

(43) International Publication Date
29 December 2016 (29.12.2016)

(51) International Patent Classification:

G06F 13/16 (2006.01) *G06F 12/02* (2006.01)
G06F 3/06 (2006.01)

(21) International Application Number:

PCT/US2016/031613

(22) International Filing Date:

10 May 2016 (10.05.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/751,760 26 June 2015 (26.06.2015) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).(72) Inventors: **KUMAR, Karthik**; 714 East Aquarius Place, Chandler, Arizona 85249 (US). **DIMITROV, Martin**; 5000 West Chandler Boulevard, Chandler, Arizona 85226 (US). **WILLHALM, Thomas**; Hardtstr. 51/2, 69207 Sandhausen BW (DE).(74) Agents: **MALLIE, Michael J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, California 94085 (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.

(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, LI, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: INSTANT RESTART IN NON VOLATILE SYSTEM MEMORY COMPUTING SYSTEMS WITH EMBEDDED PROGRAMMABLE DATA CHECKING



(57) Abstract: An apparatus is described. The apparatus includes a memory controller having a programmable component. The programmable component is to implement a data checking function. The programmable component is to receive and process partial results of the data checking function from two or more DIMM cards that are coupled to the memory controller.

Fig.1

INSTANT RESTART IN NON VOLATILE SYSTEM MEMORY COMPUTING SYSTEMS WITH EMBEDDED PROGRAMMABLE DATA CHECKING

Field of Invention

5 The field of invention pertains generally to computing systems, and, more specifically, to instant restart in non volatile system memory computing systems with embedded programmable data checking.

Background

10 Computing systems typically include system memory (or main memory) that contains data and program code of the software code that the system's processor(s) are currently executing. A pertinent bottleneck in many computer systems is the system memory. Here, as is understood in the art, a computing system operates by executing program code stored in system memory. The program code when executed reads and writes data from/to system memory. As
15 such, system memory is heavily utilized with many program codes and data reads as well as many data writes over the course of the computing system's operation. System memory is therefore ideally both high performance yet power efficient.

Figures

 A better understanding of the present invention can be obtained from the following
20 detailed description in conjunction with the following drawings, in which:

Fig. 1 shows a computing system having a multi-level memory;

Fig. 2 shows a memory controller that supports multiple memory channels and associated DIMM cards;

Fig. 3a shows a processing core, memory controller and DIMMs having
25 programmable functionality to perform application specific data checking;

Fig. 3b shows a flow diagram of a process that can be performed on the architecture of Fig. 3a;

Fig. 4 shows an exemplary depiction of a memory page;

Fig. 5 shows a first application software specific memory page data checking scenario;

Fig. 6 shows a second application software specific memory page data checking scenario;

5 **Fig. 7** shows a third application software specific memory page data checking scenario;

Fig. 8 shows various avenues for programming an application software specific memory page data checking function;

Fig. 9 shows an embodiment of a computing system.

10

Detailed Description

One way to implement a high performance system is to have a multi-tiered system memory. Fig. 1 shows an embodiment of a computing system 100 having a multi-tiered or multi-level system memory 112 with a memory side cache. A memory side cache (near memory
15 cache 113) is used to store the system memory items that are (and/or are expected to be) the items that are and/or will be more frequently called upon by the computing system. The near memory cache 113 has lower access times than the lower tiered system memory component (far memory 114). By storing the more frequently called upon items in near memory, the system memory will be observed as faster because the system will often call on items that are being kept
20 in faster near memory 113.

According to some embodiments, for example, the near memory 113 exhibits reduced access times by having a faster clock speed than the far memory 114. Here, the near memory 113 may be a faster, volatile system memory technology (e.g., high performance DRAM). By contrast, far memory 114 may be either a volatile memory technology implemented with a
25 slower clock speed (e.g., a DRAM component that receives a slower clock) or, e.g., a non volatile memory technology that is inherently slower than volatile/DRAM memory.

