

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 April 2006 (13.04.2006)

PCT

(10) International Publication Number
WO 2006/039153 A1

(51) International Patent Classification:
G06F 12/08 (2006.01) *G06F 12/12* (2006.01)

(21) International Application Number:
PCT/US2005/033671

(22) International Filing Date:
21 September 2005 (21.09.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/956,560 1 October 2004 (01.10.2004) US

(71) Applicant (for all designated States except US): **ADVANCED MICRO DEVICES, INC.** [US/US]; One AMD Place, Mail Stop 68, P.O. Box 3453, Sunnyvale, CA 94088-3453 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **GOLDEN, Michael, L.** [US/US]; 3297 Bellarmine Court, Santa Clara, CA 95051 (US). **KLASS, Richard, E.** [US/US]; 4716 Tampico Way, San Jose, CA 95118 (US).

(74) Agent: **DRAKE, Paul, S.**; Advanced Micro Devices, Inc., 5204 East Ben White Boulevard, Mail Stop 562, Austin, TX 78741 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

[Continued on next page]

(54) Title: DYNAMIC RECONFIGURATION OF CACHE MEMORY



(57) Abstract: In one embodiment, a processing node includes a plurality of processor cores each including a cache memory coupled to a cache monitor unit and to a configuration unit. Each cache monitor unit may be configured to independently monitor a current utilization of the cache memory to which it is coupled and to determine whether the current utilization is below a predetermined utilization value. The configuration unit may selectively disable one or more portions of the cache memory in response to the cache monitor unit determining that the current utilization is below the predetermined utilization value.

WO 2006/039153 A1



-
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

TITLE: DYNAMIC RECONFIGURATION OF CACHE MEMORY**Technical Field**

This invention is related to the field of microprocessors and chip multiprocessors (CMP) and, more particularly, to reconfiguration of cache memory of a processor.

Background Art

[0001] Modern microprocessors typically include on-chip cache memory. In many cases, microprocessors include an on-chip hierarchical cache structure that may include level one (L1), level two (L2) and in some cases a level three (L3) cache memory. Typical cache hierarchies may employ a small fast L1, cache that may be used to store the most frequently used cache lines. The L2 may be a larger and possibly slower cache for storing cache lines that are accessed but don't fit in the L1. The L3 cache may be used to store cache lines are accessed but do not fit in the L2 cache. Having a cache hierarchy as described above may improve processor performance by reducing the latencies associated with memory access by the processor core.

[0002] However, in certain instances such an improvement in performance may come at a cost. It is well known that modern microprocessors can consume a great deal of power and have high thermal budgets. Cache memories may consume power via static leakage even when they are not used. Thus, the larger the cache the more power consumed.

[0003] In an effort to increase efficiency and processor, chip multiprocessors (CMPs) are becoming an emerging technology that is gaining increased popularity. A CMP has two or more processor cores implemented on the same integrated circuit (IC) device. The increase in popularity may be due, at least in part, to the notion that a CMP may be a more efficient use of the millions of transistors that may be included on an integrated circuit than, for example, a more elaborate single processor. For example, by sharing processing loads between processor cores, a CMP may execute some tasks at a lower frequency than that required by some single core processors. Lower frequencies may translate in some cases to savings in power and a reduction in the thermal budget.

[0004] Since CMPs include multiple processor cores, they may also include circuitry associated with each core. For example, a CMP may include an L1 and an L2 cache memory for each processor core. Accordingly, for the reasons described above relative to single core microprocessors, the power consumption and thermal budget of a CMP may still be high. Thus, it may be desirable to find ways to reduce power consumption while maintaining processor performance.

DISCLOSURE OF INVENTION

[0005] Various embodiments of dynamic reconfiguration of a cache memory are disclosed. In one embodiment, a processor includes a cache memory, such as an L2 cache, for example coupled to a cache monitor unit and a configuration unit. The cache monitor unit may be configured to monitor a current utilization of the cache memory and to determine whether the current utilization is below a predetermined utilization value. The configuration unit may selectively disable one or more portions of the cache memory in response to the cache monitor unit determining that the current utilization is below the predetermined utilization value.

[0006] In another embodiment, a processing node includes a plurality of processor cores each including a cache memory coupled to a cache monitor unit and to a configuration unit. Each cache monitor unit may be configured to independently monitor a current utilization of the cache memory to which it is coupled and to determine whether the current utilization is below a predetermined utilization value. The configuration unit may selectably disable one or more portions of the cache memory in response to the cache monitor unit determining that a given current utilization is below the predetermined utilization value.

