(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
6 July 2006 (06.07.2006)

PCT

(10) International Publication Number
WO 2006/072058 A1

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

(21) International Application Number:
PCT/US2005/047589

(22) International Filing Date:
27 December 2005 (27.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/023,882 27 December 2004 (27.12.2004) US

(71) Applicant *(for all designated States except US)*: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **CHRYSOS, George** [US/US]; 38 Mill Pond Circle, Milford, Massachusetts 01757 (US). **MATTINA, Matthew** [US/US]; 12 Weatherstone Dr, Worcester, Massachusetts 01604 (US).

(74) Agents: **VINCENT, Lester J.** et al.; BLAKELY SOKOLOFF TAYLOR & ZAFMAN, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025 (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, KN, 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*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

[Continued on next page]

(54) **Title:** PREDICTIVE EARLY WRITE-BACK OF OWNED CACHE BLOCKS IN A SHARED MEMORY COMPUTER SYSTEM

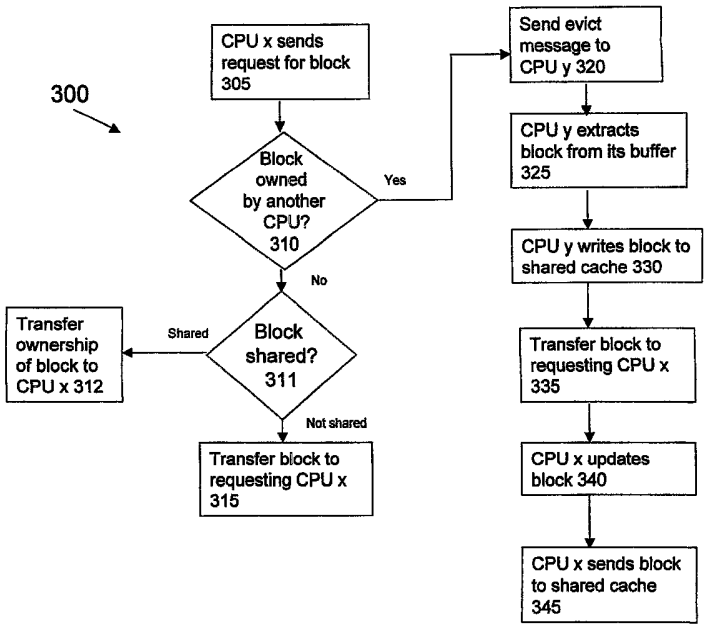

(57) **Abstract:** A method for predicting early write back of owned cache blocks in a shared memory computer system. This invention enables the system to predict which written blocks may be more likely to be requested by another CPU and the owning CPU will write those blocks back to memory as soon as possible after updating the data in the block. If another processor is requesting the data, this can reduce the latency to get that data, reducing synchronization overhead, and increasing the throughput of parallel programs.

WO 2006/072058 A1

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.

PREDICTIVE EARLY WRITE-BACK OF OWNED CACHE BLOCKS IN A SHARED MEMORY COMPUTER SYSTEM

BACKGROUND INFORMATION

[0001] Multi-processor data processing systems harness the collective computation power of multitude of processors. A memory system is central to a multi-processor system and must be scaleable in order to provide sufficient bandwidth to each processor while sharing data between the multiple processors. For certain applications, an efficient means of sharing data is critical to effective collaboration between the threads running on the multiple processors.

[0002] Most data blocks that are written by a CPU tend not to be read or written by other agents/processors shortly after ownership has been obtained. However, some blocks may be very heavily shared in short periods of time, if the blocks are being used to synchronize multiple threads of execution. Communication or synchronization latency overheads can be one factor in determining the scalability of a multi-threaded workload on a multi-processor system, like a chip multi-processor.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Various features of the invention will be apparent from the following description of preferred embodiments as illustrated in the accompanying drawings, in which like reference numerals generally refer to the same parts throughout the drawings. The drawings are not necessarily to scale, the

emphasis instead being placed upon illustrating the principles of the inventions.

[0004] Figure 1 is a block diagram of a multiprocessor system.

[0005] Figure 2 is a block diagram of a processor in the multiprocessor system of Fig. 1.

[0006] Figure 3 is a flow chart of an example method performed by the multiprocessor system of Fig. 1 to predict early write back of owned cache blocks.

[0007] Figure 4 is block diagram of a system that may provide an environment for multithreaded processors.

DETAILED DESCRIPTION

[0008] In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular structures, architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the various aspects of the invention. However, it will be apparent to those skilled in the art having the benefit of the present disclosure that the various aspects of the invention may be practiced in other examples that depart from these specific details. In certain instances, descriptions of well-known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

[0009] This invention relates to write-backs of owned cache blocks for a scalable shared memory system composed of chip multiprocessor nodes. Embodiments of the present invention may provide a semiconductor chip including processors, and a bidirectional ring interconnects to couple together the processors and the shared address space. In accordance with one

embodiment of the present invention, each processor may include multiple CPUs and the address space may include a large shared cache. Each CPU may have its own private instruction and data caches (first level caches) and all CPUs sharing a single second level cache. Cache coherence is maintained among all the caches within a chip, as well among all caches across the nodes by a protocol engine and a cache controller that may be included in each node of the system.

