



(51) International Patent Classification:

G06F 11/08 (2006.01) G06F 11/10 (2006.01)

(21) International Application Number:

PCT/US201 1/066683

(22) International Filing Date:

22 December 2011 (22.12.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, MS: RNB-4-150, Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **YIGZAW, Theodros** [US/US]; 21505 SW Chapman Rd., Sherwood, Oregon 97140 (US). **SANTHANAKRISHNAN, Geeyarpuram N.** [IN/US]; 9011 W Shorewood Dr, APT 1308, Mercer Island, Washington 98040 (US). **SRINIVASA, Ganapati N.** [US/US]; 12559 NW Hartford St., Portland, Oregon 97229 (US). **VARGAS, Jose A.** [US/US]; 3439 Foxmore Lane, Rescue, California 95672 (US). **SHAFI, Hisham** [IL/IL]; Hahagana St. 6/8; P.O. Box 189, 24314 Akko (IL). **MISHAELI, Michael** [IL/IL]; Hanotrim 19/A, 30900 Zichron Yaakov, HA (IL). **COHEN, Ehud** [IL/IL]; Dakar St. 74, 26430 Kiryat Motskin, HA (IL). **SPERBER, Zeev** [IL/IL]; 32nd Igal Alon St., 30900 Zichron Yackov, HA (IL). **RAIKIN, Shlomo** [IL/IL]; POB 56, 30900 Ofer, Z (IL).

KUMAR, Mohan J. [US/US]; 1860 SW Marko Lane, Aloha, Oregon 97007 (US). **MANDELBLAT, Julius Y.** [IL/IL]; Yaffe Nof 112/2, 34454 Haifa, HA (IL).


(74) Agents: **WEBSTER, Thomas C.** 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, 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, KM, 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, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

[Continued on nextpage]

(54) Title: APPARATUS AND METHOD FOR DETECTING AND RECOVERING FROM DATA FETCH ERRORS



(57) Abstract: An apparatus and method are described for detecting and correcting data fetch errors within a processor core. For example, one embodiment of an instruction processing apparatus for detecting and recovering from data fetch errors comprises: at least one processor core having a plurality of instruction processing stages including a data fetch stage and a retirement stage; and error processing logic in communication with the processing stages to perform the operations of: detecting an error associated with data in response to a data fetch operation performed by the data fetch stage; and responsively performing one or more operations to ensure that the error does not corrupt an architectural state of the processor core within the retirement stage.

Published:

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

**APPARATUS AND METHOD FOR DETECTING AND
RECOVERING FROM DATA FETCH ERRORS**

BACKGROUND

Field of the Invention

This invention relates generally to the field of computer processors. More particularly, the invention relates to an apparatus and method for detecting and correcting data fetch errors within a processor core.

Description of the Related Art

During normal operation, uncorrectable data errors encountered during the execution of a program are fatal to the whole system. With higher levels of integration in CPUs and ever smaller geometries, each occurrence of a fatal error detracts from the availability of the system. An additional implication is that with usage models such as virtual threading ("VT"), the "system" consists of a number of virtual machines ("VMs") which, for all intents and purposes, are running independent of each other. Therefore, when the system is brought down because of an error that occurred in one VM, it means that all VMs running on that system are brought down.

By way of example, in prior art Intel Architecture-32 (IA32) systems, uncorrected data errors encountered anywhere in the memory hierarchy during an instruction fetch or data fetch are immediately fatal. When such errors are encountered, the processor core logs a fatal error and signals the event immediately. The only action available to an OS or hypervisor at this point is to bring down the whole system.

Consequently, it would be beneficial to design a system in which uncorrectable data errors during instruction fetch and/or data fetch can be isolated to the offending VM or application so that only that portion of the system implicated in the error is brought down, leaving all the others operational.

BRIEF DESCRIPTION OF THE DRAWINGS

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

FIG. 1 illustrates one embodiment of an architecture for detecting and correcting instruction fetch errors.

FIG. 2 illustrates one embodiment of a method for detecting and correcting instruction fetch errors.

FIG. 3 illustrates one embodiment of an architecture for detecting and correcting data fetch errors.

FIG. 4 illustrates a transaction diagram showing one embodiment of a method for detecting and correcting data fetch errors.

5 **FIG. 5** illustrates a data fetch error log in accordance with one embodiment of the invention.

FIG. 6 illustrates one embodiment of a computer architecture on which embodiments of the invention may be implemented.

