



- (51) **International Patent Classification:**
G06F 9/06 (2006.01) **G06F 12/08** (2006.01)
G06F 9/30 (2006.01)
- (21) **International Application Number:**
PCT/US2013/048362
- (22) **International Filing Date:**
27 June 2013 (27.06.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/631,625 28 September 2012 (28.09.2012) US
- (71) **Applicant:** INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, M/S: RNB-4-150, Santa Clara, California 95054 (US).
- (72) **Inventors; and**
- (71) **Applicants :** MOSES, Jaideep [IN/US]; 13580 NW Stonebridge Drive, Portland, Oregon 97229 (US). IYER, Ravishankar [US/US]; 15934 NW Andalusian Way, Portland, Oregon 97229 (US). ILLIKKAL, Ramesh [US/US]; 4542 NW Silverleaf Drive, Portland, Oregon 97229 (US). SRINIVASAN, Sadagopan [IN/US]; 1470 NE Alex Way, Hillsboro, Oregon 97124 (US).
- (74) **Agents:** VINCENT, Lester J. et al.; Blakely Sokoloff Taylor & Zafman, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report (Art. 21(3))

(54) **Title:** APPARATUS AND METHOD FOR REDUCING THE FLUSHING TIME OF A CACHE




FIG. 1

(57) **Abstract:** A processor is described having cache circuitry and logic circuitry. The logic circuitry is to manage the entry and removal of cache lines from the cache circuitry. The logic circuitry includes storage circuitry and control circuitry. The storage circuitry is to store information identifying a set of cache lines within the cache that are in a modified state. The control circuitry is coupled to the storage circuitry to receive the information from the storage circuitry, responsive to a signal to flush the cache, and determine addresses of the cache therefrom so that the set of cache lines are read from the cache so as to avoid reading cache lines from the cache that are in an invalid or a clean state.

APPARATUS AND METHOD FOR REDUCING THE FLUSHING TIME OF A CACHEField of Invention

The field of invention relates generally to a computing systems, and, more specifically, to an apparatus and method for reducing the flushing time of a cache.

5 Background

Fig. 1 shows the architecture of an exemplary multi-core processor 100. As observed in Fig. 1, the processor includes: 1) multiple processing cores 101_1 to 101_N; 2) an interconnection network 102; 3) a last level caching system 103; 4) a memory controller 104 and an I/O hub 105. Each of the processing cores contain one or more instruction execution
10 pipelines for executing program code instructions. The interconnect network 102 serves to interconnect each of the cores 101_1 to 101_N to each other as well as the other components 103, 104, 105. The last level caching system 103 serves as a last layer of cache in the processor before instructions and/or data are evicted to system memory 106.

The memory controller 104 reads/writes data and instructions from/to system memory
15 106. The I/O hub 105 manages communication between the processor and “I/O” devices (e.g., non volatile storage devices and/or network interfaces). Port 107 stems from the interconnection network 102 to link multiple processors so that systems having more than N cores can be realized. Graphics processor 108 performs graphics computations. Power management circuitry 109 manages the performance and power states of the processor as a whole (“package level”) as
20 well as aspects of the performance and power states of the individual units within the processor such as the individual cores. Other functional blocks of significance (e.g., phase locked loop (PLL) circuitry) are not depicted in Fig. 1 for convenience.

The last level caching system 103 includes multiple caching agents 113_1 through 113_Z. Each caching agent is responsible for managing its own respective “slice” of cache
25 114_1 through 114_Z. According to one implementation, each system memory address in the system uniquely maps to one of the cache slices 114_1 - 114_Z. According to this particular implementation, a memory access from any of the processing cores will be directed to only one of the cache agents 113_1 - 113_Z based on a hash of the memory address.

Each cache agent is not only responsible for delivering a cache line to the requesting core
30 if there is a hit in its respective slice, but also, forward a request from a core to the memory controller 104 if there is a cache miss. Each cache agent is also responsible for implementing a cache coherence protocol (e.g., the MESI protocol or similar protocol) to ensure that the processing cores are not using stale data. Of course other processor and/or caching architectures than the particular core observed in Fig. 1 and discussed just above are possible.

As the power consumption of computing systems has become a matter of concern, most present day systems include sophisticated power management functions. A common framework is to define both “performance” states and “power” states. The entry and/or departure from any one of these states may be controlled, for example, by power management circuitry 109. A processor’s performance is its ability to do work over a set time period. The higher a processor’s performance the more work it can do over the set time period. A processor’s performance can be adjusted during runtime by changing its internal clock speeds and voltage levels. As such, a processor’s power consumption increases as its performance increases.

