

Title of the Invention:

Method for controlling cache memories, computer system, hard disk drive unit, and hard disk control unit

Detailed Description of the Invention:

[Field of the invention]

The present invention relates to a method for controlling a plurality of cache memories, as well as to a computer system, a hard disk drive unit, and a hard disk control unit used to enable those cache memories to hold data efficiently.

[Background art]

A computer system is usually composed of a CPU, a main memory, an HDD (Hard Disk Drive) unit, etc. and those devices of the computer system have different data transfer rates. Generally, a DRAM is employed as such a main memory. The data transfer rate of the DRAM is far slower than that of the CPU, so the DRAM often becomes a factor for lowering the operation speed of the CPU. The data transfer rates of such external storage devices as an HDD unit, a CD-ROM unit, etc. are far slower than that of the DRAM, so the operation speed of the DRAM is often lowered by the operation of such a device. Such an operation speed difference among devices that exchange data with each other, therefore, hinders the intended performances of those devices even when they have an ability for transferring data fast respectively, since the transfer rate of the whole system is dominated by a device having the slowest transfer rate among them.

In order to reduce such a speed difference among devices, therefore, a computer system is provided with a cache memory. The cache memory makes the most use of the localization of the object program and/or data, thereby storing part of the data stored in the low-ranking device having a slow data transfer rate. Consequently, the number of accesses to the low-ranking device is reduced, thereby improving the data transfer rate of the whole computer system.

A typical cache memory is provided between a CPU and a main

memory. If there are a plurality of cache memories, they are called the primary cache, the secondary cache, ... sequentially from the CPU side. The closer the cache is to the CPU, the smaller the storage capacity becomes. Generally, an SRAM is employed for each of those cache memories.

The main memory is used as a cache memory for an external storage. Consequently, the number of accesses to the external storage is reduced, thereby the operation speed of the main memory can be improved. Sometimes, a cache memory is also provided in such an external storage as a so-called extended card, etc. For example, each HDD unit is provided with a cache memory for holding part of data stored in a magnetic disk. With the use of such a cache memory, the number of accesses to the magnetic disk is reduced, thereby the data transfer rate of the HDD unit can be improved. A cache memory provided in the main memory, the extended card, the HDD unit, etc. in this way is referred to as a disk cache.

Generally, a data reading method referred to as a "Look-Ahead Method" is employed for those disk caches. According to this "look-ahead" data reading method, after it is ended to read requested data from an area requested by a host system, the data in the area following that of the requested data is also read together. This method can thus improve the cache hitting rate in reading data sequentially from consecutive addresses in ascending order of their numbers.

Figure 12 shows how to read data from a conventional disk cache.

In Figure 12, a host system means a unit that requests reading/writing data from/in an HDD unit. Each of the host system and the HDD unit is provided with a cache memory and data is exchanged between the cache memories of those units. In this case, the cache memory of the host system is regarded as a high-ranking cache memory and the cache memory of the HDD unit is regarded as a low-ranking cache memory. It is premised

here that data D1 requested from an application program is held in none of the host system and the HDD unit.

If an application program requests the host system for transfer of data D1, a cache mishit occurs in the host system. The HDD unit is thus requested for transfer of data D3 including both data D1 and data D2 in the address following the address of this data D1. In the HDD unit, therefore, the data D3 requested from the host system is read from the object magnetic disk and held in the cache memory to cope with this cache mishit. The data D3 is transferred to the host system. Consequently, the cache memory of the host system holds the data D3. And furthermore, in the HDD unit, the data D4 in the address following that of the transferred data D3 is read from the magnetic disk and held in the cache memory.

Such a look-ahead data reading method is disclosed in, for example, the official gazette of Published Unexamined Patent Application No. 11-110139. The gazette describes both method and apparatus for reading data in the look-ahead manner. According to the method and the apparatus described in this gazette, if a read operation is detected in the direction of consecutive addresses in descending order of their numbers, the requested data is read from an area requested from the host system, then data is also read from an address preceding the address of the requested data, thereby the cache hitting rate can be improved during reading in the reverse direction of consecutive addresses in descending order of their numbers.

[Problems to be solved by the invention]

However, according to the conventional method for holding data in a cache memory, the same data is held in both host system and HDD unit. If reading of data is requested from an application program to the host system and the host system hits the cache memory, therefore, the requested data is transferred from the host system to the high-ranking cache memory. The HDD unit is never requested to read the data. Consequently, if the

same data is held in the cache memory of the HDD unit at this time, the data will not be used effectively.

Under such circumstances, it is an object of the present invention to provide a method for controlling a plurality of cache memories, which can solve the above conventional problems and avoid a wasteful operation that common data is held in both high-ranking and low-ranking cache memories or reduce such common data to be held in both high-ranking and low-ranking cache memories, thereby holding data more efficiently in each of those cache memories, as well as to provide a computer system, a hard disk drive unit, and a hard disk control unit that are all employing the method for controlling a plurality of cache memories.

[Means for solving the problems]

A method of the present invention for controlling a plurality of cache memories controls a plurality of cache memories including a low-ranking cache memory and a high-ranking cache memory so that the high-ranking and low-ranking cache memories are operated in different swap modes.

The computer system of the present invention is provided with a low-ranking cache memory and a high-ranking cache memory connected to the low-ranking cache memory. The high-ranking and low-ranking cache memories exchange select information for selecting a swap mode respectively with each other when said computer system is started up, thereby selecting different swap modes according to the exchanged select information respectively.

The hard disk drive unit of the present invention is provided with a low-ranking cache memory for storing part of the data stored in a magnetic disk and it is connected to a host system having functions of a high-ranking cache memory.

Each of the high-ranking and low-ranking cache memories

exchanges select information with the host system so that the high-ranking and low-ranking cache memories are operated in different swap modes when the system is started up. According to the exchanged select information, the hard disk drive unit selects the swap mode of the low-ranking cache memory.

