




(86) Date de dépôt PCT/PCT Filing Date: 2013/03/01
(87) Date publication PCT/PCT Publication Date: 2013/09/19
(45) Date de délivrance/Issue Date: 2021/03/02
(85) Entrée phase nationale/National Entry: 2014/09/09
(86) N° demande PCT/PCT Application No.: JP 2013/001264
(87) N° publication PCT/PCT Publication No.: 2013/136704
(30) Priorité/Priority: 2012/03/16 (US13/422,546)

(51) Cl.Int./Int.Cl. *G06F 11/34* (2006.01)
(72) Inventeurs/Inventors:
FARRELL, MARK S., US;
GAINEY, CHARLES W., JR., US;
MITRAN, MARCEL, CA;
SHUM, CHUNG-LUNG KEVIN, US;
SLEGEL, TIMOTHY, US;
SMITH, BRIAN LEONARD, US;
STOODLEY, KEVIN A., CA
(73) Propriétaire/Owner:
INTERNATIONAL BUSINESS MACHINES
CORPORATION, US
(74) Agent: WANG, PETER

(54) Titre : COMMANDE D'OPERATION D'UN MOYEN D'INSTRUMENTATION DU TEMPS D'EXECUTION DEPUIS UN ETAT MOINS PRIVILEGIE
(54) Title: CONTROLLING OPERATION OF A RUN-TIME INSTRUMENTATION FACILITY FROM A LESSER-PRIVILEGED STATE



(57) **Abrégé/Abstract:**

Embodiments of the invention relate to enabling and disabling execution of a run-time instrumentation facility. An instruction for execution by the processor in a lesser privileged state is fetched by the processor. It is determined, by the processor, that the run-

(57) Abrégé(suite)/Abstract(continued):

time instrumentation facility permits execution of the instruction in the lesser-privileged state and that controls associated with the run-time instrumentation facility are valid. The run-time instrumentation facility is disabled based on the instruction being a run-time instrumentation facility off (RIOFF) instruction. The disabling includes updating a bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor. The run-time instrumentation facility is enabled based on the instruction being a run-time instrumentation facility on (RION) instruction. The enabling includes updating the bit in the PSW to indicate that run-time instrumentation data should be captured by the processor.

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

(19) World Intellectual Property
Organization
International Bureau



(10) International Publication Number
WO 2013/136704 A1

(43) International Publication Date
19 September 2013 (19.09.2013)

(51) International Patent Classification:
G06F 11/34 (2006.01)

(21) International Application Number:
PCT/JP2013/001264

(22) International Filing Date:
1 March 2013 (01.03.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/422,546 16 March 2012 (16.03.2012) US

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York, 10504 (US).

(71) Applicant (for MG only): **IBM JAPAN, LTD.** [JP/JP]; 19-21 Nihonbashi Hakozaiki-cho, Chuo-ku, Tokyo, 1038510 (JP).

(72) Inventors: **FARRELL, Mark S.**; IBM Corporation, 2455 South Road, Poughkeepsie, New York, 12601 (US). **GAINEY JR., Charles W.**; IBM Corporation, Intellectual Property Law, 2455 South Road, Poughkeepsie, New York, 12601 (US). **MITRAN, Marcel**; IBM Canada Ltd, 8200 Warden Avenue, Toronto Lab, Markham, Ontario, L6G1C7 (CA). **SHUM, Chung-Lung Kevin**; IBM Corporation, 2455 South Road, Poughkeepsie, New York, 12601 (US). **SLEGEL, Timothy J.**; IBM Corporation, 2455 South Road, Poughkeepsie, New York, 12601 (US). **SMITH, Brian Leonard**; IBM Corporation, Intellectual Property Law, 2455 South Road, Poughkeepsie, New York, 12601 (US). **STOODLEY, Kevin A.**; IBM Canada Ltd, 8200 Warden Avenue, Toronto Lab, Markham, Ontario, L6G1C7 (CA).

(74) Agents: **UENO, Takeshi** et al.; c/o Toyosu site, IBM Japan, Ltd., NBF TOYOSU CANAL FRONT Bldg., 6-52, Toyosu 5-chome, Koto-ku, Tokyo, 1358511 (JP).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

[Continued on next page]

(54) Title: CONTROLLING OPERATION OF A RUN-TIME INSTRUMENTATION FACILITY FROM A LESSER-PRIVILEGED STATE




FIG. 12

(57) Abstract: Embodiments of the invention relate to enabling and disabling execution of a run-time instrumentation facility. An instruction for execution by the processor in a lesser privileged state is fetched by the processor. It is determined, by the processor, that the run-time instrumentation facility permits execution of the instruction in the lesser-privileged state and that controls associated with the run-time instrumentation facility are valid. The run-time instrumentation facility is disabled based on the instruction being a run-time instrumentation facility off (RIOFF) instruction. The disabling includes updating a bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor. The run-time instrumentation facility is enabled based on the instruction being a run-time instrumentation facility on (RION) instruction. The enabling includes updating the bit in the PSW to indicate that run-time instrumentation data should be captured by the processor.

WO 2013/136704 A1

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, 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, PA, 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.

GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

Description

Title of Invention: CONTROLLING OPERATION OF A RUN-TIME INSTRUMENTATION FACILITY FROM A LESSER-PRIVILEGED STATE

Technical Field

- [0001] The present invention relates generally to processing within a computing environment, and more specifically, to controlling operation of a run-time instrumentation facility from a lesser-privileged state.

Background Art

- [0002] Computer processors execute transactions using increasingly complex branch prediction and instruction caching logic. These processes have been introduced to increase instruction throughput, and therefore processing performance. The introduction of logic for improving performance makes it difficult to predict with certainty how a particular software application will execute on the computer processor. During the software development process there is often a balance between functionality and performance. Software is executed at one or more levels of abstraction from the underlying hardware that is executing the software. When hardware is virtualized, an additional layer of abstraction is introduced. With the introduction of performance enhancing logic, and the various layers of abstraction it is difficult to have a thorough understanding of what is actually occurring at the hardware level when a program is executing. Without this information, software developers use more abstract methods, such as execution duration, memory usage, number of threads, etc., for optimizing the software application.

Summary of Invention

Technical Problem

- [0003] When hardware specific information is available, it is typically provided to a developer after the fact and it is provided in aggregate, at a high level, and/or interspersed with the activity of other programs, and the operating system, making it difficult to identify issues that may be impacting the efficiency and accuracy of the software application.

Solution to Problem

- [0004] Embodiments include a computer program product, method and system for enabling and disabling execution of a run-time instrumentation facility on a processor. An instruction for execution by the processor in a lesser privileged state is fetched by the processor. It is determined, by the processor, that the run-time instrumentation facility

permits execution of the instruction in the lesser-privileged state and that controls associated with the run-time instrumentation facility are valid. The run-time instrumentation facility is disabled based on the instruction being a run-time instrumentation facility off (RIOFF) instruction. The disabling includes updating a bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor. The run-time instrumentation facility is enabled based on the instruction being a run-time instrumentation facility on (RION) instruction. The enabling includes updating the bit in the PSW to indicate that run-time instrumentation data should be captured by the processor.

- [0005] Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention. For a better understanding of the invention with advantages and features, refer to the description and to the drawings.
- [0006] The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The forgoing and other features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0006A] In accordance with a first major embodiment, there is provided a computer program product for enabling and disabling execution of a run-time instrumentation facility on a processor, the computer program product comprising:

a non-transitory storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising:

fetching an instruction of a currently executing thread in a multi-threaded environment for execution by the processor in a first state, the instruction one of a run-time instrumentation facility off (RIOFF) instruction and a run-time instrumentation facility on (RION) instruction, the fetching by the processor;

based on determining, by the processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, executing the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor; and

enabling the run-time instrumentation facility based on the instruction being the RION instruction, the enabling including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

[0006B] In accordance with a second major embodiment, there is provided a system, comprising:

a hardware computer processor configured to enable and disable execution of a run-time instrumentation facility on the hardware computer processor, and the hardware computer processor configured to:

fetch an instruction of a currently executing thread in a multi-threaded environment for execution by the hardware computer processor in a first state, the instruction one of a run-time instrumentation facility off (RIOFF) instruction and a run-time instrumentation facility on (RION) instruction, the fetching by the hardware computer processor;

based on determining, by the hardware computer processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, execute the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the hardware computer processor to indicate that run-time instrumentation data should not be captured by the hardware computer processor; and

enable the run-time instrumentation facility based on the instruction being the RION instruction, including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the hardware computer processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

[0006C] In accordance with a third major embodiment, there is provided a method for enabling and disabling execution of a run-time instrumentation facility on a processor, the method comprising:

fetching, by the processor, an instruction of a currently executing thread in a multi-threaded environment for execution by the processor in a first state, the instruction one of a run-time instrumentation facility off (RIOFF) instruction and a run-time instrumentation facility on (RION) instruction;

based on determining, by the processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, executing the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor; and

enabling the run-time instrumentation facility based on the instruction being the RION instruction, the enabling including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

Brief Description of Drawings

- [0007] [fig.1A]FIG. 1A is a diagram depicting an example host computer system in an embodiment;
- [fig.1B]FIG. 1B is a diagram depicting an example emulation host computer system in an embodiment;
- [fig.1C]FIG. 1C is a diagram depicting an example computer system in an embodiment;
- [fig.2]FIG. 2 is a diagram depicting an example computer network in an embodiment;
- [fig.3]FIG. 3 is a diagram depicting elements of a computer system in an embodiment;
- [fig.4A]FIG. 4A depicts detailed elements of a computer system in an embodiment;
- [fig.4B]FIG. 4B depicts detailed elements of a computer system in an embodiment;
- [fig.4C]FIG. 4C depicts detailed elements of a computer system in an embodiment;
- [fig.5]FIG. 5 depicts a schematic diagram of a system for run-time instrumentation of a processor in accordance with an embodiment;
- [fig.6]FIG. 6 depicts a portion of a run-time instrumentation control block (RICCB) including controls that are settable by a privileged state in an embodiment;
- [fig.7]FIG. 7 depicts a portion of a RICCB control block when the semi-privileged bit (K) is set to 1 in an embodiment;
- [fig.8]FIG. 8 depicts a collection buffer in accordance with an embodiment;

[fig.9]FIG. 9 depicts a reporting group in accordance with an embodiment;
 [fig.10]FIG. 10 depicts a process flow for implementing a run-time instrumentation facility in accordance with an embodiment;
 [fig.11]FIG. 11 depicts a run-time instrumentation off (RIOFF) instruction in accordance with an embodiment;
 [fig.12]FIG. 12 depicts a process flow of an RIOFF instruction in accordance with an embodiment;
 [fig.13]FIG. 13 depicts a run-time instrumentation on (RION) instruction in accordance with an embodiment;
 [fig.14]FIG. 14 depicts a process flow of an RION instruction in accordance with an embodiment; and
 [fig.15]FIG. 15 illustrates a computer program product in accordance with an embodiment.

Description of Embodiments

[0008] A run-time instrumentation facility allows a processor to collect hardware event data into a rolling history buffer, sample the history buffer, and store data from the history buffer into application accessible storage. It is generally desirable to limit sampling to a subset of the execution trace of an application as there may be elements in the execution path that may not be of interest. For example, if data collected by the run-time instrumentation facility will be used to enhance a program that was compiled by a compiler, then data related to the run-time environment and the operating system services may not be useful (e.g., because the compiler may have no ability to affect the run-time environment or the operating system services). Additionally, by limiting the range of the execution path that is sampled, the run-time instrumentation facility can reduce the cost of having to manage, buffer, and reduce the collected instrumentation data, thus keeping the overhead of the instrumentation mechanism to a minimum. Embodiments described here allow the run-time instrumentation to be turned off (disabled) and on (enabled) based on instructions from an application.

[0009] As described in embodiments herein, a first instruction, referred to as a run-time instrumentation on (RION) instruction designates a starting point for a sub-trace of the execution path to be sampled. Also as described herein, a second instruction, referred to as a run-time instrumentation off (RIOFF) instruction designates an ending point of the execution path to be sampled. These instructions may be inserted into an instruction stream at transition points between the compiled code and the run-time environment. In addition, the RION and RIOFF instructions may be inserted by the compiler into the instruction stream to focus the sampling mechanism on specific areas of the compiled code.

