



- (51) International Patent Classification:
G06F 12/08 (2006.01) *G06F 9/44* (2006.01)
G06F 12/12 (2006.01)
- (21) International Application Number:
PCT/US2016/029377
- (22) International Filing Date:
26 April 2016 (26.04.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/721,625 26 May 2015 (26.05.2015) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

- (72) Inventors: ANANTARAMAN, Aravindh V.; 1550 Iron Point Road, Apt. 2611, Folsom, California 95630 (US). GREENFIELD, Zvika; 16-b Ha'emek St., 44470 Kfar Sava (IL). DIAMAND, Israel; Aderet 1039, 99850 Aderet (IL). NORI, Anant V.; 904 Maangalya Suryodaya, Near Yashomati Hospital, Bangalore 560103 (IN). RAMACHANDRAN, Pradeep; Intel Technology India Pvt Ltd, Survey # 23-56P, Outer ring road, Devarabeesanahal, Bangalore 560103 (IN). MISGAV, Nir; 5/15, Iyar Street, 4250324 Netanya (IL).
- (74) Agent: OSBORNE, David W.; THORPE NORTH & WESTERN, LLP, C/O CPA Global, P.O. Box 52050, Minneapolis, Minnesota 55402 (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, 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.

[Continued on next page]

(54) Title: MANAGING SECTORED CACHE




FIG. 1

(57) Abstract: Apparatus, systems, and methods to manage memory operations are described. In one example, a controller comprises logic to receive a first transaction to operate on a first data element in the cache memory, perform a lookup operation for the first data element in the volatile memory and in response to a failed lookup operation, to generate a cache scrub hint forward the cache scrub hint to a cache scrub engine and identify one or more cache lines to scrub based at least in part on the cache scrub hint. Other examples are also disclosed and claimed.



(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, ST, 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).

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

MANAGING SECTORED CACHE

FIELD

[0001] The present disclosure generally relates to the field of electronics. More particularly, some examples generally relate to techniques for managing sectored cache in electronic devices.

BACKGROUND

[0002] It may be useful for electronic devices to store data in a non-volatile memory such that the data remains intact when power to the electronic device is turned off or otherwise lost. The non-volatile memory device for storing digital information in an array of non-volatile memory cells may be included in a Non-Volatile Dual In-line Memory Module (NVDIMM). Digital information stored in the non-volatile memory (for example, NAND/NOR flash memory) persists in the during power loss or system failures. After power to the electronic device is restored, the electronic device can access the stored digital data from the NVDIMM.

[0003] Logic executing in an electronic device can modify data stored in non-volatile memory. For example, logic executing on a processor of the electronic device may update a data element stored in non-volatile memory. In such an instance, the logic retrieves a copy of the data element stored in non-volatile memory and stores a copy of the data element in a volatile memory, e.g., cache memory. Some cache memories may be structured as "sectored" cache, in which multiple cache lines are grouped to form a

"superline" which is identified by a single tag such that all the cache lines within that superline are identified by the single tag.

[0004] Logic executing on a processor may update data in the cache memory. Subsequent to completing any changes to the copy of the data element stored in cache memory, the logic may return, or write-back, the updated data element to non-volatile memory. Further, in some instances, e.g., prior to entry into a low-power state, logic may scrub (i.e., flush), the data stored in the cache memory to nonvolatile memory. Thus, techniques to manage sectored cache may find utility, e.g., in memory systems for electronic devices.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The detailed description is provided with reference to the accompanying figures. The use of the same reference numbers in different figures indicates similar or identical items.

[0006] Fig. 1 is a schematic, block diagram illustration of components of apparatus to implement managing sectored cache in accordance with various examples discussed herein.

[0007] Fig. 2 is a schematic block diagram of a processing environment which may be adapted to implement managing sectored cache in accordance with various examples discussed herein.

[0008] Fig. 3 is a schematic illustration of a sector cache in accordance with various examples discussed herein

[0009] Fig. 4 is a flowchart illustrating operations in a method to manage sector cache in accordance with various examples discussed herein.

[0010] Figs. 5A-5C are diagrams illustrating operations in method to manage sector cache in accordance with various examples discussed herein.

[0011] Figs. 6-10 are schematic, block diagram illustrations of electronic devices which may be adapted to manage sector cache in accordance with various examples discussed herein.

DETAILED DESCRIPTION

[0012] In the following description, numerous specific details are set forth in order to provide a thorough understanding of various examples. However, various examples may be practiced without the specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to obscure the particular examples. Further, various aspects of examples may be performed using various means, such as integrated semiconductor circuits (“hardware”), computer-readable instructions organized into one or more programs (“software”), or some combination of hardware and software. For the purposes of this disclosure reference to “logic” shall mean either hardware, software, or some combination thereof.

[0013] Described herein are techniques to manage sector cache in processing systems of electronic devices. Some examples described herein may find utility in multi-

core processor devices which include integrated cache memory. More particularly, some examples described herein may find utility in electronic devices which flush processor cache to nonvolatile memory in response to a power failure or power interruption.

[0014] In brief, in a sectored cache memory, multiple cache lines (typically 64 bytes) are grouped to form what is commonly referred to as a "superline" which is identified by a single tag.