Thus, a processor’s different performance states correspond to different clock settings and internal voltage settings so as to effect a different performance vs. power consumption tradeoff. According to the Advanced Configuration and Power Interface (ACPI) standard the different performance states are labeled with different “P numbers”: P0, P1, P2 . . . P_R, where, P0 represents the highest performance and power consumption state and P_N represents the lowest level of power consumption that a processor is able to perform work at. The term “R” in “P_R” represents the fact that different processors may be configured to have different numbers of performance states.

In contrast to performance states, power states are largely directed to defining different “sleep modes” of a processor. According to the ACPI standard, the C0 state is the only power state at which the processor can do work. As such, for the processor to enter any of the performance states (P0 through P_R), the processor must be in the C0 power state. When no work is to be done and the processor is to be put to sleep, the processor can be put into any of a number of different power states C1, C2 . . . C_S where each power state represents a different level of sleep and, correspondingly, a different amount of time needed to transition back to the operable C0 power state. Here, a different level of sleep means different power savings while the processor is sleeping.

A deeper level of sleep therefore corresponds to slower internal clock frequencies and/or lower internal supply voltages and/or more blocks of logic that receive a slower clock frequency and/or a lower supply voltage. Increasing C number corresponds to a deeper level of sleep. Therefore, for instance, a processor in the C2 power state might have lower internal supply voltages and more blocks of logic that are turned off than a processor in the C1 state. Because deeper power states corresponds to greater frequency and/or voltage swings and/or greater numbers of logic blocks that need to be turned on to return to the C0 state, deeper power states also take longer amounts of time to return to the C0 state.

A problem exists with respect to the size of the last level caching system 103 and sleep states when the last level cache is to be flushed. For example, certain “package level” power

states may reduce the supply voltage to the last level caching system 103 requiring that its cached information be saved to external system memory 106 beforehand. As last level cache sizes are becoming quite large, too much time is being expended flushing the last level cache 103 of its data when entering a sleep state that requires the last level cache to be flushed.

5 Currently, respective state machines in the cache agents of processors designed by Intel Corporation of Santa Clara, California use a WriteBackINValiDate (WBINVD) operation to effectively scroll through every location in every cache slice to flush the cache. When each cache line is read, a “dirty bit” that is kept within the cache line indicates whether the cache line has been modified or not. If it has been modified the cache line is saved externally from the
10 cache (e.g., to system memory). The time spent accessing each location consumes too much time and is becoming a performance bottleneck for sleep state entry as cache sizes increase.

Brief Description of the Drawings

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

15 Fig. 1 (prior art) shows a traditional processor design;

 Fig. 2 shows a methodology for reducing the flushing time of a cache;

 Fig. 3 shows a circuit design for reducing the flushing time of a cache;

 Fig. 4 shows a processor design that includes circuitry for reducing the flushing time of a cache;

20 Fig. 5a shows an architecture for memory devices used to keep track of which cache lines in a cache are in a modified state;

 Fig. 5b shows write and read processes to and from the memory devices of Fig. 5a;

 Fig. 6 shows a memory device approach having a smaller footprint than the approach observed in Fig. 5a.

Detailed Description

As is understood in the art, each cache line in a cache can be characterized as being clean, invalid or modified (“dirty”). If a cache line is clean, there is an identical copy of it in system memory. If a cache line is invalid it is not to be used (e.g., because a more up-to-date version of the cache line exists elsewhere in the system). If a cache line is modified it is understood to be
30 the most up-to-date version of the cache line and has no duplicate in system memory. At any instant of time, the various cache lines of a last level cache are apt to include a mixture of clean, invalid and modified cache lines.

It is pertinent to recognize that, technically speaking, only the modified cache lines need actually be saved externally in order for the system to continue operation without error. That is,
35 the clean cache lines need not be saved externally because a copy of each of them can be found

in system memory, and, the invalid cache lines need not be externally saved because by definition they are not to be used. As such, in order for a computing system to flush its cache and continue operation without error, only the cache's modified cache lines need to be saved externally – e.g., to system memory – before the cache is powered down. Upon entry to a higher
5 power state where the cache is powered up, the system need not even attempt to re-populate the cache as part of the entry to the higher power state and can simply operate out of system memory. The cache will begin to be populated with cache lines over the course of the ensuing runtime.

