality of the server system is enhanced.

[079] Further, corresponding to the steps S103 and S104, according to the obtained memory values used by each CPU during operation of “n” CPUs, all the memory values used by at least one CPU during operation of “n” CPUs are summed up to obtain an overall memory value used by at least one CPU during operation. Next, it is determined if the overall memory value used by at least one CPU during operation exceeds a preset threshold range. According to the result of the determination, either additional memory space is selected from the shared memory pool, which is allocated to each CPU in the at least one CPU, or the memory space is reallocated or freed.

[080] The steps of applying for additional memory space from the shared memory pool to allocate to each CPU in the at least one CPU are further described next.

[081] (1) First, it is determined, in turn, if there is exists a DIMM slot in the shared memory pool located in the same channel as each DIMM slot that corresponds to a pinned memory of each CPU in the at least one CPU.

[082] If there is such a DIMM slot in the shared memory pool located in the same channel as each DIMM slot that corresponds to a pinned memory of each CPU in

the at least one CPU, the corresponding memory space of the DIMM slot is then allocated to the CPU according to a certain proportion of the memory value.

[083] If there does not exist a DIMM slot in the shared memory pool locating in the same channel as a DIMM slot that corresponds to a pinned memory of each of the CPUs, a next step is initiated to determine if there exists a DIMM slot in the shared memory pool located in the same memory controller as the DIMM slot that corresponds to a pinned memory of each CPU in the at least one CPU.

[084] (2) It is determined if there is a DIMM slot in the shared memory pool located in the same memory controller as the DIMM slot that corresponds to a pinned memory of each CPU of the set of CPUs.

[085] If there is a DIMM slot in the shared memory pool located in the same memory controller as the DIMM slot that corresponds to a pinned memory of each CPU of the set of CPUs, the corresponding memory space of the DIMM slot is allocated to the CPU according to a certain proportion of the memory value.

[086] If such DIMM slot does not exist in the shared memory pool locating in the same memory controller as the DIMM slot that corresponds to a pinned memory of each CPU in the at least one CPU, the memory space in the shared memory pool is allocated to the CPU according a certain proportion of the memory value.

[087] The above-described process of determination is repeated until memory is allocated from each of the at least one CPU. Also, other embodiments are possible for implementation without limitation to the above-described process.

[088] FIG. 1B shows a flow diagram of a method of memory management according to another embodiment of the disclosure. In step S110, at least one memory is pooled to generate a memory pool. In step S130, the memory pool is divided to generate at least one memory space. In step S140, each of the memory spaces is allocated respectively to a plurality of CPUs in a one-to-one correspondence manner. In step S141, the one-to-one allocated respective memory space is allocated to a respective CPU as a pinned memory of the respective CPU. In step S142, unallocated memory space is set as a shared memory pool of the memory pool. That is, the unallocated memory space includes memory space that remains unallocated in the memory pool.

[089] In step S150, a memory value that represents usage of the respective memory space by the respective CPU during operation is obtained. In step S160, it is determined as to whether the memory value obtained in step S150 exceeds a preset threshold range. If the memory value used by the CPU during operation exceeds a preset first threshold, then step S171 is initiated to apply for additional memory space from the shared memory pool according to a certain proportion of the memory value and to allocate the additional memory space to the corresponding CPU. If the memory value used by the CPU during operation is less than the preset second threshold, in step S172, the memory space allocated to the CPU is partially released according to a certain proportion of the memory value.

[090] In FIG. 2 there is shown another embodiment of the disclosure of a device 200 of memory management that includes the following units.

[091] A memory pooling unit 210 configured to pool at least one memory to generate a memory pool.

[092] A memory divider 220 divides the memory pool to generate at least one memory space and allocates a respective memory space to a respective CPU in a one-to-one correspondence manner. The respective memory space allocated to a respective CPU is set as a pinned memory of the respective CPU. An unallocated memory space is set as a shared memory pool of the memory pool. That is, the unallocated memory space includes memory space that remains unallocated or is not pinned memory in the memory pool.