For example, the lower level 114 may be comprised of an emerging non volatile three dimensional crosspoint random access memory technology such as, to name a few possibilities, a phase change based memory (e.g., SXP or PCMS), a ferro-electric based memory (e.g., FRAM), a magnetic based memory (e.g., MRAM), a spin transfer torque based memory (e.g., STT-
5 RAM), a resistor based memory (e.g., ReRAM) or a “Memristor” based memory. Such emerging non volatile random access memories technologies typically have some combination of the following: 1) higher storage densities than DRAM (e.g., by being constructed in 3D circuit structures); 2) lower power consumption densities than DRAM (e.g., because they do not need refreshing); and/or 3) access latency that is slower than DRAM yet still faster than traditional
10 non-volatile memory technologies such as FLASH. The later characteristic in particular permits an emerging non volatile memory technology to be used in a main system memory role rather than a lower level storage role (which is the traditional architectural location of non volatile storage).

In various embodiments far memory 214 acts as a true system memory in that it
15 supports finer grained data accesses (e.g., cache lines) rather than larger blocked based accesses associated with traditional, lower layer non volatile storage, and/or, otherwise acts as an (e.g., byte) addressable memory that the program code being executed by processor(s) of the CPU operate out of.

Because near memory 113 acts as a cache, near memory 113 may not have its own
20 individual addressing space. Rather, only far memory 114 includes the individually addressable memory space of the computing system’s main memory. In various embodiments near memory 113 truly acts as a cache for far memory 114 rather than acting as a last level CPU cache. Generally, a CPU level cache is able to keep cache lines across the entirety of system memory addressing space that is made available to the processing cores 117 that are integrated on a same
25 semiconductor chip as the memory controller 116.

For example, in various embodiments, referring to Fig. 2, system memory is implemented with DIMM cards 221 plugged into a memory channel 220_1 where a single DIMM card has both DRAM and emerging non volatile memory chips and the DRAM chips effectively act as an on board cache for the non volatile memory chips. Ideally, the more frequently accessed cache lines of any particular DIMM card will be found on that DIMM card's DRAM chips rather than its non volatile memory chips. Given that multiple DIMM cards are typically plugged into a working computing system and each DIMM card is only given a section of the system memory addresses made available to the processing cores of the semiconductor chip or CPU package that the DIMM cards are coupled to, the DRAM chips are acting as a cache for the non volatile memory that they share a DIMM card with rather than a last level CPU cache.

In other configurations DIMM cards having only DRAM chips may be plugged into a same system memory channel (e.g., a DDR channel) with DIMM cards having only non volatile system memory chips. Ideally, the more frequently used cache lines of the channel will be found in the DRAM DIMM cards rather than the non volatile memory DIMM cards. Thus, again, because there are typically multiple memory channels coupled to a same semiconductor chip having multiple processing cores, the DRAM chips are acting as a cache for the non volatile memory chips that they share a same channel with rather than as a last level CPU cache.

In still yet other embodiments, the DRAM devices within system memory (e.g., disposed on DRAM DIMM cards) may be given their own unique system addressing space that is separate than the system addressing space provided to the non volatile system memory devices (e.g., disposed on non volatile system memory DIMM cards). In this case the system memory is architected to have a higher performance address region and a lower performance address region as opposed to a higher region that acts as a cache for a lower region. In still yet other embodiments, system memory only has non volatile memory devices.

Regardless of the specific system memory implementation, the above discussion indicates that non volatile memory devices may be used as true system memory devices rather than traditional lower level storage devices. The use of non volatile system memory provides for various non traditional computing system methodologies.

5 One of these is “instant restart”. Here, when a traditional computing system is powered on, software program code is loaded from lower level traditional storage into system memory. The time consumed between the initial power on of a computing system and the moment the computer is ready to be used, which can be quite lengthy, is largely a function of the software loading process.

10 With non volatile system memory, however, the loading process can be significantly diminished. Here, software program code is simply loaded into system memory where it is also kept as traditional storage. Upon power on of the computer system, the software already exists in system memory – no loading is necessary.

 Instant restart may be particularly helpful after an unforeseen power down event.

15 Here, software is presumably executing out of system memory and is at a certain “state” when the power is removed. That is, the software is along one or more particular execution paths with specific data values assigned to the software’s numerous variables. Theoretically, because of the existence of non volatile system memory, the software could “pick back up right were it left off”. That is, once the computing system is operational again, the software can begin operation again

20 from the state it was at (or just before) the unforeseen power down. Further still, because of the lack of software loading, the software could conceivably restart from this state shortly after the computer is powered on again.

 A complication is the fact that the data needs to be checked before it is deemed correct for further operation. In hardware this can be accomplished, e.g., by performing cyclic

25 redundancy checks (CRCs) on the data. Unfortunately, the software itself is often written to detect an unforeseen power down event and then verify on its own accord that the data is correct.