BRIEF DESCRIPTION OF DRAWINGS

[0007] FIG. 1 is a block diagram of one embodiment of a computer system including a multi-core processing node and configurable cache memory.

[0008] FIG. 2 is a block diagram of another embodiment of a computer system including a multi-core processing node and configurable cache memory.

[0009] FIG. 3 is a flow diagram describing the operation of one embodiment of a processing node.

[0010] FIG. 4 is a block diagram of one embodiment of microprocessor including a single processor core and configurable cache memory.

[0011] FIG. 5 is a flow diagram describing the operation of one embodiment of the microprocessor of FIG. 4.

[0012] While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the present invention as defined by the appended claims. Note, the headings are for organizational purposes only and are not meant to be used to limit or interpret the description or claims. Furthermore, note that the word "may" is used throughout this application in a permissive sense (i.e., having the potential to, being able to), not a mandatory sense (i.e., must). The term "include" and derivations thereof mean "including, but not limited to." The term "connected" means "directly or indirectly connected," and the term "coupled" means "directly or indirectly coupled."

MODE(S) FOR CARRYING OUT THE INVENTION

Chip multiprocessor with dynamically configurable cache

[0013] Turning now to FIG. 1, a block diagram of one embodiment of a computer system 10 is shown. In the illustrated embodiment, the computer system 10 includes a processing node 12 coupled to memory 14 and to peripheral devices 13A-13B. The node 12 includes processor cores 15A-15B coupled to a node controller 20 which is further coupled to a memory controller 22 and a plurality of HyperTransport™ (HT) interface circuits 24A-24C. The processor cores 15A-15B are also coupled to a shared level three (L3) cache memory 60, which is coupled to an L3 cache monitor 85 and a configuration unit 50. The HT circuit 24C is coupled to the peripheral device 16A, which is coupled to the peripheral device 16B in a daisy-chain configuration (using HT interfaces, in this embodiment). The remaining HT circuits 24A-B may be connected to other similar processing nodes (not shown) via other HT interfaces (not shown). The memory controller 22 is coupled to the memory 14. In one embodiment, node 12 may be a single integrated circuit chip comprising the circuitry shown therein in FIG. 1. That is, node 12

may be a chip multiprocessor (CMP). Any level of integration or discrete components may be used. It is noted that processing node 12 may include various other circuits that have been omitted for simplicity.

[0014] In one embodiment, node controller 20 may include various interconnection circuits (not shown) for interconnecting processor cores 15A and 15B to each other, to other nodes and to memory. Node controller 20 may also include sets of fuses (not shown) that may be selectively blown during the manufacturing process to represent a plurality of values. In some embodiments, various node properties that may be selected by the fuses. The properties include the maximum and minimum operating frequencies for the node and the maximum and minimum power supply voltages for the node. In addition, the fuses may select processor-core specific properties.

[0015] The node controller 20 may generally be configured to route communications between the processor cores 15A-15B, the memory controller 22, and the HT circuits 24A-24C dependent upon the communication type, the address in the communication, etc. In one embodiment, the node controller 20 includes a system request queue (SRQ) into which received communications are written by the node controller 20. The node controller 20 may schedule communications from the SRQ for routing to the destination or destinations among the processor cores 15A-15B, the HT circuits 24A-24C, and the memory controller 22. Operation of the node 12 and its components will be discussed in more detail below.

[0016] Generally, the processor cores 15A-15B may use the interface(s) to the node controller 20 to communicate with other components of the computer system 10 (e.g. peripheral devices 16A-16B, other processor cores (not shown), the memory controller 22, etc.). The interface may be designed in any desired fashion. Cache coherent communication may be defined for the interface, in some embodiments. In one embodiment, communication on the interfaces between the node controller 20 and the processor cores 15A-15B may be in the form of packets similar to those used on the HT interfaces. In other embodiments, any desired communication may be used (e.g. transactions on a bus interface, packets of a different form, etc.). In other embodiments, the processor cores 15A-15B may share an interface to the node controller 20 (e.g. a shared bus interface). Generally, the communications from the processor cores 15A-15B may include requests such as read operations (to read a memory location or a register external to the processor core) and write operations (to write a memory location or external register), responses to probes (for cache coherent embodiments), interrupt acknowledgements, and system management messages, etc.

