

Description

The invention relates to a computer system having the features of the first part of claim 1 (IBM-TDB, Vol. 23, No. 7B, Dec. 1980, p. 3329 and Vol. 24, No. 7A, Dec. 1981, p. 3128).

Most modern computer systems include a central processing unit (CPU) and a main memory. The speed at which the CPU can decode and execute instructions to process data has for some time exceeded the speed at which instructions and operands can be transferred from main memory to the CPU. In an attempt to reduce the problems caused by this mismatch, many computers include a cache memory or buffer between the CPU and main memory.

Cache memories are small, high-speed buffer memories used to hold temporarily those portions of the contents of main memory which are believed to be currently in use by the CPU. The main purpose of caches is to shorten the time necessary to perform memory accesses, either for data or instruction fetch. Information located in cache memory may be accessed in much less time than that located in main memory. Thus, a CPU with a cache memory needs to spend far less time waiting for instructions and operands to be fetched and/or stored. For such machines the cache memory produces a very substantial increase in execution speed.

A cache is made up of many blocks of one or more words of data, each of which is associated with an address tag that uniquely identifies which block of main memory it is a copy of. Each time the processor makes a memory reference, the cache checks to see if it has a copy of the requested data. If it does, it supplies the data; otherwise, it gets the block from main memory, replacing one of the blocks stored in the cache, then supplies the data to the processor. See, Smith, A. J., Cache Memories, ACM Computing Surveys, 14:3 (Sept. 1982), pp. 473-530.

Optimizing the design of a cache memory generally has four aspects:

- (1) Maximizing the probability of finding a memory reference's target in the cache (the hit ratio),
- (2) minimizing the time to access information that is indeed in the cache (access time),
- (3) minimizing the delay due to a miss, and
- (4) minimizing the overheads of updating main memory and maintaining multicache consistency.

All of these objectives are to be accomplished under suitable cost constraints and in view of the inter-relationship between the parameters.

When the CPU executes instructions that modify the contents of the current address space, those changes must eventually be reflected in main

memory; the cache is only a temporary buffer. There are two general approaches to updating main memory: stores can be transmitted directly to main memory (referred to as write-through or store-through), or stores can initially modify the data stored in the cache, and can later be reflected in main memory (copy-back or write-to). The choice between write-through and copy-back strategies also has implications in the choice of a method for maintaining consistency among the multiple cache memories in a tightly coupled multiprocessor system.

A major disadvantage to the write-through approach is that write-through requires a main memory access on every store. This adds significantly to the relatively slow main memory traffic load which slows the execution rate of the processor and which the cache is intended to minimize. However, when write-through is not used, the problem of cache consistency arises because main memory does not always contain an up-to-date copy of all the information in the system.

Input and output between the main memory and peripheral devices is an additional source of references to the information in main memory which must be harmonized with the operation of cache memories. It is important that an output request stream reference the most current values for the information transferred. Similarly, it is also important that input data be immediately reflected in any and all copies of those lines in memory.

There have been several approaches to solving this problem. One is to direct the I/O stream through the cache itself. This method is limited to single processor systems. Further, it interferes significantly with the processor's use of the cache, both by keeping the cache busy when the processor needs it and by displacing blocks of information currently being used by the processor with the blocks from the I/O stream. Thus it degrades both the cache access time and the hit rate. An alternate approach is to use a write-through policy and broadcast all writes so as to update or invalidate the target line wherever found. Although this method accesses main memory instead of the cache, it suffers from the disadvantages of the write-through strategy discussed above. In addition, this hardware intensive solution is expensive to implement and increases the cache access cycle time by requiring the cache to check for invalidation. This is particularly disadvantageous in multiprocessor systems because every cache memory in the system can be forced to surrender a cycle to invalidation lookup whenever any processor in the system performs a store.

