



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

(21) International Application Number: PCT/US2014/016993

(22) International Filing Date: 18 February 2014 (18.02.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 13/769,965 19 February 2013 (19.02.2013) US

(71) Applicant: **QUALCOMM INCORPORATED** [US/US];
ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, CA 92121-1714 (US).

(72) Inventor: **DONG, Xiangyu**; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

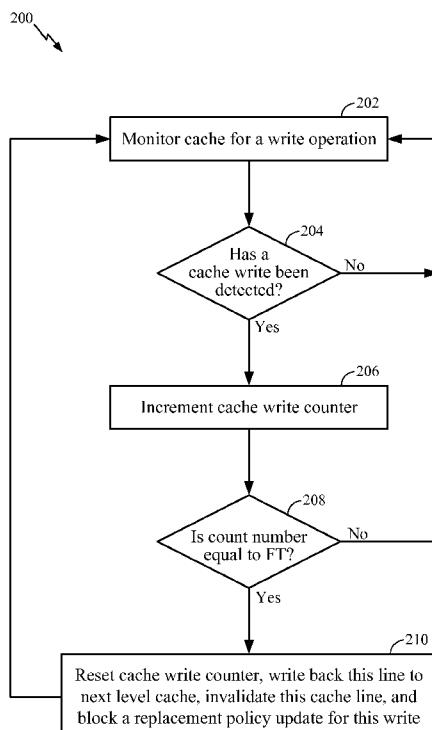
(74) Agent: **CICCOZZI, John L.**; 4000 Legato Road, Suite 310, Fairfax, Virginia 22033 (US).

(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, IR, 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, SA, 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).

[Continued on next page]

(54) Title: METHODS AND APPARATUS FOR INTRA-SET WEAR-LEVELING FOR MEMORIES WITH LIMITED WRITE ENDURANCE



(57) **Abstract:** Efficient techniques are described for extending the usable lifetime for memories with limited write endurance. A technique for wear-leveling of caches addresses unbalanced write traffic on cache lines which cause heavily written cache lines to fail much faster than other lines in the cache. A counter is incremented for each write operation to a cache array. A line affected by a current write operation which caused the counter to meet a threshold is evicted from the cache rather than writing data to the affected line. A dynamic adjustment of the threshold can be made depending on the operating program. Updates to a current replacement policy pointer are stopped due to the counter meeting the threshold.

FIG. 2



Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))* — *with international search report (Art. 21(3))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

METHODS AND APPARATUS FOR INTRA-SET WEAR-LEVELING FOR MEMORIES WITH LIMITED WRITE ENDURANCE

Field of the Disclosure

[0001] The present disclosure relates generally to aspects of processing systems and in particular to methods and apparatus to extend the usable lifetime of memories with limited write endurance.

Background

[0002] Many portable products, such as cell phones, laptop computers, personal data assistants (PDAs) and the like, utilize a processing system that executes programs, such as communication and multimedia programs. A processing system for such products may include multiple processors, complex memory systems including multi-levels of caches and memory for storing instructions and data, controllers, peripheral devices such as communication interfaces, and fixed function logic blocks configured, for example, on a single chip. At the same time, portable products have a limited energy source in the form of batteries that are often required to support high performance operations by the processing system and increasingly large memory capacities as functionality increases. To improve battery life, it is desirable to perform these operations as efficiently as possible. However, the scaling of common memory platforms, such as static random access memory (SRAM) and embedded dynamic random access memory (eDRAM), is increasingly constrained by leakage power and cell density. Such concerns extend to personal computer products which are also being developed with efficient designs to operate with reduced overall energy consumption.

[0003] A number of memory technologies, such as flash memory, magnetoresistive random access memory (MRAM), phase change memory (PCM), resistive RAM

(ReRAM), and others, have various limits on the number of write operations that can be performed to the device before memory cells begin to wear out and fail. Memories such as caches, which operate based on principles of spatial and temporal locality and at high data rates, show a wide variability in cache line accesses from program to program and may have cells that experience a very high rate of write accesses. For example, set associative caches which have a plurality of sets of data, each set divided in a plurality of selectable cache ways and each way in a set holding a cache line, experience different write access patterns within each set depending on the program in execution. Also, the write access variability from line to line in a set may be very large and may vary dynamically during system operations.

[0004] For example, an 8-way set associative 64 kbyte cache may be constructed having 256 sets of eight 32-byte cache lines per set and access one cache line for each way. Such a cache may be used in a level 1 data cache in a portable device, such as a cell phone, tablet, lap top and the like. In the eight way set associative cache for a program X, write access to a line of data in way 2 may occur multiple orders of magnitude more often than write access to a line of data in a different cache way, such as way 7. Thus, the line of memory internal to the cache for way 2 may fail much earlier than the memory line in way 7 and most other lines in the cache having write access rates lower than the write access rates of way 2. The memory wear endurance affects each level of a memory hierarchy, such as level 1, 2, and 3 caches, flash memory, and system memory, though to different degrees in each level and each device. With processors running in the gigahertz (GHz) frequency, caches experience a large number of write accesses which may be to specific lines in the cache and thus such locality of accesses may cause a system to too rapidly approach the write limits of the cache memory. Since any cache line may experience high write operations depending

on the program in execution and such a cache line hotspot is not known in advance, the cost for monitoring every cache line to determine which cache line in a cache set is affected and should be relocated to reduce wear may be prohibitive.

SUMMARY

[0005] Among its several aspects, the present disclosure recognizes that providing more efficient methods and apparatuses for intra-set wear-leveling extends the usable lifetime of memories with limited write endurance. To such ends, an embodiment of the invention addresses a method for wear-leveling of a cache. A counter is incremented for each write operation to the cache. A line affected by a current write operation which caused the counter to meet a threshold is evicted from the cache. The current write to the line affected by the current write operation is stopped due to the counter meeting the threshold.