10

DETAILED DESCRIPTION

In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the embodiments of the invention described below. It will be apparent, however, to one skilled in the art that the embodiments of the invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form to avoid obscuring the underlying principles of the embodiments of the invention.

System and Method for Recovery From Errors During Instruction Fetch

In one embodiment of the invention, an uncorrectable data error during instruction fetch is isolated to the offending virtual machine (VM) or application so that only that portion of the system that is implicated in the error is brought down, leaving all other portions operational. Corrupt data containment, referred to herein as "poisoning," is used to indicate the occurrence of an uncorrected error synchronous to the data. When an uncorrected data error is detected anywhere in the memory hierarchy, a "poison" indication is attached to the data, thereby deferring the handling of the error to the hardware that requested the data. In one embodiment, the poison indication is a single bit which is set to indicate poisoned data. As a result of these embodiments of the invention, the concept of corrupt data containment is extended into the processing core including the instruction and data fetch hardware.

30 **Figure 1** illustrates a processor architecture 150 according to one embodiment of the invention which includes one or more cores (identified as "Core 0," "Core 1," and Core N in **Figure 1**). The details of a single core (Core 0) are illustrated in **Figure 1** for simplicity. It will be understood, however, that each core shown in **Figure 1** may have the same set of logic as Core 0.

As illustrated, the core may have a dedicated Level 1 (L1) cache 112 and Level 2 (L2) cache 111 for caching instructions and data according to a specified cache management policy.

The LI cache includes a separate instruction cache 120 for storing instructions and a separate data cache 121 for storing data. The instructions and data stored within the various processor caches are managed at the granularity of cache lines which may be a fixed size (e.g., 64 or 128 Bytes in length). Each core of this exemplary embodiment has an instruction fetch unit 110 for
5 fetching instructions from main memory 100 and/or a Level 3 (L3) cache 116; a decode unit 120 for decoding the instructions (e.g., decoding program instructions into micro-operations or "i/ops"); an execution unit 130 for executing the instructions, and a retirement unit 150 for retiring the instructions.

The instruction fetch unit 110 includes various well known components including a next
10 instruction pointer 103 for storing the address of the next instruction to be fetched from memory 100 (or one of the caches); an instruction translation look-aside buffer (ITLB) for storing a map of recently used virtual-to-physical instruction addresses to improve the speed of address translation; a branch prediction unit 102 for speculatively predicting instruction branch
15 addresses; and branch target buffers (BTBs) 101 for storing branch addresses and target addresses. Once fetched, instructions are then streamed to the remaining stages of the instruction pipeline including the decode unit 130, the execution unit 140, and the retirement unit 150. The structure and function of each of these units is well understood by those of ordinary skill in the art and will not be described here in detail to avoid obscuring the pertinent aspects of the different embodiments of the invention.

As illustrated, error processing logic 105 is included within each core for implementing
20 the error detection and recovery techniques described herein. In one embodiment, when the error processing logic 105 detects an uncorrected error in a cache line at any point in the memory hierarchy, instead of signaling and logging a fatal error (as in prior implementations), it attaches a "poison" indication to the cache line with the data, thereby deferring the handling of the error
25 to the hardware that requested the data. In one embodiment, the error processing logic 105 detects an error by detecting that the cache line (or other block of data) has failed an error correction code (ECC) check such as a checksum. Each of the stages in the memory hierarchy including the main memory 100 and the various caches 111, 112, 116 may be equipped with ECC logic for determining the validity of data stored therein. In one embodiment of the
30 invention, the error processing logic 105 identifies an error and attaches a poison indication when a cache line (or other block of data) fails its ECC check at one of these levels in the memory hierarchy.

In one embodiment, the error processing logic 105 will only react to poisoned
instructions if the instructions are not due to a speculative access. If the access is on a speculative
35 path (for example, due to a mis-predicted branch) the error processing logic 105 will neither log

nor signal an error (i.e., because the error was not relevant to the current execution path). If, however, the poisoned instruction(s) within the cache line were received for a non-speculative access, the error processing logic 105 reacts by withholding the instruction(s) from consumption. In such a case, the error processing logic 105 will then clear the pipeline and the poisoned
 5 instruction(s) from the locality. It will then log and signal a recoverable instruction fetch error within an error log 106. In some cases, it is possible for the poisoned part of the cache line not to be consumed immediately. In these cases, the error processing logic 105 will move the whole cache line into the instruction cache ("Icache") 120 in a poisoned state (i.e., along with the attached poison indication). The next time the cache line is accessed, it will be discovered in the
 10 poisoned state in the Icache based on the poisoned indication. When this happens, one embodiment of the error processing logic 105 will invalidate the cache line in the Icache 120, treat the access as a cache miss and request the data from the memory hierarchy (e.g., from the main memory 100).