[093] A memory value obtaining unit 230 obtains a memory value that represents usage of the respective memory space by the respective CPU during operation. A memory value determination unit 240 determines if the memory value exceeds a preset threshold range and whether to commence a memory manager 250 if the memory value exceeds the preset threshold range.

[094] The memory manager 250 is configured for selecting additional memory space from the shared memory pool to allocate to the CPU and/or releasing or reallocate the memory space allocated to the CPU if the threshold is exceeded.

[095] Alternatively, the memory divider 220 includes a memory controller selecting sub-module configured for selecting a set of memory controllers, the number of

memory controllers is equal to the number of the set of CPUs. It also includes an allocator sub-module configured for matching a QPI port address of the CPU with a port address of the memory controller in a one-to-one correspondence manner and for allocating the corresponding memory space of at least one DIMM slot in at least one channel under the memory controller to the CPU. The corresponding memory is set as a pinned memory of the CPU and the corresponding memory address of the DIMM slot is a fixed address.

[096] Alternatively, the device 200 also includes a CPU pooling unit configured for setting each CPU in a set of CPUs as a node of the set of CPUs and for connecting all the nodes of the set of CPUs to generate a CPU pool.

[097] In a typical configuration, a computing device includes one or more CPUs, an input/output port, an Internet port, and a memory.

[098] According to the disclosure, the computer-readable media do not include the transitory media, such as the modulate data signal and carrier. The computer-readable media disclosed include, but not limit to, the phase-change memory (PRAM), static random-access memory (SRAM), dynamic random-access memory (DRAM), other types of the random-access memory (RAM), read-only memory (ROM), electrically erasable programmable read-only memory (EEPROM), flash RAM and other memory technologies, compact disc read-only memory (CD-ROM), digital video disc (DVD) and other optical storage, magnetic tape, magnetic disc, other magnetic storage, and any other non-transition media

[099] It is appreciated that the skilled in the art understands the disclosure may be implemented as the methods, systems, and/or the instructions for a computer. It is intended that the disclosure may be implemented as hardware, software, and/or hardware and software combined. The disclosure may be implemented as a computer program product utilizing one or more storage mediums including computer program instructions. The storage medium includes, but not limit to, the phase-change memory (PRAM), static random-access memory (SRAM), dynamic random-access memory (DRAM), other types of the random-access memory (RAM), read-only memory (ROM), electrically erasable programmable read-only memory (EEPROM), flash RAM and other memory technologies, compact disc read-only memory (CD-ROM), digital video disc (DVD) and

other optical storage, magnetic tape, magnetic disc, other magnetic storage, and any other storage type.

[0100] Although certain embodiments and methods have been disclosed herein, it will be apparent from the foregoing disclosure to those skilled in the art that variations and modifications of such embodiments and methods may be made without departing from the spirit and scope of the disclosure. It is intended that the disclosure be defined by the appended claims and the rules and principles of applicable law.

What is claimed is:

1. A method for memory management, comprising:

pooling at least one memory to generate a memory pool;
dividing the memory pool to generate at least one memory

5 space;

allocating a respective memory space of the at least one
memory space to a respective CPU of a plurality of CPUs in a one-to-one
correspondence manner and setting the respective memory space as a pinned
memory of the respective CPU;

10 setting an unallocated memory space as a shared memory pool
of the memory pool;

obtaining a memory value that represents usage of the
respective memory space by the respective CPU during operation;

15 determining if the memory value exceeds a preset threshold
range; and

based on whether the memory value exceeds the preset
threshold range, selecting an additional memory space from the memory pool
to allocate to the respective CPU, or reallocating at least a portion of the
respective memory space that has been allocated to the respective CPU.

20 2. The method of claim 1, further comprising:

repeating the obtaining a memory value, the determining, and
the selecting an additional memory space.

3. The method of claim 1, further comprising, in response to determining the
memory value does not exceed the preset threshold range:

25 repeating the obtaining a memory value, the determining, and
the reallocating.

