



(51) International Patent Classification:
G06F 9/30 (2006.01)

(21) International Application Number:
PCT/US2017/039577

(22) International Filing Date:
27 June 2017 (27.06.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/201,403 02 July 2016 (02.07.2016) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054-1549 (US).

(72) Inventors: **MEKKAT, Vineeth**; 305 Elan Village Lane, Apt. 117, San Jose, California 95134 (US). **LEE, Jang-haeng**; 320 Crescent Village Circle, Apt. 1371, San Jose, California 95134 (US). **WU, Youfeng**; 3658 Bryant Street, Palo Alto, California 94306 (US).

(74) Agent: **CRANDALL, Sean C.**; Patent Capital Group, 2816 Lago Vista Lane, Rockwall, Texas 75032 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

(54) Title: REGISTER RECLAMATION




FIG. 1A

(57) Abstract: In an example, there is disclosed an apparatus, including a binary translator (BT) including circuitry to: analyze a code block; determine that an architectural register mapped to a physical register in the physical register file is available for early reclamation; and insert a reclamation hint into the code block. There is also disclosed a processor to reclaim the physical register based at least in part on the reclamation hint.

Declarations under Rule 4.17:

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

Published:

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

REGISTER RECLAMATION

FIELD OF THE SPECIFICATION

[0001] This disclosure relates in general to the field of semiconductor devices, and more particularly, though not exclusively to, a system and method for register reclamation.

BACKGROUND

[0002] Multiprocessor systems are becoming more and more common. Applications of multiprocessor systems include dynamic domain partitioning all the way down to desktop computing. In order to take advantage of multiprocessor systems, code to be executed may be separated into multiple threads for execution by various processing entities. Each thread may be executed in parallel with one another. To increase the utility of a processing entity, out-of-order execute may be employed. Out-of-order execution may execute instructions as input to such instructions is made available. Thus, an instruction that appears later in a code sequence may be executed before an instruction appearing earlier in a code sequence.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The present disclosure is best understood from the following detailed description when read with the accompanying figures. It is emphasized that, in accordance with the standard practice in the industry, various features are not necessarily drawn to scale, and are used for illustration purposes only. Where a scale is shown, explicitly or implicitly, it provides only one illustrative example. In other embodiments, the dimensions of the various features may be arbitrarily increased or reduced for clarity of discussion.

[0004] FIGURE 1A is a block diagram of an exemplary computer system formed with a processor that may include execution units to execute an instruction, in accordance with embodiments of the present disclosure.

[0005] FIGURE 1B illustrates a data processing system, in accordance with embodiments of the present disclosure.

[0006] FIGURE 1C illustrates other embodiments of a data processing system for performing text string comparison operations.

[0007] FIGURE 2 is a block diagram of the micro-architecture for a processor that may include logic circuits to perform instructions, in accordance with embodiments of the present disclosure.

[0008] FIGURE 3A illustrates various packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure.

[0009] FIGURE 3B illustrates possible in-register data storage formats, in accordance with embodiments of the present disclosure.

[0010] FIGURE 3C illustrates various signed and unsigned packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure.

[0011] FIGURE 3D illustrates an embodiment of an operation encoding format.

[0012] FIGURE 3E illustrates another possible operation encoding format having forty or more bits, in accordance with embodiments of the present disclosure.

[0013] FIGURE 3F illustrates yet another possible operation encoding format, in accordance with embodiments of the present disclosure.

[0014] FIGURE 4A is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline, in accordance with embodiments of the present disclosure.

[0015] FIGURE 4B is a block diagram illustrating an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor, in accordance with embodiments of the present disclosure.

[0016] FIGURE 5A is a block diagram of a processor, in accordance with embodiments of the present disclosure.

[0017] FIGURE 5B is a block diagram of an example implementation of a core, in accordance with embodiments of the present disclosure;

[0018] FIGURE 6 is a block diagram of a system, in accordance with embodiments of the present disclosure.

[0019] FIGURE 7 is a block diagram of a second system, in accordance with embodiments of the present disclosure.

[0020] FIGURE 8 is a block diagram of a third system in accordance with embodiments of the present disclosure.

[0021] FIGURE 9 is a block diagram of a system-on-a-chip, in accordance with embodiments of the present disclosure.

[0022] FIGURE 10 illustrates a processor containing a central processing unit and a graphics processing unit which may perform at least one instruction, in accordance with embodiments of the present disclosure.

[0023] FIGURE 11 is a block diagram illustrating the development of IP cores, in accordance with embodiments of the present disclosure.

[0024] FIGURE 12 illustrates how an instruction of a first type may be emulated by a processor of a different type, in accordance with embodiments of the present disclosure.

[0025] FIGURE 13 illustrates a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set, in accordance with embodiments of the present disclosure.

[0026] FIGURE 14 is a block diagram of an instruction set architecture of a processor, in accordance with embodiments of the present disclosure.

[0027] FIGURE 15 is a more detailed block diagram of an instruction set architecture of a processor, in accordance with embodiments of the present disclosure.

[0028] FIGURE 16 is a block diagram of an execution pipeline for an instruction set architecture of a processor, in accordance with embodiments of the present disclosure.

[0029] FIGURE 17 is a block diagram of an electronic device for utilizing a processor, in accordance with embodiments of the present disclosure.

[0030] FIGURE 18 is an illustration of a system 1800 with a binary translator (BT), in accordance with embodiments of the present disclosure.

[0031] FIGURE 19 illustrates a logical representation of elements of a BT system with total store elimination, in accordance with embodiments of the present disclosure.

[0032] FIGURE 20 is a block diagram illustration of register mapping according to embodiments of the present disclosure.

[0033] FIGURE 21 is a block diagram of a program's life cycle in accordance with embodiments of the present disclosure.

[0034] FIGURE 22 is a flow chart of an example method of optimizing a program according to embodiments of the present disclosure.

[0035] FIGURE 23 is a flow chart of an example method performed by a processor according to embodiments of the present disclosure.

SUMMARY

[0036] In an example, there is disclosed an apparatus, including a binary translator (BT) including circuitry to: analyze a code block; determine that an architectural register mapped to a physical register in the physical register file is available for early reclamation; and insert a reclamation hint into the code block. There is also disclosed a processor to reclaim the physical register based at least in part on the reclamation hint.

EMBODIMENTS OF THE DISCLOSURE

[0037] The following disclosure provides many different embodiments, or examples, for implementing different features of the present disclosure. Specific examples of components and arrangements are described below to simplify the present disclosure. These are, of course, merely examples and are not intended to be limiting. Further, the present disclosure may repeat reference numerals and/or letters in the various examples. This repetition is for the purpose of simplicity and clarity and does not in itself dictate a relationship between the various embodiments and/or configurations discussed. Different embodiments may have different advantages, and no particular advantage is necessarily required of any embodiment.

[0038] Some modern processors execute instructions out-of-order by removing anti/output dependencies between instructions. Those dependencies can be eliminated by mapping architectural registers to physical registers. In other words, the processor may provide a number of user-addressable architectural registers ($R_0 \dots R_N$) in a “register file.” But the processor may include a greater number of physical registers ($P_0 \dots P_K$) (e.g., $K > N$). When a program accesses an architectural register (for example, EAX), the architectural register may be mapped to a physical register (for example, P_6) when an instruction writes the result to the architectural register. Because one architectural register can be renamed to several physical registers, the number of physical registers is larger than that of architectural registers. However, it is desirable to limit the number of physical registers as much as possible, as a larger physical register bank consumes more power.

[0039] Because the number of physical registers are finite, they may need to be “reclaimed” so that they can be used again. Reclamation includes marking the physical register as available for mapping to a new architectural register. For example, instruction 1 writes to register R_1 . This is mapped to P_7 . Later, instruction 100 writes to R_1 . Now R_1 may be mapped to a different physical register, such as P_8 . Once instruction 100 retires, P_7 may be reclaimed.

[0040] Table 1 illustrates a register reclamation according to embodiments of this specification. Architectural register R_2 , is renamed to P_2 for the first instruction, and P_2 is reclaimed when the tenth instruction retires because the instruction updates the same architectural register R_2 .

Op #	Before Rename	After Rename	Comment
1	$R_2 = \text{LD } [R_1]$	$P_2 = \text{LD } [P_1]$	R_2 renamed to P_2
2	$R_3 = \text{ADD } R_2, R_0$	$P_3 = \text{ADD } P_2, P_0$	Last use of R_2
...	...	...	
10	$R_5 = \text{MUL } R_0, R_3$	$P_5 = \text{MUL } P_0, P_3$	P_2 is reclaimed after this instruction retires

Table 1. Register Reclamation

[0041] In this example, hardware does not reclaim P_2 until the instruction 10 retires, thus ensuring that all user instructions that need R_2 (mapped to P_2) have finished their executions.

[0042] But P_2 could have been safely reclaimed after instruction 2 retires, as this is the last instruction that uses R_2 . Thus, it is advantageous to provide a facility, such as a binary translator (BT) that is configured to “look ahead” and determine that a register is no longer needed, and thus allow hardware to reclaim the physical register it is mapped to. Early reclamation of physical registers means that a design may require fewer physical registers for the same performance gain. This in turn may reduce power consumed by the processor.

[0043] In some examples, hardware alone may have difficulty determining whether a register is still “live” (should not be reclaimed), or “dead” (may be reclaimed). In some cases, hardware alone could reclaim temporary registers used in a small set of instructions, but this may not realize the performance gains of a more comprehensive register reclamation regime. Such hardware-only reclamation may also require complex hardware support that could in fact nullify (or even exacerbate) the power consumption savings of early reclamation.