In one embodiment, when the error processing logic 105 logs a recoverable instruction
 15 fetch error in the error log 106, it does so with a specific signature. The error log 106 also includes the system address of the request that generated the error. The system address is essential for error handling software to isolate the error to the process or VM (in a VT environment) that is affected by this error. In one embodiment, system software can then read the error log 106 and carry out recovery by terminating the affected processes or VM and
 20 returning to the rest, without bringing down the whole system.

Figure 2 illustrates one embodiment of a process for detecting and processing errors during instruction fetch. The process may be implemented by the error processing logic 105 shown in **Figure 1** but is not necessarily limited to any particular processor architecture.

At 301, a determination is made as to whether the error is associated with a speculative
 25 access (e.g., a speculated branch instruction). If so, then at 210 no error is logged or signaled and program execution continues (i.e., because the error was not relevant to the current execution path). If the error is associated with a non-speculative access then steps are taken to ensure that the poisoned instruction is not consumed. In one embodiment, each cache line is comprised of two 32-Byte "chunks." Thus, at 202, a determination is made as to whether the poisoned
 30 instruction is associated with a critical chunk (i.e., one which will be used). If so, then the data is withheld at 204 and, at 205, instruction execution is stalled. A recoverable instruction fetch error is logged and a machine check exception (MCE) is signaled at 206 (i.e., to notify the MCA of a machine check event). As is well known by those of ordinary skill in the art, MCA refers to a mechanism in which the CPU reports hardware errors to the operating system.

If, however, at 202, it is determined that the poison is not on a critical chunk, then at 203 a determination is made as to whether the non-critical chunk will be used immediately. If so, then the process returns to operation 204. If not, then at 207, the cache line is stored in the instruction cache in a poisoned state. The next time the cache line is accessed, at 208, the cache line is discovered in the poisoned state and the cache line is invalidated. Once invalidated, the access to the cache line is treated as a cache miss and the data is requested from further up the memory hierarchy.

System and Method for Recovery From Errors During Data Fetch

Figure 3 illustrates one embodiment of the invention in which the error processing logic 105 uses data poisoning techniques during data fetch operations. While not illustrated together in a single figure, this embodiment may be employed within the same processor architecture as the embodiments used to recover from errors during instruction fetch described above with respect to **Figures 1-2**.

As mentioned, data poisoning is a technique employed in one embodiment of the invention so that an error is isolated and not permitted to corrupt the architectural state of the processor (which would affect the entire computer system). In prior systems, an uncorrected data error detected in one of a processor's cores would cause the hardware to broadcast an error immediately. However, the corrupted data would still be sent to its destination where it would be consumed, causing the computer system's architectural state to be corrupted (generally referred to as "Processor Context Corrupt" in the Machine Check Architecture (MCA) error log).

Corrupt data containment as defined herein starts with an uncorrectable data error discovered anywhere in the system (e.g., within the cache hierarchy or memory or any other temporary storage of data or device). As mentioned above, the error may be detected at each stage of the memory hierarchy using various known error correction code (ECC) techniques. In one embodiment, when a data error is detected, the error processing logic 105 attaches a poison indication to the data, which is then communicated along with the data down the processor pipeline wherever and whenever the data is transmitted.

Figure 3 illustrates one embodiment of a processor architecture which includes a plurality of cores, identified as core 0, core 1, . . . core N. The details of a single core (Core 0) with respect to one embodiment of the invention are illustrated in **Figure 3** for simplicity. It will be understood, however, that each core shown in **Figure 3** may have the same set of logic as Core 0 with respect to the embodiments of the invention. Additionally, it will be understood that non-homogeneous cores may be employed while still complying with the underlying principles of the invention.

A data fetch unit 301 is illustrated for fetching data from a cache/memory subsystem comprising a level 1 (L1) cache 112, Level 2 (L2) cache 111, Level 3 (L3) cache 116 and main memory 100. The L1 cache includes a dedicated instruction cache (ICache) 120 for storing instructions and a dedicated data cache 121 for storing data. As mentioned, the instructions and data stored within the various processor caches may be managed at the granularity of cache lines of a fixed size (e.g., 64 or 128 Bytes in length). A decode unit 130 is illustrated for decoding instructions (e.g., decoding program instructions into micro-operations or "i/ops"); an execution unit 140 for executing the instructions; and a retirement unit 150 for retiring the instructions (e.g., updating sets of the processor's architectural state registers following the execution of each instruction).