[0010] All specific quantities (such as number of processors, number of nodes, number of interfaces) disclosed in this document, are provided solely for purposes of explaining the operation of one particular embodiment. These quantities will typically vary, sometimes significantly, from one implementation of the invention to another.

[0011] Referring to Fig. 1, there is shown a multiprocessor system 100 including a multiplicity of processor nodes 102. Each processor node 102 is preferably implemented as a single chip multiprocessor. In a preferred embodiment, each processor node 102 may have eight processor cores 104; however, other embodiments may have two to sixteen processor cores 104. The processor cores 104, which may be comprised of a central processing unit, are processor cores since their caches and other supporting circuitry are shown separately (see Fig. 2).

[0012] Each core 104 is connected to a logically shared backing cache 106. Coupled to the shared backing cache 106 is memory 108. The memory 108 may represent any type of memory device used to store data and instructions that may have been or will be used by the processors.

[0013] Typically, though the invention is not limited in this respect, the

memory 108 may consist of dynamic random access memory (DRAM). In another embodiment the memory 108 may include a semiconductor memory. In yet another embodiment, memory 108 may include a magnetic storage device such as a disk drive. The present invention, however, is not limited to the examples of memory mentioned here.

[0014] Multiple nodes 102 may be linked by a subsystem including a router 109 and an interconnect 110. The subsystem may also include an input queue, an output queue, a packet switch (all not shown). The router109 may send and receive packets to and from other nodes via the interconnect 110. The interconnect 110 may physically link the multiple nodes 102. A system control module 111 takes care of the miscellaneous maintenance-related functions (e.g., system configuration, initialization, interrupt distribution).

[0015] While the present invention is described in the context of a multi-processor, shared memory system, the present invention is also applicable to single processor computer systems.

[0016] The embodiment of the computing system of Fig. 1 may include a plurality of processors and a plurality of caches. These processors and caches comprise a multiprocessor system in which the caches are kept coherent with one another through a cache coherence mechanism. The cache coherence protocol may be implemented in the processor caches and memory interconnects to keep the caches coherent.

[0017] Cache coherency must be addressed in multi-processor systems with shared memory. Cache coherence protocols address the issue of ensuring that no processors in the system are using stale data in the local caches. In general, stale cache entries can be eliminated by either invalidating

in the caches all but the most recent updated cache data or updating the caches with the most recent data. In a system using the invalidation protocol, an attempt to access an invalidated memory location from cache will cause the processor to read a copy of the most recent data either from another cache or from main memory. In the update protocol, following a write operation all the caches having a cached version of the data are updated with the most recent data. Thus, the most recent data is available in the caches.

[0018] Fig. 2 is a block diagram of a processor in the multiprocessor system of Fig. 1 predicting early write back of owned cache blocks. Each core 104 contains a private cache 112 which is a subset to the shared cache 106. These caches 112 may be write-through caches.

[0019] When a CPU has to do a write or a store to a memory address, the CPU needs to request ownership from the lower levels of the cache hierarchary of that block and not commit to write that block until its granted ownership. This has to occur to maintain correct cache coherency. When an ownership request is sent by CPU x, invalidate signals are driven to potential sharers, meaning other CPUs that have copies of this block. The coherence protocol will send an acknowledgment message, to CPU x that requested ownership, notifying them of their ownership of that block. After the acknowledgement message is sent, the write may now occur.

[0020] To minimize the store bandwidth that the write-through caches 112 would generate, each CPU104 may have a coalescing merge buffer 114 to aggregate recent store data to the same cache block. The merge buffer 114 may contain 16 cache blocks total. Generally, the merge buffer 114 may attempt to keep owned blocks as long as possible to maximize the store data

coalescing, and minimize write-through traffic back to the cache.

[0021] The write-back maybe caused by setting a "purgeMe" bit in the merge buffer 114 entry for which the original write request was sent. A purgeMe bit exists for every entry in the merge buffer 114. For each cycle, control logic associated with the merge buffer 114, attempts to select an entry with its purgeMe bit set, and write the block back to the shared cache 106. If no entries have their purgeMe bit set, then no entries are selected.

[0022] The purgeMe bit may be set for various reasons. One reason is because the merge buffer 114 fills up with valid entries and crosses a "purge threshold". If the merge buffer 114 crosses this purge threshold, the control logic associated with the merge buffer 114 may then set the purgeMe bit for one or more entries. By setting the purgeMe bit of the entries, they can now be written back to the shared cache 106 to make room for newly arriving committed stores. For system 100, the purgeMe bit may be set for an entry if it is predicted that the corresponding data is going to be utilized by another CPU or cache coherent agent.