Another alternative is to implement a directory to keep track of the location and status of all copies of each block of data. The directory can be central-

ized in main memory or distributed among the caches, I/O channels and main memory. This system insures that at any time only one processor or I/O channel is capable of modifying any block of data. See, Tang, C.K., Cache Design in the Tightly Coupled Multiprocessor System, AFPIS Proc., N.C.C., vol. 45, pp. 749-53 (1976). The major disadvantage of the directory control system is the complexity and expense of the additional hardware it requires.

Finally, if a processor fails, for instance because of a power interruption, the memory system must assure that the most current copies of information are stored in main memory, so that recovery can be more easily accomplished.

The problem underlying the invention is to provide a system for maintaining the memory integrity and consistency in a computer system having cache memories, to place the burden of maintaining integrity on the software, thus allowing the hardware to remain relatively simple, cheap and fast, and to minimize the impact of the overhead for maintaining memory integrity and consistency of the operation of the cache memories, so that the cache access time and miss ration can be minimized.

These and other objects of the invention are accomplished by a computer system as claimed in claim 1.

By invoking different combinations of these explicit instructions, the operating system software can maintain memory integrity in the event of transfer of data or instructions into or out of memory via the I/O channel, removal or destruction of a page or a segment in a virtual memory system, modification of an instruction in a cache, a processor failure.

Each block of data stored in the caches has two one-bit status flags, a valid bit to indicate whether the block contains up-to-date information, and a dirty bit to indicate whether the data in the block has been stored to by the processor since it entered the cache. The instruction set includes instructions for removing a block with a particular address from the cache and writing it back to memory if necessary, for removing a block without writeback to main memory, for suspending execution of instructions until pending cache control operations are completed, and for efficiently removing and writing back to main memory all "dirty" blocks in the cache in case of a processor failure. The operating system software invokes these instructions in situations which could result in inconsistent or stale data in the cache memories.

Figure 1 is a schematic block diagram of a computer system which incorporates the invention.

Figure 2 is a schematic illustration of a cache memory constructed in accordance with the inven-

tion.

Figure 3 is a schematic illustration of an alternative form of cache memory constructed in accordance with the invention.

A computer system which operates according to the invention is schematically illustrated in Figure 1. The main processor 11, often referred to as the CPU, communicates with main memory 13 and input/output channel 15 via memory bus 17. The main processor includes a processor 19 which fetches, decodes and executes instructions to process data. Data and instructions are stored in main memory 13, transferred to processor 19 when they are requested during the execution of a program or routine and returned to main memory 13 after the program or routine has been completed.

Access to main memory 13 is relatively slow compared with the operation of processor 19. If processor 19 had to wait for main memory access to be completed each time an instruction or data was needed, its execution rate would be reduced significantly. In order to provide access times which more closely match the needs of the processor, cache 21, which may be referred to as a buffer memory, stores a limited number of instructions and data. Since cache 21 is much smaller than main memory 13 it can be economically built to have higher access rates.

The operating system software for the computer, rather than the hardware of the component units, is responsible for maintaining the integrity and consistency of the memory. In order to accomplish this, the operating system invokes explicit control instructions included in the computer's instruction set.

To explain the system of the invention more completely, an understanding of the structure of cache memory 21 is necessary. The entries of the array in cache memory 21 are illustrated in Figure 2. Cache 21 comprises an array of locations labeled with an index 31 which store data 33 and a physical page tag 35 which corresponds to the physical page number of the location of the copy of the data in main memory.

In addition to the data 33 and tags 35 stored in the cache, each block has associated with it two one-bit status flags, "valid" and "dirty". The valid bit 37 is set if and only if that block has valid data, i.e., up-to-date data.

The dirty bit 39 is set if the processor has stored to the address since it has been brought into the cache. Unless cache 21 updates main memory 13 every time processor 19 does a store (write-through), the cache has more up-to-date data for a block than main memory has. Dirty bit 39 serves to indicate that main memory 13 must be updated by writing the data in the block in cache 21 back to main memory 13 when the block is

swapped out of the cache.