[0015] Implementing write-back operations in a sectored cache presents certain issues. In a sectored write-back cache a complete superline must be evicted if it is chosen as a replacement (i.e., victim) candidate. Since multiple cache lines within the victim superline may be dirty (i.e., not yet written to main memory), a single eviction of a superline can cause a burst of dirty cache lines to be written back to main memory. In an extreme case, every single cache line in the victim superline may have to be written back to main memory, causing a sudden spurt of demand on memory bandwidth, which results in memory latency. Further, disparities in latency between the high bandwidth data path between the processor and the cache and the relatively lower bandwidth data path to the main memory may result in slow rate of evictions from the cache, which in turn can cause fill data to be dropped because the eviction was not timely. Dropping fill data causes subsequent accesses to these addresses that would have been hits to be misses.

[0016] Another issue with sectored cache is the amount of time required to flush the cache when the electronic device which comprises the cache is moved to a low-power state. This large flush time penalty can raises difficulties in making a quick transition to a low-power state. In order to enter a lower power state in which cache

contents cannot be retained, dirty data in the cache must be flushed to main memory before power can be turned off. In some examples as much as 80% of a cache memory may be dirty. Thus, in some examples a 128MB cache as much as approximately 100MB may be dirty across a wide range of workloads. It can take several milliseconds to flush 100MB of data to main memory, which can cause a significant delay in transitioning to a low-power state.

[0017] Described herein are techniques to manage sectored cache in a way that addresses these and other issues. In a first aspect, miss detection is decoupled from victim selection. In a second aspect, a proactive opportunistic scrubbing mechanism may be implemented to increase the likelihood of finding clean victim superlines. Additional structural features and operations to implement managing sectored cache will be described with reference to Figs. 1-5, below.

[0018] Fig. 1 is a schematic, block diagram illustration of components of apparatus to manage sectored cache in accordance with various examples discussed herein. Referring to Fig. 1, in some examples a processor 100 may comprise one or more processors 110 coupled to a control unit 120 and a local memory 130. Control unit 120 comprises a memory controller 122 and a memory interface 124.

[0019] Memory interface 124 is coupled to one or more remote memory devices 140 by a communication bus 160. Memory device 140 may comprise a controller 142 and one or more memory banks 150. In various examples, at least some of the memory banks 150 may be implemented using nonvolatile memory, e.g., , ferroelectric random-access memory (FeTRAM), nanowire-based non-volatile memory, memory that incorporates memristor technology, a static random access memory (SRAM), three

dimensional (3D) cross point memory such as phase change memory (PCM), spin-transfer torque memory (STT-RAM) or NAND memory. In some examples the memory device(s) 140 may comprise one or more nonvolatile direct in-line memory modules (NVDIMMs) coupled to a memory channel 144 which provides a communication link to controller 142. The specific configuration of the memory bank(s) 150 in the memory device(s) 140 is not critical.

[0020] Fig. 2 is a schematic block diagram of a processing environment which may be adapted to manage sector cache in accordance with various examples discussed herein. Referring to Fig. 2, a processor environment 200 may comprise at least one processor 210 coupled to a power supply 260 and to nonvolatile memory 250. An energy storage device 262 such as one or more capacitors stores a portion of the power provided by power supply 260. As described above, in the event of a power failure such as a condition in which the power supply 260 no longer outputs power in a proper voltage range to power processor environment 200, the energy stored in energy storage resource 262 continues to provide appropriate power to processor 210 for at least a limited amount of holdup time.

[0021] Processor 210 may comprise one or more processor units 220-1, 220-2, which may be referred to herein collectively by reference numeral 220. Processor 210 may further comprise a memory controller 230 having a buffer 232 and a controller 240. Each processor unit 220 may comprise a core 222, cache memory 224, and a cache controller 226, which in turn may include logic which defines a scrub manager 228 and a scrub engine 229.

[0022] The processor core(s) 222 execute logic instructions to implement one or more processing threads. Core(s) 222 may work independently or cooperatively to define a multi-core processing environment. Cache 224 may be implemented as multi-level fast access cache defined in volatile memory, typically on the same die as the core(s) 222.

[0023] Fig. 3 is a schematic illustration of a sectored cache in accordance with various examples discussed herein. Referring to Fig. 3, a sectored cache may comprise two or more ways, indicated in Fig. 3 as WAY0, WAY1, up to WAY N. Each cache superline in each way may comprise multiple cache lines indicated in Fig. 3 as CL0, CL1, CL2, CL3. Each cache superline is identified by a single tag. Cache 224 may be configured as a level 1 (L1) or level 2 (L2) cache or any other level of cache used in a similar memory system. Furthermore, the cache 224 may be configured as an instruction cache, a data cache, or a unified cache for storing both instructions and data.

[0024] In some embodiments, the cache may be integrated with a processor, as illustrated in Fig. 2. In operation, core 222 is configured to communicate addresses and data with cache 224. Cache controller 226 is configured to receive address information from core 222. Cache controller 226 is also configured to receive data from nonvolatile memory 250 and/or core 222. Cache controller 222 is configured to use the address information to select a cache storage line within which to retrieve or store the corresponding data bytes. For example, the cache controller 226 may be configured to access a plurality of tags that track which entries are present in the cache 224.

[0025] The cache controller 226 may be configured to partition each set into one or more sectors, where each sector includes two or more ways, and to maintain one or

more usage status indicators that indicate the relative usage status of these sectors. For example, a least recently used (LRU) usage status pointer may point to the LRU sector in the set. Alternately, status bits and/or counters may be used to track the relative usage status of each sector in the set.

