

REPLACED SHEETS		SG
✓ Description :	paged 1... to paged 18.....	
✓ Date(s) :	paged 19... to paged 23.....	
✓ Drawing(s) :	paged 25... to paged 25.....	
✓ Amount :	paged 24... to paged 2.....	
DATE OF RECEIPT		

ORIGINAL

**METHODS AND APPARATUS FOR CONTROLLING
A MULTI-SEGMENT CACHE MEMORY**

5 The present invention relates to a method and apparatus for controlling a multi-segment cache memory.

 Cache memories are used within computer storage systems. In a magnetic storage system, for example, the cache memory is located within the disk controller which controls operation of a magnetic disk drive. The disk
10 controller is configured for connection with a host computer through a data bus, such as a small computer systems interface (SCSI). The disk controller passes data back and forth between the disk drive and the host computer in response to memory instructions or commands
15 generated by the host computer.

 The magnetic disk drive stores the data on the surface of a rotating magnetic disk. The disk surface includes a plurality of generally concentric data tracks. Each data track is divided into sectors or
20 blocks. In one typical disk drive, each data track is divided into thirty five blocks, each block including 512 bytes of data. A byte of data is represented by an individual address in the storage system.

 The memory instructions cause the disk controller
25 to initiate read and write operations within the disk drive. Each memory instruction includes a starting address, a block length (the number of blocks required) and an instruction type, such as read or write. After the disk controller receives a read instruction, for
30 example, the controller signals the disk drive to retrieve the required number of blocks beginning at the starting address. After the data has been retrieved, the disk drive controller passes the data back to the host computer.



The cache memory improves storage system performance by reducing the number of times the controller must go to the disk to retrieve data. The cache memory stores the most recently accessed data in a random access memory (RAM). When the host computer runs a particular software application, the input/output pattern seen by the disk drive controller tend to repeat. Therefore, there is a greater chance that data obtained from the disk will be used again. When the host computer requests retrieval of data that is in the cache memory, the controller does not have to retrieve the information from the disk but can retrieve it directly from the cache memory. Because the cache memory has a much faster data access time than the disk drive, retrieving the data directly from the cache memory substantially reduces the data access time seen by the host computer.

The cache memory includes a cache buffer and a cache table. In one version, the cache buffer can store up to 192K bytes of data. The cache table maintains a list of the data present in the cache buffer by recording the starting addresses of the data retrieved from the disks. When the controller receives a write instruction that corresponds to addresses present in the cache memory, it up-dates the cache table such that it no longer indicates the data is present within the buffer. This prevents old data from being read from the cache buffer.

When the controller receives a read instruction, it refers to the cache table to determine whether the requested data is present in the cache buffer. If the data is not present in the cache buffer (cache miss), the controller retrieves the data from the disk and relays the data to the host computer. The controller also




loads the data into the cache buffer and up-dates the cache table. A cache hit occurs when the data requested by the host computer is present in the cache buffer. In this situation, the controller receives the information directly from the buffer and transmits the data to the host computer. The ratio of cache hits to cache misses is a measure of cache memory performance.

In a multi-segment cache, the cache buffer is divided into segments. The number of segments in the cache buffer is known as the segmentation level. By segmenting the cache buffer, the controller can perform "look ahead" operations by retrieving excess blocks ~~to~~ ^{from} ~~fill~~ an entire segment. Because memory operations tend to form repetitive and somewhat predictable patterns, the host computer ~~will~~ ^{is} likely request data from the next consecutive blocks. The segmentation level determines the number of blocks in each segment and therefore determines the number of excess blocks retrieved. The greater the segmentation level, the smaller the number of blocks in each segment. Performance of the cache memory may be controlled by controlling the number of excess blocks retrieved in a read operation.

The optimum segmentation level depends upon the particular software application running on the host computer. Because the host computer frequently switches between various applications, a selected segmentation level may not achieve the desired performance over a wide range of applications. Therefore, a cache having an adaptive segment control is needed.

Heuristic programming is one method of controlling segmentation level. A statistical performance analysis is performed on the cache memory and these results are applied to a set of rules. The variables applied to the rules include, among others, the number of blocks in



each read that are a cache hit and the number of blocks that are a cache miss. However, there are simply not enough information in ^apast statistical performance to predict a preferred segment^{ation} level. Because there are many exceptions to the rules, heuristic programs do not provide accurate, dependable results. Therefore, there is a continuing need for a multi-segment cache having improved segment control.

It is an object of the present invention to provide an adaptive segment control offering improved control of the performance of a multi-segment cache in a storage system.

According to a first aspect of the present invention, there is provided a method of controlling a cache memory in a storage system, the method comprising segmenting the cache memory to operate at a selected working segmentation level, and being characterised by segmenting at least one virtual cache table to operate at at least one further segmentation level, monitoring instructions received from a host computer, determining which of the working and further segmentation levels results in the highest hit ratio as a function of the monitored instructions, and selecting the segmentation level resulting in the highest hit ratio to be a new working segmentation level for the cache memory.

According to a second aspect of the present invention, there is provided a storage system a storage device comprising a cache memory operable at a selected working segmentation level, and characterised by at least one virtual cache table operable at at least one further segmentation level, a store for instructions received from a host computer, and means for determining which of the working and further segmentation levels results in the highest hit ratio as a function of the





stored instructions and for selecting the segmentation level resulting in the highest hit ratio to be a new working segmentation level for the cache memory.

