

Description

The present invention relates I/O caching in computer systems.

Computer systems generally have several different storage layers defining a memory hierarchy. These layers typically include registers in a central processing unit, main system memory, and mass storage, such as disks. The performance of these different levels is quite different; the processor registers are much faster than system memory, which in turn is significantly faster than access to mass storage.

In order to improve system performance, a cache, consisting of a small, fast memory, is often used between the central processor and system memory. The cache takes advantage of the locality of reference phenomenon within computer programs to store data which is likely to be reused into the fast memory.

A similar bottleneck occurs between system memory and the much slower mass storage devices and other input/output devices which need to transfer data to and from the system memory. Any system design which makes data transfer between input/output devices, which term as used herein includes mass storage devices, more efficient will have a positive impact on overall system performance. However, traditional caching technology such as used between a central processor and system memory does not work well with input/output devices.

This is due, in large part, to differences in access patterns found with a single central processing unit and a plurality of input/output devices. A central processor is a single unit accessing memory, and makes a large number of relatively random accesses to different locations throughout system memory. Memory accesses by input/output devices have a different pattern of use. Each device usually accesses only a few memory locations, but different devices access different memory locations. In addition, the different devices access memory in unpredictable patterns.

Since some input/output devices access memory relatively infrequently, a cache hit would seldom occur for those devices. This is especially true when some input/output devices, such as disks, read or write a large number of consecutive memory locations. After such a transfer, a cache would tend to be filled with data transferred by that device. Providing a separate cache memory for each input/output device is not generally feasible.

If there are some portions of the main system memory which are accessible by both the input/output devices and the central processor, the central processor may have to access such memory through the input/output cache in order to avoid a cache coherency problem with a normal CPU cache. This tends to interfere with use of an input/output cache by the input/output devices.

Another problem which must be considered by designers of computer systems which support multi-user environments is that of access protection. Concurrent processes must be prevented from accessing memory locations used by other processes, and those used for special purposes by the operating system. A large number of access control schemes are known to those skilled in the art. Mechanisms which implement access control in hardware give superior performance to software based solutions.

In such a multi-user system, input/output from various devices must somehow be accommodated. Input/output accesses must be prevented from occurring to certain areas of memory in order to implement the memory protection schemes required for multi-user systems. Attempting to devise a solution to this problem has been difficult.

The present invention provides a cache for use in a computer system, the cache being interposed between a system memory and a plurality of input/output devices and comprising: a buffer memory coupled to the input/output devices and to said system memory, said buffer memory being addressable as a plurality of lines; cache control means for storing the addresses of the data contained in each line of said buffer memory; means for receiving a data request from one of said input/output devices, said data request including the identity of the input/output device and the address of the requested data; means for comparing said data request with the addresses in said cache control means to determine whether or nor said data is currently contained within said buffer memory; means for accessing the requested data from said buffer memory if contained therein; and characterised by a directory having at least one entry associated with each input/output device, each directory entry including an identification of a corresponding line in said buffer memory, whereby in response to the identity of the input/output device making the data request the directory selects from the cache control means the address of the data for said corresponding line in buffer memory, the selected address then being used by said comparing means to determine whether or not the requested data is currently contained within said buffer memory; and means for ensuring that an input/output device can only access the line or lines in said system memory corresponding to directory entries associated with that input/output device.

Such an arrangement can provide an input/output interface device to provide hardware memory access protection for each device and an efficient input/output buffer to system memory which performs virtual memory to real memory address translations. It can also provide a system memory to input/output device interface which allows separate input/output devices access to reserved portions of an interface buffer.

In one such arrangement disclosed in greater detail hereinafter, an input/output cache has a plurality of cache lines. Each input/output device has a reserved register containing a directory field. The directory field is used to select one or more cache lines associated with each input/output device. The reserved registers also contain access authority information for each input/output device. This authority is checked against access information contained in a virtual page table before an input/output device is allowed to access a virtual page.

The present invention will be described further by way of example with reference to an embodiment thereof as illustrated in the accompanying drawings, in which

FIGURE 1 is a block diagram of a portion of a computer system;

FIGURE 2 is a block diagram of an input/output cache controller; and

FIGURE 3, in section a - d, is a flowchart describing operation of the cache controller of Figure 2.

Referring to Figure 1, a system memory 10 is connected to an input/output (I/O) bus 12 through an input/output cache controller 14 (IOCC). A system bus 15 connects the system memory 10 and IOCC 14, and contains address, data, and control lines. The IOCC 14 is connected to the I/O bus 12 through address, data, and control lines 16.

One input/output device 17, of potentially plural such, is shown connected to the I/O bus 12. Input/output devices can include, for example, video displays, disk and tape drive controllers, and serial and parallel I/O ports. As will be described below, bus memory 18 which is addressable as system memory is also be attached to the I/O bus 12. All data transfers between devices attached to the I/O bus 12 and the system memory 10 occurs through the IOCC 14.

In order to improve the performance of the system memory 10, it may be interleaved as known in the art, or may be otherwise optimised for high speed block transfers. Data is transferred along signal lines 16 using high speed, block transfer communications, periodically freeing the IOCC 14 to communicate with the devices on the I/O bus 12 as necessary.

The IOCC 14 performs the functions of data caching and virtual memory to real memory address translation for the I/O bus 12. In addition, the IOCC 14 performs access authority checking on all I/O requests. Data writes are not written directly through to the system memory 10. Instead, they are retained in the IOCC 14 until such time as the written data is removed from the IOCC 14.