[0026] In some examples controller 226 implements cache management operations to facilitate efficient scrubbing techniques for a sectored cache such as cache 224. Cache management operations implemented by controller 226 will be described with reference to Figs. 4 and 5. In some examples the operations depicted in Figs. 4 and 5 may be implemented by the cache controller, and more particularly by the scrub manager 228 and/or by the scrub engine 229.

[0027] Referring first to Fig. 4, at operation 410 a cache transaction is received in cache controller 226. By way of example, the cache transaction may be a cache read operation. At operation 415 the cache controller 226 performs a lookup operation in cache 224 to locate the data identified in the cache transaction

[0028] At operation 420 the cache controller 226 determines whether the lookup operation for the cache transaction received in operation 410 is resulted in a cache miss (i.e., when the cache 224 does not contain the data requested in the transaction request). If, at operation 420, it is determined that the lookup in operation 415 resulted in a cache miss, then control passes to operation 425 and the cache controller 226 generates a cache scrub hint. In some examples the scrub manager 228 may select one or more cache ways that are dirty and sends these ways as a cache scrub hint to the cache scrub engine. At operation 430 the cache scrub engine 229 generates cacheline writebacks for the superline(s) associated with the cache scrub hint in order to write the data in the

superline(s) back to nonvolatile memory 250. Control then passes to operation 435, where the scrub manager 228 identifies one or more tag(s) for the superlines(s) that were flushed in operation 430 as clean.

[0029] By contrast, if at operation 420 the lookup for the cache transaction did not result in a cache miss (i.e., if the transaction resulted in a cache hit) then control passes to operation 440 and the cache controller 226 determines whether an active scrub feature is enabled. If, at operation 440, an active scrub feature is not enabled then control passes to operation 450 and the results of the cache transaction are sent to the memory controller 230 for main memory 250. At operation 455 the cache controller 226 receives a completion signal from the memory controller 230 for main memory 250, and at operation 460 the cache controller 226 returns a completion signal to the originator of the cache transaction received in operation 410.

[0030] By contrast, if at operation 440 an active scrub feature is enabled then control passes to operation 445 and the cache controller 226 determines whether one or more conditions are met to generate a scrub hint. In some examples the controller may proactively implement cache scrubbing operations in response to one or more conditions. The conditions may correspond to an event (e.g., a cache hit or the elapse of a specified time period) or may correspond to operating conditions of the cache controller 226 (e.g., if the operating capacity of the cache controller 226 is below a threshold). Various conditions will be described below with reference to Figs. 5B and 5C.

[0031] If, at operation 445, one or more conditions are met then control passes to operations 425-435, as described above. By contrast, if at operation 445 one or more

scrub hint conditions are not met then control passes to operations 450-460, as described above.

[0032] Figs. 5A-5C are diagrams illustrating the timing of operations in method to manage sectored cache in accordance with various examples discussed herein. Fig. 5A illustrates operation of the cache controller 226 in the context of a cache miss scenario. Referring to Fig. 5A, at operation 510 the cache manager 228 detects a cache miss (i.e., a failed lookup operation) at a first point in time and, in response to the failed lookup operation the controller sends a cache scrub hint to the cache scrub engine 229. In some examples the cache hint may identify one or more superlines to be scrubbed. At operation 515 the cache scrub engine 229 generates cache writeback operations for each cache line within the superline(s) identified for scrubbing. Multiple writeback operations may be required to clean the superline(s).

[0033] At operation 520 the last dirty cache line in the superline(s) is scrubbed, resulting in a clean cache superline. At operation 525 the cache controller 226 receives fill data from the memory controller 230 and selects a clean victim superline(s) for replacement. Thus, in the context of a cache miss the operations described herein delay the operation of victim selection until after fill data is returned from the memory controller 230, thereby providing a cleaning time window during which the scrub manager 228 and the scrub engine 229 can cooperate to clean one or more superline(s) in the cache.

[0034] Fig. 5B illustrates operation of the cache controller 226 in the context of a cache hit scenario. Referring to Fig. 5B, at operation 530 the cache manager 228 detects a cache hit (i.e., a successful lookup operation) at a first point in time and, in

response to the successful lookup operation the scrub manager 228 generates a scrub hint to send to the scrub engine 229. In one example the scrub manager 228 generates a scrub hint only if the number of dirty ways in a set meets or exceeds a threshold number, which may be fixed or programmable. The scrub manager 228 may use one or more techniques to select a cache hint. For example, scrub manager 228 may select as a cache hint the least recently used (LRU) way in a set that received a cache miss. Alternatively, the scrub manager may select as a cache hint the least recently used (LRU) of all the modified ways in the set. The scrub manager 228 sends the cache scrub hint to the cache scrub engine 229. In some examples the cache hint may identify one or more superlines to be scrubbed.

[0035] At operation 535 the cache scrub engine 229 generates cache writeback operations for each cache line within the superline(s) identified for scrubbing. Multiple writeback operations may be required to clean the superline(s). At operation 540 the last dirty cache line in the superline(s) is scrubbed, resulting in a clean cache superline.