A transaction diagram illustrating the operation of one embodiment of the invention is illustrated in **Figure 4A**. In this embodiment, an instruction has been decoded by the decode unit 130, resulting in a plurality of micro-operations, one of which is a load micro-operation serviced by the data fetch unit 452. In response, the data fetch unit 452 attempts to fetch the requested data from the memory hierarchy 451. In the example shown in **Figure 4A**, the error processing component 105 of the memory hierarchy 451 returns data with a poison indication 402 to the data fetch unit 301. For example, the error processing logic 105 may detect an ECC failure associated with the data when retrieving data from the various caches 111-112, 116 and/or the main memory 100.

In response, the error processing component 105 of the data fetch unit 301 ensures that the load operation does not retire (get committed to architectural state) by logging the error within the error log 106 and transmitting a poison notification 403 to the retirement unit 150. In one embodiment, the "notification" is the update to the error log 106 which is accessible by the retirement unit 150. For example, the retirement unit 150 may subsequently read the error log 106 to determine if an error is associated with the load operation. If an error is detected by the retirement unit 150 then, in one embodiment, it will not commit the load operation to update the architectural state of the core. Regardless of how the notification is implemented, as indicated in **Figure 4A**, the notification to the retirement unit 150 is the first step to ensure that the error does not result in the modification of the architectural state of the core (i.e., the error's effect has been limited to the process, application or VM that tried to consume the bad data and thus encountered the error).

In one embodiment, after the load operation is denied retirement, the next step is to ensure that no architectural state is modified by the instruction that spawned the micro-operation which attempted to load corrupted data (i.e., which generated the load data request micro-

operation 401). In one embodiment, this is accomplished by preserving the architectural state which existed before attempting to execute the instruction that encountered corrupted data.

In one embodiment, to ensure that the error is linked to a specific instruction and a specific system address, indications of both the specific instruction and system address are stored by the error processing logic 105 within the error log 106. Linking the error to a specific instruction and system address is necessary for isolation because the system address is what the error handling software will use to decide the scope of the error (e.g., which processes/VM are affected by the error). Thus, in one embodiment, when the error processing logic 105 logs the error it includes the system address for the load that experienced the error.

Finally, the hardware must communicate the nature of the error to the error handling software, which is responsible for taking the recovery action. As illustrated in **Figure 4A**, in one embodiment, this is accomplished via a Machine Check Event 404 transmitted to the Machine Check Architecture 450 of the system, which is accessible by software. The assignee of the present application has developed a Machine Check Architecture (MCA) for its family of processor architectures including, by way of example and not limitation, the Pentium™, Xeon™, Itanium™, Itanium™, and Intel Core™ family of processors. The MCA provides a mechanism for detecting and reporting hardware (machine) errors, such as system bus errors, ECC errors, parity errors, cache errors, and translation lookaside buffer errors. It includes a set of model-specific registers (MSRs) that are used to set up machine checking and additional banks of MSRs used for recording errors that are detected. Thus, in one embodiment, the MSRs of the MCA store the system address where the error occurred which is then read by the software. The software may then determine if the error is recoverable and take an appropriate action. For example, the software may then use the system address to determine which process, application or virtual machine encountered the error, and can then take action by either terminating the offending process, application or virtual machine, and allowing all others to continue. In one embodiment, the software may also re-initiate the request for the data so that a "clean" copy (i.e., one without errors) of the data is fetched from memory or disk.

[0001] Returning to the transaction diagram in **Figure 4A**, following the registration of the machine check event 404 to the MCA 450, the MCA transmits a command to the data fetch unit 301 to amend the error log based on the results of the machine check operations. In addition, a machine check error signal is propagated to the programmable interrupt controller (PIC) 454 at 406, which broadcasts a machine check error indication at 407 to the logic outside of the core in which the error occurred - commonly referred to as the uncore 455.