Referring to Figure 2, a detailed diagram of the IOCC 14 is shown. I/O data buffer 20 is connected to a cache buffer 22. This cache buffer 22 is preferably organised as 16 lines of 64 bytes each. The cache buffer 22 is in turn connected to the system memory 10. Data is transferred between the cache buffer 22

and the system memory 10 using a high speed, block data transfer. The precise nature of this block transfer does not form part of the present invention.

The system memory 10 is organised into pages, and further into lines within pages. Although the sizes of the pages and lines do not form a part of the present invention, the page size is taken to be 4K bytes, and the line size, 64 bytes. There are therefore 64 lines per page. Transfers to and from the system memory 10 are done in minimum sized units of one line.

Data is transferred between the I/O data buffer 20 and the cache buffer 22 according to memory locations addressed by 4 bit address line 24 and 6 bit address line 26. The value on line 24 selects one of the 16 lines within the cache buffer 22, while the value on line 26 selects one of the 64 bytes within the selected line.

Addresses from the I/O bus 12 are stored in I/O address buffer 28. The six Least Significant Bits (LSBs) of the I/O address buffer 28 are used as the byte address value for the cache buffer 22 supplied by line 26. The next six more significant bits of the I/O address buffer 28 are connected through signal lines 30 to a comparator 32, and through lines 34 to the system memory 10. Within the system memory 10, the signals on lines 34 are used to select one line out of 64 within a page.

The 20 Most Significant Bits (MSBs) of the I/O address buffer 28 are connected through lines 36 to a comparator 38 and a Translation Control Word (TCW) table 40. The TCW table 40 is also known in the art as a page table. It is a fast, dedicated memory forming part of the IOCC 14, and does not reside in the address space of the system.

Should a large enough dedicated memory for storing the entire TCW table 40 not be able to be incorporated, the dedicated memory can be used as a cache for the table, with the table actually stored in the main address space of the system, but such an arrangement is generally undesirable from a performance viewpoint. However, in many systems, the I/O devices will have access to only a small part of the system address space. Only that part of the address space to which the I/O devices have access need be represented in the TCW table 40. In such systems, the dedicated memory for the TCW table 40 need not be extremely large, and it should be possible to incorporate one which can hold the entire TCW table 40 needed for the I/O devices. The organisation of the TCW memory 40 is described below.

A file of cache control registers 42 contains 16 separate registers. One register of the file 42 is active at any given time, with the active register being selected by a 4 bit signal on lines 44. The 20 MSBs of the currently selected cache control register 42 are connected to the comparator 38 through signal lines 46. The next six less significant bits of the selected

cache control register 42 are connected to comparator 32 through signal lines 48. The three LSBs are labelled D, C, and I. These bits are a dirty bit (D-bit), current bit (C-bit), and invalid bit (I-bit). Use of these bits is described in connection with Figure 3.

A file 50 of 16 registers defines a TCW cache, with one active register being selected by the value on signal lines 52. The value on lines 52 is always the same as the value on lines 44. The 20 MSBs of the selected TCW cache register 50 are connected to the system memory 10 through address lines 54. This 20 bit address is used to address one real page within the system memory 10.

The next less significant bits of the TCW cache 50 define a DIR field which is 4 bits wide. The purpose of the DIR field the function of the DIR signal will be described below. The DIR signal is connected to a select circuit 56 through signal lines 58. A 4-bit signal on lines 60 is also connected to the select circuit 56, and contains the same value as found on signal lines 52.

The 8 LSBs of the TCW cache 50 are divided into two 4-bit fields. The first four LSBs are control bits having functions which will be described below. The next four bits define a KEY field, which is used to determine access authorisation for each page in the system memory 10. The KEY bits and the lower two control bits are connected through signal lines 62 and 64 respectively to a DECODE/SELECT circuit 66.

A file of channel status registers 68 contains 16 registers, one of which is active depending on the value of a four bit signal on lines 70. The signal lines 70 are driven by a DEVICE SELECT signal as will be described below.

The 16 MSBs of the channel status registers 68 are connected to the DECODE/SELECT circuitry 66 through signal line 72. The four LSBs define a DIR field, and are communicated from the active channel status register 68 on line 74, which is connected to the signal lines 44, 52, and 60 as described above.

The TCW table 40 is a memory containing, at each location, an entry having an identical format to that of the TCW cache 50. Thus, each entry in the TCW table 40 contains a real page number, a DIR entry, a KEY, and four control bits. When an entry is read from the TCW table 40, it is communicated to a holding register 76 along signal lines 78. Four bits of an entry in the TCW table 40, corresponding to the DIR field, are communicated directly to the DIR field of the active channel status register 68 on signal lines 80. When an entry from the TCW table 40 is loaded into a TCW cache register 50, it is communicated from the holding register 76 along signal lines 82.

Each device attached to the I/O bus 12 is represented by one of the 16 channel status registers 68. The DEVICE SELECT signal connected to line 70 uniquely identifies one of the devices attached to the I/O bus 12. This signal indicates which device is performing an I/O operation. The device select signal

may be generated in any desired manner. When the IOCC 14 is used with a computer system using an IBM microchannel bus, the DEVICE SELECT signal is preferably the same as the bus arbitration signal. On the microchannel bus, each device has a unique arbitration level, and this identification can be used as the DEVICE SELECT signal.