[0036] Fig. 5C illustrates operation of the cache controller 226 in the context of a free-roaming or event-driven scrubbing operation. Referring to Fig. 5C, at operation 550 the cache manager 228 evaluates the degree to which the cache is dirty in order to determine whether to send a scrub hint to send to the scrub engine 229. In one example the scrub manager 228 generates a scrub hint only if the number of dirty ways in a set meets or exceeds a threshold number, which may be fixed or programmable. The scrub manager 228 may use one or more techniques to select a cache hint. For example, scrub manager 228 may select as a cache hint the least recently used (LRU) way in a set that received a cache miss. Alternatively, the scrub manager may select as a cache hint the

least recently used (LRU) of all the modified ways in the set. The scrub manager 228 sends the cache scrub hint to the cache scrub engine 229. In some examples the cache hint may identify one or more superlines to be scrubbed.

[0037] At operation 555 the cache scrub engine 229 generates cache writeback operations for each cache line within the superline(s) identified for scrubbing. Multiple writeback operations may be required to clean the superline(s). At operation 560 the cache scrub engine scrubs one or more superlines in the cache 224. At operation 565 the last dirty cache line in the superline(s) is scrubbed, resulting in a clean cache superline.

[0038] In some examples it may be useful to throttle a free-roaming or event-driven scrubbing operation. An unconstrained scrubbing operation can consume large amounts of power by executing excessive writes and may decrease performance by consuming excessive bandwidth in the memory system. Thus, in some examples the scrubbing operation may be throttled using various techniques. In a first technique, the degree to which a cache is dirty may be used as a triggering mechanism. For example, the number of dirty lines in the cache may be tracked using a simple counter mechanism. The counter may be incremented every time the state of a cache is changed from clean (I/E) to dirty (M) and it is decremented every time a cache line is changed from dirty (M) to clean (E/I). The scrubbing operation is engaged only when the number of dirty lines exceeds a certain programmable threshold. In a second technique the scrubbing operation will be engaged only if there is sufficient spare memory bandwidth available. This ensures that the scrubbing operation does not oversubscribe memory bandwidth due to scrub requests.

[0039] Thus, the operations depicted in Figs. 4-5 enable a processing unit such as cache controller 226 to manage sectored cache.

[0040] As described above, in some examples the electronic device may be embodied as a computer system. Fig. 6 illustrates a block diagram of a computing system 600 in accordance with an example. The computing system 600 may include one or more central processing unit(s) (CPUs) 602 or processors that communicate via an interconnection network (or bus) 604. The processors 602 may include a general purpose processor, a network processor (that processes data communicated over a computer network 603), or other types of a processor (including a reduced instruction set computer (RISC) processor or a complex instruction set computer (CISC)). Moreover, the processors 602 may have a single or multiple core design. The processors 602 with a multiple core design may integrate different types of processor cores on the same integrated circuit (IC) die. Also, the processors 602 with a multiple core design may be implemented as symmetrical or asymmetrical multiprocessors. In an example, one or more of the processors 602 may be the same or similar to the processors 102 of Fig. 1. For example, one or more of the processors 602 may include the control unit 120 discussed with reference to Figs. 1-3. Also, the operations discussed with reference to Figs. 3-5 may be performed by one or more components of the system 600.

[0041] A chipset 606 may also communicate with the interconnection network 604. The chipset 606 may include a memory control hub (MCH) 608. The MCH 608 may include a memory controller 610 that communicates with a memory 612 (which may be the same or similar to the memory 130 of Fig. 1). The memory 412 may store data, including sequences of instructions, that may be executed by the CPU 602, or any

other device included in the computing system 600. In one example, the memory 612 may include one or more volatile storage (or memory) devices such as random access memory (RAM), dynamic RAM (DRAM), synchronous DRAM (SDRAM), static RAM (SRAM), or other types of storage devices. Nonvolatile memory may also be utilized such as a hard disk. Additional devices may communicate via the interconnection network 604, such as multiple CPUs and/or multiple system memories.

[0042] The MCH 608 may also include a graphics interface 614 that communicates with a display device 616. In one example, the graphics interface 614 may communicate with the display device 616 via an accelerated graphics port (AGP). In an example, the display 616 (such as a flat panel display) may communicate with the graphics interface 614 through, for example, a signal converter that translates a digital representation of an image stored in a storage device such as video memory or system memory into display signals that are interpreted and displayed by the display 616. The display signals produced by the display device may pass through various control devices before being interpreted by and subsequently displayed on the display 616.

[0043] A hub interface 618 may allow the MCH 608 and an input/output control hub (ICH) 620 to communicate. The ICH 620 may provide an interface to I/O device(s) that communicate with the computing system 600. The ICH 620 may communicate with a bus 622 through a peripheral bridge (or controller) 624, such as a peripheral component interconnect (PCI) bridge, a universal serial bus (USB) controller, or other types of peripheral bridges or controllers. The bridge 624 may provide a data path between the CPU 602 and peripheral devices. Other types of topologies may be utilized. Also, multiple buses may communicate with the ICH 620, e.g., through multiple bridges

or controllers. Moreover, other peripherals in communication with the ICH 620 may include, in various examples, integrated drive electronics (IDE) or small computer system interface (SCSI) hard drive(s), USB port(s), a keyboard, a mouse, parallel port(s), serial port(s), floppy disk drive(s), digital output support (e.g., digital video interface (DVI)), or other devices.