A hard disk control unit of the present invention is provided with the first connection terminal connected to an extended connection terminal of a peripheral device provided for the computer system, the second connection terminal connected to the hard disk drive unit having a low-ranking cache memory for storing part of the data stored in a magnetic disk, and a high-ranking cache memory for holding part of the data stored in the hard disk drive unit connected to the second connection terminal. This hard disk control unit controls the hard disk drive unit connected to the second connection terminal and enables the high-ranking and low-ranking cache memories to exchange select information with the hard disk drive unit connected to the second connection terminal when the computer system connected to the first connection terminal is started up so that the high-ranking and low-ranking cache memories are operated in different swap modes. The hard disk control unit then selects the swap mode of the high-ranking cache memory according to the exchanged select information.

Another hard disk control unit of the present invention is provided with the first connection terminal connected to an extended connection terminal of a peripheral device provided for a computer system provided with a host system having functions of a high-ranking cache memory, the second connection terminal connected to the hard disk drive unit, and a low-ranking cache memory for holding part of the data stored in the hard disk drive unit connected to the second connection terminal. This hard disk control unit controls the hard disk drive unit connected to the second connection terminal and enables both high-ranking and low-ranking cache memories to exchange select information with the host system so that they can be operated in different swap modes when the computer

system connected to the first connection terminal is started up. The hard disk control unit then selects the swap mode of the low-ranking cache memory according to the exchanged select information.

Another method of the present invention for controlling cache memories comprises the step of erasing data from a low-ranking cache memory after the data is transferred to a host system in response to a request of data transfer issued from the host system to the low-ranking cache memory.

Further another method of the present invention for controlling cache memories comprises the step of holding part of data held in a low-ranking cache memory in a high-ranking cache memory; wherein data transferred from the low-ranking cache memory to the high-ranking cache memory in response to a request of data transfer issued from the high-ranking cache memory to the low-ranking cache memory is erased from the low-ranking cache memory.

Another computer system of the present invention includes: a low-ranking cache memory; and a high-ranking cache memory for holding part of data stored in the low-ranking cache memory; wherein the low-ranking cache memory, when a request of data transfer is issued from the high-ranking cache memory to the low-ranking cache memory, erases the data transferred from the low-ranking cache memory to the high-ranking cache memory in response to the request.

Another hard disk drive unit of the present invention is provided with a cache memory for storing part of the data stored in a magnetic disk. The cache memory is characterized in that the cache memory erases the data transferred to the host system in response to a request of data transfer issued from the host system connected to this hard disk drive unit.

Another hard disk control unit of the present invention is provided with the first connection terminal connected to an

extended connection terminal of a peripheral device provided for the computer system, the second connection terminal connected to the hard disk drive unit, and a cache memory for holding part of the data stored in the hard disk drive unit connected to the second connection terminal. This control unit controls the hard disk drive unit connected to the second connection terminal and the cache memory, when receiving a request of data transfer from the host system through the first connection terminal, erases the data transferred to the host system in response to the request.

[Preferred embodiments]

Hereunder, the embodiments of the present invention will be described with reference to the accompanying drawings with respect to a method for controlling cache memories, as well as a computer system.

First Embodiment

A method for controlling cache memories in a first embodiment of the present invention enables both high-ranking and low-ranking cache memories to be operated in different swap modes. The "swap" means replacing old data held in a cache memory with new data when the cache memory has no empty space.

Main swap modes to be employed for the method of the present invention are the LRU (Least Recently Used) mode, the LFU (Least Frequently Used) mode, and the FIFO (First-In First-Out) mode. The LRU mode is a mode for swapping out data that is not accessed for the longest time. The LFU mode is a mode for swapping out data that is accessed least frequently. The FIFO mode is a mode for swapping out data that is held first sequentially. The swap modes are not limited only to those; for example, various other swap modes including a method for swapping out a segment having the least amount of data may also be employed.

If both high-ranking and low-ranking cache memories are operated in the same swap mode, both cache memories come to hold data of the same characteristic. For example, if both high-ranking and low-ranking cache memories are operated in the LRU mode, the data that is not accessed for the longest time is swapped out. Consequently, the data that is accessed most recently is held in those cache memories. And, if both high-ranking and low-ranking cache memories are operated in the LFU mode, the data that is accessed least frequently is swapped out. Consequently, the data that is accessed frequently is held in those cache memories. Furthermore, if both high-ranking and low-ranking cache memories are operated in the FIFO mode, then the data that is held first is swapped out sequentially from those cache memories. Consequently, the data that is accessed most recently is held in those cache memories.

On the contrary, if, for example, the high-ranking cache memory is operated in the LFU mode and the low-ranking cache memory is operated in the FIFO mode, then the data that is accessed frequently is held in the high-ranking cache memory and the data that is accessed most recently is held in the low-ranking cache memory. In this way, data items of different characteristics are held in those cache memories, and thus it is possible to reduce the possibility that the same data is held in both cache memories. In other words, it is possible to reduce the possibility that the data held in the high-ranking cache memory is also held in the low-ranking cache memory, so data can be held in each cache memory efficiently.

Preferably, those high-ranking and low-ranking cache memories should be operated in different swap modes of different characteristics. This is because the possibility that different data items are held in those high-ranking and low-ranking cache memories becomes higher than when those cache memories are operated in swap modes of similar characteristics, thereby each cache memory can be used more efficiently. However, it should be avoided to use any swap mode that might lower the cache hitting rate. This is because such a cache

memory cannot give full play to its function for making the most use of data localization.

Next, a description will be made of a computer system for realizing this cache memory controlling method.

Figure 1 shows a schematic block diagram of a computer system in the first embodiment of the present invention.

As shown in Figure 1, this computer system comprises a CPU 11, a secondary cache 13, a first bridge circuit 14, and a main memory 15. The CPU 11 is provided with a primary cache 12. The CPU 11 is connected to the secondary cache 13 through an external bus 10a. The CPU 11 is connected to the first bridge circuit 14 through an external bus 10b. The first bridge circuit 14 is connected to the main memory 15 through an external bus 10c. Each of the primary cache 12 and the secondary cache 13 is composed of, for example, an SRAM. The primary cache 12 is provided with a storage capacity of several tens of kilobytes to several hundreds of kilobytes. The secondary cache 13 is provided with a storage capacity of, for example, several hundreds of kilobytes to several megabytes. The main memory 15 is composed of, for example, a DRAM. The main memory 15 has a storage capacity of, for example, several tens of megabytes to several hundreds of megabytes.

