

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2016/077491 A1

(43) International Publication Date

19 May 2016 (19.05.2016)

(51) International Patent Classification:

G06F 1/32 (2006.01) G11B 19/04 (2006.01)

(21) International Application Number:

PCT/US2015/060213

(22) International Filing Date:

11 November 2015 (11.11.2015)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

201410635703.1 12 November 2014 (12.11.2014) CN

(71) Applicant: **ALIBABA GROUP HOLDING LIMITED** [—/US]; Fourth Floor, One Capital Place, P.O. Box 847, George Town, Grand Cayman (KY).

(72) Inventors: **YANG, Chunxin**; Alibaba Group Legal Department, 5/f, Building 3, No. 969 West Wen Yi Road, Yu Hang District, Hangzhou 311121 (CN). **JIANG, Xiong-wei**; 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.

(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))

(54) Title: METHOD AND APPARATUS FOR CONTROLLING HARD DRIVE POWER CONSUMPTION AND CONTROLLING A MANAGEMENT SERVICE OF HARD DRIVE POWER CONSUMPTION

DETECTING WHETHER A FAULT EXISTS IN A HARD DRIVE IN A NODE HARD DRIVE GROUP

S101

IF THERE IS NO FAULT DETECTED IN THE HARD DRIVE IN THE NODE HARD DRIVE GROUP, CONTROLLING THE HARD DRIVE IN THE NODE HARD DRIVE GROUP TO CONVERT THE HARD DRIVE BETWEEN PRESET MODES OF POWER CONSUMPTION USING A HARD DRIVE POWER CONSUMPTION MANAGEMENT SERVICE

S102

FIG. 1

(57) Abstract: A method and apparatus for controlling hard drive power consumption and controlling a hard drive power consumption management service are provided. The method includes detecting whether a fault exists in a hard drive in a node hard drive group. The method further includes, if there is no fault detected in the hard drive in the node hard drive group, controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption by a hard drive power consumption management service.

WO 2016/077491 A1

METHOD AND APPARATUS FOR CONTROLLING HARD DRIVE POWER CONSUMPTION  
AND CONTROLLING A MANAGEMENT SERVICE OF HARD DRIVE POWER  
CONSUMPTION

CROSS REFERENCE TO RELATED APPLICATION

[001] This application claims priority to and the benefit of Chinese Patent Application No. 201410635703.1, filed on November 12<sup>th</sup>, 2014, which is incorporated herein by reference in its entirety.

TECHNICAL FIELD

[002] The disclosure relates generally to the field of computer, and more particularly, to method and apparatus for controlling hard drive power consumption and controlling a management service of hard drive power consumption.

BACKGROUND

[003] Under the scenarios of storing and applying data which is not frequently accessed (known as cold data), the accessing mode to a hard drive is mainly in the form of writing to the hard drive, the hard drive after being written to, is generally not accessed or read again. After the hard drive is full, it operates in an “idle mode” during which it consumes considerable power.

[004] During such idle mode, a method of controlling the hard drive power consumption is usually performed by changing the power consumption mode of the hard drives of the node hard drive group from a high power consumption state to a low power consumption state through a programming of commands of reading and writing on the disk. However, this method of controlling the hard drive power consumption does not address power consumption states of changing the mode of the hard drive to perform the conversion between the high power consumption state and the low power consumption state.

[005] Generally, when controlling hard drive power consumption, the conversion states of controlling power consumption of hard drive include four modes of controlling, which are described next.

[006] An active mode: a hard drive is under a full speed working state, a disk of the hard drive is running at a high speed and a response time to a user request is the shortest in this mode.

[007] An idle mode: a disk of a hard drive maintains in operating state, a head arm in the hard drive stops running, most other electronic components are under inactivate state, and power consumption is lower in this mode than in an active mode.

[008] A standby mode: a hard drive stops running, the hard drive is under low power consumption state.

[009] A sleep mode: a hard drive is under inactive state and power consumption is the lowest in this mode compared to other modes.

[010] When converting a hard drive from a low power consumption state to high power consumption state, the lower the power consumption is of the hard drive, the higher the power loss is and the longer the time it takes to transition between states which additionally results in delays during this transition. Even though power loss is occurring, it is not always possible to perform conversion directly between the low power consumption state and the high power consumption state to immediately reduce the power loss occurring. The reason is because a hard drive needs to transit through certain modes such as a standby mode when changing from a sleep mode to an active mode. The transition of modes causes time delay, as well as can result in high power loss.

[011] Therefore, a conventional method of controlling hard drive power consumption has drawbacks of a time delay during conversion because of necessary transitions, resulting in wasted resources.

## SUMMARY OF THE DISCLOSURE

[012] It would be advantageous to provide a method and an apparatus for controlling hard drive power consumption to save time and power.

[013] Embodiments of the disclosure provide methods and apparatus for controlling hard drive power consumption and controlling a hard drive power consumption management service. The disclosure addresses the issues of the conventional implementations by providing desirable techniques to address the shortcomings of power losses during converting modes of hard drives and the controlling of power consumption of hard drive during operation.

[014] Accordingly, an embodiment according to the disclosure provides a process to control hard drive power consumption. The process includes detecting if a fault exists in a hard drive in a node hard drive group. If there is no fault in any hard drive in the node hard drive group, the process further includes controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption by a hard drive power consumption management service.

[015] Accordingly, an embodiment according to the disclosure provides a process to control a hard drive power consumption management service. The process includes detecting an operating environment of a hard drive power consumption management service. Further, the process includes, if, the operating environment of the hard drive power consumption management service meets preconditions, controlling a start/stop operation the hard drive power consumption management service.

[016] Accordingly, an embodiment according to the disclosure provides a mechanism to control hard drive power consumption. There is a hard drive detector configured for detecting fault in a hard drive of a node hard drive group. A power consumption controller module actuates, if there is no fault in the hard drive in the node hard drive group, the hard drive of the node hard drive group to convert between preset modes of power consumption.

[017] Accordingly, an embodiment according to the disclosure provides another mechanism to control a hard drive power consumption management service. The mechanism includes an environment detector for detecting an operating environment of a hard drive power consumption management service. Further, the mechanism includes a process controller to control a start/stop operation of the hard drive power consumption management service in accordance with parameters sent from an operating system if the operating environment of the hard drive power consumption management service meets preconditions.

[018] Accordingly, an embodiment according to the disclosure provides a system to control hard drive power consumption. The system includes instructions to detect if there is a fault in a hard drive in a node hard drive group. The system further includes instructions to control the hard drive in the node hard drive group to convert between preset modes of power consumption through a hard drive power consumption management service if there is no fault in the hard drive in the node hard drive group.

[019] Accordingly, an embodiment according to the disclosure provides a system to control a management service of hard drive power consumption. The system includes instructions to

detect an operating environment of a hard drive power consumption management service. The system further includes instructions to control a start/stop operation of the hard drive power consumption management service if the operating environment of the hard drive power consumption management service meets preconditions.

[020] 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

[021] 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.

[022] Figure 1 is a flow diagram of a method of controlling power consumption of hard drive in accordance with an embodiment.