[0044] The bus 622 may communicate with an audio device 626, one or more disk drive(s) 628, and a network interface device 630 (which is in communication with the computer network 603). Other devices may communicate via the bus 622. Also, various components (such as the network interface device 630) may communicate with the MCH 608 in some examples. In addition, the processor 602 and one or more other components discussed herein may be combined to form a single chip (e.g., to provide a System on Chip (SOC)). Furthermore, the graphics accelerator 616 may be included within the MCH 608 in other examples.

[0045] Furthermore, the computing system 600 may include volatile and/or nonvolatile memory (or storage). For example, nonvolatile memory may include one or more of the following: read-only memory (ROM), programmable ROM (PROM), erasable PROM (EPROM), electrically EPROM (EEPROM), a disk drive (e.g., 628), a floppy disk, a compact disk ROM (CD-ROM), a digital versatile disk (DVD), flash memory, a magneto-optical disk, or other types of nonvolatile machine-readable media that are capable of storing electronic data (e.g., including instructions).

[0046] Fig. 7 illustrates a block diagram of a computing system 700, according to an example. The system 700 may include one or more processors 702-1 through 702-N (generally referred to herein as “processors 702” or “processor 702”). The processors

702 may communicate via an interconnection network or bus 704. Each processor may include various components some of which are only discussed with reference to processor 702-1 for clarity. Accordingly, each of the remaining processors 702-2 through 702-N may include the same or similar components discussed with reference to the processor 702-1.

[0047] In an example, the processor 702-1 may include one or more processor cores 706-1 through 706-M (referred to herein as “cores 706” or more generally as “core 706”), a shared cache 708, a router 710, and/or a processor control logic or unit 720. The processor cores 706 may be implemented on a single integrated circuit (IC) chip. Moreover, the chip may include one or more shared and/or private caches (such as cache 708), buses or interconnections (such as a bus or interconnection network 712), memory controllers, or other components.

[0048] In one example, the router 710 may be used to communicate between various components of the processor 702-1 and/or system 700. Moreover, the processor 702-1 may include more than one router 710. Furthermore, the multitude of routers 710 may be in communication to enable data routing between various components inside or outside of the processor 702-1.

[0049] The shared cache 708 may store data (e.g., including instructions) that are utilized by one or more components of the processor 702-1, such as the cores 706. For example, the shared cache 708 may locally cache data stored in a memory 714 for faster access by components of the processor 702. In an example, the cache 708 may include a mid-level cache (such as a level 2 (L2), a level 3 (L3), a level 4 (L4), or other levels of cache), a last level cache (LLC), and/or combinations thereof. Moreover,

various components of the processor 702-1 may communicate with the shared cache 708 directly, through a bus (e.g., the bus 712), and/or a memory controller or hub. As shown in Fig. 7, in some examples, one or more of the cores 706 may include a level 1 (L1) cache 716-1 (generally referred to herein as “L1 cache 716”). In one example, the control unit 720 may include logic to implement the operations described above with reference to the memory controller 122 in Fig. 2.

[0050] Fig. 8 illustrates a block diagram of portions of a processor core 706 and other components of a computing system, according to an example. In one example, the arrows shown in Fig. 8 illustrate the flow direction of instructions through the core 706. One or more processor cores (such as the processor core 706) may be implemented on a single integrated circuit chip (or die) such as discussed with reference to Fig. 7. Moreover, the chip may include one or more shared and/or private caches (e.g., cache 708 of Fig. 7), interconnections (e.g., interconnections 704 and/or 112 of Fig. 7), control units, memory controllers, or other components.

[0051] As illustrated in Fig. 8, the processor core 706 may include a fetch unit 802 to fetch instructions (including instructions with conditional branches) for execution by the core 706. The instructions may be fetched from any storage devices such as the memory 714. The core 706 may also include a decode unit 804 to decode the fetched instruction. For instance, the decode unit 804 may decode the fetched instruction into a plurality of uops (micro-operations).

[0052] Additionally, the core 706 may include a schedule unit 806. The schedule unit 806 may perform various operations associated with storing decoded instructions (e.g., received from the decode unit 804) until the instructions are ready for dispatch,

e.g., until all source values of a decoded instruction become available. In one example, the schedule unit 806 may schedule and/or issue (or dispatch) decoded instructions to an execution unit 808 for execution. The execution unit 808 may execute the dispatched instructions after they are decoded (e.g., by the decode unit 804) and dispatched (e.g., by the schedule unit 806). In an example, the execution unit 808 may include more than one execution unit. The execution unit 808 may also perform various arithmetic operations such as addition, subtraction, multiplication, and/or division, and may include one or more an arithmetic logic units (ALUs). In an example, a co-processor (not shown) may perform various arithmetic operations in conjunction with the execution unit 808.

[0053] Further, the execution unit 808 may execute instructions out-of-order. Hence, the processor core 706 may be an out-of-order processor core in one example. The core 706 may also include a retirement unit 810. The retirement unit 810 may retire executed instructions after they are committed. In an example, retirement of the executed instructions may result in processor state being committed from the execution of the instructions, physical registers used by the instructions being de-allocated, etc.

[0054] The core 706 may also include a bus unit 714 to enable communication between components of the processor core 706 and other components (such as the components discussed with reference to Fig. 8) via one or more buses (e.g., buses 804 and/or 812). The core 706 may also include one or more registers 816 to store data accessed by various components of the core 706 (such as values related to power consumption state settings).