The prior art WBINVD process operates according to these principles. However,
10 whether a cache line is modified or not (e.g., a dirty bit) is kept within the cache line itself. Therefore, each cache line needs to be read from the cache according to the prior art WBINVD process to understand whether it is modified. This “walkthrough” of each cache line consumes too much time as cache sizes are becoming very large.

As such, Fig. 2 shows an improved process for managing entry to a power state that
15 requires a cache flush. According to the process of Fig. 2, before entry into the power state is envisioned (e.g., during normal runtime when the cache is fully powered and operational), logic circuitry associated with the cache monitors the accesses made to the cache and tracks not only how many cache lines are modified, but also, where the modified cache lines exist in the cache
201.

Subsequently, a decision is made to enter a power state that requires the cache to be
20 flushed (e.g., because it is being powered down) 202. In response to the decision that the power state is being entered, the logic circuitry refers to the tracked data and attempts to access the modified cache lines and save them externally to system memory. By so doing, the cache circuitry purposefully avoids reading clean or invalid data from the cache 203 only to discover
25 that they are not modified. Said another way, by pinpointing the accesses to cache for only the modified cache lines, the walk through time of reading non modified cache lines is avoided which greatly speeds up the process by which dirty cache lines are identified. This “cuts down” the total number of accesses made to the cache during the walkthrough process as compared to the prior art approach.

Fig. 3 shows reduced cache flush time logic circuitry 330 designed to perform the
30 methodology described above. The logic circuitry of Fig. 3 may be embedded in logic circuitry, such as cache agent logic circuitry 313, that includes logic circuitry that identifies cache line states and reads/writes from/to a cache or cache slice 315. According to the design of Fig. 3, the cache read/write logic 316 is enhanced with: 1) increment logic 317 that increments a counter
35 321 each time a cache line is written to so as to cause the cache line to transition to a modified

state (e.g., from a clean state); 2) decrement logic 318 that decrements the counter 321 each time a modified cache line is evicted or declared invalid; and, 3) location tracking logic 319 that tracks where the cache lines within the cache 314 having a modified state reside within the cache. Here, cache 314 may be an entire cache or a segment of cache such as a cache slice. For simplicity the term cache will be used to refer to an entire cache or a segment of a cache.

In an embodiment, the increment logic 317 is further designed to ignore multiple writes to a same cache line. That is, once a cache line location is written to, the increment logic will not increment the counter if that same cache line is written to again. Here, two cache lines are deemed to be the “same” if they have the same system memory address. As such, multiple writes to the same cache line will not falsely increase the count of modified cache lines within the cache. In the situation where a modified cache line is evicted or declared invalid and then overwritten or otherwise replaced with a same or different cache line, the increment logic 317 will increment again because the counter will have been properly decremented with the eviction and/or invalidation.

In an embodiment, the location tracking logic 319 is coupled to one or more memories 323 or other type of storage circuitry that store information that identifies where the cache lines having a modified state reside within the cache. State machine logic circuitry 320 is also coupled to power management circuitry 309. Power management circuitry 309 recognizes a power state that requires the flushing of the cache and sends a signal along communication wiring 321 to state machine logic 320 indicating that the cache is to be flushed.

In response, in an embodiment, the state machine logic 320 refers to: 1) the counter 321 to understand how many modified cache lines exist within the cache 314; and, 2) location tracking circuitry 319 (which refers to storage circuitry 323) to understand where these cache lines are located within the cache 314. The state machine logic circuitry 320 proceeds to cause the modified cache lines to be read from the cache by presenting specific read addresses to read/write circuitry 316, where, the addresses correspond to cache lines in the modified state in the cache 314. (The term “tag” may also be used to refer to an address for a cache. For convenience the term “address” will be used). As such, largely if not entirely, only modified cache lines are read from the cache 314 and saved externally from it. By avoiding accesses to clean and invalid cache lines, the total time spent accessing the cache 314 during the flushing process should be noticeably reduced compared to prior art approaches that scanned the entire cache.