[023] Figure 2 is a flow diagram of another method of controlling power consumption of hard drive in accordance with an embodiment.

[024] Figure 3A is a flow diagram of a method of controlling a hard drive power consumption management service in accordance with an embodiment.

[025] Figure 3B is a flow diagram of a method of controlling a hard drive power consumption management service in accordance with an embodiment.

[026] Figure 4 is a block diagram of an apparatus for controlling power consumption of hard drive in accordance with an embodiment.

[027] Figure 5 is a block diagram of a service controlling module for controlling power consumption in accordance with an embodiment.

[028] Figure 6 is a block diagram of an apparatus for controlling a hard drive power consumption management service in accordance with an embodiment.

## DETAILED DESCRIPTION

[029] 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.

[030] In an embodiment according to the disclosure, terminal devices, service network devices and trusted parties each include one or more central processing units (CPU), input/output interfaces, network interfaces, and memories.

[031] The memories may include multiple types of computer-readable mediums, such as, a volatile memory, a random access memory (RAM), and/or a non-volatile memory. For example, a non-volatile memory may be a read only memory (ROM) or a flash random access memory (flash RAM). A memory is one example of a computer-readable medium.

[032] A computer-readable medium includes permanent, non-permanent, removable, and/or non-removable types. Information storage can be achieved by any means or technique. The information can be computer readable instructions, data structures, program sub-units or other type of data. Examples of a computer readable medium include, but not limit to, phase-change RAM (PRAM), static random access memory (SRAM), dynamic random access memory (DRAM), other type of RAMs, read only memory (ROM), electrically erasable programmable read only memory (EEPROM), flash memory, other memory technologies, CD-ROM, DVD, other optical mediums, magnetic tape cassette, magnetic drive storage or other magnetic storages, or any other non-transitory mediums that can be used as a storage and can be accessed by computing devices. According to disclosure, the computer-readable medium do not include a transitory medium, such as a modulated data signal and a carrier wave.

[033] FIG. 1 shows a process for controlling hard drive power consumption. In step S101, it is detected whether there is a fault in a hard drive in a node hard drive group. The node hard drive group includes at least two hard drives.

[034] When no fault is detected in any hard drive in the node hard drive group, in step S102, a hard drive power consumption management service is actuated to make the hard drives convert between the preset power consumption states.

[035] The pre-configured power consumption states include one high power consumption state and one low power consumption state. A high power consumption state may be the power consumption state when the hard drive is under active mode. A low power consumption state may be the power consumption state when the hard drive is under standby mode.

[036] FIG. 2 shows another embodiment of the disclosure describing a process of controlling hard drive power consumption. At step S201, detection takes place for a fault in a hard drive in a node hard drive group. If there is any fault detected in the hard drives of the node hard drive group or a hard drive of the node hard drive group has not been mounted, it may not be possible to control the hard drive in the node hard drive group of converting between the preset states through a hard drive power consumption management service provided by an operating system. Therefore, in order to ensure the reliable operation of the hard drive power consumption management service, it is first detected whether there is any fault in the hard drives of a node hard drive group.

[037] If there is no fault detected in the hard drive of the node hard drive group, step S202 is executed. However, if there is any fault in the hard drive in the node hard drive group or a hard drive of the node hard drive group has not been mounted, an error warning message can be prompted for necessary countermeasures to overcome the detected faults.

[038] In step S202, a command is received to convert a power consumption state in the hard drive in the node hard drive group. The command to convert the power consumption state is sent by the hard drive power consumption management service, which is installed in the operating system.

[039] In an embodiment according to the disclosure, a hard drive in a node hard drive group can be controlled by the hard drive power consumption management service through a hard drive programming interface, which can prevent conflicts with the operation of the hard drive programming when sending a command to the hard drive directly through an Input/output (I/O) port.

This hard drive programming interface is configured to convert a hard drive between preset high power consumption states to a preset low power consumption state or vice versa.

[040] In an embodiment according to the disclosure, the high power consumption state is the power consumption state when the hard drive in the node hard drive group is under an active mode. The low power consumption state is the power consumption state when the hard drive in the node hard drive group is under a standby mode.

[041] When the command is received to convert the hard drive in the node hard drive group from a high power consumption state to a low power consumption state, steps S203 to S207 are executed. Alternately, when the command is received to convert the hard drive in the node hard drive group from a low power consumption state to a high power consumption state, steps S208 to S209 are executed.

[042] In step S203, data of a cache of the operating system is synchronized to the hard drive in the node hard drive group.

[043] In an embodiment according to the disclosure, when the command sent by the hard drive power consumption management service to convert the hard drive in the node hard drive group from a high power consumption state to a low power consumption state is received, in order to avoid data loss, the data in the cache of the operating system is synchronized to the hard drive in the node hard drive group. Thereafter, the hard drive is converted to sleep mode for a short period (e.g., 10 seconds) for the completion of the data synchronization.

[044] In step S204, the hard drive in the node hard drive group is mounted in read-only mode.

[045] In an embodiment according to the disclosure, in order to avoid the operation of both read and write operations to the hard drive in the node hard drive group, the hard drive in the node hard drive group can be mounted in read-only mode.

[046] In step S205, the hard drive in the node hard drive group is converted from a high power consumption state to a low power consumption state, when the hard drive in the node hard drive group is mounted in read-only mode.

[047] In an embodiment according to the disclosure, the hard drive power consumption management service can send a set of commands of converting a hard drive of a node hard drive group from a high power consumption state to a low power consumption state in parallel, which saves time for converting multiple hard drives in the node hard drive group from high power consumption states to low power consumption states.

[048] If there is a failure in the conversion operation, such as converting a hard drive in a node hard drive group from a high power consumption state to a low power consumption state, the hard drive may be placed in a sleep mode for 2 seconds. Thereafter, the conversion operation will be retried again. If the conversion operation fails again, a notification of this failure can be issued, e.g., releasing warning information of the fault that the conversion of hard drive power consumption has failed to convert states.

[049] In order to finish the conversion of a hard drive in a node hard drive group from a high power consumption state to a low power consumption state, embodiments according to the disclosure includes the further steps of setting preset time intervals as below.

[050] In step S206, the hard drive in the node hard drive group is detected periodically according to a preset time interval to determine if it is under a low power consumption state.

[051] By traversing the hard drives in the node hard drive group periodically according to a preset time interval, it can be determined whether the hard drives are under a low power consumption state. The lengths of preset time intervals can be set according to operational needs. For example, preset time intervals can be set as 10 seconds, 50 seconds, or 2 hours, etc.

[052] In step S207, when a hard drive in the node hard drive group is not under a low power consumption state, the hard drive is converted to a low power consumption state.

[053] In an embodiment according to the disclosure, when a hard drive in the node hard drive group is detected to be operating in state that is not a low power consumption state, in order to keep the hard drive in the node hard drive group in low power consumption state, the hard drive power consumption management service can convert the hard drive which is not operating in a low power consumption state to a low power consumption state by a hard drive programming interface.

[054] When the conversion operation of a hard drive from a high power consumption state to a low power consumption state fails, the hard drive power consumption management service can issue notification of the failure by releasing warning information of a fault indicating a failure of the conversion.