[0017] The memory 14 may include any suitable memory devices. For example, a memory 14 may comprise one or more RAMBUS DRAMs (RDRAMs), synchronous DRAMs (SDRAMs), double data rate (DDR) SDRAM, static RAM, etc. The memory controller 22 may comprise control circuitry for interfacing to the memories 14. Additionally, the memory controller 22 may include request queues for queuing memory requests, etc.

[0018] The HT circuits 24A-24C may comprise a variety of buffers and control circuitry for receiving packets from an HT link and for transmitting packets upon an HT link. The HT interface comprises unidirectional links for transmitting packets. Each HT circuit 24A-24C may be coupled to two such links (one for transmitting and one for receiving). A given HT interface may be operated in a cache coherent fashion (e.g. between processing nodes) or in a non-coherent fashion (e.g. to/from peripheral devices 16A-16B). In the illustrated embodiment, the HT circuits 24A-24B are not in use, and the HT circuit 24C is coupled via non-coherent links to the peripheral devices 16A-16B.

[0019] The peripheral devices 16A-16B may be any type of peripheral devices. For example, the peripheral

devices 16A-16B may include devices for communicating with another computer system to which the devices may be coupled (e.g. network interface cards, circuitry similar to a network interface card that is integrated onto a main circuit board of a computer system, or modems). Furthermore, the peripheral devices 16A-16B may include video accelerators, audio cards, hard or floppy disk drives or drive controllers, SCSI (Small Computer Systems Interface) adapters and telephony cards, sound cards, and a variety of data acquisition cards such as GPIB or field bus interface cards. It is noted that the term "peripheral device" is intended to encompass input/output (I/O) devices.

[0020] Generally, a processor core 15A-15B may include circuitry that is designed to execute instructions defined in a given instruction set architecture. That is, the processor core circuitry may be configured to fetch, decode, execute, and store results of the instructions defined in the instruction set architecture. For example, in one embodiment, processor cores 15A-15B may implement the x86 architecture. The processor cores 15A-15B may comprise any desired configurations, including superpipelined, superscalar, or combinations thereof. Other configurations may include scalar, pipelined, non-pipelined, etc. Various embodiments may employ out of order speculative execution or in order execution. The processor cores may include microcoding for one or more instructions or other functions, in combination with any of the above constructions. Various embodiments may implement a variety of other design features such as caches, translation lookaside buffers (TLBs), etc.

[0021] It is noted that, while the present embodiment uses the HT interface for communication between nodes and between a node and peripheral devices, other embodiments may use any desired interface or interfaces for either communication. For example, other packet based interfaces may be used, bus interfaces may be used, various standard peripheral interfaces may be used (e.g., peripheral component interconnect (PCI), PCI express, etc.), etc.

Configurable cache control

[0022] In the illustrated embodiment, processor core 15A includes an L2 cache 17A. Likewise, processor core 15B includes an L2 cache 17B. The respective L2 caches may be representative of any L2 cache found in a microprocessor. However, in one embodiment, the L2 caches 17A-17B may be implemented using a number of independently controllable memory blocks, as indicated by the dashed lines. In one implementation, each of the blocks may include an independently controllable power grid. As such the power (VDD) supplied to each of the memory blocks may be independently removed or in some cases the voltage supplied to the blocks by the individual power grids may be independently lowered.

[0023] Like the L2 caches described above, in the illustrated embodiment, the L3 cache 60 may also be implemented using a number of independently controllable memory blocks, as indicated by the dashed lines, in which each of the blocks may include an independently controllable power grid. As such the power (VDD) supplied to each of the memory blocks may be independently removed or in some cases VDD supplied to the blocks by the individual power grids may be independently lowered.

[0024] In the illustrated embodiment, similar to the L3 cache, each of the respective L2 cache memories 17A-17B are coupled to an L2 cache monitor unit 18A-18B, respectively and to a configuration unit 19A-19B, respectively. In one embodiment, each cache monitor unit (e.g., 18A, 19A, 85) may be configured to monitor the utilization of the respective cache memory connected to it. For example, L2 cache monitor unit 18A may monitor the utilization of L2 cache 17A. In the illustrated embodiment, each cache monitor unit includes circuits (e.g., 40A, 40B, 86) that may determine the current cache utilization and determine whether the current utilization is within a given utilization limit. More particularly, by way of example, cache monitor unit 18A may determine if the current

utilization of L2 cache 17A is below a predetermine threshold value. If so, the cache monitor unit 18A may notify configuration unit 19A.