- [0010] A bit in a program status word (PSW) is used, in accordance with embodiments, to designate the state of the run-time instrumentation facility. The bit in the PSW indicates whether the current execution point is inside of a sampling interval (i.e., the run-time instrumentation facility is turned on and collecting data) or outside of a sampling interval (i.e., the run-time instrumentation facility is turned off and not collecting data) of the executing trace. As the PSW is saved and restored across context switches, the state of the sampling mechanism is naturally maintained across dispatches of the executing thread.
- [0011] An embodiment of the present invention is a hardware based run-time instrumentation facility for managed run-times. As used herein the term "managed run-time" refers to an environment that encapsulates a state and manages resources used to execute a program or application (e.g., Java(R) virtual machine or "JVM", operating system, middleware, etc.). Embodiments of the run-time instrumentation facility enable a program to collect information about program execution, including central processing unit (CPU) data. The collected information allows the manager run-time environment to acquire insights about the program from which the information is collected. Embodiments of the run-time instrumentation facility include a hardware facility for collecting sequences of events (e.g., taken branches, register values, etc.) in a collection buffer. The collection buffer (or a subset of the collection buffer containing the most recent records) is copied into a program buffer in the application's address space (for example the address space of a JVM) upon a programmable set of sample triggering events such as, but not limited to: a software directive in the form of an instruction inserted into the instruction stream; an interval of executed instructions are completed, a given elapsed time since the last sample expires, and/or a given hardware event such as data or instruction cache miss is observed.
- [0012] Dynamic compilers exploit runtime information, such as that collected by the hardware based run-time instrumentation facility described herein to perform online feedback directed optimizations. For example, information about important execution paths, profiled values and preferred branch directions can be used by a dynamic compiler to perform optimizations that specialize or version code, direct in-lining, reorder execution paths, and straighten branches. Data generated by embodiments described herein may also be used for other forms of optimization such as data reorganization.
- [0013] FIG. 1A, depicts the representative components of a host computer system 50 in an embodiment. Other arrangements of components may also be employed in a computer system. The representative host computer system 50 comprises one or more processors 1 in communication with main store (computer memory) 2 as well as I/O interfaces to storage devices 11 and networks 10 for communicating with other computers or SANs

and the like. The processor 1 is compliant with an architecture having an architected instruction set and architected functionality. The processor 1 may have dynamic address translation (DAT) 3 for transforming program addresses (virtual addresses) into a real address in memory. A DAT 3 typically includes a translation lookaside buffer (TLB) 7 for caching translations so that later accesses to the block of computer memory 2 do not require the delay of address translation. Typically a cache 9 is employed between the computer memory 2 and the processor 1. The cache 9 may be hierarchical having a large cache available to more than one CPU and smaller, faster (lower level) caches between the large cache and each CPU. In some embodiments, the lower level caches are split to provide separate low level caches for instruction fetching and data accesses. In an embodiment, an instruction is fetched from the computer memory 2 by an instruction fetch unit 4 via the cache 9. The instruction is decoded in an instruction decode unit 6 and dispatched (with other instructions in some embodiments) to instruction execution units 8. Typically several instruction execution units 8 are employed, for example an arithmetic execution unit, a floating point execution unit and a branch instruction execution unit. The instruction is executed by the instruction execution unit 8, accessing operands from instruction specified registers or the computer memory 2 as needed. If an operand is to be accessed (loaded or stored) from the computer memory 2, the load store unit 5 typically handles the access under control of the instruction being executed. Instructions may be executed in hardware circuits or in internal microcode (firmware) or by a combination of both.

[0014] In FIG. 1B, depicts an emulated host computer system 21 is provided that emulates a host computer system of a host architecture, such as the host computer system 50 of FIG. 1. In the emulated host computer system 21, a host processor (CPU) 1 is an emulated host processor (or virtual host processor) 29, and comprises a native processor 27 having a different native instruction set architecture than that of the processor 1 of the host computer system 50. The emulated host computer system 21 has memory 22 accessible to the native processor 27. In an embodiment, the memory 22 is partitioned into a computer memory 2 portion and an emulation routines memory 23 portion. The computer memory 2 is available to programs of the emulated host computer system 21 according to the host computer architecture. The native processor 27 executes native instructions of an architected instruction set of an architecture other than that of the emulated processor 29, the native instructions obtained from the emulation routines memory 23, and may access a host instruction for execution from a program in the computer memory 2 by employing one or more instruction(s) obtained in a sequence & access/decode routine which may decode the host instruction(s) accessed to determine a native instruction execution routine for emulating the function of the host instruction accessed. Other facilities that are defined for the host computer

system 50 architecture may be emulated by architected facilities routines, including such facilities as general purpose registers, control registers, dynamic address translation and input/output (I/O) subsystem support and processor cache for example. The emulation routines may also take advantage of function available in the native processor 27 (such as general registers and dynamic translation of virtual addresses) to improve performance of the emulation routines. Special hardware and off-load engines may also be provided to assist the native processor 27 in emulating the function of the host computer system 50.

- [0015] In a mainframe, architected machine instructions are used by programmers, usually today "C" programmers often by way of a compiler application. These instructions stored in the storage medium may be executed natively in a z/Architecture IBM Server, or alternatively in machines executing other architectures. They can be emulated in the existing and in future IBM mainframe servers and on other machines of IBM (e.g. pSeries(R) Servers and xSeries(R) Servers). They can be executed in machines running Linux on a wide variety of machines using hardware manufactured by IBM(R), Intel(R), AMDTM, Sun Microsystems and others. Besides execution on that hardware under a Z/Architecture(R), Linux can be used as well as machines which use emulation by Hercules, UMX, Fundamental Software, Inc. (FSI) or Platform Solutions, Inc. (PSI), where generally execution is in an emulation mode. In emulation mode, emulation software is executed by a native processor to emulate the architecture of an emulated processor.
- [0016] One or more of the components of the emulated host computer system 21 are further described in "IBM(R) z/Architecture Principles of Operation, " Publication No. SA22-7832-08, 9th Edition, August, 2010. IBM is a registered trademark of International Business Machines Corporation, Armonk, New York, USA. Other names used herein may be registered trademarks, trademarks or product names of International Business Machines Corporation or other companies.
- [0017] The native processor 27 typically executes emulation software stored in the emulation routines memory 23 comprising either firmware or a native operating system to perform emulation of the emulated processor. The emulation software is responsible for fetching and executing instructions of the emulated processor architecture. The emulation software maintains an emulated program counter to keep track of instruction boundaries. The emulation software may fetch one or more emulated machine instructions at a time and convert the one or more emulated machine instructions to a corresponding group of native machine instructions for execution by the native processor 27. These converted instructions may be cached such that a faster conversion can be accomplished. The emulation software maintains the architecture

rules of the emulated processor architecture so as to assure operating systems and applications written for the emulated processor operate correctly. Furthermore the emulation software provides resources identified by the emulated processor architecture including, but not limited to control registers, general purpose registers, floating point registers, dynamic address translation function including segment tables and page tables for example, interrupt mechanisms, context switch mechanisms, time of day (TOD) clocks and architected interfaces to I/O subsystems such that an operating system or an application program designed to run on the emulated processor 29, can be run on the native processor 27 having the emulation software.

[0018] A specific instruction being emulated is decoded, and a subroutine called to perform the function of the individual instruction. An emulation software function emulating a function of an emulated processor 29 is implemented, for example, in a "C" subroutine or driver, or some other method of providing a driver for the specific hardware as will be within the skill of those in the art after understanding the description of the preferred embodiment.

[0019] In an embodiment, the invention may be practiced by software (sometimes referred to licensed internal code, firmware, micro-code, milli-code, pico-code and the like, any of which would be consistent with the present invention). Referring to FIG. 1A, software program code which embodies the present invention is accessed by the processor also known as a CPU (Central Processing Unit) 1 of the host computer system 50 from the storage device 11 such as a long-term storage media, a CD-ROM drive, tape drive or hard drive. The software program code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from the computer memory 2 or storage of one computer system over a network 10 to other computer systems for use by users of such other systems.

[0020] Alternatively, the program code may be embodied in the computer memory 2, and accessed by the processor 1 using a processor bus (not shown). Such program code includes an operating system which controls the function and interaction of the various computer components and one or more application programs. Program code is normally paged from a dense media such as the storage device 11 to computer memory 2 where it is available for processing by the processor 1. The techniques and methods for embodying software program code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to electronic memory modules (RAM), flash memory, compact discs (CDs), DVDs, Magnetic Tape and the like is often referred to as a "computer program product. " The computer program product medium is typically readable by a

processing circuit preferably in a computer system for execution by the processing circuit.

[0021] FIG. 1C illustrates a representative workstation or server hardware system in which the present invention may be practiced. The system 100 of FIG. 1C comprises a representative base computer system 101, such as a personal computer, a workstation or a server, including optional peripheral devices. The base computer system 101 includes one or more processors 106 and a bus (not shown) employed to connect and enable communication between the one or more processors 106 and the other components of the base computer system 101 in accordance with known techniques. The bus connects the processor 106 to memory 105 and long-term storage 107 which may include a hard drive (including any of magnetic media, CD, DVD and Flash Memory for example) or a tape drive for example. The base computer system 101 may also include a user interface adapter, which connects the one or more processors 106 via the bus to one or more interface devices, such as a keyboard 104, a mouse 103, a printer/scanner 110 and/or other interface devices, which may be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus also connects the one or more processors to a display device 102, such as an LCD screen or monitor via a display adapter.

[0022] The base computer system 101 may communicate with other computers or networks of computers by way of a network adapter capable of communicating 108 with a network 109. Example network adapters are communications channels, token ring, Ethernet or modems. Alternatively, the base computer system 101 may communicate using a wireless interface, such as a cellular digital packet data (CDPD) card. The base computer system 101 may be associated with such other computers in a local area network (LAN) or a wide area network (WAN), or the base computer system 101 may be a client in a client/server arrangement with another computer, etc.

[0023] FIG. 2 illustrates a data processing network 200 in which the present invention may be practiced. The data processing network 200 may include a plurality of individual networks, such as a wireless network and a wired network, each of which may include a plurality of individual workstations 201, 202, 203, 204 and or the base computer system 101 of FIG. 1C. Additionally, as those skilled in the art will appreciate, one or more LANs may be included, where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

[0024] Programming code 111 may be embodied in the memory 105, and accessed by the processor 106 using the processor bus. Such programming code includes an operating system which controls the function and interaction of the various computer components and one or more application programs 112. Program code is normally paged from long-term storage 107 to high-speed memory 105 where it is available for

processing by the processor 106. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape and the like is often referred to as a "computer program product". The computer program product medium is typically readable by a processing circuit preferably in a computer system for execution by the processing circuit.

[0025] The cache that is most readily available to the processor (normally faster and smaller than other caches of the processor) is the lowest (L1 or level one) cache and main store (main memory) is the highest level cache (L3 if there are 3 levels). The lowest level cache is often divided into an instruction cache (I-Cache) holding machine instructions to be executed and a data cache (D-Cache) holding data operands.

[0026] Still referring to FIG. 2, the networks may also include mainframe computers or servers, such as a gateway computer (client server) 206 or application server (remote server) 208 which may access a data repository and may also be accessed directly from a workstation 205. A gateway computer 206 serves as a point of entry into each network 207. A gateway is needed when connecting one networking protocol to another. The gateway computer 206 may be preferably coupled to another network (the Internet 207 for example) by means of a communications link. The gateway computer 206 may also be directly coupled to the one or more workstations 101, 201, 202, 203, and 204 using a communications link. The gateway computer may be implemented utilizing an IBM eServer™ zSeries(R) z9(R) Server available from International Business Machines Corporation.

[0027] In an embodiment, software programming code which embodies the present invention is accessed by the processor 106 of the base computer system 101 from long-term storage media, such as the long-term storage 107 of FIG. 1C. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users 210 and 211 from the memory or storage of one computer system over a network to other computer systems for use by users of such other systems.

[0028] Referring to FIG. 3, an exemplary processor embodiment is depicted for processor 106. One or more levels of cache 303 are employed to buffer memory blocks in order to improve the performance of the processor 106. The cache 303 is a high speed buffer holding cache lines of memory data that are likely to be used. Typical cache lines are 64, 128 or 256 bytes of memory data. In an embodiment, separate caches are employed for caching instructions than for caching data. Cache coherence (synchronization of

copies of lines in memory and the caches) is often provided by various "snoop" algorithms well known in the art. Main storage, such as memory 105 of a processor system is often referred to as a cache. In a processor system having 4 levels of cache 303 memory 105 is sometimes referred to as the level 5 (L5) cache since it is typically faster and only holds a portion of the non-volatile storage (DASD, Tape etc) that is available to a computer system. Memory 105 "caches" pages of data paged in and out of the memory 105 by the operating system.