Cache 21 can also be divided into two sections, one for data and another for instructions, as illustrated in Figure 3. For many computer architectures, this split cache provides performance advantages. Both the instruction cache 41 and the data cache 51 have structures similar to that of the unified cache described above. Instruction cache 41 has an array of locations labeled with an index 43. Each location stores an entry comprising: a physical page tag 45, an instruction 46 and a valid bit 47. Data cache 51 has an array of locations labelled with an index 53. Each location stores an entry comprising: a physical tag 55, a block of data 56, a valid bit 57 and a dirty bit 58. Although this cache organization provides certain advantages, it also requires additional control instructions. In particular, instructions may be modified and copies may then appear in both sections of the cache. The operating system must therefore flush blocks from the instruction cache 41 and from data cache 51 back to main memory 13 to insure consistency.

The operating system performs the required memory maintenance functions using six instructions: Flush Data Cache, Purge Data Cache, Flush Instruction Cache, Flush Data Cache Entry, Flush Instruction Cache Entry and Synchronize Caches.

The Flush Data Cache (FDC) instruction sets the addressed data cache valid bit to "invalid" if the data address hits the data cache. The block of data at the given address is removed from the cache and written back to the main memory if the dirty bit is set.

The Purge Data Cache (PDC) instruction sets the addressed data cache valid bit to "invalid" if the data address hits the cache. The block of data at the given address is removed from the cache and no write-back is performed.

The Flush Instruction Cache (FIC) instruction sets the addressed instruction cache valid bit to "invalid" if the address hits the cache. The instruction at the given address is removed from the cache.

The Flush Data Cache Entry (FDCE) instruction is a special kind of flush that can be used in a routine to flush the entire cache, for example in the event of a processor failure. This routine is implementation dependent. For a multiple-set cache, the routine steps through the index range of cache once for each set. The FDCE instruction flushes a block of data and sets the addressed data cache valid bit to "invalid" whether or not there is a hit at the cache index. That is, the block is written back to main memory if and only if it is valid and dirty, without comparing the cache tag to any requested address.

The Flush Instruction Cache Entry (FICE) instruction accomplishes the same function in the

instruction cache as the FDCE instruction accomplishes in the data cache.

The Synchronize Caches (SYNC) instruction suspends instruction execution by the processor until the completion of all instruction cache and data cache operations. This guarantees that any reference to data will await the completion of the cache operations required to ensure the integrity of that data.

The operation of the system is illustrated by the following examples. The operating system controls access to main memory by the processor and by the peripheral devices attached to I/O channel 15.

When data is to be read into main memory 13 from an external device through I/O channel 15, the operating system must insure that the addresses into which or from which the data is transferred do not overlap areas mapped into either data or instruction caches. In order to clear any stale data out of the caches, before the I/O is performed, the system broadcasts to each cache the FDC and FIC instruction over the range of addresses into which the input data is to be mapped.

When data is to be read out of main memory to an external device through I/O channel 15, the operating system must insure that the addresses from which the data is transferred do not overlap areas mapped into data caches, so that the most up-to-date data is transferred. In order to update main memory with the data in the caches that has been modified by the processors, the system broadcasts to each cache the FDC instruction for the range of addresses from which the output data is to be read. The FDC instruction causes the cache to write any dirty blocks back to main memory.

In a virtual memory system, whenever a page or segment is moved from main memory 13 to a peripheral memory (eg., a disc memory) connected to I/O channel 15, the data from the page or segment must be flushed from all caches. The operating system broadcasts to the caches the FDC and FIC instruction over the range of addresses included in the page or segment. When a page or segment is destroyed, for example because of program termination, the data must be removed from the cache but need not be stored. In this instance, the operating system uses the PDC and FIC instructions. No flush or purge operations are needed when a page or segment is created or brought in from a peripheral memory because the addresses into which it is mapped will have just been flushed or purged during the removal of the previous page or segment to make room for the new page or segment.