[0025] In one embodiment, configuration unit 19A includes one or more storages designated 42A such as programmable registers, for example, that may store values which correspond to various configuration preferences and states of the L2 cache 17A. Depending upon the values stored within storages 42A, configuration unit 19A may be programmed to selectably enable or disable one or more of the independently controllable memory blocks of L2 cache 17A in response to the notification from the L2 cache monitor 18A. In addition, the values stored within storages 42A may determine whether the power grid for a given block may be turned off or the voltage lowered. For example, in certain cases such as when processor core performs a context switch, the L2 cache may not be used for the new thread or context. However, it may be desirable for the data stored within L2 cache 17A to be ready when the context is switched back. In such cases, simply lowering the voltage to one or more blocks may allow the data stored therein to be retained but further access may be prohibited or disabled until the blocks are re-enabled and the voltage is brought back to normal.

[0026] Configuration unit 50 may be further configured to allow processor core 15A to share the L2 cache 17B of processor core 15B and vice versa dependent upon the configuration values stored within it. For example, in response to cache monitor 40A determining that the current utilization of L2 cache 17A is above an upper threshold value, configuration unit 42A may request access to L2 cache 17B. In such cases, configuration unit 42B, may allow the shared access if L2 cache 17B has blocks that are disabled and therefore sharable.

[0027] In one embodiment, the storages 42A, 42B and 51 of configuration units 19A, 19B and 50, respectively may include default preference values. Alternatively, preference values may be programmed into the storages. These preference values may be used to determine which action may be taken in response to receiving a notification from the corresponding cache monitor unit. For example, as described above, the configuration units may use these values to determine whether to disable the portions of the L2 cache in the corresponding processor core and the L3 cache by removing power or lowering the voltage. In addition, in one embodiment, the configuration units may use the values to determine which, if any, portions of the L2 or L3 cache memories should be disabled at power-up.

[0028] In one embodiment, storages 42A, 42B and 51 may be programmed using software. For example, special instructions may be used to access storages 42A, 42B and 51. In another embodiment, storages 42A, 42B and 51 may be programmed more permanently using hard fuses as described above. In such an embodiment, the values are read from blown fuses and stored within the storages. In one embodiment, the hard fuses may be overridden by external programming. In such an embodiment, storages 42A, 42B and 51 may be programmed via configuration ports 44A, 44B and 54, respectively, via an external port such as the joint test action group (JTAG) port, for example. Overriding the hard fuses configuration in this way is sometimes referred to as programming via soft fuses.

[0029] In one embodiment, cache monitors 18A, 18B and 85 may determine the current utilization using the parameters such as how many entries of a respective cache are being used and how frequently the respective cache is accessed. In one implementation, the number of entries in a given cache that are being used may be determined without using probes using a counter (not shown explicitly, part of 40A, 40B and 86) to count writes to the cache and by monitoring whether a write generates a victim or not. If a victim is not generated, then the write must have resulted in the allocation of a new entry. Otherwise, the number of entries stays the same. Thus if a victim is

generated the write may not be counted. In addition, in one implementation, the frequency in which a cache is accessed may be determined using a saturation counter, for example (not explicitly shown, part of circuit 40A). In one embodiment, the saturation counter may be incremented each time the cache is accessed. The counter may be decremented at some clock interval, but never below zero. If the count value stays above some predetermined value, it may be assumed that the cache access frequency is high; thereby indicating a high utilization. From these two parameters a lower predetermined utilization value and an upper predetermined utilization value may be calculated. It is noted that in other embodiments, other mechanisms may be employed to determine the cache utilization as desired.

[0030] It is noted that, while the computer system 10 illustrated in FIG. 1 includes one processing node 12, other embodiments may implement any number of processing nodes. Similarly, a processing node such as node 12 may include any number of processor cores, in various embodiments. Various embodiments of the computer system 10 may also include different numbers of HT interfaces per node 12, and differing numbers of peripheral devices 16 coupled to the node, etc.