The first bridge circuit 14 is connected to an expansion bus 20 used for an expanded peripheral device. The expansion bus 20 is composed of, for example, a PCI (Peripheral Component Interconnect) bus. The first bridge circuit 14 has functions of, for example, a clock generator, a CPU interface, a secondary cache controller, a main memory controller, a PCI bus interface, etc. that are not illustrated.

An HDC (Hard Disk Controller) card 21 is connected to the expansion bus 20. In addition, an HDD (Hard Disk Drive) unit 22 and a CD-ROM unit 23 are connected to the HDC card 21 through a bus 24a and a bus 24b respectively. Each of the

buses 24a and 24b is, for example, an ATA (AT Attachment)/ATAPI (ATA Packet Interface) bus. The HDC card 21 is used to add another HDD unit and/or expand the controlling functions of the HDD unit.

The HDC card 21 is not limited only to that used for expanding the controlling functions of the HDD unit 22; it may be any card if it is provided with a cache memory. For example, it may be an SCSI (Small Computer System Interface) card. The expansion bus 20 is not limited only to the PCI bus and the buses 24a and 24b are not limited only to the ATA/ATAPI bus.

Figure 2 shows a schematic block diagram of the HDD unit shown in Figure 1.

As shown in Figure 2, the HDD unit 22 is provided with a magnetic disk 41 having a data recording surface and a head slider 43 (not shown) having a magnetic head for both recording and regenerating.

The magnetic disk 41 is attached fixedly to the rotary shaft of a spindle motor 42 and rotated by the spindle motor 42. The head slider 43 is fixed to a tip portion of the head arm 44. The head arm 44 is attached to an actuator 45 and supports the head slider 43 elastically. The head arm 44 is rotated by the actuator 45 above the recording surface of the magnetic disk 41. Consequently, the head slider 43 moves approximately in the radial direction of the magnetic disk 41, thereby data is accessed at a given position on the recording surface of the magnetic disk 41.

Herein, one or more magnetic disks may be used as the magnetic disk 41. A recording surface is formed at one side or both sides of the magnetic disk 41. The head slider 43 and the head arm 44 are provided so as to be corresponded to each recording surface of the magnetic disk 41.

A driving mechanism, which is composed of the magnetic disk 41,

the spindle motor 42, the head slider 43, the head arm 44, and the actuator 45, is controlled by a control circuit 50. The control circuit 50 is provided with an HDC (Hard Disk Controller) 51, a control memory 53, a disk cache 54, and a host I/F 55. Those components are connected to each another through a bus 56.

The HDC 51 controls the whole HDD unit 22 according to a control program and/or control data stored in the control memory 53. The HDC 51 executes an arithmetic operation for controlling the error to occur in servo controlling and/or in recording/regenerating, thereby driving the spindle motor 42 and the actuator 45 so as to record data on the magnetic disk 41 through the magnetic head for recording, as well as regenerate and read data from the magnetic disk 41 through the magnetic head for regeneration. The HDC 51 is provided with a cache controller 52 for enabling the disk cache 54 to hold part of the data recorded on the magnetic disk 41.

The control memory 53 stores a control program executed by the HDC 51 and control data used for the control program. The disk cache 54 stores write data to be recorded on the magnetic disk 41 temporarily. The disk cache 54 also has functions of a buffer memory for storing data read from the magnetic disk 41 temporarily. The HDC card 21 is connected to the host I/F 55 through a bus 21a. The host I/F 55 is an interface circuit for exchanging data with the HDC card 21.

The disk cache 54 has a cache memory function for holding part of the data recorded on the magnetic disk 41 as described above. The disk cache 54 is equivalent to a low-ranking cache memory. The disk cache 54 is composed of, for example, a DRAM and has a storage capacity of several megabytes to several tens of megabytes. The cache controller 52 controls this disk cache. The cache controller 52 supports swap modes, for example, the FIFO mode and the LRU mode.

The HDD unit 22 exchanges setup information with the HDC card

(host system) 21 when this computer system is started up, and the HDD unit 22 is initialized according to the exchanged setup information. At this time, the HDD unit 22 transfers the swap mode (FIFO or LRU mode) of the cache memory that it can select to the HDC card 21 and selects the swap mode transferred from the HDC card 21 in response to the transfer as the swap mode of its cache memory.

Figure 3 shows a schematic block diagram of the HDC card 21 shown in Figure 1.

As shown in Figure 3, the HDC card 21 is provided with an HDC 31, a control memory 33, a disk cache 34, an I/O port 35, an I/O port 36, and a host I/F 37. The HDC 31, the control memory 33, the disk cache 34, and the host I/F 37 are connected to each another through a bus 38c. The I/O port 35 is connected to the HDC 31 through a bus 38a and the I/O port 36 is connected to the HDC 31 through a bus 38b respectively.

The HDC 31 controls the whole HDC card 21 according to the control program and control data stored in the control memory 33. The HDC 31 is provided with a cache controller 32 for enabling the disk cache 34 to hold part of the data stored in the HDD unit 22.

The control memory 33 stores a control program executed by the HDC 31 and control data used for the control program. The disk cache 34 stores write data to be stored in the HDD unit 22 temporarily. The disk cache 34 is also provided with functions of a buffer memory for storing data read from the HDD unit 22 temporarily. The I/O port 35 is connected to the HDD unit 22 through a bus 24a and the I/O port 36 is connected to the CD-ROM unit 23 through a bus 24b. The host I/F 37 is connected to an expansion bus 20. The host I/F 37 is an interface circuit for sending/receiving data through the expansion bus 20.

The disk cache 34 is provided with functions of a cache memory

for holding part of the data stored in the HDD unit 22 as described above. The disk cache 34 is thus equivalent to a high-ranking cache memory with respect to the disk cache 54 provided in the HDD unit 22. The disk cache 34 is composed of, for example, a DRAM and has a storage capacity of several megabytes to several tens of megabytes. The cache controller 32 controls the disk cache 34. The cache controller 32 supports, for example, FIFO, LRU, and LFU modes.