5 A preferred embodiment of the invention described below features a controller ~~z~~ which monitors memory instructions transmitted by a host computer and received by the storage system. The cache memory is configured to operate at a selected initial working segmentation level. A plurality of virtual cache tables are
10 connected to receive the monitored memory instructions and are configured such that one table operates at the working segmentation level and the other tables operate at different segmentation levels. The monitored memory instructions are used to determine which virtual cache
15 table ^{will} obtains the highest hit ratio with its respective segmentation level. The working segmentation level is adjusted to equal the segmentation level of the virtual cache table having the highest hit ratio.

Any number of virtual cache tables may be used in
20 accordance with the present invention. In the described embodiment, there are two virtual cache tables. The first virtual cache table is segmented to operate at the selected working segmentation level. The second virtual cache table is segmented to operate at a selected test
25 segmentation level, the test segmentation level being different from the working segmentation level. The test segmentation level is selected by indexing between entries within a look up table which includes all selectable segmentation levels. In the preferred
30 embodiment, the selectable segmentation levels range from 1 to 16.

Preferably, the controller includes an instruction queue for storing the memory instructions received by the storage system. To ensure that the instructions



stored in the instruction queue are current and representative of the application being executed by the host computer, the instruction queue is flushed after the number of instructions stored is equal to a queue limit, such as thirty instructions. After the instruction queue has been flushed, the controller continues to store instructions transmitted in the storage system.

Hit ratio simulations are preferably performed on the first and second virtual cache tables by executing a selected number of instructions stored in the instruction queue. The storage system returns to an idle state where all instructions in the instruction queue have been simulated, and resumes hit ratio simulations after the next received instruction is stored in the instruction queue. This process is repeated until the selected number of instructions have been executed on the first and second virtual cache tables. For example, in one embodiment, one hundred instructions are simulated in the hit ratio simulations.

Hit ratio values are maintained during the hit ratio simulations and are compared after execution of the selected number of stored instructions. When the hit ratio value for the second virtual cache table is greater than the hit ratio value for the first virtual cache table, the working segmentation level is adjusted to equal the test segmentation level. The cache memory is then re-configured to operate at the new working segmentation level. Thereafter, the controller resumes the hit ratio simulations with the new working segmentation level and selects a new test segmentation level.

In one embodiment, the hit ratio simulations are only performed when the storage system is in an idle



state. When the storage system receives a memory instruction from the host computer, it operates in a busy state by performing the requested memory operation. Simulation of an individual stored instruction is interrupted upon receipt of the instruction from the host computer. This prevents the hit ratio simulations from interfering with the memory operations. After completion of the memory operation, the storage system returns to the idle state. Thereafter, the interrupted simulation is resumed upon completion of the memory operation. The overhead added by the segment control is therefore minimal because simulation is performed only when the storage system is idle.

The present invention improves performance of a multi-segment cache memory by analysing the instruction pattern received from the host computer. Because the instruction pattern is representative of the application being executed by the host computer, the controller adapts the segmentation level of the cache to fit the current application. The present invention therefore provides an adaptive segment controls that is both accurate and dependable over a wide range of software applications.

The present invention is described further, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a diagram of a memory controller connected between a host computer and a memory device in accordance with the present invention;

Figure 2 is a diagram of the memory controller shown in Figure 1, including an adaptive segment control unit in accordance with the present invention;

Figures 3a and 3b, when joined, form a flow chart, which illustrates operation of the memory controller

This Sheet Has Been Replaced



shown in Figures 1 and 2; and

Figure 4 is a flow chart of a simulation routine performed by the adaptive segment control unit shown in Figure 2.

5 The present invention provides an adaptive segment control unit (ASCU) for a multi-segment cache memory in a memory controller. The ASCU performs hit ratio simulations on virtual cache tables to determine the
10 cache memory segmentation level that matches the software application being run on a connected host computer. The hit ratio simulations performs actual instructions received from the host computer on the virtual cache table.

15 Figure 1 is a diagram of a memory controller 2 connected between a host computer 4 and a memory device 6. The memory controller 2 and the memory device 6 form a storage system for the host computer 4. The host computer 4 transmits memory instructions to the controller 2 through a data bus 8. The data bus 8 can
20 be any known communications interface, such as a small computer systems interface (SCSI). The controller 2 includes a micro-processor 10, a read only memory (ROM) 11, a random access memory (RAM) 12, a cache memory 13, and an adaptive segment control unit (ASCU) 14. The ROM
25 11 is connected to the micro-processor 10 through data and address lines 15. The ROM 11 holds control software which operates the micro-processor 10 to perform memory operations in response to the memory instructions received from the host computer 4. The RAM 12 is
30 connected to the micro-processor 10 through data and address lines 16. The cache memory 13 is connected to the micro-processor 10 through data and control lines 17. The adaptive segment control unit 14 is connected to the micro-processor 10 through data and control lines

This Sheet Has Been Replaced



18. The control of software also operates the micro-processor 10 to control operation of the cache memory 13 and the adaptive segment control unit 14.

5 The memory device 6 is connected to the micro-processor 10 through control lines 19 and a data bus 20. The micro-processor 10 operates the memory device 6 to perform memory operations in response to memory instructions received from the host computer 4. During a read operation, the host computer 4 generates a read
10 instruction, which is decoded by the micro-processor 10. The micro-processor 10 transmits control signals through the control lines 19 to the memory device 6, which causes the memory device to retrieve the requested data. The micro-processor 10 receives the data through the
15 data bus 20 and returns the data to the host computer 4 through the data bus 8. Similarly, during a write operation, the host computer 4 generates a write instruction which is decoded by the micro-processor 10. The micro-processor 10 transmits write data to the
20 memory device 6 through the data bus 20. The micro-processor 10 operates the memory device 6 through the control lines 19 to write the data to the desired memory locations within the memory device.