When one of the channel status registers 68 is selected by the signal on lines 70, the four LSBs are used as a DIR (directory) signal to select a cache control register 42, a TCW cache register 50, and one line of the cache buffer 22. Except for one situation described below, the select circuitry 56 causes the signal on lines 60 to be connected to lines 24 in order to select a line of the cache buffer 22. The selected cache control register 42, TCW cache register 50, and line in the cache buffer 22 always correspond to describe the virtual page number, line number, and real page number of the selected line within the cache buffer 22. The virtual page number of the selected line is stored in the 20 MSBs of the selected cache control register, while the system memory line number corresponding to the line number of the selected line in the cache buffer 22 is stored in the six LSBs of the cache control register 42. The corresponding real page number is stored in the 20 MSBs of the TCW cache register 50.

The virtual page number and line number from the cache control register 42 are compared in comparators 38 and 32, respectively, to the virtual page number and line number in the I/O address buffer 28. If both of these comparisons are a match, the signals TCW VALID and CACHE LINE VALID indicate this fact. When these signals indicate a match, the requested virtual address in the I/O address buffer 28 currently resides in the cache buffer 22 at the line number indicated on lines 24.

To provide an access protection mechanism, each page is assigned a subclassification, also referred to as a subject. Processes and devices which are authorised to have access to a particular subject are allowed access to a virtual page containing that subject. The 16 MSBs of the channel status register define, for each device attached to the I/O bus 12, those subjects which that device is authorised to access.

This authorisation is implemented as a bit map defining an authorisation mask. Of 16 possible subjects, those to which a device is authorised access will have a 1 in the corresponding bit position, and a 0 resides in those bit positions corresponding to subjects for which access is not authorised.

The KEY field of the TCW cache 50 defines the subject for the corresponding page, and is used to select a bit position from the authorisation mask signal present on line 72. For example, if the signal on line 62 is F₁₆, the left most bit of the activated channel status register 68 must be a 1 for the current device to be authorised access to the current page. Each page

has exactly one associated access authority classification, so that one page is the minimum granularity for access classification mapping.

In addition to subject matter access, read and write access is indicated by the 2 LSBs of the TCW cache register 50. These bits also indicate whether there is a page fault (this page is not in system memory 10, having been paged out), or whether the accessed memory page is found in a bus memory instead of system memory 10. These signals are also connected to the DECODE/SELECT circuit 66, as are read/write signals (not shown) from the I/O bus 12. These signals are used to ensure that a device is not, for example, attempting to perform a write operation into a read-only virtual page. The ACCESS OK signal out of the DECODE/SELECT circuit 66 determines whether the attempted access is valid. The meanings of the 2 LSBs in the TCW cache 50 are described in Table I.

Table I - 2 LSB control bits

00 - Bus Memory

01 - Page Fault

10 - Read Only

11 - Read/Write

As to bus memory 18 connected to the I/O bus 12, portions of its bus memory address space can be mapped to a part of the system address space. Thus, to devices on the I/O bus 12, and to the system central processor, bus memory appears the same as system memory 10. However, the cache buffer 22 is not used to cache the bus memory. The differences in handling bus memory are shown in Figure 3. Bus memory can be used to attach I/O devices with a large memory requirement, such as video display controllers.

The remaining two bits of the control field of the selected TCW cache register 50 are labelled R and C, meaning Referenced and Changed, and are preferably used by a virtual memory manager to determine which virtual pages should be swapped to disk when a page fault occurs. Use of the RC bits for memory management is known in the art.

The design of the cache controller of Figure 2 is such that it can be accessed by a system central processor as well as the I/O devices attached to the I/O bus 12. The flow of control described in Figure 3 is basically the same regardless of whether the IOCC 14 is being used by a central processor or an I/O device. The first few steps of the control description are used

only when a system central processor is accessing memory through the IOCC 14. The remaining steps are used by both the central processor and I/O devices.

The system central processor, presumed but not shown, which is performing a load or store operation through the IOCC 14 can be operating in one of two modes. If the central processor is operating in a supervisor mode, it always has access to all memory locations, and access checking is not performed. If the central processor is operating in non-supervisory mode, the access checks are made as described in Figure 3. Not all accesses to devices attached to the I/O bus 12 are routed through the IOCC 14. Only those memory accesses which are to memory locations in the normal system memory address space invoke operation of the IOCC 14. Accesses to a separate I/O address space, such as supported by Intel's 80x86 family of processors, bypass the IOCC 14 and are routed directly to the I/O bus 12. For accesses of this type, a bypass flag is set to indicate that the IOCC 14 is to be bypassed.

Detailed operation of the IOCC 14 will now be described in connection with Figure 3. The central processor appears to the IOCC 14 as merely another I/O device, having its own reserved channel status register 68 as described in connection with Figure 2.

The devices connected to the I/O bus 12, these logically including the system central processor if it is attempting to perform a load/store operation through the IOCC 14, arbitrate to determine which device will have access to the IOCC 14. As described above, the winner of such arbitration will have a unique arbitration level identifier, which is used as the signal DEVICE SELECT.

At several points within the control flow, error checking will be performed. If an error is detected, control will jump to step 94, which causes the error to be logged. Preferably, a 4-bit error indicator and the page and line from the buffer 28 on which the error occurred are placed into the channel status register 68 for the selected device. The register 68 is reloaded with the authorisation mask and DIR field after the error is handled. The error is then reported to the device 96.

The signal DEVICE SELECT is used to select a device register -step 106-, and the appropriate cache line in the cache buffer 22, and corresponding registers in the cache control register file 42 and the TCW cache register file 50, are selected using the DIR signal -step 108-. Also at this time, the device winning the arbitration battle for the I/O bus 12 asserts the desired address -step 110-.