In order to accommodate programs with self-modifying code, the operating system must remove

from the caches any stale copies of the modified instruction to guarantee that only the new version of the instruction will be executed. After the modification of the instruction has been done in data cache 51, the operating system uses the FDC instruction to force the modified copy out to main memory, uses the FIC instruction to remove any stale copy of the instruction from instruction cache 41, then executes the SYNC instruction to insure that the modified instruction is not invoked until the FDC and FIC instructions have been completed.

In the event of a processor failure, for example caused by a power failure, the modified blocks of data residing in the caches must be written back to main memory. The operating system can accomplish this in a minimal amount of time with the FDCE and FICE instructions. A routine using the FDCE and FICE instructions flushes the caches quickly because by stepping through the index range of the caches rather than using the address space which is much larger. As the routine steps through the caches, only the blocks that are valid and dirty are written back to main memory 13.

Claims

1. A computer system having a multi-level memory hierarchy and means for maintaining the integrity of the blocks of information stored at different levels in the hierarchy, comprising:
 a processor (19) for executing instructions and processing data;
 memory (13) for storing instructions and data;
 an I/O channel (15) connected to the memory (13) and to I/O devices for transferring data and instructions into and out of the memory (13);
 a cache (21) connected between the processor (19) and the memory (13) for storing selected blocks of information from the memory (13) for use by the processor (19), and comprising hit/miss indicating means and having associated with each stored block a valid status flag and a dirty status flag;
 an operating system controlling access to said memory by the processor and the I/O devices;
characterized by
 a set of instructions for providing explicit control of the removal on invalidation of blocks of data from the cache (21) and of the updating of the memory from the cache before I/O operations are executed, the instruction set comprising Flush Data Cache, Purge Data Cache, Flush Instruction Cache, Flush Data Cache Entry, Flush Instruction Cache Entry and Synchronize Caches instructions; and
 said operating system capable of causing the execution of certain of the instructions from the

instruction set to ensure the consistency of the information stored in the cache (21) with the information transferred into and out of memory (13) in the event of transfer of data or instructions into or out of memory (13) via the I/O channel (15); modification of an instruction in the cache (21), or failure of processor (19), whereby on a hit condition the valid bit is set to "invalid", and depended on the value of the dirty bit information is removed from the cache either with or without copy-back to the memory.

2. Computer system according to claim 1, **characterized** in that prior to transfer of data or instructions into or out of memory via the I/O channel, the operating system broadcasts the Flush Data Cache and Flush Instruction Cache instructions to the cache over the range of addresses into which or out of which data is transferred.
3. Computer system according to claim 2, **characterized** in that virtual memory (13) is used; in that, when a page or a segment is removed from memory (13), the operating system broadcasts the Flush Data Cache and Flush Instruction Cache instructions to the cache (21) over the range of addresses in the page or segment; and in that, when a page or segment is destroyed, the operating system broadcasts the Purge Data Cache and Flush Instruction Cache instructions to the cache (21) over the range of addresses in the page or segment.
4. Computer system according to claim 2 or 3, **characterized** in that the cache (21) is divided into two segments (41,51), a data cache (51) for storing data and an instruction cache (41) for storing instructions; that instructions can be modified when stored as part of a block of data in the data cache (51); and in that after modification of an instruction, the operating system issues the Flush Data Cache instruction to the data cache (51) for the address of the block including the modified instruction, issues the Flush Instruction Cache instruction to the instruction cache (41) for the address of the modified instruction and then executes the Synchronize Caches instruction.
5. Computer system according to any of claims 2 to 4, **characterized** in that in the event of a processor failure, the operating system executes a routine including the Flush Data Cache Entry and Flush Instruction Cache Entry instructions over the index range for the data cache (51) and for the instruction cache (41).



from the caches any stale copies of the modified instruction to guarantee that only the new version of the instruction will be executed. After the modification of the instruction has been done in data cache 51, the operating system uses the FDC instruction to force the modified copy out to main memory, uses the FIC instruction to remove any stale copy of the instruction from instruction cache 41, then executes the SYNC instruction to insure that the modified instruction is not invoked until the FDC and FIC instructions have been completed.