4. The method of claim 1, further comprising:

selecting a number of memory controllers, wherein the
number of memory controllers is equal to a number of the respective CPUs;

30 matching in a one-to-one correspondence manner port
addresses of the respective CPUs to port addresses of the memory controllers;
and

allocating a corresponding memory space of at least one dual
in-line memory module (DIMM) slot of the memory pool in at least one
35 channel associated with the memory controllers to the respective CPU as the
pinned memory, wherein the allocated corresponding memory space of the
DIMM slot includes a fixed address.

5. The method of claim 4, wherein the obtaining a memory value comprises:

obtaining the memory values of the respective CPUs from the
40 plurality of CPUs.

6. The method of claim 5, wherein the determining if the memory value exceeds a
preset threshold range comprises:

determining if at least one of the memory values exceeds a
preset first threshold; and

45 if the memory value exceeds the preset first threshold,
selecting the additional memory space from the memory pool according to a
certain proportion of the memory value and allocating the additional memory
space to the respective CPU.

7. The method of claim 6, further comprising:

50 determining, if the memory value does not exceed a preset
first threshold, if the memory value is lower than a preset second threshold;
and

if the memory value does not exceed the preset first threshold
and if the memory value is lower than the preset second threshold,

55 reallocating at least a portion of the respective memory space allocated to the
respective CPU according to a certain proportion of the memory value.

8. The method of claim 7, further comprising:

if the memory value is not lower than the preset second
threshold, repeating the obtaining a memory value, the determining if the
60 memory value exceeds the preset threshold range, and the based on whether

the memory value exceeds the preset threshold range, selecting an additional memory space from the memory pool to allocate to the respective CPU or reallocating the respective memory space allocated to the respective CPU.

9. The method of claim 6, wherein the selecting additional memory space from the memory pool comprises:

determining if there exists a second DIMM slot of the memory pool located in the same channel as the DIMM slot that corresponds to the pinned memory of the respective CPU; and

if the second DIMM slot exists, allocating a second corresponding memory space of the second DIMM slot to the respective CPU according to the certain proportion of the memory value.

10. The method of claim 9, further comprising:

if the second DIMM slot does not exist, determining if there exists a third DIMM slot in the memory pool located in the memory controllers as the DIMM slot that corresponds to the pinned memory of the respective CPU;

if the third DIMM slot exists, allocating a third corresponding memory space of the third DIMM slot to the respective CPU according to the certain proportion of the memory value; and

if the third DIMM slot does not exist, allocating the additional memory space from the memory pool to the respective CPU according to another certain proportion of the memory value.

11. The method of claim 5, wherein the obtaining the memory values comprises:

summing the memory values of the respective CPUs from the plurality of CPUs to obtain a total memory value.

12. The method of claim 11, wherein the determining if the memory value exceeds a preset threshold range comprises:

determining if the total memory value exceeds a preset first threshold; and

if the total memory value exceeds the preset first threshold, selecting the additional memory space from the memory pool according to a

certain proportion of the memory value and allocating the additional memory space to the respective CPU.

13. The method of claim 12, further comprising:

95 determining, if the total memory value does not exceed a preset first threshold, if the total memory value is lower than a preset second threshold;

 if the total memory value does not exceed the preset first threshold and if the total memory value is lower than the preset second
100 threshold, reallocating at least a portion of the respective memory space allocated to the respective CPU according to a certain proportion of the memory value; and

 if the memory value is not lower than the preset second
105 threshold, repeating the obtaining a memory value, the determining if the memory value exceeds the preset threshold range, and the based on whether the memory value exceeds the preset threshold range, selecting an additional memory space from the memory pool to allocate to the respective CPU or reallocating the respective memory space allocated to the respective CPU.

14. The method of claim 12, wherein the selecting additional memory space from the
110 memory pool comprises:

 determining if there exists a second DIMM slot of the shared memory pool located in the same channel as the DIMM slot that corresponds to the pinned memory of the respective CPU;

 if the second DIMM slot exists, allocating a second
115 corresponding memory space of the second DIMM slot to the respective CPU according to the certain proportion of the memory value;

 if the second DIMM slot does not exist, determining if there
120 exists a third DIMM slot in the shared memory pool located in the respective memory controller as the DIMM slot that corresponds to the pinned memory of the respective CPU;

if the third DIMM slot exists, allocating a third corresponding memory space of the third DIMM slot to the respective CPU according to the certain proportion of the memory value; and

125 if the third DIMM slot does not exist, allocating the additional memory space from the memory pool to the respective CPU according to the certain proportion of the memory value.