[0023] Figure 3 is a flow chart of an example method 300 performed by the multiprocessor system of Fig. 1 to predict early write back of owned cache blocks. The prediction of blocks to write-back is based upon messages in the cache coherence protocol in response to the originating CPU's ownership request.

[0024] If an ownership request is sent from CPU x to the shared cache 305, the shared cache has to first determine if the requested block is owned by another CPU 310. If the requested block is not owned by another CPU, then the system determines if the requested block is shared 311 by another CPU or

may other CPUs. If the block does not have any other sharers of the block, an AckNonShared message is sent to grant ownership of the block and declare that no other CPUs had a copy of the block 315. However, if the block is not owned but shared, an AckShared message is sent and ownership is transferred to CPU x 312.

[0025] If the block was owned by another CPU, such as CPY y, an evict message is sent to CPU y 320. If the evict message is received, CPU y updates the block with its recent store data, and sets the purgeMe bit, causing the extraction of the block from its merge buffer 325. Once the block is extracted from the merge buffer, CPU y writes the block back to memory, relinquishing ownership of it 330.

[0026] The updated block is now sent to the requesting CPU x 335. CPU x may now update the block with the current data 340 and then send the block immediately to the shared cache 345. The updated block is sent immediately to the shared cache because the block of data was previously owned by another CPU, CPU y.

[0027] Now assuming CPU x needs to write the block again, another ownership request will be sent to the shared cache 305. This time, if no other CPUs have read the block in the interim, meaning block is not owned 310 and is not shared 311, an AckNonShared message is sent to CPU x 315. CPU x may now hold the block until it is evicted for capacity reasons (or is requested by another CPU).

[0028] If, however another CPU has read the block in the interim, meaning block is not owned 310 but is shared 311, the AckWasShared message is sent in to CPU x 312, and again the block is written back early.

[0029] If CPU y now requests the block, it may already have been written back to the shared cache by last CPU to request the block, or if not, it may be in transit. If the block had already been written back to the shared cache, this would reduce a 4-hop transaction to a 2-hop transaction, potentially cutting the synchronization latency in half.

[0030] Therefore, if other CPUs had (or might have had) copies of the block when the requesting CPU makes an ownership request, it writes it back as soon as possible. If other CPUs had copies of the block, but are no longer accessing the block, the first write to it may cause an expedient write-back of the block, but subsequent writes may not.

[0031] Previously, a processor did not distinguish acknowledgement of a block that was shared and one that was not shared. With the method described above, a processor may now determine how quickly it should send a block back to the shared backing cache. This enables the system to predict when fast sharing is necessary.

[0032] Advantageously, this method enables the system to predict which written blocks may be more likely to be requested by another CPU and the owning CPU will write those blocks back to memory as soon as possible after updating the data in the block. If another processor is requesting the data, this can reduce the latency to get that data, reducing synchronization overhead, and increasing the throughput of parallel programs.

[0033] Referring now to Fig, 4, the system 400 includes processors supporting predictive early write-back of owned cache blocks. The system 400 generally shows a system where processors, memory, and input/output devices are interconnected by a number of point-to-point interfaces. The system 400

may also include several processors, of which only two, processors 405, 410 are shown for clarity. Each processor 405, 410 may each include a processor core 407, 412, respectively. Processors 405, 410 may each include a local memory controller hub (MCH) 415, 420 to connect with memory 425, 430. Processors 405, 410 may exchange data via a point-to-point interface 435 using point-to-point interface circuits 440, 445. Processors 405, 410 may each exchange data with a chipset 450 via individual point-to-point interfaces 455, 460 using point to point interface circuits 465, 470, 475, 480. Chipset 450 may also exchange data with a high-performance graphics circuit 485 via a high-performance graphics interface 490.

[0034] The chipset 450 may exchange data with a bus 416 via a bus interface 495. In either system, there may be various input/output I/O devices 414 on the bus 416, including in some embodiments low performance graphics controllers, video controllers, and networking controllers. Another bus bridge 418 may in some embodiments be used to permit data exchanges between bus 416 and bus 420. Bus 420 may in some embodiments be a small computer system interface (SCSI) bus, an integrated drive electronics (IDE) bus, or a universal serial bus (USB) bus. Additional I/O devices may be connected with bus 420. These may include keyboard and cursor control devices 422, including mouse, audio I/O 424, communications devices 426, including modems and network interfaces, and data storage devices 428. Software code 430 may be stored on data storage device 428. In some embodiments, data storage device 428 may be a fixed magnetic disk, a floppy disk drive, an optical disk drive, a magneto-optical disk drive, a magnetic tape, or non-volatile memory including flash memory.

[0035] In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular structures, architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the various aspects of the invention. However, it will be apparent to those skilled in the art having the benefit of the present disclosure that the various aspects of the invention may be practiced in other examples that depart from these specific details. In certain instances, descriptions of well-known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

WHAT IS CLAIMED IS:

1. A method comprising:

requesting a block by a first processor;

determining if a second processor has the requested block;

transferring block to the first processor if the second processor does not have the requested block; and

immediately writing updated block to a memory by the first processor.