Thus, certain embodiments of this specification provide for block-level hints from the BT in a BT-based processor, such as Intel® processors from the Pentium® family and newer. This provides hardware with register liveness information on a block-level granularity so that dead registers may be reclaimed earlier, thus providing for power savings that may be realized for example through gating.

[0044] In an example, the BT annotates the beginning of each basic block (or a group of small basic blocks) with information about which registers are dead and thus may be reclaimed. This may be instantiated in different ways. In one example, the information is provided to hardware in the form of a new instruction. In another example, the information may be provided by existing or new BT overhead instructions. Hardware can then use this information to reclaim the indicated physical registers to improve energy efficiency through power gating.

[0045] Advantageously, a BT may be able to identify dead registers far in advance of some hardware solutions, and may do so more efficiently in terms of power consumption and processor complexity. In particular, the BT has visibility over a much larger code window than bare hardware, with a view that is essentially the entire checkpoint-to-checkpoint translation. Further advantageously, the BT's liveness analysis may also blend into existing register allocation analyses without introducing additional compile or runtime overhead.

[0046] In an example, the BT provides "reclamation hints" to indicate registers that the BT has determined are dead and may thus be reclaimed. In some embodiments, reclamation hints may be provided dynamically at the instruction level. Hints may also be provided by a bit, flag, or register with "dead hints" for the registers used. In some examples, this may lead to earlier reclamation than block-level hints in a linear code flow. However, this may cause additional delay in register reclamation in a code with control flow. Such flag-based hints may also require changes to the instruction set architecture, adding to each instruction two or three bits for dead hints for each registers used in the instruction.

[0047] In some examples, static reclamation hints may also be provided by a compiler. Compared to static hints, dynamic hints may, however, better handle reclamation across branching operations. As used throughout this specification, a "branching operation" expressly includes all operations that non-linearly change program flow, including by way of nonlimiting example, "branch," "branch" (with conditions, such as equal, not equal, greater

than, less than, overflow, negative, zero, not zero, borrow, or not borrow by way of nonlimiting example) “jump,” “jump” (with conditions), “loop,” “goto,” or any equivalents thereof.

[0048] A system and method for register reclamation will now be described with more particular reference to the attached FIGURES. It should be noted that throughout the FIGURES, certain reference numerals may be repeated to indicate that a particular device or block is wholly or substantially consistent across the FIGURES. This is not, however, intended to imply any particular relationship between the various embodiments disclosed. In certain examples, a genus of elements may be referred to by a particular reference numeral (“widget 10”), while individual species or examples of the genus may be referred to by a hyphenated numeral (“first specific widget 10-1” and “second specific widget 10-2”).

[0049] Although the below examples describe instruction handling and distribution in the context of execution units and logic circuits, other embodiments of the present disclosure may be accomplished by way of a data or instructions stored on a machine-readable, tangible medium, which when performed by a machine cause the machine to perform functions consistent with at least one embodiment of the disclosure. In one embodiment, functions associated with embodiments of the present disclosure are embodied in machine-executable instructions. The instructions may be used to cause a general-purpose or special-purpose processor that may be programmed with the instructions to perform the steps of the present disclosure. Embodiments of the present disclosure may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform one or more operations according to embodiments of the present disclosure. Furthermore, steps of embodiments of the present disclosure might be performed by specific hardware components that contain fixed-function logic for performing the steps, or by any combination of programmed computer components and fixed-function hardware components.

[0050] Instructions used to program logic to perform embodiments of the present disclosure may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions may be distributed via a network or by way of other computer-readable media. Thus a machine-readable medium may include

any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium may include any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

[0051] A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as may be useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, designs, at some stage, may reach a level of data representing the physical placement of various devices in the hardware model. In cases wherein some semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. A memory or a magnetic or optical storage such as a disc may be the machine-readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or retransmission of the electrical signal is performed, a new copy may be made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present disclosure.

[0052] In modern processors, a number of different execution units may be used to process and execute a variety of code and instructions. Some instructions may be quicker to

complete while others may take a number of clock cycles to complete. The faster the throughput of instructions, the better the overall performance of the processor. Thus it would be advantageous to have as many instructions execute as fast as possible. However, there may be certain instructions that have greater complexity and require more in terms of execution time and processor resources, such as floating point instructions, load/store operations, data moves, etc.

[0053] As more computer systems are used in internet, text, and multimedia applications, additional processor support has been introduced over time. In one embodiment, an instruction set may be associated with one or more computer architectures, including data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O).

[0054] In one embodiment, the instruction set architecture (ISA) may be implemented by one or more micro-architectures, which may include processor logic and circuits used to implement one or more instruction sets. Accordingly, processors with different micro-architectures may share at least a portion of a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale CA implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. Similarly, processors designed by other processor development companies, such as ARM Holdings, Ltd., MIPS, or their licensees or adopters, may share at least a portion a common instruction set, but may include different processor designs. For example, the same register architecture of the ISA may be implemented in different ways in different micro-architectures using new or well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB) and a retirement register file. In one embodiment, registers may include one or more registers, register architectures, register files, or other register sets that may or may not be addressable by a software programmer.

[0055] An instruction may include one or more instruction formats. In one embodiment, an instruction format may indicate various fields (number of bits, location of bits, etc.) to specify, among other things, the operation to be performed and the operands on

which that operation will be performed. In a further embodiment, some instruction formats may be further defined by instruction templates (or sub-formats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields and/or defined to have a given field interpreted differently. In one embodiment, an instruction may be expressed using an instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and specifies or indicates the operation and the operands upon which the operation will operate.

[0056] Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) may require the same operation to be performed on a large number of data items. In one embodiment, Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data elements. SIMD technology may be used in processors that may logically divide the bits in a register into a number of fixed-sized or variable-sized data elements, each of which represents a separate value. For example, in one embodiment, the bits in a 64-bit register may be organized as a source operand containing four separate 16-bit data elements, each of which represents a separate 16-bit value. This type of data may be referred to as 'packed' data type or 'vector' data type, and operands of this data type may be referred to as packed data operands or vector operands. In one embodiment, a packed data item or vector may be a sequence of packed data elements stored within a single register, and a packed data operand or a vector operand may be a source or destination operand of a SIMD instruction (or 'packed data instruction' or a 'vector instruction'). In one embodiment, a SIMD instruction specifies a single vector operation to be performed on two source vector operands to generate a destination vector operand (also referred to as a result vector operand) of the same or different size, with the same or different number of data elements, and in the same or different data element order.

[0057] SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, ARM processors, such as the ARM Cortex® family of processors having an instruction set including the Vector Floating Point (VFP) and/or NEON

instructions, and MIPS processors, such as the Loongson family of processors developed by the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences, has enabled a significant improvement in application performance (Core™ and MMX™ are registered trademarks or trademarks of Intel Corporation of Santa Clara, Calif.).

[0058] In one embodiment, destination and source registers/data may be generic terms to represent the source and destination of the corresponding data or operation. In some embodiments, they may be implemented by registers, memory, or other storage areas having other names or functions than those depicted. For example, in one embodiment, “DEST1” may be a temporary storage register or other storage area, whereas “SRC1” and “SRC2” may be a first and second source storage register or other storage area, and so forth. In other embodiments, two or more of the SRC and DEST storage areas may correspond to different data storage elements within the same storage area (e.g., a SIMD register). In one embodiment, one of the source registers may also act as a destination register by, for example, writing back the result of an operation performed on the first and second source data to one of the two source registers serving as a destination registers.

[0059] FIGURE 1A is a block diagram of an exemplary computer system formed with a processor that may include execution units to execute an instruction, in accordance with embodiments of the present disclosure. System 100 may include a component, such as a processor 102 to employ execution units including circuits with logic to perform algorithms for process data, in accordance with the present disclosure, such as in the embodiment described herein. System 100 may be representative of processing systems based on the PENTIUM® III, PENTIUM® 4, Xeon™, Itanium®, XScale™ and/or StrongARM™ microprocessors available from Intel Corporation of Santa Clara, California, although other systems (including PCs having other microprocessors, engineering workstations, set-top boxes and the like) may also be used. In one embodiment, sample system 100 may execute a version of the WINDOWS™ operating system available from Microsoft Corporation of Redmond, Washington, although other operating systems (UNIX and Linux for example), embedded software, and/or graphical user interfaces, may also be used. Thus, embodiments of the present disclosure are not limited to any specific combination of hardware circuitry, programmable circuitry, and software.

[0060] Embodiments are not limited to computer systems. Embodiments of the present disclosure may be used in other devices such as handheld devices and embedded applications. Some examples of handheld devices include cellular phones, Internet Protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications may include a micro controller, a digital signal processor (DSP), system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that may perform one or more instructions in accordance with at least one embodiment.

[0061] Computer system 100 may include a processor 102 that may include one or more execution units 108 to perform an algorithm to perform at least one instruction in accordance with one embodiment of the present disclosure. One embodiment may be described in the context of a single processor desktop or server system, but other embodiments may be included in a multiprocessor system. System 100 may be an example of a 'hub' system architecture. System 100 may include a processor 102 for processing data signals. Processor 102 may include a complex instruction set computer (CISC) microprocessor, a reduced instruction set computing (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, a processor implementing a combination of instruction sets, or any other processor device, such as a digital signal processor, for example. In one embodiment, processor 102 may be coupled to a processor bus 110 that may transmit data signals between processor 102 and other components in system 100. The elements of system 100 may perform conventional functions that are well known to those familiar with the art.

[0062] In one embodiment, processor 102 may include a Level 1 (L1) internal cache memory 104. Depending on the architecture, the processor 102 may have a single internal cache or multiple levels of internal cache. In another embodiment, the cache memory may reside external to processor 102. Other embodiments may also include a combination of both internal and external caches depending on the particular implementation and needs. Register file 106 may store different types of data in various registers including integer registers, floating point registers, status registers, and instruction pointer register.