[0029] A program counter (instruction counter) 311 keeps track of the address of the current instruction to be executed. A program counter in a z/Architecture processor is 64 bits and may be truncated to 31 or 24 bits to support prior addressing limits. A program counter is typically embodied in a program status word (PSW) of a computer such that it persists during context switching. Thus, a program in progress, having a program counter value, may be interrupted by, for example, the operating system (i.e., the current context switches from the program environment to the operating system environment). The PSW of the program maintains the program counter value while the program is not active, and the program counter (in the PSW) of the operating system is used while the operating system is executing. In an embodiment, the program counter is incremented by an amount equal to the number of bytes of the current instruction. Reduced Instruction Set Computing (RISC) instructions are typically fixed length while Complex Instruction Set Computing (CISC) instructions are typically variable length. Instructions of the IBM z/Architecture are CISC instructions having a length of 2, 4 or 6 bytes. The program counter 311 is modified by either a context switch operation or a branch taken operation of a branch instruction for example. In a context switch operation, the current program counter value is saved in the PSW along with other state information about the program being executed (such as condition codes), and a new program counter value is loaded pointing to an instruction of a new program module to be executed. A branch taken operation is performed in order to permit the program to make decisions or loop within the program by loading the result of the branch instruction into the program counter 311.

[0030] In an embodiment, an instruction fetch unit 305 is employed to fetch instructions on behalf of the processor 106. The instruction fetch unit 305 either fetches the "next sequential instructions," the target instructions of branch taken instructions, or the first instructions of a program following a context switch. In an embodiment, the instruction fetch unit 305 employs prefetch techniques to speculatively prefetch instructions based on the likelihood that the prefetched instructions might be used. For example, the instruction fetch unit 305 may fetch 16 bytes of instructions that include the next sequential instruction and additional bytes of further sequential instructions.

[0031] The fetched instructions are then executed by the processor 106. In an embodiment,

the fetched instruction(s) are passed to a decode/dispatch unit 306 of the instruction fetch unit 305. The decode/dispatch unit 306 decodes the instruction(s) and forwards information about the decoded instruction(s) to appropriate execution units 307, 308, and/or 310. An execution unit 307 receives information about decoded arithmetic instructions from the instruction fetch unit 305 and will perform arithmetic operations on operands according to the operation code (opcode) of the instruction. Operands are provided to the execution unit 307 either from the memory 105, architected registers 309, or from an immediate field of the instruction being executed. Results of the execution, when stored, are stored either in memory 105, architected registers 309 or in other machine hardware (such as control registers, PSW registers and the like).

[0032] A processor 106 typically has one or more execution units 307, 308, and 310 for executing the function of the instruction. Referring to FIG. 4A, an execution unit 307 may communicate with the architected registers 309, the decode/dispatch unit 306, the load/store unit 310 and other processor units 401 by way of interfacing logic 407. The execution unit 307 may employ several register circuits 403, 404, and 405 to hold information that the arithmetic logic unit (ALU) 402 will operate on. The ALU 402 performs arithmetic operations such as add, subtract, multiply and divide as well as logical function such as and, or and exclusive-or (xor), rotate and shift. In an embodiment, the ALU supports specialized operations that are design dependent. Other circuits may provide other architected facilities 408 including condition codes and recovery support logic for example. Typically the result of an ALU operation is held in an output register circuit 406 which can forward the result to a variety of other processing functions. In other embodiments, there are many arrangements of processor units, the present description is only intended to provide a representative understanding of one embodiment.

[0033] An ADD instruction for example would be executed in an execution unit 307 having arithmetic and logical functionality while a floating point instruction for example would be executed in a floating point execution unit (not shown) having specialized floating point capability. Preferably, an execution unit operates on operands identified by an instruction by performing an opcode defined function on the operands. For example, an ADD instruction may be executed by an execution unit 307 on operands found in two architected registers 309 identified by register fields of the instruction.

[0034] The execution unit 307 performs the arithmetic addition on two operands and stores the result in a third operand where the third operand may be a third register or one of the two source registers. The execution unit 307 preferably utilizes an arithmetic logic unit (ALU) 402 that is capable of performing a variety of logical functions such as shift, rotate, and, or and XOR as well as a variety of algebraic functions including any of add, subtract, multiply, divide. Some ALUs 402 are designed for scalar operations

and some for floating point. In embodiments, data may be big endian (where the least significant byte is at the highest byte address) or little endian (where the least significant byte is at the lowest byte address) depending on architecture. The IBM z/Architecture is big endian. Signed fields may be sign and magnitude, 1's complement or 2's complement depending on architecture. A 2's complement number is advantageous in that the ALU does not need to design a subtract capability since either a negative value or a positive value in 2's complement requires only an addition within the ALU. Numbers are commonly described in shorthand, where a 12 bit field defines an address of a 4,096 byte block and is commonly described as a 4 Kbyte (Kilo-byte) block for example.

[0035] Referring to FIG. 4B, Branch instruction information for executing a branch instruction is typically sent to a branch unit 308 which employs a branch prediction algorithm such as a branch history table 432 to predict the outcome of the branch before other conditional operations are complete. The target of the current branch instruction will be fetched and speculatively executed before the conditional operations are complete. When the conditional operations are completed the speculatively executed branch instructions are either completed or discarded based on the conditions of the conditional operation and the speculated outcome. A typical branch instruction may test condition codes and branch to a target address if the condition codes meet the branch requirement of the branch instruction, a target address may be calculated based on several numbers including ones found in register fields or an immediate field of the instruction for example. In an embodiment, the branch unit 308 may employ an ALU 426 having a plurality of input register circuits 427, 428, and 429 and an output register circuit 430. The branch unit 308 may communicate with general registers, decode/dispatch unit 306 or other circuits 425 for example.

[0036] The execution of a group of instructions may be interrupted for a variety of reasons including a context switch initiated by an operating system, a program exception or error causing a context switch, an I/O interruption signal causing a context switch or multi-threading activity of a plurality of programs (in a multi-threaded environment) for example. In an embodiment, a context switch action saves state information about a currently executing program and then loads state information about another program being invoked. State information may be saved in hardware registers or in memory for example. State information includes a program counter value pointing to a next instruction to be executed, condition codes, memory translation information and architected register content. A context switch activity may be exercised by hardware circuits, application programs, operating system programs or firmware code (microcode, pico-code or licensed internal code (LIC) alone or in combination.

[0037] A processor accesses operands according to instruction defined methods. The in-

struction may provide an immediate operand using the value of a portion of the instruction, may provide one or more register fields explicitly pointing to either general purpose registers or special purpose registers (floating point registers for example). The instruction may utilize implied registers identified by an opcode field as operands. The instruction may utilize memory locations for operands. A memory location of an operand may be provided by a register, an immediate field, or a combination of registers and immediate field as exemplified by the z/Architecture long displacement facility wherein the instruction defines a base register, an index register and an immediate field (displacement field) that are added together to provide the address of the operand in memory. Location herein implies a location in main memory (main storage) unless otherwise indicated.

[0038] Referring to FIG. 4C, a processor accesses storage using a load/store unit 310. The load/store unit 310 may perform a load operation by obtaining the address of the target operand in memory through the cache/memory interface and loading the operand in an architected register 309 or another memory location, or may perform a store operation by obtaining the address of the target operand in memory and storing data obtained from an architected register 309 or another memory location in the target operand location in memory. The load/store unit 310 may be speculative and may access memory in a sequence that is out-of-order relative to the instruction sequence; however the load/store unit 310 maintains the appearance to programs that instructions were executed in order. A load/store unit 310 may communicate with architected registers 309, decode/dispatch unit 306, cache/memory interface or other elements 455 and comprises various register circuits, ALUs 458 and control logic 463 to calculate storage addresses and to provide pipeline sequencing to keep operations in-order. Some operations may be out of order but the load/store unit provides functionality to make the out of order operations appear to the program as having been performed in order as is well known in the art.

[0039] Preferably addresses that an application program "sees" are often referred to as virtual addresses. Virtual addresses are sometimes referred to as "logical addresses" and "effective addresses. " These virtual addresses are virtual in that they are redirected to physical memory location by one of a variety of DAT technologies such as the DAT 312 of FIG. 3, including, but not limited to prefixing a virtual address with an offset value, translating the virtual address via one or more translation tables, the translation tables including at least a segment table and a page table alone or in combination, preferably, the segment table having an entry pointing to the page table. In z/Architecture, a hierarchy of translations is provided including a region first table, a region second table, a region third table, a segment table and an optional page table. The performance of the address translation is often improved by utilizing a translation

look-aside buffer (TLB) which comprises entries mapping a virtual address to an associated physical memory location. The entries are created when DAT 312 translates a virtual address using the translation tables. Subsequent use of the virtual address can then utilize the entry of the fast TLB rather than the slow sequential translation table accesses. The TLB content may be managed by a variety of replacement algorithms including least recently used (LRU).

- [0040] In the case where the processor 106 is a processor of a multi-processor system, each processor has responsibility to keep shared resources such as I/O, caches, TLBs and Memory interlocked for coherency. In an embodiment, "snoop" technologies will be utilized in maintaining cache coherency. In a snoop environment, each cache line may be marked as being in any one of a shared state, an exclusive state, a changed state, an invalid state and the like in order to facilitate sharing.
- [0041] The I/O units 304 of FIG. 3 provide the processor 106 with means for attaching to peripheral devices including tape, disc, printers, displays, and networks for example. The I/O units 304 are often presented to the computer program by software drivers. In mainframes such as the z/Series from IBM, channel adapters and open system adapters are I/O units of the mainframe that provide the communications between the operating system and peripheral devices.
- [0042] Instrumentation data is data related to the operations of the processor 106. In an embodiment, access to instrumentation data and other system level metrics may be restricted, or unavailable. A computer processor operates under a privileged state (or supervisor state), and a lesser-privileged state (or problem state). In the privileged state, a program may have access to all system resources via privileged operations (e.g., access to all control registers and the supervisor memory space). The privileged state is also referred to as privileged mode or supervisor mode. An operating system executing on the computer processor may be operating in the privileged state. The lesser-privileged state is a non-privileged state where access to system resources is limited. For example, application programs running in lesser-privileged state may have limited or no access to control registers and may access only user memory space assigned to the application program by the operating system. The lesser-privileged state is typically assigned to application programs executed under control of an operating system, and no privileged operations can be performed in the lesser-privileged state. The lesser-privileged state is also known as a problem state, problem mode or user mode.
- [0043] One such restricted resource that is not write accessible to a program executing in the lesser-privileged state is the program status word (PSW). The PSW may comprise a program counter of the next instruction to be executed, a condition code field usable by branch instructions, an instrumentation control field for indicating whether instrumentation is enabled or disabled, and other information used to control instruction se-

quencing and to determine the state of the computer processor including the privilege state assigned to the program. In a multithreaded processing environment, multiple programs share, or time slice, the available computer processor capacity. Each of the programs has context information including an associated PSW, an origin address of an address translation table for accessing main storage assigned to the program, a set of general purpose register current values, control registers, floating point registers, etc. The currently active, or controlling PSW, is called the current PSW. It governs the program currently being executed. The computer processor has an interruption capability, which permits the computer processor to context switch rapidly to another program in response to exception conditions and external stimuli. When an interruption occurs, the computer processor places the current PSW in an assigned storage location, called the old-PSW location, for the particular class of interruption. The computer processor fetches a new PSW from a second assigned storage location. This new context determines the next program to be executed. In an embodiment, these storage locations are located in a memory location accessible to the computer processor. When the computer processor has finished processing the interruption, the program handling the interruption may reload the old context including the old PSW, making it again the current PSW, so that the interrupted program can continue.

- [0044] The fields of the PSW may be referenced either explicitly (e.g., when instruction execution reads part of the PSW bits), or implicitly (e.g., in instructions fetching, operand fetching, address generation calculations, address generation sources, etc.). The explicit reference is generally performed at execution time, whereas the implicit reference is generally performed at different stages of the pipeline during instruction execution (i.e., instruction fetch, instruction decode, execution time and completion time). Individual fields in the PSW may be referenced or updated independently of each other.
- [0045] In an embodiment, by manipulating the context, an operating system controls computer processing resources, including enabling run-time-instrumentation by the computer processor. The run-time-instrumentation may be enabled or disabled during the execution of the operating system, as well as by any software applications executed by the operating system. The enabled/disabled state of run-time-instrumentation is saved as context information in the PSW associated with a program.
- [0046] A run-time-instrumentation (RI) facility may be incorporated on models implementing z/Architecture. When the RI facility is installed and enabled, data is collected during program execution into one or more collection buffers within the CPU and then reported to a program buffer. Each unit of information stored is called a reporting group. The contents of a reporting group consist of multiple records whose contents represent events recognized by the CPU during program execution.