The HDC card 21 sends/receives setup information to/from the HDD unit (low-ranking unit) 22 when this computer system is started up, and the HDC card 21 is initialized according to the sent/received setup information. At this time, for example, the HDC card 21 selects a swap mode that the HDD unit 22 should select, from among the swap modes transferred from the HDD unit 22, then transfers the selected swap mode to the HDD unit 22. And furthermore, it selects a swap mode different from that transferred to the HDD unit 22 as the swap mode of its cache memory. The priority for selecting the object swap mode is set beforehand.

Figure 4 shows an example of an operation sequence of a cache memory for selecting a swap mode.

The computer system is started up, then the HDD unit 22 transfers selectable swap modes (FIFO and LRU modes) of the cache memory to the HDC card 21 (step S1). Receiving the swap modes from the HDD unit 22 (step S2), the HDC card 21 selects a swap mode that should be selected by the HDD unit 22 from the received swap modes according to the preset priority (step S3) and responses the selected swap mode to the HDD unit 22 (step S4).

Receiving the swap mode from the HDC card 21 (step S5), the HDD unit 22 selects it as the swap mode of its cache memory (step S6). On the other hand, the HDC card 21 selects a swap mode (LRU mode) different from the swap mode of the HDD unit 22 as the swap mode of its cache memory according to the preset

priority (step S7).

The sequence shown in Figure 5 can also be used to select a swap mode of an object cache memory. It is premised in this case that a default swap mode (FIFO mode) is preset in the HDD unit 22.

The default swap mode (FIFO mode) preset in the HDD unit 22 is transferred to the HDC card 21 from the HDD unit 22 (step S11) when the computer system is started up. Receiving the swap mode from the HDD unit 22 (step S12), the HDC card 21 selects a swap mode (LFU mode) different from that of the HDD unit 22 as the swap mode of its cache memory (step S13) according to the preset priority.

Furthermore, if the HDC card 21 does not have any swap mode different from the default swap mode of the HDD unit 22, then it is possible for the HDC card 21 to notify the HDD unit 22 of the effect so that the HDD unit 22 selects a swap mode different from the default swap mode.

In Figure 4, the HDC card 21 can decide the object swap mode, but it is also possible to enable the HDD unit 22 to decide the object swap mode. In such a case, it is only required that the unit that cannot decide the swap mode has at least one swap mode. Also in Figure 5, the HDC card 21 may be replaced with the HDD unit 22 in position so that the HDC card 21 transfers the default swap mode to the HDD unit 22.

In this way, swap modes can be exchanged between the cache memory (high-ranking cache memory) of the HDC card 21 and the cache memory (low-ranking cache memory) of the HDD unit 22 when the computer system is started up, thereby those cache memories can select different swap modes. Consequently, each of those cache memories can select its optimized swap mode.

The swap mode selecting operation is not limited only to that described above; other various methods can also be employed.

For example, it is also possible that one cache memory transfers its unselectable swap modes to the other cache memory and the other cache memory selects its swap mode from those transferred swap modes.

Second Embodiment

The method for controlling cache memories in a second embodiment of the present invention prevents the high-ranking and low-ranking cache memories from holding common data. Concretely, the data transferred from the low-ranking cache memory to the high-ranking cache memory is erased from the low-ranking cache memory. With this method, the low-ranking cache memory can hold data more efficiently.

Figure 6 shows a schematic block diagram of a computer system in the second embodiment of the present invention.

In Figure 6, the same symbols are given to the same components as those shown in the first embodiment shown in Figure 1, avoiding redundant description.

As shown in Figure 6, the second bridge circuit 71 is connected to the expansion bus 20. And, a peripheral device (not illustrated), an HDD unit 72, and a CD-ROM unit 73 can be connected to the second bridge circuit 71 through a bus 70, a bus 74a, and a bus 74b respectively. The bus 70 is, for example, an ISA (Industry Standard Architecture) bus. The buses 74a and 74b are, for example, an ATA (AT Attachment)/ATAPI (ATA Packet Interface) bus.

The second bridge circuit 71 is used to connect a peripheral device to an expansion bus whose specification is different from that of the expansion bus 20. The second bridge circuit 71 is provided with functions of a PCI bus interface, an interruption controller, a DMA (Direct Memory Access) controller, an ISA bus interface, etc. The bus 70 is not limited only to the ISA bus and the buses 74a and 74b are not

limited only to the ATA/ATAPI bus.

The main memory 15 exchanges data with the HDD unit 72 according to the instructions issued from the CPU 11. This main memory 15 is provided with a disk cache 16 for holding part of the data stored in the HDD unit 72. The disk cache 16 is equivalent to a high-ranking cache memory with respect to the disk cache 75 provided in the HDD unit 72. The disk cache 16 has a storage capacity of, for example, several megabytes to several tens of megabytes. The main memory controller 17 provided in the first bridge circuit 14 controls the disk cache 16.

The disk cache 16 of the main memory 15 (precisely, the main memory controller 17), when it holds no data requested from the CPU 11, requests the HDD unit 72 for transfer of the data including both data requested from the CPU 11 and data predicted to be requested next time from the CPU 11, then holds the data transferred from the HDD unit 72 in response to the request. The data predicted to be requested next time is, for example, the data requested previously from the CPU 11 and/or the look-ahead data of the data requested from the CPU 11.

The configuration of the HDD unit 72 shown in Figure 6 is the same as that of the HDD 22 shown in Figure 1. Although the HDD unit 72 is provided with a disk cache 75 having functions of a low-ranking cache memory, the cache controlling of the HDD unit 72 is different from that of the disk cache 54 of the HDD unit 22. The disk cache 75 is composed so as to erase the data read from the HDD unit 72. The disk cache 75 has a storage capacity of, for example, several megabytes to several tens of megabytes.