[0063] Execution unit 108, including circuits with logic to perform integer and floating point operations, also resides in processor 102. Processor 102 may also include a microcode (ucode) ROM that stores microcode for certain macroinstructions. In one

embodiment, execution unit 108 may include circuits with logic to handle a packed instruction set 109. By including the packed instruction set 109 in the instruction set of a general-purpose processor 102, along with associated circuitry to execute the instructions, the operations used by many multimedia applications may be performed using packed data in a general-purpose processor 102. Thus, many multimedia applications may be accelerated and executed more efficiently by using the full width of a processor's data bus for performing operations on packed data. This may eliminate the need to transfer smaller units of data across the processor's data bus to perform one or more operations one data element at a time.

[0064] Embodiments of an execution unit 108 may also be used in micro controllers, embedded processors, graphics devices, DSPs, and other types of logic circuits. System 100 may include a memory 120. Memory 120 may be implemented as a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, or other memory device. Memory 120 may store instructions and/or data represented by data signals that may be executed by processor 102.

[0065] A system logic chip 116 may be coupled to processor bus 110 and memory 120. System logic chip 116 may include a memory controller hub (MCH). Processor 102 may communicate with MCH 116 via a processor bus 110. MCH 116 may provide a high bandwidth memory path 118 to memory 120 for instruction and data storage and for storage of graphics commands, data and textures. MCH 116 may direct data signals between processor 102, memory 120, and other components in system 100 and to bridge the data signals between processor bus 110, memory 120, and system I/O 122. In some embodiments, the system logic chip 116 may provide a graphics port for coupling to a graphics controller 112. MCH 116 may be coupled to memory 120 through a memory interface 118. Graphics card 112 may be coupled to MCH 116 through an Accelerated Graphics Port (AGP) interconnect 114.

[0066] System 100 may use a proprietary hub interface bus 122 to couple MCH 116 to I/O controller hub (ICH) 130. In one embodiment, ICH 130 may provide direct connections to some I/O devices via a local I/O bus. The local I/O bus may include a high-speed I/O bus for connecting peripherals to memory 120, chipset, and processor 102. Examples may include the audio controller, firmware hub (flash BIOS) 128, wireless transceiver 126, data storage 124, legacy I/O controller containing user input and keyboard interfaces, a serial expansion port

such as Universal Serial Bus (USB), and a network controller 134. Data storage device 124 may comprise a hard disk drive, a floppy disk drive, a CD-ROM device, a flash memory device, or other mass storage device.

[0067] For another embodiment of a system, an instruction in accordance with one embodiment may be used with a system on a chip. One embodiment of a system on a chip comprises of a processor and a memory. The memory for one such system may include a flash memory. The flash memory may be located on the same die as the processor and other system components. Additionally, other logic blocks such as a memory controller or graphics controller may also be located on a system on a chip.

[0068] FIGURE 1B illustrates a data processing system 140 which implements the principles of embodiments of the present disclosure. It will be readily appreciated by one of skill in the art that the embodiments described herein may operate with alternative processing systems without departure from the scope of embodiments of the disclosure.

[0069] Computer system 140 comprises a processing core 159 for performing at least one instruction in accordance with one embodiment. In one embodiment, processing core 159 represents a processing unit of any type of architecture, including but not limited to a CISC, a RISC or a VLIW type architecture. Processing core 159 may also be suitable for manufacture in one or more process technologies and by being represented on a machine-readable media in sufficient detail, may be suitable to facilitate said manufacture.

[0070] Processing core 159 comprises an execution unit 142, a set of register files 145, and a decoder 144. Processing core 159 may also include additional circuitry (not shown) which may be unnecessary to the understanding of embodiments of the present disclosure. Execution unit 142 may execute instructions received by processing core 159. In addition to performing typical processor instructions, execution unit 142 may perform instructions in packed instruction set 143 for performing operations on packed data formats. Packed instruction set 143 may include instructions for performing embodiments of the disclosure and other packed instructions. Execution unit 142 may be coupled to register file 145 by an internal bus. Register file 145 may represent a storage area on processing core 159 for storing information, including data. As previously mentioned, it is understood that the storage area may store the packed data might not be critical. Execution unit 142 may be coupled to decoder 144. Decoder 144 may decode instructions received by processing core 159 into

control signals and/or microcode entry points. In response to these control signals and/or microcode entry points, execution unit 142 performs the appropriate operations. In one embodiment, the decoder may interpret the opcode of the instruction, which will indicate what operation should be performed on the corresponding data indicated within the instruction.

[0071] Processing core 159 may be coupled with bus 141 for communicating with various other system devices, which may include but are not limited to, for example, synchronous dynamic random access memory (SDRAM) control 146, static random access memory (SRAM) control 147, burst flash memory interface 148, personal computer memory card international association (PCMCIA)/compact flash (CF) card control 149, liquid crystal display (LCD) control 150, direct memory access (DMA) controller 151, and alternative bus master interface 152. In one embodiment, data processing system 140 may also comprise an I/O bridge 154 for communicating with various I/O devices via an I/O bus 153. Such I/O devices may include but are not limited to, for example, universal asynchronous receiver/transmitter (UART) 155, universal serial bus (USB) 156, Bluetooth wireless UART 157 and I/O expansion interface 158.

[0072] One embodiment of data processing system 140 provides for mobile, network and/or wireless communications and a processing core 159 that may perform SIMD operations including a text string comparison operation. Processing core 159 may be programmed with various audio, video, imaging and communications algorithms including discrete transformations such as a Walsh-Hadamard transform, a fast Fourier transform (FFT), a discrete cosine transform (DCT), and their respective inverse transforms; compression/decompression techniques such as color space transformation, video encode motion estimation or video decode motion compensation; and modulation/demodulation (MODEM) functions such as pulse coded modulation (PCM).

[0073] FIGURE 1C illustrates other embodiments of a data processing system that performs SIMD text string comparison operations. In one embodiment, data processing system 160 may include a main processor 166, a SIMD coprocessor 161, a cache memory 167, and an input/output system 168. Input/output system 168 may optionally be coupled to a wireless interface 169. SIMD coprocessor 161 may perform operations including instructions in accordance with one embodiment. In one embodiment, processing core 170 may be

suitable for manufacture in one or more process technologies and by being represented on a machine-readable media in sufficient detail, may be suitable to facilitate the manufacture of all or part of data processing system 160 including processing core 170.

[0074] In one embodiment, SIMD coprocessor 161 comprises an execution unit 162 and a set of register files 164. One embodiment of main processor 165 comprises a decoder 165 to recognize instructions of instruction set 163 including instructions in accordance with one embodiment for execution by execution unit 162. In other embodiments, SIMD coprocessor 161 also comprises at least part of decoder 165 to decode instructions of instruction set 163. Processing core 170 may also include additional circuitry (not shown) which may be unnecessary to the understanding of embodiments of the present disclosure.

[0075] In operation, main processor 166 executes a stream of data processing instructions that control data processing operations of a general type including interactions with cache memory 167, and input/output system 168. Embedded within the stream of data processing instructions may be SIMD coprocessor instructions. Decoder 165 of main processor 166 recognizes these SIMD coprocessor instructions as being of a type that should be executed by an attached SIMD coprocessor 161. Accordingly, main processor 166 issues these SIMD coprocessor instructions (or control signals representing SIMD coprocessor instructions) on the coprocessor bus 166. From coprocessor bus 166, these instructions may be received by any attached SIMD coprocessors. In this case, SIMD coprocessor 161 may accept and execute any received SIMD coprocessor instructions intended for it.

[0076] Data may be received via wireless interface 169 for processing by the SIMD coprocessor instructions. For one example, voice communication may be received in the form of a digital signal, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples representative of the voice communications. For another example, compressed audio and/or video may be received in the form of a digital bit stream, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples and/or motion video frames. In one embodiment of processing core 170, main processor 166, and a SIMD coprocessor 161 may be integrated into a single processing core 170 comprising an execution unit 162, a set of register files 164, and a decoder 165 to recognize instructions of instruction set 163 including instructions in accordance with one embodiment.

[0077] FIGURE 2 is a block diagram of the micro-architecture for a processor 200 that may include logic circuits to perform instructions, in accordance with embodiments of the present disclosure. In some embodiments, an instruction in accordance with one embodiment may be implemented to operate on data elements having sizes of byte, word, doubleword, quadword, etc., as well as datatypes, such as single and double precision integer and floating point datatypes. In one embodiment, in-order front end 201 may implement a part of processor 200 that may fetch instructions to be executed and prepares the instructions to be used later in the processor pipeline. Front end 201 may include several units. In one embodiment, instruction prefetcher 226 fetches instructions from memory and feeds the instructions to an instruction decoder 228 which in turn decodes or interprets the instructions. For example, in one embodiment, the decoder decodes a received instruction into one or more operations called “micro-instructions” or “micro-operations” (also called micro op or uops) that the machine may execute. In other embodiments, the decoder parses the instruction into an opcode and corresponding data and control fields that may be used by the micro-architecture to perform operations in accordance with one embodiment. In one embodiment, trace cache 230 may assemble decoded uops into program ordered sequences or traces in uop queue 234 for execution. When trace cache 230 encounters a complex instruction, microcode ROM 232 provides the uops needed to complete the operation.