[0055] Furthermore, even though Fig. 7 illustrates the control unit 720 to be coupled to the core 706 via interconnect 812, in various examples the control unit 720 may be located elsewhere such as inside the core 706, coupled to the core via bus 704, etc.

[0056] In some examples, one or more of the components discussed herein can be embodied as a System On Chip (SOC) device. Fig. 9 illustrates a block diagram of an SOC package in accordance with an example. As illustrated in Fig. 9, SOC 902 includes one or more Central Processing Unit (CPU) cores 920, one or more Graphics Processor Unit (GPU) cores 930, an Input/Output (I/O) interface 940, and a memory controller 942. Various components of the SOC package 902 may be coupled to an interconnect or bus such as discussed herein with reference to the other figures. Also, the SOC package 902 may include more or less components, such as those discussed herein with reference to the other figures. Further, each component of the SOC package 902 may include one or more other components, e.g., as discussed with reference to the other figures herein. In one example, SOC package 902 (and its components) is provided on one or more Integrated Circuit (IC) die, e.g., which are packaged into a single semiconductor device.

[0057] As illustrated in Fig. 9, SOC package 902 is coupled to a memory 960 (which may be similar to or the same as memory discussed herein with reference to the other figures) via the memory controller 942. In an example, the memory 960 (or a portion of it) can be integrated on the SOC package 902.

[0058] The I/O interface 940 may be coupled to one or more I/O devices 970, e.g., via an interconnect and/or bus such as discussed herein with reference to other figures. I/O device(s) 970 may include one or more of a keyboard, a mouse, a touchpad,

a display, an image/video capture device (such as a camera or camcorder/video recorder), a touch screen, a speaker, or the like.

[0059] Fig. 10 illustrates a computing system 1000 that is arranged in a point-to-point (PtP) configuration, according to an example. In particular, Fig. 10 shows a system where processors, memory, and input/output devices are interconnected by a number of point-to-point interfaces. The operations discussed with reference to Fig. 2 may be performed by one or more components of the system 1000.

[0060] As illustrated in Fig. 10, the system 1000 may include several processors, of which only two, processors 1002 and 1004 are shown for clarity. The processors 1002 and 1004 may each include a local memory controller hub (MCH) 1006 and 1008 to enable communication with memories 1010 and 1012. MCH 1006 and 1008 may include the memory controller 120 and/or logic 125 of Fig. 1 in some examples.

[0061] In an example, the processors 1002 and 1004 may be one of the processors 702 discussed with reference to Fig. 7. The processors 1002 and 1004 may exchange data via a point-to-point (PtP) interface 1014 using PtP interface circuits 1016 and 1018, respectively. Also, the processors 1002 and 1004 may each exchange data with a chipset 1020 via individual PtP interfaces 1022 and 1024 using point-to-point interface circuits 1026, 1028, 1030, and 1032. The chipset 1020 may further exchange data with a high-performance graphics circuit 1034 via a high-performance graphics interface 1036, e.g., using a PtP interface circuit 1037.

[0062] As shown in Fig. 10, one or more of the cores 106 and/or cache 108 of Fig. 1 may be located within the processors 902 and 904. Other examples, however, may exist in other circuits, logic units, or devices within the system 900 of Fig. 9.

Furthermore, other examples may be distributed throughout several circuits, logic units, or devices illustrated in Fig. 9.

[0063] The chipset 920 may communicate with a bus 940 using a PtP interface circuit 941. The bus 940 may have one or more devices that communicate with it, such as a bus bridge 942 and I/O devices 943. Via a bus 944, the bus bridge 943 may communicate with other devices such as a keyboard/mouse 945, communication devices 946 (such as modems, network interface devices, or other communication devices that may communicate with the computer network 803), audio I/O device, and/or a data storage device 948. The data storage device 948 (which may be a hard disk drive or a NAND flash based solid state drive) may store code 949 that may be executed by the processors 902 and/or 904.

[0064] The following pertains to further examples.

[0065] Example 1 is an electronic device comprising at least one processor, a cache memory communicatively coupled to the at least one processor, a controller communicatively coupled to the cache memory and comprising logic, at least partially including hardware logic, to receive a first transaction to operate on a first data element in the cache memory, perform a lookup operation for the first data element in the volatile memory, and in response to a failed lookup operation, to generate a cache scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0066] In Example 2, the subject matter of Example 1 can optionally include logic, at least partially including hardware logic, to scrub a section of the cache memory associated with the one or more tags.

[0067] In Example 3, the subject matter of any one of Examples 1–2 can optionally include logic , at least partially including hardware logic, to designate the section of the cache memory as a potential victim for a replacement operation.

[0068] In Example 4, the subject matter of any one of Examples 1–3 can optionally include logic , at least partially including hardware logic, to implement a replacement operation to overwrite at least a portion of the data in the potential victim.

[0069] In Example 5, the subject matter of any one of Examples 1–4 can optionally include logic , at least partially including hardware logic, to determine whether an active cache scrub capability is activated, and in response to a determination that the active cache scrub capability is activated, to generate a cache scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0070] In Example 6, the subject matter of any one of Examples 1–5 can optionally include logic, at least partially including hardware logic, to locate a least recently used (LRU) cache block and assign the LRU cache block as a cache hint.

[0071] In Example 7, the subject matter of any one of Examples 1–6 can optionally include logic, at least partially including hardware logic, to locate a cache block associated with a cache hit and assign cache block as a cache hint.