When the state machine 320 determines the flushing of the cache is complete it notifies the power management unit by way of communication wiring 321. In response the power management circuitry 309 sends a signal to voltage control circuitry 322 which, in turn, powers

down the cache 314. Here, the unread clean and invalid cache lines are “lost” from the cache 314 as a consequence of the power down, but the system, as discussed above, does not need these cache lines to continue proper operation going forward. State machine 320 could alternatively be implemented with a controller, micro-controller or other type of circuit, all of which, including a state machine, may be referred to as a control circuit. More globally, any of the functions of the reduced flush time logic 330 could be implemented with executable code (e.g., by a micro-controller), dedicated logic circuitry or a mixture thereof.

As is understood in the art, different types of caches exist. On one end of the spectrum is a “direct mapped” cache in which any particular cache line in system memory can be stored in only one section of the cache. The section is specified by an index value. Said another way, each cache line in system memory can only have one associated index. The cache is divided up into different sections each having its own unique index value. Two different cache lines can map to different or same index values depending on the particular address of each and the (e.g., hashing) algorithm used to determine the index from the address.

On the other end of the spectrum is a fully associative cache. In the case of a fully associative cache, a cache line can be mapped to any cache location. In a sense, any cache line can map to any index value.

A set associative cache is a kind of hybrid between these two extremes. In the case of at least one type of set-associative cache, a cache line in system memory can map to a “set” but the set has more than one “way” where the cache line can be stored. Thus, for example, if each set in the cache has two ways, a cache line in system memory can map to two different locations (ways) within the set that the cache line maps to.

Fig. 4 shows a more detailed implementation of a processor having a set-associative last level cache 403. The processor 400 of Fig. 4 includes the same elements as the processor 100 of Fig. 1, with the addition of reduced cache flush time logic 430_1 to 430_Z in each of the cache agents 413_1 to 413_Z. According to the embodiment of Fig. 4, each cache slice 414_1 to 414_Z is composed of Q “sets”, each set having M “ways” of storage locations. According to one embodiment, each system memory address maps to a particular set within a particular cache slice unit. Each cache line within system memory can be stored in any of the ways of the particular set that it maps to.

Consistent with the discussion of Figs. 2 and 3 above, the cache flush time reduction circuitry 430_1 to 430_Z of each of the cache agents 413_1 to 413_Z track which ones of their stored cache lines is in the modified state, and, upon receiving a cache flush command from the power management unit 409, reads the modified cache lines but refrains from reading invalid or clean cache lines from its cache slice prior to power down of the cache slice. The “saved”

modified cache lines are written elsewhere, such as within a system memory 406 that is external to the processor 400. The cache flush time reduction circuitry 430_1 to 430_Z of each cache agent includes storage circuitry to track where the modified cache lines within its respective cache slice reside. In an embodiment, each of the storage circuits record where each of its
 5 respective cache slice's modified cache lines are stored in terms of a specific set and specific way.

Fig. 5a shows storage circuit 523_1 to 523_M structured as M different storage matrices where M is the number of ways associated with each set in a caching system. Storage circuits 523_1 to 523_M may be viewed as an embodiment of storage circuit(s) 323 of Fig. 3. The size
 10 of the matrices are defined by X rows and Y columns where $X \times Y = Q$ = the number of sets within the cache slice. Each storage location in the matrix need only store one bit to indicate whether the cache line corresponding to the location is in a modified state. Other bits could be kept at each location to store, for instance, additional information for each cache line in the cache slice.

Here, particularly with a large cache, the number of sets stored by a cache slice is apt to be much larger than the number of ways established for each set. For example, a possible implementation for a cache slice would include storage resources for 8,192 sets each having 16 associated ways. As such, the storage circuitry 523_1 to 523_M maintained by the cache agent for the cache slice would include $M = 16$ matrices each having storage capacity for 8,192
 20 different locations. In a further embodiment, the 8,192 storage locations of each matrix are identifiable through one of 128 rows and one of 64 columns (i.e., $X = 128$; $Y = 64$; $X \times Y = 128 \times 64 = 8192 = Q$ = number of sets kept by the cache slice). Note that, in the case of 64 byte cache line sizes, this architecture corresponds to a cache slice size of $8192 \times 16 \times 64 = 8,388,608$ megabytes (MB) which is large compared to traditional designs.

It should be apparent that stepping through such a large number of locations to individually determine the modified state of each cache line would consume a noticeable amount of time and dampen system performance owing to the time wasted waiting for the cache to be flushed. The prior art WBINVD approach effectively steps through each cache line location to flush the cache.