15. The method of claim 1, further comprising:

130 setting each CPU in the plurality of CPUs as a node of the plurality of CPUs and connecting the nodes of the plurality of CPUs to generate a CPU pool.

16. An apparatus for memory management, comprising:

a memory pooling unit configured to pool at least one memory to generate a memory pool;

135 a memory divider configured to divide the memory pool to generate at least one memory space, to allocate a respective memory space to a respective CPU of a plurality of CPUs in a one-to-one correspondence manner, wherein the respective memory space allocated to the respective CPU is set as a pinned memory of the respective CPU, and to set an unallocated memory space as a shared memory pool of the memory pool;

140 a memory value obtaining unit configured to obtain a memory value that represents usage of the respective memory space by the respective CPU during operation; and

145 a memory value determination unit configured to determine if the respective memory value exceeds a preset threshold range to commence a manager,

wherein the manager is configured to select additional memory space from the memory pool to allocate to the respective CPU or to reallocate at least a portion of the respective memory space allocated to the respective CPU.

150 17. The apparatus of claim 16, wherein the memory divider comprises:

a selector configured to select a set of memory controllers,
wherein a number of the set of memory controllers is equal to the number of
the plurality of CPUs; and

155 an allocator configured to match a port address of the
respective CPU with the port address of the memory controller in a one-to-
one correspondence manner and to allocate the respective memory space that
corresponds to at least one dual in-line memory module (DIMM) slot in at
least one channel controlled by the memory controller to the respective CPU,
wherein the respective memory space is set as a pinned memory of the
160 respective CPU, and wherein a memory address that corresponds to the
DIMM slot includes a fixed address.

18. The apparatus of claim 16, further comprising:

a CPU pooling unit configured to set each CPU of the plurality
of CPUs as a node of the plurality of CPUs and to connect the nodes of the
165 plurality of CPUs to generate a CPU pool.

19. A non-transitory computer readable storage medium having embedded therein
program instructions, when executed by one or more processors of a device,
manages memory space, comprising:

170 pooling at least one memory to generate a memory pool;
dividing the memory pool to generate at least one memory
space;