[0078] Some instructions may be converted into a single micro-op, whereas others need several micro-ops to complete the full operation. In one embodiment, if more than four micro-ops are needed to complete an instruction, decoder 228 may access microcode ROM 232 to perform the instruction. In one embodiment, an instruction may be decoded into a small number of micro ops for processing at instruction decoder 228. In another embodiment, an instruction may be stored within microcode ROM 232 should a number of micro-ops be needed to accomplish the operation. Trace cache 230 refers to an entry point programmable logic array (PLA) to determine a correct micro-instruction pointer for reading the micro-code sequences to complete one or more instructions in accordance with one embodiment from micro-code ROM 232. After microcode ROM 232 finishes sequencing micro-ops for an instruction, front end 201 of the machine may resume fetching micro-ops from trace cache 230.

[0079] Out-of-order execution engine 203 may prepare instructions for execution. The out-of-order execution logic has a number of buffers to smooth out and re-order the flow of instructions to optimize performance as they go down the pipeline and get scheduled for execution. The allocator logic allocates the machine buffers and resources that each uop needs in order to execute. The register renaming logic renames logic registers onto entries in a register file. The allocator also allocates an entry for each uop in one of the two uop queues, one for memory operations and one for non-memory operations, in front of the instruction schedulers: memory scheduler, fast scheduler 202, slow/general floating point scheduler 204, and simple floating point scheduler 206. Uop schedulers 202, 204, 206, determine when a uop is ready to execute based on the readiness of their dependent input register operand sources and the availability of the execution resources the uops need to complete their operation. Fast scheduler 202 of one embodiment may schedule on each half of the main clock cycle while the other schedulers may only schedule once per main processor clock cycle. The schedulers arbitrate for the dispatch ports to schedule uops for execution.

[0080] Register files 208, 210 may be arranged between schedulers 202, 204, 206, and execution units 212, 214, 216, 218, 220, 222, 224 in execution block 211. Each of register files 208, 210 perform integer and floating point operations, respectively. Each register file 208, 210, may include a bypass network that may bypass or forward just completed results that have not yet been written into the register file to new dependent uops. Integer register file 208 and floating point register file 210 may communicate data with the other. In one embodiment, integer register file 208 may be split into two separate register files, one register file for low-order thirty-two bits of data and a second register file for high order thirty-two bits of data. Floating point register file 210 may include 128-bit wide entries because floating point instructions typically have operands from 64 to 128 bits in width.

[0081] Execution block 211 may contain execution units 212, 214, 216, 218, 220, 222, 224. Execution units 212, 214, 216, 218, 220, 222, 224 may execute the instructions. Execution block 211 may include register files 208, 210 that store the integer and floating point data operand values that the micro-instructions need to execute. In one embodiment, processor 200 may comprise a number of execution units: address generation unit (AGU) 212, AGU 214, fast ALU 216, fast ALU 218, slow ALU 220, floating point ALU 222, floating point move unit 224. In another embodiment, floating point execution blocks 222, 224, may

execute floating point, MMX, SIMD, and SSE, or other operations. In yet another embodiment, floating point ALU 222 may include a 64-bit by 64-bit floating point divider to execute divide, square root, and remainder micro-ops. In various embodiments, instructions involving a floating point value may be handled with the floating point hardware. In one embodiment, ALU operations may be passed to high-speed ALU execution units 216, 218. High-speed ALUs 216, 218 may execute fast operations with an effective latency of half a clock cycle. In one embodiment, most complex integer operations go to slow ALU 220 as slow ALU 220 may include integer execution hardware for long-latency type of operations, such as a multiplier, shifts, flag logic, and branch processing. Memory load/store operations may be executed by AGUs 212, 214. In one embodiment, integer ALUs 216, 218, 220 may perform integer operations on 64-bit data operands. In other embodiments, ALUs 216, 218, 220 may be implemented to support a variety of data bit sizes including sixteen, thirty-two, 128, 256, etc. Similarly, floating point units 222, 224 may be implemented to support a range of operands having bits of various widths. In one embodiment, floating point units 222, 224, may operate on 128-bit wide packed data operands in conjunction with SIMD and multimedia instructions.

[0082] In one embodiment, uops schedulers 202, 204, 206, dispatch dependent operations before the parent load has finished executing. As uops may be speculatively scheduled and executed in processor 200, processor 200 may also include circuits with logic to handle memory misses. If a data load misses in the data cache, there may be dependent operations in flight in the pipeline that have left the scheduler with temporarily incorrect data. A replay mechanism tracks and re-executes instructions that use incorrect data. Only the dependent operations might need to be replayed and the independent ones may be allowed to complete. The schedulers and replay mechanism of one embodiment of a processor may also be designed to catch instruction sequences for text string comparison operations.

[0083] The term “registers” may refer to the on-board processor storage locations that may be used as part of instructions to identify operands. In other words, registers may be those that may be usable from the outside of the processor (from a programmer’s perspective). However, in some embodiments registers might not be limited to a particular type of circuit. Rather, a register may store data, provide data, and perform the functions

described herein. The registers described herein may be implemented by circuitry within a processor using any number of different techniques, such as dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. In one embodiment, integer registers store 32-bit integer data. A register file of one embodiment also contains eight multimedia SIMD registers for packed data. For the discussions below, the registers may be understood to be data registers designed to hold packed data, such as 64-bit wide MMX™ registers (also referred to as ‘mm’ registers in some instances) in microprocessors enabled with MMX technology from Intel Corporation of Santa Clara, California. These MMX registers, available in both integer and floating point forms, may operate with packed data elements that accompany SIMD and SSE instructions. Similarly, 128-bit wide XMM registers relating to SSE2, SSE3, SSE4, or beyond (referred to generically as “SSEx”) technology may hold such packed data operands. In one embodiment, in storing packed data and integer data, the registers do not need to differentiate between the two data types. In one embodiment, integer and floating point may be contained in the same register file or different register files. Furthermore, in one embodiment, floating point and integer data may be stored in different registers or the same registers.

[0084] In the examples of the following figures, a number of data operands may be described. FIGURE 3A illustrates various packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure. FIGURE 3A illustrates data types for a packed byte 310, a packed word 320, and a packed doubleword (dword) 330 for 128-bit wide operands. Packed byte format 310 of this example may be 128 bits long and contains sixteen packed byte data elements. A byte may be defined, for example, as eight bits of data. Information for each byte data element may be stored in bit 7 through bit 0 for byte 0, bit 15 through bit 8 for byte 1, bit 23 through bit 16 for byte 2, and finally bit 120 through bit 127 for byte 15. Thus, all available bits may be used in the register. This storage arrangement increases the storage efficiency of the processor. As well, with sixteen data elements accessed, one operation may now be performed on sixteen data elements in parallel.

[0085] Generally, a data element may include an individual piece of data that is stored in a single register or memory location with other data elements of the same length. In

packed data sequences relating to SSEx technology, the number of data elements stored in a XMM register may be 128 bits divided by the length in bits of an individual data element. Similarly, in packed data sequences relating to MMX and SSE technology, the number of data elements stored in an MMX register may be 64 bits divided by the length in bits of an individual data element. Although the data types illustrated in FIGURE 3A may be 128 bits long, embodiments of the present disclosure may also operate with 64-bit wide or other sized operands. Packed word format 320 of this example may be 128 bits long and contains eight packed word data elements. Each packed word contains sixteen bits of information. Packed doubleword format 330 of FIGURE 3A may be 128 bits long and contains four packed doubleword data elements. Each packed doubleword data element contains thirty-two bits of information. A packed quadword may be 128 bits long and contain two packed quad-word data elements.

[0086] FIGURE 3B illustrates possible in-register data storage formats, in accordance with embodiments of the present disclosure. Each packed data may include more than one independent data element. Three packed data formats are illustrated; packed half 341, packed single 342, and packed double 343. One embodiment of packed half 341, packed single 342, and packed double 343 contain fixed-point data elements. For another embodiment one or more of packed half 341, packed single 342, and packed double 343 may contain floating-point data elements. One embodiment of packed half 341 may be 128 bits long containing eight 16-bit data elements. One embodiment of packed single 342 may be 128 bits long and contains four 32-bit data elements. One embodiment of packed double 343 may be 128 bits long and contains two 64-bit data elements. It will be appreciated that such packed data formats may be further extended to other register lengths, for example, to 96-bits, 160-bits, 192-bits, 224-bits, 256-bits, 512-bits or more.

[0087] FIGURE 3C illustrates various signed and unsigned packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure. Unsigned packed byte representation 344 illustrates the storage of an unsigned packed byte in a SIMD register. Information for each byte data element may be stored in bit 7 through bit 0 for byte 0, bit 15 through bit 8 for byte 1, bit 23 through bit 16 for byte 2, and finally bit 120 through bit 127 for byte 15. Thus, all available bits may be used in the register. This storage arrangement may increase the storage efficiency of the processor. As well, with

sixteen data elements accessed, one operation may now be performed on sixteen data elements in a parallel fashion. Signed packed byte representation 345 illustrates the storage of a signed packed byte. Note that the eighth bit of every byte data element may be the sign indicator. Unsigned packed word representation 346 illustrates how word seven through word zero may be stored in a SIMD register. Signed packed word representation 347 may be similar to the unsigned packed word in-register representation 346. Note that the sixteenth bit of each word data element may be the sign indicator. Unsigned packed doubleword representation 348 shows how doubleword data elements are stored. Signed packed doubleword representation 349 may be similar to unsigned packed doubleword in-register representation 348. Note that the necessary sign bit may be the thirty-second bit of each doubleword data element.