As such, in order to reduce "walkthrough" scanning times, the storage circuits 523_1 to 523_M are architected such that one dimension (e.g., the 128 rows along the X dimension) corresponds to addresses of different memory words of a particular storage circuit 523, and, the other dimension of the matrix (e.g., the 64 columns along the Y dimension) corresponds to different bit positions in a memory word of the particular storage circuit 523. In this case, each
 35 storage circuit matrix 523_1 to 523_M could be realized with a 1024 KB memory having 128

separately addressable 64 bit words. Each time a word from the memory is read, the modified state of 64 different cache lines is provided. This has the effect of reducing the scan time by a factor of 64 as compared to a process that must read each cache line individually. Of course, some additional time will be spent actually fetching the modified cache lines from the cache slice
5 once their location is known. However, the reading of the storage circuits 523_1 through 523_M could be performed, at least partially, in parallel with the fetching of modified cache lines from the cache slice.

Fig. 5b shows a more detailed embodiment of how the modified tracking logic circuitry 319 of Fig. 3 may operate with the storage circuitry 523_1 – 523_M of Fig. 5. Fig. 5b shows
10 both a write process and a read process. According to the write process, a determination is made 550 that a cache line is changing state to a modified state (e.g., from a clean state) or is otherwise recognized as being within a modified state. The detection may be made, for instance, in response to a determination that a clean cache line is being overwritten in the cache with new data. In the embodiment of Fig. 3, the determination may be made by circuitry 315 and
15 forwarded to tracking circuitry 319. In response to the determination that a cache line is in a modified state, the cache line's address is analyzed 551 to determine an address and word location (e.g., bit position) of a memory word in a memory. The address and bit position correspond to the cache line's location in the cache. A write operation is then performed 552
20 into the memory with the address to set a bit in the bit position of the memory word to indicate that the cache line is in a modified state. Concurrently, the cache line itself may be stored in the cache.

Subsequently another similar process could be used to clear the same bit in response to a later detection that the same cache line was later invalidated.

According to the read process of Fig. 5b, a signal is received (e.g., from a state machine
25 and/or power control unit) to flush a cache 560. In response, the memory is read at a first address 561. The memory word that is returned is analyzed 562 to determine which of the multiple different cache lines in the cache represented by the word are presently in a modified state. The cache lines that are determined to be in a modified state are then reported to the state machine 320 which fetches them from the cache. The next address in the memory is determined
30 and the corresponding word analyzed to repeat the process. Eventually the memory is fully read and the appropriate cache lines are read from the cache. If a cache agent has multiple instances of the memory (e.g., one memory for each of M ways in the cache agent's cache slice as observed in Fig. 5a), the methodologies of Fig. 5b can be performed for each memory associated with the cache agent. At least write processes to different memories can be performed in parallel
35 and read memories from different memories can be performed in parallel.

Those of ordinary skill will appreciate that other architectures are possible. For example, in a more centralized approach, a single memory could be instantiated at a single cache agent rather than multiple memories. In an even further centralized approach, a single memory could be used for an entire cache rather than just a cache slice. Also, the memories themselves may be architected in various ways. For instance, word lengths may be greater or less than 64 bits. Registers may be used as storage cells. For purposes of this application the term “memory” may be used to refer to a random access memory or a plurality of registers.

The exemplary embodiment discussed above with respect to Fig. 5a and 5b could also be provided with an additional potential speed up by storing additional words to each of the storage circuits 523_1 to 523_M that indicate which of the memory words in the memory do not have any indication of a modified cache line. For example, recall that the exemplary memory matrices included 128 separately addressable words of 64 bits each. Each word represents 64 different cache lines in the cache slice. If the size of each memory circuit was expanded to include 130 separately addressable words rather than 128 separately addressable words, the two extra words could be used to indicate which of the 128 words do not contain any indication of a modified cache line.

For example, the 64 bits of the first extra word could be used to indicate which of separately addressable words 1 through 64 do not contain any indication of a modified cache line, and, the 64 bits of the second extra word could be used to indicate which of separately addressable words 65 through 128 do not contain any indication of a modified cache line. Upon an indication of a need to flush the cache slice, the state machine first reads the two extra words to understand which, if any, of separately addressable words 1 through 128 do not contain any indication of a modified cache line. In response to digesting this information, the state machine then proceeds, in scanning separately addressable words 1 through 128, to skip over any of those words that do not contain any indication of a cache line that is in a modified state. If more than two words are skipped over, a speed up is realized (owing to the read of the extra two words).