25 Figure 2 is a block diagram of parts of the memory controller 2 including the cache memory 13 and the adaptive segment control unit 14. The cache memory 13 stores the data most recently accessed from the memory device 6. Because the instruction pattern generated by the host computer 4 tends to repeat while it runs a
30 particular software application, data stored in the cache memory 13 is likely to be used again by the host computer. If the data is requested again by the host computer 4, it can be retrieved directly from the cache memory 13 rather than from the memory device 6. Since

This Sheet Has Been Replaced



10

data can be retrieved from the cache memory 13 much quicker than from the memory device 6, the cache memory can significantly increase storage system performance by reducing the average data access time.

5 The cache memory 13 includes a cache buffer 26 and cache table 28. The cache buffer 26 includes a plurality of storage locations 40 and 42 for storing data read from the memory device 6. In one embodiment, the cache buffer 26 can store up to 192K bytes of data.
10 The cache memory 13 is a multi-segment cache. In other words, the cache buffer 26 is divided into segments. The number of segments is known as the segmentation level. For example, in the embodiment shown in Figure 2, the cache memory 13 has a segmentation level of 2,
15 the cache buffer 26 being divided into segments 30 and 32. The segment 30 includes the storage locations 40. Similarly, the segment 32 includes the storage locations 42. The cache memory 13 can be configured with any segmentation level.

20 The memory device 6 (shown in Figure 1) reads data in blocks or sectors. In one embodiment, a block includes 512 bytes of data. A read instruction generated by the host computer 4 includes a starting address and a number of blocks requested. In response
25 to each read instruction, the micro-processor 10 retrieves the requested number of blocks of data from the memory device 6, beginning at the starting address, and enough additional blocks to fill an entire segment of the cache buffer 26. Each storage location 40 and 42
30 in the buffer 26 can store a block of data. The segmentation level determines the number of blocks in each segment. The retrieval of additional blocks is known as "look ahead" operation. Because of particular application run by the host computer 4 tends to form a

This Sheet Has Been Replaced



pattern of read and write instructions, the additional blocks stored in the cache buffer 26 have a relatively high probability of being used. Therefore, the look ahead operation further reduces the average data access time of the storage system.

5

The cache table 28 maintains a list of the data present in the cache buffer 26. The cache table 28 includes a plurality of storage locations 44 for storing the starting addresses of each segment of data stored in the cache buffer 26. The cache table 28 also stores the number of blocks of data within each segment in the cache buffer 26. The micro-processor 10 can determine from the cache table 28 whether a particular block of data is present in the cache buffer 26 by looking at the starting addresses and the number of blocks in each segment (or segmentation level). For a segmentation level of 2, the cache table 28 includes two starting addresses in locations 44. For a segmentation level of 3 (not shown), the cache table 28 includes three starting addresses.

10

15

20

When the micro-processor 10 receives a read instruction, it refers to the cache table 28 to determine whether the requested data is already present in the cache buffer 26. If the requested data is present in the cache buffer 26 (a cache hit), the controller 2 retrieves the requested data from the cache buffer and transmits the data to the host computer 4 through the data bus 8. If the cache table 28 indicates that the requested data is not present within the cache buffer 26, (a cache miss), the micro-processor 10 retrieves the requested data from the memory device 6 through the data bus 20. After the data has been retrieved, the micro-processor 10 stores the data in the cache buffer 26 and up-dates the cache table 28.

25

30

For the embodiment shown in Figure 2, a first read instruction causes the segment 30 of the cache buffer 26 to be filled with blocks of data from the memory device 6, beginning at the starting address specified in the instructions. A second read instruction causes the segment 32 to be filled, beginning at the starting address specified in the second read instruction. The starting addresses for each read instruction are stored in the cache table 28, along with the segmentation levels.

The adaptive segment control unit 14 improves cache memory performance by dynamically adjusting the segmentation level of the cache buffer 26. The optimum segmentation level varies, depending upon the particular software application being run by the host computer 4. Because the host computer 4 runs a variety of software applications, a selected segmentation level may not be optimal for all the applications. The adaptive segment control unit 14 monitors the instructions received by the micro-processor 10 and adapts the segmentation level of the cache memory 13 to fit the pattern of instructions received from the host computer 4.

The adaptive segment control unit 14 includes an instruction queue 50 and virtual cache tables 52 and 54. The instruction queue 50 includes a plurality of storage locations 56, each for storing an instruction received from the host computer 4. In one embodiment, the instruction queue 50 can store up to thirty consecutively received instructions. When the controller 2 is in an idle state, the adaptive segment control unit 14 performs hit ratio simulations on the virtual cache tables 52 and 54, based on the stored instructions. The results of the hit ratio simulations are used to adjust the working segmentation level of the

This Sheet Has Been Replaced



13

cache memory 13 to fit the current software application being run by the host computer 4.

Prior to performing the hit ratio simulations, the virtual cache table 52 is segmented to the working segmentation level of the cache memory 13. In the embodiment shown in Figure 2, the working segmentation level of the cache memory 13 is equal to 2. Therefore, the virtual cache table 52 includes two storage locations 60 for storing two starting addresses. Similarly, the virtual cache table 54 includes storage locations 62 for storing starting addresses. The virtual cache table 54 is segmented to a selected test segmentation level which is different from the working segmentation level.

