

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

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2014/025920 A1

(43) International Publication Date

13 February 2014 (13.02.2014)

WIPO | PCT

(51) International Patent Classification:

G06F 15/78 (2006.01) G06F 12/08 (2006.01)

(21) International Application Number:

PCT/US2013/054004

(22) International Filing Date:

7 August 2013 (07.08.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/571,426 10 August 2012 (10.08.2012) US

(71) Applicant: QUALCOMM INCORPORATED [US/US];  
Attn: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121 (US).

(72) Inventors: KANG, Seung H.; 5775 Morehouse Drive, San Diego, California 92121 (US). ZHU, Xiaochun; 5775 Morehouse Drive, San Diego, California 92121 (US). WU, Xiaoxia; 5775 Morehouse Drive, San Diego, California 92121 (US).

(74) Agent: TALPALATSKY, Sam; 5775 Morehouse Drive, San Diego, California 92121 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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).

**Declarations under Rule 4.17:**

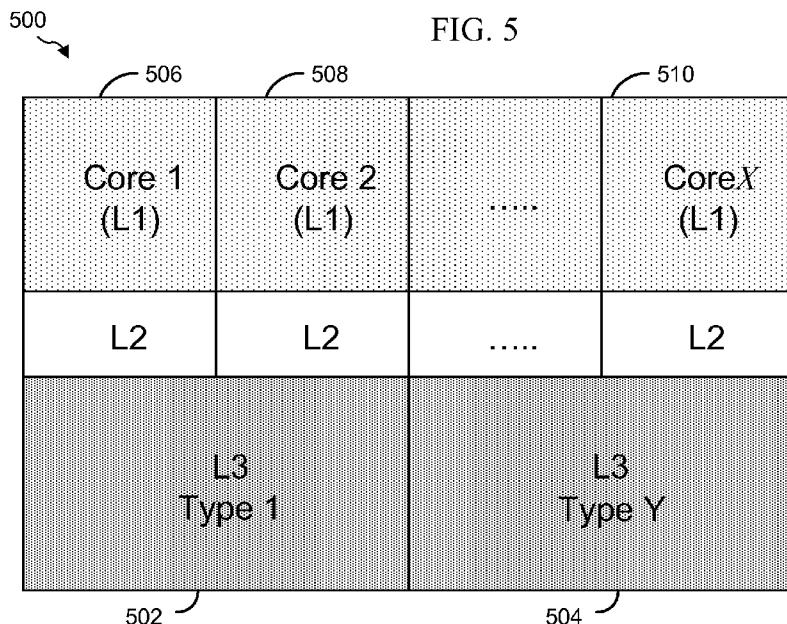
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

**Published:**

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

[Continued on next page]

(54) Title: TUNABLE MULTI-TIERED STT-MRAM CACHE FOR MULTI-CORE PROCESSORS



(57) Abstract: A multi-core processor is presented. The multi-core processor includes a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache associated with a first core of the multi-core processor and tuned according to first attributes and a second STT-MRAM cache associated with a second core of the multi-core processor and tuned according to second attributes.

WO 2014/025920 A1



---

- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

## TUNABLE MULTI-TIERED STT-MRAM CACHE FOR MULTI-CORE PROCESSORS

### TECHNICAL FIELD

**[0001]** The present disclosure relates generally to spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache. More specifically, the disclosure relates to associating and tuning STT-MRAM cache in a multi-core processor.

### BACKGROUND

**[0002]** Typically, multi-core processors, such as the Intel® Core™ i7 Processor, have a multi-tier memory hierarchy. Each core in the multi-core processor may have a dedicated hierarchy of cache (Level 1 (L1) - Level 2 (L2)) and may also share a lower level cache (Level 3 (L3)). The lower level cache tends to have an increased chip area, and therefore, the lower level cache tends to increase the costs associated with multi-core processors.

**[0003]** In some prior art systems, the embedded cache (e.g., L1 - L3) tend to utilize static random access memory (SRAM). SRAM may be desirable because it is logic compatible and fast. Still, SRAM has a large size area that results in a high cost, and furthermore, SRAM tends to leak power. Therefore, the use of SRAM for the lower level cache of a multi-core processor may be undesirable due to, at least, an increased cost and power consumption.

**[0004]** Other prior art systems may have a hybrid cache (e.g., through silicon stacking (TSS) or package on package (PoP)). Hybrid cache refers to a system that utilizes different types of memory in the cache levels. The hybrid cache may be configured as an inter-level hybrid cache, intra-level hybrid cache, or three-dimensional (3D) hybrid cache.

**[0005]** FIGURE 1 illustrates an example of a prior art inter-level hybrid cache 100. As illustrated in FIGURE 1, the core 102 (which includes L1) and L2 cache 104 may use SRAM, and the L3 cache 106 may use a different type of memory, such as embedded dynamic random access memory (eDRAM), MRAM, or phase-change random access memory (PRAM). The inter-level hybrid cache 100 is not limited to the memory types discussed in FIGURE 1 and may utilize various memory types for each cache layer 102-106.

**[0006]** FIGURE 2 illustrates an example of a prior art intra-level hybrid cache 200. As illustrated in FIGURE 2, the core 202 (with L1) and a first L2 cache 204 may use SRAM. Specifically, the first L2 cache 204 may be configured for high speed memory, and therefore, the first L2 cache 204 may utilize SRAM. Moreover, a second L2 cache 206 may be configured for slow memory, and therefore, the second L2 cache 206 may be a type of memory that is different from the type of memory used in the first L2 cache 204. For example the second L2 cache 206 may use eDRAM, MRAM, or PRAM. The intra-level hybrid cache 200 is not limited to the memory types illustrated in FIGURE 2 and may utilize various memory types for each cache layer 202-206.

**[0007]** The inter-level or intra-level hybrid cache may use various memory technologies. Still, the inter-level cache and intra-level hybrid cache are complex and have increased fabrication costs due to the monolithic integration of various memory technologies. Therefore, the use of an inter-level or intra-level hybrid cache for the lower level cache of a multi-core processor is undesirable due to the increased costs and complexity.