- [0047] When the run-time-instrumentation facility is installed in a configuration, a PSW field (RI bit) enables run-time-instrumentation. Validity of the run-time-instrumentation controls determines the capability of turning on the RI bit, but when RI is one, the CPU controls are valid and run-time-instrumentation is enabled. The run-time-instrumentation facility may include the following instructions: load run-time-instrumentation controls, modify run-time-instrumentation controls, run-time-instrumentation emit, run-time-instrumentation next, run-time-instrumentation off, run-time-instrumentation on, store run-time-instrumentation controls, and test run-time-instrumentation controls.
- [0048] The load run-time-instrumentation controls (LRIC) instruction initializes the run-time-instrumentation controls that govern run-time-instrumentation. The modify run-time-instrumentation controls (MRIC) instruction modifies all or a subset of the run-time-instrumentation controls originally established by LRIC. The run-time-instrumentation emit (RIEMIT) instruction collects the value of a general register by storing it into a collection buffer. The run-time-instrumentation next (RINEXT) instruction performs directed sampling of the next, sequential instruction (NSI) after RINEXT. The run-time-instrumentation off (RIOFF) instruction disables run-time-instrumentation. The run-time-instrumentation on (RION) instruction enables run-time-instrumentation. The store run-time-instrumentation controls (STRIC) instruction places the current values of the run-time-instrumentation controls into a specified storage location. The test run-time-instrumentation controls (TRIC) instruction examines the run-time-instrumentation controls. If valid, the state of a controls-altered indicator is set.
- [0049] The run-time-instrumentation facility includes the ability for making a measurement-alert external interruption pending. Some of the information collected by run-time-instrumentation and reported to a program buffer is model-dependent and thus not defined. Samples and data provided by the run-time-instrumentation facility are intended for statistical estimation of performance characteristics, are substantially accurate, and may not be repeatable. For example, regardless of sampling mode, it is unpredictable if a sample instruction that caused an exception or is associated with certain system internal activities would result in the store of a reporting group and, if stored, whether the model-dependent data included in run-time-instrumentation data is affected.
- [0050] A collection buffer is used to capture a set of records whose contents report on events recognized by the processor during program execution. Examples are: execution of one or more taken branches, transactional-execution abort events, instruction-fetch cache misses, data fetch or store cache misses, and an operand of the RIEMIT instruction. Execution of the RIEMIT instruction collects the value of a general register by storing

it into the collection buffer. Additional data can be collected and/or stored in other buffers, such as an instruction-data buffer.

[0051] Reporting is subject to reporting controls. When a sample instruction is identified, each reporting control enables the checking of a corresponding condition. If a corresponding condition exists, a reporting group is formed and stored. A reporting group is not stored when no reporting control is enabled or the corresponding condition does not exist for an enabled reporting control. Data reported about a sample instruction may be acquired from the instruction-data buffer and other model-dependent sources, and then used to create the contents of one or more records of the reporting group, one such record being an instruction record.

[0052] Record types that may be captured in the reporting group store include: filler, extra, begin, timestamp, instruction, emit, TX abort, call, return, and transfer. A filler record is used in a reporting group when the number of valid records in the collection buffer is not sufficient to fill a reporting group of the current reporting-group size. An extra record may be used in the extra section of a reporting group. A begin record is the first record of the first reporting group. A timestamp record is stored as record 0 of every reporting group other than the first reporting group. An instruction record is created when a reporting group is stored for a sample instruction as the last record of the reporting group. An emit record is created by successful execution of RIEMIT. A transaction-execution (TX) mode abort record is created by either an implicit abort or by execution of a transaction abort instruction. A call record is created by execution of a branch instruction which is categorized as a call-type branch instruction. A return record is created by execution of a return-type branch instruction which is categorized as a return instruction. A transfer record is created by execution of a branch instruction which meets certain condition code criteria.

[0053] FIG. 5 depicts a schematic diagram of a system for run-time-instrumentation of a processor that may be implemented in an embodiment. In an embodiment, the system 500 includes a central processing unit (CPU) such as the processor 106 of FIG. 1. In an embodiment, the processor 106 is a single processor. In an alternate embodiment, the processor 106 is a single processing core of a multi-core processor. In an embodiment, the processor 106 is capable of operating at varying speeds.

[0054] In an embodiment, the processor 106 further includes a register 510. The register 510 is a hardware register capable of storing words of data for use by the processor 106. The register 510 includes one or more latches for storing bits of data that are accessible by the processor 106. The register 510 may include general purpose registers and control registers for example. The processor 106 additionally includes an instrumentation module 506 that is in communication with the register 510. The instrumentation module 506 is a processing circuit that controls the instrumentation of the

processor 106. The instrumentation module 506 is configured to collect instrumentation data, such as the execution path of one or more taken branches, transactional execution abort events, various runtime operands, timestamp information, etc. directly from the processor 106. The instrumentation module 506 collects the instrumentation data from the processor 106, and stores the instrumentation data in a collection buffer 508. In an embodiment, the collection buffer 508 is a circular buffer that collects data received from the instrumentation module 506, and when the circular buffer is filled it overwrites the oldest data with new data.

- [0055] The processor 106 executes one or more operating systems 516 and one or more applications 518. The one or more operating systems 516 and one or more applications 518 are stored in a storage 520, such as a hard drive, CD-ROM, flash memory, etc. and are loaded into a main memory 514 in a runtime memory 504 area reserved for storing one or more active pieces of the currently executing operating system and/or application, called pages, which are loaded from the storage 520 into runtime memory 504 as needed. In an embodiment, each of the operating systems execute as a virtual machine managed by a hypervisor (not shown) and executed by the processor 106.
- [0056] In an embodiment the processor 106 loads a PSW 512 in the register 510 from PSW data 512 in the main memory 514 for the currently executing operating system or application from the main memory 514 and sets one or more processor settings in, for example, the register 510. In an embodiment, the PSW in the register 510, includes one or more bits for enabling and controlling the instrumentation module 506.
- [0057] The one or more applications 518 include software applications compiled to execute on a specific operating system, interpreted code executing on an interpreter (e.g., Java), or operating system support threads (e.g., process management, daemons, etc.). Each of the one or more operating systems 516 and or the one or more applications 518 may execute an instruction to trigger the instrumentation module 506 to start, or to stop, the collecting instrumentation data.
- [0058] In an embodiment, one of the one or more applications 518 executes an instruction that has been determined to be a sample instruction, thereby creating a sample point at the completion of execution of the sample instruction and that then causes the instrumentation module 506 to move the application's collected data from the collection buffer 508, to a program buffer 522 in main memory 514 that is accessible to the application. The main memory 514 may be any addressable memory known in the art. In an embodiment, the main memory 514 may include a fast-access buffer storage, sometimes called a cache. Each CPU may have an associated cache. In an additional embodiment, the main memory 514 is dynamic random access memory (DRAM). In a yet another embodiment, the main memory is a storage device, such as a computer hard drive, or flash memory accessible by an application.

- [0059] To configure run-time instrumentation controls, the processor 106 supports a load run-time instrumentation controls (LRIC) instruction. Beyond the specific LRIC fields described further herein, it will be understood that additional fields can be defined to support other functionality. The LRIC instruction can be used to load and initially configure run-time instrumentation and is supported by instrumentation module 506 of FIG. 5. In an embodiment, the instrumentation module 506, also referred to as run-time instrumentation module 506, implements run-time-instrumentation controls and reporting controls. A current state of run-time instrumentation controls can be stored from register 510 of FIG. 5 into main memory 514 using the store run-time controls (STRIC) instruction. The definition of various fields of a control block loadable as an operand of the LRIC instruction is also used herein to refer to the state of corresponding values of the run-time-instrumentation controls.
- [0060] FIG. 6 depicts a portion of a run-time-instrumentation controls control block (RICCB) including controls that are settable by a privileged state in an embodiment. The control block portion 600 may include additional values other than those described in reference to FIG. 6. Modification to the control block portion 600 may be performed by an LRIC instruction.
- [0061] The control block portion includes a validity bit 602 (V bit). The validity bit 602 indicates the validity of the set of run-time-instrumentation controls in the processor, as they were previously set by an LRIC instruction.
- [0062] The control block also includes an S bit 604, which is used to determine if the lesser-privileged state program is allowed to execute an MRIC instruction. The K bit 606 indicates if the lesser-privileged state program is permitted to execute in a semi-privileged state with regard to the run-time-instrumentation controls, such as the origin address, and the limit address of the run-time-instrumentation controls. The H bit 608 determines whether the address controls (i.e., the origin address, limit address, and current address) refer to a primary virtual address space or a home virtual address space. The 0 bit 610 is ignored and treated as a 0.
- [0063] A lesser-privileged state sample reporting control bit 612 (Ps bit) is used in conjunction with lesser-privileged state programs. When in the lesser-privileged state and the Ps bit 612 in the run-time-instrumentation controls is zero, the reporting controls of the run-time-instrumentation controls are ignored when run-time-instrumentation is enabled, and thus do not cause a reporting group to be stored. When in the lesser-privileged state and the Ps bit 612 in the run-time-instrumentation controls is one, the reporting controls are checked and used according to their defined function.
- [0064] A supervisor-state sample reporting control bit 614 (Qs bit) is used in conjunction with supervisor-state programs. When in the supervisor state and the Qs bit 614 in the run-time-instrumentation controls is zero, the reporting controls of the run-

time-instrumentation controls are ignored when run-time-instrumentation is enabled, and thus do not cause a reporting group to be stored. When in the supervisor state and the Qs bit 614 in the run-time-instrumentation controls is one, the reporting controls are checked and used according to their defined function.

- [0065] The lesser-privileged state collection buffer control bit 616 (Pc bit) controls updates to the collection buffer 508 of FIG. 5. When in lesser-privileged state and the Pc bit 616 in the run-time-instrumentation controls is zero, collection buffer controls of the run-time-instrumentation controls are ignored when run-time-instrumentation is enabled and updates of the collection buffer 508 are prevented. When in the lesser-privileged state and the Pc bit 616 in the run-time-instrumentation controls is one, the collection buffer controls are checked and used according to their defined function.
- [0066] The supervisor-state collection buffer control bit 618 (Qc bit) controls updates to the collection buffer 508. When in supervisor state and the Qc bit 618 in the run-time-instrumentation controls is zero, collection buffer controls of the run-time-instrumentation controls are ignored when run-time-instrumentation is enabled and the updates to the collection buffer 508 are prevented. When in supervisor state and the Qc bit 618 in the run-time-instrumentation controls is one, the indicated collection-buffer controls are checked and used according to their defined function.
- [0067] The G bit 620 is the pending control of a run-time-instrumentation-halted interruption, also called a halted interruption. When the G bit 620 is zero, a halted interruption is not pending. When the G bit 602 is one, a halted interruption is pending. When the first reporting group in a program buffer 522 is written, the G bit 620 is set to zero. That is, when run-time-instrumentation program-buffer origin address (ROA) 702 equals a run-time-instrumentation program buffer current address (RCA) 706 of FIG. 7, the G bit 620 is set to zero. When an attempt to store other than the first reporting group in program buffer 522 is made, the G bit 620 is set to zero if the run-time-instrumentation-halted condition does not exist, and the reporting group is stored. When an attempt to store other than the first reporting group in program buffer 522 is made, the G bit 620 is set to one if the run-time-instrumentation-halted condition does exist, and the reporting group is not stored.
- [0068] The U bit 622 is the enablement control for a buffer-full interruption and a halted interruption. When U bit 622 is zero, generation of an interruption request is disabled and, if pending, remains pending.
- [0069] The L bit 624 is the pending control of a buffer-full interruption. When L bit 624 is zero, a buffer-full interruption is not pending. When L bit 624 is one, a buffer-full interruption is pending.
- [0070] The key field 626 is a 4-bit unsigned integer whose value is used as a storage-protect key for the store of a reporting group. A store of a reporting group is permitted only

when the storage key matches the access key associated with the request for storage access, and a fetch is permitted when the storage key matches the access key or when a fetch-protection bit of the storage key is zero. The keys match when the four access control bits of the storage key are equal to the access key, or when the access key is zero.