2. The method of claim 1 wherein the transferring block to the first processor further comprises sending an first signal to the first processor.

3. The method of claim 1 wherein the determining if the second processor has the requested block further comprises:

determining if the block is owned by the second processor; and

determining if the block is shared by the second processor, if the block is not owned by the second processor.

4. The method of claim 3, wherein the determining if the block is shared by the second processor further comprises sending an second signal to the first processor.

5. The method of claim 3 wherein the determining if the block is owned by the second processor further comprises sending an evict signal to the second processor.

6. The method of claim 5 further comprises updating block with recent stored data by the second processor.

7. The method of claim 6 further comprising:

setting a bit for an entry of the block in a buffer; and

extracting the updated block from the buffer.

8. The method of claim 7 further comprises writing the updated block to a memory by the second processor.

9. The method of claim 8 further comprising:

sending the updated block to the first processor by the memory; and

updating the updated block by the first processor.

10. The method of claim 9 further comprising:

requesting the updated block again by the first processor;

determining if the second processor has requested the updated block since the updated block was sent to the memory by the first processor; and

transferring updated block to the first processor.

11. The method of claim 10 further comprising sending the first signal to the first processor if the second processor has not requested the updated block since the updated block was sent to memory by the first processor.

12. The method of claim 11 further comprising:

updating the requested updated block by the first processor; and

holding the new updated block by the first processor.

13. The method of claim 10 further comprising sending a third signal to the first processor if the second processor had requested the updated block since the updated block was sent to the memory.

14. The method of claim 13 further comprising:

updating the requested updated block by the first processor; and

immediately writing the new updated block to the memory.

15. A first and second processors comprising:

a shared cache coupled to the first and second processors, wherein the shared cache determines ownership of a block of data;

a buffer located within each of the first and second processor, wherein the buffer stores most recent data for the block; and

wherein the first processor requests the block, if the block is shared by the second processor, the first processor obtains the block, updates the block and immediately writes the updated block to the shared cache.

16. The first and second processors of claim 15 wherein the buffer includes a bit for every block within the buffer.

17. The first and second processor of claim 16 wherein the block in the buffer is extracted when the bit is set.

18. The first and second processors of claim 15 wherein if the block requested by the first processor is shared by the second processor, the shared cache sends the second signal to the first processor.

19. The first and second processors of claim 15 wherein if the block requested by the first processor is not shared by the second processor, the shared cache sends the first signal to the first processor.

20. The first and second processors of clam 19 further comprising holding the updated block by the first processor.

21. A multiprocessor computer system comprising a plurality of nodes, each node from said plurality of nodes comprising:

a router to communicate with the plurality of nodes in the multiprocessor system;

at least first and second processors, wherein the first and second processor comprising:

a shared cache coupled to the first and second processors, wherein the shared cache determines ownership of a block of data,

a buffer located within each of the first and second processor,

wherein the buffer stores most recent data for the block, and

wherein the first processor requests the block, if the block is shared by the second processor, the first processor obtains the block, updates the block and immediately writes the updated block to the shared cache,

a first interface to input/output devices; and

an audio input/output device coupled to the first interface.

22. The multiprocessor system of claim 21 wherein the buffer in the first and second processors includes a bit for every block within the buffer.

23. The multiprocessor system of claim 22 wherein the block in the buffer is extracted when the bit is set.

24. The multiprocessor system of claim 21 wherein if the block requested by the first processor is shared by the second processor, the shared cache sends the second signal to the first processor.

25. The multiprocessor system of claim 21wherein if the block requested by the first processor is not shared by the second processor, the shared cache sends the first signal to the first processor.