**[0008]** FIGURE 3 illustrates an example of a prior art 3D hybrid cache 300. As illustrated in FIGURE 3, the core 302 (with L1) and a first L2 cache 304 may use SRAM. Moreover, a second L2 cache 306 may be configured to utilize slow memory, and therefore, the second L2 cache 306 may utilize a different type of memory, such as eDRAM, MRAM, or PRAM. Furthermore, the core 302, first L2 cache 304, and second L2 cache 306 may all be defined on a single chip. The L3 cache 308, may use a memory such as PRAM, and may be defined on a separate die that is connected to the chip including the core 302, first L2 cache 304, and second L2 cache 306. That is, the L3 cache 308 is on a layer (e.g., chip) that is different from the layer including the core 302, first L2 cache 304, and second L2 cache 306.

**[0009]** The 3D hybrid cache may be desirable due to the use of heterogeneous memories in multiple layers and high-density memories in known good die. Still, a 3D hybrid cache specifies a multi-die solution that is stacked. Accordingly, the 3D hybrid cache may increase costs due to the stacked die. Furthermore, the stacking of the die also increases the overhead (e.g., cost, reliability).

## SUMMARY

**[0010]** In accordance with an aspect of the disclosure, a multi-core processor is presented. The multi-core processor includes a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache associated with a first core of the multi-core processor and tuned according to first attributes. The multi-core processor further includes a second STT-MRAM cache associated with a second core of the multi-core processor and tuned according to second attributes.

**[0011]** In accordance with another aspect, a multi-core processor is presented. The multi-core processor includes a first storage means associated with a first core of the multi-core processor and tuned according to first attributes. The multi-core processor also includes a second storage means associated with a second core of the multi-core processor and tuned according to second attributes.

**[0012]** In accordance with still yet another aspect, a method for associating caches in a multi-core processor is presented. The method includes associating a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache with a first core of the multi-core processor and tuned according to first attributes. The method also includes associating a second STT-MRAM cache with a second core of the multi-core processor and tuned according to second attributes.

**[0013]** In accordance with another aspect, a method for fabricating caches for a multi-core processor is presented. The method includes tuning a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache according to first attributes. The method also includes tuning a second STT-MRAM cache according to second attributes.

**[0014]** This has outlined, rather broadly, the features and technical advantages of the present disclosure in order that the detailed description that follows may be better understood. Additional features and advantages of the disclosure will be described below. It should be appreciated by those skilled in the art that this disclosure may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present disclosure. It should also be realized by those skilled in the art that such equivalent constructions do not depart from the teachings of the disclosure as set forth in the appended claims. The novel features, which are believed to be characteristic of the disclosure, both as to its organization and method of operation, together with further objects and advantages, will

be better understood from the following description when considered in connection with the accompanying figures. It is to be expressly understood, however, that each of the figures is provided for the purpose of illustration and description only and is not intended as a definition of the limits of the present disclosure.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0015]** The features, nature, and advantages of the present disclosure will become more apparent from the detailed description set forth below when taken in conjunction with the drawings.

**[0016]** FIGURES 1-3 illustrate prior art memory systems.

**[0017]** FIGURES 4-5 illustrate a multi-tiered multi-core processor utilizing an STT-MRAM cache according to an aspect of the present disclosure.

**[0018]** FIGURE 6 illustrates a block diagram for a method of associating and tuning an STT-MRAM cache in a multi-core processor according to an aspect of the present disclosure.

**[0019]** FIGURE 7 illustrates an exemplary wireless communication system in which an embodiment of the disclosure may be advantageously employed.

**[0020]** FIGURE 8 is a block diagram illustrating a design workstation used for circuit, layout, and logic design of a semiconductor component according to one aspect of the present disclosure.

## DETAILED DESCRIPTION

**[0021]** The detailed description set forth below, in connection with the appended drawings, is intended as a description of various configurations and is not intended to represent the only configurations in which the concepts described herein may be practiced. The detailed description includes specific details for the purpose of providing a thorough understanding of the various concepts. However, it will be apparent to those skilled in the art that these concepts 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 such concepts.

**[0022]** One aspect of the disclosure provides a lower level cache that may be tunable for speed, power, and density, while decreasing costs and complexity. In one aspect of the disclosure, a multi-tiered STT-MRAM cache for a multi-core processor is described.

**[0023]** Unlike conventional random access memory (RAM) chip technologies, in magnetoresistive RAM (MRAM) data is not stored as electric charge, but is instead stored by magnetic polarization of storage elements. The storage elements are formed from two ferromagnetic layers separated by a tunneling layer. One of the two ferromagnetic layers, which is referred to as the fixed layer or pinned layer, has a magnetization that is fixed in a particular direction. The other ferromagnetic magnetic layer, which is referred to as the free layer, has a magnetization direction that can be altered to represent either “1” when the free layer magnetization is anti-parallel to the fixed layer magnetization or “0” when the free layer magnetization is parallel to the fixed layer magnetization. One such device having a fixed layer, a tunneling layer, and a free layer is a magnetic tunnel junction (MTJ). The electrical resistance of an MTJ depends on whether the free layer magnetization and fixed layer magnetization are parallel or anti-parallel with each other. A memory device such as MRAM is built from an array of individually addressable MTJs.

**[0024]** STT-MRAM is a type of MRAM. The free layer magnetization of STT-MRAM may be switched by an electrical current that passes through an MTJ. Thus, STT-MRAM is differentiated from conventional MRAM which uses a magnetic field. STT-MRAM is tunable for speed, power, density, and cost. STT-MRAM may be tailored as an alternative to memories utilized for lower level cache. Furthermore, STT-MRAM cells and macros can be fabricated in multiple configurations (e.g., multi-tiered) on a monolithic die without incurring extra process steps and cost. Thus, it may be desirable to utilize the STT-MRAM as a multi-tiered lower level cache for each processor core. It should be noted that the STT-MRAM is not limited to the lower level (L3) cache and may be utilized for other cache layers (e.g., L1 or L2) as desired.

**[0025]** In a multi-core processor, each processor core may process different workloads, and therefore, each processor core may specify a different type of memory. For example, one processor core may specify a memory tuned for high performance while a

second processor core may specify a memory tuned for low power. In a typical multi-core processor, the lower level cache is the same circuitry and is not tuned for each core.

According to the present aspect, the multi-tiered STT-MRAM cache may be tuned (e.g., customized) for each processor core so that the tailored multi-tiered STT-MRAM cache may improve the efficiency of each processor core, and thereby, deliver improved chip-level power-performance. According to some aspects, each multi-tiered STT-MRAM may be shared by at least two processor cores. In the present disclosure, the term processor core may sometimes be referred to as a core.