[0071] FIG. 7 depicts a portion of an RICCB control block when MRIC is permitted to execute in semi-privileged mode (i.e., K bit is one). The control block 700 can also be an operand of an LRIC instruction for initialization of run-time-instrumentation controls. The control block 700 may include additional values other than those described in reference to FIG. 7. In an embodiment, sections of the MRIC instruction operand that are not otherwise designated are inaccessible by a lesser-privileged state program. When the semi-privileged mode is permitted, a run-time-instrumentation program-buffer origin address (ROA) 702 and a run-time-instrumentation program-buffer limit address (RLA) 704 are set with the MRIC instruction by the lesser-privileged state program. The ROA 702 is the location of the first byte of the program buffer 522 of FIG. 5. The RLA 704 indicates the location of the last byte of the program buffer 522.

[0072] In an embodiment, a run-time-instrumentation program buffer current address (RCA) 706 may be updated by the MRIC instruction. The RCA 706 is the location in the program buffer 522 of a next reporting group to be stored. The RCA 706 examines the reporting group size field 744 (RGS field) and affects the number of significant bit positions used to form the address of the program buffer 522. The 64-bit RCA 706 is word 0, bit positions 0 through 26-RGS of word 1, and RGS+5 binary zeros appended on the right. This is the starting location in the program buffer 522 of FIG. 5 of a subsequent reporting group that will be stored in the program buffer 522. The reporting group is a unit of information that is created by the instrumentation module 506, and subsequently stored in the program buffer 522. In an embodiment, when the RGS field 744 specified by the RCA 706 is not equal to the run-time-instrumentation control's current reporting group size (i.e., the RCA 706 would change the RGS field 744) then the RCA 706 is set to the ROA 702.

[0073] A remaining sample interval count field 742 (RSIC field) may be updated by the lesser-privileged program using the MRIC instruction. The RSIC field 742 includes a 64-bit unsigned integer that indicates a remaining sample interval count. When the value of the RSIC field 742 in the run-time-instrumentation controls is zero or equal to the value in a scaling factor field 740 (SF field), and run-time-instrumentation is enabled, then the next sample interval is a full interval based on the sampling mode 708 (M) and SF field 740 values. When RSIC field 742 is nonzero and less than the SF field 740 and run-time-instrumentation is enabled, the next sample interval is a partial

interval. When the RSIC field 742 is nonzero and greater than the SF field 740 value and run-time-instrumentation is enabled, the next sample interval is an extended interval. When an extended interval expires, the next interval is based on the SF field 740 value. When the RSIC field 742 is set to a nonzero value, it is subject to the same model-dependent maximum limit to which the SF field 740 is also subject. When the original value of the RSIC field 742 is zero, the sampling mode will dictate whether the RSIC field 742 is set to the value in the SF field 740 during execution of LRIC and MRIC instructions, or whether it continues to show as zero until run-time-instrumentation is enabled.

[0074] The SF field 740 contains a 64-bit unsigned integer whose value is a scaling factor count of units. The dimension of the units is determined from the mode field 708 (M field). When the value in the RSIC field 742 is zero, the SF field 740 provides an initial value of the RSIC field 742 that is decremented to zero at which point the current instruction is recognized as a sample instruction, and the interval count is refreshed from the SF field 740 value. A valid value of the SF field 740 is in the range one to $2^{64} - 1$. If zero is specified, a value of one is assumed. However, each model may have both a minimum and a maximum value of the SF field 740. The minimum and maximum values may also be different based on the mode field 708. If a value less than the minimum is specified, the model-dependent minimum value is loaded. If a value greater than the maximum value is specified, the model-dependent maximum value is loaded.

[0075] The DC control field 736 is a 4-bit unsigned integer whose value designates a cache-latency level associated with a data fetch or store cache miss. That is, the sample instruction encountered a data access cache miss. Unless prohibited by another run-time-instrumentation control, an attempt is made to store a reporting group representing the sample instruction whose data access recognized a miss at a cache-latency level numerically greater than or equal to the level designated by the value of the DC control field 736. The cache structure and cache-latency level for data access is model dependent. For an instruction with multiple or long operands, it is model dependent which, if any, operand access is used for reporting control. Model-dependent behavior may ignore the value of the DC control field 736 and thus not use it as a reason to store a reporting group.

[0076] The IC field 734 is a 4-bit unsigned integer whose value designates a cache-latency level associated with an instruction-fetch cache miss. That is, the fetch of the sample instruction encountered an instruction-fetch cache miss. For both the IC field 734 and DC control field 736, a cache-latency level is an abstraction of how far a certain cache level access is from the observing processor. The latency level depends on the combination of the amount of nested cache levels between the processor and main storage,

and how such cache levels are shared among multiple processors. A larger latency level generally corresponds to a more time-consuming access. Values in the IC field 734 and DC control field 736 may be thought of as zero-origin identification of a cache-latency level. For example, a value of zero corresponds to an L1 cache (i.e., the cache that is closest to the processor). A value of one is therefore the next layer of cache which may be known as an L2 cache, or even an L1.5 cache in some machines. Values of 2-15 designate the logical progression of additional cache-latency layers until main memory is reached, but not including main memory itself. Generally, cache structures do not go as deep as fifteen layers. Therefore, a value of 15 in the IC field 734 and DC control field 736 is interpreted as a special case, meaning that a cache miss on instruction fetch or data access, respectively and regardless of cache-latency level, is not recognized for the purpose of generating the store of a reporting group. Unless prohibited by another run-time-instrumentation control, an attempt is made to store a reporting group representing the sample instruction whose fetch recognized a miss at a cache-latency level numerically greater than or equal to the level designated by the value of the IC field 734. The cache structure and cache-latency level for instruction fetching is model dependent. Model-dependent behavior may ignore the value of the IC field 734 and thus not use it as a reason to store a reporting group.

- [0077] The cache-latency-level-override reporting control bit 732 (F bit) is for non-branch instructions and for branch-prediction controls. When the F bit 732 in the run-time-instrumentation controls is zero, the cache-reporting controls (IC field 734 and DC control field 736) of the run-time-instrumentation controls are checked and used according to their defined function. The branch-prediction controls (BPxn 722, BPxt 724, BPti 726, and BPni 728 bits) of the run-time-instrumentation controls are checked and used according to their defined function. When the F bit 732 is one, these same controls are ignored and a reporting group is stored unless prohibited by another control.
- [0078] The data-cache-miss control bit 730 (D bit) indicates if a reporting group is to be stored. If the D bit 730 is one, an extra type record may or may not be placed in the extra section of the reporting group which contains model dependent data about the sample instruction.
- [0079] The MRIC instruction includes branch-prediction (BP) reporting controls (BPxn 722, BPxt 724, BPti 726, and BPni 728). If a BP reporting control bit in the run-time-instrumentation controls is zero, the corresponding condition is not checked. If a BP reporting-control bit is one and the corresponding branch-prediction condition exists, and a reporting group is stored.
- [0080] The BPxn bit 722, when one, enables checking of branch-prediction information. Thus, if the sample branch is incorrectly predicted to be taken but is not taken, a

reporting group is stored.

- [0081] The BPxt bit 724, when one, enables checking of the branch-prediction information. Thus, if the sample branch is incorrectly predicted to be not taken but is taken, a reporting group is stored.
- [0082] The BPti bit 726, when one, enables checking of the branch-prediction information. Thus, if the sample branch is correctly predicted to be taken, and is taken, but the branch target is incorrectly predicted, a reporting group is stored.
- [0083] The BPni bit 728, when one, enables checking of the branch-prediction information. Thus, if the sample branch is correctly predicted to not be taken, and is not taken, and the branch target is incorrectly predicted, a reporting group is stored.
- [0084] The enablement control of transactional-execution-mode records bit 720 (X bit) controls the collection of transactional-execution-mode abort records. When the X bit 720 in the run-time-instrumentation controls is zero, transactional-execution-mode abort records are not collected. When the X bit 720 is one, transactional-execution mode abort records are collected and placed in the collection buffer 508 of FIG. 5. If a model does not have a transactional-execution facility installed, the X bit 720 is ignored.
- [0085] The RIEMIT instruction control bit 718 (E bit) controls the execution of the RIEMIT instruction. When the E bit 718 in the run-time-instrumentation controls is zero or ignored and treated as zero when run-time-instrumentation is enabled, RIEMIT executes a no-operation. When E bit 718 is one, and not otherwise ignored, RIEMIT is enabled to execute its defined function.
- [0086] The J bit 746 when zero, specifies that the branch on condition (BC) instruction is in the other-type branch category, regardless of mask value. If the J bit 746 is one, the BC instruction which specifies a mask of 15 is in the return-type branch category. When the BC instruction specifies a mask of 1- 14, it is not affected by the J bit 746 and is always in the other type branch category. When in the return-type branch category, the R bit 716 controls inclusion into the collection buffer 508 of FIG. 5. When in the other type branch category, the B bit 748 controls inclusion into the collection buffer 508. The other-type branch category may also be indicated as the transfer-type branch category.
- [0087] The instruction address code bit 714 (C bit) controls the enablement of call type branches. If the C bit 714 in the run-time-instrumentation controls is one and the instruction is a call-type branch, the collection buffer 508 is updated. If model-dependent detection of both call-type and return-type branches is combined, the C bit 714 operates on both types and the R bit 716 is not effective.
- [0088] The R bit 716 is the enablement control of return-type branches. If the R bit 716 in the run-time-instrumentation controls is one and the instruction is a return-type branch,

then the collection buffer 508 is updated.

- [0089] The B bit 748 is the enablement control of branches other than call-type and return-type branches. If the B bit 748 in the run-time-instrumentation controls is one and the instruction is an other-type branch recognized by run-time-instrumentation, then the collection buffer 508 is updated.
- [0090] The maximum-address exceeded bit 712 (MAE bit), if set to 1, indicates that, one or more reporting groups have been stored that have an instruction address code (C field) set to one. Once the MAE bit 712 is set to one, continuing execution of run-time-instrumentation does not set it back to zero. Execution of the LRIC instruction or the MRIC instruction which specifies the MAE bit 712 as zero will set the MAE bit 712 to zero.
- [0091] The run-time-instrumentation next (RINEXT) control bit 710 (N bit) controls the enablement of the run-time-instrumentation next instruction, which controls the execution of a sample instruction. When the N bit 710 in the run-time-instrumentation controls is zero or ignored and treated as zero, RINEXT executes a no-operation. When the N bit 710 is one, and not otherwise ignored, RINEXT is enabled to execute its defined function.
- [0092] The mode field 708 (M field) is a 4-bit unsigned integer whose value in the run-time-instrumentation controls specifies the sampling mode for the run-time-instrumentation controls. Supported sampling modes, may include sampling based on counting CPU cycles, counting instructions, or be directed to sample in response to a sample instruction, such as RINEXT.
- [0093] The reporting group size field 744 (RGS) is a 3-bit unsigned integer whose value specifies the number of records of a reporting group (R_{RG}). The number of records in a reporting group may vary from two records, including a begin/timestamp record and an instruction last record, up to two hundred fifty-six records. In an embodiment, the upper limit may be model dependent. The number of 16-byte records placed into a reporting group is $2^{(RGS+1)}$.
- [0094] The primary-CPU capability suppression control bit 738 (Y bit) and the secondary-CPU capability suppression control bit 739 (Z bit) are collectively referred to as the suppression control. Suppression of the storing of a reporting group means that an attempt to store is not performed. The suppression control is not effective and no suppression occurs when the CPU capability of all CPUs in the configuration is the same. In a configuration, if the CPU capability of a CPU differs from the capability of another CPU, the suppression control is in effect, and at least one CPU is said to be operating at the CPU capability or primary-CPU capability while at least one other CPU is said to be operating at the secondary-CPU capability. The primary and secondary CPU capabilities are different operating speeds. When Y bit 738 and Z bit

739 are both zero, suppression does not occur. When Y bit 738 is zero and Z bit 739 is one, suppression occurs if the CPU, e.g., processor 106, is operating at the secondary-CPU capability. When Y bit 738 is one and Z bit 739 is zero, suppression occurs if the CPU, e.g., processor 106, is operating at the primary-CPU capability. When Y bit 738 and Z bit 739 are both one, suppression occurs.