The disk cache 75 of the HDD unit 72 (precisely, it is a cache controller (not illustrated)), when not holding the data requested from the main memory 15, reads the data from the magnetic disk 41. The data includes both data requested from the main memory 15 and data predicted to be requested next time

from the main memory 15. The disk cache 75 then holds the data requested from the main memory 15 and transfers the held data to the main memory 15. When the requested data is transferred to the main memory 15, the disk cache 75 erases the data and holds the data predicted to be requested next time from the main memory 15. The data predicted to be requested next time includes, for example, look-ahead data of the data requested from the main memory 15.

"Erase" in the present specification means not only erasing of data transferred to the main memory, but also updating of the control data (tag, valid flag, etc.) of the data held in the memory management area together with the data and other processings related to the erasure of the data, for example, specification of those areas for the data items as empty areas.

Figure 7 shows an example of the operation of a cache memory for holding data.

It is premised here that neither the disk cache 16 of the main memory 15 (cache memory of the host system) nor the disk cache 75 of the HDD unit 72 (cache memory of the low-ranking unit) holds the data requested from the CPU 11 (application program).

At first, an application program requests the main memory 15 for transfer of data. Then, a cache mishit occurs in the disk cache 16 of the main memory 15. The HDD unit 72 is thus requested for transfer of data A including both data requested from the application program and the data predicted to be requested next time from the application program. As a result, a cache mishit occurs in the disk cache 75 of the HDD unit 72. The data A requested from the main memory 15 is thus read from the magnetic disk 41 and held in the disk cache 75, and then it is transferred to the main memory 15. Consequently, the data A is held in the disk cache 16 of the main memory 15.

Furthermore, when the data A is transferred to the main memory 15, the HDD unit 72 erases the data A from its disk cache 75.

After that, the host system reads the data B in the address following that of the transferred data A from the magnetic disk 41 and holds the data B in its disk cache 75.

According to the conventional data reading method shown in Figure 12, the cache memory of the conventional HDD unit holds both data D3 requested from the host system and data D4 (look-ahead data) in the address following that of the data D3. On the contrary, the cache memory of the HDD unit in this second embodiment does not hold the data A requested from the host system, but holds only the data B (look-ahead data) in the address following that of the data A.

Consequently, if the storage capacity is the same in both cache memories, the cache memory of the HDD unit in the first embodiment can hold look-ahead data much more than that of the conventional HDD unit. In Figure 12, the data B1 means the look-ahead data D4 that can be held in the cache memory of the conventional HDD unit and the data B2 means the look-ahead data following the data B1 read from the magnetic disk after the data A is erased from the cache memory. Consequently, it is possible to make the most use of the localization of the data space so as to improve the hitting rate of the cache memory in the first embodiment. This is why the cache memory in this embodiment can hold data more efficiently.

There are two types for data localization; time localization and space localization. The time localization means a high possibility of accessing data in the same address repetitively. The space localization means a high possibility of accessing data in addresses close to the accessed data.

Figures 8 through 11 show another example of the operation of the cache memory for holding data.

As shown in Figure 8, the disk cache 16 of the main memory 15 employs data A2 requested last from an application program as the data predicted to be requested next time from the

application program. This is why the time localization of data can be used so as to improve the hitting rate of the disk cache 16.

As shown in Figure 9, the disk cache 16 of the main memory 15 employs the look-ahead data A2 of the data A1 requested from an application program as the data predicted to be requested next time from the application program. Consequently, it is possible to make the most use of the space localization of data so as to improve the hitting rate of the disk cache 16.

In the example shown in Figure 10, the disk cache 16 of the main memory 15 has a plurality of segments, for example, three segments of segment 1, segment 2, and segment 3. In those segments 1 to 3 are held data items requested to the HDD unit 72 repetitively according to the commands issued sequentially from an application program.

In the segment 1 are held data A1 requested by the command 1 from the application program and the look-ahead data A2 of the data A1. In the segment 2 are held data A3 requested by the command 2 from the application program and look-ahead data A4 of the data A3. And, in the segment 3 are held data A5 requested by the command 3 from the application program and look-ahead data A6 of the data A5. If a cache memory is divided into a plurality of segments in this way, therefore, it is possible to make the most use of the time localization of data so as to improve the hitting rate of the disk cache 16.

On the other hand, the disk cache 75 of the HDD unit 72 holds the latest data A5 requested from the main memory 15 and the look-ahead data B of the data A6. Consequently, it is possible to make the most use of the space localization of data so as to improve the hitting rate of the disk cache 75.

In the example shown in Figure 11, the disk cache 75 of the HDD unit 72 also has a plurality of segments, for example, three segments of segments 1 to 3. In those segments 1 to 3 of the

disk cache 16 of the main memory 15 are held data items; data A1, data A2, and data A3 requested to an HDD unit 72 according to command 1, command 2, and command 3 requested sequentially from an application program. In the segments 1 to 3 of the disk cache 75 of the HDD unit 72 are also held data items; look-ahead data C1, C2, and C3 of the data A1, A2, and A3 requested from the main memory 15 sequentially. Consequently, it is possible to make the most use of the time localization of data so as to improve the hitting rate of the disk cache 75.

The method for disposing a plurality of segments in a cache memory is not limited only to that described above; for example, other conventional well-known methods, such as direct mapping, full associative, set associative, etc. can also be employed. The size of each segment is not limited only to a fixed one; it may be variable.

In the second embodiment, there are two stages for the storage layer; the disk cache 16 (high-ranking cache memory) of the main memory 15 and the disk cache 75 (low-ranking cache memory) of the HDD unit 72. However, it is avoided that the low-ranking cache memory holds the data transferred to the high-ranking cache memory. On the contrary, if there are three storage layers of main memory, extended card, and HDD unit as described in the first embodiment, it can be specified that the cache memory of the extended card does not hold the data transferred from the extended card to the main memory and the cache memory of the HDD unit does not hold the data transferred from the HDD unit to the extended card.

In the first and second embodiments, the HDD unit is defined as the lowest-ranking storage unit. However, the lowest-ranking storage unit is not limited only to such an HDD unit; for example, such a disk unit as a magnetic tape unit, an optical magnetic disk unit, etc. may also be employed.