A page out of range test -step 112- is made to determine whether the asserted address is within that portion of the virtual memory space which can be accessed by I/O devices. If the page is out of range, control passes to step 94 and an error is logged and re-

ported. If the page is within the appropriate range, a test is made to see if there is a page hit -step 114- between the asserted address and the cache buffer 22. This determination is made in comparator 38, and is indicated by the signal TCW VALID.

If the previously accessed page stored in the cache control register 42 matches the page within the I/O buffer 28, there is a page hit and control passes to step 116. A test is made for a page fault -step 116-, which is indicated by the 2 LSBs of the selected TCW cache register 50 using the bit pattern shown in Table I. If there is a page fault, an error is logged and reported to the processor.

If there is no page fault, indicating that the requested virtual page currently resides in system memory 10 or in a bus memory mapped to the system address space, a check is made to see whether the requesting device is authorised to access that page -step 118-. In addition to the subject matter check made by comparing the KEY field of the TCW cache register 50 to the authorisation mask field of the channel status register 68, a check is made to ensure that a device is not attempting to write into a read only page. If the signal ACCESS OK indicates that the attempted access is not valid, an error is logged and reported to the processor.

If the device does have valid access to the requested page, a check is made to see whether the selected cache line in the buffer 22 is invalid. This check is made by testing the I-bit of the selected cache control register 42. When the I-bit is set, data exists in the selected cache line within the buffer 22 which might not be accessible to the device making access. This is a security feature which prevents processes and devices from accessing data from each other's memory space. The I-bit is set when an I/O transfer using the cache line is completed, and the data is not erased from the cache line within the buffer 22. Whenever a cache line is written from the buffer 22 to system memory, that cache line is filled with zeros, effectively deleting the data which was previously stored therein. However, when data is being read from system memory 10 to an I/O device, the last line of data loaded into the cache buffer 22 remains untouched. When an I/O transfer completes, a system supervisor writes a 1 into the corresponding I-bit unless that cache line was written to system memory 10 and zeroed out.

If the cache line invalid test 120 gives a yes result, the test of step 122 is skipped. Step 122 is a test to determine whether the cache line was a hit. If it is, control drops to step 124. If there is not a cache line hit, the page is correct but the wrong line from that page currently resides in the cache buffer 22. If that cache line is not dirty, as indicated by the D-bit in the cache control register 42, it does not need to be saved. If the cache line is dirty 126, the cache line is written to memory 128 and the cache line zeroed out

as described above. If an error occurs in the transfer -step 130-, the error is logged and reported to the processor. If no error occurs, the D-bit and C-bit are both reset to zero -step 132-.

Next, a test is made to see whether this is an operation to or from bus memory -step 124-. This test is made on the 2 LSBs of the selected TCW cache register 50 for the pattern defined in Table I. If the operation is a bus memory operation, control skips directly ahead to the data transfer and bypasses the remaining cache operations.

If the operation is to system memory, the result of test 124 will be no, and another test is made to see if the page was previously referenced -step 126-. This test is made of the R-bit in the selected TCW cache register 50, which, if set, indicates that the selected page was previously referenced. If the page was referenced, a test is made to see if the current operation is a read operation -step 128-. If it is not a read operation, a test is made to see if this page has previously been altered -step 130-. If it has, the RC bits are unchanged, as is the case if the current operation is a read operation. If the page has not previously been altered, and this is a write operation, the RC bits in the TCW table 40, and in the selected TCW cache register 50, are updated -step 132-. They are also updated if the test of step 126 indicates that this page has not previously been referenced.

Next, another test is made to see if the cache line is invalid -step 134-. As described above, this consists of a test of the I-bit. If the cache line is invalid, the next few steps are skipped, which forces a new line to be read from system memory 10 into the cache buffer -step 136-. This destroys the data previously residing on that cache line. Reading a new line into the cache buffer 22 -step 136- causes the corresponding I-bit to be reset to zero -step 137-.

If the cache line was not invalid at step 134, a check is again made to see if the current operation is a read operation -step 136-. If so, a check is made to see whether the C-bit is set to 1 -step 138-. If the current operation is a write operation, the other possibility, the result of step 136 is no, and a check is made to see whether the requested location is aligned with a cache address -step 140-. This means that the requested location is aligned with a line boundary, which is the case when the 6 LSBs of the I/O address buffer 28 are all 0. If the requested location is not aligned, control passes to step 138. If it is aligned, and if the central processor is the device currently using the I/O bus -step 141-, a test is made to see if a current transfer is less than the 64 byte cache buffer size -step 142-. If so, control also passes to step 138. If a transfer larger than 64 bytes is being made, the result of step 142 is no, and control skips around step 136. The effect of reaching this point in the flow of control is that a large write operation, one which will fill every byte of the selected line of the cache buffer 22, is be-

ing made. Under these circumstances, it is a waste of effort to read in a line from system memory 10 since every single byte of that line is about to be overwritten by the current write operation. Thus, the read from system memory -step 138- is skipped.

If the central processor was not using the bus -step 141-, a test is made to see whether the I/O device is operating in sequential mode -step 143-. Sequential mode is assigned to devices which read and write large blocks of data, such as disk drives. Sequential mode for a device can be indicated by a bit (not shown) in the channel status register 68. If a device is writing in sequential mode, it is also unnecessary to load the cache line. Thus, if the result of step 143 is yes, control skips step 136. Otherwise, control passes to step 138.