- [0095] The above fields and bits of FIG. 7 are an example of the placement and naming of the fields and are provided herein for purposes of clarity. It will be understood that in other embodiments the only a subset of the fields may be used, fields may be in any order or position, and/or may be signified by different names.
- [0096] When run-time instrumentation is installed and enabled, a number of events and data can be captured in collection buffer 508. The collection buffer 508 is used to capture a set of records whose contents report on events recognized by the processor 106 during program execution. Examples are: execution of one or more taken branches, transactional-execution abort events, cache-misses, and an operand of a run-time instrumentation emit instruction. The IC and DC controls fields 734 and 736 set a level at which the program would be interested in taking some corrective action to improve instruction or data pre-fetch behavior. Execution of the RIEMIT instruction collects the value of a general register by storing it into the collection buffer 508. Additional data can be collected and/or stored in other buffers, such as an instruction-data buffer (IDB) (not depicted) used to collect model-dependent sample-instruction data to construct a run-time-instrumentation instruction record.
- [0097] Collected run-time-instrumentation information is reported on a sampling basis. Instructions from the instruction stream are sampled. The instruction that is sampled is called the sample instruction. A number of modes for determining a sample instruction are defined as follows when run-time instrumentation is enabled. In cycle-count mode, a count is the number of CPU cycles specified in either SF 740 or RSIC 742, whichever is used to provide the count for the current interval. The count is adjusted responsive to an event associated with the sampling mode. For example, the count may be decremented when the processor 106 is in the operating state. When the count is decremented to threshold value, such as zero, the current instruction is recognized as a sample instruction, and the count is reinitialized to the SF 740 value and begins to be decremented with the next cycle. When execution of the sample instruction completes, reporting is performed, if appropriate.
- [0098] In instruction-count mode, a count is specified in either SF 740 or RSIC 742, whichever is used to provide the count for the current interval. For an instruction which consists of a single unit of operation, the count is decremented at the completion of the instruction as an event used to adjust the count. The instruction is a sample instruction when the count is decremented to a threshold value, such as zero. For an in-

struction which consists of multiple units-of-operation, the count may be decremented in one of the following ways:

- a. For an interruptible instruction, all units of operation through partial completion represent one counted unit for which the count is decremented.
- b. For an interruptible instruction, all units of operation since the most-recent partial completion through final completion represent one counted unit for which the count is decremented.
- c. For an instruction that completes after performing a CPU-determined subportion of the processing specified by the parameters of the instruction, the completion represents one counted unit for which the count is decremented.
- d. For an instruction that completes after performing multiple units of operation but not in categories a-c above, completion of the last unit of operation represents one counted unit for which the count is decremented.

An instruction is a sample instruction when the count is decremented to zero for any counted unit of the instruction. When a threshold value is reached, such as zero, the count is reinitialized to the SF 740 value and begins to count down as described in a-d above. In all cases of the count modes, reporting, if appropriate, occurs after completion of the last unit of operation of the sample instruction.

- [0099] In directed-sampling mode, directed sampling occurs when the N-bit 710 is one and the RINEXT instruction is executed successfully. The sample instruction is the next, sequential instruction (NSI) after the RINEXT instruction. If the next, sequential instruction is an execute-type instruction, the sample instruction is the target instruction of the execute-type instruction. Directed sampling may occur when in the cycle-count or instruction-count mode. Count sampling continues in conjunction with directed sampling and any of its resulting actions, and is not otherwise affected, except that if the sample instruction determined from count sampling is the same instruction determined by directed sampling, two reporting groups are not stored.
- [0100] Whatever the sampling mode is, when a sample instruction is identified by execution of the RINEXT instruction, a reporting group is stored. However, the run-time-instrumentation controls Y 738, Z 739, Qs 614, and Ps 612 continue to be effective.
- [0101] Cycle-count and instruction-count sampling each determine an approximate interval which is subject to an amount of variability based on internal system events and exception conditions. The countdown begins when run-time instrumentation transitions from disabled to enabled. Directed sampling is subject to a lesser amount of variability, depending on any event that can be interposed between completion of RINEXT and the NSI. Of note, an interruption can cause what was thought to be the NSI to no longer be the NSI.

- [0102] Sampling, regardless of the mode, identifies a sample instruction. Once a sample instruction is identified, collection stops upon completion of execution of the sample instruction and reporting begins. The various reporting controls that govern reporting then apply. Collection resumes when store of the reporting group is made pending.
- [0103] When not in the transactional-execution mode, store of a reporting group becomes pending upon completion of execution of a sample instruction. When in the transactional-execution mode, upon completion of execution of a sample instruction, store of a reporting group is deferred until the transaction ends and then becomes pending. When the store of a reporting group is deferred or pending, it may be purged if any of the following interruptions is recognized: 1) program interruption; 2) exigent machine-check interruption; 3) restart interruption; and 4) supervisor-call interruption.
- [0104] Any pending I/O, external, and repressible machine-check interruption remains pending until either the reporting group has been stored or the run-time-instrumentation controls determine that a reporting group is not to be stored.
- [0105] Each mode may or may not allow a different set of reporting controls. When the sampling mode is either instruction count or cycle count, but directed sampling is also used, it is possible for the same sample instruction to be identified by multiple sampling methods. When this occurs, and the reporting controls to be used differ according to the sampling mode, the reporting controls associated with directed sampling apply.
- [0106] Precise determination of an interval meant to sample a particular instruction is generally not feasible, due to asynchronous and unsolicited system events that may occur. Instead, the RINEXT instruction can be used to more-closely designate a sample instruction.
- [0107] When in cycle-count mode or instruction-count mode, the RINEXT instruction can be issued in too close a proximity to the sample instruction identified from instruction-count or cycle-count sampling. The contents of the associated reporting group are as if the sample instruction were identified as the NSI of the RINEXT instruction and not as if a cycle-count or instruction-count identification of the sample instruction applied.
- [0108] Execution of RINEXT may execute as a no-operation if any one or more of the following exception conditions is met:
1. Run-time-instrumentation controls are not valid.
 2. In the problem state, Ps 612 of the current run-time-instrumentation controls is zero, indicating that problem-state reporting is not permitted.
 3. In the supervisor state, Qs 614 of the current run-time-instrumentation controls is zero, indicating that supervisor-state reporting is not permitted.
 4. The N-bit 710 of the current run-time-instrumentation controls is zero, indicating that the RINEXT instruction itself is not permitted.

5. Storage is suppressed.
6. A field in the current PSW indicates that run-time instrumentation is disabled.
7. A model-dependent threshold would be exceeded. The number of times RINEXT has been issued in a period of time has exceeded a model-dependent limit.
8. A program-buffer-full condition exists.
9. A run-time-instrumentation-halted condition exists.
10. The next, sequential instruction is a start interpretive execution instruction.
11. The next, sequential instruction is a supervisor call instruction.

[0109] Turning to FIG. 8, an embodiment of collection buffer 508 is generally shown. As described previously, when run-time instrumentation is enabled during program execution, run-time-instrumentation data is collected within the processor 106. In an embodiment, the place where data is collected within the processor 106 is the collection buffer 508, and optionally an instruction-data buffer. In an embodiment, the collection buffer 508 is an internal buffer of the processor 106 that is used to save the most recent records collected. When a sample trigger point is detected, the records are copied from the collection buffer 508 into the program buffer 522 as part of a reporting group that is written to the program buffer 522. In an embodiment, the records are copied from the collection buffer in a non-destructive manner.

[0110] The collection buffer 508 may be referred to as a "hardware collection buffer" because the collection buffer 508 is located in the processor and in an embodiment implemented as an array of register pairs representing instruction address 802 and event metadata 804 of a given event. In an embodiment, the instruction-data buffer is also implemented by an array of register pairs. An example of an event is a taken branch for which the register pair may hold the instruction address of the branch, and the metadata may hold the target of the branch as well as information regarding the historic behavior of the branch. In an embodiment, the registers pairs are ordered and updated sequentially as events occur in the instruction stream. A counter is maintained to indicate the index of the most recently updated entry in the array. In an embodiment the collection buffer 508 is a circular buffer, and when the collection buffer 508 is full, the next event overwrites the first entry in the array, and sequential updating of the array's register pairs re-starts on subsequent events. As such, assuming an array CB[0] to CB[N-1] and a counter *i* indicating the latest updated index, the trace of events captured would be represented by the sequence CB[*i*], CB[*i*-1] ... CB[1], CB[0], CB[N-1], CB[N-2] ... CB[*i*+1]. In another embodiment, two pointers are used: a head pointer pointing to the oldest entry in the buffer, and a tail/current pointer pointing to the newest entry in the buffer.

[0111] Events that represent a state of the processor 106 at any given execution point are captured sequentially in the collection buffer 508. The collection buffer 508 is used to

capture a set of records whose contents report on events recognized by the processor 106 during program execution (e.g., execution of one or more taken branches, transactional-execution abort events, the operand of a RIEMIT instruction, etc.). In an embodiment the events recognized depend on the contents of the RICCB shown in FIG. 7. Entries in the embodiment of the collection buffer 508 shown in FIG. 8 include an event instruction address 802 and other relevant event metadata 804. Examples of event metadata 804 include, but are not limited to: the instruction address of a taken branch and its target including some information about the historic behavior of the branch; the instruction address of a RIEMIT instruction and a respective register value; and the address of a transaction abort instruction and a respective transaction recovery entry point.

- [0112] The embodiment of the collection buffer 508 shown in FIG. 8 is capable of storing up to thirty-two entries (i.e., information about thirty-two events), with each instruction address 802 specified by sixty-four bits (e.g., bits 0:63), and event metadata 804 by sixty-four bits (e.g., bits 64:127). The size of the collection buffer (R_{CB}) is a model dependent count, representing a number of records. In the embodiment of the collection buffer 508 shown in FIG. 8, the byte size of the collection buffer is a multiple of the sixteen byte record size. In an embodiment, the size of the collection buffer is a number of records greater than or equal to the difference between the count of the largest reporting group (R_{RG}) of the model and the count of the records in a reporting group that are not acquired from the collection buffer (R_{NC}). Thus, in an embodiment, the size of the collection buffer is expressed as:

$$R_{CB} \geq (R_{RG} - R_{NC}).$$

- [0113] In an embodiment, contents of the collection buffer 508 and the instruction data buffer (if one is used) are purged or otherwise affected by the following events: (1) an interruption; (2) the PSW bit that turns on and off the run-time instrumentation facility (e.g., bit 24) changes from a one to a zero; and (3) when a sample instruction is identified when the run-time instrumentation facility is in a transactional-execution mode (in this case, further update of the collection data buffer 508 and instruction-data buffer stops and resumes when the transaction ends, at which time, a store of the reporting group is pending and the collection buffer 508 and instruction-data buffers are purged).
- [0114] In an embodiment, such as the emulated host computer system shown in FIG. 1B, the collection buffer 508 is implemented using registers and/or memory. In this embodiment, the optional instruction-data buffer, if present, is also implemented using registers and/or memory.
- [0115] In embodiments, additional capabilities can affect data collection and may be viewed

as providing additional data-collection points while not substantially disturbing the regular instruction-count or cycle-count sampling described previously. These include execution of a RIEMIT instruction, which collects the value of a general register by storing it into the collection buffer 508. In addition, the data-collection control bits in the run-time instrumentation controls described previously can be used to customize the types of data collected (e.g., the E, C, R, and B control bits). In this manner, the type of data collected is programmable.

- [0116] In an embodiment, an instruction-data buffer is implemented to collect model dependent sample instruction data that is used to construct a run-time-instrumentation instruction record. The instruction-data buffer collects data from an instruction in anticipation of being available when the instruction is identified as a sample instruction. In an embodiment, the instruction-data buffer is a hardware buffer/storage location in the processor where information about an instruction that would become the trigger as a sample point is saved, so that during the log out process, it can be written out together with data from the collection buffer 508. Similar to the collection buffer 508 it includes the instruction address, and meta-data associated with that instruction. The metadata in the instruction-data buffer is often machine dependent and may include, but is not limited to: cache miss related information, and branch prediction related information.
- [0117] In accordance with embodiments, other data collected may not be from the collection buffer 508 and not from the instruction-data buffer. Examples include data used to form parts of the following: (1) the first record of a reporting group: timestamp or begin record; and (2) additional types of records may be created for every reporting group and thus not stored in the collection buffer 508, such records, when present, may be placed in the extra or machine-dependent section of a reporting group. These records are referred to herein as "system information records."
- [0118] FIG. 9 depicts a high-level example of a reporting group 900 stored to program buffer 522 at a sample point. The size of a reporting group in records is represented by R_{RG} , equals $2^{(RGS+1)}$, where RGS is the reporting group size as an exponent. A model-dependent number of records (R_{NC}) copied from a location other than the collection buffer 508 may or may not be copied non-destructively when used in a reporting group. In the example of FIG. 9, $R_{RG} = 8$, $R_{GS} = 2$, and $R_{NC} = 4$. The example reporting group 900 shown in FIG. 9 includes a header section 902, a body section 904, an extra records section 906, and a footer section 908.
- [0119] The header section 902 may include a begin record or a timestamp record to hold status, tracking, and/or timing information. A begin record is stored in the header section 902 for the first reporting group stored in a program buffer (i.e., when the RCA 706 is equal to the ROA 702). In an embodiment, the begin record includes a record

type field of "02", a number of reporting groups (NRG) field for indicating how many reporting groups are currently stored in the program buffer, a RGS field to indicate the size of the reporting groups, a stopped (S) field for indicating whether or not the program buffer 522 is full, a halted (H) field for indicating whether the run-time instrumentation is halted, and a time of day (TOD) clock field for indicating when the begin record was written. In an embodiment, at least a subset of the fields in the begin record are sourced from the RI control block (e.g., RICCB). An embodiment of the timestamp record has a record type of "03" and includes a TOD clock field for indicating when the record was stored. In an embodiment, a timestamp record is stored in the header section 902 for each reporting group other than the first reporting group.