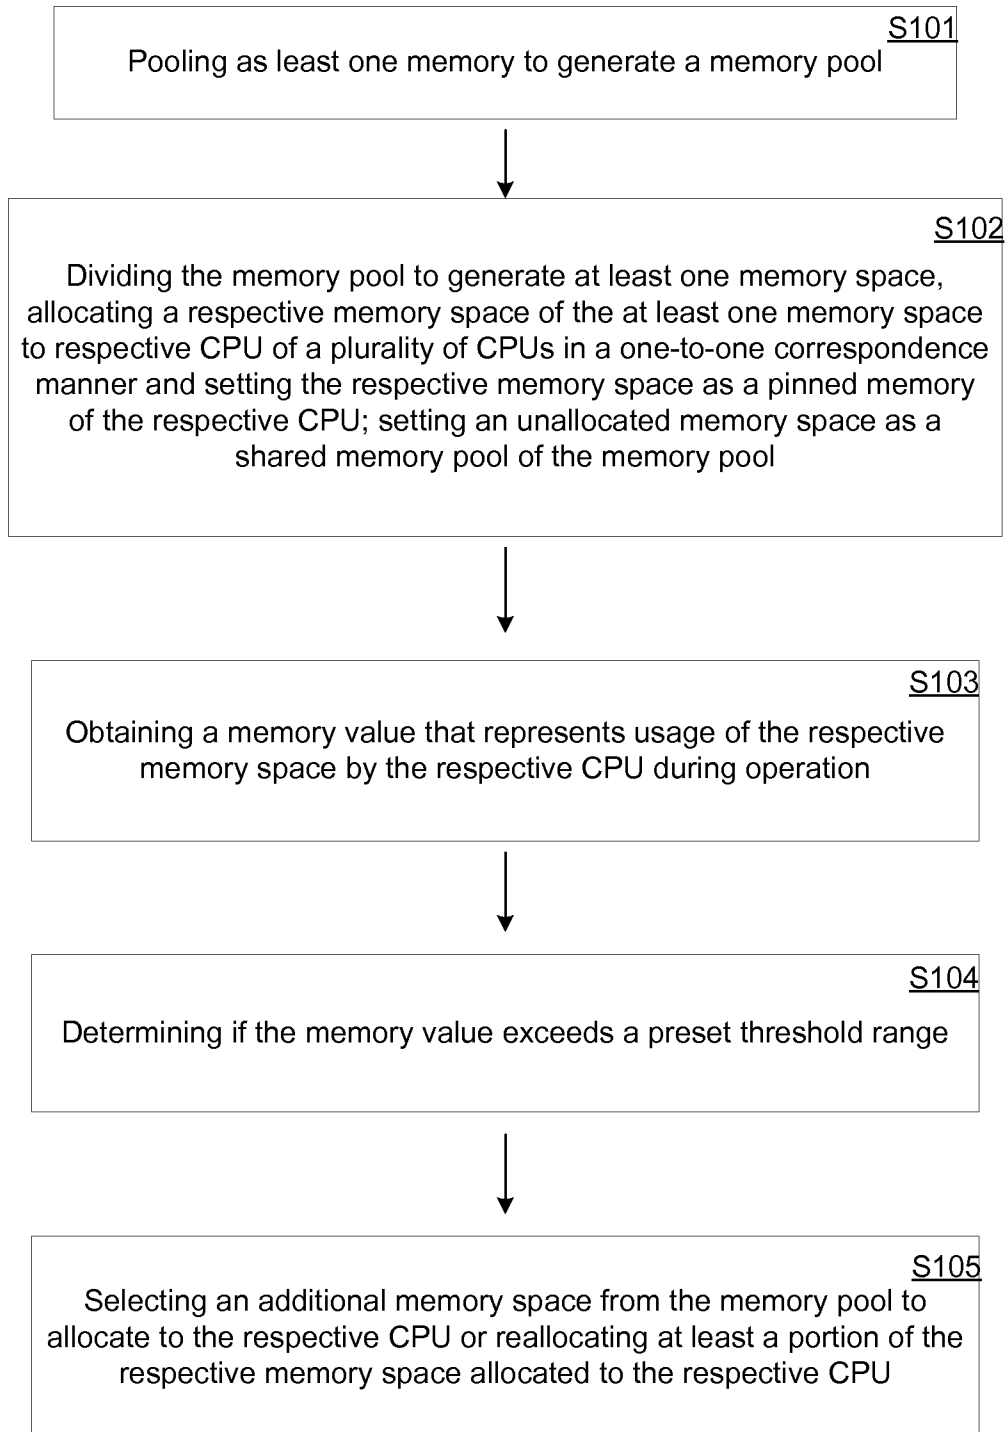
allocating a respective memory space of the at least one
memory space to a respective CPU of a plurality of CPUs in a one-to-one
175 correspondence manner and setting the respective memory space as a pinned
memory of the respective CPU;