[0006] Another embodiment addresses an apparatus for wear-leveling a cache. A cache having a plurality of ways is coupled to a processor complex which writes data to the cache. A write threshold circuit and a counter circuit are configured to count write operations to the cache and generate an indication when the counter circuit has reached a threshold based on a comparison in the write threshold circuit. A line flush circuit is configured to evict from the cache a line affected by the current write operation in response to the indication. A replacement policy circuit is configured to stop updates to change a current line replacement policy pointer in response to the indication.

[0007] Another embodiment addresses a computer readable non-transitory medium encoded with computer readable program data and code. A counter is incremented for each write operation to a cache. A line affected by a current write operation which caused the counter to meet a threshold is evicted from the cache. The current write to

the line affected by the current write operation is stopped due to the counter meeting the threshold.

[0008] A further embodiment addresses apparatus for wear-leveling a cache. Means is utilized for writing data to a cache having a plurality of ways coupled to a processor complex. Means is utilized for counting write operations to the cache and generating an indication when the counter has reached a threshold. Means is utilized for evicting from the cache a line affected by the current write operation in response to the indication.

Means is utilized for stopping updates to change a current line replacement policy pointer in response to the indication.

[0009] It is understood that other embodiments of the present invention will become readily apparent to those skilled in the art from the following detailed description, wherein various embodiments of the invention are shown and described by way of illustration. As will be realized, the invention is capable of other and different embodiments and its several details are capable of modification in various other respects, all without departing from the spirit and scope of the present invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not as restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Various aspects of the present invention are illustrated by way of example, and not by way of limitation, in the accompanying drawings, wherein:

[0011] FIG. 1 illustrates an exemplary cache system in which an embodiment of the invention may be advantageously employed;

[0012] FIG. 2 illustrates an advantageous process for wear-leveling; and

[0013] FIG. 3 illustrates a particular embodiment of a portable device having a processor complex that is configured to extend wear endurance in the portable devices memory system.

DETAILED DESCRIPTION

[0014] The detailed description set forth below in connection with the appended drawings is intended as a description of various exemplary embodiments of the present invention and is not intended to represent the only embodiments in which the present invention may be practiced. The detailed description includes specific details for the purpose of providing a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced without these specific details. In some instances, well known structures and components are shown in block diagram form in order to avoid obscuring the concepts of the present invention.

[0015] FIG. 1 illustrates an exemplary memory system 100 in which an embodiment of the invention may be advantageously employed. The memory system 100 includes a cache 102, for example a data cache, a next level cache 103, a system memory 105, and a processor complex 106. The system memory 105 provides access for instructions and data that are not found in the cache 102 or in the next level cache 103. It is noted that the cache 102 and the next level cache 103 may be integrated with processor complex 106. The system memory 105 may also be integrated with the processor complex 106 or separately coupled to the processor complex 106. The processor complex 106 may include one or more processors, such as a dual or a quad processor system. The cache 102 is illustrated as an N-way set associative cache that is representative of a level 1 data cache and in this case, the next level cache 103 would be a level 2 (L2) cache. It is

noted that the wear-leveling techniques described herein are also applicable to an L2 cache and a level 3 (L3) cache depending on the type of memory used at each level of the memory hierarchy. The system memory 105 generally experiences a different form of wear-leveling than that experienced in caches. The cache 102 comprises a cache control circuit 108 having a cache write counter 110, a write threshold circuit 111, an address tag selector circuit 114, a line flush circuit 115, and a replacement policy circuit 116 for implementing wear-leveling in the cache 102. The cache 102 also comprises way 1 118₁, way 2 118₂, ..., way N 118_N. Each way is configured with sets such as set i 120 comprising a tag 122 and line 123 in way 1 118₁ and each way 118₂ ... 118_N is similarly configured.

[0016] A technique for wear-leveling of caches addresses unbalanced write traffic on cache lines which cause heavily written cache lines to fail much faster than other lines in the cache. It is also noted that write traffic to blocks of memory in a bulk memory, such as system memory 105 having limited write endurance, may also be unbalanced and cause heavily written memory blocks to fail sooner than other memory blocks in the bulk memory. To address this problem of wear-leveling in caches with limited write endurance, means are provided to count write operations to the cache and generate an indication when a write threshold has been reached. For example, the cache write counter 110 is configured to increment a count for every cache write operation under control of the cache control circuit 108. Upon reaching or exceeding a write count threshold, also referred to as a line-flush threshold (FT), the cache line under the current write operation is evicted, for example by means to flush the line from the cache such as provided by the line flush circuit 115. Generally, a line is flushed from a cache when the line is indicated to be a least recently used (LRU) line and thus is available to be rewritten. In the case of reaching the FT, means are provided to stop an update to a

cache line replacement operation for this current write operation. The cache write counter 110 is then reset to initialize the counter in preparation for the next cache write monitoring period. The evicted cache line is written to the next level cache and whenever that cache line is accessed again it is reloaded into a line in the cache based on the replacement policy which indicates a cache line different from the line it was evicted from.

[0017] In particular, when a cache line is evicted from the lower level cache, that cache line is marked invalid indicating that the invalid line can be filled with new content. A cache line may be marked invalid by setting a valid bit to an invalid state in the tag associated with the cache line. A cache line that is evicted in standard operations is generally identified by the replacement policy, such as a least recently used (LRU) policy, to be available to be rewritten. After invalidation, it is possible for a cache line that is used frequently, to have the frequently used cache line be rewritten on a subsequent miss on the same cache line. To avoid this rewriting problem for evicted cache lines, when the cache write counter has reached or exceeded the count threshold, an update to a current line replacement policy pointer is blocked and the pointer is not updated. Since the replacement policy pointer is not updated, the current LRU pointer, for example, points to a different line than was evicted. Thus, the cache line that was evicted remains empty until that line really becomes least recently used.