[055] When the command is received to convert the hard drive in the node hard drive group from a low power consumption state to a high power consumption state, steps S208 to S209 are executed.

[056] In step S208, a hard drive in the node hard drive group is mounted in read-and-write mode.

[057] In an embodiment according to the disclosure, hard drives in the node hard drive group can be mounted in read-and-write mode in parallel to control the conversion of the hard drives in the node hard drive group when transitioning from a low power consumption states to a high power consumption states, which can decrease the time used for awakening hard drives.

[058] When the conversion of a hard drive from a low power consumption state to a high power consumption state fails, the hard drive power consumption management service can issue warning information of a fault during the conversion process.

[059] In step S209, the hard drive in the node hard drive group is accessed to awake the hard drives in the node hard drive group.

[060] In an embodiment according to the disclosure, access operation may be performed to the hard drive to ensure the success of the conversion of a hard drive from a low power consumption state to a high power consumption state; and the hard drive may be accessed in a number of ways. For example, one way of accessing the hard drives is by creating an empty hidden file under a mount directory of the hard drive in the node hard drive group. After the successful creation of the hidden file, the hidden file is deleted from the mount directory.

[061] FIG. 3A shows an embodiment according to the disclosure which illustrates a process of controlling by a hard drive power consumption management service.

[062] In step S310, the operating environment of a hard drive power consumption management service is detected.

[063] In an embodiment according to the disclosure, by detecting the operating environment of the hard drive power consumption management service, it can be determined in step S330 whether the operating environment meets a preset condition.

[064] In an embodiment according to the disclosure, the preset condition or precondition determined in step S330 may include the following: the existence of log files of the hard drive power consumption management service, the existence of configuration files of the hard drive in the node hard drive group, and the existence of files of analyzing service and files of preprocessing modes. Preconditions may also include that the configuration files of the hard drives in the node hard drive group are found not to be empty, and the files of the analyzing service and the files of preprocessing modes have execution permissions.

[065] In an embodiment according to the disclosure, to implement the above preconditions, the method includes the following:

[066] First, in step S331, it is detected if there are available log files of the hard drive power consumption management service. If there is no such log file, an empty log file is created.

[067] After checking the availability of the log files, then the method proceeds step S332 to detect if the configuration files of the hard drive are available. If there are no configuration files of the hard drive, warning information that no configuration files of the hard drive can be returned. If there is a configuration file of the hard drive, it is determined if a configuration file of the hard drive in the node hard drive group is empty. If the configuration file of the hard drive in the node hard drive group is empty, warning information that the configuration file of the hard drive in the node hard drive group is empty is returned.

[068] Then, in step S333 it is detected if files of analyzing service or files of preprocessing associated with modes are available. If there are no such files, warning information that no files of analyzing service or files for preprocessing is returned. If there are such files, then it is determined whether the files of analyzing service or files of preprocessing associated to modes have execution permissions. If the files of analyzing service and files of preprocessing associated to modes do not have execution permissions, warning information of non-executable files of the analyzing services and non-executable files of preprocessing associated to modes is returned.

[069] Finally, in step S334, it is determined if commands of “awk” and “hdparm” may be executed under current operating environment. If the commands are not executable, warning information that these commands do not exist is returned. The “awk” command applies to text processing, and analyzing outputs of other commands, while the “hdparm” command applies to querying and setting hard drive parameters.

[070] In step S350, when an operating environment of the hard drive power consumption management service meets a set of preconditions, the on/off (or start/stop) operation of the hard drive power consumption management service is controlled according to the parameters sent by the operating system. In some embodiments, the start operation and the stop operation are executed separately.

[071] In an embodiment according to the disclosure, the hard drive power consumption management service is initiated when a first parameter of the operating system, e.g. a “starting parameter” is received. When receiving a second parameter (such as a “stopping parameter”) sent by the operating system, certain operations may occur including: first, it is determined if the hard drive power consumption management service is running. If so, then, the hard drive power consumption management service is turned off.

[072] If the termination of the hard drive power consumption management service fails, warning information may be issued showing the failure of the termination of the hard drive power consumption management service.

[073] It should be noted that in embodiments according to disclosure, hard drive power consumption management service installed in the operation system is used to control the hard drive for user's cold data storage to convert between a high power consumption state and a low power consumption state. There are other ways for the conversion of the hard drive between a high power consumption state and a low power consumption state, e.g., conversion of the hard drive may be controlled by using a function "IOCTL" that manages an I/O passageway of a device in a driver program.

[074] FIG. 3B is a flow chart of a method of controlling a hard drive power consumption management service according to an embodiment of the disclosure. In step S301, an operating environment of a hard drive power consumption management service is detected. Next in step S302, if the detected operating environment of the hard drive power consumption management service meets preconditions, a start/stop operation of the hard drive power consumption management service is controlled in accordance with parameters sent from the operating system.

[075] FIG. 4 shows a block diagram of a device for controlling hard drive power consumption according to an embodiment of the disclosure.

[076] The device 400 includes a hard drive detector 410 configured for detecting a fault in a hard drive of a node hard drive group. The preset modes of the power consumption include at least one high power consumption state and at least one low power consumption state

[077] The device 400 also includes a power consumption controller 420 configured for converting a hard drive in the node hard drive group between preset modes of power consumption through a hard drive power consumption management service if there is no fault detected in any hard drive in the node hard drive group .

[078] FIG. 5 illustrates a block diagram of a power consumption controller 500 according to an embodiment of the disclosure,. The power consumption controller 500 includes a mounting subunit 510 which is configured to mount a hard drive in a node hard drive group in read-only mode when a command of converting a hard drive in a node hard drive group from a high power consumption state to a low power consumption state is received from a hard drive power consumption management service.

[079] The power consumption controller 500 also includes a convertor subunit 520 which is configured to convert the hard drive in the node hard drive group from the high power consumption state to the low power consumption state.

[080] In embodiments according to disclosure, the power consumption controller 500 may also include a power consumption detector subunit 530 which is configured to detect periodically according to a preset time interval if the hard drives in a node of the hard drive group are operating in a low power consumption state.

[081] In embodiments according to disclosure, the power consumption controller 500 may also include a power consumption convertor subunit 540 which is configured to convert, if a hard drive in the node hard drive group is not operating in a low power consumption state, the hard drive to the low power consumption state.

[082] In embodiments according to disclosure, the power consumption controller 500 may also include a synchronizer subunit 550 configured for synchronizing data in a cache of an operating system to a hard drive in a node hard drive group prior to converting the hard drive in the node hard drive group from a high power consumption state to a low power consumption state.

[083] In embodiments according to disclosure, the power consumption controller 500 may also include a second mounting subunit 560 configured to mount a hard drive operation in a node hard drive group in read-and-write mode when a command is received from the management service for converting the hard drive from a low power consumption state to a high power consumption state.

[084] In embodiments according to disclosure, the power consumption controller 500 may also include an awakening subunit 570 configured for awakening a hard drive in a node hard drive group by accessing the hard drive.