[0031] Turning to FIG. 2 a block diagram of another embodiment of a computer system 20 is shown. Components that correspond to those shown in FIG. 1 are numbered identically for clarity and simplicity. Similar to the computer system 10 of FIG. 1, computer system 20 of FIG. 2 includes a processing node 22 coupled to a memory 14 and to peripheral devices 13A-13B. Processing node 22 includes similar functionality to node 12 of FIG. 1. Accordingly, differing functionality will be described while description of functionality common to node 12 will be omitted here for brevity. It is noted that processing node 12 of FIG. 2 may include various other circuits that have been omitted for simplicity.

[0032] In the embodiment illustrated in FIG. 2, like processing node 12, processing node 22 also includes processor cores 15A-15B coupled to a node controller 20 which is further coupled to a memory controller 22 and a plurality of HyperTransport™ (HT) interface circuits 24A-24C. In addition, the processor cores 15A-15B are also coupled to a shared level three (L3) cache memory 60. However, L3 cache 60 is coupled to a cache monitor 95 and a configuration unit 65. In one embodiment, node 22 may be a single integrated circuit chip comprising the circuitry shown therein in FIG. 2. That is, node 22 may be a chip multiprocessor (CMP).

[0033] In the illustrated embodiment, and in contrast to the embodiment of FIG. 1, each of the L2 cache memories 17A-17B are also coupled to cache monitor unit 95 and to configuration unit 65. In one embodiment, cache monitor unit 95 may be configured to monitor the utilization of all cache memories connected to it (e.g., L2 cache 17A-17B and L3 cache 60) using functionality as described above in conjunction with the description of FIG. 1.

[0034] In one embodiment, similar to the description above configuration unit 65 includes one or more storages designated 61 such as programmable registers, for example, that may store values which correspond to various configuration preferences and states of the L2 cache 17A-17B and L3 cache 60. Depending upon the values stored within storages 61, configuration unit 65 may be programmed to selectably enable or disable one or more of the independently controllable memory blocks of L2 cache 17A-17B and L3 cache 60 in response to the notification from the cache monitor unit 95. Likewise, the values stored within storages 61 may determine whether the power grid for a given block may be turned off or the voltage lowered. It is noted that storage 61 may be programmed similar to storages 42A, 42B and 51 described above.

[0035] Further, configuration unit 60 may be configured to allow processor core 15A to share the L2 cache 17B of processor core 15B and vice versa dependent upon the configuration values stored within it. For example, in response to cache monitor 95 determining that the current utilization of L2 cache 17A is above an upper threshold value, configuration unit 65 may enable access to L2 cache 17B by processor core 15A. In such cases, configuration unit 65, may allow the shared access if L2 cache 17B has blocks that are disabled and therefore sharable.

[0036] FIG. 3 illustrates a flow diagram describing the operation of one embodiment of processing node 12. Referring collectively to FIG. 1 and FIG. 3, in one embodiment, when node 12 is first powered up or brought out of reset, the configuration values stored within configuration unit 42A may be read and used to configure L2 cache 17A. Likewise, the configuration values stored within configuration unit 42B may be read and used to configure L2 cache 17B and the configuration values stored within configuration unit 50 may be read and used to configure L3 cache 60. During operation of node 12, cache monitor units 18A, 18B and 85 are configured to monitor the current utilization of their respective cache memories (block 300). For example, as described above, cache monitor units 18A, 18B and 85 may track cache access frequency and cache entry use to determine the current utilization for their respective caches.

[0037] In block 305, based upon the current utilization, cache monitor unit 85 may determine whether the respective current utilization is within predetermined limits. In one embodiment, cache monitor unit 85 may compare the current utilization of L3 cache 60 to an upper threshold value and a lower threshold value.

[0038] If the current utilization is within the predetermined threshold values, cache monitor unit 85 continues to monitor the utilization of L3 cache 60 as described above in block 300. However, if the current utilization is not within the predetermined threshold values and is below the lower threshold value (block 305), cache monitor unit 85 may provide notification to configuration unit 50 (block 310). Configuration unit 50 may selectably disable one or more blocks of L3 cache 60 (block 315). In one embodiment, configuration unit 50 may disable blocks according to configuration preferences stored within storages 51 and in response to the notification. For example, as described above, power may be removed from the selected blocks or it may be reduced to a lower voltage. In addition, in one embodiment, the blocks selected first may be the blocks furthest away from the processor cores 15A-15B, since these blocks may represent the longest path delay. Once the selected blocks have been disabled, cache monitor 85 continues to monitor the utilization as described above in block 300.