**[0026]** FIGURE 4 illustrates a multi-tiered multi-core processor 400 utilizing an STT-MRAM cache according to an aspect of the present disclosure. As illustrated in FIGURE 4, each core 408-412 (Core 1-Core X) of a multi-core processor is associated with tiers of a cache (L1-L3). The L1 and L2 caches may be dedicated to each core. It should be noted that X is equal to the number of cores and is equal to or greater than the number of STT-MRAM types that may be utilized for the cache of each core. Furthermore, Y is equal to the number to STT-MRAM cache types and Y is equal to or less than X.

**[0027]** As illustrated in FIGURE 4, each lower level STT-MRAM cache 402-406 may be tuned according to specific attributes (Type 1-Type Y). According to one aspect, the Type 1 STT-MRAM cache 402 may be tuned with a first set of attributes for fast access (e.g., increased performance). That is, the Type 1 STT-MRAM cache 402 may be tuned to include a short latency and a fast cycle time. The size of the Type 1 STT-MRAM cache 402 cell may also be tuned to be greater (e.g., 2 transistors (T)/1 MTJ) than a typical cell size of other types of an STT-MRAM cache. According to some aspects, the size of the MTJ of the Type 1 STT-MRAM cache 402 may be smaller than the size of the MTJ of other types of STT-MRAM cache. The decreased size of the MTJ may increase the speed of a STT-MRAM cache due to an increased current density.

**[0028]** Furthermore, the Type Y STT-MRAM cache 404 may be tuned with a second set of attributes. In this example, the Type Y STT-MRAM cache 404 may be tuned for a higher density. That is, the Type Y STT-MRAM cache 404 may be tuned to have a density that is greater than a density of other types of STT-MRAM cache. For example, the Type 1 STT-MRAM cache 402 may have a density of 8 megabits (Mb) and the Type Y STT-MRAM cache 404 may have a density of 16-32 Mb. The increased density of the Type Y STT-MRAM cache 404 increases the cache size, thereby, increasing a cache hit rate and

decreasing a cache miss rate. To increase the density of the Type  $Y$  STT-MRAM cache 404, the cell size of the Type  $Y$  STT-MRAM cache 404 may be less than the size of other types of STT-MRAM. For example, the cell size of the Type  $Y$  STT-MRAM cache 404 may be 1 transistor-1 MTJ (1T-1J). Because the Type  $Y$  STT-MRAM cache 404 has an increased density, the Type  $Y$  STT-MRAM cache 404 may store data such as programming code for power-down and instant-on (e.g., portions of the operating system (OS)).

**[0029]** According to some aspects, the attributes may be similar for all or some STT-MRAM cache types. For example, the Type 2 STT-MRAM cache 406 may have the same attributes as the Type 1 STT-MRAM cache 402 or the Type  $Y$  STT-MRAM cache 404. Alternatively, the attributes may be different for each type of STT-MRAM. For example, each of the attributes of the Type 1 STT-MRAM cache 402, Type 2 STT-MRAM cache 406, and Type  $Y$  STT-MRAM cache 404 may be different. Furthermore, the tuning of the types of STT-MRAM cache is not limited to the configurations described for FIGURE 4. That is, the STT-MRAM may be tuned for various configurations, such as, for example, security, performance, reliability, data type, power, usage, density, and other configurations associated with various memory types.

**[0030]** FIGURE 5 illustrates a multi-tiered multi-core processor 500 utilizing STT-MRAM cache according to another aspect of the present disclosure. As illustrated in FIGURE 5, each core 506-510 (Core 1-Core  $X$ ) of a multi-core processor is associated with a high level cache (L1) and mid-level cache (L2), furthermore, a lower level cache (L3) is shared between the cores. It should be noted that  $X$  is equal to the number of cores and the STT-MRAM cache may be shared between the cores when  $X$  is greater than two. Furthermore,  $Y$  is equal to the number to STT-MRAM cache types and  $Y$  is equal to or less than  $X$ .

**[0031]** As illustrated in FIGURE 5, each lower level STT-MRAM cache may be tuned according to specific attributes (Type 1-Type  $Y$ ). Furthermore, the Type 1 STT-MRAM cache 502 may be a shared cache for Core 1 506 and Core 2 510 and the Type  $Y$  STT-MRAM cache 504 may be a shared cache for all other cores. According to one aspect, the Type 1 STT-MRAM cache 502 may be tuned according to a first set of attributes and the Type  $Y$  STT-MRAM cache 504 may be tuned according to a second set of attributes. Furthermore, the sharing of the STT-MRAM cache is not limited to the aspect illustrated in FIGURE 5, the STT-MRAM cache may be shared by more than two cores and some cores

may not share an STT-MRAM cache in the shared configuration. That is, some types of STT-MRAM cache may only be associated with one core while other types of STT-MRAM cache may be shared amongst cores.

**[0032]** As previously discussed, the types of STT-MRAM may be tuned according to specific attributes. The attributes may include latency, cache hit/miss rate, energy, energy-delay product (e.g., combination of performance and energy consumption), data utilization (e.g., duty factor: read, write, or hold/store), and data reliability and security (e.g., strong data retention and protection). The aforementioned attributes are examples of various attributes that may be configured to tune the types of STT-MRAM cache, the present disclosure is not limited to the aforementioned attributes and is contemplated for other attributes. It should be noted that in the present disclosure attributes include parameters.

**[0033]** The various attributes allow for each STT-MRAM cache to be tailored for performance (e.g., latency, cache hit/miss rate), data type usage, reliability, power, and density. Furthermore, the STT-MRAM cache utilizes a homogeneous lower level process integration resulting in no extra process overhead because the different tiers of the lower level cache are simultaneously fabricated. That is, the same fabrication processes may be used for each type of STT-MRAM cache with only slight variation in the overall process flow. For example, different mask layouts may be used within the same process flow to create the differently tuned caches.

**[0034]** Moreover, the use of an STT-MRAM cache may decrease chip costs. That is, the STT-MRAM cache may use less space in comparison to the SRAM cache. For example, according to some aspects, the STT-MRAM cache may save approximately two to three times the area in comparison to an SRAM lower level cache of the same density.

**[0035]** FIGURE 6 illustrates a block diagram of a method for associating and tuning STT-MRAM caches 600. As illustrated in FIGURE 6, as shown in block 602, a first STT-MRAM cache is associated with a first core. As shown in block 604, the first STT-MRAM is tuned according to first attributes. Furthermore, as shown in block 606, a second STT-MRAM cache is associated with a second core. Additionally, the second STT-MRAM is tuned according to second attributes, as shown in block 608. The method may also include associating at least one other core with the first STT-MRAM cache or the second STT-MRAM, as shown in block 610.