26. The multiprocessor system of claim 25 further comprising holding the updated block by the first processor.

27. The multiprocessor system of claim 21 wherein the first interface is a point to point interface.

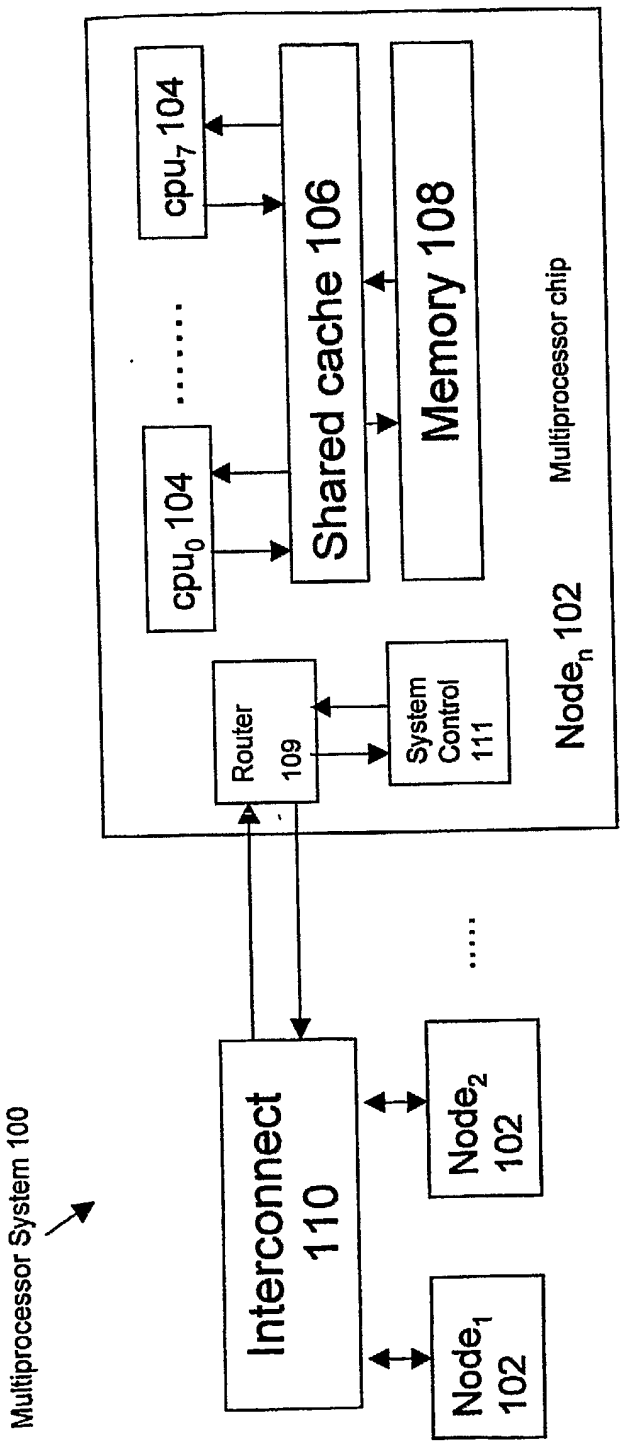
Multiprocessor System 100
Interconnect 110
Node1 102
Node2 102
.....
Noden 102
Multiprocessor chip
Router 109
System Control 111
cpu0 104
cpu7 104
.....
Shared cache 106
Memory 108

Fig. 1

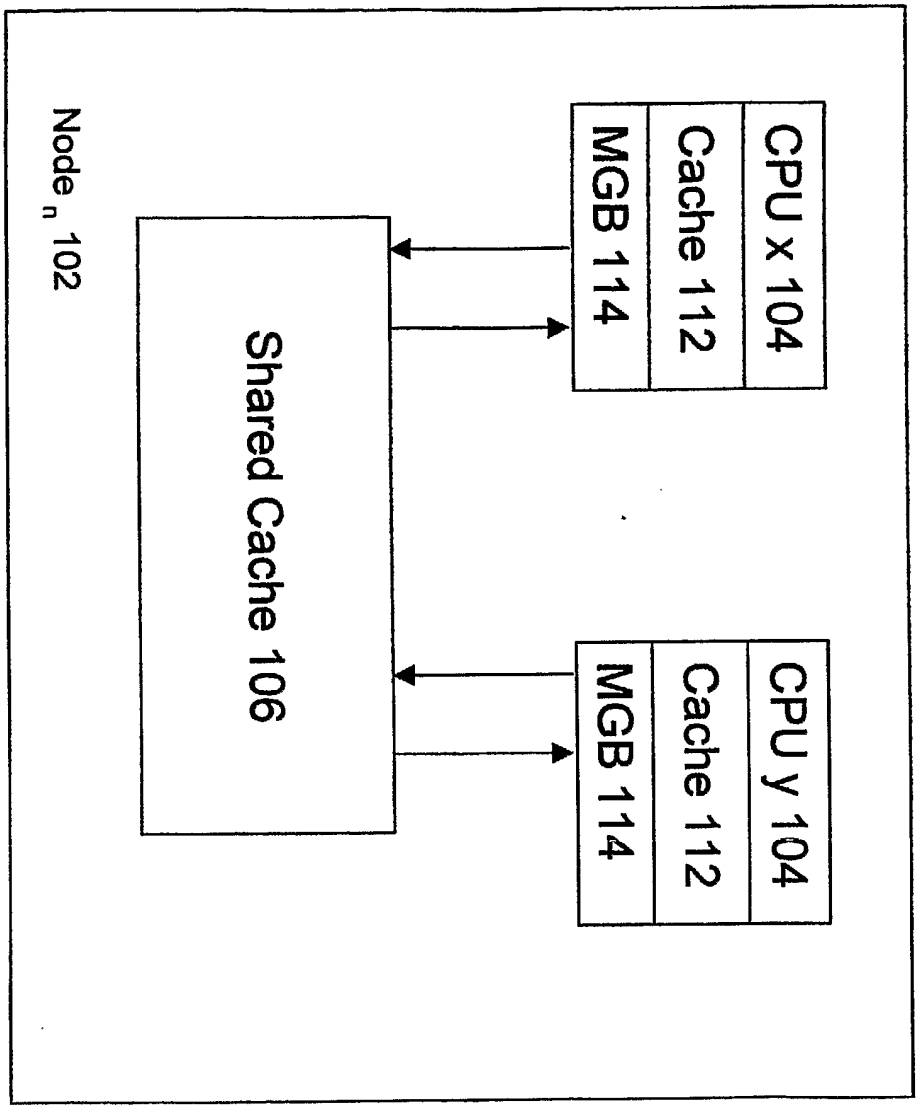
Node n 102
CPU x 104
Cache 112
MGB 114
Shared Cache 106
CPU y 104
Cache 112
MGB 114