In the event of a processor failure, for example caused by a power failure, the modified blocks of data residing in the caches must be written back to main memory. The operating system can accomplish this in a minimal amount of time with the FDCE and FICE instructions. A routine using the FDCE and FICE instructions flushes the caches quickly because by stepping through the index range of the caches rather than using the address space which is much larger. As the routine steps through the caches, only the blocks that are valid and dirty are written back to main memory 13.

Claims

1. A computer system having a multi-level memory hierarchy and means for maintaining the integrity of the blocks of information stored at different levels in the hierarchy, comprising:
 - a processor (19) for executing instructions and processing data;
 - memory (13) for storing instructions and data;
 - an I/O channel (15) connected to the memory (13) and to I/O devices for transferring data and instructions into and out of the memory (13);
 - a cache (21) connected between the processor (19) and the memory (13) for storing selected blocks of information from the memory (13) for use by the processor (19), and comprising hit/miss indicating means and having associated with each stored block a valid status flag and a dirty status flag;
 - an operating system controlling access to said memory by the processor and the I/O devices;
 characterized by
 - a set of instructions for providing explicit control of the removal on invalidation of blocks of data from the cache (21) and of the updating of the memory from the cache before I/O operations are executed, the instruction set comprising Flush Data Cache, Purge Data Cache, Flush Instruction Cache, Flush Data Cache Entry, Flush Instruction Cache Entry and Synchronize Caches instructions; and
 - said operating system capable of causing the execution of certain of the instructions from the

instruction set to ensure the consistency of the information stored in the cache (21) with the information transferred into and out of memory (13) in the event of transfer of data or instructions into or out of memory (13) via the I/O channel (15); modification of an instruction in the cache (21), or failure of processor (19), whereby on a hit condition the valid bit is set to "invalid", and depended on the value of the dirty bit information is removed from the cache either with or without copy-back to the memory.

2. Computer system according to claim 1, characterized in that prior to transfer of data or instructions into or out of memory via the I/O channel, the operating system broadcasts the Flush Data Cache and Flush Instruction Cache instructions to the cache over the range of addresses into which or out of which data is transferred.
3. Computer system according to claim 2, characterized in that virtual memory (13) is used; in that, when a page or a segment is removed from memory (13), the operating system broadcasts the Flush Data Cache and Flush Instruction Cache instructions to the cache (21) over the range of addresses in the page or segment; and in that, when a page or segment is destroyed, the operating system broadcasts the Purge Data Cache and Flush Instruction Cache instructions to the cache (21) over the range of addresses in the page or segment.
4. Computer system according to claim 2 or 3, characterized in that the cache (21) is divided into two segments (41, 51), a data cache (51) for storing data and an instruction cache (41) for storing instructions; that instructions can be modified when stored as part of a block of data in the data cache (51); and in that after modification of an instruction, the operating system issues the Flush Data Cache instruction to the data cache (51) for the address of the block including the modified instruction, issues the Flush Instruction Cache instruction to the instruction cache (41) for the address of the modified instruction and then executes the Synchronize Caches instruction.
5. Computer system according to any of claims 2 to 4, characterized in that in the event of a processor failure, the operating system executes a routine including the Flush Data Cache Entry and Flush Instruction Cache Entry instructions over the index range for the data cache (51) and for the instruction cache (41).