As described above, the present invention is especially suitable for peripheral devices provided with a cache memory

respectively. Because the primary cache, the secondary cache, and the main memory are usually installed on the mother board of the object computer system, the method of data transfer among those memories is designed optimally as an integrated system. On the contrary, because such conventional peripheral devices as an extended card, an HDD unit, etc. are usually designed independently, those devices are not always designed optimally with respect to the high-ranking storage layer. And, according to the present invention, it is avoided that the low-ranking cache memory holds the data held in the high-ranking cache memory as much as possible, thereby the hitting rate of the low-ranking cache memory can be improved with the most use of the localization of data. Consequently, the method of data transfer among those memories can be designed optimally with respect to the high-ranking storage layer.

Furthermore, the present invention is most effective when both high-ranking and low-ranking cache memories have approximately the same storage capacity. This is because if the low-ranking cache memory has a storage capacity extremely larger than the high-ranking cache memory, the amount of the data becomes small for the storage capacity of the low-ranking cache memory even when the same data is held in both high-ranking and low-ranking cache memories.

On the contrary, if both high-ranking and low-ranking cache memories have almost the same storage capacity, the amount of the data becomes large for the storage capacity of the low-ranking cache memory when the same data is held in both high-ranking and low-ranking cache memories. Consequently, it is possible for the low-ranking cache memory to hold such data as look-ahead data instead of the data transferred to the high-ranking cache memory, thereby the hitting rate of the low-ranking cache memory is relatively improved.

[Advantages of the invention]

According to the present invention, therefore, both high-ranking and low-ranking cache memories can be operated in different swap modes. It is thus possible to reduce the possibility that the low-ranking cache memory holds the same data as that held in the high-ranking cache memory, thereby each of those cache memories can hold data more efficiently.

Furthermore, according to the present invention, if a request of data transfer is issued from the host system to the low-ranking cache memory, the low-ranking cache memory erases the data transferred to the host system in response to the request. If a request of data transfer is issued from the high-ranking cache memory to the low-ranking cache memory, the low-ranking cache memory erases the data transferred to the high-ranking cache memory in response to the request. Consequently, the low-ranking cache memory can hold data more efficiently.

Brief Description of the Drawings:

Figure 1 is a schematic block diagram of a computer system in a first embodiment of the present invention;

Figure 2 is a schematic block diagram of the HDD unit shown in Figure 1;

Figure 3 is a schematic block diagram of the HDC card 21 shown in Figure 1;

Figure 4 shows an example of an operation sequence of a cache memory for selecting a swap mode;

Figure 5 shows another example of an operation sequence of a cache memory for selecting a swap mode;

Figure 6 is a schematic block diagram of a computer system in a

second embodiment of the present invention;

Figure 7 shows an example of the operation of a cache memory for holding data;

Figure 8 shows another example of the operation of the cache memory for holding data;

Figure 9 shows further another example of the operation of the cache memory for holding data;

Figure 10 shows further another example of the operation of the cache memory for holding data;

Figure 11 shows further another example of the operation of the cache memory for holding data; and

Figure 12 shows how data is read from a conventional disk cache.

[Description of symbols]

10a, 10b, 10c ... External Bus
 11 ... CPU
 12 ... Primary Cache
 13 ... Secondary Cache
 14 ... First Controller
 15 ... Main Memory
 16 ... Cache Area (High-ranking Cache memory)
 17 ... Cache Controller
 20 ... Expansion Bus
 21 ... HDC Card (Hard Disk Control Unit)
 22 ... HDD Unit
 23 ... CD-ROM Unit
 24a, 24b ... Bus
 31 ... HDC
 32 ... Cache Controller
 33 ... Control Memory

- 34 ... Buffer Memory (High-ranking or Low-ranking Cache Memory)
- 35 ... I/O Port
- 36 ... I/O Port
- 37 ... Host I/F
- 38a, 38b ... Bus
- 41 ... Magnetic Disk
- 42 ... Spindle Motor
- 43 ... Head Slider
- 44 ... Head Arm
- 45 ... Actuator
- 51 ... HDC
- 52 ... Cache Controller
- 53 ... Control Memory
- 54 ... Buffer Memory (Low-ranking Cache Memory)
- 55 ... Host I/F
- 56 ... Bus
- 70 ... Bus
- 71 ... Second Controller
- 72 ... HDD Unit
- 73 ... CD-ROM Unit
- 74a, 74b ... Bus
- 75 ... Buffer Memory (Low-ranking Cache Memory)

Claims:

- (1) A method for controlling a plurality of cache memories comprising a low-ranking cache memory and a high-ranking cache memory,

wherein said high-ranking cache memory and said low-ranking cache memory are operated in different swap modes.

- (2) The method for controlling a plurality of cache memories according to Claim 1, comprising the steps of:

exchanging select information for selecting a swap mode of each cache memory between said high-ranking cache memory and said low-ranking cache memory when a system is started up; and

by each cache memory, selecting a swap mode different from one of the other cache memory based on the exchanged select information.

- (3) The method for controlling a plurality of cache memories according to Claim 2,

wherein one of said high-ranking and low-ranking cache memories transfers select information for selecting a swap mode of the other cache memory to the other cache memory; and



said other cache memory selects its swap mode according to said select information transferred from said one cache memory.

- (4) The method for controlling a plurality of cache memories according to Claim 3,

said one cache memory selects its swap mode and transfers its selected swap mode to said other cache memory; and

said other cache memory selects a swap mode different from said transferred swap mode as its swap mode.

- (5) The method for controlling a plurality of cache memories according to Claim 3,

said one cache memory transfers a plurality of swap modes that it cannot select by itself to said other cache memory; and

said other cache memory selects one of said transferred swap modes.

- (6) The method for controlling a plurality of cache memories according to Claim 2,

wherein one of said high-ranking and low-ranking cache memories transfers a plurality of swap modes that it can select by itself to the other cache memory;

said other cache memory selects one of said transferred swap modes, which should be selected by said one cache memory and responds said selected swap mode to said one cache memory, then selects a swap mode different from said responded swap mode as its swap mode; and

said one cache memory selects said responded swap mode as its swap mode.