[0072] Example 8 is controller communicatively coupled to a cache memory and comprising logic, at least partially including hardware logic, to receive a first transaction to operate on a first data element in the cache memory, perform a lookup operation for the first data element in the volatile memory, and in response to a failed lookup

operation, to generate a cache scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0073] In Example 9, the subject matter of Example 8 can optionally include logic, at least partially including hardware logic, to scrub a section of the cache memory associated with the one or more tags.

[0074] In Example 10, the subject matter of any one of Examples 8-9 can optionally include logic, at least partially including hardware logic, to designate the section of the cache memory as a potential victim for a replacement operation.

[0075] In Example 11, the subject matter of any one of Examples 8-10 can optionally include logic, at least partially including hardware logic, to implement a replacement operation to overwrite at least a portion of the data in the potential victim.

[0076] In Example 12, the subject matter of any one of Examples 8-11 can optionally include logic, at least partially including hardware logic, to determine whether an active cache scrub capability is activated, and in response to a determination that the active cache scrub capability is activated, to generate a cache scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0077] In Example 13, the subject matter of any one of Examples 8-12 can optionally include logic, at least partially including hardware logic, to locate a least recently used (LRU) cache block and assign the LRU cache block as a cache hint.

[0078] In Example 14, the subject matter of any one of Examples 8-13 can optionally include logic, at least partially including hardware logic, to locate a cache block associated with a cache hit and assign cache block as a cache hint.

[0079] Example 15 is an electronic device, comprising at least one processor, a cache memory communicatively coupled to the at least one processor, a controller communicatively coupled to the cache memory and comprising logic, at least partially including hardware logic, to generate a cache scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0080] In Example 16, the subject matter of Example 15 can optionally include logic, at least partially including hardware logic, to scrub a section of the cache memory associated with the one or more tags.

[0081] In Example 17, the subject matter of any one of Examples 15-16 can optionally include logic, at least partially including hardware logic, to monitor a number of dirty cache lines in the cache memory and generate a cache scrub hint only when the number of dirty cache lines exceeds a threshold.

[0082] In Example 18, the subject matter of any one of Examples 15-17 can optionally include logic, at least partially including hardware logic, to implement a replacement operation to overwrite at least a portion of the data in the potential victim.

[0083] Example 18 is a controller communicatively coupled to a cache memory and comprising logic, at least partially including hardware logic, to generate a cache

scrub hint, forward the cache scrub hint to a cache scrub engine, and identify one or more cache lines to scrub based at least in part on the cache scrub hint.

[0084] In Example 19, the subject matter of Example 18 can optionally include logic, at least partially including hardware logic, to scrub a section of the cache memory associated with the one or more tags.

[0085] In Example 20 the subject matter of any one of Examples 18-19 can optionally include logic, at least partially including hardware logic, to monitor a number of dirty cache lines in the cache memory and generate a cache scrub hint only when the number of dirty cache lines exceeds a threshold.

[0086] In various examples, the operations discussed herein, e.g., with reference to Figs. 1-10, may be implemented as hardware (e.g., circuitry), software, firmware, microcode, or combinations thereof, which may be provided as a computer program product, e.g., including a tangible (e.g., non-transitory) machine-readable or computer-readable medium having stored thereon instructions (or software procedures) used to program a computer to perform a process discussed herein. Also, the term “logic” may include, by way of example, software, hardware, or combinations of software and hardware. The machine-readable medium may include a storage device such as those discussed herein.

[0087] Reference in the specification to “one example” or “an example” means that a particular feature, structure, or characteristic described in connection with the example may be included in at least an implementation. The appearances of the phrase “in one example” in various places in the specification may or may not be all referring to the same example.

[0088] Also, in the description and claims, the terms “coupled” and “connected,” along with their derivatives, may be used. In some examples, “connected” may be used to indicate that two or more elements are in direct physical or electrical contact with each other. “Coupled” may mean that two or more elements are in direct physical or electrical contact. However, “coupled” may also mean that two or more elements may not be in direct contact with each other, but may still cooperate or interact with each other.

[0089] Thus, although examples have been described in language specific to structural features and/or methodological acts, it is to be understood that claimed subject matter may not be limited to the specific features or acts described. Rather, the specific features and acts are disclosed as sample forms of implementing the claimed subject matter.

CLAIMS

1. An electronic device comprising:
 - at least one processor;
 - a cache memory communicatively coupled to the at least one processor;
 - a controller communicatively coupled to the cache memory and comprising logic, at least partially including hardware logic, to:
 - receive a first transaction to operate on a first data element in the cache memory;
 - perform a lookup operation for the first data element in the volatile memory; and
 - in response to a failed lookup operation, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.
2. The electronic device of claim 1, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - scrub a section of the cache memory associated with the one or more tags.
3. The electronic device of claim 2, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - designate the section of the cache memory as a potential victim for a replacement operation.
4. The electronic device of claim 3, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - implement a replacement operation to overwrite at least a portion of the data in the potential victim.