Revendications

1. Un système d'ordinateur comportant une hiérarchie de mémoires à niveaux multiples et des moyens pour maintenir l'intégrité des blocs d'informations mémorisés à différents niveaux de la hiérarchie, comprenant:
 - 5 un processeur (19) pour exécuter des instructions et traiter des données;
 - 10 une mémoire (13) pour mémoriser des instructions et des données; un canal entrée/sortie (ou I/O) (15) relié à la mémoire (13) et à des dispositifs I/O pour transférer des données et des instructions dans la mémoire (13) et à partir de celle-ci;
 - 15 un cache (21) relié entre le processeur (19) et la mémoire (13) pour mémoriser, à partir de la mémoire (13), des blocs choisis d'informations en vue d'une utilisation par le processeur (19) et comprenant des moyens indiquant réussi/manqué et comprenant un drapeau de statut valide et un drapeau de statut maculé associés à chaque bloc mémorisé;
 - 20 un système d'exploitation commandant l'accès à ladite mémoire par le processeur et les dispositifs I/O;
 2. Système d'ordinateur selon la revendication 1, **caractérisé** en ce que, avant un transfert de données ou d'instructions dans la mémoire ou à partir de celle-ci par l'intermédiaire du canal I/O, le système d'exploitation émet des instructions Vider Cache de Données et Vider Cache d'Instructions vers le cache sur la plage d'adresses dans lesquelles, ou à partir desquelles, des données sont transférées.
 3. Système d'ordinateur selon la revendication 2, **caractérisé** en ce qu'il est utilisé une mémoire virtuelle (13); en ce que, lorsqu'une page ou un segment est enlevé de la mémoire (13), le système d'exploitation émet les instructions Vider Cache de Données et Vider Cache d'Instructions vers le cache (21) sur la plage d'adresses de la page ou du segment; et en ce que, lorsqu'une page ou un segment est détruit, le système d'exploitation émet les instructions Vider Cache de Données et Vider Cache d'Instructions vers le cache (21) sur la plage d'adresses de la page ou du segment.
 4. Système d'ordinateur selon la revendication 2 ou 3, **caractérisé** en ce que le cache (21) est divisé en deux segments (41, 51), un cache de données (51) pour mémoriser des données et un cache d'instructions (41) pour mémoriser des instructions; en ce que des instructions peuvent être modifiées lorsqu'elles sont mémorisées en tant que parties d'un bloc de données dans le cache de données (51); et en ce que, après modification d'une instruction, le système d'exploitation envoie l'instruction Vider Cache de Données vers le cache (51) de données pour l'adresse du bloc comprenant l'instruction modifiée, envoie l'instruction Vider Cache d'Instructions vers le cache d'Instructions (41) pour l'adresse d'une instruction modifiée et exécute ensuite l'instruction Synchroniser Caches.
 5. Système d'ordinateur selon l'une quelconque des revendications 2 à 4, **caractérisé** en ce que, dans le cas d'une défaillance du processeur, le système d'exploitation exécute un programme comprenant les instructions Vider Entrée de Cache de Données et Vider Entrée de Cache d'Instructions sur la plage d'index pour le cache de données (51) et le cache d'Instructions (41).
- Patentansprüche**
1. Rechnersystem mit einer mehrstufigen Speicherhierarchie und einer Vorrichtung zum Aufrechterhalten der Unversehrtheit der in verschiedenen Hierarchiestufen gespeicherten Informationsblöcke, mit:



einem Prozessor (19) zum Ausführen von Befehlen und Verarbeiten von Daten;

einem Speicher (13) zum Speichern von Befehlen und Daten;

einem mit dem Speicher (13) und Ein-
Ausgabegeräten verbundenen Ein-
Ausgabekanal (15) zum Übertragen von Daten
und Befehlen in den und aus dem Speicher
(13);

einem zwischen Prozessor (19) und Speicher
(13) geschalteten Cache-Speicher (21) zum
Speichern ausgewählter Informationsblöcke
aus dem Speicher (13), die für den Prozessor
(19) nutzbar sind, und

mit einer Treffer/Fehlreffer-Anzeige,

wobei jedem gespeicherten Block eine zulässige
Zustandsmarke (valid status flag) und eine
unzulässige Zustandsmarke (dirty status flag)
zugeordnet ist, sowie

mit einem Betriebssystem, das den Zugriff auf
den Speicher durch den Prozessor und die
Ein-/Ausgabegeräte steuert,

gekennzeichnet durch

einen Befehlssatz für die explizite Steuerung
des Löschvorgangs oder der Annullierung von
Datenblöcken aus dem Cache-Speicher (21)
und der Aktualisierung des Speichers durch
den Cache-Speicher, bevor Ein-