This involves calling large amounts of data from system memory and having the data executed on by the processing cores of the system. The loading of the data from system memory and its forwarding into the caching levels of the processing cores can consume large amounts of time and defeat the instant restart potential of non volatile system memory.

5 Fig. 3a shows an improved hardware architecture that specially includes programmable logic circuitry to execute the data checking functionality of the application software. As described in more detail below, the data checking functionality of an application software program is “ported” into the hardware as one of various combinations of logic functions $H()$, $G()$ and $F()$. Here, $H()$ can be viewed as programmable circuitry that is resident in a
10 processor core 317 (or other architectural location “north” of a memory controller 316). Likewise, $G()$ can be viewed as programmable circuitry that is resident in a memory controller 316 and $F()$ can be viewed as programmable circuitry that is loaded into one or more DIMM cards or other system memory device package.

Thus, the overall environment is arranged to follow the process illustrated in Fig. 3b.
15 First the data checking functionality of the application software is defined 301 (e.g., by being extracted from the application software or by being provided as a separable or sibling component of it). Then, the data checking functionality is compiled or otherwise converted into information that programs the underlying hardware to perform the data checking functionality 302. As described in more detail below, the information may be partitioned into different architectural
20 components of the hardware (e.g., processing core $H()$, memory controller $G()$, DIMM card $H()$)) based on any interleaving of the application software’s memory pages within system memory. Here, data checking sequences are typically easily partitionable because they have strong commutative properties. That is, the order in which the data is processed is irrelevant. For instance, the data checking function may be implemented as one large XOR function partitioned
25 between $F()$, $G()$ and $H()$. The data checking function can be an arbitrarily complex software operation. However, for cases in which the data checking function is not commutative, the unit

of data being checked should reside at the level of the function being used. For example if the entire data fits inside a single DIMM, only F() is used. Alternately, if the entire data fits within a single channel (memory controller) only G() is used.

Then, upon a restart sequence, the data is checked by the programmed hardware rather than by the software 303. Because the data is checked in the hardware, much of which may be checked “beneath” the processing core(s) (e.g., within a memory controller G() and/or on a DIMM card H()), the transfer of large amounts of data from system memory up into the processing core(s) is avoided. As such, the instant restart is closer to actual realization.

As is understood in the art, a page is a collection of data for an application software program that can be switched into memory and switched out of memory as a unit. Here, one or more pages of information are called into memory by a software program, where, there is typically some limit on the size of the pages and/or the number of pages that can be called into memory by the software program and/or an underlying operating system instant and/or virtual machine monitor. In operation, to effect efficient use of memory, the pages that are called into memory are pages that are needed or believed to be needed, the pages that are switched out of memory are the pages that are not needed or not believed to be needed.

Each page is also viewed as being broken down into smaller individual addressable data units such as cache lines. Fig. 4 shows an exemplary depiction of a page having 8 cache lines labeled a through h. Of course this is only exemplary. Memory pages and their associated data units are apt to be of various sizes resulting in a wide variety of embodiments in terms of number of units per memory page.

Memory interleaving is a technique by which the data units of a page may be spread out across different physical resources, such as across different DIMMs or different memory controllers. Here, software programs may have often call on data on a same page in approximately the same timeframe. By spreading the data units of a same page across multiple physical resources the system memory accesses to these data units may be of a more parallelized

nature than a serialized one. Said another way, if all the data units of a same page were stored in a same memory chip and the application software instantaneously wanted all of the data units, the system would have no choice but to wait while each data unit was accessed individually in succession.

As mentioned above, any interleaving of the data within the software application's memory pages of system memory may be a factor in determining which programmable partition H(), G(), F() of Fig. 3 performs which function. Fig. 5 shows a first possible scenario where an entire page is stored on the same DIMM 521. In this case, the entire data checking function for the software application can be performed by the F() function on the DIMM card. If the memory footprint for the entire application is identically architected, the entire data checking function of all memory pages of the application can be performed by the one or more DIMM cards of the system where the pages of the software application are kept. On each DIMM card, after a first page is checked from the DIMM a next page is checked until all pages of the application on the DIMM are checked.