**[0036]** FIGURE 7 shows an exemplary wireless communication system 700 in which an embodiment of the disclosure may be advantageously employed. For purposes of illustration, FIGURE 7 shows three remote units 720, 730, and 750 and two base stations 740. It will be recognized that wireless communication systems may have many more remote units and base stations. Remote units 720, 730, and 750 include multi-core processors with lower level STT-MRAM cache 725A, 725B, 725C. FIGURE 7 shows forward link signals 770 from the base stations 740 and the remote units 720, 730, and 750 and reverse link signals 710 from the remote units 720, 730, and 750 to base stations 740.

**[0037]** In FIGURE 7, the remote unit 720 is shown as a mobile telephone, remote unit 730 is shown as a portable computer, and remote unit 750 is shown as a fixed location remote unit in a wireless local loop system. For example, the remote units may be cell phones, hand-held personal communication systems (PCS) units, a set top box, a music player, a video player, an entertainment unit, a navigation device, portable data units, such as personal data assistants, or fixed location data units such as meter reading equipment. Although FIGURE 7 illustrates remote units, which may employ multi-core processors with lower level STT-MRAM cache 725A, 725B, 725C according to the teachings of the disclosure, the disclosure is not limited to these exemplary illustrated units. For instance, multi-core processors with lower level STT-MRAM cache according to aspects of the present disclosure may be suitably employed in any device.

**[0038]** FIGURE 8 is a block diagram illustrating a design workstation used for circuit, layout, and logic design of a semiconductor component, such as the multi-core processor with lower level STT-MRAM cache disclosed above. A design workstation 800 includes a hard disk 801 containing operating system software, support files, and design software such as Cadence or OrCAD. The design workstation 800 also includes a display 802 to facilitate design of a circuit 810 or a semiconductor component 812 such as the multi-core processor with lower level STT-MRAM. A storage medium 804 is provided for tangibly storing the circuit design 810 or the semiconductor component 812. The circuit design 810 or the semiconductor component 812 may be stored on the storage medium 804 in a file format such as GDSII or GERBER. The storage medium 804 may be a CD-ROM, DVD, hard disk, flash memory, or other appropriate device. Furthermore, the design workstation 800 includes a drive apparatus 803 for accepting input from or writing output to the storage medium 804.

**[0039]** Data recorded on the storage medium 804 may specify logic circuit configurations, pattern data for photolithography masks, or mask pattern data for serial write tools such as electron beam lithography. The data may further include logic verification data such as timing diagrams or net circuits associated with logic simulations. Providing data on the storage medium 804 facilitates the design of the circuit design 810 or the semiconductor component 812 by decreasing the number of processes for designing semiconductor wafers.

**[0040]** In one configuration, the memory apparatus includes a first storage means associated with a first core of the multi-core processor and tuned according to first attributes. The memory apparatus also includes a second storage means associated with a second core of the multi-core processor and tuned according to second attributes. The storage means may be a lower level STT-MRAM cache 402-406 502-504, a level two (L2) cache, and/or other memory types configured to perform the functions recited by the identifying means.

**[0041]** Although specific circuitry has been set forth, it will be appreciated by those skilled in the art that not all of the disclosed circuitry is required to practice the disclosed embodiments. Moreover, certain well known circuits have not been described, to maintain focus on the disclosure.

**[0042]** The methodologies described herein may be implemented by various means depending upon the application. For example, these methodologies may be implemented in hardware, firmware, software, or any combination thereof. For a hardware implementation, the processing units may be implemented within one or more application specific integrated circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), processors, controllers, micro-controllers, microprocessors, electronic devices, other electronic units designed to perform the functions described herein, or a combination thereof.

**[0043]** For a firmware and/or software implementation, the methodologies may be implemented with modules (e.g., procedures, functions, and so on) that perform the functions described herein. Any machine or computer readable medium tangibly embodying instructions may be used in implementing the methodologies described herein. For example, software code may be stored in a memory and executed by a processor. When executed by the processor, the executing software code generates the operational environment that implements the various methodologies and functionalities of the different aspects of the

teachings presented herein. Memory may be implemented within the processor or external to the processor. As used herein, the term “memory” refers to any type of long term, short term, volatile, nonvolatile, or other memory and is not to be limited to any particular type of memory or number of memories, or type of media upon which memory is stored.

**[0044]** The machine or computer readable medium that stores the software code defining the methodologies and functions described herein includes physical computer storage media. A storage medium may be any available medium that can be accessed by a computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to store desired program code in the form of instructions or data structures and that can be accessed by a computer. As used herein, disk and/or disc includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and blu-ray disc where *disks* usually reproduce data magnetically, while *discs* reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer readable media.

**[0045]** In addition to storage on computer readable medium, instructions and/or data may be provided as signals on transmission media included in a communication apparatus. For example, a communication apparatus may include a transceiver having signals indicative of instructions and data. The instructions and data are configured to cause one or more processors to implement the functions outlined in the claims.

**[0046]** Although the present teachings and their advantages have been described in detail, it should be understood that various changes, substitutions and alterations can be made herein without departing from the technology of the teachings as defined by the appended claims. Moreover, the scope of the present application is not intended to be limited to the particular aspects of the process, machine, manufacture, composition of matter, means, methods and steps described in the specification. As one of ordinary skill in the art will readily appreciate from the disclosure, processes, machines, manufacture, compositions of matter, means, methods, or steps, presently existing or later to be developed that perform substantially the same function or achieve substantially the same result as the corresponding aspects described herein may be utilized according to the present teachings. Accordingly, the appended claims are intended to include within their scope such processes, machines, manufacture, compositions of matter, means, methods, or steps.

PAGE INTENTIONALLY LEFT BLANK

## CLAIMS

What is claimed is:

1. A multi-core processor, comprising:  
a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache associated with a first core of the multi-core processor and tuned according to first attributes; and  
a second STT-MRAM cache associated with a second core of the multi-core processor and tuned according to second attributes.
2. The multi-core processor of claim 1, in which the first STT-MRAM cache and the second STT-MRAM cache are a lower level cache.
3. The multi-core processor of claim 1, in which the first STT-MRAM cache and the second STT-MRAM cache are a mid-level cache.
4. The multi-core processor of claim 1, in which the first STT-MRAM cache and the second STT-MRAM cache are embedded or monolithically integrated with each core.
5. The multi-core processor of claim 1, in which the first attributes and second attributes comprise at least latency, cache hit/miss rate, energy, energy-delay product, data utilization, area size, capacity size, and/or data reliability.
6. The multi-core processor of claim 1, in which each core is associated with one type of STT-MRAM cache.
7. The multi-core processor of claim 1, in which at least one other core is associated with the first STT-MRAM cache or the second STT-MRAM cache.
8. The multi-core processor of claim 1, further comprising a third STT-MRAM cache associated with a third core of the multi-core processor and tuned according to third attributes.