[0088] FIGURE 3D illustrates an embodiment of an operation encoding (opcode). Furthermore, format 360 may include register/memory operand addressing modes corresponding with a type of opcode format described in the "IA-32 Intel Architecture Software Developer's Manual Volume 2: Instruction Set Reference," which is available from Intel Corporation, Santa Clara, CA on the world-wide-web (www) at intel.com/design/litcentr. In one embodiment, an instruction may be encoded by one or more of fields 361 and 362. Up to two operand locations per instruction may be identified, including up to two source operand identifiers 364 and 365. In one embodiment, destination operand identifier 366 may be the same as source operand identifier 364, whereas in other embodiments they may be different. In another embodiment, destination operand identifier 366 may be the same as source operand identifier 365, whereas in other embodiments they may be different. In one embodiment, one of the source operands identified by source operand identifiers 364 and 365 may be overwritten by the results of the text string comparison operations, whereas in other embodiments identifier 364 corresponds to a source register element and identifier 365 corresponds to a destination register element. In one embodiment, operand identifiers 364 and 365 may identify 32-bit or 64-bit source and destination operands.

[0089] FIGURE 3E illustrates another possible operation encoding (opcode) format 370, having forty or more bits, in accordance with embodiments of the present disclosure. Opcode format 370 corresponds with opcode format 360 and comprises an optional prefix byte 378. An instruction according to one embodiment may be encoded by one or more of

fields 378, 371, and 372. Up to two operand locations per instruction may be identified by source operand identifiers 374 and 375 and by prefix byte 378. In one embodiment, prefix byte 378 may be used to identify 32-bit or 64-bit source and destination operands. In one embodiment, destination operand identifier 376 may be the same as source operand identifier 374, whereas in other embodiments they may be different. For another embodiment, destination operand identifier 376 may be the same as source operand identifier 375, whereas in other embodiments they may be different. In one embodiment, an instruction operates on one or more of the operands identified by operand identifiers 374 and 375 and one or more operands identified by operand identifiers 374 and 375 may be overwritten by the results of the instruction, whereas in other embodiments, operands identified by identifiers 374 and 375 may be written to another data element in another register. Opcode formats 360 and 370 allow register to register, memory to register, register by memory, register by register, register by immediate, register to memory addressing specified in part by MOD fields 363 and 373 and by optional scale-index-base and displacement bytes.

[0090] FIGURE 3F illustrates yet another possible operation encoding (opcode) format, in accordance with embodiments of the present disclosure. 64-bit single instruction multiple data (SIMD) arithmetic operations may be performed through a coprocessor data processing (CDP) instruction. Operation encoding (opcode) format 380 depicts one such CDP instruction having CDP opcode fields 382 and 389. The type of CDP instruction, for another embodiment, operations may be encoded by one or more of fields 383, 384, 387, and 388. Up to three operand locations per instruction may be identified, including up to two source operand identifiers 385 and 390 and one destination operand identifier 386. One embodiment of the coprocessor may operate on eight, sixteen, thirty-two, and 64-bit values. In one embodiment, an instruction may be performed on integer data elements. In some embodiments, an instruction may be executed conditionally, using condition field 381. For some embodiments, source data sizes may be encoded by field 383. In some embodiments, Zero (Z), negative (N), carry (C), and overflow (V) detection may be done on SIMD fields. For some instructions, the type of saturation may be encoded by field 384.

[0091] FIGURE 4A is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline, in accordance with embodiments of

the present disclosure. FIGURE 4B is a block diagram illustrating an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor, in accordance with embodiments of the present disclosure. The solid lined boxes in FIGURE 4A illustrate the in-order pipeline, while the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline. Similarly, the solid lined boxes in FIGURE 4B illustrate the in-order architecture logic, while the dashed lined boxes illustrates the register renaming logic and out-of-order issue/execution logic.

[0092] In FIGURE 4A, a processor pipeline 400 may include a fetch stage 402, a length decode stage 404, a decode stage 406, an allocation stage 408, a renaming stage 410, a scheduling (also known as a dispatch or issue) stage 412, a register read/memory read stage 414, an execute stage 416, a write-back/memory-write stage 418, an exception handling stage 422, and a commit stage 424.

[0093] In FIGURE 4B, arrows denote a coupling between two or more units and the direction of the arrow indicates a direction of data flow between those units. FIGURE 4B shows processor core 490 including a front end unit 430 coupled to an execution engine unit 450, and both may be coupled to a memory unit 470.

[0094] Core 490 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. In one embodiment, core 490 may be a special-purpose core, such as, for example, a network or communication core, compression engine, graphics core, or the like.

[0095] Front end unit 430 may include a branch prediction unit 432 coupled to an instruction cache unit 434. Instruction cache unit 434 may be coupled to an instruction translation lookaside buffer (TLB) 436. TLB 436 may be coupled to an instruction fetch unit 438, which is coupled to a decode unit 440. Decode unit 440 may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which may be decoded from, or which otherwise reflect, or may be derived from, the original instructions. The decoder may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read-only memories (ROMs), etc. In one embodiment,

instruction cache unit 434 may be further coupled to a level 2 (L2) cache unit 476 in memory unit 470. Decode unit 440 may be coupled to a rename/allocator unit 452 in execution engine unit 450.

[0096] Execution engine unit 450 may include rename/allocator unit 452 coupled to a retirement unit 454 and a set of one or more scheduler units 456. Scheduler units 456 represent any number of different schedulers, including reservations stations, central instruction window, etc. Scheduler units 456 may be coupled to physical register file units 458. Each of physical register file units 458 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, etc., status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. Physical register file units 458 may be overlapped by retirement unit 154 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using one or more reorder buffers and one or more retirement register files, using one or more future files, one or more history buffers, and one or more retirement register files; using register maps and a pool of registers; etc.). Generally, the architectural registers may be visible from the outside of the processor or from a programmer's perspective. The registers might not be limited to any known particular type of circuit. Various different types of registers may be suitable as long as they store and provide data as described herein. Examples of suitable registers include, but might not be limited to, dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. Retirement unit 454 and physical register file units 458 may be coupled to execution clusters 460. Execution clusters 460 may include a set of one or more execution units 462 and a set of one or more memory access units 464. Execution units 462 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. Scheduler units 456, physical register file units 458, and execution clusters 460 are shown as being possibly plural because certain embodiments

create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file unit, and/or execution cluster – and in the case of a separate memory access pipeline, certain embodiments may be implemented in which only the execution cluster of this pipeline has memory access units 464). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

[0097] The set of memory access units 464 may be coupled to memory unit 470, which may include a data TLB unit 472 coupled to a data cache unit 474 coupled to a level 2 (L2) cache unit 476. In one exemplary embodiment, memory access units 464 may include a load unit, a store address unit, and a store data unit, each of which may be coupled to data TLB unit 472 in memory unit 470. L2 cache unit 476 may be coupled to one or more other levels of cache and eventually to a main memory.

[0098] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement pipeline 400 as follows: 1) instruction fetch 438 may perform fetch and length decoding stages 402 and 404; 2) decode unit 440 may perform decode stage 406; 3) rename/allocator unit 452 may perform allocation stage 408 and renaming stage 410; 4) scheduler units 456 may perform schedule stage 412; 5) physical register file units 458 and memory unit 470 may perform register read/memory read stage 414; execution cluster 460 may perform execute stage 416; 6) memory unit 470 and physical register file units 458 may perform write-back/memory-write stage 418; 7) various units may be involved in the performance of exception handling stage 422; and 8) retirement unit 454 and physical register file units 458 may perform commit stage 424.

[0099] Core 490 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA).

[0100] It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads) in a variety of manners. Multithreading support may be performed by, for example, including time sliced multithreading,

simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof. Such a combination may include, for example, time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology.

[0101] While register renaming may be described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor may also include a separate instruction and data cache units 434/474 and a shared L2 cache unit 476, other embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that may be external to the core and/or the processor. In other embodiments, all of the cache may be external to the core and/or the processor.

[0102] FIGURE 5A is a block diagram of a processor 500, in accordance with embodiments of the present disclosure. In one embodiment, processor 500 may include a multicore processor. Processor 500 may include a system agent 510 communicatively coupled to one or more cores 502. Furthermore, cores 502 and system agent 510 may be communicatively coupled to one or more caches 506. Cores 502, system agent 510, and caches 506 may be communicatively coupled via one or more memory control units 552. Furthermore, cores 502, system agent 510, and caches 506 may be communicatively coupled to a graphics module 560 via memory control units 552.

[0103] Processor 500 may include any suitable mechanism for interconnecting cores 502, system agent 510, and caches 506, and graphics module 560. In one embodiment, processor 500 may include a ring-based interconnect unit 508 to interconnect cores 502, system agent 510, and caches 506, and graphics module 560. In other embodiments, processor 500 may include any number of well-known techniques for interconnecting such units. Ring-based interconnect unit 508 may utilize memory control units 552 to facilitate interconnections.

[0104] Processor 500 may include a memory hierarchy comprising one or more levels of caches within the cores, one or more shared cache units such as caches 506, or external memory (not shown) coupled to the set of integrated memory controller units 552.

Caches 506 may include any suitable cache. In one embodiment, caches 506 may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof.

[0105] In various embodiments, one or more of cores 502 may perform multi-threading. System agent 510 may include components for coordinating and operating cores 502. System agent unit 510 may include for example a power control unit (PCU). The PCU may be or include logic and components needed for regulating the power state of cores 502. System agent 510 may include a display engine 512 for driving one or more externally connected displays or graphics module 560. System agent 510 may include an interface for communications busses for graphics. In one embodiment, the interface may be implemented by PCI Express (PCIe). In a further embodiment, the interface may be implemented by PCI Express Graphics (PEG) 514. System agent 510 may include a direct media interface (DMI) 516. DMI 516 may provide links between different bridges on a motherboard or other portion of a computer system. System agent 510 may include a PCIe bridge 518 for providing PCIe links to other elements of a computing system. PCIe bridge 518 may be implemented using a memory controller 520 and coherence logic 522.