[085] In an embodiment according to the disclosure, the awakening subunit 570 can create an empty hidden file under a mount directory of a hard drive of a node of a hard drive group. The awakening subunit 570 deletes the hidden file after the file is created successfully.

[086] FIG. 6 illustrates a block diagram of a device 600 for controlling a hard drive power consumption management service according to an embodiment of the disclosure.

[087] The device 600 includes an environment detector 610 configured for detecting an operating environment of a hard drive power consumption management service.

[088] The device 600 also includes a service controller 620 configured for controlling a start/stop operation of the hard drive power consumption management service in accordance with

parameters sent from an operating system, if the operating environment of hard drive power consumption the management service meets certain preconditions.

[089] The preconditions may include: the existence of log files of the hard drive power consumption management service, existence of configuration files of the hard drive in a node hard drive group, and/or existence of files of analyzing service and files preprocessing associated with modes. Preconditions may also include that the configuration files of the hard drives in the node hard drive group are found not to be empty, and the files of analyzing service and files of preprocessing modes have execution permissions.

[090] The service controller 620 may include a starting subunit 621 which is configured to activate the hard drive power consumption management service a first parameter is received from the operating system.

[091] Additionally, the service controller 620 may include a stopping subunit 630 which is configured for shutting down the hard drive power consumption management service when a second parameter is received from the operating system.

[092] It is appreciated that those skilled in the art can understand that steps or part of steps described in the descriptions may be implemented by using a program/software to command corresponding hardware. The program may be stored in a computer-readable medium including a ROM, a RAM, a hard drive, and/or a CD.

[093] 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 controlling hard drive power consumption, comprising:
  - detecting whether a fault exists in a hard drive in a node hard drive group; and
  - if there is no fault detected in the hard drive in the node hard drive group, controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service.
2. The method of claim 1, wherein:
  - the hard drive power consumption management service is installed in an operating system.
3. The method of claim 2, wherein preset modes of power consumption comprise at least one high power consumption state and at least one low power consumption state.
4. The method of claim 3, wherein, if there is no fault detected in the hard drive in the node hard drive group, the controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service comprises:
  - mounting the hard drive in the node hard drive group in a read only mode when a first command is received from the hard drive power consumption management service to convert the hard drive in the node hard drive group from a high power consumption state to a low power consumption state; and
  - converting the hard drive in the node hard drive group from the high power consumption state to the low power consumption state at same time the hard drive in the node hard drive group is mounted in the read only mode.
5. The method of claim 4, wherein, if there is no fault detected in the hard drive in the node hard drive group, the controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service further comprises:
  - detecting, after a preset time interval, if every hard drive in the node hard drive group is operating in a low power consumption state; and
  - if the hard drive in the node hard drive group is not operating in the low power consumption state, converting the hard drive to the low power consumption state.
6. The method of claim 5, further comprising:

synchronizing data in a cache of the operating system with data of the hard drive in the node hard drive group prior to converting the hard drive from a high power consumption state to a low power consumption state.

7. The method of claim 3, wherein, if there is no fault detected in the hard drive in the node hard drive group, the controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service comprises:

mounting the hard drive in the node hard drive group in a read and write mode when a second command is received from the hard drive power consumption management service in order to convert the hard drive from a low power consumption state to a high power consumption state.

8. The method of claim 7, wherein, if there is no fault detected in the hard drive in the node hard drive group, the controlling the hard drive in the node hard drive group to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service further comprises:

45 awakening the hard drive in the node hard drive group by accessing the hard drive in the node hard drive group.

9. A method for controlling a hard drive power consumption management service, comprising:  
detecting an operating environment of a hard drive power consumption management service; and

50 if the operating environment of the hard drive power consumption management service meets preconditions, controlling a start/stop operation of the hard drive power consumption management service in accordance with parameters sent from an operating system.

10. The method of claim 9, wherein the preconditions comprise at least one of an existing log file of the management service of the hard drive power consumption, a configuration file of a hard drive in a node hard drive group, or files of analyzing service and files of preprocessing modes.

55 11. The method of claim 9, wherein the preconditions comprise configuration files of a hard drive in a node hard drive group, wherein the configuration files are not empty.

12. The method of claim 9, wherein the preconditions comprise files of analyzing service and files of the preprocessing modes which hold execution permissions.

60 13. The method of claim 9, wherein the controlling a start/stop operation of the hard drive power consumption management service in accordance with parameters sent from an operating system comprises:

activating the hard drive power consumption management service when a first parameter is received from the operating system and

65 terminating the hard drive power consumption management service when a second parameter is received from the operating system.

14. An apparatus for controlling hard drive power consumption, comprising:

a hard drive detector configured for detecting whether there is a fault in a hard drive in a node hard drive group; and

70 a power consumption controller configured for controlling, if no fault detected in any of the hard drives in the node hard drive group, the hard drive to convert the hard drive between preset modes of power consumption by a hard drive power consumption management service associated with the hard drive power consumption.

15. The apparatus of claim 14, wherein the preset modes of power consumption include at least one high power consumption state and at least one low power consumption state.

75 16. The apparatus of claim 15, wherein:

the hard drive power consumption management service is installed in an operating system.

17. The apparatus of claim 16, wherein the power consumption controller comprises:

80 a mounting subunit configured for mounting the hard drive in the node hard drive group in a read only mode when a command is received from the hard drive power consumption management service to convert the hard drive in the node hard drive group from a high power consumption state to a low power consumption state; and

85 a convertor subunit configured for converting the hard drive in the node hard drive group from the high power consumption state to the low power consumption state after the read only mode is selected for the hard drive in the node hard drive group.

18. The apparatus of claim 17, wherein the power consumption controller further comprises:

a power consumption detector subunit configured for detecting, after a preset time interval, if every hard drive in the node hard drive group is operating in a low power consumption state; and

90 a power consumption convertor subunit configured for converting, if the hard drive in the node hard drive group is not operating in the low power consumption state, the hard drive to the low power consumption state.

19. The apparatus of claim 18, wherein the power consumption controller further comprises:

95 a synchronizer subunit configured for synchronizing data in a cache of the operating system with a data of the hard drive in the node hard drive group before converting the hard drive from a high power consumption state to a low power consumption state.

20. The apparatus of claim 16, wherein the power consumption controller further comprises:

100 a second mounting subunit configured for mounting the hard drive in the node hard drive group in a read and write mode when a command is received from the hard drive power consumption management service to convert the hard drive in the node hard drive group from a low power consumption state to a high power consumption state.

21. The apparatus of claim 20, wherein the power consumption controller further comprises:

105 an awakening subunit configured for awakening the hard drive in the node hard drive group by accessing the hard drive in the node hard drive group.

22. An apparatus for controlling a hard drive power consumption management service, comprising:

110 an environment detector configured for detecting an operating environment of the hard drive power consumption management service; and

115 a service controller configured for controlling a start/stop operation of the hard drive power consumption management service in accordance with parameters sent from an operating system, if the operating environment of the hard drive power consumption management service meets preconditions.

23. The apparatus of claim 22, wherein the preconditions comprise at least one of an existing log file of the hard drive power consumption management service, a configuration file of a hard drive in a node hard drive group, or a file of analyzing service and a file of preprocessing modes.