Here, the read process of Fig. 5b would have additional procedures to write to the appropriate extra word when a word in the memory is written to so as to indicate a cache line is in a modified state. Likewise, the write process of Fig. 5b would have additional procedures to initially read the extra words to understand which if any of the normally scanned words could be eliminated. The number of extra words could of course vary from embodiment to embodiment. In an embodiment, the processor is designed to enable the use of the extra words at the option of the user. For example, software may write to model specific register space (MSR) to indicate whether the extra words and associated extra processes are to be enabled. The setting may be influenced by the type of application software program being run and its associated data (e.g.,

the extra words are enabled if the data is expected to have large portions of cached cache lines that are not in a modified state).

Note that the approaches discussed above for Figs 5a and 5b do not require the counter 321 of Fig. 3b to actually be utilized. That is, the state machine 320 can determine the number of cache lines that are in a modified state from the memory circuit(s) 323/523. In an embodiment the counter 321 is nevertheless utilized as a check against the information stored in the memory circuit(s) 323/523. If a discrepancy arises between the counter value 321 and what is gleaned from the memory circuit(s) 323/523, the state machine may raise an interrupt or other type of flag to signify some kind error.

Fig. 6 shows another alternative approach that attempts to reduce the size of the footprint of memory circuit(s) 323. That is, the approach of Fig. 6 is designed to reduce the size of the memory 623 used to keep track of which cache lines are in a modified state. Here, some accuracy is sacrificed compared to the approach described just above in that there is only one bit reserved to represent the modified/not modified status of multiple cache lines in the cache. Hash logic 601 implements a hashing algorithm to map different addresses of different cache lines that are in a modified state to a single bit within memory 623. Here, if any cache lines in a group that map to a same bit position in the memory 623 change state to a modified state or are otherwise recognized as being in the modified state, the bit representing the group is set in memory 623 to indicate a modified state. This means that upon a cache flush operation, all cache lines in the group will be read from the cache before power down of the cache even though some of them may not be in a modified state.

As any of the logic processes taught by the discussion above may be performed with a controller, micro-controller or similar component, such processes may be program code such as machine-executable instructions that cause a machine that executes these instructions to perform certain functions. Processes taught by the discussion above may also be performed by (in the alternative to the execution of program code or in combination with the execution of program code) by electronic circuitry designed to perform the processes (or a portion thereof).

It is believed that processes taught by the discussion above may also be described in source level program code in various object-orientated or non-object-orientated computer programming languages. An article of manufacture may be used to store program code. An article of manufacture that stores program code may be embodied as, but is not limited to, one or more memories (e.g., one or more flash memories, random access memories (static, dynamic or other)), optical disks, CD-ROMs, DVD ROMs, EPROMs, EEPROMs, magnetic or optical cards or other type of machine-readable media suitable for storing electronic instructions. Program code may also be downloaded from a remote computer (e.g., a server) to a requesting computer

(e.g., a client) by way of data signals embodied in a propagation medium (e.g., via a communication link (e.g., a network connection)).

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and
5 changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

Claims

1. A processor, comprising:

a) cache circuitry;

b) logic circuitry to manage the entry and removal of cache lines from said cache

5 circuitry, said logic circuitry, comprising:

i) storage circuitry to store information identifying a set of cache lines within said cache that are in a modified state;

10 ii) control circuitry coupled to said storage circuitry to receive said information from said storage circuitry, responsive to a signal to flush said cache, and determine addresses of said cache therefrom so that said set of cache lines are read from said cache so as to avoid reading cache lines from said cache that are in an invalid or a clean state.

2. The processor of claim 1 wherein said logic circuitry is part of a cache agent and said cache is a cache slice.

15 3. The processor of claim 1 wherein said storage circuitry is a memory to store words having locations corresponding to specific locations within said cache.

4. The processor of claim 2 wherein said specific locations within said cache correspond to a specific set and way of said cache.

5. The processor of claim 1 further comprising multiple processing cores coupled to an interconnection network, said logic circuitry coupled to said interconnection network.

20 6. The processor of claim 1 further comprising power management circuitry coupled to said logic circuitry, said power management circuitry to provide said flush signal.

7. The processor of claim 1 further comprising voltage control circuitry to cause a voltage applied to said cache to be lowered after said set of cache lines have been read from said cache.