[0106] Cores 502 may be implemented in any suitable manner. Cores 502 may be homogenous or heterogeneous in terms of architecture and/or instruction set. In one embodiment, some of cores 502 may be in-order while others may be out-of-order. In another embodiment, two or more of cores 502 may execute the same instruction set, while others may execute only a subset of that instruction set or a different instruction set.

[0107] Processor 500 may include a general-purpose processor, such as a Core™ i3, i5, i7, 2 Duo and Quad, Xeon™, Itanium™, XScale™ or StrongARM™ processor, which may be available from Intel Corporation, of Santa Clara, Calif. Processor 500 may be provided from another company, such as ARM Holdings, Ltd, MIPS, etc. Processor 500 may be a special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, co-processor, embedded processor, or the like. Processor 500 may be implemented on one or more chips. Processor 500 may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0108] In one embodiment, a given one of caches 506 may be shared by multiple ones of cores 502. In another embodiment, a given one of caches 506 may be dedicated to one of cores 502. The assignment of caches 506 to cores 502 may be handled by a cache controller or other suitable mechanism. A given one of caches 506 may be shared by two or more cores 502 by implementing time-slices of a given cache 506.

[0109] Graphics module 560 may implement an integrated graphics processing subsystem. In one embodiment, graphics module 560 may include a graphics processor. Furthermore, graphics module 560 may include a media engine 565. Media engine 565 may provide media encoding and video decoding.

[0110] FIGURE 5B is a block diagram of an example implementation of a core 502, in accordance with embodiments of the present disclosure. Core 502 may include a front end 570 communicatively coupled to an out-of-order engine 580. Core 502 may be communicatively coupled to other portions of processor 500 through cache hierarchy 503.

[0111] Front end 570 may be implemented in any suitable manner, such as fully or in part by front end 201 as described above. In one embodiment, front end 570 may communicate with other portions of processor 500 through cache hierarchy 503. In a further embodiment, front end 570 may fetch instructions from portions of processor 500 and prepare the instructions to be used later in the processor pipeline as they are passed to out-of-order execution engine 580.

[0112] Out-of-order execution engine 580 may be implemented in any suitable manner, such as fully or in part by out-of-order execution engine 203 as described above. Out-of-order execution engine 580 may prepare instructions received from front end 570 for execution. Out-of-order execution engine 580 may include an allocate module 1282. In one embodiment, allocate module 1282 may allocate resources of processor 500 or other resources, such as registers or buffers, to execute a given instruction. Allocate module 1282 may make allocations in schedulers, such as a memory scheduler, fast scheduler, or floating point scheduler. Such schedulers may be represented in FIGURE 5B by resource schedulers 584. Allocate module 1282 may be implemented fully or in part by the allocation logic described in conjunction with FIGURE 2. Resource schedulers 584 may determine when an instruction is ready to execute based on the readiness of a given resource's sources and the availability of execution resources needed to execute an instruction. Resource schedulers 584

may be implemented by, for example, schedulers 202, 204, 206 as discussed above. Resource schedulers 584 may schedule the execution of instructions upon one or more resources. In one embodiment, such resources may be internal to core 502, and may be illustrated, for example, as resources 586. In another embodiment, such resources may be external to core 502 and may be accessible by, for example, cache hierarchy 503. Resources may include, for example, memory, caches, register files, or registers. Resources internal to core 502 may be represented by resources 586 in FIGURE 5B. As necessary, values written to or read from resources 586 may be coordinated with other portions of processor 500 through, for example, cache hierarchy 503. As instructions are assigned resources, they may be placed into a reorder buffer 588. Reorder buffer 588 may track instructions as they are executed and may selectively reorder their execution based upon any suitable criteria of processor 500. In one embodiment, reorder buffer 588 may identify instructions or a series of instructions that may be executed independently. Such instructions or a series of instructions may be executed in parallel from other such instructions. Parallel execution in core 502 may be performed by any suitable number of separate execution blocks or virtual processors. In one embodiment, shared resources—such as memory, registers, and caches—may be accessible to multiple virtual processors within a given core 502. In other embodiments, shared resources may be accessible to multiple processing entities within processor 500.

[0113] Cache hierarchy 503 may be implemented in any suitable manner. For example, cache hierarchy 503 may include one or more lower or mid-level caches, such as caches 572, 574. In one embodiment, cache hierarchy 503 may include an LLC 595 communicatively coupled to caches 572, 574. In another embodiment, LLC 595 may be implemented in a module 590 accessible to all processing entities of processor 500. In a further embodiment, module 590 may be implemented in an uncore module of processors from Intel, Inc. Module 590 may include portions or subsystems of processor 500 necessary for the execution of core 502 but might not be implemented within core 502. Besides LLC 595, Module 590 may include, for example, hardware interfaces, memory coherency coordinators, interprocessor interconnects, instruction pipelines, or memory controllers. Access to RAM 599 available to processor 500 may be made through module 590 and, more specifically, LLC 595. Furthermore, other instances of core 502 may similarly access module 590. Coordination of the instances of core 502 may be facilitated in part through module 590.

[0114] FIGURES 6–8 may illustrate exemplary systems suitable for including processor 500, while FIGURE 9 may illustrate an exemplary system on a chip (SoC) that may include one or more of cores 502. Other system designs and implementations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, may also be suitable. In general, a huge variety of systems or electronic devices that incorporate a processor and/or other execution logic as disclosed herein may be generally suitable.

[0115] FIGURE 6 illustrates a block diagram of a system 600, in accordance with embodiments of the present disclosure. System 600 may include one or more processors 610, 615, which may be coupled to graphics memory controller hub (GMCH) 620. The optional nature of additional processors 615 is denoted in FIGURE 6 with broken lines.

[0116] Each processor 610,615 may be some version of processor 500. However, it should be noted that integrated graphics logic and integrated memory control units might not exist in processors 610,615. FIGURE 6 illustrates that GMCH 620 may be coupled to a memory 640 that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one embodiment, be associated with a non-volatile cache.

[0117] GMCH 620 may be a chipset, or a portion of a chipset. GMCH 620 may communicate with processors 610, 615 and control interaction between processors 610, 615 and memory 640. GMCH 620 may also act as an accelerated bus interface between the processors 610, 615 and other elements of system 600. In one embodiment, GMCH 620 communicates with processors 610, 615 via a multi-drop bus, such as a frontside bus (FSB) 695.

[0118] Furthermore, GMCH 620 may be coupled to a display 645 (such as a flat panel display). In one embodiment, GMCH 620 may include an integrated graphics accelerator. GMCH 620 may be further coupled to an input/output (I/O) controller hub (ICH) 650, which may be used to couple various peripheral devices to system 600. External graphics device 660 may include be a discrete graphics device coupled to ICH 650 along with another peripheral device 670.

[0119] In other embodiments, additional or different processors may also be present in system 600. For example, additional processors 610, 615 may include additional processors that may be the same as processor 610, additional processors that may be heterogeneous or asymmetric to processor 610, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor. There may be a variety of differences between the physical resources 610, 615 in terms of a spectrum of metrics of merit including architectural, micro-architectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst processors 610, 615. For at least one embodiment, various processors 610, 615 may reside in the same die package.

[0120] FIGURE 7 illustrates a block diagram of a second system 700, in accordance with embodiments of the present disclosure. As shown in FIGURE 7, multiprocessor system 700 may include a point-to-point interconnect system, and may include a first processor 770 and a second processor 780 coupled via a point-to-point interconnect 750. Each of processors 770 and 780 may be some version of processor 500 as one or more of processors 610, 615.

[0121] While FIGURE 7 may illustrate two processors 770, 780, it is to be understood that the scope of the present disclosure is not so limited. In other embodiments, one or more additional processors may be present in a given processor.

[0122] Processors 770 and 780 are shown including integrated memory controller units 772 and 782, respectively. Processor 770 may also include as part of its bus controller units point-to-point (P-P) interfaces 776 and 778; similarly, second processor 780 may include P-P interfaces 786 and 788. Processors 770, 780 may exchange information via a point-to-point (P-P) interface 750 using P-P interface circuits 778, 788. As shown in FIGURE 7, IMCs 772 and 782 may couple the processors to respective memories, namely a memory 732 and a memory 734, which in one embodiment may be portions of main memory locally attached to the respective processors.

[0123] Processors 770, 780 may each exchange information with a chipset 790 via individual P-P interfaces 752, 754 using point to point interface circuits 776, 794, 786, 798. In one embodiment, chipset 790 may also exchange information with a high-performance graphics circuit 738 via a high-performance graphics interface 739.

[0124] A shared cache (not shown) may be included in either processor or outside of both processors, yet connected with the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

[0125] Chipset 790 may be coupled to a first bus 716 via an interface 796. In one embodiment, first bus 716 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present disclosure is not so limited.

[0126] As shown in FIGURE 7, various I/O devices 714 may be coupled to first bus 716, along with a bus bridge 718 which couples first bus 716 to a second bus 720. In one embodiment, second bus 720 may be a low pin count (LPC) bus. Various devices may be coupled to second bus 720 including, for example, a keyboard and/or mouse 722, communication devices 727 and a storage unit 728 such as a disk drive or other mass storage device which may include instructions/code and data 730, in one embodiment. Further, an audio I/O 724 may be coupled to second bus 720. Note that other architectures may be possible. For example, instead of the point-to-point architecture of FIGURE 7, a system may implement a multi-drop bus or other such architecture.

[0127] FIGURE 8 illustrates a block diagram of a third system 800 in accordance with embodiments of the present disclosure. Like elements in FIGURES 7 and 8 bear like reference numerals, and certain aspects of FIGURE 7 have been omitted from FIGURE 8 in order to avoid obscuring other aspects of FIGURE 8.