[0039] Referring back to block 305, if the current utilization is not within the predetermined threshold values and is above the upper threshold value, cache monitor unit 85 may provide notification to configuration unit 50 (block 320). Configuration unit 50 may selectably enable one or more disabled blocks of L3 cache 60 (block 315) if there are disabled blocks. In one embodiment, configuration unit 50 may enable blocks according to configuration preferences stored in storages 51. Once the selected blocks have been re-enabled, cache monitor 85 continues to monitor the utilization as described above in block 300.

[0040] In block 330, based upon the current utilization, cache monitor units 18A and 18B may determine whether the respective current utilization is within predetermined limits. More particularly, in one embodiment cache monitor units 18A, 18B may compare the current utilization of L2 caches 17A-17B to an upper threshold value and a lower threshold value.

[0041] If the current utilization is within the predetermined threshold values, cache monitor units 18A, 18B continue to monitor the utilization of L2 caches 17A-17B as described above in block 300. However, if the current utilization of either cache is not within the predetermined threshold values and is below the lower threshold value (block 330), the appropriate cache monitor unit 18A-18B may provide notification to configuration unit 19A-19B (block 335). Configuration units 19A-19B may selectably disable one or more blocks of the corresponding L2 caches 17A-17B (block 340). In one embodiment, configuration units 19A-19B may disable blocks according to configuration preferences stored within storages 42A-42b, respectively and in response to the notification. For example, as described above, power may be removed from the selected blocks or it may be reduced to a lower voltage. In addition, in one embodiment, the blocks selected first may be the blocks furthest away from the CPU logic of processor cores 15A-15B, since these blocks may represent the longest path delay. Once the selected blocks have been disabled, cache monitor units 18A-18B continue to monitor the utilization as described above in block 300.

[0042] Referring back to block 330, if the current utilization is not within the predetermined threshold values and is above the upper threshold value, cache monitor units 18A-18B may provide notification to configuration unit 19A-19B (block 345). If any L2 blocks are disabled in the affected L2 cache (block 350), configuration unit 19A-19B may selectably enable one or more blocks of L2 caches 17A-17B (block 355). In one embodiment, configuration units 19A-19B may disable blocks according to configuration preferences stored within storages 42A-42B, respectively and in response to the notification. For example, as described above, power may be removed from the selected blocks or it may be reduced to a lower voltage. In addition, in one embodiment, the blocks selected first may be the blocks furthest away from the processor cores 15A-15B, since these blocks may represent the longest path delay. Once the selected blocks have been disabled, cache monitor units 18A-18B continue to monitor the utilization as described above in block 300.

[0043] However, if there are no disabled blocks in the affected L2 cache (block 350), the configuration unit of the affected L2 cache (i.e., the L2 cache that is over-utilized) may request from the configuration unit of the other processor core, access permission to share the L2 cache of the other core. For example, if cache monitor unit 18B determines that the current utilization of L2 cache 17B is over the upper threshold, configuration unit 19B may request access from configuration unit 19A to L2 cache 17A. In response to the request, configuration unit 19A may determine if L2 cache 17A has disabled blocks to share (block 360). If so, configuration unit 19A may grant access to processor core 15B to share the disabled blocks of L2 cache 17A (block 365). The disabled blocks may be re-enabled. Operation proceeds as described above in block 300.

[0044] The descriptions of FIG. 4 and FIG. 5, below, provide additional embodiments of configurable cache memory in the context of a microprocessor including a single processor core.

Single microprocessor with dynamically configurable cache

[0045] Referring now to FIG. 4, a block diagram of one embodiment of a microprocessor including a single processor core and configurable cache memory is shown. Microprocessor 415 includes processor core logic 450 which is coupled to an L1 cache 475, an L2 cache 417 and an L3 cache 460. Microprocessor 415 also includes a cache monitor unit 418 that is coupled to the L2 cache 417 and the L3 cache 460. Further, microprocessor 415 includes a configuration unit 419 that is coupled to cache monitor 418 and to each of L1 cache 475, L2 cache 417

and L3 cache 460. It is noted that microprocessor 415 may include various other circuits that have been omitted for simplicity.