Fig. 6 shows another scenario where the cache lines of a same page are spread out over different DIMMs. Specifically, a first DIMM card 621_1 has cache lines a,c,e,g and a second DIMM card 621_2 has cache lines b,d,f,h. As such the F() function on DIMM card 621_1 can perform data checking on cache lines a,c,e,g and the F() function on DIMM card 621_2 can perform data checking on cache lines b,d,f,h. The mathematical result of both the F() functions on both DIMMs are akin to partial products or partial results that still need to be combined to complete the data checking function for the entire page. As such, DIMM card 621_1 sends the partial result from cache lines a,c,e,g to the G() function of memory controller 616 and DIMM card 621_2 sends the partial results from cache lines b,d,f,h to the G() function of memory controller 616. The G() function of memory controller then completes the data checking for the entire page by performing calculations with both of the received partial results.

Again, if all the memory pages of the software application are similarly architected, all the memory pages of the software application can be checked according to the same process.

After all DIMM cards have reported all partial results of all memory pages and the memory controller has performed calculations on all partial products from a same page the application's memory pages will have been completely checked. If the memory pages of the application are implemented in system memory as a combination of the scenarios of Figs. 5 and 6, then the application's memory pages can be completely checked through a combination of the processes described above with respect to Figs. 5 and 6.

The scenario of Fig. 6 is directed to an environment where the different DIMM cards 621_1, 621_2 are coupled to a same memory controller 616. As depicted, the different DIMM cards 621_1, 621_2 are coupled to different memory channels that emanate from the same memory controller 616. In other embodiments the different DIMM cards 621_1, 621_2 may be coupled to a same memory channel. Additionally, the scenario of Fig. 6 assumes only two different DIMM cards hold all the cache lines of a single memory page. In alternative embodiments more than two DIMM cards may store the cache lines of a single memory page.

The scenario of Fig. 6 shows only one $G()$ function per memory controller. However, various implementations may include more than one $G()$ function instance per memory controller. For example, a memory controller that supports eight memory channels may have two $G()$ function instances (one per four memory channels). Likewise, more than one $F()$ function instance may reside on any DIMM card or other system memory device package.

Fig. 7 shows yet another scenario where the data units of a same memory page are not only spread out over multiple DIMM cards but are additionally spread out over multiple memory controllers 716_1, 716_2 as well as multiple DIMM cards per memory controller. Specifically, DIMM cards 721_1 and 721_2 are coupled to memory controller 716_1 and DIMM cards 721_3 and 721_4 are coupled to memory controller 716_2.

As observed in Fig. 7, the memory page is interleaved such that DIMM card 721_1 holds cache lines a and e, DIMM card 721_2 holds cache lines c and g, DIMM card 721_3 holds cache lines b and f and DIMM card 721_4 holds cache lines d and h. As such, partial data checking results are calculated by the $F()$ function on DIMM card 721_1 for cache lines a and e and by the $F()$ function on DIMM card 721_2 for cache lines c and g. These partial data checking results are forwarded to the $G()$ function on memory controller 716_1 which creates a combined partial data checking result that accounts for cache lines a, e, c and g.

Similarly, partial data checking results are calculated by the $F()$ function on DIMM card 721_3 for cache lines b and f and by the $F()$ function on DIMM card 721_4 for cache lines d and h. These partial data checking results are forwarded to the $G()$ function on memory controller 716_2 which creates a combined partial data checking result that accounts for cache lines b, f, d and h. The combined partial data checking results from memory controllers 716_1 and 716_2 are then forwarded to a processing core whose $H()$ function performs the final data checking for the memory page. Final checking for a memory page may be performed by comparing the output of the entire checking function against a stored value. The value may be stored in register space or memory space. In various embodiments, the checking of the function output against a stored value may be performed with logic (hardware and/or software) at any of the $F()$, $G()$ and $H()$ levels (for example, if the entire data checking function is realized at the $G()$ level, the checking may be performed at the memory controller, or, at the higher processor level).

In various embodiments, the multiple memory controllers 716_1, 716_2 may exist on a same die (e.g., a set of processor cores are coupled to multiple system memory controllers on a same semiconductor die) or at least within a same CPU package (e.g., multiple system-on-chip (SOC) semiconductor die, each having a memory controller and corresponding set of processor cores). In the case of the later, memory controllers 716_1 and 716_2 are on different

semiconductor die and the H() function of one of the cores of one of the die is used to perform the final data checking calculation for a memory page.

Thus, a data checking sequence for a memory page can be partitioned in hardware depending on the manner in which the memory page has been interleaved in the hardware.