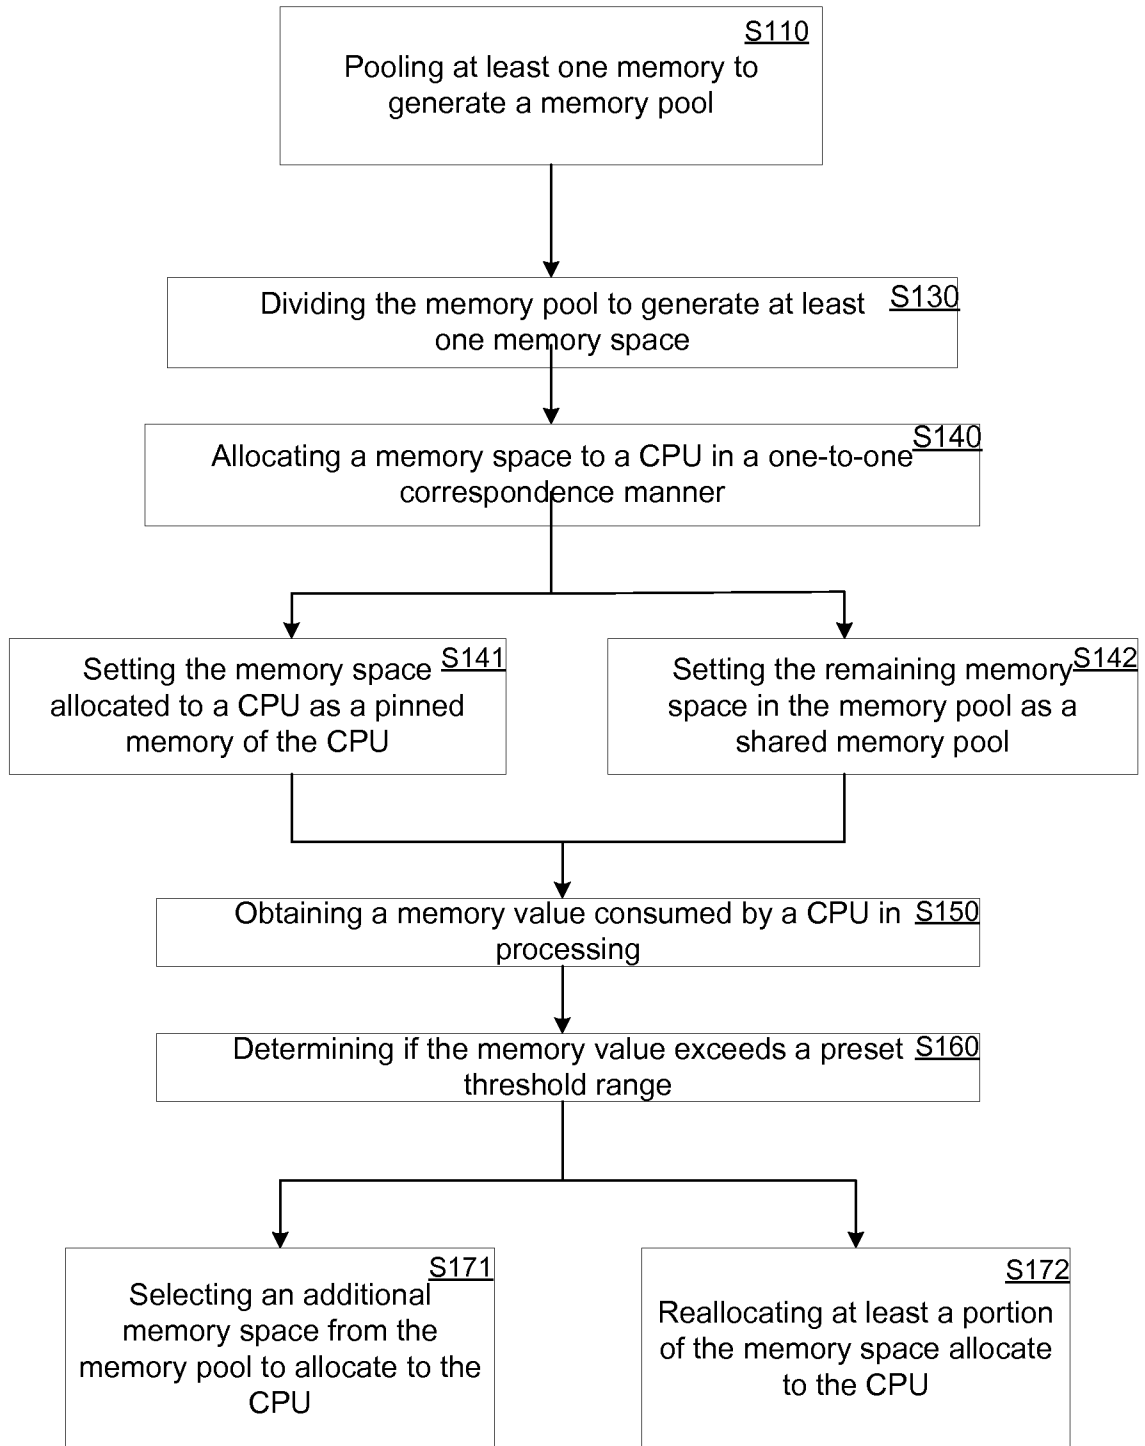
setting an unallocated memory space as a shared memory pool
of the memory pool;