The hit ratio simulations are performed by simulating a selected number of stored instructions (simulation count) on the virtual cache tables 52 and 54. While simulating the instructions, the adaptive segment control unit 14 maintains a count of the number cache hits and the number of cache misses for each table 52 and 54. Because the tables 52 and 54 are segmented to different segmentation levels, the hit ratios (cache hits/cache misses) of each table will be different after the instructions have been simulated. If the hit ratio of the virtual cache table 54 is greater than the hit ratio of the virtual cache table 52 after the simulation count, the working segmentation level of the cache memory 13 is adjusted to equal the test segmentation level. After the hit ratio simulation is completed, the virtual cache table 52 is re-configured to the new working segmentation level and the virtual cache table 54 is re-configured to a new test segmentation level. The hit ratio simulations are then repeated with a new, current set of stored instructions.

The test segmentation level is selected from a

This Sheet Has Been Replaced



look up table (not shown). The look up table contains all of the selectable segmentation levels. In one embodiment, the selectable segmentation levels range from 1 to 16. The adaptive segment control unit 14 indexes between elements in the look up table from one hit ratio simulation to the next. The entries can be indexed in any pre-selected order or in a random order.

In an alternative embodiment, the adaptive segment control unit 14 includes a plurality of virtual cache tables, as indicated by broken lines 64. In this embodiment, each table is configured to a different segmentation level. After the simulations are completed, the cache memory 13 is re-configured to a segmentation level equal to the segmentation level of the virtual cache table having the highest hit ratio.

Figures 3a and 3b, when joined, form a flow chart, which illustrates the operation of the controller 2 in greater detail. At step 70, an instruction is received from the host computer 4 through the data bus 8. At step 72, micro-processor 10 decodes the instruction into a starting address, a block length and an instruction type. The different types of instructions include, among others, read and write instructions.

At step 72, the adaptive segment control unit 14 determines whether the instruction queue 50 is full. If not, the received instruction is stored in the instruction queue 50 at step 76. Next, the micro-processor 10 executes the received instruction by performing a memory operation at step 78. The micro-processor 10 operates in a busy state (normal operation) while it executes the received instruction to prevent interference from operation of the adaptive segment control unit 14.

If the instruction queue 50 is full at step 74, the adaptive segment control unit 14 flushes the

This Sheet Has Been Replaced



15

instruction queue at step 80. Preferably, individual instructions are executed when the micro-processor 10 is in the idle state. If the micro-processor 10 is too busy to process the instruction simulations, the
5 instruction queue 50 will fill up to the queue limit. By flushing the instruction queue 50 when it is full, the adaptive segment control unit 14 ensures that the instructions stored in the queue are current and representative of the particular application being run
10 by the host computer 4.

After the instruction queue 50 has be flushed, the micro-processor 10 executes the received instruction at step 78. The adaptive segment control unit 14 waits until after a selected number of instructions (wait
15 count) has been subsequently received to begin refilling the instruction queue 50. This allows the micro-processor 10 to return to a relatively inactive state before the received instructions are stored to further ensure that the instructions in the queue 50 are
20 current.

At step 82, the adaptive segment control unit 14 determines whether an instruction simulation has been previously interrupted at step 90 (discussed in detail below). If not, the micro-processor 10 returns to an
25 idle loop at step 74, in which it operates in the idle state. If an instruction is received from the host computer 4 while the micro-processor 10 is in the idle loop, the micro-processor 10 returns to step 72 to decode the received instruction.

30 If no instruction is received from the host computer 4, the adaptive segment control unit 14 determines whether there is an instruction to simulate in the queue 50 (step 86). If the queue 50 is empty, the micro-processor 10 returns to the idle loop at step

This Sheet Has Been Replaced




84 and waits for an instruction from the host computer
4. If the instruction queue 50 is not empty and the
micro-processor 10 is in the idle state, the adaptive
segment control unit 14 simulates one of the instruc-
5 tions stored in the queue on the virtual cache tables 52
and 54 (step 88). After one instruction has been
simulated, the adaptive segment control unit 14 stores
the resulting cache hit and cache miss information for
the virtual cache tables 52 and 54 (step 89). The
10 micro-processor 10 returns to the step 86 to determine
whether there are any instructions left in the
instruction queue 50. If so, the next instruction is
simulated at step 88. This process is repeated until
the number of simulated instructions is equal to the
15 simulation count, such as one hundred.

If an instruction is received from the host
computer 4 during simulation of a stored instruction,
the micro-processor 10 interrupts the simulation at step
90. The adaptive segment control unit 14 saves the
20 state of the simulation so that it can be resumed after
the received instruction is executed. The micro-
processor 10 returns to step 72 to decode and execute
the received instruction. After the instruction has
been executed at step 78, the adaptive segment control
25 unit 14 determines, at step 82, that the simulation has
been interrupted. The state of the simulation which was
stored at step 90 is restored at step 92 and the
simulation is resumed at step 88. In this manner, the
normal execution of instructions receives the highest
30 priority so that response time is not affected by
operation at the adaptive segment control unit 14. The
overhead added by the adaptive segment control unit
simulation is minimal therefore.

Figure 4 is a flow chart which describes in

This Sheet Has Been Replaced



17

greater detail the operation of a simulation routine performed by the adaptive segment control unit 14. At step 100, the cache memory 13 is segmented to a selected working segmentation level. The virtual cache table 52 is also segmented to the working segmentation level (step 102). At step 104, the adaptive segment control unit 14 selects the test segmentation level from the look up table. At step 106 the virtual cache table 54 is segmented to the test segmentation level. Next, the adaptive segment control unit 14 simulates one hundred instructions on the virtual cache tables 52 and 54 (step 108).