If the cache buffer 22 load of step 136 is made, the I-bit is set to zero -step 137-, and a check is made to see if there were any errors in the transfer -step 144-. If so, the error is logged and reported to the processor. If not, the C-bit and D-bit in the cache control register 42 are updated -step 146-. The C-bit is set to 1. The D-bit is set to the logical OR of the current value of the D-bit and a signal indicating whether this is a write operation. That is, if the cache line was already dirty, or is now being written into, the dirty bit is set. At this time, the IOCC 14 is properly set up, and the data transfer can begin -step 148-.

A single word of data is transferred -step 150-. This word can be a byte, two bytes, four bytes, or some other value as established by the hardware configuration of the system. In any event, the transfer of step 150 is a single transfer accomplished in one bus cycle. If there is an error in the transfer -step 152-, it is logged and reported to the processor as described above. If the central processor is currently using the I/O bus -step 153-, a test is made to see whether the transfer is complete -step 154-. If so, nothing further is done.

If the transfer is not complete, a test is made to see whether this operation is a bus memory operation -step 156-. As described above, a bus memory operation is indicated by the 2 LSBs of the TCW cache register 50. If this is a system memory operation, a test is made to see whether the next address is a cache line hit -step 158-. If so, control returns to step 150 to transfer the next word. If not, a cache line boundary has been crossed, and control drops to step 160. Step 160 is a test to see whether the current I/O operation has been pre-empted by another device. If not, control returns to step 114, and the entire cache validation process is repeated. If the current operation is pre-empted in step 160, this operation is completed for now. The central processor must then arbitrate for access to finish its transfer.

If the test of step 156 indicates that the current operation is a bus memory operation, a test is made to see whether the current transfer has been pre-empt-

ed -step 166-. If not, and if the page of the next address matches the page of the current address indicating a page hit -step 168-, control returns to step 150 to transfer the next word. If step 168 does not result in a page hit, control returns to step 114 as described above.

The net effect of the control flow of steps 156 through 168 is to allow operations to bus memory to be pre-empted at any time, while those to cached system memory locations can be pre-empted only when a line boundary is crossed.

If the result of step 153 was no and if the transfer is not complete -step 162-, a check is made for a cache line hit -step 164-. If there is a hit, control returns to step 150, and the transfer continues. If not, control returns to step 114. In the disclosed arrangement, I/O devices cannot be pre-empted so that they must give up the I/O bus voluntarily.

Returning to Figure 3a, the operations to be undertaken if a page miss is detected in step 114 will now be described. If step 114 determines that there was not a page hit, the virtual page currently being accessed by the device is different from the virtual page previously accessed by that device. If this is the case, a next TCW entry is fetched -step 170- from the TCW table 40 and placed into holding register 76. If there is a parity error -step 171-, it is logged and reported. The next TCW entry is that entry addressed by the 20 MSBs of the I/O buffer 28. The 4 LSBs of the channel status register 68 are updated -step 172- using the DIR signals present on line 80. This signal immediately becomes available on line 74, causing a new line in the cache buffer 22, and its corresponding registers in the cache control register file 42 and the TCW cache register file 50 to be selected -step 174-. Comparator 38 then compares the newly selected previous page available on line 46 with the page within the I/O address buffer 28. If there is now a TCW hit -step 176-, the currently requested page was accessed by the same device at a time prior to the immediately preceding access, and that page is still available in the cache control and TCW cache registers 42, 50. Whenever a new TCW entry is loaded into the TCW cache 50, the cache line addressed by the old entry must first be flushed to memory 10 if it was dirty. This is performed in circuit 56 using the D-bit of cache control register 42.

If the second compare -step 176- also results in a page miss, it is necessary to load a new entry from the TCW table 40. First, a check is made to see if the cache line is invalid -step 178- using the I-bit. If not, a check is made to see if the currently selected cache line is dirty -step 180-. This check is made of the D-bit within the cache control register 42. If so, that cache line is written to memory -step 182- as described above. A test is made to see if there was an error in the transfer -step 184-, and if so it is reported and logged. If there is no error, the D-bit and C-bit are both

reset to 0 -step 186-.

Whether or not the cache line was dirty, the next step is to load the new TCW entry into the TCW cache -step 188-. This entry currently resides in the holding register 76, and can be loaded into the TCW cache 50 without the necessity of again accessing the TCW table 40. The 26 MSBs of the current I/O address are then loaded into the cache control register -step 190-, and control passes to step 116.

In most cases, each device on the I/O bus will be allocated a single line in the cache buffer 22. For these devices, the DIR fields in the TCW table 40 should all be set to the same value as the value of the DEVICE SELECT which corresponds to each device. In some cases, it will be desired for some devices to have more than one line in the cache buffer 22. For these devices, the DIR fields in the entries in the TCW table 40 are set to point to the appropriate lines in the cache buffer 22. As described above, when there is a TCW miss, the next entry in the TCW table 40 determines which line of the cache buffer 22 is used next. By setting the DIR fields of these entries to point to different lines of the cache buffer 22 which are reserved for a single device, that device can use more than one line in the cache buffer 22.

The DIR fields in the TCW table 40 are allocated by the user, which will typically be an operating system program. Different systems, having different usage statistics, can be reconfigured through the operating system by adjusting the values in the DIR fields of the TCW table 40 to make most efficient use of the cache buffer 22. Manipulation of DIR fields in a cache, although the use thereof is different that of the present invention, is described in U.S. Patent No. 4,719,568, by Carrubba et al, titled HIERARCHICAL MEMORY SYSTEM INCLUDING SEPARATE CACHE MEMORIES FOR STORING DATA AND INSTRUCTIONS.

When a device is opened, often at system power up, it is allocated some memory space. At that time, the line or lines in cache buffer 22 which it will use are determined. The DIR entries in the TCW table 40 corresponding to the allocated memory space are set to point to the selected line or lines in cache buffer 20. When the device is deallocated, the lines in the cache buffer 20 are freed for use by other device.