[0018] For example in Fig. 1, with the cache 102 configured for a write operation, the processor complex 106 issues an address 130 to the cache 102. In response to the write operation an address tag 131 is generated and compared in tag compare circuit 124 to the tags in each way, such as tag 122 in way 1 118₁. If a tag compare matches, the tag compare circuit 124 causes a cache line select circuit 125 to select a path to a way, such as associated with line 123 in way 1 118₁ and data 132 is written to the selected

line 123. A dirty bit may also be written to the tag 122 indicating data in the line 123 has been modified. For each write operation, the cache write counter 110 is incremented. A write operation may be determined and the cache write counter incremented by a write signal internal to the cache control circuit 108 that is generated when data is to be written to a particular way and line in the cache. Alternatively, the cache control circuit 108 may use an internal command decode operation to set a bit or internal signal state that indicates whether the commanded operation will cause a write operation or a read operation.

[0019] The count value from the cache write counter 110 is compared to a preset threshold, such as FT in the write threshold circuit 111. Upon reaching or exceeding the threshold, the line flush circuit 115 is notified and the cache line for the current write operation is flushed, rather than writing the data to the selected line. The evicted line, such as the line 123 in way 1 118₁ is written to the next level cache 103, such as an L2 cache if the cache 102 is a level 1 data cache. In an alternative embodiment, the evicted line is written to an L3 cache if the cache 102 is an L2 cache and the processor complex 106 includes a level 1 data cache. The replacement policy circuit 116 is also notified and the replacement policy, such as least recently used (LRU) is not updated. At a later time when the evicted cache line 123 is rewritten to the cache 102, the cache control circuit 108 ensures that the evicted cache line 123 is written to a different location in set i 120, such as a line in one of the other ways, way 2 118₂ through way N 118_N.

[0020] The more frequently a program writes to a cache line, the more likely that cache line will be selected when the cache write counter saturates by reaching or exceeding a preset threshold value. Over an operating period, this wear-leveling

technique is probabilistic, evicts hotspot cache lines, and relocates these lines to different cache locations providing a significant degree of wear-leveling across the cache array. Thus, by using a global counter of cache write accesses a significant degree of wear-leveling is achieved at a low cost without requiring a costly monitoring of each cache line. Also, by allowing the write count threshold to be settable, a dynamic adjustment of the threshold can be made depending on the operating program. A smaller write count threshold results in a higher frequency of line invalidations, thus the threshold can be dynamically tuned to provide a more balanced intra-set write distribution that is adaptable to a particular program or set of programs. For example, when a new program is executed that has an increased number of hotspots in a cache compared to a previous operating program, the write counter threshold for that cache is reduced to improve the odds of reducing write operations to the hotspots and thereby improve the wear endurance of that cache.

[0021] In an alternative embodiment, rather than marking the evicted cache line invalid and blocking an update to the replacement policy due to the evicted line, an additional bit may be added to each cache line's tag which indicates this cache line is a probabilistic invalidated line and reuse of the evicted line should be delayed. A replacement policy, such as LRU, will select a different line for eviction if the selected LRU line has this bit set. The probabilistic invalidated line bit may be reset after a predetermined number of successful cache line write operations to that cache set or reset after a predetermined time period under control of a timer, for example.

[0022] The cache unit 102, as a stand-alone cache chip or fixed-integrated module, may receive controls 134 to load a threshold value to the write threshold circuit 111. The controls 134 may include a serial input path that is used to load the threshold value on initialization or at a time determined by a program operating on the processor

complex 106. The threshold value could initially be a default threshold value hardcoded on the chip or fixed-integrated module. Then, the default threshold value as loaded in the write threshold circuit 111 may be adjusted according to the program in operation. In an alternative embodiment, monitors, such as a performance drop monitor or a wear effectiveness monitor are implemented in hardware external of the cache 102 or in software coded in a program that operates on the processor complex 106. The internal cache control circuit 108 then receives a signal from one or more such monitors that indicates a threshold value in the write threshold circuit 111 should be increased, decreased, or allowed to stay the same as previously set.

[0023] FIG. 2 illustrates an advantageous process 200 for wear-leveling. At block 202, cache operations are monitored for a write operation. At decision block 204, a determination is made whether a cache write has been detected. A cache write operation may be determined, for example, in the cache control circuit 108 of FIG. 1. If no cache write operation has been identified, the process 200 returns to block 202. If a cache write operation has been identified, then the process 200 proceeds to block 206. At block 206, the cache write counter is incremented. At block 208, a determination is made whether the count value in the cache write counter is equal to or in an alternate embodiment has exceeded a preset threshold (FT). If it is determined that the count value is not equal to or has not exceeded the preset threshold FT, the process 200 returns to block 202. If the count value is equal to or has exceeded the preset threshold FT, the process 200 proceeds to block 210. At block 210, the cache write counter is reset, the line that would have been written to the monitored cache is instead written to the next level cache, the line in the monitored cache is invalidated, and a replacement policy update is blocked for this write. The process 200 returns to block 202.