Fig. 2

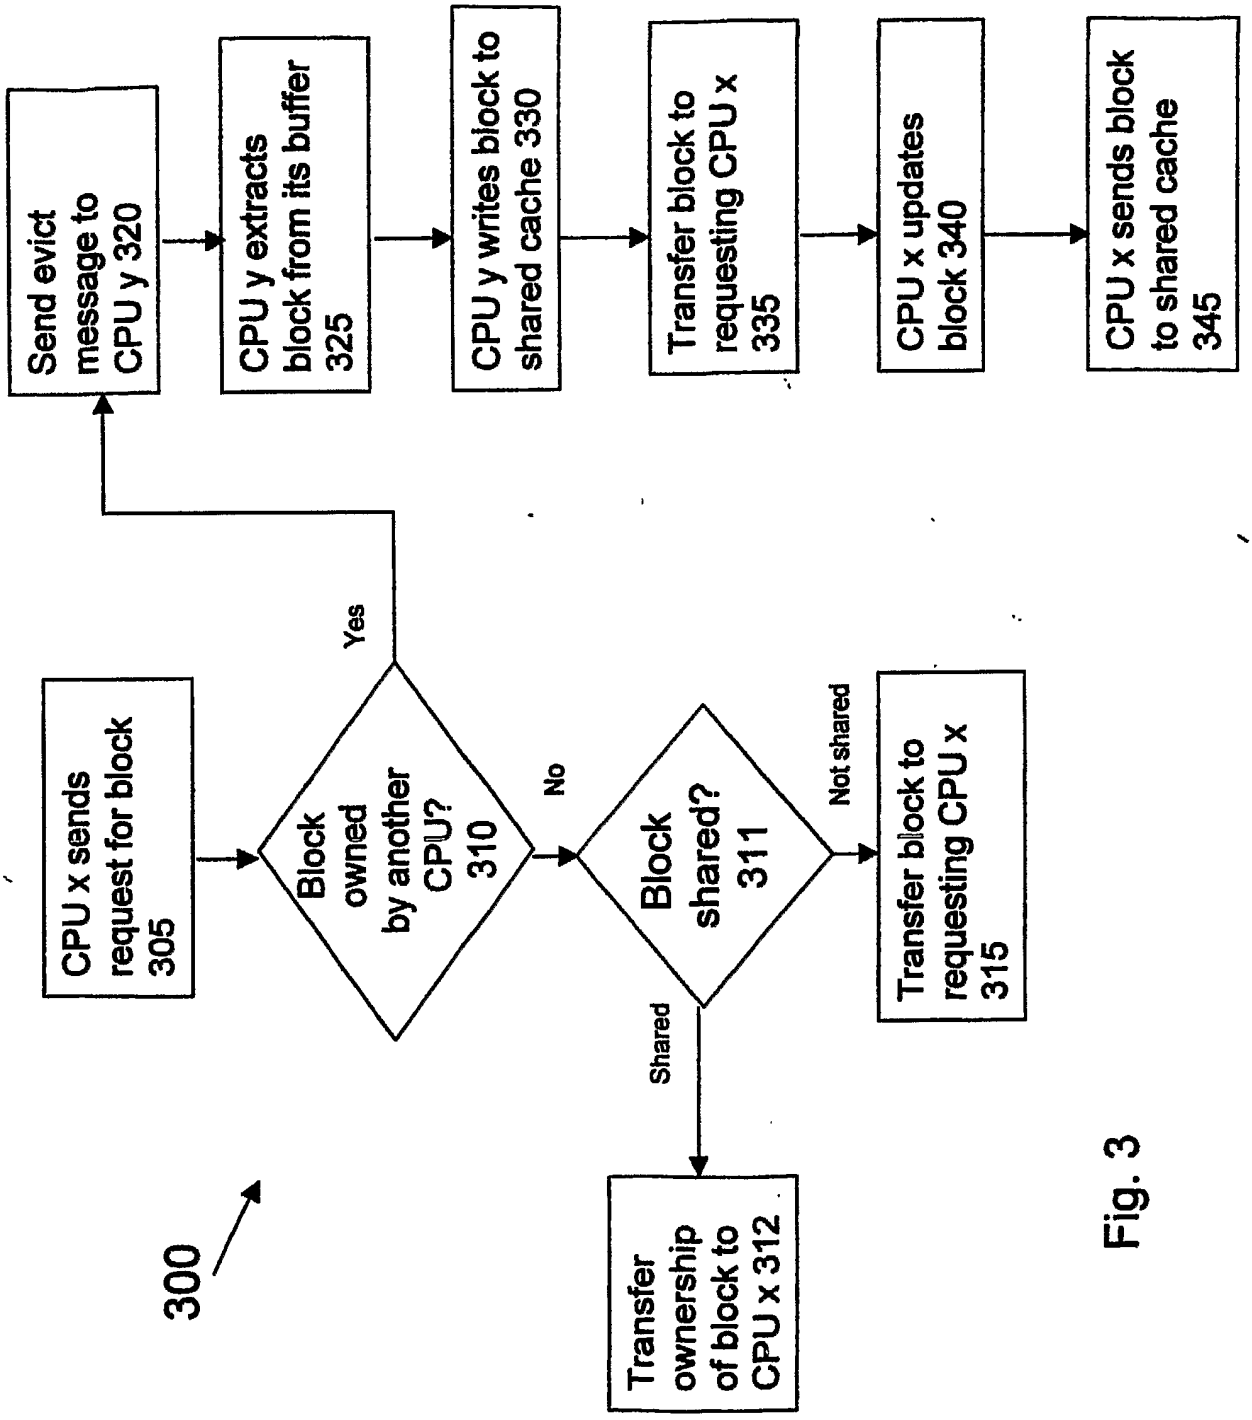
CPU x sends request for block 305
Block owned by another CPU? 310
Yes
No
Send evict message to CPU y 320
CPU y extracts block from its buffer 325
CPU y writes block to shared cache 330
Transfer block to requesting CPU x 335
CPU x updates block 340
CPU x sends block to shared cache 345
Block shared? 311
Not shared
Transfer block to requesting CPU x 315
Shared
Transfer ownership of block to CPU x 312
300