As will be appreciated by those skilled in the art, the input/output cache controller described above provides an efficient interface between I/O devices attached to an I/O bus 12 and system memory 10. The use of a DIR field allows more than one 64 bit line within the cache buffer 22 to be allocated to a single input/output device. The described buffer 22 and controller 14 appear to the user as a 16-way set associative cache.

The described input/output cache controller also provides virtual to real page translation, and automatic hardware access authorisation checking. Access is allocated on a virtual page basis, with every access

by an input/output device checked to see whether that device is authorised to perform the operation. Since the authorisation checking is performed in hardware, it is fast and cannot be bypassed by software without breaking a trusted error handler within the operating system.

It is noted that the description hereof is common to that of co-pending application EP-A-0 377 971, though claiming a different invention.

Claims

1. A cache (14) for use in a computer system, the cache being interposed between a system memory (10) and a plurality of input/output devices (17) and comprising:

a buffer memory (22) coupled to the input/output devices and to said system memory, said buffer memory being addressable as a plurality of lines;

cache control means (42) for storing the addresses of the data contained in each line of said buffer memory;

means (28, 70) for receiving a data request from one of said input/output devices, said data request including the identity of the input/output device and the address of the requested data;

means (32, 38) for comparing said data request with the addresses in said cache control means to determine whether or not said data is currently contained within said buffer memory;

means (24, 26) for accessing the requested data from said buffer memory if contained therein;

and characterised by a directory (70, 68) having at least one entry associated with each input/output device, each directory entry including an identification of a corresponding line in said buffer memory, whereby in response to the identity of the input/output device making the data request the directory selects from the cache control means the address of the data for said corresponding line in buffer memory, the selected address then being used by said comparing means to determine whether or not the requested data is currently contained within said buffer memory;

and means (56) for ensuring that an input/output device can only access the line or lines in said system memory corresponding to directory entries associated with that input/output device.

2. A cache as claimed in Claim 1, wherein at least one input/output device has more than one directory entry associated therewith.

3. A cache as claimed in either preceding claim,

reset to 0 -step 186-.

Whether or not the cache line was dirty, the next step is to load the new TCW entry into the TCW cache -step 188-. This entry currently resides in the holding register 76, and can be loaded into the TCW cache 50 without the necessity of again accessing the TCW table 40. The 26 MSBs of the current I/O address are then loaded into the cache control register -step 190-, and control passes to step 116.

In most cases, each device on the I/O bus will be allocated a single line in the cache buffer 22. For these devices, the DIR fields in the TCW table 40 should all be set to the same value as the value of the DEVICE SELECT which corresponds to each device. In some cases, it will be desired for some devices to have more than one line in the cache buffer 22. For these devices, the DIR fields in the entries in the TCW table 40 are set to point to the appropriate lines in the cache buffer 22. As described above, when there is a TCW miss, the next entry in the TCW table 40 determines which line of the cache buffer 22 is used next. By setting the DIR fields of these entries to point to different lines of the cache buffer 22 which are reserved for a single device, that device can use more than one line in the cache buffer 22.

The DIR fields in the TCW table 40 are allocated by the user, which will typically be an operating system program. Different systems, having different usage statistics, can be reconfigured through the operating system by adjusting the values in the DIR fields of the TCW table 40 to make most efficient use of the cache buffer 22. Manipulation of DIR fields in a cache, although the use thereof is different that of the present invention, is described in U.S. Patent No. 4,719,568, by Carrubba et al, titled HIERARCHICAL MEMORY SYSTEM INCLUDING SEPARATE CACHE MEMORIES FOR STORING DATA AND INSTRUCTIONS.

When a device is opened, often at system power up, it is allocated some memory space. At that time, the line or lines in cache buffer 22 which it will use are determined. The DIR entries in the TCW table 40 corresponding to the allocated memory space are set to point to the selected line or lines in cache buffer 20. When the device is deallocated, the lines in the cache buffer 20 are freed for use by other device.

As will be appreciated by those skilled in the art, the input/output cache controller described above provides an efficient interface between I/O devices attached to an I/O bus 12 and system memory 10. The use of a DIR field allows more than one 64 bit line within the cache buffer 22 to be allocated to a single input/output device. The described buffer 22 and controller 14 appear to the user as a 16-way set associative cache.

The described input/output cache controller also provides virtual to real page translation, and automatic hardware access authorisation checking. Access is allocated on a virtual page basis, with every access

by an input/output device checked to see whether that device is authorised to perform the operation. Since the authorisation checking is performed in hardware, it is fast and cannot be bypassed by software without breaking a trusted error handler within the operating system.

It is noted that the description hereof is common to that of co-pending application EP-A-0 377 971, though claiming a different invention.

Claims

1. A cache (14) for use in a computer system, the cache being interposed between a system memory (10) and a plurality of input/output devices (17) and comprising:

a buffer memory (22) coupled to the input/output devices and to said system memory, said buffer memory being addressable as a plurality of lines;

cache control means (42) for storing the addresses of the data contained in each line of said buffer memory;

means (28, 70) for receiving a data request from one of said input/output devices, said data request including the identity of the input/output device and the address of the requested data;

means (32, 38) for comparing said data request with the addresses in said cache control means to determine whether or not said data is currently contained within said buffer memory;

means (24, 26) for accessing the requested data from said buffer memory if contained therein;

and characterised by a directory (70, 68) having at least one entry associated with each input/output device, each directory entry including an identification of a corresponding line in said buffer memory, whereby in response to the identity of the input/output device making the data request the directory selects from the cache control means the address of the data for said corresponding line in buffer memory, the selected address then being used by said comparing means to determine whether or not the requested data is currently contained within said buffer memory;

and means (56) for ensuring that an input/output device can only access the line or lines in said system memory corresponding to directory entries associated with that input/output device.