[0024] FIG. 3 illustrates a particular embodiment of a device 300 having a cache controller circuit 328 that is configured to extend the usable lifetime of the memories having limited wear endurance in the memory system of the device 300. The device 300 may be a wireless electronic portable device and include the processor complex 310 coupled to a system memory 312 having software instructions 318. The system memory 312 may include the system memory 105 of FIG. 1. The processor complex 310 may include a processor 311, an integrated memory subsystem 314 having a level 1 data cache (L1 Dcache) 322, a level 1 instruction cache (L1 Icache) 324, and the cache controller circuit 328. The processor complex 310 may include the processor complex 106 of FIG. 1. The integrated memory subsystem 314 may also include a level 2 (L2) unified cache 326. The L1 Dcache 322 may include the cache 102 of FIG. 1 and the L2 unified cache 326 may include the next level cache 103 of FIG. 1, each of which may have limited wear endurance.

[0025] The integrated memory subsystem 314 may be included in the processor complex 310 or may be implemented as one or more separate devices or circuitry (not shown) external to the processor complex 310. In an illustrative example, the processor complex 310 operates according to the software instructions 318 and in accordance with any of the embodiments illustrated in or associated with FIGS. 1 and 2. For example, as shown in FIG. 3, the L1 Icache 326, the L1 Dcache 322, and the cache controller circuit 328 are accessible within the processor complex 310, and the processor 311 is configured to access data or program instructions stored in the memories of the integrated memory subsystem 314 or in the system memory 312.

[0026] A camera interface 334 is coupled to the processor complex 310 and also coupled to a camera, such as a video camera 336. A display controller 340 is coupled to the processor complex 310 and to a display device 342. A coder/decoder (CODEC) 344

may also be coupled to the processor complex 310. A speaker 346 and a microphone 348 may be coupled to the CODEC 344. A wireless interface 350 may be coupled to the processor complex 310 and to a wireless antenna 352 such that wireless data received via the antenna 352 and wireless interface 350 can be provided to the dual processor 311.

[0027] The processor 311 may be configured to execute software instructions 318 accessed from the system memory 312 and stored in a non-transitory computer-readable medium that are executable to cause a computer, such as the processor 311, to execute a program, such as the process 200 of FIG. 2. The software instructions 318 are further executable to cause the processor 311 to process instructions that read and write access the memories of the integrated memory subsystem 314 and the system memory 312.

[0028] In a particular embodiment, the processor complex 310, the display controller 340, the system memory 312, the CODEC 344, the wireless interface 350, and the camera interface 334 are included in a system-in-package or system-on-chip device 304. In a particular embodiment, an input device 356 and a power supply 358 are coupled to the system-on-chip device 304. Moreover, in a particular embodiment, as illustrated in FIG. 3, the display device 342, the input device 356, the speaker 346, the microphone 348, the wireless antenna 352, the video camera 336, and the power supply 358 are external to the system-on-chip device 304. However, each of the display device 342, the input device 356, the speaker 346, the microphone 348, the wireless antenna 352, the video camera 336, and the power supply 358 can be coupled to a component of the system-on-chip device 304, such as an associated interface or controller.

[0029] The device 300 in accordance with embodiments described herein may be incorporated in a variety of electronic devices, such as a set top box, an entertainment

unit, a navigation device, a communications device, a personal digital assistant (PDA), a fixed location data unit, a mobile location data unit, a mobile phone, a cellular phone, a computer, a portable computer, tablets, a monitor, a computer monitor, a television, a tuner, a radio, a satellite radio, a music player, a digital music player, a portable music player, a video player, a digital video player, a digital video disc (DVD) player, a portable digital video player, any other device that stores or retrieves data or computer instructions, or any combination thereof.

[0030] The various illustrative logical blocks, modules, circuits, elements, or components described in connection with the embodiments disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic components, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. The processor 311 may be a general purpose processor, or in the alternative, the processor 311 may be a combination of a conventional processor, a digital signal processor, a microcontroller, or a state machine. The processor 311 may also be implemented as a combination of computing components, for example, a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration appropriate for a desired application.

[0031] The methods described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, hard disk, a removable disk, a CD-ROM, or any other form of non-transitory storage medium

known in the art. A non-transitory storage medium may be coupled to the processor such that the processor can read information from, and write information to, the non-transitory storage medium. In the alternative, the non-transitory storage medium may be integral to the processor.

[0032] The processor complex 106 of FIG. 1 or the processor 311 of FIG. 3, for example, may be configured to execute instructions including instructions selected from a wear-leveling program stored on a computer readable non-transitory storage medium either directly associated locally with the processor 311, such as may be available through an instruction cache, or accessible through input device 356 or through the wireless interface 350. The input device 356 or the wireless interface 350 also may access data residing in a memory device either directly associated locally with the processor, such as the L1 Icache 324, the L1 Dcache 322, the L2 cache 326 or accessible from another processor's memory. The computer readable non-transitory storage medium may include random access memory (RAM), dynamic random access memory (DRAM), synchronous dynamic random access memory (SDRAM), flash memory, read only memory (ROM), programmable read only memory (PROM), erasable programmable read only memory (EPROM), electrically erasable programmable read only memory (EEPROM), compact disk (CD), digital video disk (DVD), other types of removable disks, or any other suitable non-transitory storage medium.

[0033] While the invention is disclosed in the context of illustrative embodiments for use in processor systems, it will be recognized that a wide variety of implementations may be employed by persons of ordinary skill in the art consistent with the above discussion and the claims which follow below. For example, a fixed function implementation may also utilize various embodiments of the present invention.

CLAIMS**WHAT IS CLAIMED IS:**

1. A method for wear-leveling of a cache, the method comprising:
 - incrementing a counter for each write operation to the cache;
 - evicting from the cache a line affected by a current write operation which caused the counter to meet a threshold; and
 - stopping the current write to the line affected by the current write operation due to the counter meeting the threshold.
2. The method of claim 1 further comprising:
 - stopping updates to a current replacement policy pointer due to the counter meeting the threshold.
3. The method of claim 1, wherein the threshold is a program settable value that is compared to the count value output of the counter which provides an indication when the count value matches the threshold.
4. The method of claim 1, wherein the line when next accessed is written to a cache line that is different from the cache line it was previously evicted from.