5 The mechanism by which the data checking routine of the software is ported into the hardware may vary from embodiment. Referring to Fig. 8, as is understood in the art, a software application is typically written in source code (e.g., C++). The source code representation of the application's data checking routine 801 may be compiled into lower forms of code such as abstract code 802 or object code 803. Abstract code 802 is a form of code that can be executed
10 by an interpreter or virtual machine that "runs" on a hardware processor. Object code 803 runs directly on a processor.

In various implementations, any of the H(), G() or F() functions is implemented as a micro-controller, processing core or digital signal processor (all of which may be referred to as a "processor") that executes program code out of a local memory (e.g., a memory that is directly or
15 otherwise associated with the micro-controller or processing core). Here, the source code representation of the software application's data checking routine 801 may be compiled into abstract code 802 that is executed by an interpreter or virtual machine that runs on the processor, or, may be compiled into object code that is directly executed by the processor. H() functions may be implemented by the standard processing core of an SOC, but note that the scenarios of
20 Figs. 5, 6 and 7 reduce the amount of information that is directed to the H() function as compared to an approach that performs all data checking with the processing core.

In yet another approach, any of the H(), G() or F() functions may be implemented as programmable logic hardware such as a programmable logic device (PLD) or field programmable gate array (FPGA). Programmable logic hardware does not execute program
25 code like a processor does. Rather, the logic functions to be performed by the hardware are programmed into hardware circuits of the hardware. For example, fuses may be blown to

decouple routes to/from certain logic gates, or, embedded memory (e.g., SRAM) may be loaded with look-up table information that implements the desired logic functionality. In cases where any of the $H()$, $G()$ or $F()$ functions are implemented with programmable logic hardware, the source code of the application software's data checking routine 801 may be compiled into a hardware circuit description format 804 that may be used to ultimately program a programmable logic hardware device. Various implementations may employ a combination of processor(s) and programmable logic hardware devices to implement any of the $H()$, $G()$ and $F()$ functions of a system.

Fig. 9 shows a depiction of an exemplary computing system 900 such as a personal computing system (e.g., desktop or laptop) or a mobile or handheld computing system such as a tablet device or smartphone. The computing system observed in Fig. 9 may include any of the $H()$, $G()$ and $F()$ discussed above.

As observed in Fig. 9, the basic computing system may include a central processing unit 901 (which may include, e.g., a plurality of general purpose processing cores and a main memory controller disposed on an applications processor or multi-core processor), system memory 902, a display 903 (e.g., touchscreen, flat-panel), a local wired point-to-point link (e.g., USB) interface 904, various network I/O functions 905 (such as an Ethernet interface and/or cellular modem subsystem), a wireless local area network (e.g., WiFi) interface 906, a wireless point-to-point link (e.g., Bluetooth) interface 907 and a Global Positioning System interface 908, various sensors 909_1 through 909_N (e.g., one or more of a gyroscope, an accelerometer, a magnetometer, a temperature sensor, a pressure sensor, a humidity sensor, etc.), a camera 910, a battery 911, a power management control unit 912, a speaker and microphone 913 and an audio coder/decoder 914. The display 903 may be an LED display that is driven by an LED driver circuit as described herein.

An applications processor or multi-core processor 950 may include one or more general purpose processing cores 915 within its CPU 901, one or more graphical processing units

916, a memory management function 917 (e.g., a memory controller) and an I/O control function 918. The general purpose processing cores 915 typically execute the operating system and application software of the computing system. The graphics processing units 916 typically execute graphics intensive functions to, e.g., generate graphics information that is presented on the display 903. The memory control function 917 interfaces with the system memory 902. The system memory 902 may be a multi-level system memory. During operation, data and/or instructions are typically transferred between a lowest level non volatile (e.g., “disk”) storage 920 and system memory 902. The power management control unit 912 generally controls the power consumption of the system 900.

Each of the touchscreen display 903, the communication interfaces 904 – 907, the GPS interface 908, the sensors 909, the camera 910, and the speaker/microphone codec 913, 914 all can be viewed as various forms of I/O (input and/or output) relative to the overall computing system including, where appropriate, an integrated peripheral device as well (e.g., the camera 910). Depending on implementation, various ones of these I/O components may be integrated on the applications processor/multi-core processor 950 or may be located off the die or outside the package of the applications processor/multi-core processor 950.