24. The apparatus of claim 22, wherein the power consumption controller further comprises:

120 a starting subunit configured for activating the hard drive power consumption management service when a first parameter is received from the operating system; and

125 a stopping subunit configured for terminating the hard drive power consumption management service when a second parameter is received from the operating system.

25. A non-transitory computer readable storage medium having embedded therein program instructions, when executed by one or more processors of a device, controls hard drive power consumption comprising:

detecting whether a fault exists in a hard drive in a node hard drive group; and

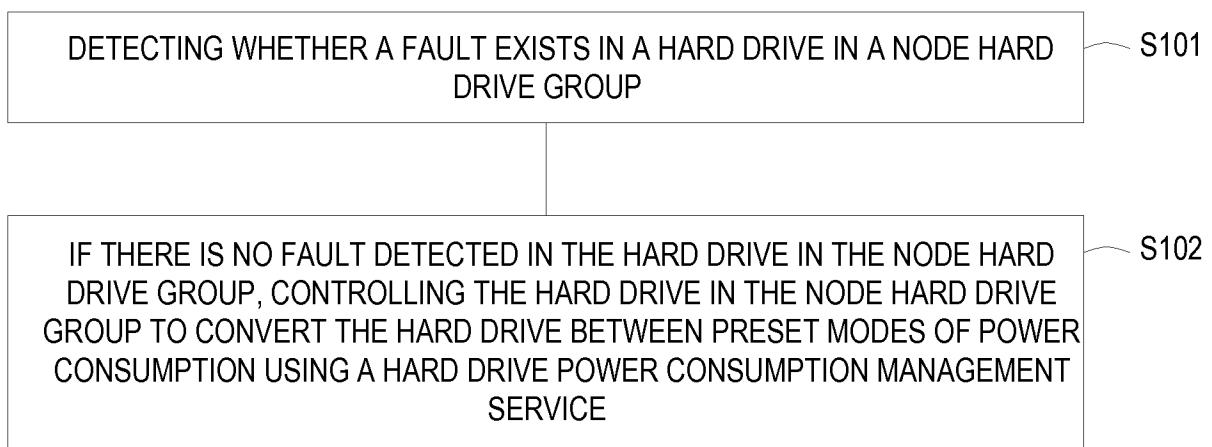
125 if no fault detected in the hard drive in the node hard drive group, controlling the hard drive to convert the hard drive between preset modes of power consumption using a hard drive power consumption management service.

26. A non-transitory computer readable storage medium having embedded therein program instructions, when executed by one or more processors of a device, controls a hard drive power 130 consumption management service, comprising:

detecting an operating environment of the hard drive power consumption management service; and

if the operating environment of the hard drive power consumption management service meets preconditions, controlling a start/stop operation of the hard drive power consumption 135 management service in accordance with parameters sent from an operating system.