[0128] FIGURE 8 illustrates that processors 870, 880 may include integrated memory and I/O control logic ("CL") 872 and 882, respectively. For at least one embodiment, CL 872, 882 may include integrated memory controller units such as that described above in connection with FIGURES 5 and 7. In addition, CL 872, 882 may also include I/O control logic. FIGURE 8 illustrates that not only memories 832, 834 may be coupled to CL 872, 882, but also that I/O devices 814 may also be coupled to control logic 872, 882. Legacy I/O devices 815 may be coupled to chipset 890.

[0129] FIGURE 9 illustrates a block diagram of a SoC 900, in accordance with embodiments of the present disclosure. Similar elements in FIGURE 5 bear like reference numerals. Also, dashed lined boxes may represent optional features on more advanced SoCs.

An interconnect units 902 may be coupled to: an application processor 910 which may include a set of one or more cores 902A-N and shared cache units 906; a system agent unit 910; a bus controller units 916; an integrated memory controller units 914; a set or one or more media processors 920 which may include integrated graphics logic 908, an image processor 924 for providing still and/or video camera functionality, an audio processor 926 for providing hardware audio acceleration, and a video processor 928 for providing video encode/decode acceleration; an static random access memory (SRAM) unit 930; a direct memory access (DMA) unit 932; and a display unit 940 for coupling to one or more external displays.

[0130] FIGURE 10 illustrates a processor containing a central processing unit (CPU) and a graphics processing unit (GPU), which may perform at least one instruction, in accordance with embodiments of the present disclosure. In one embodiment, an instruction to perform operations according to at least one embodiment could be performed by the CPU. In another embodiment, the instruction could be performed by the GPU. In still another embodiment, the instruction may be performed through a combination of operations performed by the GPU and the CPU. For example, in one embodiment, an instruction in accordance with one embodiment may be received and decoded for execution on the GPU. However, one or more operations within the decoded instruction may be performed by a CPU and the result returned to the GPU for final retirement of the instruction. Conversely, in some embodiments, the CPU may act as the primary processor and the GPU as the co-processor.

[0131] In some embodiments, instructions that benefit from highly parallel, throughput processors may be performed by the GPU, while instructions that benefit from the performance of processors that benefit from deeply pipelined architectures may be performed by the CPU. For example, graphics, scientific applications, financial applications and other parallel workloads may benefit from the performance of the GPU and be executed accordingly, whereas more sequential applications, such as operating system kernel or application code may be better suited for the CPU.

[0132] In FIGURE 10, processor 1000 includes a CPU 1005, GPU 1010, image processor 1015, video processor 1020, USB controller 1025, UART controller 1030, SPI/SDIO controller 1035, display device 1040, memory interface controller 1045, MIPI controller 1050, flash memory controller 1055, dual data rate (DDR) controller 1060, security engine 1065, and

I²S/I²C controller 1070. Other logic and circuits may be included in the processor of FIGURE 10, including more CPUs or GPUs and other peripheral interface controllers.

[0133] One or more aspects of at least one embodiment may be implemented by representative data stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as “IP cores” may be stored on a tangible, machine-readable medium (“tape”) and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor. For example, IP cores, such as the Cortex™ family of processors developed by ARM Holdings, Ltd. and Loongson IP cores developed the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences may be licensed or sold to various customers or licensees, such as Texas Instruments, Qualcomm, Apple, or Samsung and implemented in processors produced by these customers or licensees.

[0134] FIGURE 11 illustrates a block diagram illustrating the development of IP cores, in accordance with embodiments of the present disclosure. Storage 1130 may include simulation software 1120 and/or hardware or software model 1110. In one embodiment, the data representing the IP core design may be provided to storage 1130 via memory 1140 (e.g., hard disk), wired connection (e.g., internet) 1150 or wireless connection 1160. The IP core information generated by the simulation tool and model may then be transmitted to a fabrication facility where it may be fabricated by a 3rd party to perform at least one instruction in accordance with at least one embodiment.

[0135] In some embodiments, one or more instructions may correspond to a first type or architecture (e.g., x86) and be translated or emulated on a processor of a different type or architecture (e.g., ARM). An instruction, according to one embodiment, may therefore be performed on any processor or processor type, including ARM, x86, MIPS, a GPU, or other processor type or architecture.

[0136] FIGURE 12 illustrates how an instruction of a first type may be emulated by a processor of a different type, in accordance with embodiments of the present disclosure. In FIGURE 12, program 1205 contains some instructions that may perform the same or substantially the same function as an instruction according to one embodiment. However the instructions of program 1205 may be of a type and/or format that is different from or

incompatible with processor 1215, meaning the instructions of the type in program 1205 may not be able to execute natively by the processor 1215. However, with the help of emulation logic, 1210, the instructions of program 1205 may be translated into instructions that may be natively be executed by the processor 1215. In one embodiment, the emulation logic may be embodied in hardware. In another embodiment, the emulation logic may be embodied in a tangible, machine-readable medium containing software to translate instructions of the type in program 1205 into the type natively executable by processor 1215. In other embodiments, emulation logic may be a combination of fixed-function or programmable hardware and a program stored on a tangible, machine-readable medium. In one embodiment, the processor contains the emulation logic, whereas in other embodiments, the emulation logic exists outside of the processor and may be provided by a third party. In one embodiment, the processor may load the emulation logic embodied in a tangible, machine-readable medium containing software by executing microcode or firmware contained in or associated with the processor.

[0137] FIGURE 13 illustrates a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set, in accordance with embodiments of the present disclosure. In the illustrated embodiment, the instruction converter may be a software instruction converter, although the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. FIGURE 13 shows a program in a high level language 1302 may be compiled using an x86 compiler 1304 to generate x86 binary code 1306 that may be natively executed by a processor with at least one x86 instruction set core 1316. The processor with at least one x86 instruction set core 1316 represents any processor that may perform substantially the same functions as an Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. x86 compiler 1304 represents a compiler that may be operable to generate x86 binary code 1306 (e.g., object code) that may, with or without additional linkage processing, be executed on the processor with at least one x86 instruction

set core 1316. Similarly, FIGURE 13 shows the program in high level language 1302 may be compiled using an alternative instruction set compiler 1308 to generate alternative instruction set binary code 1310 that may be natively executed by a processor without at least one x86 instruction set core 1314 (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, CA and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, CA). Instruction converter 1312 may be used to convert x86 binary code 1306 into code that may be natively executed by the processor without an x86 instruction set core 1314. This converted code might not be the same as alternative instruction set binary code 1310; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, instruction converter 1312 represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute x86 binary code 1306.

[0138] FIGURE 14 is a block diagram of an instruction set architecture 1400 of a processor, in accordance with embodiments of the present disclosure. Instruction set architecture 1400 may include any suitable number or kind of components.

[0139] For example, instruction set architecture 1400 may include processing entities such as one or more cores 1406, 1407 and a graphics processing unit 1415. Cores 1406, 1407 may be communicatively coupled to the rest of instruction set architecture 1400 through any suitable mechanism, such as through a bus or cache. In one embodiment, cores 1406, 1407 may be communicatively coupled through an L2 cache control 1408, which may include a bus interface unit 1409 and an L2 cache 1410. Cores 1406, 1407 and graphics processing unit 1415 may be communicatively coupled to each other and to the remainder of instruction set architecture 1400 through interconnect 1410. In one embodiment, graphics processing unit 1415 may use a video code 1420 defining the manner in which particular video signals will be encoded and decoded for output.

[0140] Instruction set architecture 1400 may also include any number or kind of interfaces, controllers, or other mechanisms for interfacing or communicating with other portions of an electronic device or system. Such mechanisms may facilitate interaction with, for example, peripherals, communications devices, other processors, or memory. In the

example of FIGURE 14, instruction set architecture 1400 may include a liquid crystal display (LCD) video interface 1425, a subscriber interface module (SIM) interface 1430, a boot ROM interface 1435, a synchronous dynamic random access memory (SDRAM) controller 1440, a flash controller 1445, and a serial peripheral interface (SPI) master unit 1450. LCD video interface 1425 may provide output of video signals from, for example, GPU 1415 and through, for example, a mobile industry processor interface (MIPI) 1490 or a high-definition multimedia interface (HDMI) 1495 to a display. Such a display may include, for example, an LCD. SIM interface 1430 may provide access to or from a SIM card or device. SDRAM controller 1440 may provide access to or from memory such as an SDRAM chip or module. Flash controller 1445 may provide access to or from memory such as flash memory or other instances of RAM. SPI master unit 1450 may provide access to or from communications modules, such as a Bluetooth module 1470, high-speed 3G modem 1475, global positioning system module 1480, or wireless module 1485 implementing a communications standard such as 802.11.

[0141] FIGURE 15 is a more detailed block diagram of an instruction set architecture 1500 of a processor, in accordance with embodiments of the present disclosure. Instruction architecture 1500 may implement one or more aspects of instruction set architecture 1400. Furthermore, instruction set architecture 1500 may illustrate modules and mechanisms for the execution of instructions within a processor.

[0142] Instruction architecture 1500 may include a memory system 1540 communicatively coupled to one or more execution entities 1565. Furthermore, instruction architecture 1500 may include a caching and bus interface unit such as unit 1510 communicatively coupled to execution entities 1565 and memory system 1540. In one embodiment, loading of instructions into execution entities 1564 may be performed by one or more stages of execution. Such stages may include, for example, instruction prefetch stage 1530, dual instruction decode stage 1550, register rename stage 155, issue stage 1560, and writeback stage 1570.