Embodiments of the invention may include various processes as set forth above. The processes may be embodied in machine-executable instructions. The instructions can be used to cause a general-purpose or special-purpose processor to perform certain processes. Alternatively, these processes may be performed by specific hardware components that contain hardwired logic for performing the processes, or by any combination of programmed computer components and custom hardware components.

Elements of the present invention may also be provided as a machine-readable medium for storing the machine-executable instructions. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, FLASH memory, ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards,

propagation media or other type of media/machine-readable medium suitable for storing electronic instructions. For example, the present invention may be downloaded as a computer program which may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

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

Claims

5

What is claimed is:

1. An apparatus, comprising:

10 a memory controller comprising a programmable component, said programmable component to implement a data checking function, said programmable component to receive and process partial results of said data checking function from two or more DIMM cards that are coupled to said memory controller.

15 2. The apparatus of claim 1 wherein said programmable component comprises a processor.

3. The apparatus of claim 1 wherein said programmable component comprises a programmable logic hardware component.

20 4. The apparatus of claim 3 wherein said programmable logic hardware component comprises a programmable logic device.

5. The apparatus of claim 3 wherein said programmable hardware component comprises a field programmable gate array.

25

6. The apparatus of claim 1 wherein said two or more DIMMs are coupled to said memory controller through at least two memory channels.

7. The apparatus of claim 1 wherein said memory controller is coupled to a processing core
30 having a second programmable component to receive and process a partial result of said data checking function from said programmable component.

8. A computing system, comprising:

 one or more DIMM cards having a respective first programmable component to
35 implement a data checking function;

a memory controller having a second programmable component to implement said data checking function, said memory controller to receive and process partial results of said data checking function from said one or more DIMM cards.

5 9. The computing system of claim 8 wherein said memory controller includes a third programmable component to receive and process other partial results of said data checking function from one or more other DIMM cards that said memory controller is coupled to.

10 10. The computing system of claim 9 further comprising a second memory controller having a third programmable component to implement said data checking function, said memory controller to receive and process partial results of said data checking function from one or more other DIMM cards that said second memory controller is coupled to.

15 11. The computing system of claim 10 wherein said first and second memory controllers are coupled to a processing core having a fourth programmable component to receive and process partial results of said data checking function from said first and second memory controllers.

20 12. The computing system of claim 10 wherein said first and second memory controllers are disposed on different semiconductor die.

13. The apparatus of claim 8 wherein said first programmable component is any of:
a processor;
a programmable logic hardware component.

25 14. The apparatus of claim 8 wherein said second programmable component is any of:
a processor;
a programmable logic hardware component.

30 15. An apparatus, comprising:
means for programming a first component of a data checking function into a DIMM card;
means for programming a second component of said data checking function into a memory controller to which said DIMM card is coupled;

means for executing said first component of said data checking function on said DIMM card to produce a partial result of said data checking function and forwarding said partial result to a memory controller to which said DIMM card is coupled;

5 means for executing said second component of said data checking function with said partial result on said memory controller.

16. The apparatus of claim 15 wherein said means for programming of said first component includes any of:

10 means for loading abstract program code onto a processor on said DIMM card;
means for loading object code onto a processor on said DIMM card;
means for loading a hardware description into a programmable logic hardware component on said DIMM card.

15 17. The apparatus of claim 16 wherein said means for programming of said second component includes any of:

means for loading abstract program code onto a processor on said DIMM card;
means for loading object code onto a processor on said DIMM card;
means for loading a hardware description into a programmable logic hardware component on said DIMM card.

20

18. The apparatus of claim 16 further comprising

means for programming a respective first component of said data checking function into a second DIMM card;

25 means for executing said respective first component of said data checking function on said second DIMM card to produce a second partial result of said data checking function and forwarding said second partial result to a memory controller.

19. The apparatus of claim 16 further comprising:

30 means for programming a respective first component of said data checking function into a second DIMM card;

means for executing said respective first component of said data checking function on said second DIMM card to produce a second partial result of said data checking function and forwarding said second partial result to a second memory controller.

20. The apparatus of claim 19 wherein said first and second memory controllers are within a same CPU package.

21. An apparatus, comprising:

- 5 a DIMM card having a programmable component to implement a memory page data checking function.