1/5

**FIG. 1**

2/5

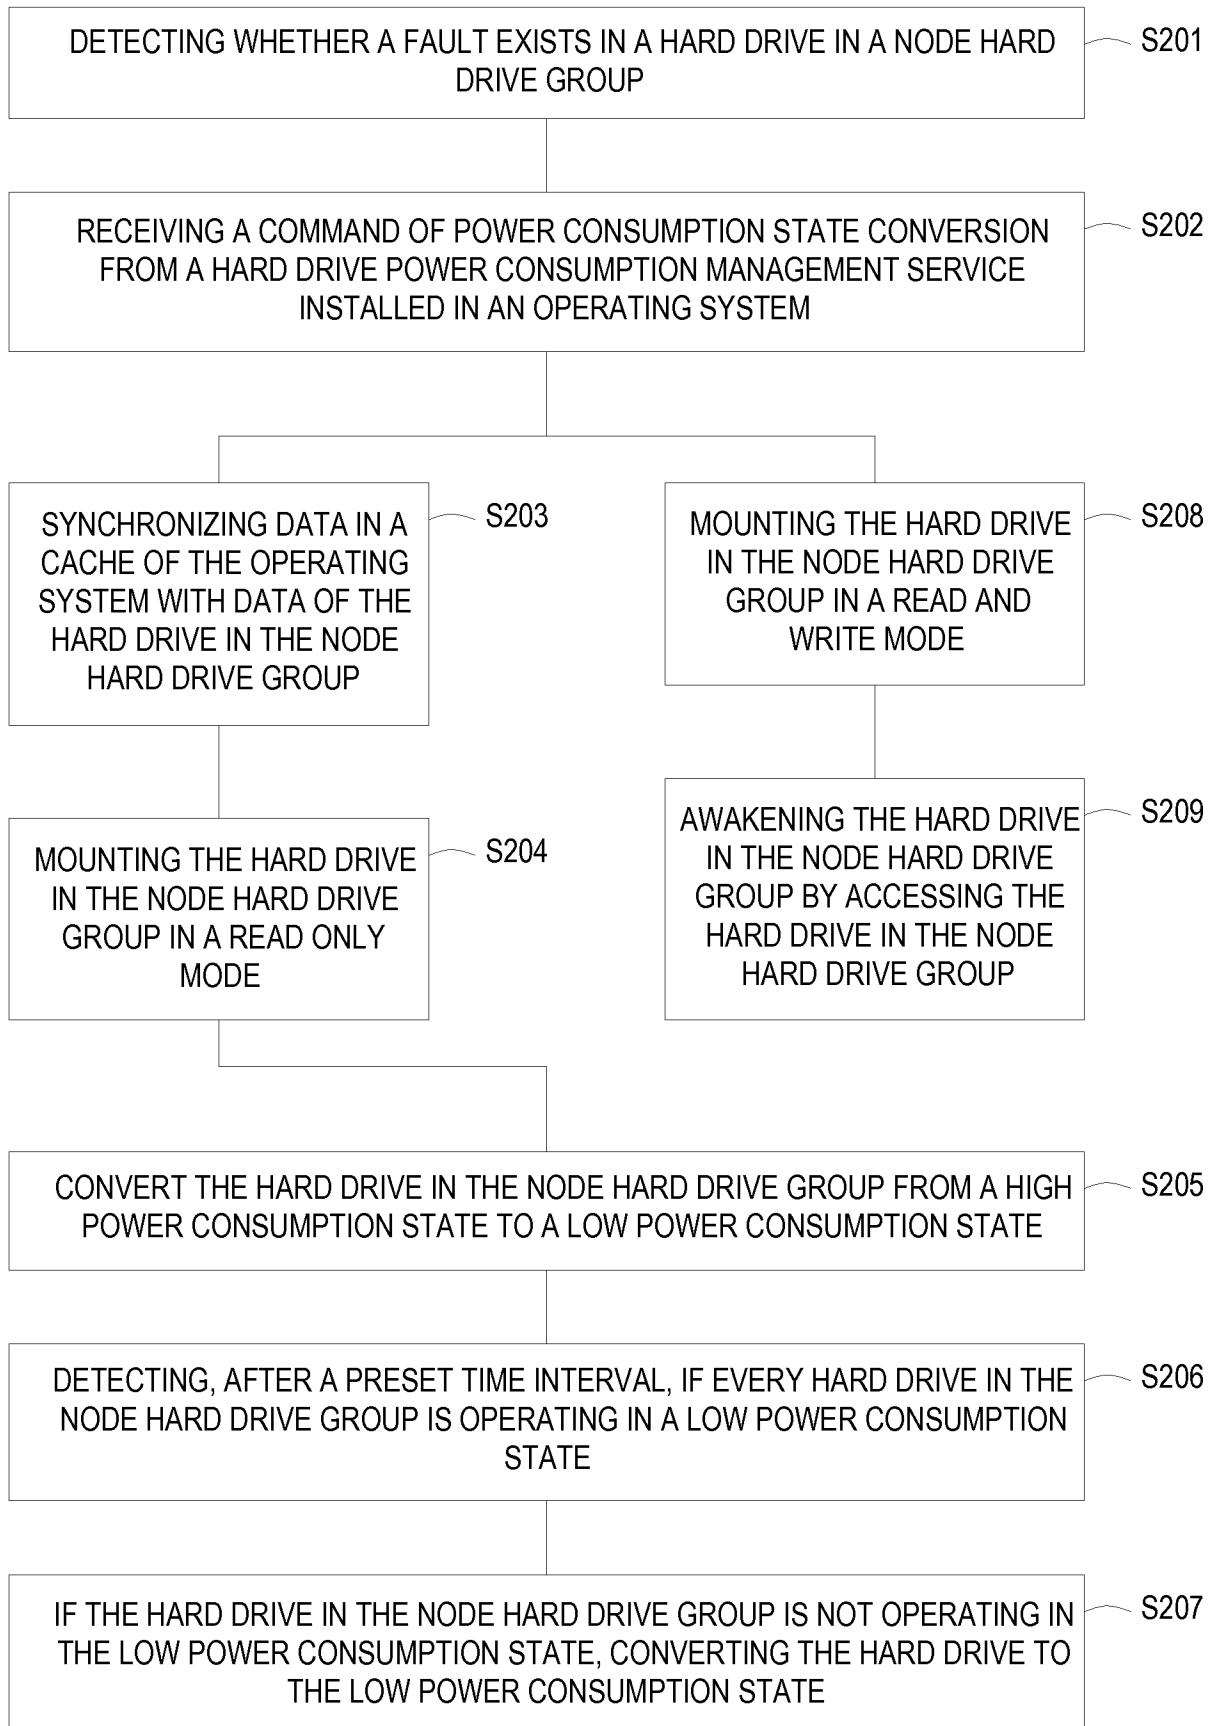


FIG. 2

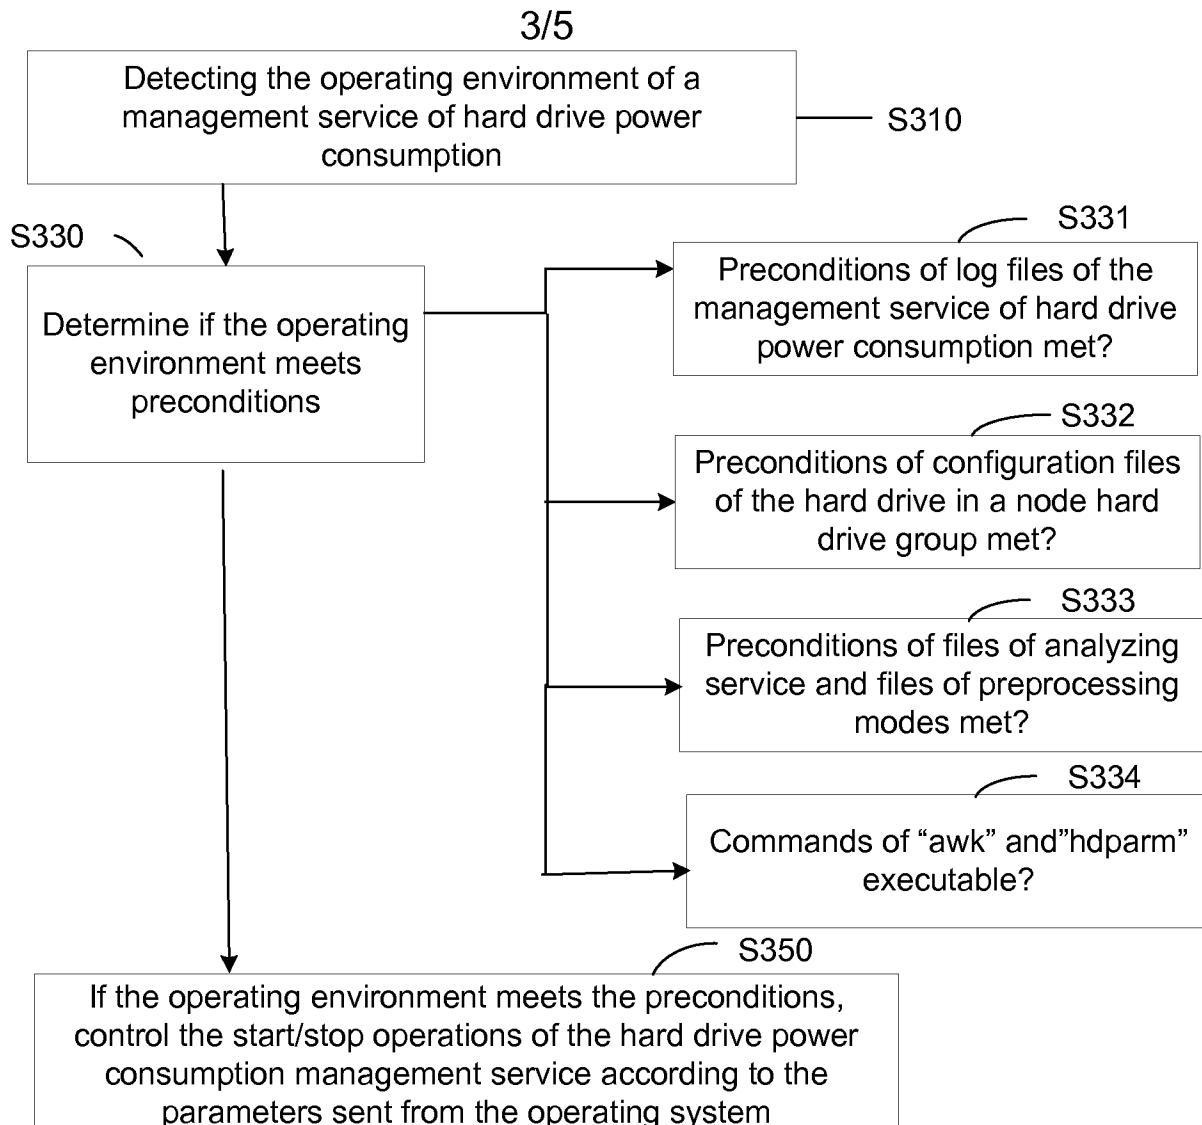


FIG. 3A