The uncore 455 may include, for example, the communication pathways between each of the cores, to the main memory 100 and input/output (I/O) interfaces (not shown). In one

embodiment, at 408, the machine check error is broadcast to the other cores in the processor 150 and/or to other components located off of the processor package (e.g., memory and/or I/O devices) so that they may take appropriate action to handle the error. In one embodiment, these components may determine that there is a recoverable error by reading the error log 106. In the exemplary error log shown in **Figure 4B**, a recoverable error is indicated if the processor context is not corrupt as indicated by MCI_STATUS.PCC=0, the system address is available as indicated by MCI_STATUS.ADDRV=1 and MCI_MISC[8:6]='010. In one embodiment, the MCACOD field and MCG_STATUS fields are different to distinguish between for the instruction fetch vs. the data fetch error logs. The following fields have the following meanings in one embodiment of the error log:

In MCLSTATUS:

MCACOD = architecturally specified error code that describes the type of error encountered.

UC = Uncorrected. When '1 the error was of an uncorrected type. When '0 the error was of a corrected type.

PCC = Processor Context Corrupt. When set it means the processor context was corrupt. When clear it means it was not.

S = signaling. When set the error logged was signaled via a machine check exception. When clear it was signaled via a CMCI.

AR = Action Required. When set, SW must take immediate action. When clear SW may delay the action until a later time.

ADDRV - Address Valid. When set it indicates that a valid address was logged in MCI_ADDR MSR. When clear it indicates no valid address exists.

MISCV= MCLMISC Valid. When set it indicates that there is valid content in the MCLMISC MSR. When clear it indicates that there is no valid content in MCI_MISC.

In MCG_STATUS:

RIPV = Restart IP Valid. When set it indicates that program execution can be restarted reliably at the instruction pointed to by the instruction pointer pushed on the stack by the machine check exception handler call. When clear, the program cannot be reliably restarted at the pushed instruction pointer.

EIPV = Error IP Valid. Indicates (when set) that the instruction pointed to by the instruction pointer pushed onto the stack when the machine check exception is generated is directly associated with the error. When this flag is cleared, the instruction pointed to may not be associated with the error.

MCI_ADDR - contains the address associated with the error.

MCi_MISC[5:0] - contains the LSB of the address captured in MCi_ADDR.

MCi_MISC[8:6] - contains the address type (one of segment offset, linear address, system address, memory address or generic) captured in MCi_ADDR.

Of course, the underlying principles of the invention are not limited to any particular organization of data stored in the error log.

Referring now to **Figure 5**, shown is a block diagram of another computer system 500 in accordance with one embodiment of the present invention. The system 500 may include one or more processing elements 510, 515, which are coupled to graphics memory controller hub (GMCH) 520. The optional nature of additional processing elements 515 is denoted in **Figure 5** with broken lines.

Each processing element may be a single core or may, alternatively, include multiple cores. The processing elements may, optionally, include other on-die elements besides processing cores, such as integrated memory controller and/or integrated I/O control logic. Also, for at least one embodiment, the core(s) of the processing elements may be multithreaded in that they may include more than one hardware thread context per core.

Figure 5 illustrates that the GMCH 520 may be coupled to a memory 540 that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one embodiment, be associated with a non-volatile cache.

The GMCH 520 may be a chipset, or a portion of a chipset. The GMCH 520 may communicate with the processor(s) 510, 515 and control interaction between the processor(s) 510, 515 and memory 540. The GMCH 520 may also act as an accelerated bus interface between the processor(s) 510, 515 and other elements of the system 500. For at least one embodiment, the GMCH 520 communicates with the processor(s) 510, 515 via a multi-drop bus, such as a frontside bus (FSB) 595.

Furthermore, GMCH 520 is coupled to a display 540 (such as a flat panel display). GMCH 520 may include an integrated graphics accelerator. GMCH 520 is further coupled to an input/output (I/O) controller hub (ICH) 550, which may be used to couple various peripheral devices to system 500. Shown for example in the embodiment of **Figure 5** is an external graphics device 560, which may be a discrete graphics device coupled to ICH 550, along with another peripheral device 570.

Alternatively, additional or different processing elements may also be present in the system 500. For example, additional processing element(s) 515 may include additional processors(s) that are the same as processor 510, additional processor(s) that are heterogeneous or asymmetric to processor 510, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processing element. There

can be a variety of differences between the physical resources 510, 515 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst the processing elements 510, 515. For at least one embodiment, the various processing elements 510, 515 may reside in the same die package.

Figure 6 is a block diagram illustrating another exemplary data processing system which may be used in some embodiments of the invention. This data processing system may be a handheld computer, a personal digital assistant (PDA), a mobile telephone, a portable gaming system, a portable media player, a tablet or a handheld computing device which may include a mobile telephone, a media player, and/or a gaming system. As another example, the data processing system 500 may be a network computer or an embedded processing device within another device.