The instructions are simulated by removing a stored instruction from the instruction queue 50 and processing that instruction on the virtual cache table 52 and 54. For example, with a read instruction, the adaptive segment control unit 14 determines how many of the requested blocks of data are a cache hit and how many are a cache miss in each virtual cache table 52 and 54. Right instructions are also processed on the virtual cache tables 52 and 54. Any addresses written by the right instructions are flushed in each cache table. The number of cache hits and cache misses are summed over the simulation count of one hundred instructions.

Upon completion of the simulation, the ratio of cache hits to cache misses of the virtual cache table 52 (hit ratio value HR1) is compared at step 110 with that of the virtual cache table 54 (hit ratio value HR2). If the hit ratio value HR2 is greater than the hit ratio value HR1, the working segmentation level is adjusted to equal the test segmentation level at step 112. The simulation routine then returns to steps 100 and 102 to segment the cache memory 13 and the virtual cache table

This Sheet Has Been Replaced



52 to the new working segmentation level. A new test segmentation level is selected from the look up table at step 104. The virtual cache table 54 is re-configured to the new test segmentation level at step 106.

5 If the hit ratio value HR2 is not greater than the hit ratio value HR1, the working segmentation level is not changed. The cache memory 13 and the virtual cache table 52 remain at the previous working segmentation level. A new test segmentation level is selected from
10 the look up table at step 104 and the virtual cache table 54 is re-configured to the new test segmentation level at step 106.

 The simulation sub-routine shown in Figure 4 is continuously repeated so that if the instruction pattern
15 produced by the host computer 4 changes, the adaptive segment control of the present invention will adjust the working segmentation level accordingly. As a result, the performance of the cache memory 13 is improved by
20 segmentation regardless of the particular software application being run by the host computer 4. The adaptive segment control of the present invention is also more accurate and reliable than other segment control
25 approaches, such as Heuristic programming, because the present invention uses a simulation approach on actual, current instructions, rather than pre-selected
30 statistical rules. Further, the overhead added by the present invention is minimal because simulation is performed while the memory controller is idle and does not require any attention from the host computer.



* C L M S *

ORIGINAL
replaced

This Sheet Has Been Replaced



19

claim
7/5/93

C L A I M S

1. A method of controlling a cache memory in a storage system, the method comprising segmenting the cache memory (13) to operate at a selected working segmentation level, and being characterised by segmenting at least one virtual cache table (52, 54) to operate at at least one further segmentation level, monitoring instructions received from a host computer (4), determining which of the working and further segmentation levels results in the highest hit ratio as a function of the monitored instructions, and selecting the segmentation level resulting in the highest hit ratio to be a new working segmentation level for the cache memory.

10/letter see
Jan 18/93

2. A method according to claim 1, ^{therein} characterised in ~~that the~~ ^{said} step of segmenting at least one virtual cache table comprises segmenting a first virtual cache table (52) to operate at the selected working segmentation level; ~~also~~ segmenting a second virtual cache table (54) to operate at a selected test segmentation level, different from the working segmentation level; and ~~is~~ ^{said} ~~that the~~ step of determining comprises performing hit ratio simulations on the first and second virtual cache tables by executing a selected number of the monitored instructions.

3. A method according to claim 2, ^{therein} characterised in ~~that the~~ ^{said} step of segmenting the second virtual cache table comprises selecting the test segmentation level from entries within a look up table, which includes all selectable segmentation levels.

4. A method according to claim 2 or 3 ^{therein} further characterised by repeating the steps of segmenting the first and second virtual cache tables, ^{as regards} performing hit



20

are performed
ratio simulations, ^{by} and selecting a new working segmentation level, ^{is selected}

5. A method according to any of claims 2 to 4 characterised in that the step of monitoring comprises storing the instructions in an instruction queue (50).

6. A method according to claim 5 ~~further characterised by the steps of flushing the instruction queue~~ *is flushed* after the number of stored instructions is equal to a queue limit, ~~and repeating the step of storing the instructions~~ ^{is repeated} after the number of instructions received subsequently from the host computer is equal to a wait count.

7. A method according to claim 5 or 6 ~~characterised in that the step of performing hit ratio simulations~~ *is performed* comprises ~~the steps of~~ executing one of the stored instructions on the first and second virtual cache tables when the instruction queue contains at least one stored instruction and the storage system is in an idle state, repeating the step of executing one of the stored instructions until the selected number of stored instructions has been executed, and returning the storage system to the idle state when the instruction queue is empty.

8. A method according to claim 7 ~~further characterised by the steps of suspending execution of the stored instructions~~ ^{is suspended} upon receipt of an instruction from the host computer, ~~performing a memory operation~~ ^{is performed} on the storage system in response to the received instruction, and ~~resuming the execution of the stored instructions~~ ^{is resumed} upon completion of the memory operation.

9. A storage device ~~comprising~~ ^{having} a cache memory (13) operable at a selected working segmentation level, and characterised by ~~at least one virtual cache table~~ (52, 54) operable ~~at~~ at least one further segmentation level,

in accordance with

This Sheet Has Been Replaced



5 //a store (50) for instructions received from a host computer; and means (10) for determining which of the working and further segmentation levels results in the highest hit ratio, as a function of the stored instructions, and for selecting the segmentation level resulting in the highest hit ratio to be a new working segmentation level for the cache memory.

10 10. A storage device according to claim 9 ~~characterised~~ *character* in that the at least one virtual cache table comprises a first virtual cache table (52) arranged to be operated at the selected working segmentation level and a second virtual cache table (54) arranged to be operated at a selected test segmentation level different from the working segmentation level, and in that the means for determining and selecting comprise means for performing hit ratio simulations on the first and second virtual cache tables by executing a selected number of the stored instructions.