5. The method of claim 1, wherein the evicted line is selected from a way of an addressed set of ways in the cache and the line when next accessed is written to a different line in the addressed set of ways.
6. The method of claim 1 further comprising:
resetting the counter to initialize the counter in preparation for a succeeding cache write monitoring period.
7. The method of claim 1 further comprising:
writing the evicted line to the next level cache, wherein the cache is a level 1 cache and the next level cache is a level 2 cache.
8. The method of claim 1 further comprising:
writing the evicted line to the next level cache, wherein the cache is a level 2 cache and the next level cache is a level 3 cache.
9. The method of claim 1 further comprising:
marking the line affected by the current write operation which caused the counter to meet the threshold invalid, wherein the invalid mark causes the line to be evicted from the cache.
10. An apparatus for wear-leveling a cache, the apparatus comprising:
a cache having a plurality of ways coupled to a processor complex which controls write operations to write data to the cache;

a write threshold circuit and a counter circuit configured to count the write operations to the cache and to generate an indication when the counter circuit has reached a threshold based on a comparison of a count output of the counter circuit with a threshold value;

a line flush circuit configured to evict from the cache a line affected by the current write operation in response to the indication; and

a replacement policy circuit is configured to stop updates to change a current line replacement policy pointer in response to the indication.

11. The apparatus of claim 10, wherein the threshold value is a program settable value that is compared to the count value output of the counter which provides the indication when the count value matches the threshold.

12. The apparatus of claim 10, wherein the threshold value is a default value that is stored in the write threshold circuit and adjusted according to a program in operation.

13. The apparatus of claim 10, wherein the line when next accessed is written to a cache line that is different from the cache line it was previously evicted from.

14. The apparatus of claim 10, wherein the evicted line is selected from a way of an addressed set of ways in the cache and the line when next accessed is written to a different line in the addressed set of ways.

15. The apparatus of claim 10, wherein the counter circuit is reset to initialize the counter in preparation for a succeeding cache write monitoring period.

16. The apparatus of claim 10 further comprising:
 - a level 1 cache configured as the cache; and
 - a level 2 cache configured as a next level cache, wherein the evicted line is written to the level 2 cache.
17. The apparatus of claim 10 further comprising:
 - a level 2 cache configured as the cache; and
 - a level 3 cache configured as a next level cache, wherein the evicted line is written to the level 3 cache.
18. The apparatus of claim 10, wherein an invalidated line bit is included in each cache line's tag to indicate this cache line is a probabilistic invalidated line and reuse of the evicted line should be delayed
19. A computer readable non-transitory medium encoded with computer readable program data and code, the program data and code when executed operable to:
 - increment a counter for each write operation to a cache;
 - evict from the cache a line affected by a current write operation which caused the counter to meet a threshold; and
 - stop the current write to the line affected by the current write operation due to the counter meeting the threshold.
20. An apparatus for wear-leveling a cache, the apparatus comprising:

means for writing data to a cache having a plurality of ways coupled to a processor complex;

means for counting write operations to the cache and generating an indication when the counter has reached a threshold;

means for evicting from the cache a line affected by the current write operation in response to the indication; and

means for stopping updates to change a current line replacement policy pointer in response to the indication.