5. The electronic device of claim 1 wherein the controller further comprises logic, at least partially including hardware logic, to:
 - determine whether an active cache scrub capability is activated, and;
 - in response to a determination that the active cache scrub capability is activated, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.
6. The electronic device of claim 5, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - locate a least recently used (LRU) cache block; and
 - assign the LRU cache block as a cache hint.
7. The electronic device of claim 5, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - locate a cache block associated with a cache hit; and
 - assign cache block as a cache hint.
8. A controller communicatively coupled to a cache memory and comprising logic, at least partially including hardware logic, to:
 - receive a first transaction to operate on a first data element in the cache memory;
 - perform a lookup operation for the first data element in the volatile memory; and
 - in response to a failed lookup operation, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.

9. The controller of claim 8, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - scrub a section of the cache memory associated with the one or more tags.
10. The controller of claim 9, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - designate the section of the cache memory as a potential victim for a replacement operation.
11. The controller of claim 10, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - implement a replacement operation to overwrite at least a portion of the data in the potential victim.
12. The controller of claim 8, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - determine whether an active cache scrub capability is activated, and;
 - in response to a determination that the active cache scrub capability is activated, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.
13. The controller of claim 12, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - locate a least recently used (LRU) cache block; and
 - assign the LRU cache block as a cache hint.
14. The controller of claim 12, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - locate a cache block associated with a cache hit; and
 - assign cache block as a cache hint.

15. An electronic device comprising:
 - at least one processor;
 - a cache memory communicatively coupled to the at least one processor;
 - a controller communicatively coupled to the cache memory and comprising logic, at least partially including hardware logic, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.
16. The electronic device of claim 15, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - scrub a section of the cache memory associated with the one or more tags.
17. The electronic device of claim 15, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - monitor a number of dirty cache lines in the cache memory; and
 - generate a cache scrub hint only when the number of dirty cache lines exceeds a threshold.
18. A controller communicatively coupled to a cache memory and comprising logic, at least partially including hardware logic, to:
 - generate a cache scrub hint;
 - forward the cache scrub hint to a cache scrub engine; and
 - identify one or more cache lines to scrub based at least in part on the cache scrub hint.
19. The controller of claim 19, wherein the controller further comprises logic, at least partially including hardware logic, to:
 - scrub a section of the cache memory associated with the one or more tags.

20. The controller of claim 19, wherein the controller further comprises logic, at least partially including hardware logic, to:
- monitor a number of dirty cache lines in the cache memory; and
 - generate a cache scrub hint only when the number of dirty cache lines exceeds a threshold.

1/10




FIG. 1




FIG. 2




FIG. 3

4/10




FIG. 4

5/10




FIG. 5A




FIG. 5B





FIG. 5C

6/10



7/10




FIG. 7

8/10




FIG. 8




FIG. 9

10/10




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/029377

A. CLASSIFICATION OF SUBJECT MATTER INV. G06F12/08 G06F12/12 G06F9/44 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, INSPEC, COMPENDEX, IBM-TDB		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	WO 2013/095437 A1 (INTEL CORP [US]; NACHIMUTHU MURUGASAMY K [US]; KUMAR MOHAN [US]) 27 June 2013 (2013-06-27) abstract page 13, line 15 - page 14, line 26 page 34, line 16 - page 37, line 3 page 41, line 24 - page 44, line 7 figures 1, 3, 9-11 -----	1-5, 8-12, 15, 16, 18, 19 6, 7, 13, 14, 17, 20
X Y A	US 2015/039835 A1 (AGARWAL VINEET [IN] ET AL) 5 February 2015 (2015-02-05) abstract paragraph [0013] - paragraph [0027] figure 1 ----- -/-	1, 8, 15, 18 17, 20 2-7, 9-14, 16, 19
☒ 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 7 June 2016		Date of mailing of the international search report 17/06/2016
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 Mandato, Davide

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2016/029377

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 0 529 217 A1 (MOTOROLA INC [US])	6,13
A	3 March 1993 (1993-03-03) abstract page 2, line 40 - line 41 page 5, line 18 - page 6, line 20 -----	1,2,8,9, 15,16, 18,19
Y	EP 0 557 884 A1 (MOTOROLA INC [US])	7,14
A	1 September 1993 (1993-09-01) abstract figure 6 -----	1,2,8,9, 15,16, 18,19

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/029377

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 2013095437	A1	27-06-2013	CN 104115129 A	22-10-2014
			DE 112011106013 T5	02-10-2014
			GB 2514023 A	12-11-2014
			KR 20140098220 A	07-08-2014
			US 2014297919 A1	02-10-2014
			WO 2013095437 A1	27-06-2013

US 2015039835	A1	05-02-2015	NONE	

EP 0529217	A1	03-03-1993	DE 69210714 D1	20-06-1996
			DE 69210714 T2	12-12-1996
			EP 0529217 A1	03-03-1993
			HK 1000866 A1	01-05-1998
			JP 3504282 B2	08-03-2004
			JP H05225061 A	03-09-1993
			US 5586293 A	17-12-1996

EP 0557884	A1	01-09-1993	DE 69326491 D1	28-10-1999
			DE 69331448 D1	14-02-2002
			DE 69331448 T2	20-06-2002
			DE 69332663 D1	27-02-2003
			DE 69332663 T2	19-02-2004
			EP 0557884 A1	01-09-1993
			EP 0838762 A2	29-04-1998
			EP 0895163 A1	03-02-1999
			EP 1286269 A2	26-02-2003
			JP 3285644 B2	27-05-2002
			JP H0612327 A	21-01-1994
			US 5375216 A	20-12-1994