20 11. A storage device according to claim 10, ~~further characterised~~ *further* by a look up table, ~~which includes~~ *which* all selectable segmentation levels, and from which the test segmentation level is selected.

25 12. A storage device according to any of claims 9 to 11 ~~characterised in that~~ *the* the store ~~comprises~~ *comprises* an instruction queue, and ~~in that~~ means (10) are provided for flushing the instruction queue, after the number of stored instructions ~~is equal to~~ *is* a queue limit, and storing further instructions after the number of instructions received subsequently from the host computer is equal to a wait count.

30 13. A method of controlling ~~performance of~~ a cache in a storage system, characterised by segmenting the cache (13) to operate at a selected working segmentation level; monitoring instructions transmitted by a host

This Sheet Has Been Replaced



22



computer (4) and received by the storage system;
//segmenting a plurality of virtual cache tables (52, 54)
such that one table operates at the working segmentation
level and the other tables operate at different segmen-
5 tation levels; //determining which virtual cache table
obtains the highest hit ratio with its respective
segmentatuion level as a function of the monitored
instructions; and //adjusting the working segmentation
level to equal the segmentation level of the virtual
10 cache table having the highest hit ratio.

17 14. A method of controlling performance of a cache in
a disk drive controller, characterised by segmenting the
cache (13) to operate at a selected working segmentation
level, receiving instructions transmitted by a host
15 computer (4), performing memory operations with the disk
drive controller in response to the received instruc-
tions, storing the received instructions in an
instruction queue (50), segmenting a first virtual cache
table (52) to operate at the selected working segmen-
20 tation level, segmenting a second virtual cache table (54)
to operate at a selected test segmentation level being
different from the working segmentation level,
performing hit ratio simulations on the first and second
virtual cache tables by executing a selected number of
25 stored instructions, maintaining first and second hit
ratio values for the first and second virtual cache
tables respectively, comparing the first and second hit
ratio values after execution of the selected number of
stored instructions, adjusting the working segmentation
30 level to equal the test segmentation level when the
second hit ratio value is greater than the first hit
ratio value, and repeating the steps of storing the
received instructions, segmenting the first and second
virtual tables, performing hit ratio simulations,

Let see
18 Jan 93



maintaining first and second hit ratio values, and comparing the first and second hit ratio values, after completion of the step of adjusting the working segmentation level.

5 15. An adaptive segment control assembly for
controlling performance of a multi-segment cache in a
storage system, the assembly being characterised by
means (10) for segmenting the cache (13) at a selected
working segmentation level, means (10) for monitoring
10 instructions transmitted by a host computer and received
by the storage system, a plurality of virtual cache
tables (52, 54) connected to the monitoring means, means
(10) for segmenting the plurality of virtual cache
tables such that one table operates at the working
segmentation level and the other tables operate at
different segmentation levels, means (10) for deter-
mining which virtual cache table obtains the highest hit
ratio with its respective segmentation level as a
function of the monitored instructions, and means (10)
20 for adjusting the working segmentation level to equal
the segmentation level of the virtual cache table having
the highest hit ratio.

25

This Sheet Has Been Replaced



30

See
letter
18 Jun 93

A B S T R A C T**METHODS AND APPARATUS FOR CONTROLLING
A MULTI-SEGMENT CACHE MEMORY**

5

10

15

20

The present invention relates to the control of a cache memory (13) in a storage system (2). According to the invention, the cache memory is segmented to operate at a selected working segmentation level, and at least two virtual cache tables (52, 54) are segmented to operate, respectively, at the working segmentation level and at least one different segmentation level. Instructions received from a host computer (4) are stored in an instruction queue (50). Then, while the storage system is in an idle state, hit ratio simulations are performed on the virtual cache tables by executing a selected number of instructions stored in the instruction queue. The working segmentation level is subsequently adjusted to equal the segmentation level of the virtual cache table having the highest hit ratio.

[Figure 2]

25

30

This Sheet Has Been Replaced



1/5





Fig. 1



2/5-






Fig. 3a



4/5





Fig. 3b



5/5-





Fig. 4





-3-


REPLACEMENT SHEETS		SG
☒ Description :	paged 3-5, 58, 9, 17	paged
☒ Page(s) :	paged 11	to paged 22
☒ Page(s) :	paged	to paged
☒ Index :	paged 23	to paged
DATE OF RECEIPT ... 30/7/99		

loads the data into the cache buffer and up-dates the cache table. A cache hit occurs when the data requested by the host computer is present in the cache buffer. In this situation, the controller receives the information directly from the buffer and transmits the data to the host computer. The ratio of cache hits to cache misses is a measure of cache memory performance.

In a multi-segment cache, the cache buffer is divided into segments. The number of segments in the cache buffer is known as the segmentation level. By segmenting the cache buffer, the controller can perform "look ahead" operations by retrieving excess blocks to fill an entire segment. Because memory operations tend to form repetitive and somewhat predictable patterns, the host computer is likely to request data from the next consecutive blocks. The segmentation level determines the number of blocks in each segment and therefore determines the number of excess blocks retrieved. The greater the segmentation level, the smaller the number of blocks in each segment. Performance of the cache memory may be controlled by controlling the number of excess blocks retrieved in a read operation.

The optimum segmentation level depends upon the particular software application running on the host computer. Because the host computer frequently switches between various applications, a selected segmentation level may not achieve the desired performance over a wide range of applications. Therefore, a cache having an adaptive segment control is needed.