According to one embodiment of the invention, the exemplary architecture of the data processing system may be used for the mobile devices described above. The data processing system 600 includes the processing system 620, which may include one or more microprocessors and/or a system on an integrated circuit. The processing system 620 is coupled with a memory 610, a power supply 625 (which includes one or more batteries) an audio input/output 640, a display controller and display device 660, optional input/output 650, input device(s) 670, and wireless transceiver(s) 630. It will be appreciated that additional components, not shown in **Figure 6**, may also be a part of the data processing system 600 in certain embodiments of the invention, and in certain embodiments of the invention fewer components than shown in **Figure 6** may be used. In addition, it will be appreciated that one or more buses, not shown in **Figure 6**, may be used to interconnect the various components as is well known in the art.

The memory 610 may store data and/or programs for execution by the data processing system 600. The audio input/output 640 may include a microphone and/or a speaker to, for example, play music and/or provide telephony functionality through the speaker and microphone. The display controller and display device 660 may include a graphical user interface (GUI). The wireless (e.g., RF) transceivers 630 (e.g., a WiFi transceiver, an infrared transceiver, a Bluetooth transceiver, a wireless cellular telephony transceiver, etc.) may be used to communicate with other data processing systems. The one or more input devices 670 allow a user to provide input to the system. These input devices may be a keypad, keyboard, touch panel, multi touch panel, etc. The optional other input/output 650 may be a connector for a dock.

Other embodiments of the invention may be implemented on cellular phones and pagers (e.g., in which the software is embedded in a microchip), handheld computing devices (e.g., personal digital assistants, smartphones), and/or touch-tone telephones. It should be noted,

however, that the underlying principles of the invention are not limited to any particular type of communication device or communication medium.

Embodiments of the invention may include various steps, which have been described above. The steps may be embodied in machine-executable instructions which may be used to
5 cause a general-purpose or special-purpose processor to perform the steps. Alternatively, these steps may be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

Elements of the present invention may also be provided as a computer program product
10 which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic device) to perform a process. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, propagation media or other type of media/machine-readable medium suitable for storing
15 electronic instructions. For example, the present invention may be downloaded as a computer program product, wherein the program 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).

Throughout this detailed description, for the purposes of explanation, numerous specific
20 details were set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the invention may be practiced without some of these specific details. In certain instances, well known structures and functions were not described in elaborate detail in order to avoid obscuring the subject matter of the present
25 invention. Accordingly, the scope and spirit of the invention should be judged in terms of the claims which follow.

Referring now to **Figure 5**, shown is a block diagram of another computer system 500 in
accordance with one embodiment of the present invention. The system 500 may include one or
more processing elements 510, 515, which are coupled to graphics memory controller hub
30 (GMCH) 520. The optional nature of additional processing elements 515 is denoted in **Figure 5** with broken lines.

Each processing element may be a single core or may, alternatively, include multiple cores. The processing elements may, optionally, include other on-die elements besides processing cores, such as integrated memory controller and/or integrated I/O control logic. Also,

for at least one embodiment, the core(s) of the processing elements may be multithreaded in that they may include more than one hardware thread context per core.

Figure 5 illustrates that the GMCH 520 may be coupled to a memory 540 that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one
5 embodiment, be associated with a non-volatile cache.

The GMCH 520 may be a chipset, or a portion of a chipset. The GMCH 520 may communicate with the processor(s) 510, 515 and control interaction between the processor(s) 510, 515 and memory 540. The GMCH 520 may also act as an accelerated bus interface between the processor(s) 510, 515 and other elements of the system 500. For at least one
10 embodiment, the GMCH 520 communicates with the processor(s) 510, 515 via a multi-drop bus, such as a frontside bus (FSB) 595.

Furthermore, GMCH 520 is coupled to a display 540 (such as a flat panel display). GMCH 520 may include an integrated graphics accelerator. GMCH 520 is further coupled to an input/output (I/O) controller hub (ICH) 550, which may be used to couple various peripheral
15 devices to system 500. Shown for example in the embodiment of **Figure 5** is an external graphics device 560, which may be a discrete graphics device coupled to ICH 550, along with another peripheral device 570.

Alternatively, additional or different processing elements may also be present in the system 500. For example, additional processing element(s) 515 may include additional
20 processors(s) that are the same as processor 510, additional processor(s) that are heterogeneous or asymmetric to processor 510, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processing element. There can be a variety of differences between the physical resources 510, 515 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption
25 characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst the processing elements 510, 515. For at least one embodiment, the various processing elements 510, 515 may reside in the same die package.