1/3

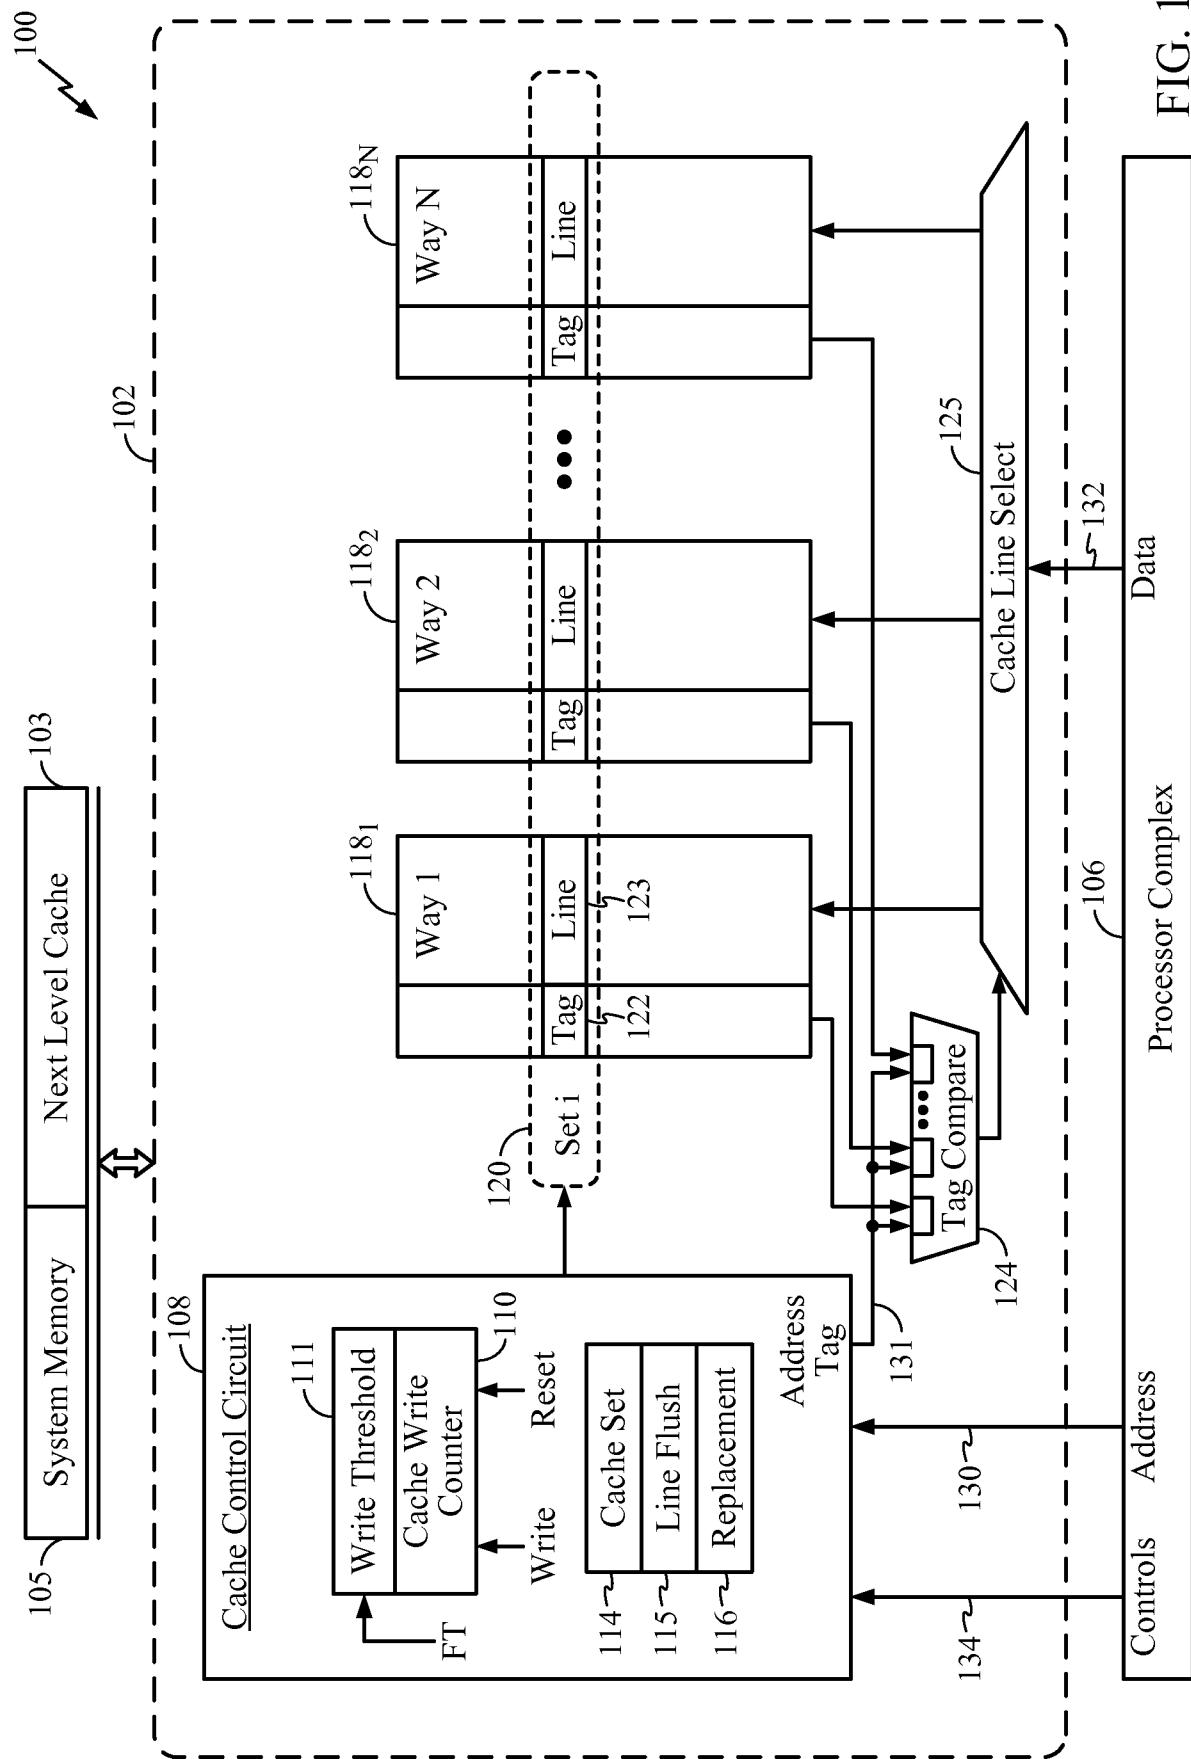


FIG. 1

2/3

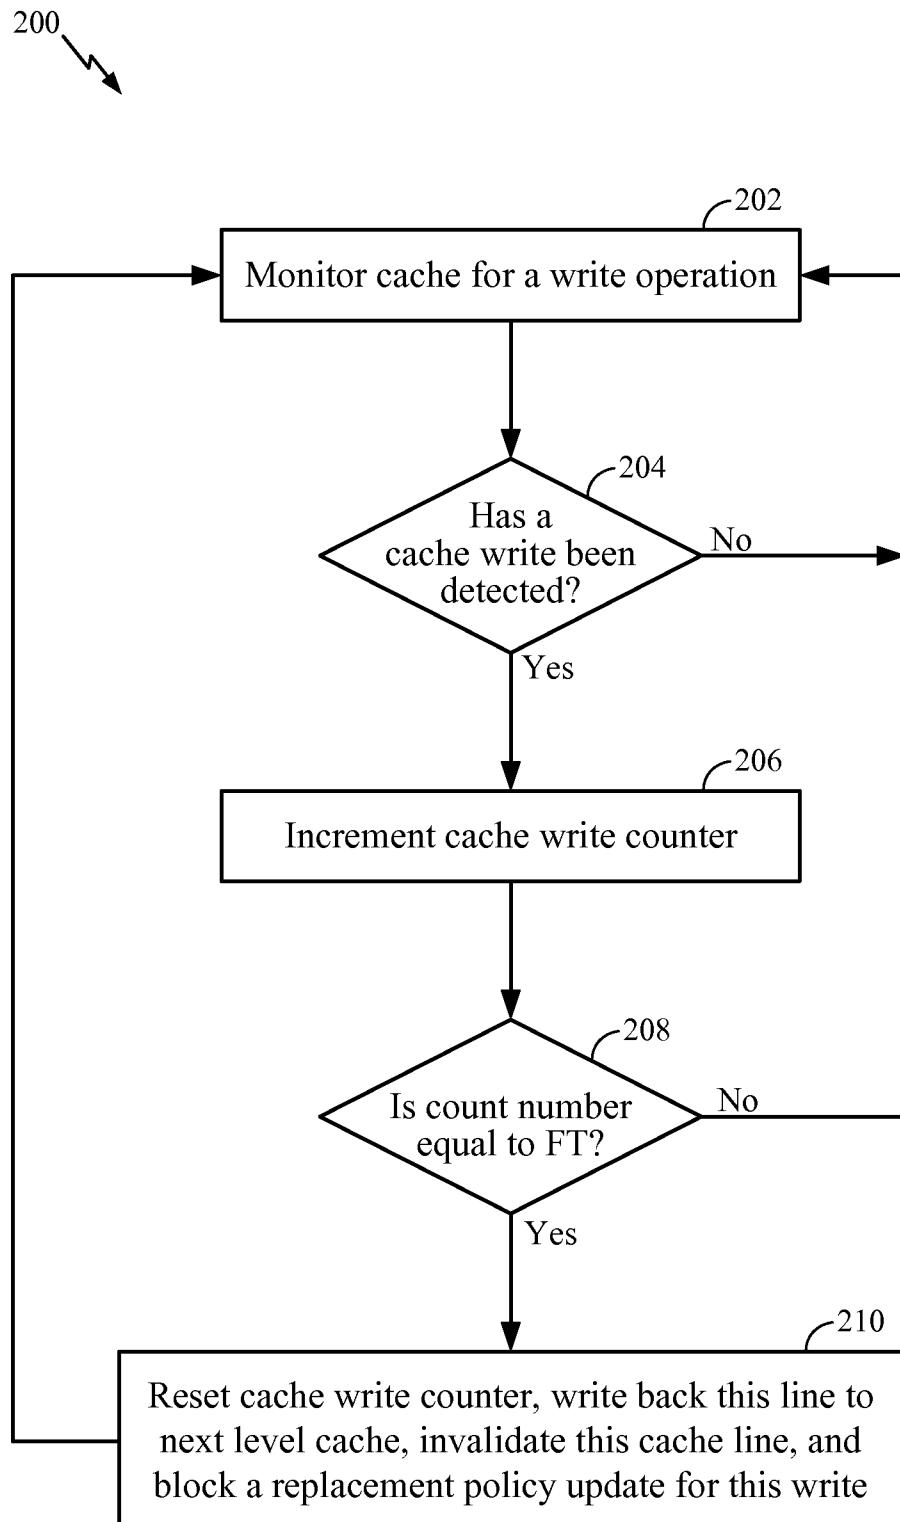


FIG. 2

3/3

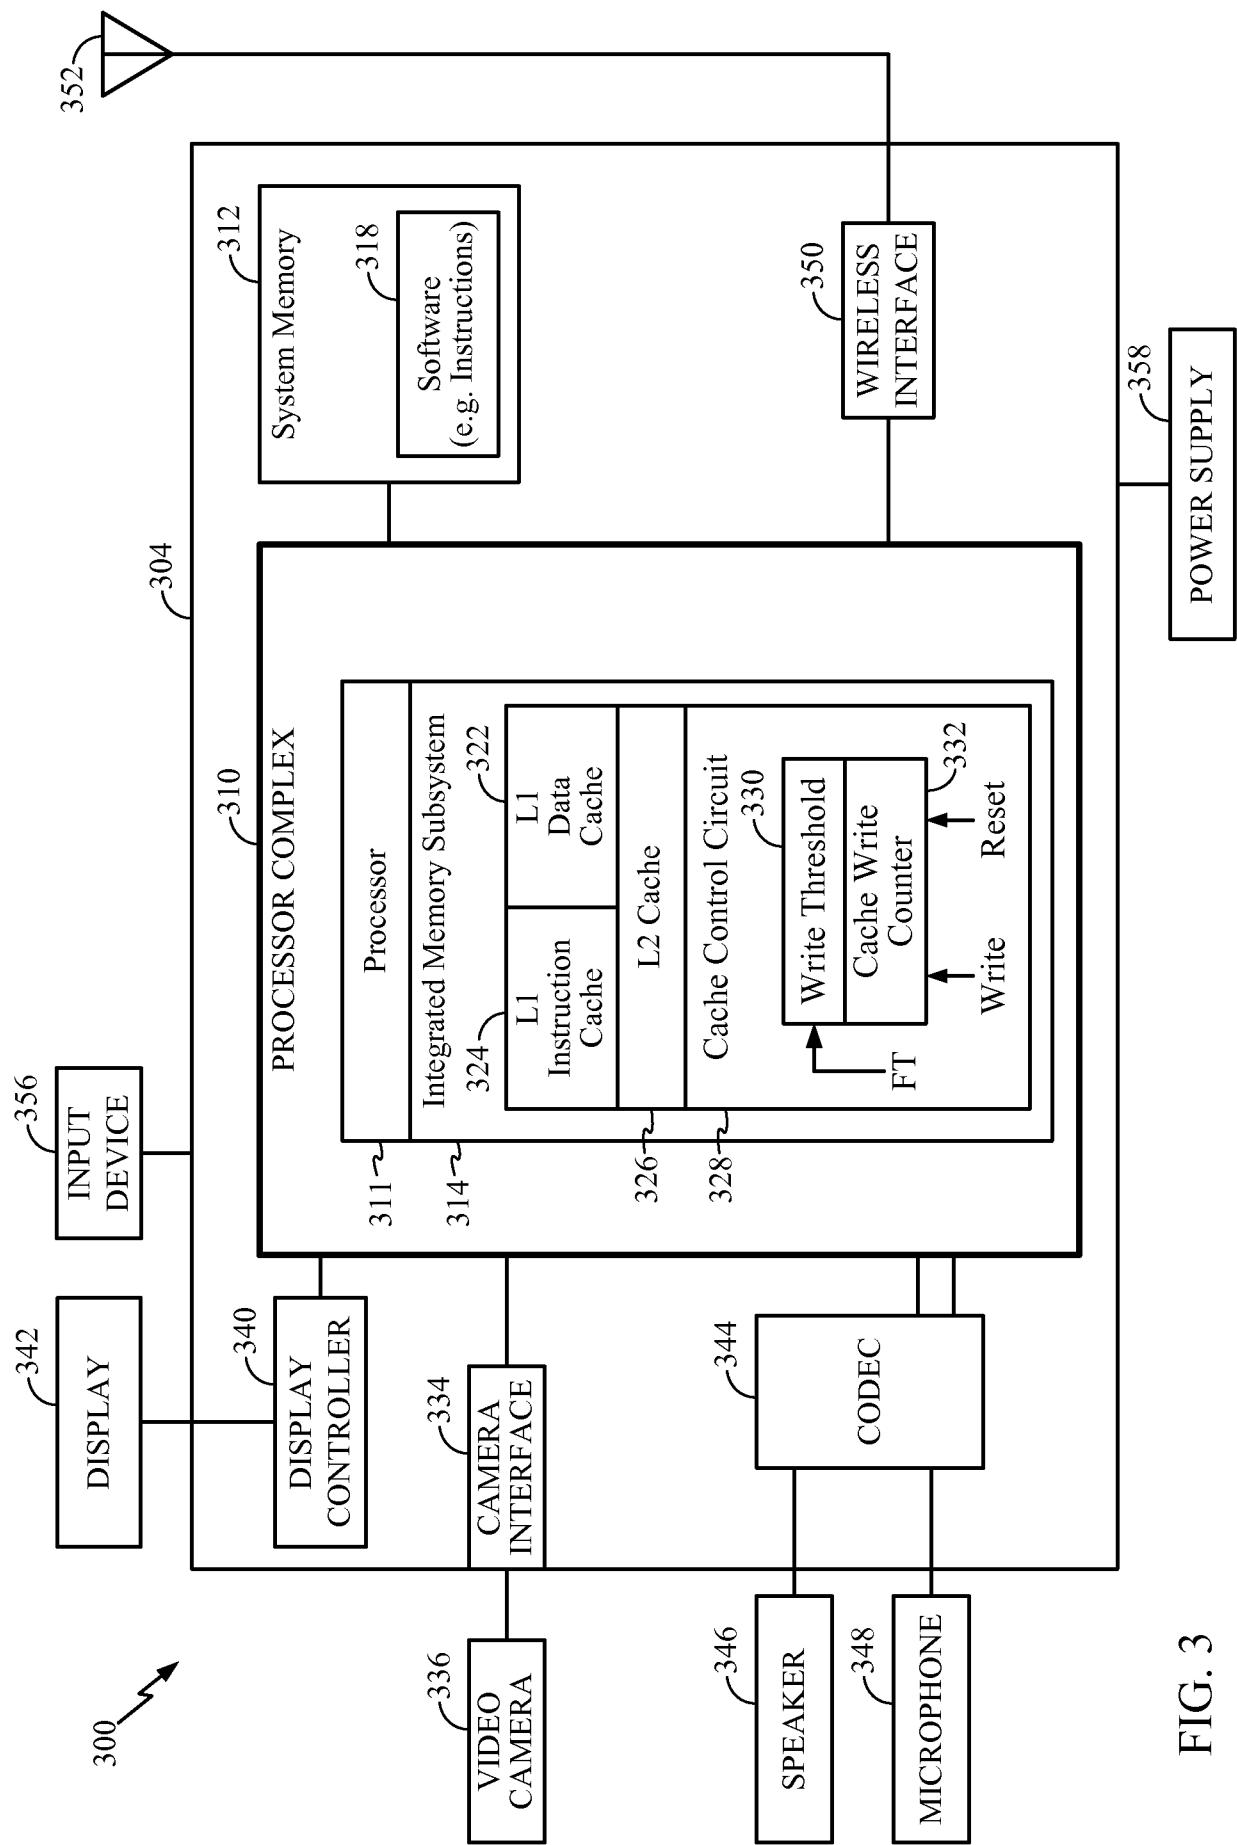


FIG. 3

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/016993

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F12/08
ADD.

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 practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/114930 A1 (SANVIDO MARCO [US] ET AL KULKARNI ANAND KRISHNAMURTHI [US] ET AL) 15 May 2008 (2008-05-15) the whole document -----	1-17,19, 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 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
23 May 2014	04/06/2014
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Filip, Liviu

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/016993

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008114930	A1 15-05-2008	NONE	