- [0120] The body section 904 of the reporting group may include a variety of records for events and information sampled from collection buffer 508. Events and information may represent, for example, state information captured by an emit instruction, a transactional-execution abort, a call, a return, a branch, and filler.
- [0121] In an embodiment, an emit record is created and stored in the collection buffer 508 upon a successful execution of a RIEMIT instruction. An embodiment of the emit record includes a record type field of "10", an instruction address code field to indicate how the instruction address bit positions of the current PSW are represented in the emit record, an instruction address field which varies depending on the addressing mode (e.g., 64, 31 or 24 bit) and contains the instruction address of the RIEMIT instruction or execute type instruction if the RIEMIT was the target of an execute type instruction, and an emit data field for storing the data from the general register specified by the RIEMIT instruction.
- [0122] In an embodiment, a transactional execution mode abort record is created and stored in the collection buffer 508 by either an implicit abort or by execution of a transaction abort instruction. An embodiment of the abort record includes a record type field of "11", an instruction address code field to indicate how the instruction address bit positions of the current PSW are represented in the transactional-execution abort record, an instruction address field which varies depending on the addressing mode (e.g., 64, 31 or 24 bit) and contains the instruction address of the aborted instruction or execute type instruction if the aborted instruction was the target of an execute type instruction, and a field for any model dependent data associated with the abort.
- [0123] In an embodiment, a call record is created by execution of a call type branch instruction, such as: BRANCH AND SAVE (BASR) when the R2 field is nonzero, BRANCH AND SAVE (BAS), BRANCH RELATIVE AND SAVE LONG, BRANCH RELATIVE AND SAVE, BRANCH AND LINK (BALR) when the R2 field is nonzero, BRANCH AND LINK (BAL), and BRANCH AND SAVE AND SET MODE when the R2 field is nonzero. An embodiment of the call record includes a

record type field of "12", an instruction address code field to indicate how the instruction address bit positions of the current PSW are represented in the call record, an instruction address field which varies depending on the addressing mode (e.g., 64, 31 or 24 bit) and contains the address of the branch instruction or execute type instruction if the branch instruction was the target of an execute type instruction, and a well behaved field for indicating whether or not the branch was correctly predicted, and a target address field containing the branch target address (also referred to as the "called location").

- [0124] Return records and transfer records may have the same format as the call records. In an embodiment, a return record has a record type field of "13" and is created by execution of a return type branch instruction such as a BRANCH ON CONDITION (BCR) when the R2 field is nonzero and the mask is 15. For the return record, the instruction address field contains the address of the branch instruction or execute type instruction if the branch is the target of an execute type instruction, and the target address field contains the return location.
- [0125] In an embodiment, a transfer record has a record type field of "14" and is created by execution of a return type branch instruction such as: a. BRANCH ON CONDITION (BCR) when the R2 field is nonzero and the mask is in the range 1-14; b. BRANCH ON CONDITION (BC) when the J bit is zero or the mask is in the range 1-14; c. BRANCH ON COUNT (BCT, BCTR, BCTG, BCTGR); d. BRANCH ON INDEX HIGH (BXH, BXHG); e. BRANCH ON INDEX LOW OR EQUAL (BXLE, BXLEG); f. BRANCH RELATIVE ON CONDITION (BRC); g. BRANCH RELATIVE ON CONDITION LONG (BRCL); h. BRANCH RELATIVE ON COUNT (BRCT, BRCTG); i. BRANCH RELATIVE ON COUNT HIGH (BRCTH); j. BRANCH RELATIVE ON INDEX HIGH (BRXH, BRXHG); k. BRANCH RELATIVE ON INDEX LOW OR EQUAL (BRXLE, BRXLG); l. COMPARE AND BRANCH (CRB, CGRB); m. COMPARE AND BRANCH RELATIVE (CRJ, CGRJ); n. COMPARE IMMEDIATE AND BRANCH (CIB, CGIB); o. COMPARE IMMEDIATE AND BRANCH RELATIVE (CIJ, CGIJ); p. COMPARE LOGICAL AND BRANCH (CLRB, CLGRB); q. COMPARE LOGICAL AND BRANCH RELATIVE (CLRJ, CLGRJ); r. COMPARE LOGICAL IMMEDIATE AND BRANCH (CLIB, CLGIB); and s. COMPARE LOGICAL IMMEDIATE AND BRANCH RELATIVE (CLIJ, CLGIJ). The transfer record is created when the branch is taken. For the transfer record, the instruction address field contains the address of the branch instruction or execute type instruction if the branch is the target of an execute type instruction, and the target address field contains the return location.
- [0126] A filler record is used in a reporting group when the number of valid records in the collection buffer 508 is not sufficient to fill a reporting group of the current RGS. An

embodiment of a filler record includes record type field of "00" to indicate that the record is a filler record and the remaining bytes are undefined.

- [0127] The extra records section 906, when present, may contain model-dependent records. In an embodiment, the format of an extra record is similar to the filler record except for the record type is set to "01" to indicate that the record is an extra record and the remaining bytes of the extra record may contain model dependent data.
- [0128] The footer section 908 can include an instruction record containing information about execution of a sample instruction. An instruction record is created when a reporting group is stored for a sample instruction. An embodiment of the instruction record includes a record type field of "04", an instruction address code field to indicate how the instruction address bit positions of the current PSW are represented in the instruction record, an instruction address field which varies depending on the addressing mode (e.g., 64, 31 or 24 bit) and contains the instruction address of the sample instruction or execute type instruction if the sample instruction was the target of an execute type instruction, and an instruction-data buffer (IDB) field containing any model dependent data collected from the IDB.
- [0129] As described previously, the run-time-instrumentation function is a new facility that may be used in not only in a laboratory environment, or for off-line analysis, but also in live software environments within programs at runtime, and under program control. Initially, a privileged state may set controls of a processor 106 to manage run-time-instrumentation. In an embodiment, the set of controls are originally loaded by successful execution of the load run-time-instruction control (LRIC) instruction by the privileged state. The flexibility of the run-time-instrumentation facility is enhanced by providing the ability to turn the run-time instrumentation facility on and off from a lesser-privileged state. In this manner, the data collected by the run-time instrumentation facility can be targeted to particular instructions.
- [0130] In accordance with embodiments, both the RIOFF instruction and the RION instruction may be executed by an application executing in a lesser-privileged state. If the execution of the RIOFF instruction is successful, then the run-time instrumentation is disabled when execution of the RIOFF instruction successfully completes, otherwise, if the execution of the RIOFF instruction is unsuccessful, then the RIOFF instruction will have no impact on the status (e.g., enabled or disabled) of the run-time instrumentation facility. Similarly, if the execution of the RION instruction is successful, then the run-time instrumentation is enabled when execution of the RION instruction successfully completes, otherwise, if the execution of the RION instruction is unsuccessful, then the RION instruction will have no impact on the status (e.g., enabled or disabled) of the run-time instrumentation facility.
- [0131] FIG. 11 depicts an RIOFF instruction in accordance with an embodiment. As shown

in FIG. 11, the RIOFF instruction 1100 includes an operation code 1102 and 1104 (also referred to as "opcode" or "split opcode" in this particular case). The opcode 1102 and 1104 identifies the RIOFF instruction 1100 to the processor, such as processor 106 of FIG. 5. FIG. 11 depicts one embodiment of the RIOFF instruction, it will be understood by those of ordinary skill in the art that the RIOFF instruction may be formatted differently and/or contain different operands and opcodes in other embodiments.

[0132] FIG. 12 depicts a process flow of an RIOFF instruction in accordance with an embodiment. In an embodiment, the process flow of FIG. 12 is executed by the instrumentation module 506 of FIG. 5. At block 1202, an RIOFF instruction issued by a lesser-privileged state program (e.g. an application 518 executing in the problem state) is fetched (e.g., from an instruction stream). In an embodiment, a bit of the PSW is examined (e.g., bit 15) to determine whether the program is executing in a supervisory state or a lesser privileged state (also referred to herein as a "problem state"). At block 1204 it is determined if execution of the RIOFF instruction is permitted. The run-time-instrumentation S bit in the RICCB (controlled only by LRIC) determines if the lesser-privileged state program is allowed to execute the RIOFF instruction. In an embodiment, if the run-time-instrumentation S bit in the current run-time instrumentation controls (e.g., in the RICCB) is set to 1, then the lesser privileged state program is allowed to execute the RIOFF instruction and processing continues at block 1208. Alternatively, if the S bit is set to 0, then the lesser privileged state program is not allowed to execute the RIOFF instruction, and processing continues at block 1206. At block 1206, the RIOFF instruction does not change the prior settings, and the runtime controls and PSW remain at their prior values. In an embodiment, at block 1206, a condition code is set to two to indicate that lesser-privileged state execution is not enabled.

[0133] At block 1208, it is determined if the validity bit (also referred to as the V bit) of the current run-time-instrumentation controls (e.g., in the RICCB) is set to 1. The validity bit indicates the validity of the set of run-time instrumentation controls in the processor, as they were previously set by an LRIC instruction. If the current run-time instrumentation controls are not valid, (i.e. the previous LRIC instruction was invalid), then processing continues at block 1201. In an embodiment, at block 1210, a condition code is set to three to indicate that the current run-time instrumentation controls are not valid.

[0134] If the validity bit is set to 1, indicating that the run-time-instrumentation controls are valid, then processing continues at block 1212, where the decrementing of the current run-time RSIC is suppressed and the value preserved in anticipation of a subsequent RION instruction. Processing continues at block 1214, where the RI bit in the PSW

(e.g., bit 24) is set to zero. In an embodiment, if bit 24 of the PSW is already zero, then execution of RIOFF completes and a condition code is set to zero. In an embodiment, an EXTRACT PSW instruction is executed to determine a current value of PSW bit 24.

[0135] FIG. 13 depicts a RION instruction in accordance with an embodiment. As shown in FIG. 13, the RION instruction 1300 includes an operation code 1302 and 1304 (also referred to as "opcode" or "split opcode" in this particular case). The opcode 1302 and 1304 identifies the RION instruction 1300 to the processor, such as processor 106 of FIG. 5. FIG. 13 depicts one embodiment of the RION instruction, it will be understood by those of ordinary skill in the art that the RION instruction may be formatted differently and/or contain different operands and opcodes in other embodiments.

[0136] FIG. 14 depicts a process flow of an RION instruction in accordance with an embodiment. In an embodiment, the process flow of FIG. 14 is executed by the instrumentation module 506 of FIG. 5. At block 1402, an RION instruction issued by a lesser-privileged state program (e.g. an application 518 executing in the problem state) is fetched (e.g., from an instruction stream). In an embodiment, a bit of the PSW is examined (e.g., bit 15) to determine whether the program is executing in a supervisory state or a lesser privileged state (also referred to herein as a "problem state"). At block 1404 it is determined if execution of the RION instruction is permitted. The run-time-instrumentation S bit (controlled only by LRIC) determines if the lesser-privileged state program is allowed to execute the RION instruction. In an embodiment, if the run-time-instrumentation S bit in the current run-time instrumentation controls (e.g., in the RICCB) is set to 1, then the lesser privileged state program is allowed to execute the RION instruction and processing continues at block 1408. Alternatively, if the S bit is set to 0, then the lesser privileged state program is not allowed to execute the RION instruction, and processing continues at block 1406. In an embodiment, at block 1406, the RION instruction does not change the prior settings, and the runtime controls and PSW remain at their prior values. In an embodiment, at block 1406, a condition code is set to two to indicate that lesser-privileged state execution is not enabled.