Figure 6 is a block diagram illustrating another exemplary data processing system which may be used in some embodiments of the invention. This data processing system may be a
30 handheld computer, a personal digital assistant (PDA), a mobile telephone, a portable gaming system, a portable media player, a tablet or a handheld computing device which may include a mobile telephone, a media player, and/or a gaming system. As another example, the data processing system 500 may be a network computer or an embedded processing device within another device.

According to one embodiment of the invention, the exemplary architecture of the data processing system may be used for the mobile devices described above. The data processing system 600 includes the processing system 620, which may include one or more microprocessors and/or a system on an integrated circuit. The processing system 620 is coupled with a memory 610, a power supply 625 (which includes one or more batteries) an audio input/output 640, a display controller and display device 660, optional input/output 650, input device(s) 670, and wireless transceiver(s) 630. It will be appreciated that additional components, not shown in **Figure 6**, may also be a part of the data processing system 600 in certain embodiments of the invention, and in certain embodiments of the invention fewer components than shown in **Figure 6** may be used. In addition, it will be appreciated that one or more buses, not shown in **Figure 6**, may be used to interconnect the various components as is well known in the art.

The memory 610 may store data and/or programs for execution by the data processing system 600. The audio input/output 640 may include a microphone and/or a speaker to, for example, play music and/or provide telephony functionality through the speaker and microphone. The display controller and display device 660 may include a graphical user interface (GUI). The wireless (e.g., RF) transceivers 630 (e.g., a WiFi transceiver, an infrared transceiver, a Bluetooth transceiver, a wireless cellular telephony transceiver, etc.) may be used to communicate with other data processing systems. The one or more input devices 670 allow a user to provide input to the system. These input devices may be a keypad, keyboard, touch panel, multi touch panel, etc. The optional other input/output 650 may be a connector for a dock.

Other embodiments of the invention may be implemented on cellular phones and pagers (e.g., in which the software is embedded in a microchip), handheld computing devices (e.g., personal digital assistants, smartphones), and/or touch-tone telephones. It should be noted, however, that the underlying principles of the invention are not limited to any particular type of communication device or communication medium.

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

Elements of the present invention may also be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic device) to perform a process. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs,

and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet 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 product, wherein the program 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).

Throughout this detailed description, for the purposes of explanation, numerous specific details were set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the invention may be practiced without some of these specific details. In certain instances, well known structures and functions were not described in elaborate detail in order to avoid obscuring the subject matter of the present invention. Accordingly, the scope and spirit of the invention should be judged in terms of the claims which follow.

CLAIMS

What is claimed is: 1. A method for detecting and recovering from data fetch errors within a processor core comprising:

detecting an error associated with data in response to a data fetch operation; and

5 responsively performing one or more operations to ensure that the error does not corrupt an architectural state of the processor core.

2. The method as in claim 1 wherein responsively performing one or more operations comprises attaching a poison indication to the data.

3. The method as in claim 2 further comprising:

10 transmitting the poison indication with the data down the processor pipeline to provide the poison indication to stages and units within the processor pipeline.

4. The method as in claim 3 wherein the poison indication is transmitted to a retirement unit to ensure that the retirement unit does not retire the data fetch operation which resulted in the error.

15 5. The method as in claim 1 wherein the error is detected in a chunk of data equal to a specified fraction of a cache line.

6. The method as in claim 5 wherein the chunk of data comprises one half of a cache line.

20 7. The method as in claim 6 wherein the cache line is 64 Bytes and the chunk of data comprises 32 Bytes.

8. The method as in claim 1 wherein detecting an error comprises executing an error correction code (ECC) operation on the data.

9. A data processing apparatus for detecting and recovering from data fetch errors comprising:

25 at least one processor core having a plurality of instruction processing stages including a data fetch stage and a retirement stage; and

error processing logic in communication with the processing stages to perform the operations of:

30 detecting an error associated with data in response to a data fetch operation performed by the data fetch stage; and

responsively performing one or more operations to ensure that the error does not corrupt an architectural state of the processor core within the retirement stage.

10. The apparatus as in claim 9 wherein responsively performing one or more operations comprises attaching a poison indication to the data.

11. The apparatus as in claim 10 wherein the error processing logic performs the additional operation of:

transmitting the poison indication with the data down the processor pipeline to provide the poison indication to stages and units within the processor pipeline.