9. The multi-core processor of claim 1, in which the first STT-MRAM cache and the second STT-MRAM cache are fabricated utilizing a homogenous lower level integration process.

10. The multi-core processor of claim 1, integrated in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

11. A multi-core processor, comprising:

a first storage means associated with a first core of the multi-core processor and tuned according to first attributes; and

a second storage means associated with a second core of the multi-core processor and tuned according to second attributes.

12. The multi-core processor of claim 11, integrated in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

13. A method for associating caches in a multi-core processor, comprising:

a step of associating a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache with a first core of the multi-core processor and tuned according to first attributes; and

a step of associating a second STT-MRAM cache with a second core of the multi-core processor and tuned according to second attributes.

14. The method of claim 13, further comprising a step of integrating the multi-core processor in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

15. A method for associating caches in a multi-core processor, comprising:  
associating a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache with a first core of the multi-core processor and tuned according to first attributes; and  
associating a second STT-MRAM cache with a second core of the multi-core processor and tuned according to second attributes.

16. The method of claim 15, in which the first STT-MRAM cache and the second STT-MRAM cache are a lower level cache.

17. The method of claim 15, in which the associating comprises embedding or monolithically integrating the first STT-MRAM cache and the second STT-MRAM cache with each core.

18. The method of claim 15, further comprising integrating the multi-core processor in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

19. A method for fabricating caches for a multi-core processor, the method comprising:

tuning a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache according to first attributes; and  
tuning a second STT-MRAM cache according to second attributes.

20. The method of claim 19, in which the first attributes and second attributes comprise at least latency, cache hit/miss rate, energy, energy-delay product, data utilization, area size, capacity size, and/or data reliability.

21. The method of claim 20, further comprising utilizing a homogenous lower level integration process when tuning the first STT-MRAM cache and the second STT-MRAM cache.

22. The method of claim 19, further comprising integrating the multi-core processor in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

23. A method for fabricating caches for a multi-core processor, the method comprising:

a step of tuning a first spin transfer torque magnetoresistive random-access memory (STT-MRAM) cache according to first attributes; and  
a step of tuning a second STT-MRAM cache according to second attributes.

24. The method of claim 23, further comprising a step of integrating the multi-core processor in a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and/or a fixed location data unit.

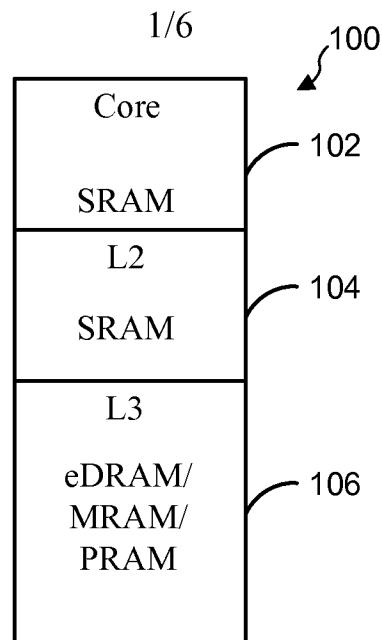


FIG. 1  
(PRIOR ART)

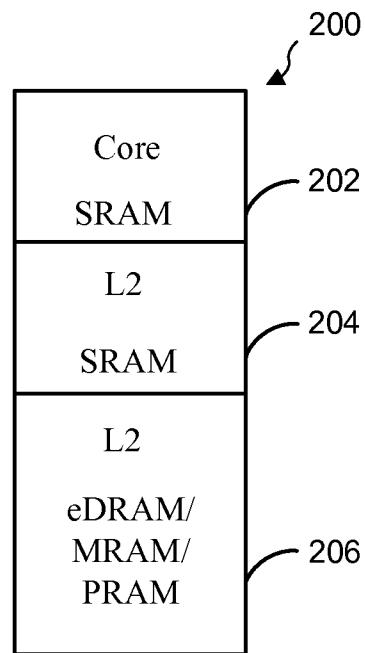


FIG. 2  
(PRIOR ART)

2/6

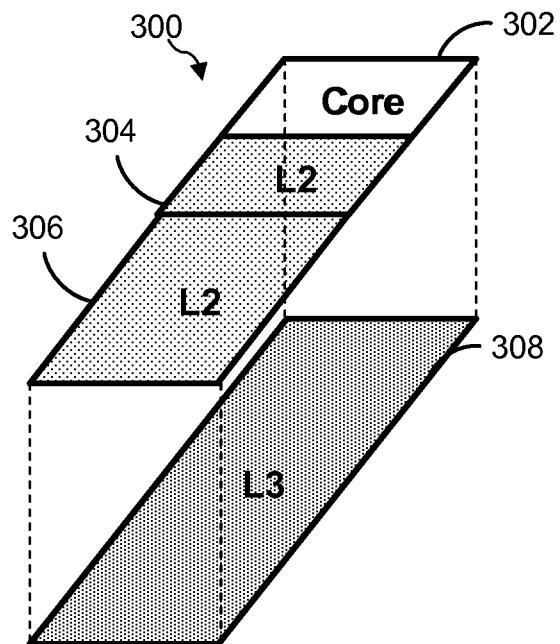


FIG. 3  
(PRIOR ART)