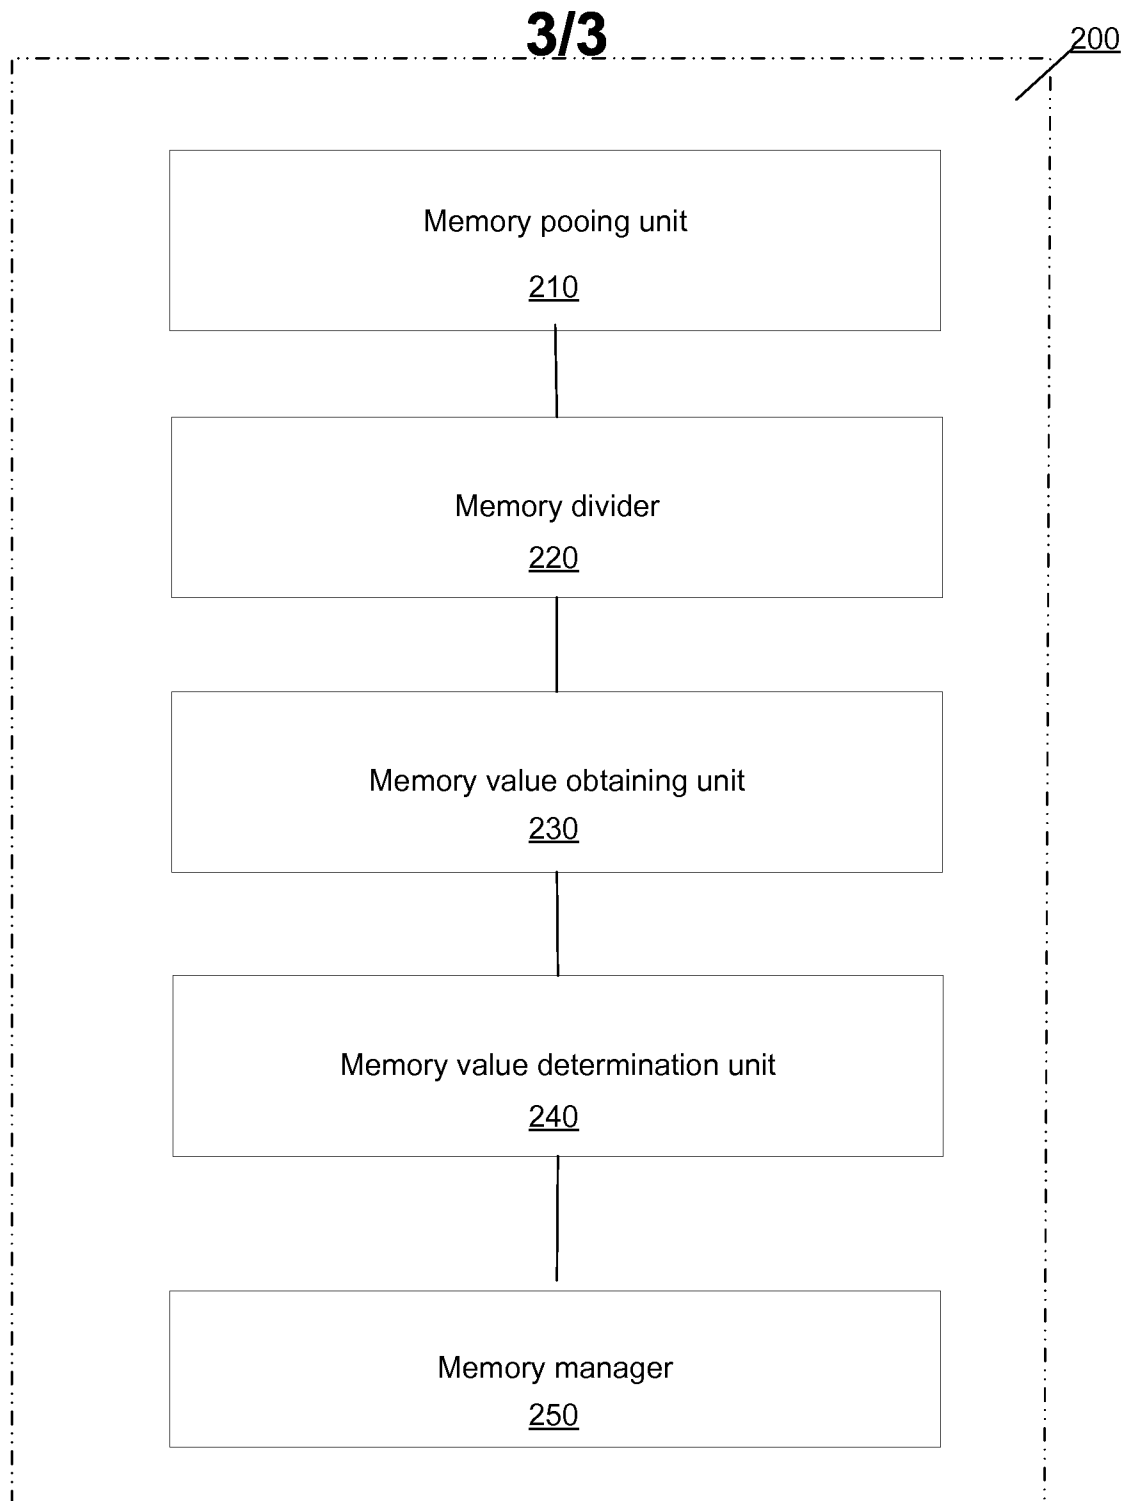
180 obtaining a memory value that represents usage of the
respective memory space by the respective CPU during operation;