Heuristic programming is one method of controlling segmentation level. A statistical performance analysis is performed on the cache memory and these results are applied to a set of rules. The variables applied to the rules include, among others, the number of blocks in each read that are a cache hit and the number of blocks that are a cache miss. One example of heuristic programming to control segmentation level is given in an Article entitled "Intelligent Controllers for Disk Drives Boosts Performance of Micros" from Computer Technology Review, Summer 1986, No. 3 at pages 133 to



139. The article assesses the advantages of segmented cache buffers and suggests that it should be possible to maintain a record of the operations carried out by a storage system, carry out a statistical analysis of these operations and then alter the number and size of segments in the cache memory in dependence upon the statistics.

However, there is simply not enough information in a past statistical performance to predict a preferred segmentation level. Given that there are many exceptions to the rules, heuristic programs do not provide accurate and dependable results. Therefore, there is a continuing need for a multi-segment cache having improved segment control.

It is an object of the present invention to provide an adaptive segment control offering improved control of the performance of a multi-segment cache in a storage system.

According to a first aspect of the present invention, there is provided a method of controlling a cache memory in a storage system, comprising:

- segmenting the cache memory into a number of segments of equal size to operate at a selected working segmentation level, the segmentation level corresponding to the number of segments;

- receiving instructions transmitted by a host computer; characterised by:

- storing the received instructions in an instruction queue;

- configuring a first virtual cache table to operate at the selected working segmentation level;

- configuring at least one further virtual cache table to operate at a selected test segmentation level, the test segmentation level(s) being different from the working segmentation level;

- performing hit ratio simulations on the first and further virtual cache tables by executing a selected number of stored instructions, to determine which virtual cache table obtains a highest hit ratio based upon the stored instructions; and

- adjusting the working segmentation level to equal the segmentation level of the virtual cache table having the greater/greatest hit ratio value.

According to a second aspect of the present invention, there is provided an adaptive segment control assembly for controlling a multi-segment cache memory in a storage system, the assembly comprising:

means for segmenting the cache memory into a number of segments of equal size at a selected working segmentation level, where the segmentation level corresponds to the number of segments;

means for receiving instructions transmitted by a host computer; characterised by:

means for storing the received instructions in an instruction queue;


means for configuring a first virtual cache table to operate at the selected working segmentation level;

means for configuring at least one further virtual cache table to operate at a selected test segmentation level, the test segmentation level(s) being different from the working segmentation level;

means for performing hit ratio simulations on the first and further virtual cache tables by executing a selected number of stored instructions to determine which virtual cache table obtains a highest hit ratio based upon the stored instructions; and


means for adjusting the working segmentation level to equal the segmentation level of the virtual cache table having the greater/greatest hit ratio value.


A preferred embodiment of the invention, described below, features a controller which monitors memory instructions transmitted by a host computer and received by the storage system. The cache memory is configured to operate at a selected initial working segmentation level. A plurality of virtual cache tables are connected to receive the monitored memory instructions and are configured such that one table operates at the working segmentation level and the other tables operate at different segmentation levels. The monitored memory instructions are used to determine which virtual cache tables would obtain the highest hit ratio with its respective segmentation level. The working segmentation level is adjusted to equal the segmentation level of the virtual cache table having the highest hit ratio.



Any number of virtual cache tables may be used in accordance with the present invention. In the described embodiment, there are two virtual cache tables. The first virtual cache table is segmented to operate at the selected working segmentation level. The second virtual cache table is segmented to operate at a selected test segmentation level, the test segmentation level being different from the working segmentation level. The test segmentation level is selected by indexing between entries within a look up table which includes all selectable segmentation levels. In the preferred embodiment, the selectable segmentation levels range from 1 to 16.

Preferably, a controller includes an instruction queue for storing the memory instructions received by the storage system. To ensure that the instructions _____





18. The control software also operates the micro-processor 10 to control operation of the cache memory 13 and the adaptive segment control unit 14.

The memory device 6 is connected to the micro-processor 10 through control lines 19 and a data bus 20. The micro-processor 10 operates the memory device 6 to perform memory operations in response to memory instructions received from the host computer 4. During a read operation, the host computer 4 generates a read instruction, which is decoded by the micro-processor 10. The micro-processor 10 transmits control signals through the control lines 19 to the memory device 6, which causes the memory device to retrieve the requested data. The micro-processor 10 receives the data through the data bus 20 and returns the data to the host computer 4 through the data bus 8. Similarly, during a write operation, the host computer 4 generates a write instruction which is decoded by the micro-processor 10. The micro-processor 10 transmits write data to the memory device 6 through the data bus 20. The micro-processor 10 operates the memory device 6 through the control lines 19 to write the data to the desired memory locations within the memory device.


Figure 2 is a block diagram of parts of the memory controller 2 including the cache memory 13 and the adaptive segment control unit 14. The cache memory 13 stores the data most recently accessed from the memory device 6. Because the instruction pattern generated by the host computer 4 tends to repeat while it runs a particular software application, data stored in the cache memory 13 is likely to be used again by the host computer. If the data is requested again by the host computer 4, it can be retrieved directly from the cache memory 13 rather than from the memory device 6. Since

30 JUL 1999

greater detail the operation of a simulation routine performed by the adaptive segment control unit 14. At step 100, the cache memory 13 is segmented to a selected working segmentation level. The virtual cache table 52 is also segmented to the working segmentation level (step 102). At step 104, the adaptive segment control unit 14 selects the test segmentation level from the look up table. At step 106 the virtual cache table 54 is segmented to the test segmentation level. Next, the adaptive segment control unit 14 simulates one hundred instructions on the virtual cache tables 52 and 54 (step 108).