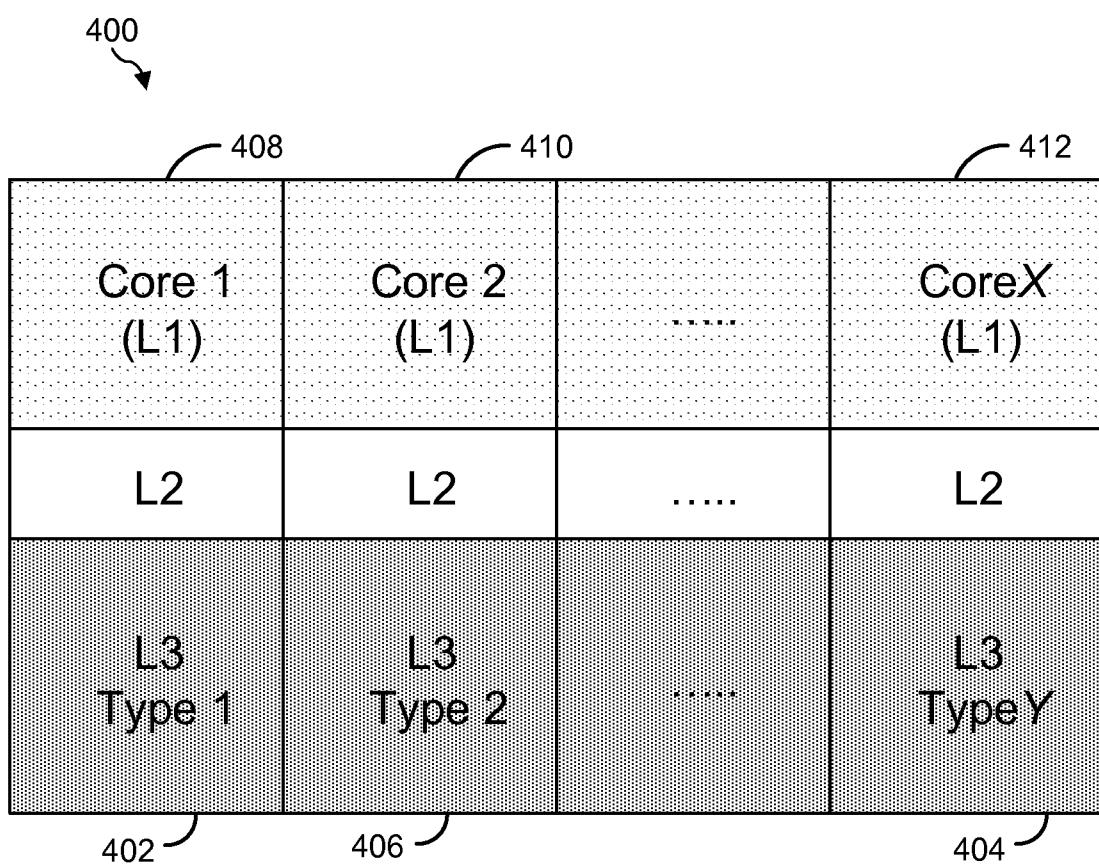


FIG. 4