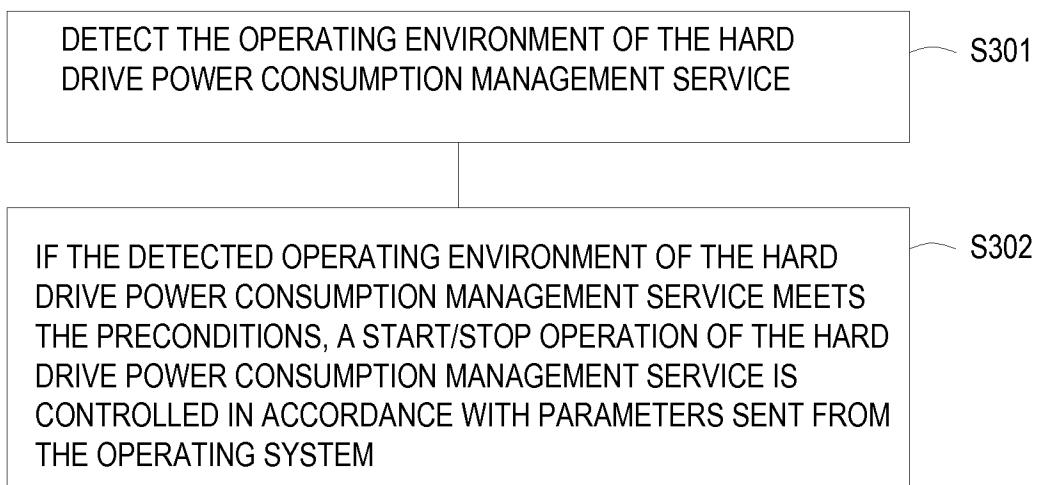
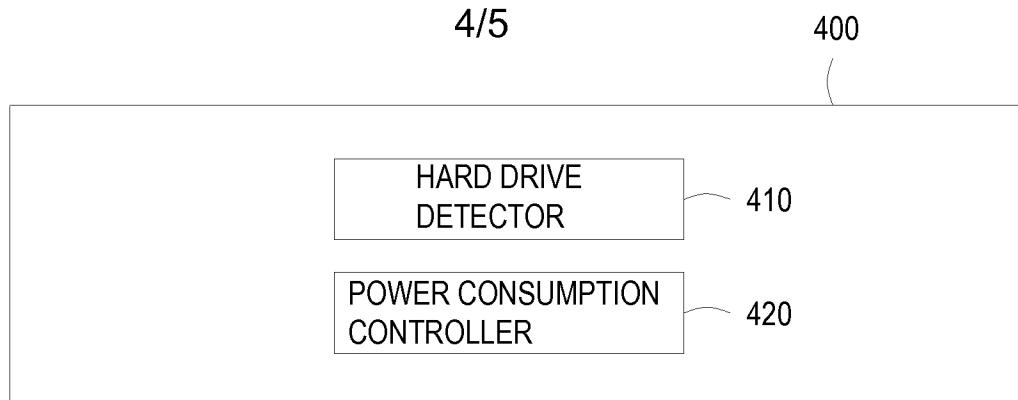
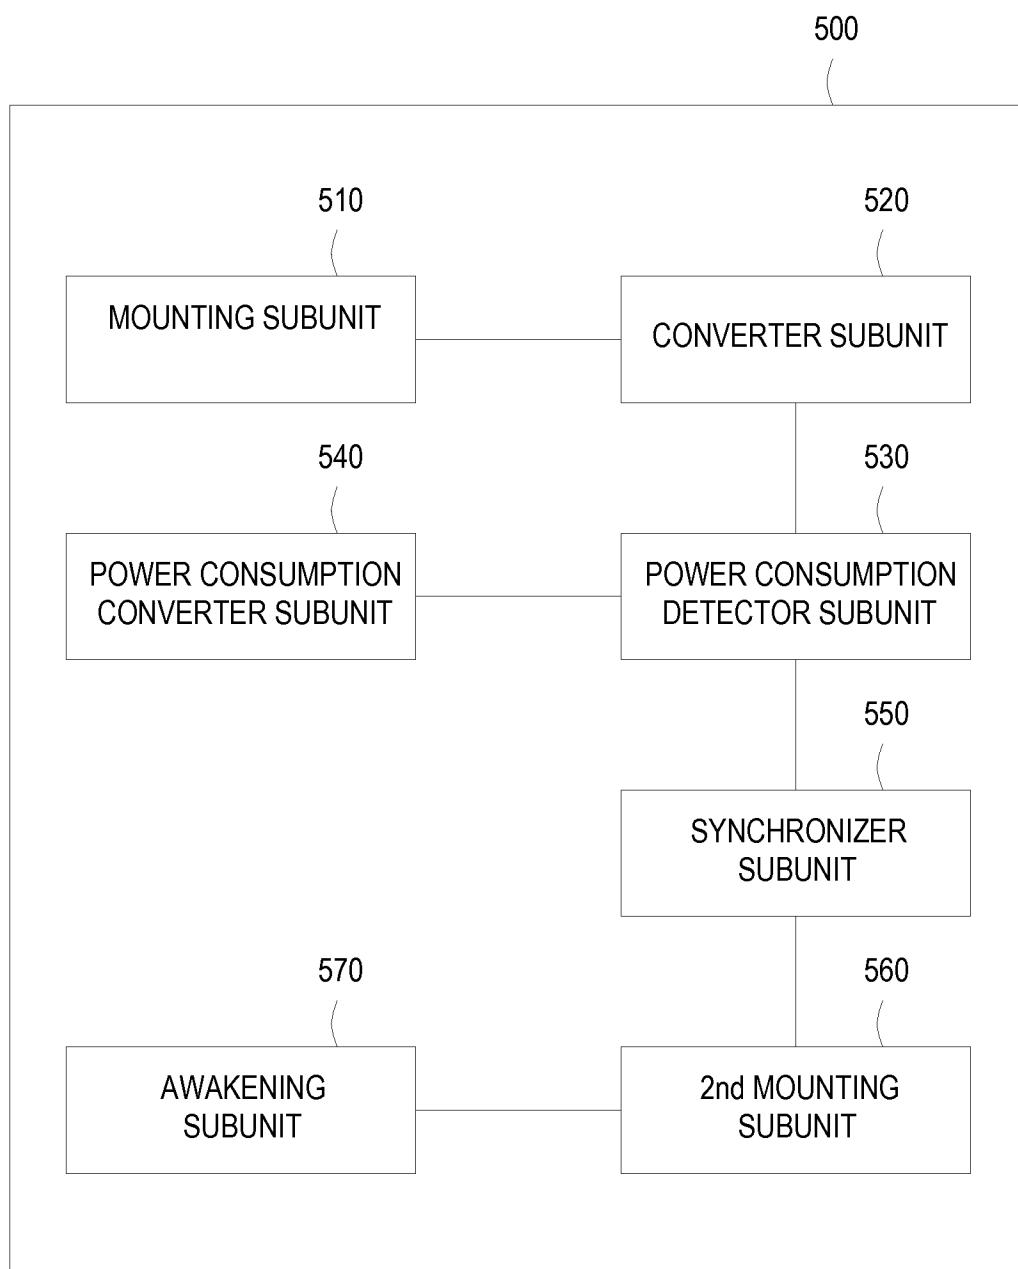