[0046] In the illustrated embodiment, the functionality of cache monitor 418 and configuration unit 419 is similar to that of L2 cache monitor 95 and configuration unit 65 described above in conjunction with FIG. 2 and FIG. 3. Accordingly, their description has been omitted here for brevity.

[0047] Turning to FIG. 5, a flow diagram describing the operation of one embodiment of the microprocessor of FIG. 4 is shown. Referring collectively to FIG. 4 and FIG. 5, in one embodiment, when microprocessor 415 is first powered up or brought out of reset, the configuration values stored within storages 442 of configuration unit 419 may be read and used to configure L2 cache 417. Likewise, the configuration values stored within configuration unit 419 may be read and used to configure L3 cache 460. During operation of microprocessor 415, cache monitor unit 418 is configured to monitor the current utilization of each of the cache memories (block 500). For example, as described above, cache monitor unit 418 may track cache access frequency and cache entry use to determine the current utilization for each cache.

[0048] In block 505, based upon the current utilization, cache monitor unit 418 may determine whether the respective current utilization of each cache is within predetermined limits. In one embodiment, cache monitor unit 418 may compare the current utilization of L2 cache 417 and the current utilization of L3 cache 460 to an upper threshold value and a lower threshold value.

[0049] If the current utilization is within the predetermined threshold values, cache monitor unit 418 continues to monitor the utilization of L2 cache 417 and L3 cache 460 as described above in block 500. However, if the current utilization is not within the predetermined threshold values and is below the lower threshold value (block 505), cache monitor unit 418 may provide notification to configuration unit 419 (block 510). Configuration unit 419 may selectably disable one or more blocks of either L2 cache 417 and L3 cache 460 (block 515). In one embodiment, configuration unit 419 may disable blocks according to configuration preferences stored within storages 442 and in response to the notification. For example, as described above, power may be removed from the selected blocks or it may be reduced to a lower voltage. In addition, in one embodiment, the blocks selected first may be the blocks furthest away from the processor core logic 450, since these blocks may represent the longest path delay. Once the selected blocks have been disabled, cache monitor 418 continues to monitor the utilization as described above in block 500.

[0050] Referring back to block 505, if the current utilization is not within the predetermined threshold values and is above the upper threshold value, cache monitor unit 418 may provide notification to configuration unit 419 (block 320). Configuration unit 419 may selectably enable one or more disabled blocks of either L2 cache 417 and L3 cache 460 (block 525) if there are disabled blocks. In one embodiment, configuration unit 419 may enable blocks according to configuration preferences stored in storages 442. Once the selected blocks have been re-enabled, cache monitor 418 continues to monitor the utilization as described above in block 500.

[0051] Similar to the description above, in one embodiment, storages 440 and 442 may be programmed using software. For example, special instructions may be used to access storages 440 and 442. In another embodiment, storages 440 and 442 may be programmed more permanently using hard fuses as described above. In such an embodiment, the values are read from blown fuses and stored within the storages. In one embodiment, the hard fuses may be overridden by external programming. In such an embodiment, storages 440 and 442 may be

programmed via configuration port 44A via an external port such as the joint test action group (JTAG) port, for example. Overriding the hard fuses configuration in this way is sometimes referred to as programming via soft fuses.

5 [0052] It is noted that in the embodiments described above, an L3 cache memory is shown. However, it is contemplated that other embodiments may not include an L3 cache. The absence of an L3 cache does not in any way detract from the configurability of the L2 caches.

[0053] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

10 Industrial Applicability

[0054] This invention may generally be applicable to reconfiguration of cache memory for microprocessors and chip multiprocessors.

WHAT IS CLAIMED IS:

1. A method for configuring a cache memory (417, 460) of a processor (415) characterized in that, the method comprising:

5 monitoring a current utilization of the cache memory;
determining whether the current utilization is below a predetermined utilization value; and
selectably disabling one or more portions of the cache memory in response to determining that the current utilization is below the predetermined utilization value.

10 2. The method as recited in claim 1, further comprising storing the predetermined utilization value within a configurable storage (442).

3. The method as recited in claim 1, wherein determining whether the utilization is below a predetermined utilization value includes comparing the predetermined utilization value with the current utilization.

15 4. The method as recited in claim 1, further comprising determining the current utilization value by determining a number of entries being used in the cache memory and determining how frequently the cache memory is accessed.

20 5. The method as recited in claim 4, wherein determining a number of entries being used in the cache memory includes monitoring a number of writes to the cache memory that cause a victim to be generated.