Ausgabeoperationen ausgeführt werden, wobei
der Befehlssatz folgende Befehle umfaßt:
Daten-Cache-Speicher spülen (Flush Data Cache),
Daten-Cache-Speicher löschen (Purge Data Cache),
Befehls-Cache-Speicher spülen (Flush Instruction Cache),
Daten-Cache-Speichereingang spülen (Flush Data Cache Entry),
Befehls-Cache-Speichereingang spülen (Flush Instruction Cache Entry) und
Cache-Speicher synchronisieren (Synchronize Caches); und

das Betriebssystem, das bestimmte Befehle
aus dem Befehlssatz ausführen kann, um die
Übereinstimmung von in dem Cache-Speicher
(21) gespeicherten Informationen mit in den
und aus dem Speicher (13) übertragenen Informationen
bei der Übertragung von Daten oder
Befehlen in den oder aus dem Speicher (13)
über den Ein-/Ausgabekanal (15) sicherzustellen;
sowie durch

Verändern eines Befehles in dem Cache-Speicher
(21) oder Ausfall des Prozessors (19),
wobei bei einer Treffer-Bedingung das zulässige
Bit (valid bit) auf "unzulässig" (invalid) gesetzt
wird und abhängig von dem unzulässigen
Bit (dirty bit) Information, mit oder ohne Rück-
schreiben in den Speicher, aus dem Cache-
Speicher entfernt wird.

2. Rechnersystem nach Anspruch 1, dadurch **gekennzeichnet**, daß vor dem Übertragen von

Daten oder Befehlen in den oder aus dem
Speicher über den Ein-/Ausgabekanal das Betriebssystem
die Befehle Flush Data Cache
und Flush Instruction Cache für den Adressbereich,
in welchen oder aus welchem Daten
übertragen werden sollen, zum Cache-Speicher
sendet.

3. Rechnersystem nach Anspruch 2, dadurch **gekennzeichnet**, daß ein virtueller Speicher (13) eingesetzt ist, daß beim Entfernen einer Seite oder eines Segmentes aus dem Speicher (13) das Betriebssystem die Befehle Flush Data Cache und Flush Instruction Cache für den Adressbereich in der Seite oder dem Segment zum Cache-Speicher (21) sendet; und daß beim Löschen einer Seite oder eines Segmentes das Betriebssystem die Befehle Purge Data Cache und Flush Instruction Cache für den Adressbereich in der Seite oder dem Segment zum Cache-Speicher (21) sendet.

4. Rechnersystem nach Anspruch 2 oder 3, dadurch **gekennzeichnet**, daß der Cache-Speicher (21) in zwei Segmente (41, 51) aufgeteilt ist, einen Daten-Cache-Speicher (51) zum Speichern von Daten und einen Befehls-Cache-Speicher (41) zum Speichern von Befehlen, daß als Teil eines Datenblocks im Daten-Cache-Speicher (51) gespeicherte Befehle verändert werden können; und daß nach dem Verändern eines Befehles das Betriebssystem den Befehl Flush Data Cache an den Daten-Cache-Speicher (51) für die Adresse des Blocks mit dem veränderten Befehl, den Befehl Flush Instruction Cache an den Befehls-Cache-Speicher (41) für die Adresse des veränderten Befehles ausgibt und dann den Befehl Synchronize Caches ausführt.

5. Rechnersystem nach einem der Ansprüche 2 bis 4, dadurch **gekennzeichnet**, daß bei einem Ausfall des Prozessors das Betriebssystem eine Routine ausführt, die die Befehle Flush Data Cache Entry und Flush Instruction Cache Entry für den Indexbereich des Daten-Cache-Speichers (51) und des Befehls-Cache-Speichers (41) beinhaltet.