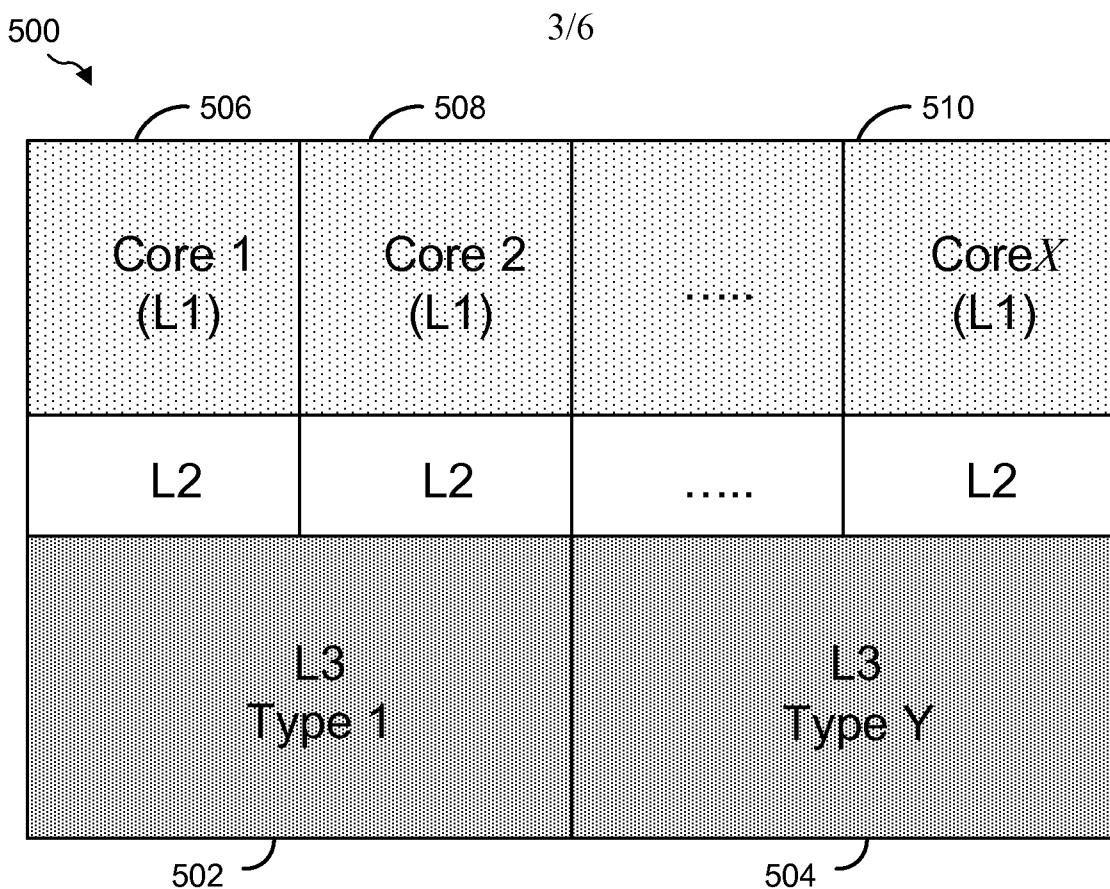
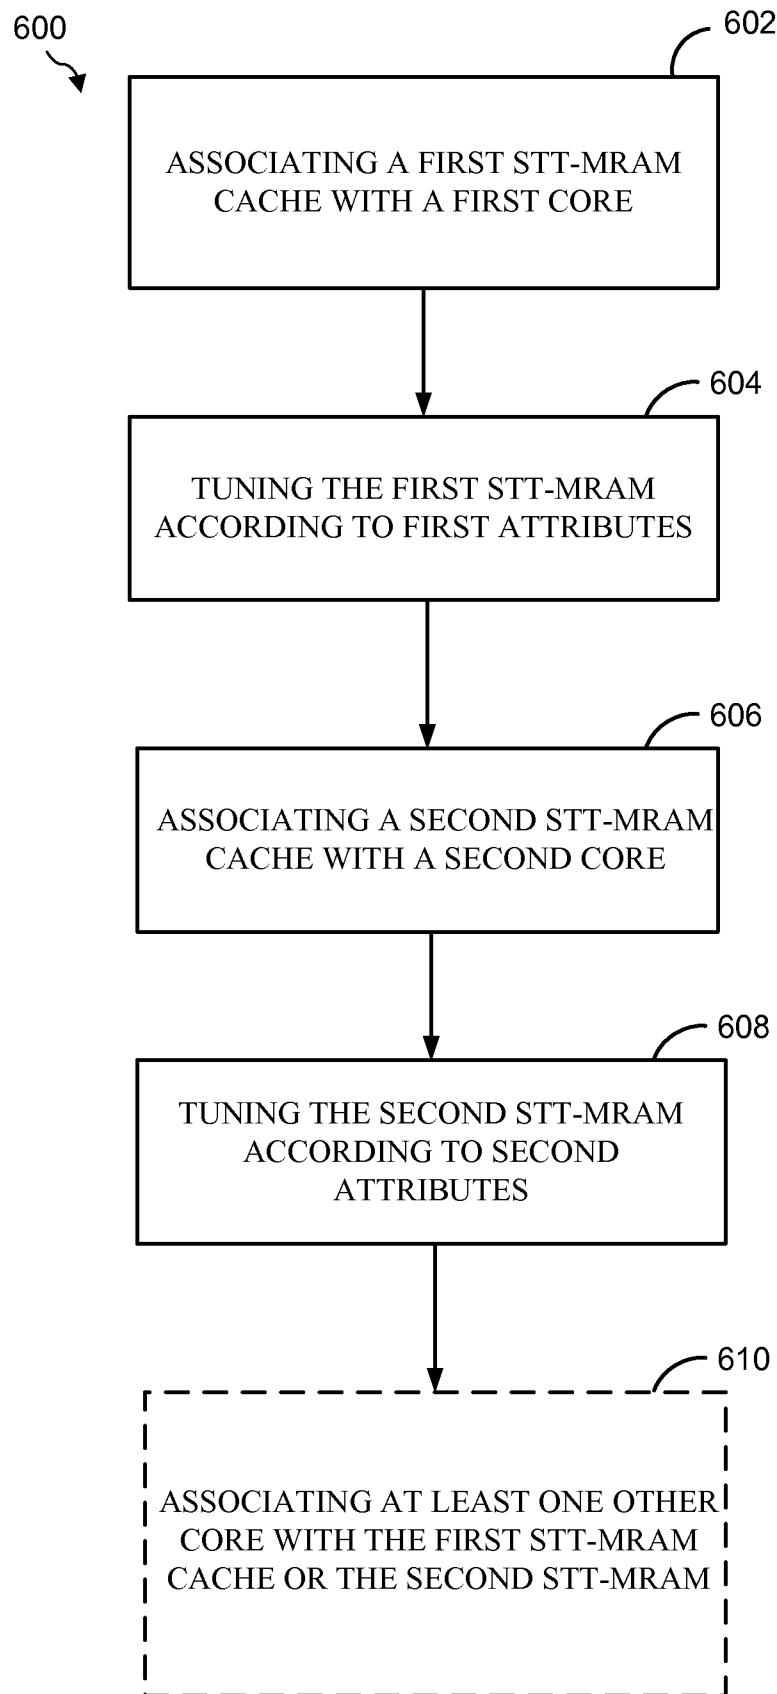