2. A cache as claimed in Claim 1, wherein at least one input/output device has more than one directory entry associated therewith.

3. A cache as claimed in either preceding claim,



wherein at least two input/output devices have a single directory entry associated therewith.

4. A cache as claimed in any preceding Claim, further comprising for each line in said buffer memory, an indicator of that line's corresponding locations in said system memory and, should it exist, an indicator of another buffer memory line accessible by the input/output device currently accessing such line. 5
10
5. A cache as claimed in any preceding Claim, further comprising:
for each line in said buffer memory, a protection indicator of an access protection class thereof; 15
for each input/output device, an authorization indicator of access protection classes to which that device is authorised access; and
an access protection circuit (66) coupled to said protection indicators and to said authorization indicators for limiting input/output device access to those lines having an access protection class to which an input/output device is authorised access. 20
25
6. A cache as claimed in Claim 5, wherein the system is divided into regions having different access protection classes, and wherein each line in said buffer memory is assigned an access protection class which is the same as that assigned to the region of system memory cached in that line. 30
7. A cache as claimed in any preceding Claim, wherein each input device has a unique identifier and a register file is provided which contains a register corresponding to each input/output device containing a current directory for its corresponding input/output device. 35
40
8. A cache as claimed in Claim 7, wherein said register file further includes the authorization indicator for each device. 45

Patentansprüche

1. Ein Cachespeicher (14) zur Benutzung in einem Computer-system, in welchem der Cachespeicher zwischen einem Systemspeicher (10) und einer Vielzahl von Ein-/Ausgabeeinheiten (17) liegt, wobei der Cachespeicher enthält:
einen Pufferspeicher (22), der mit den Ein-/Ausgabeeinheiten und diesem Systemspeicher verbunden ist, wobei der Pufferspeicher in einer Vielzahl von Zeilen adressierbar ist; 50
Cachespeicher-Steuermittel (42), um die Adres- 55

sen von Daten, die in jeder Zeile dieses Pufferspeichers enthalten sind, zu speichern;
Mittel (28, 70), um eine Datenanforderung von einer der Ein-/Ausgabeeinheiten zu empfangen, wobei die Datenanforderung die Identität der Ein-/Ausgabeeinheit und die Adresse der angeforderten Daten enthält;

Mittel (32, 38), um diese Datenanforderung mit den Adressen in den Cache-Speichersteuermitteln zu vergleichen, um zu bestimmen, ob diese Daten aktuell innerhalb des Pufferspeichers enthalten sind oder nicht;

Mittel (24, 26), um auf die angeforderten Daten von diesem Pufferspeicher aus zuzugreifen, sofern diese darin enthalten sind;

und dadurch gekennzeichnet ist, daß ein Verzeichnis (70, 68) wenigstens eine Eingabe hat, die mit jeder Ein-/Ausgabeeinheit verbunden ist, wobei jedes Verzeichnis eine Identifikation von einer entsprechenden Zeile in dem Pufferspeicher enthält, wobei als Reaktion auf die Identität der Ein-/Ausgabeeinheit, welche die Daten anfordert, das Verzeichnis aus den Cachespeicher-Steuermitteln die Adresse der Daten für die entsprechende Zeile im Pufferspeicher auswählt, wobei die ausgewählte Adresse anschließend von den Vergleichsmitteln benutzt wird, um zu bestimmen, ob die angeforderten Daten aktuell in dem Pufferspeicher enthalten sind oder nicht;

und Mittel (56), um sicherzustellen, daß eine Ein-/Ausgabeeinheit nur auf die Zeile oder die Zeilen in diesem Systemspeicher zugreifen kann, welche den Verzeichniseingaben, die zu der Ein-/Ausgabeeinheit gehören, entspricht bzw. entsprechen.

2. Ein Cachespeicher wie in Anspruch 1 angemeldet, in dem wenigstens eine Ein-/Ausgabeeinheit mehr als eine Verzeichniseingabe hat, die damit verbunden ist.
3. Ein Cachespeicher wie in irgendeinem der vorhergehenden Ansprüche angemeldet, in dem wenigstens zwei Ein-/Ausgabeeinheiten eine einzelne Verzeichniseingabe haben, die damit verbunden ist.
4. Ein Cachespeicher wie in irgendeinem vorhergehenden Anspruch angemeldet, der außerdem für jede Zeile in dem Pufferspeicher ein Anzeiger von der Zeile enthält, die den Speicherplätzen in diesem Systemspeicher entspricht, und sofern vorhanden, ein Anzeiger von einer anderen Pufferspeicherzeile, auf die von der Ein-/Ausgabeeinheit zugegriffen werden kann, die aktuell auf diese Zeile zugreift.
5. Ein Cachespeicher wie in irgendeinem vorherge-



henden Anspruch angemeldet, der außerdem enthält:
 für jede Zeile in dem Pufferspeicher einen Schutzanzeiger einer hierfür bestimmten Schutzklasse;
 für jede Ein-/Ausgabeeinheit einen Berechtigungsanzeiger der Berechtigungsklassen, für welche die Einheit zugriffsberechtigt ist; und
 einen Zugriffsschutz-Schaltkreis (66), der mit den Schutzanzeigern und den Berechtigungsanzeigern gekoppelt ist, um den Zugriff der Ein-/Ausgabeeinheit auf jene Zeilen zu beschränken, die eine Zugriffsschutzklasse haben, für die eine Ein-/Ausgabeeinheit zugriffsberechtigt ist.