determining if the memory value exceeds a preset threshold range; and

185 based on whether the memory value exceeds the preset threshold range, selecting an additional memory space from the memory pool to allocate to the respective CPU or reallocating at least a portion of the respective memory space that has been allocated to the respective CPU.

1/3**FIG. 1A**

2/3**FIG. 1B**

**FIG. 2**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US15/62801

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 12/02 (2016.01)

CPC - G06F 12/0263, 9/5016, 9/5055

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8): G06F 12/00, 12/02, 12/06, 9/50 (2016.01)

CPC: G06F 12/00, 12/02, 12/0284, 12/06, 12/0623, 9/50, 9/5016, 9/5055

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, IN, CA, INPADOC Data); ProQuest; Google Patents; Google Scholar; IEEE
 Keywords: memory, management, pool, threshold, exceed, greater, allocate, processor, CPU, match, map, correlate, DIMM, controller, port, address, card, slot

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2005/0257020 A1 (KASHYAP, S. et al.) 17 November 2005; figures 4, 5; paragraphs [0010], [0031], [0032], [0036], [0041].	1-19
Y	US 2013/0138840 A1 (KEGEL, A. et al.) 30 May 2013; figure 3; paragraphs [0073], [0093], [0147].	1-19
Y	US 2013/0275707 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 17 October 2013; figures 2, 3, 4; paragraphs [0049], [0050].	4-15, 17-18
Y	US 2004/0243884 A1 (VU, P. et al.) 02 December 2004; figure 1; paragraphs [0008], [0021], [0024], [0034]; claim 1.	9-10, 14



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

12 January 2016 (12.01.2016)

Date of mailing of the international search report

03 FEB 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300
 PCT OSP: 571-272-7774