- (7) The method for controlling a plurality of cache memories according to Claim 1,

wherein said high-ranking and low-ranking cache memories are operated in different swap modes selected respectively from among the LRU (Least Recently Used) mode, the LFU (Least Frequently Used) mode, the FIFO (First-In First-Out) mode, and a mode for swapping out one of a plurality of segments, which holds the least amount of data.

- (8) The method for controlling a plurality of cache memories according to Claim 1,

wherein said high-ranking cache memory is a disk cache provided for a hard disk control unit for controlling a hard disk drive unit and said low-ranking cache memory is a disk cache provided for said hard disk drive unit.

- (9) A computer system, comprising:

a low-ranking cache memory; and

a high-ranking cache memory connected to said low-ranking cache memory;

wherein said high-ranking and low-ranking cache memories exchange select information for selecting a swap mode respectively with each other when said computer system is started up and select different swap modes according to said exchanged select information respectively.

- (10) The computer system according to Claim 9,

wherein one of said high-ranking and low-ranking cache memories has a transferring device for transferring select information for selecting a swap mode of the other cache memory to said other cache memory; and



said other cache memory has a selecting device for selecting its swap mode according to said select information received from said one cache memory.

- (11) The computer system according to Claim 10,

wherein said one cache memory has a mode selecting device for selecting its swap mode and said transferring device transfers a swap mode selected by said selecting device to the other cache memory; and

said selecting device of said other cache memory selects a swap mode different from said swap mode received from said one cache memory as a swap mode of said other cache memory.

- (12) The computer system according to Claim 10,

wherein said transferring device of said one cache memory transfers a plurality of swap modes that said one cache memory cannot select by itself, to the other cache memory;

said selecting device of said other cache memory selects one of a plurality of said swap modes received from said one cache memory as its swap mode.

- (13) The computer system according to Claim 9,

wherein one of said high-ranking and low-ranking cache memories has a candidate mode transferring device for transferring a plurality of swap modes that it can select by itself to the other cache memory and a mode selecting device for selecting said swap mode received from said other cache memory as its swap mode; and

said other cache memory has;

a partner's mode selecting device for selecting a swap



mode that said one cache memory should select, from among a plurality of said swap modes received from said one cache memory;

a partner's mode transferring device for transferring a swap mode selected by itself to said one cache memory; and

a self-mode selecting device for selecting a swap mode different from said swap mode selected by said partner's mode selecting device as a swap mode of said other cache memory.

(14) The computer system according to Claim 9,

wherein said high-ranking cache memory and said low-ranking cache memory are operated in different swap modes selected from LRU (Least Recently Used) mode, LFU (Least Frequently Used) mode, FIFO (First-In First-Out) mode, and a mode for swapping out a segment holding the least amount of data among a plurality of segments.

(15) The computer system according to Claim 9,

wherein said high-ranking cache memory is a disk cache provided for said hard disk control unit for controlling said hard disk drive unit; and

said low-ranking cache memory is a disk cache provided for said hard disk drive unit.

(16) The computer system, comprising:

a hard disk drive unit having a low-ranking cache memory for holding part of data stored in a magnetic disk; and

a hard disk control unit having a high-ranking cache memory for holding part of data stored in said hard disk drive unit and used to control said hard disk drive unit;



wherein said high-ranking and low-ranking cache memories exchange select information for selecting a swap mode respectively with each other when said computer system is started up and select different swap modes according to exchanged select information respectively.

(17) The computer system, comprising:

a hard disk drive unit provided with a cache memory for holding part of data stored in a magnetic disk; and

a main memory for holding part of data stored in a hard disk drive unit and sending/receiving data to/from said hard disk drive unit according to an instruction issued from a central processing unit;

wherein said cache memory and said main memory exchange select information for selecting a swap mode respectively with each other when said computer system is started up and select different swap modes according to said exchanged select information respectively.

(18) A hard disk drive unit, comprising:

a low-ranking cache memory for storing part of data stored in a magnetic disk and connected to a host system having functions of a high-ranking cache memory;

wherein each of said high-ranking and low-ranking cache memories exchanges select information for selecting a swap mode with said host system when a system is started up so that said high-ranking and low-ranking cache memories are operated in different swap modes; and

said host system selects a swap mode of said low-ranking cache memory according to said exchanged select information.



(19) A hard disk control unit, comprising:

a first connection terminal connected to an extended connection terminal of a peripheral device provided for a computer system;

a second connection terminal connected to a hard disk drive unit having a low-ranking cache memory for storing part of data stored in a magnetic disk; and

a high-ranking cache memory for holding part of data stored in said hard disk drive unit connected to said second connection terminal;

wherein said control unit controls said hard disk drive unit connected to said second connection terminal;

said high-ranking and said low-ranking cache memories exchange select information for selecting a swap mode of each cache memory with said hard disk drive unit connected to said second connection terminal respectively so as to be operated in different swap modes when said computer system connected to said first connection terminal is started up; and

said control unit selects a swap mode of said high-ranking cache memory according to said exchanged select information.

(20) A hard disk control unit, comprising:

a first connection terminal connected to an extended connection terminal of a peripheral device provided for a computer system provided with a host system having functions of a high-ranking cache memory;

a second connection terminal connected to a hard disk drive unit; and



a low-ranking cache memory for holding part of data stored in said hard disk drive unit connected to said second connection terminal;

wherein said hard disk control unit controls said hard disk drive unit connected to said second connection terminal;

each of said high-ranking cache memory and said low-ranking cache memory exchanges select information for selecting a swap mode with a host system so as to be operated in different swap modes when said computer system connected to said first connection terminal is started up; and .

said hard disk control unit selects a swap mode of said low-ranking cache memory according to said exchanged select information.

- (21) A method for controlling a plurality of cache memories by erasing data from a low-ranking cache memory after said data is transferred to a host system in response to a request of data transfer issued from said host system to said low-ranking cache memory.