6. Ein Cachespeicher wie in Anspruch 5 angemeldet, wobei das System in Bereiche aufgeteilt ist, die unterschiedliche Zugriffsschutzklassen haben, und wobei jede Zeile in dem Pufferspeicher einer Zugriffsschutzklasse zugeordnet wird, welche die gleiche ist wie diejenige, die dem Bereich des Systemspeichers zugewiesen wird, der in dieser Zeile gepuffert wird.
7. Ein Cachespeicher wie in irgendeinem vorhergehenden Anspruch angemeldet, wobei jede Eingabeeinheit eine einzige Identifizierung hat und eine Registerdatei geliefert wird, welche ein Register enthält, das jeder Ein-/Ausgabeeinheit entspricht, die ein aktuelles Verzeichnis für die entsprechende Ein-/Ausgabeeinheit hat.
8. Ein Cachespeicher wie in Anspruch 7 angemeldet, wobei die Registerdatei außerdem den Berechtigungsanzeiger für jede Einheit enthält.

Revendications

1. Antémémoire (14) pour utilisation dans un système ordinateur, l'antémémoire étant interposée entre une mémoire de système (10) et une pluralité de dispositifs d'entrée/sortie (17) et comprenant :
 une mémoire tampon (22) connectée aux dispositifs d'entrée/sortie et à ladite mémoire de système, ladite mémoire tampon pouvant être adressée comme une pluralité de lignes,
 des moyens de contrôle d'antémémoire (42) pour emmagasiner les adresses des données contenues dans chaque ligne de ladite mémoire tampon,
 des moyens (28, 70) pour recevoir une requête de données d'un desdits dispositifs d'entrée/sortie, ladite requête de données comportant l'identité du dispositif d'entrée/sortie et l'adresse des données demandées,
 des moyens (32, 38) pour comparer ladite


requête de données aux adresses dans lesdits moyens de contrôle d'antémémoire et déterminer si lesdites données sont contenues dans ladite mémoire tampon,

des moyens (24, 26) pour accéder aux données demandées à partir de ladite mémoire tampon si elles y sont contenues,

et caractérisée par un répertoire (70, 68) ayant au moins une entrée associée à chaque dispositif d'entrée/sortie, chaque entrée de répertoire comprenant une identification correspondante dans ladite mémoire tampon, de sorte qu'en réponse à l'identité du dispositif d'entrée/sortie émettant la requête de données, le répertoire sélectionne à partir des moyens de contrôle d'antémémoire l'adresse des données pour ladite ligne correspondante dans la mémoire tampon, l'adresse sélectionnée étant alors utilisée par lesdits moyens de comparaison pour déterminer si les données demandées sont contenues dans ladite mémoire tampon,

et des moyens (56) pour assurer qu'un dispositif d'entrée/sortie peut seulement avoir accès à la ligne ou aux lignes dans ladite mémoire de système correspondant aux entrées du répertoire associées à ce dispositif d'entrée/sortie.

2. Antémémoire telle que revendiquée dans la revendication 1, dans laquelle au moins un dispositif d'entrée/sortie a plus d'une entrée de répertoire qui lui est associée.
3. Antémémoire telle que revendiquée dans une des revendications précédentes, dans laquelle au moins deux dispositifs d'entrée/sortie ont une seule entrée de répertoire associée.
4. Antémémoire telle que revendiquée dans une des revendications précédentes, comprenant en outre pour chaque ligne dans ladite mémoire tampon, un indicateur des emplacements correspondant à cette ligne dans ladite mémoire de système et, s'il existe, un indicateur d'une autre ligne de mémoire tampon accessible par le dispositif d'entrée/sortie ayant accès à cette ligne.
5. Antémémoire telle que revendiquée dans une des revendications précédentes, comprenant en outre
 pour chaque ligne dans ladite mémoire tampon, un indicateur de protection d'une classe de protection d'accès pour la mémoire tampon,
 pour chaque dispositif d'entrée/sortie, un indicateur d'autorisation de classes de protection d'accès pour lesquelles l'accès est autorisé pour ce dispositif, et
 un circuit de protection d'accès (66) connecté auxdits indicateurs de protection et



auxdits indicateurs d'autorisation pour limiter l'accès des dispositifs d'entrée/sortie aux lignes ayant une classe de protection d'accès pour laquelle l'accès est autorisé pour un dispositif d'entrée/sortie.

5

6. Antémémoire telle que revendiquée dans la revendication 5, dans laquelle le système est divisé en régions ayant des classes de protection d'accès différentes, et dans laquelle à chaque ligne dans ladite mémoire tampon est affectée une classe de protection d'accès qui est la même que celle affectée à la région de la mémoire de système dont l'antémémoire correspond à cette ligne.

10

15

7. Antémémoire telle que revendiquée dans une des revendications précédentes, dans laquelle chaque dispositif d'entrée comporte un identificateur unique et un ensemble de registres est fourni pour contenir un registre correspondant à chaque dispositif d'entrée/ sortie contenant un répertoire courant pour son dispositif d'entrée/sortie correspondant.

20

25

8. Antémémoire telle que revendiquée dans la revendication 7, dans laquelle ledit ensemble de registres comprend en outre l'indicateur d'autorisation pour chaque dispositif.

30

35

40

45

50

55

11