FIG. 3B



**FIG. 4**



**FIG. 5**

5/5

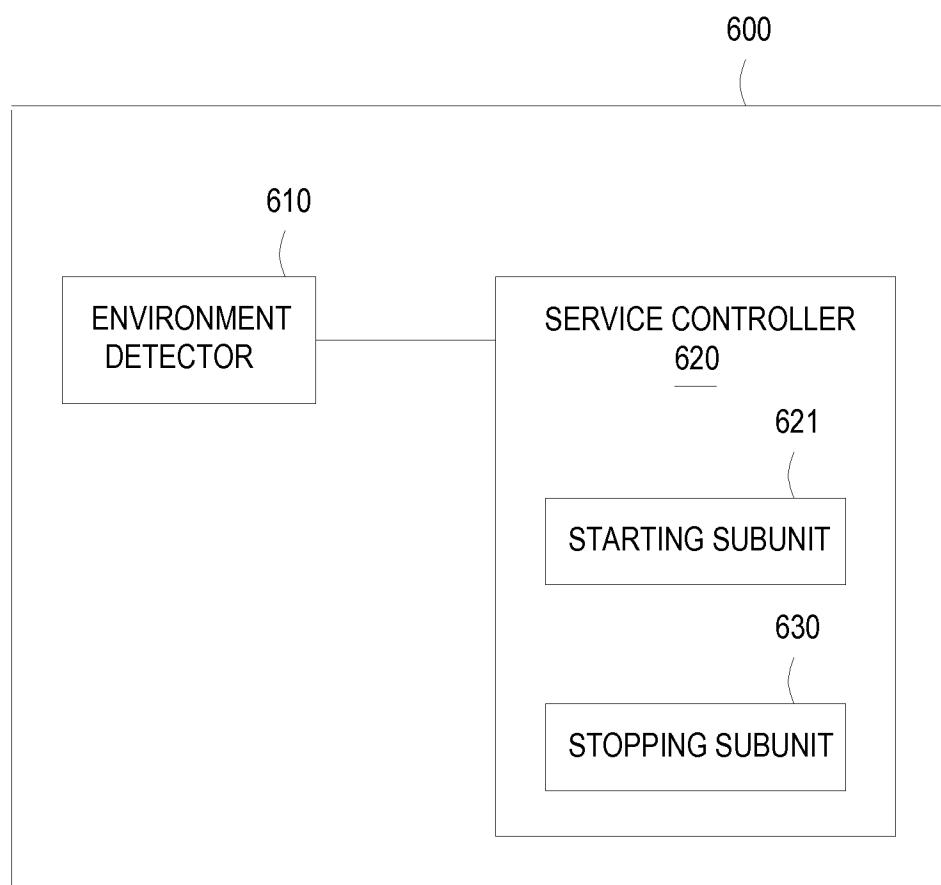


FIG. 6

**INTERNATIONAL SEARCH REPORT**

International application No.  
PCT/US15/60213

<p><b>A. CLASSIFICATION OF SUBJECT MATTER</b>            IPC(8) - G06F 1/32; G11B 19/04 (2015.01)            CPC - G06F 1/3221, 1/3268; G11B 19/041            According to International Patent Classification (IPC) or to both national classification and IPC</p>																						
<p><b>B. FIELDS SEARCHED</b></p> <p>Minimum documentation searched (classification system followed by classification symbols)            IPC(8): G06F 1/32, 3/06; G11B 5/55, 19/00, 19/04, 19/06 (2015.01)            CPC: G06F 1/3221, 1/3268, 3/0601; G11B 5/55, 19/00, 19/041, 19/06; Y02B 60/00, 60/1246, 60/32</p>																						
<p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched</p>																						
<p>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); EBSCO; IEEE; Google/Google Scholar; hard, drive*, power*, consum*, detect, fault*, node*, preset*, mode*</p>																						
<p><b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%;">Category*</th> <th style="width: 80%;">Citation of document, with indication, where appropriate, of the relevant passages</th> <th style="width: 10%;">Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>X</td> <td>CN 103268147 A (BEIJING BAIDU NETWORK TECHNOLOGY CO., LTD.) 28 August 2013, (see machine translation).</td> <td>1-3, 9, 13-16, 22, 24-26</td> </tr> <tr> <td>Y</td> <td></td> <td>10-11, 23</td> </tr> <tr> <td>Y</td> <td>US 2012/0265959 A1 (LE, B et al.) 18 October 2012, Paragraphs [0231]-[0233], [0351], [0382].</td> <td>10-11, 23</td> </tr> <tr> <td>A</td> <td>US 5682273 A (HETZLER, S) 28 October 1997, entire document.</td> <td>1-26</td> </tr> <tr> <td>A</td> <td>US 5345347 A (HOPKINS, D et al.) 06 September 1994, entire document.</td> <td>1-26</td> </tr> <tr> <td>A</td> <td>US 6624962 B1 (KODAMA, T et al.) 23 September 2003, entire document.</td> <td>1-26</td> </tr> </tbody> </table>		Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	X	CN 103268147 A (BEIJING BAIDU NETWORK TECHNOLOGY CO., LTD.) 28 August 2013, (see machine translation).	1-3, 9, 13-16, 22, 24-26	Y		10-11, 23	Y	US 2012/0265959 A1 (LE, B et al.) 18 October 2012, Paragraphs [0231]-[0233], [0351], [0382].	10-11, 23	A	US 5682273 A (HETZLER, S) 28 October 1997, entire document.	1-26	A	US 5345347 A (HOPKINS, D et al.) 06 September 1994, entire document.	1-26	A	US 6624962 B1 (KODAMA, T et al.) 23 September 2003, entire document.	1-26
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																				
X	CN 103268147 A (BEIJING BAIDU NETWORK TECHNOLOGY CO., LTD.) 28 August 2013, (see machine translation).	1-3, 9, 13-16, 22, 24-26																				
Y		10-11, 23																				
Y	US 2012/0265959 A1 (LE, B et al.) 18 October 2012, Paragraphs [0231]-[0233], [0351], [0382].	10-11, 23																				
A	US 5682273 A (HETZLER, S) 28 October 1997, entire document.	1-26																				
A	US 5345347 A (HOPKINS, D et al.) 06 September 1994, entire document.	1-26																				
A	US 6624962 B1 (KODAMA, T et al.) 23 September 2003, entire document.	1-26																				
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.																						
<p>* Special categories of cited documents:</p> <p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p> <p>“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</p> <p>“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</p> <p>“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</p> <p>“&amp;” document member of the same patent family</p>																						
<p>Date of the actual completion of the international search</p> <p>23 December 2015 (23.12.2015)</p>																						
<p>Date of mailing of the international search report</p> <p>27 JAN 2016</p>																						
<p>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</p>																						
<p>Authorized officer</p> <p align="center">Shane Thomas</p> <p>PCT Helpdesk: 571-272-4300            PCT OSP: 571-272-7774</p>																						