- (22) The method for controlling a plurality of cache memories according to Claim 21,

wherein said low-ranking cache memory holds data predicted to be requested next time from said host system according to data requested from said host system.

- (23) The method for controlling a plurality of cache memories according to Claim 22,

wherein said data predicted to be requested next time from said host system is look-ahead data of said data requested by said host system.



- (24) A method for controlling a plurality of cache memories to hold part of data held in a low-ranking cache memory in a high-ranking cache memory,

wherein data transferred from said low-ranking cache memory to said high-ranking cache memory in response to a request of data transfer issued from said high-ranking cache memory to said low-ranking cache memory is erased from said low-ranking cache memory.

- (25) The method for controlling a plurality of cache memories according to Claim 24,

wherein said high-ranking cache memory requests said low-ranking cache memory for transfer of data comprising both data requested from said host system and first data predicted to be requested next time from said host system according to said data requested from said host system when said high-ranking cache memory does not hold said data requested from said host system; and

said low-ranking cache memory transfers said requested data to said high-ranking cache memory and holds second data predicted to be requested next time from said high-ranking cache memory according to said data requested from said high-ranking cache memory.

- (26) The method for controlling a plurality of cache memories according to Claim 25,

wherein said first predicted data is data requested previously from said host system; and

said second predicted data is look-ahead data of said data requested from said high-ranking cache memory.

- (27) The method for controlling a plurality of cache memories according to Claim 25,



wherein said first predicted data is look-ahead data of data requested from said host system; and

said second predicted data is look-ahead data of said data requested from said high-ranking cache memory.

- (28) The method for controlling a plurality of cache memories according to Claim 24,

wherein said high-ranking cache memory is a disk cache provided for said hard disk control unit for controlling said hard disk drive unit; and

said low-ranking cache memory is a disk cache provided for said hard disk drive unit.

- (29) A computer system, comprising:

a low-ranking cache memory; and

a high-ranking cache memory for holding part of data stored in said low-ranking cache memory;

wherein said low-ranking cache memory, when a request of data transfer is issued from said high-ranking cache memory to said low-ranking cache memory, erases said data transferred from said low-ranking cache memory to said high-ranking cache memory in response to said request.

- (30) The computer system according to Claim 29,

wherein said high-ranking cache memory, when not holding data requested from said host system, requests said low-ranking cache memory for transfer of data comprising both data requested from said host system and first data predicted to be requested next time from said host system according to said data requested from said host system; and



said low-ranking cache memory holds second data predicted to be requested next time from said high-ranking cache memory according to said data requested from said high-ranking cache memory.

- (31) The computer system according to Claim 29,

wherein said high-ranking cache memory is a disk cache provided for said hard disk control unit for controlling said hard disk drive unit; and

said low-ranking cache memory is a disk cache provided for said hard disk drive unit.

- (32) A computer system, comprising:

a hard disk drive unit having a low-ranking cache memory for holding part of data stored in a magnetic disk; and

a hard disk control unit having a high-ranking cache memory for holding part of data stored in said hard disk drive unit and used to control said hard disk drive unit;

wherein said low-ranking cache memory, when said hard disk control unit requests said hard disk drive unit for transfer of data, erases said data transferred from said hard disk drive unit to said hard disk control unit in response to said request.

- (33) A computer system, comprising:

a hard disk drive unit having a cache memory for holding part of data stored in a magnetic disk; and

a main memory for holding part of data stored in said hard disk drive unit and exchanging data with said hard disk drive unit in response to an instruction issued from a central processing unit;



wherein both of said cache memory and said main memory, when requested for transfer of data from said main memory to said hard disk drive unit, erase said data transferred from said hard disk drive unit to said main memory in response to said request.

(34) A hard disk drive unit, comprising:

a cache memory for storing part of data stored in a magnetic disk;

wherein said cache memory, when said hard disk drive unit is requested for transfer of data from a host system connected to said hard disk drive unit, erases said data transferred to said host system in response to said request.

(35) A hard disk control unit, comprising:

a first connection terminal connected to an extended connection terminal of a peripheral device provided for a computer system;

a second connection terminal connected to a hard disk drive unit; and

a cache memory for holding part of data stored in said hard disk drive unit connected to said second connection terminal;

wherein said control unit controls said hard disk drive unit connected to said second connection terminal; and

said cache memory, when requested by said host system for transfer of data through said first connection terminal, erases said data after said data is transferred to said host system in response to said request.



METHOD FOR CONTROLLING CACHE MEMORIES, COMPUTER SYSTEM, HARD
DISK DRIVE UNIT, AND HARD DISK CONTROL UNIT

ABSTRACT

A computer system is provided with an HDC card (21) connected to an expansion bus (20) and an HDD unit (22) connected to the HDC card (21). The HDC card (21) is provided with a disk cache (high-ranking cache memory) and the HDD unit (22) is provided with a disk cache (54) (low-ranking cache memory). The HDC card (21) and the HDD unit (22) exchange select information for selecting a swap mode of each cache memory when the system is started up, thereby selecting different swap modes according to the exchanged select information respectively. The object is to avoid holding common data in both the high-ranking and low-ranking cache memories, or to reduce the data held commonly, thereby having each of the cache memories hold data efficiently.

FIG 1

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6





Figure 7

Main memory disk cache

data A = Data requested from application program



HDD unit disk cache

data A = Data (to be erased) requested from main memory

data B = Look-ahead data of data A in HDD unit

= Look-ahead data B1 held in conventional HDD unit + Look-ahead data B2 held as a result of erasure of data A




Figure 8




Figure 9




Figure 10

Main memory disk cache

data A1 = Data requested by command 1 from application program
 data A2 = Look-ahead data of data A1
 data A3 = Data requested by command 2 from application program
 data A4 = Look-ahead data of data A3 program
 data A5 = Data requested by command 3 from application
 data A6 = Look-ahead data of data A5

Segment 1	data A1	data A2
Segment 2	data A3	data A4
Segment 3	data A5	data A6

HDD unit disk cache

data A = Data requested most recently from main memory = data A5 + data A6
 data B = Look-ahead data of data A held in HDD unit



Figure 11




Figure 12