6. The method as recited in claim 4, wherein determining how frequently the cache memory is accessed includes counting a number of accesses to the cache memory using a counter.

25 7. A processor (415) comprising:

a cache memory (417) (460);

a cache monitor unit (418) coupled to the cache memory and configured to monitor a current utilization of the cache memory and to determine whether the current utilization is below a predetermined utilization value; and

30 a configuration unit (419) coupled to the cache monitor unit and to the cache memory, wherein the cache configuration unit is configured to selectably disable one or more portions of the cache memory in response to the cache monitor unit determining that the current utilization is below the predetermined utilization value.

35 8. A method for dynamically configuring a processing node (12, 22) including a plurality of processor cores (15A, 15B), each processor core including a cache memory, the method comprising:

independently monitoring a respective current utilization corresponding to each of the cache memories;
determining whether the respective current utilization is below a predetermined utilization value;

in response to determining that the respective current utilization is below the predetermined utilization value selectably disabling one or more portions of the corresponding cache memory.

9. The method as recited in claim 1, further comprising determining the current utilization value by

5 determining a number of entries being used in the cache memory and determining how frequently the cache memory is accessed.

10. A processing node (12, 22) comprising:

a plurality of processor cores (15A, 15B) each including:

10 a cache memory (17A, 17B, 60);

a cache monitor unit (18A, 18B, 85, 95) coupled to the cache memory, wherein each cache monitor unit is configured to independently monitor a current utilization of the cache memory to which each cache monitor unit is coupled and to determine whether the current utilization is below a predetermined utilization value; and

15 a configuration unit (19A, 19B, 50, 65) coupled to the cache monitor unit and to the cache memory, wherein the cache configuration unit is configured to selectably disable one or more portions of the corresponding cache memory to which the configuration unit is coupled in response to the cache monitor unit determining that the current utilization is below the predetermined utilization value.

20

1/5




FIG. 1

2/5




FIG. 2

3/5




FIG. 3

4/5




FIG. 4

5/5




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/033671

A. CLASSIFICATION OF SUBJECT MATTER

G06F12/08 G06F12/12

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, IBM-TDB, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 860 106 A (DOMEN ET AL) 12 January 1999 (1999-01-12) figures 4a-c figures 8a,8b column 1, line 64 - column 2, line 3 column 2, line 53 - column 3, line 4 column 4, line 30 - line 55 column 5, line 52 - column 6, line 14 column 8, line 24 - column 9, line 23 -----	1-10
A	US 6 795 896 B1 (HART FRANK ET AL) 21 September 2004 (2004-09-21) abstract figure 2 column 1, line 38 - line 43 column 2, line 54 - column 3, line 24 ----- -/-	1-10

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

* & * document member of the same patent family

Date of the actual completion of the international search

6 February 2006

Date of mailing of the international search report

13/02/2006

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Jardon, S

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/033671

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	UNSAI O S ET AL: "Cool-cache for hot multimedia" MICROARCHITECTURE, 2001. MICRO-34. PROCEEDINGS. 34TH ACM/IEEE INTERNATIONAL SYMPOSIUM ON DEC. 1-5, 2001, PISCATAWAY, NJ, USA, IEEE, 1 December 2001 (2001-12-01), pages 274-283, XP010583691 ISBN: 0-7965-1369-7 the whole document	1-10
A	EP 1 182 567 A (TEXAS INSTRUMENTS FRANCE; TEXAS INSTRUMENTS INCORPORATED) 27 February 2002 (2002-02-27) paragraph '0066! paragraph '0080! - paragraph '0108!	1-10
A	US 5 752 045 A (CHEN ET AL) 12 May 1998 (1998-05-12) abstract column 2, line 66 - column 3, line 35	1-10
A	"POWER MANAGED SECOND-LEVEL CACHE CONTROL" IBM TECHNICAL DISCLOSURE BULLETIN, IBM CORP. NEW YORK, US, vol. 39, no. 4, 1 April 1996 (1996-04-01), pages 79-82, XP000587428 ISSN: 0018-8689 the whole document	1-10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2005/033671

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5860106	A	12-01-1999	NONE	
US 6795896	B1	21-09-2004	NONE	
EP 1182567	A	27-02-2002	US 2002065992 A1	30-05-2002
US 5752045	A	12-05-1998	NONE	