The instructions are simulated by removing a stored instruction from the instruction queue 50 and processing that instruction on the virtual cache table 52 and 54. For example, with a read instruction, the adaptive segment control unit 14 determines how many of the requested blocks of data are a cache hit and how many are a cache miss in each virtual cache table 52 and 54. Write instructions are also processed on the virtual cache tables 52 and 54. Any addresses written by the write instructions are flushed in each cache table. The number of cache hits and cache misses are summed over the simulation count of one hundred instructions.

Upon completion of the simulation, the ratio of cache hits to cache misses of the virtual cache table 52 (hit ratio value HR1) is compared at step 110 with that of the virtual cache table 54 (hit ratio value HR2). If the hit ratio value HR2 is greater than the hit ratio value HR1, the working segmentation level is adjusted to equal the test segmentation level at step 112. The simulation routine then returns to steps 100 and 102 to segment the cache memory 13 and the virtual cache table



19

Claims



1. A method of controlling a cache memory (13) in a storage system, comprising:
segmenting the cache memory into a number of segments (30) of equal size to operate at a selected working segmentation level, the segmentation level corresponding to the number of segments;
receiving instructions transmitted by a host computer (4); characterised by:
storing the received instructions in an instruction queue;
configuring a first virtual cache table to operate at the selected working segmentation level;
configuring at least one further virtual cache table to operate at a selected test segmentation level, the test segmentation level(s) being different from the working segmentation level;
performing hit ratio simulations on the first and further virtual cache tables by executing a selected number of stored instructions, to determine which virtual cache table obtains a highest hit ratio based upon the stored instructions; and
adjusting the working segmentation level to equal the segmentation level of the virtual cache table having the greater/greatest hit ratio value.
2. A method according to claim 1, wherein determining which virtual cache table obtains a highest hit ratio comprises:
determining a hit ratio for each virtual cache table based upon the stored instructions; and
comparing the hit ratios for the virtual cache tables.
3. A method according to claim 1 or claim 2, wherein configuring the at least one further virtual cache table comprises selecting the test segmentation level from entries within a look up table, which includes all selectable segmentation levels.



4. A method according to claim 3, wherein the steps of configuring the first and further virtual cache tables, performing hit ratio simulations, and selecting a new working segmentation level, are repeated.

5. A method according to any one of the preceding claims, wherein the instruction queue is flushed after the number of stored instructions becomes equal to a queue limit and storing the instructions is repeated after the number of instructions received subsequently from the host computer becomes equal to a wait count.

6. A method according to claim 5, wherein performing hit ratio simulations comprises:

executing one of the stored instructions on the first and further virtual cache tables when the instruction queue contains at least one stored instruction and the storage system is in an idle state;

repeating the step of executing one of the stored instructions until the selected number of stored instructions has been executed; and

returning the storage system to the idle state when the instructions queue is empty.

7. A method according to claim 6, wherein

execution of the stored instructions is suspended upon receipt of an instruction from the host computer;

a memory operation is performed on the storage system in response to the received instructions; and

execution of the stored instructions is resumed upon completion of the memory operation.

REPLACEMENT SHEET

30 JUL 1999



8. An adaptive segment control assembly for controlling a multi-segment cache memory (13) in a storage system, the assembly comprising:

means (10) for segmenting the cache memory (13) into a number of segments (30) of equal size at a selected working segmentation level, where the segmentation level corresponds to the number of segments;

means for receiving instructions transmitted by a host computer (4);
characterised by:

means for storing the received instructions in an instruction queue;

means for configuring a first virtual cache table to operate at the selected working segmentation level;

means for configuring at least one further virtual cache table to operate at a selected test segmentation level, the test segmentation level(s) being different from the working segmentation level;


means for performing hit ratio simulations on the first and further virtual cache tables by executing a selected number of stored instructions to determine which virtual cache table obtains a highest hit ratio based upon the stored instructions; and

means for adjusting the working segmentation level to equal the segmentation level of the virtual cache table having the greater/greatest hit ratio value.

9. An adaptive segment control assembly as claimed in claim 8, wherein the means for performing hit ratio simulations determines a hit ratio for each virtual cache table based upon the stored instructions; and

further comprising means connected to the virtual cache tables for determining which virtual cache table obtains a higher/highest hit ratio.

10. An adaptive segment control assembly as claimed in claim 9, including a look-up table having all selectable segmentation levels, from which said test segmentation level(s) is/are selected.





11. An adaptive segment control assembly as claimed in any one of claims 8 to 10, further comprising means (10) for flushing said instruction queue, after the number of stored instructions equals a queue limit, and storing further instructions after the number of instructions received subsequently from the host computer is equal to a wait count.



A B S T R A C T**METHODS AND APPARATUS FOR CONTROLLING
A MULTI-SEGMENT CACHE MEMORY**

The present invention relates to the control of a cache memory (13) in a storage system (2). According to the invention, the cache memory is segmented to operate at a selected working segmentation level, and at least two virtual cache tables (52, 54) are segmented to operate, respectively, at the working segmentation level and at least one different segmentation level. Instructions received from a host computer (4) are stored in an instruction queue (50). Then, while the storage system is in an idle state, hit ratio simulations are performed on the virtual cache tables by executing a selected number of instructions stored in the instruction queue. The working segmentation level is subsequently adjusted to equal the segmentation level of the virtual cache table having the highest hit ratio.

Figure 2