Fig.1




Fig. 2




Fig. 3a




Fig. 3b

a
b
c
d
e
f
g
h

Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8




Fig. 9

A. CLASSIFICATION OF SUBJECT MATTER**G06F 13/16(2006.01)i, G06F 3/06(2006.01)i, G06F 12/02(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 13/16; G06F 12/06; G11C 7/10; G06F 12/00; G11C 7/22; G06F 12/08; G06F 13/12; G06F 3/06; G06F 12/02

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: memory controller, DIMM(dual in-line memory module), programmable component, checking function, and similar terms.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014-0040550 A1 (BILL NALE et al.) 06 February 2014 See paragraphs [0002], [0063], [0065], [0068], [0083]-[0084], [0107], [0135], [0138]-[0139], [0186], and [0214]; and figures 3-4, 7A, and 9.	1-7, 15
Y		8-14, 16-21
Y	WO 2014-143056 A1 (INTEL CORPORATION et al.) 18 September 2014 See paragraphs [0001], [0035]-[0037], and [0095]; and figure 1.	8-14, 16-21
Y	US 8,364,867 B2 (VIJAY KARAMCHETI et al.) 29 January 2013 See column 2, line 27 - column 3, line 49; and figure 1.	9-12, 18-20
A	US 2009-0031078 A1 (LIDIA WARNES et al.) 29 January 2009 See paragraphs [0022]-[0027] and figures 5A-5B.	1-21
A	US 2007-0220228 A1 (KUN-HUNG HUANG) 20 September 2007 See paragraphs [0022]-[0027] and figure 2.	1-21



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

19 August 2016 (19.08.2016)

Date of mailing of the international search report

19 August 2016 (19.08.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/031613

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0040550 A1	06/02/2014	CN 104025060 A	03/09/2014
		EP 2761472 A1	06/08/2014
		EP 2761472 A4	06/05/2015
		TW 201329994 A	16/07/2013
		TW 201604889 A	01/02/2016
		TW I512748 B	11/12/2015
		US 9342453 B2	17/05/2016
		WO 2013-048493 A1	04/04/2013
WO 2014-143056 A1	18/09/2014	CN 104981872 A	14/10/2015
		EP 2973571 A1	20/01/2016
		KR 10-2015-0120359 A	27/10/2015
		US 2015-0149735 A1	28/05/2015
		US 2016-0098195 A1	07/04/2016
		US 2016-0098366 A1	07/04/2016
		US 2016-0099044 A1	07/04/2016
		US 2016-0132269 A1	12/05/2016
		US 2016-0147678 A1	26/05/2016
		US 2016-0148653 A1	26/05/2016
US 8364867 B2	29/01/2013	TW 200825751 A	16/06/2008
		TW 200834585 A	16/08/2008
		TW 200839517 A	01/10/2008
		TW 200839767 A	01/10/2008
		TW 201511025 A	16/03/2015
		TW 201535402 A	16/09/2015
		TW I446171 B	21/07/2014
		TW I454923 B	01/10/2014
		TW I471861 B	01/02/2015
		TW I500031 B	11/09/2015
		US 2008-0082731 A1	03/04/2008
		US 2008-0082732 A1	03/04/2008
		US 2008-0082733 A1	03/04/2008
		US 2008-0082734 A1	03/04/2008
		US 2008-0082766 A1	03/04/2008
		US 2010-0274956 A1	28/10/2010
		US 2010-0274957 A1	28/10/2010
		US 2010-0274958 A1	28/10/2010
		US 2010-0274959 A1	28/10/2010
		US 7761623 B2	20/07/2010
		US 7761624 B2	20/07/2010
		US 7761625 B2	20/07/2010
		US 7761626 B2	20/07/2010
		US 8051253 B2	01/11/2011
		US 8370547 B2	05/02/2013
		US 8370548 B2	05/02/2013
		US 8380898 B2	19/02/2013
		WO 2008-039885 A2	03/04/2008
		WO 2008-039885 A3	14/08/2008

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/031613

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		WO 2008-039886 A2	03/04/2008
		WO 2008-039886 A3	14/08/2008
		WO 2008-040028 A2	03/04/2008
		WO 2008-040028 A3	23/10/2008
		WO 2008-051940 A2	02/05/2008
		WO 2008-051940 A3	14/08/2008
US 2009-0031078 A1	29/01/2009	US 8892942 B2	18/11/2014
US 2007-0220228 A1	20/09/2007	None	