25 8. The processor of claim 1 wherein said storage circuitry includes a separate memory for each way of said cache.

9. A method, comprising:

in response to a determination that a cache line is in a modified state, writing information to a memory that indicates said cache line is in a modified state;

30 writing said cache line, while in said modified state, to a cache;

writing one or more clean or invalid cache lines to said cache;

in response to a signal to flush said cache, reading said information from said memory to identify said cache line;

35 reading said cache line from said cache but not any of said one or more clean or invalid cache lines;

after said reading, reducing a supply voltage to said cache.

10. The method of claim 9 wherein said cache is a cache slice.

11. The method of claim 10 wherein said information is stored in a location of said memory that corresponds to a set and way of said cache slice where said cache line is stored.

5 12. The method of claim 11 further comprising determining an address for said cache line in said cache from said set and way information, and, applying that address to said cache to read said cache line.

13. The method of claim 9 further comprising power management circuitry generating said signal.

10 14. The method of claim 9 further comprising, prior to said reading of said information, reading second information from said storage circuitry to identify at least one region of said storage circuitry where no information concerning a modified cache line is found.

15. The method of claim 9 further comprising saving said cache line to second storage circuitry after reading said cache line from said cache.

15 16. A processor, comprising:

a cache agent coupled to a cache slice, said cache agent comprising:

i) cache line state determination circuitry;

ii) cache line state tracking circuitry coupled to said cache line state determination circuitry;

20 iii) storage circuitry coupled to said cache line state tracking circuitry, said cache line state tracking circuitry to write information into said storage circuitry that identifies which of said cache slice's cache lines are in a modified state, said cache line state determination circuitry to inform said cache line state tracking circuitry which of said cache slice's cache lines are in a modified state.

25 17. The processor of claim 15 further comprising multiple processing cores coupled to an interconnection network, said cache agent coupled to said interconnection network.

18. The processor of claim 15 further comprising power management circuitry coupled to said cache agent, said power management circuitry to provide a flush signal to said cache agent to cause said cache agent to refer to information within said storage circuitry in order to
30 read modified cache lines from said cache slice and avoid reading clean or invalid cache lines from said cache slice.

19. The processor of claim 15 further comprising voltage control circuitry to cause a voltage applied to said cache slice to be lowered after said cache slice's modified cache lines have been read from said cache.

20. The processor of claim 15 wherein said storage circuitry is a memory whose memory words having different locations each correlated to a different set and way combination within said cache slice.










FIG. 1

**FIG. 2**

**FIG. 3**

**FIG. 4**

**FIG. 5A****FIG. 5B**

**FIG. 6**

A. CLASSIFICATION OF SUBJECT MATTER**G06F 9/06(2006.01)i, G06F 9/30(2006.01)i, G06F 12/08(2006.01)i**

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 9/06; G06F 12/02; G06F 12/00; G06F 13/24; G06F 12/08; G06F 9/30

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: flush, time, reduce, cache line, processor;

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 8180968 B2 (O'KRAFKA BRIAN W. et al.) 15 May 2012 See column 4, line 33 - column 6, line 67; claims 1, 9, and figures 1-7.	1-20
A	US 7934069 B2 (COCHRAN ROBERT A et al.) 26 April 2011 See column 2, line 54 - column 5, line 28; claims 1,15, and figures 1-4.	1-20
A	US 2012-0166731 A1 (MACIOCCO CHRISTIAN et al.) 28 June 2012 See paragraphs [0012] - [0018]; claims 1, 10, 18, and figures 1-2.	1-20
A	US 6347379 B1 (DAI, XIA et al.) 12 February 2002 See column 3, line 36 - column 6, line 37; claim 1, and figures 1-5.	1-20



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 application or patent 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 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

27 August 2013 (27.08.2013)

Date of mailing of the international search report

29 August 2013 (29.08.2013)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

BOK Jin Yo

Telephone No. +82-42-481-5113



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/048362

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8180968 B2	15/05/2012	US 2008-244185 A1	02/10/2008
US 7934069 B2	26/04/2011	US 2008-0270692 A1	30/10/2008
US 2012-0166731 A1	28/06/2012	TW 201239609 A	01/10/2012
		WO 2012-087655 A2	28/06/2012
		WO 2012-087655 A3	16/08/2012
US 6347379 B1	12/02/2002	US 2002-0026597 A1	28/02/2002
		US 6704880 B2	09/03/2004