FIG. 5

4/6

**FIG. 6**

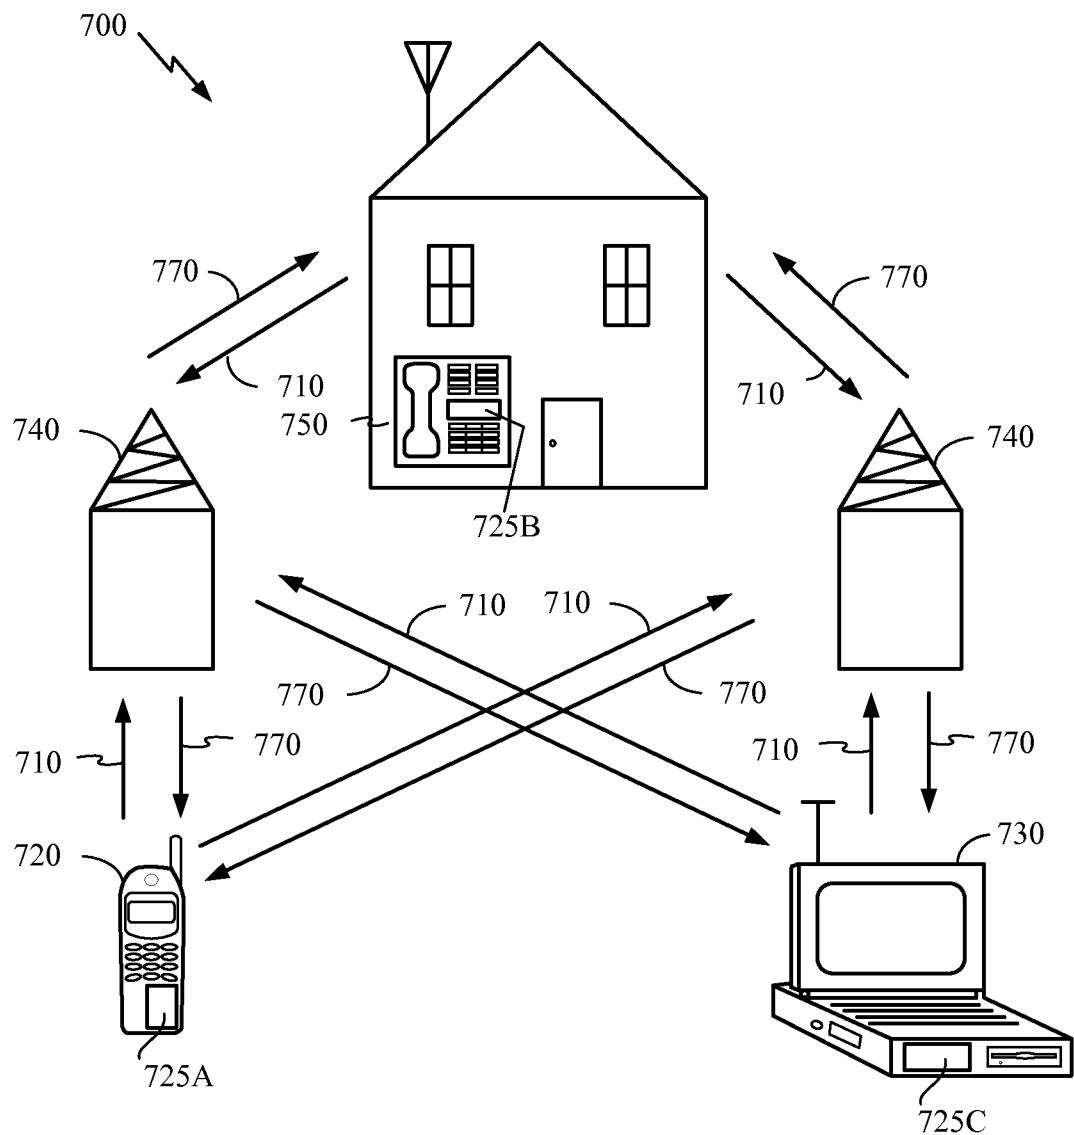


FIG. 7

6/6

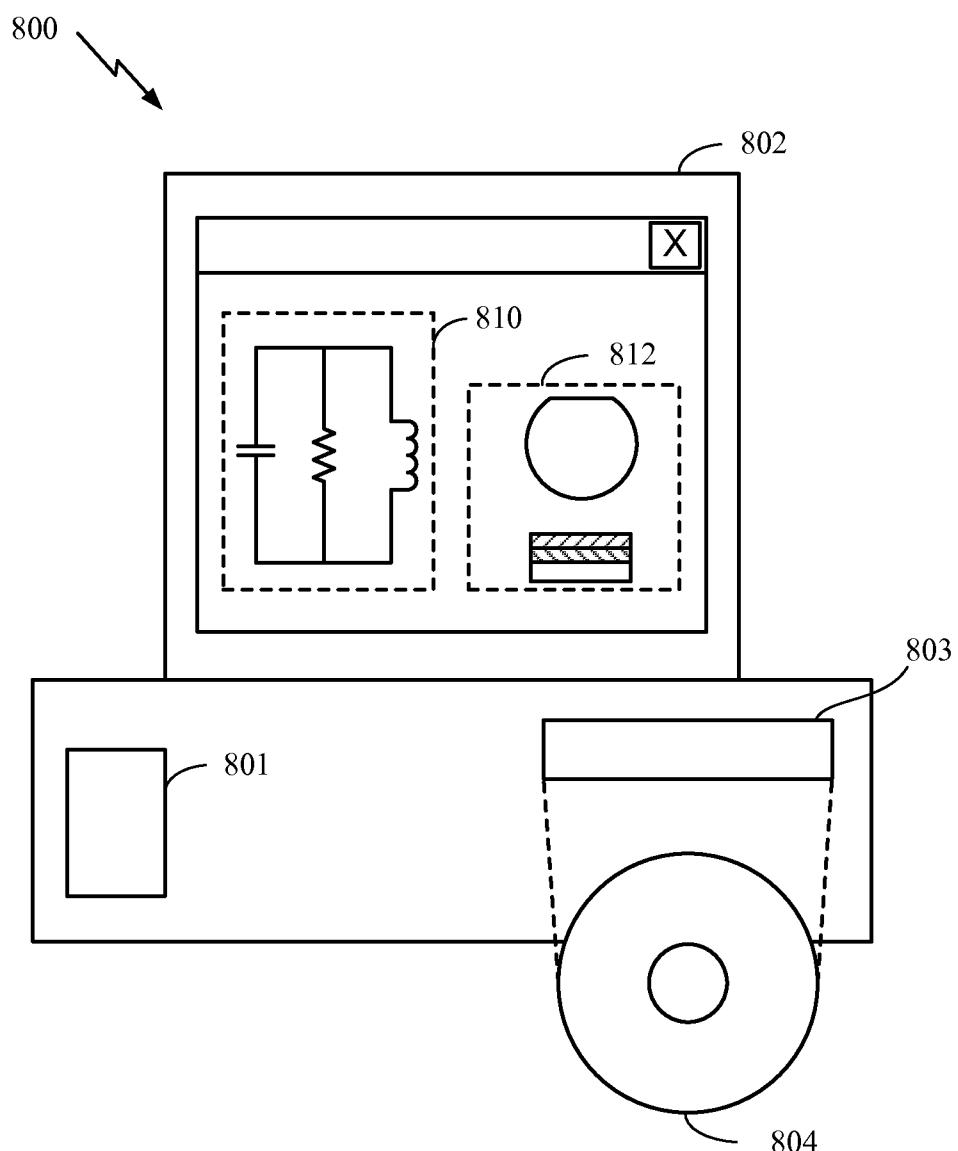


FIG. 8

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2013/054004

**A. CLASSIFICATION OF SUBJECT MATTER**  
INV. G06F15/78 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

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>Texas Instruments Incorporated: "TMS320DM814x DaVinci Video Processors", , 1 July 2011 (2011-07-01), pages 1-7, XP055094683, Retrieved from the Internet: URL:<a href="http://web.archive.org/web/20110830152923/http://www.ti.com/lit/ds/symlink/tms320dm8148.pdf">http://web.archive.org/web/20110830152923/http://www.ti.com/lit/ds/symlink/tms320dm8148.pdf</a> [retrieved on 2014-01-02] figures 1-1-512Kb page 1</p> <p style="text-align: center;">----- -/-</p>	1-24



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
6 January 2014	21/01/2014
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  Bosch Vivancos, P

## INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/054004

## C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ADWAIT JOG ET AL: "Cache revive: Architecting Volatile STT-RAM Caches for Enhanced Performance in CMPs", PROCEEDINGS OF THE 49TH ANNUAL DESIGN AUTOMATION CONFERENCE ON, DAC '12, 1 January 2012 (2012-01-01), page 243, XP055092913, New York, New York, USA DOI: 10.1145/2228360.2228406 ISBN: 978-1-45-031199-1 page 1, left-hand column, line 30 - line 40 page 1, column 2, line 4 - line 9 -----	1,11,13, 15,19,23
A	TONY GIVARGIS ET AL: "Platune: A Tuning Framework for System-on-a-Chip Platforms", IEEE TRANSACTIONS ON COMPUTER AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 21, no. 11, 1 November 2002 (2002-11-01), XP011070650, ISSN: 0278-0070 the whole document -----	1-24