Fig. 3

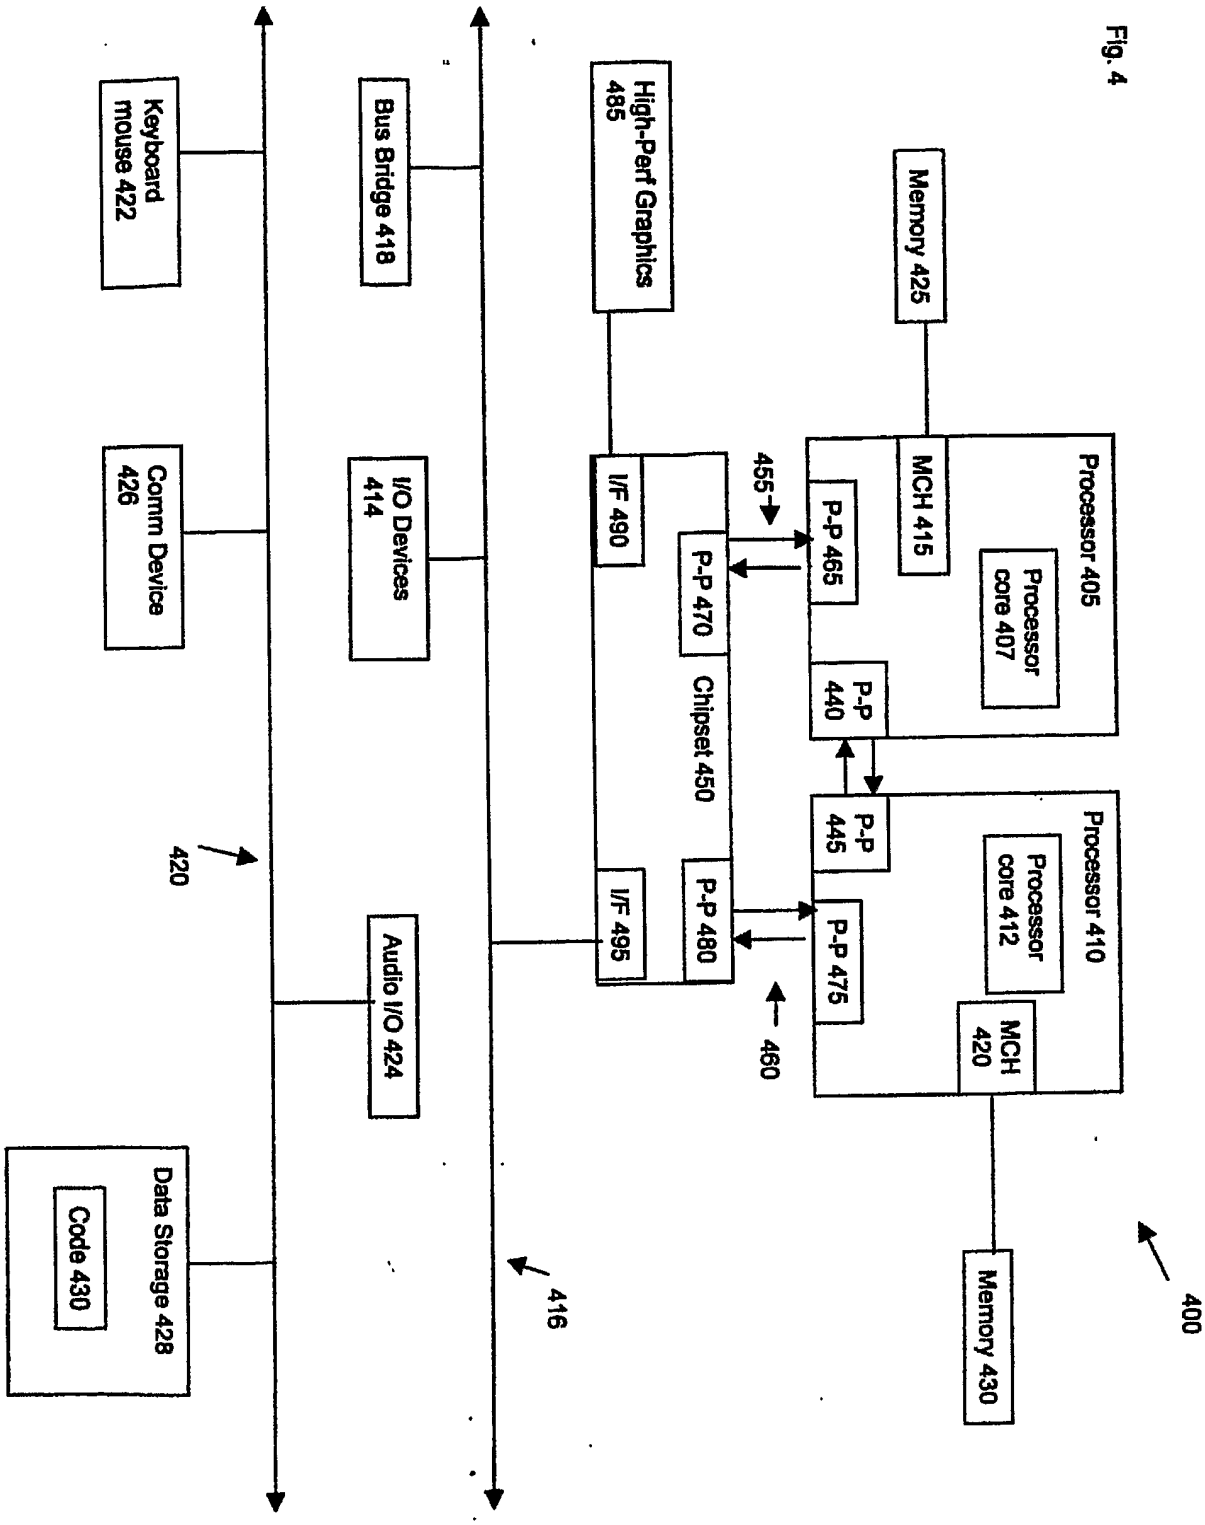
Memory 425
Processor 405
MCH 415
Processor core 407
P-P 465
P-P 440
455
High-Perf Graphics 485
I/F 490
P-P 470
Chipset 450
P-P 445
Processor 410
Processor core 412
MCH 420
P-P 475
I/F 495
P-P 480
460
Memory 430
400
416
420
Keyboard mouse 422
Bus Bridge 418
Comm Device 426
I/O Devices 414
Audio I/O 424
Data Storage 428
Code 430

Fig. 4

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/047589

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F12/08

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/124144 A1 (GHARACHORLOO KOUROSH ET AL) 5 September 2002 (2002-09-05) figures 1,10A paragraph [0006] paragraph [0050] paragraph [0061] paragraph [0138] paragraph [0141] paragraph [0150]	1-27
A	US 5 926 829 A (HAGERSTEN ET AL) 20 July 1999 (1999-07-20) the whole document	1-27

☐ 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	Date of mailing of the international search report
28 April 2006	08/05/2006

Name and mailing address of the ISA/	Authorized officer
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	Filip, L

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2005/047589

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2002124144 A1	05-09-2002	US 2002129208 A1	12-09-2002
		US 2002083274 A1	27-06-2002
		US 2002087806 A1	04-07-2002
		US 2002087807 A1	04-07-2002
		US 2003023814 A1	30-01-2003
		US 2002007443 A1	17-01-2002
US 5926829 A	20-07-1999	EP 0780769 A1	25-06-1997
		JP 9185549 A	15-07-1997
		SG 74577 A1	22-08-2000
		US 5710907 A	20-01-1998