12. The apparatus as in claim 11 wherein the poison indication is transmitted to a retirement unit to ensure that the retirement unit does not retire the data fetch operation which resulted in the error.

13. The apparatus as in claim 9 wherein the error is detected in a chunk of data equal to a specified fraction of a cache line.

14. The apparatus as in claim 13 wherein the chunk of data comprises one half of a cache line.

15. The apparatus as in claim 14 wherein the cache line is 64 Bytes and the chunk of data comprises 32 Bytes.

16. The apparatus as in claim 9 wherein detecting an error comprises executing an error correction code (ECC) operation on the data.

17. A computer system comprising:

a central processing unit including at least one processor core having a plurality of instruction processing stages including a data fetch stage and a retirement stage;

error processing logic in communication with the processing stages to perform the operations of:

detecting an error associated with data in response to a data fetch operation performed by the data fetch stage; and

responsively performing one or more operations to ensure that the error does not corrupt an architectural state of the processor core within the retirement stage; and

at least one input/output interface communicatively coupling the central processing unit to one or more input/output devices.

18. The computer system as in claim 17 wherein responsively performing one or more operations comprises attaching a poison indication to the data.

19. The computer system as in claim 18 wherein the error processing logic performs the additional operation of:

transmitting the poison indication with the data down the processor pipeline to provide the poison indication to stages and units within the processor pipeline.


20. The computer system as in claim 19 wherein the poison indication is transmitted to a retirement unit to ensure that the retirement unit does not retire the data fetch operation which resulted in the error.


5 21. The computer system as in claim 17 wherein the error is detected in a chunk of data equal to a specified fraction of a cache line.


22. The computer system as in claim 20 wherein the chunk of data comprises one half of a cache line.

23. The computer system as in claim 13 wherein the cache line is 64 Bytes and the chunk of data comprises 32 Bytes.

10 24. The computer system as in claim 17 wherein detecting an error comprises executing an error correction code (ECC) operation on the data.



**Fig. 2**






Fig. 4A

Error Log 106



Errors	MCI_STATUS							MCG_STATUS				MCI_Addr	MCI_MISC
								Errored Thread		Other Threads			
								RIP V	EIP V	RIP V	EIP V		
Data Fetch Recoverable	MCACOD	UC	PCC	S	AR	ADDRV	MISCV	0	1	1	0	PA[45:6]	[5:0]='000110 [8:6]='010

Fig. 4B





Fig. 5

**Fig. 6**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2011/066683**A. CLASSIFICATION OF SUBJECT MATTER****G06F 11/08(2006.01)i, G06F 11/10(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 11/08; H04L 1/22; H03M 13/00; G06F 12/00; G11C 029/00

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: "Fetch, error, recover, processor"

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	US 2009-0204766 A1 (JACOBI CHRISTIAN et al.) 13 August 2009 See abstract; paragraphs [0017]-[0038] and figures 1,2.	1, 5-9, 13-17, 21-24 2-4, 10-12, 18-20
Y A	US 2005-0149781 A1 (LEMPEL ODED et al.) 07 July 2005 See abstract; paragraphs [0010]-[0028] and figures 1,2.	2-4, 10-12, 18-20 1, 5-9, 13-17, 21-24
A	US 7353445 B1 (JAMA I. BARREH et al.) 01 April 2008 See abstract; column 5 line 19 - column 18 line 42, and figures 2-5.	1-24
A	US 2003-0159086 A1 (RICHARD LOUIS ARNDT) 21 August 2003 See abstract; paragraphs [0062]-[0067] and figures 6,7.	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 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

26 SEPTEMBER 2012 (26.09.2012)

Date of mailing of the international search report

27 SEPTEMBER 2012 (27.09.2012)

Name and mailing address of the ISA/KR

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

Facsimile No. 82-42-472-7140

Authorized officer

Lee Byoung Soo

Telephone No. 82-42-481-5697



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2011/066683

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009-0204766 A1	13.08.2009	US 7987384 B2	26.07.2011
US 2005-0149781 A1	07.07.2005	US 7290179 B2	30.10.2007
US 7353445 B1	01.04.2008	None	
US 2003-0159086 A1	21.08.2003	US 2003-0139896 A1	24.07.2003
		US 6640202 B1	28.10.2003
		US 6658591 B1	02.12.2003
		US 6836855 B2	28.12.2004
		US 6957164 B2	18.10.2005