[0137] At block 1408, it is determined if the validity bit (also referred to as the V bit) of the current run-time-instrumentation controls (e.g., in the RICCB) is set to 1. The validity bit indicates the validity of the set of run-time instrumentation controls in the processor, as they were previously set by an LRIC instruction. If the current run-time instrumentation controls are not valid, (i.e. the previous LRIC instruction was invalid), then processing continues at block 1410. In an embodiment, at block 1410, a condition code is set to three to indicate that the current run-time instrumentation controls are not valid.

- [0138] If the validity bit is set to 1, indicating that the run-time-instrumentation controls are valid, then processing continues at block 1412, where the contents of the collection buffer, and the instruction-data buffer if present, are purged. In an embodiment, the buffers are purged at block 1412 only when the RI bit in the PSW was previously set to zero. Processing continues at block 1414, where the RI bit in the PSW (e.g., bit 24) is set to one. In an embodiment, if bit 24 of the PSW is already one, then execution of the RION instruction completes and a condition code is set to zero. In an embodiment, an EXTRACT PSW instruction is executed to determine a current value of PSW bit 24.
- [0139] In an embodiment, if the run-time instrumentation facility is executing in sampling modes 0 or 1 (as determined, for example by the RCCB), then a non-zero RSIC is used to continue the remainder of the interval. If the RSIC is zero, then a new sampling interval is initialized to the value of the scaling factor and the A bit in the RICCB is set to one.
- [0140] As described above, embodiments can be embodied in the form of computer-implemented processes and apparatuses for practicing those processes. An embodiment may include a computer program product 1500 as depicted in FIG. 15 on a computer readable/usable medium 1502 with computer program code logic 1504 containing instructions embodied in tangible media as an article of manufacture. Exemplary articles of manufacture for computer readable/usable medium 1502 may include floppy diskettes, CD-ROMs, hard drives, universal serial bus (USB) flash drives, or any other computer-readable storage medium, wherein, when the computer program code logic 1504 is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. Embodiments include computer program code logic 1504, for example, whether stored in a storage medium, loaded into and/or executed by a computer, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the computer program code logic 1504 is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. When implemented on a general-purpose microprocessor, the computer program code logic 1504 segments configure the microprocessor to create specific logic circuits.
- [0141] Technical effects and benefits include the ability to limit sampling by a run-time instrumentation facility to a subset of the execution trace of an application. This may lead to a more focused data set from the instrumentation. In addition, by limiting the sampling to a subset of the execution trace, the cost of having to manage, buffer, and reduce the instrumentation data is decreased.
- [0142] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the

context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising, " when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

- [0143] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.
- [0144] As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit, " "module" or "system. " Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.
- [0145] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or

in connection with an instruction execution system, apparatus, or device.

- [0146] A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.
- [0147] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.
- [0148] Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).
- [0149] Aspects of the present invention are described above with reference to flowchart illustrations and/or schematic diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.
- [0150] These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the

computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

- [0151] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.
- [0152] As described above, embodiments can be embodied in the form of computer-implemented processes and apparatuses for practicing those processes. In embodiments, the invention is embodied in computer program code executed by one or more network elements. Embodiments include a computer program product on a computer usable medium with computer program code logic containing instructions embodied in tangible media as an article of manufacture. Exemplary articles of manufacture for computer usable medium may include floppy diskettes, CD-ROMs, hard drives, universal serial bus (USB) flash drives, or any other computer-readable storage medium, wherein, when the computer program code logic is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. Embodiments include computer program code logic, for example, whether stored in a storage medium, loaded into and/or executed by a computer, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the computer program code logic is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. When implemented on a general-purpose microprocessor, the computer program code logic segments configure the microprocessor to create specific logic circuits.
- [0153] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and com-

binations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

CLAIMS

What is claimed is:

1. A computer program product for enabling and disabling execution of a run-time instrumentation facility on a processor, the computer program product comprising:

a non-transitory storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising:

fetching an instruction of a currently executing thread in a multi-threaded environment for execution by the processor in a first state, the instruction one of a run-time instrumentation facility off instruction (RIOFF instruction) and a run-time instrumentation facility on instruction (RION instruction), the fetching by the processor;

based on determining, by the processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, executing the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor; and

enabling the run-time instrumentation facility based on the instruction being the RION instruction, the enabling including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

2. The computer program product of claim 1, wherein the method further comprises:

fetching, by the processor, a run-time instrumentation control block (RICCB) that includes a problem state execution control bit, which was previously set to a value by a program executing in a privileged state, wherein the determining that the run-time instrumentation facility permits execution of the instruction in the first state is based on the value of the problem state execution control bit, wherein the first state is different than the privileged state.

3. The computer program product of claim 1, wherein the method further comprises:

fetching, by the processor, a run-time instrumentation control block (RICCB) that includes a validity bit, which was previously set to a value by a program executing in a privileged state, wherein the determining that controls associated with the run-time instrumentation facility are valid is based on the value of the validity bit.

4. The computer program product of claim 1, wherein the method further comprises, based on the enabling:

capturing, by the processor, the run-time instrumentation data based on an instruction stream of instructions of an application program executing on the processor, the capturing comprising storing the run-time instrumentation data in a collection buffer of the processor;

detecting, by the processor, a run-time instrumentation sample point trigger; and

copying contents of the collection buffer into a program buffer as a reporting group based on the detecting the run-time instrumentation sample point trigger, the program buffer located in main storage in an address space that is accessible by the application program.

5. The computer program product of claim 4, wherein the method further comprises:

capturing, in the collection buffer of the processor, instruction addresses and metadata corresponding to events detected during the executing of the instruction stream of instructions.

6. The computer program product of claim 4, wherein:

the reporting group includes a predetermined number of one or more instrumentation records comprising the contents of the collection buffer and system information records.

7. The computer program product of claim 4, wherein:

the copying includes copying the reporting group into the program buffer starting at a current address of the program buffer, the program buffer stored at a program buffer

origin address specified by an instruction accessible control block that also specifies an address of a last byte in the program buffer and the current address in the program buffer.

8. The computer program product of claim 1, wherein the method further comprises:

saving a remaining sample interval count field (RSIC field) based on the disabling;
fetching, by the processor, the RION instruction; and
restoring the RSIC field.

9. A system, comprising:

a hardware computer processor configured to enable and disable execution of a run-time instrumentation facility on the hardware computer processor, and the hardware computer processor configured to:

fetch an instruction of a currently executing thread in a multi-threaded environment for execution by the hardware computer processor in a first state, the instruction one of a run-time instrumentation facility off instruction (RIOFF instruction) and a run-time instrumentation facility on instruction (RION instruction), the fetching by the hardware computer processor;

based on determining, by the hardware computer processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, execute the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the hardware computer processor to indicate that run-time instrumentation data should not be captured by the hardware computer processor; and

enable the run-time instrumentation facility based on the instruction being the RION instruction, including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the hardware computer processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for

the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

10. The system of claim 9, wherein the hardware computer processor is also configured to:

fetching, by the hardware computer processor, a run-time instrumentation control block (RICCB) that includes a problem state execution control bit, which was previously set to a value by a program executing in a privileged state, wherein the determining that the run-time instrumentation facility permits execution of the instruction in the first state is based on the value of the problem state execution control bit, wherein the first state is different than the privileged state.

11. The system of claim 9, wherein the hardware computer processor is also configured to:

fetching, by the hardware computer processor, a run-time instrumentation control block (RICCB) that includes a validity bit, which was previously set to a value by a program executing in a privileged state, wherein the determining that controls associated with the run-time instrumentation facility are valid is based on the value of the validity bit.

12. The system of claim 9, wherein the hardware computer processor is also configured to, based on the enabling:

capturing, by the hardware computer processor, the run-time instrumentation data based on an instruction stream of instructions of an application program executing on the hardware computer processor, the capturing comprising storing the run-time instrumentation data in a collection buffer of the hardware computer processor;

detecting, by the hardware computer processor, a run-time instrumentation sample point trigger; and

copying contents of the collection buffer into a program buffer as a reporting group based on the detecting the run-time instrumentation sample point trigger, the program buffer located in main storage in an address space that is accessible by the application program.

13. A method for enabling and disabling execution of a run-time instrumentation facility on a processor, the method comprising:

fetching, by the processor, an instruction of a currently executing thread in a multi-threaded environment for execution by the processor in a first state, the instruction one of a run-time instrumentation facility off instruction (RIOFF instruction) and a run-time instrumentation facility on instruction (RION instruction);

based on determining, by the processor, that the run-time instrumentation facility permits execution of the instruction in the first state and that controls associated with the run-time instrumentation facility are valid, executing the instruction, the executing comprising any one of: disabling the run-time instrumentation facility based on the instruction being the RIOFF instruction, the disabling including updating a run-time instrumentation facility state bit in a program status word (PSW) of the processor to indicate that run-time instrumentation data should not be captured by the processor; and

enabling the run-time instrumentation facility based on the instruction being the RION instruction, the enabling including updating the run-time instrumentation facility state bit in the PSW to indicate that the run-time instrumentation data should be captured by the processor, wherein the PSW is associated with the currently executing thread, and a state of the run-time instrumentation facility for the currently executing thread is maintained in the PSW across dispatches of the currently executing thread.

14. The method of claim 13, further comprising:

fetching, by the processor, a run-time instrumentation control block (RICCB) that includes a problem state execution control bit, which was previously set to a value by a program executing in a privileged state, wherein the determining that the run-time instrumentation facility permits execution of the instruction in the first state is based on the value of the problem state execution control bit, wherein the first state is different than a privilege state.

15. The method of claim 13, further comprising:

fetching, by the processor, a run-time instrumentation control block (RICCB) that includes a validity bit, which was previously set to a value by a program executing in a

privileged state, wherein the determining that controls associated with the run-time instrumentation facility are valid is based on the value of the validity bit.

16. The method of claim 13, further comprising, based on the enabling:

capturing, by the processor, the run-time instrumentation data based on an instruction stream of instructions of an application program executing on the processor, the capturing comprising storing the run-time instrumentation data in a collection buffer of the processor;

detecting, by the processor, a run-time instrumentation sample point trigger; and

copying contents of the collection buffer into a program buffer as a reporting group based on the detecting the run-time instrumentation sample point trigger, the program buffer located in main storage in an address space that is accessible by the application program.

17. The method of claim 16, further comprising:

capturing, in the collection buffer of the processor, instruction addresses and metadata corresponding to events detected during the executing of the instruction stream of instructions.

18. The method of claim 16, wherein:

the reporting group includes a predetermined number of one or more instrumentation records comprising the contents of the collection buffer and system information records.

19. The method of claim 16, wherein:

the copying includes copying the reporting group into the program buffer starting at a current address of the program buffer, the program buffer stored at a program buffer origin address specified by an instruction accessible control block that also specifies an address of a last byte in the program buffer and the current address in the program buffer.

20. The method of claim 13, further comprising:

saving a remaining sample interval count field (RSIC field) based on the disabling;
fetching, by the processor, the RION instruction; and
restoring the RSIC field.

[Fig. 1A]




FIG. 1A

[Fig. 1B]




FIG. 1B

[Fig. 1C]




FIG. 1C

[Fig. 2]




FIG. 2

[Fig. 3]





FIG. 3

[Fig. 4A]



[Fig. 4B]





FIG. 4B

[Fig. 4C]



[Fig. 5]


500

FIG. 5

[Fig. 6]


600

FIG. 6

[Fig. 7]




FIG. 7

[Fig. 8]




FIG. 8

[Fig. 9]

900

Record No.	Record Type(s)	Section
0	Begin, Timestamp	Header 902
1	Emit, TX Abort, Call, Return, Branch, Filler	Body 904
2		
3		
$R_{RG}-R_{NG} (=4)$		
$R_{RG}-R_{NG}+1 (=5)$	Extra, Model-Dependent	Extra 906
$R_{RG}-2 (=6)$		
$R_{RG}-1$	Instruction	Footer 908

FIG. 9

[Fig. 10]




FIG. 10

[Fig. 11]




FIG. 11

[Fig. 12]




FIG. 12

[Fig. 13]




FIG. 13

[Fig. 14]




FIG. 14

[Fig. 15]




FIG. 15