[0143] In one embodiment, memory system 1540 may include an executed instruction pointer 1580. Executed instruction pointer 1580 may store a value identifying the oldest, undispached instruction within a batch of instructions. The oldest instruction may correspond to the lowest Program Order (PO) value. A PO may include a unique number of an

instruction. Such an instruction may be a single instruction within a thread represented by multiple strands. A PO may be used in ordering instructions to ensure correct execution semantics of code. A PO may be reconstructed by mechanisms such as evaluating increments to PO encoded in the instruction rather than an absolute value. Such a reconstructed PO may be known as an “RPO.” Although a PO may be referenced herein, such a PO may be used interchangeably with an RPO. A strand may include a sequence of instructions that are data dependent upon each other. The strand may be arranged by a binary translator at compilation time. Hardware executing a strand may execute the instructions of a given strand in order according to PO of the various instructions. A thread may include multiple strands such that instructions of different strands may depend upon each other. A PO of a given strand may be the PO of the oldest instruction in the strand which has not yet been dispatched to execution from an issue stage. Accordingly, given a thread of multiple strands, each strand including instructions ordered by PO, executed instruction pointer 1580 may store the oldest—illustrated by the lowest number—PO in the thread.

[0144] In another embodiment, memory system 1540 may include a retirement pointer 1582. Retirement pointer 1582 may store a value identifying the PO of the last retired instruction. Retirement pointer 1582 may be set by, for example, retirement unit 454. If no instructions have yet been retired, retirement pointer 1582 may include a null value.

[0145] Execution entities 1565 may include any suitable number and kind of mechanisms by which a processor may execute instructions. In the example of FIGURE 15, execution entities 1565 may include ALU/multiplication units (MUL) 1566, ALUs 1567, and floating point units (FPU) 1568. In one embodiment, such entities may make use of information contained within a given address 1569. Execution entities 1565 in combination with stages 1530, 1550, 1555, 1560, 1570 may collectively form an execution unit.

[0146] Unit 1510 may be implemented in any suitable manner. In one embodiment, unit 1510 may perform cache control. In such an embodiment, unit 1510 may thus include a cache 1525. Cache 1525 may be implemented, in a further embodiment, as an L2 unified cache with any suitable size, such as zero, 128k, 256k, 512k, 1M, or 2M bytes of memory. In another, further embodiment, cache 1525 may be implemented in error-correcting code memory. In another embodiment, unit 1510 may perform bus interfacing to other portions of a processor or electronic device. In such an embodiment, unit 1510 may thus include a bus

interface unit 1520 for communicating over an interconnect, intraprocessor bus, interprocessor bus, or other communication bus, port, or line. Bus interface unit 1520 may provide interfacing in order to perform, for example, generation of the memory and input/output addresses for the transfer of data between execution entities 1565 and the portions of a system external to instruction architecture 1500.

[0147] To further facilitate its functions, bus interface unit 1520 may include an interrupt control and distribution unit 1511 for generating interrupts and other communications to other portions of a processor or electronic device. In one embodiment, bus interface unit 1520 may include a snoop control unit 1512 that handles cache access and coherency for multiple processing cores. In a further embodiment, to provide such functionality, snoop control unit 1512 may include a cache-to-cache transfer unit that handles information exchanges between different caches. In another, further embodiment, snoop control unit 1512 may include one or more snoop filters 1514 that monitors the coherency of other caches (not shown) so that a cache controller, such as unit 1510, does not have to perform such monitoring directly. Unit 1510 may include any suitable number of timers 1515 for synchronizing the actions of instruction architecture 1500. Also, unit 1510 may include an AC port 1516.

[0148] Memory system 1540 may include any suitable number and kind of mechanisms for storing information for the processing needs of instruction architecture 1500. In one embodiment, memory system 1504 may include a load store unit 1530 for storing information such as buffers written to or read back from memory or registers. In another embodiment, memory system 1504 may include a translation lookaside buffer (TLB) 1545 that provides look-up of address values between physical and virtual addresses. In yet another embodiment, bus interface unit 1520 may include a memory management unit (MMU) 1544 for facilitating access to virtual memory. In still yet another embodiment, memory system 1504 may include a prefetcher 1543 for requesting instructions from memory before such instructions are actually needed to be executed, in order to reduce latency.

[0149] The operation of instruction architecture 1500 to execute an instruction may be performed through different stages. For example, using unit 1510 instruction prefetch stage 1530 may access an instruction through prefetcher 1543. Instructions retrieved may be

stored in instruction cache 1532. Prefetch stage 1530 may enable an option 1531 for fast-loop mode, wherein a series of instructions forming a loop that is small enough to fit within a given cache are executed. In one embodiment, such an execution may be performed without needing to access additional instructions from, for example, instruction cache 1532. Determination of what instructions to prefetch may be made by, for example, branch prediction unit 1535, which may access indications of execution in global history 1536, indications of target addresses 1537, or contents of a return stack 1538 to determine which of branches 1557 of code will be executed next. Such branches may be possibly prefetched as a result. Branches 1557 may be produced through other stages of operation as described below. Instruction prefetch stage 1530 may provide instructions as well as any predictions about future instructions to dual instruction decode stage.

[0150] Dual instruction decode stage 1550 may translate a received instruction into microcode-based instructions that may be executed. Dual instruction decode stage 1550 may simultaneously decode two instructions per clock cycle. Furthermore, dual instruction decode stage 1550 may pass its results to register rename stage 1555. In addition, dual instruction decode stage 1550 may determine any resulting branches from its decoding and eventual execution of the microcode. Such results may be input into branches 1557.

[0151] Register rename stage 1555 may translate references to virtual registers or other resources into references to physical registers or resources. Register rename stage 1555 may include indications of such mapping in a register pool 1556. Register rename stage 1555 may alter the instructions as received and send the result to issue stage 1560.

[0152] Issue stage 1560 may issue or dispatch commands to execution entities 1565. Such issuance may be performed in an out-of-order fashion. In one embodiment, multiple instructions may be held at issue stage 1560 before being executed. Issue stage 1560 may include an instruction queue 1561 for holding such multiple commands. Instructions may be issued by issue stage 1560 to a particular processing entity 1565 based upon any acceptable criteria, such as availability or suitability of resources for execution of a given instruction. In one embodiment, issue stage 1560 may reorder the instructions within instruction queue 1561 such that the first instructions received might not be the first instructions executed. Based upon the ordering of instruction queue 1561, additional branching information may be

provided to branches 1557. Issue stage 1560 may pass instructions to executing entities 1565 for execution.

[0153] Upon execution, writeback stage 1570 may write data into registers, queues, or other structures of instruction set architecture 1500 to communicate the completion of a given command. Depending upon the order of instructions arranged in issue stage 1560, the operation of writeback stage 1570 may enable additional instructions to be executed. Performance of instruction set architecture 1500 may be monitored or debugged by trace unit 1575.

[0154] FIGURE 16 is a block diagram of an execution pipeline 1600 for an instruction set architecture of a processor, in accordance with embodiments of the present disclosure. Execution pipeline 1600 may illustrate operation of, for example, instruction architecture 1500 of FIGURE 15.

[0155] Execution pipeline 1600 may include any suitable combination of steps or operations. In 1605, predictions of the branch that is to be executed next may be made. In one embodiment, such predictions may be based upon previous executions of instructions and the results thereof. In 1610, instructions corresponding to the predicted branch of execution may be loaded into an instruction cache. In 1615, one or more such instructions in the instruction cache may be fetched for execution. In 1620, the instructions that have been fetched may be decoded into microcode or more specific machine language. In one embodiment, multiple instructions may be simultaneously decoded. In 1625, references to registers or other resources within the decoded instructions may be reassigned. For example, references to virtual registers may be replaced with references to corresponding physical registers. In 1630, the instructions may be dispatched to queues for execution. In 1640, the instructions may be executed. Such execution may be performed in any suitable manner. In 1650, the instructions may be issued to a suitable execution entity. The manner in which the instruction is executed may depend upon the specific entity executing the instruction. For example, at 1655, an ALU may perform arithmetic functions. The ALU may utilize a single clock cycle for its operation, as well as two shifters. In one embodiment, two ALUs may be employed, and thus two instructions may be executed at 1655. At 1660, a determination of a resulting branch may be made. A program counter may be used to designate the destination to which the branch will be made. 1660 may be executed within a single clock cycle. At 1665,

floating point arithmetic may be performed by one or more FPUs. The floating point operation may require multiple clock cycles to execute, such as two to ten cycles. At 1670, multiplication and division operations may be performed. Such operations may be performed in four clock cycles. At 1675, loading and storing operations to registers or other portions of pipeline 1600 may be performed. The operations may include loading and storing addresses. Such operations may be performed in four clock cycles. At 1680, write-back operations may be performed as required by the resulting operations of 1655-1675.

[0156] FIGURE 17 is a block diagram of an electronic device 1700 for utilizing a processor 1710, in accordance with embodiments of the present disclosure. Electronic device 1700 may include, for example, a notebook, an ultrabook, a computer, a tower server, a rack server, a blade server, a laptop, a desktop, a tablet, a mobile device, a phone, an embedded computer, or any other suitable electronic device.

[0157] Electronic device 1700 may include processor 1710 communicatively coupled to any suitable number or kind of components, peripherals, modules, or devices. Such coupling may be accomplished by any suitable kind of bus or interface, such as I²C bus, system management bus (SMBus), low pin count (LPC) bus, SPI, high definition audio (HDA) bus, Serial Advance Technology Attachment (SATA) bus, USB bus (versions 1, 2, 3), or Universal Asynchronous Receiver/Transmitter (UART) bus.

[0158] Such components may include, for example, a display 1724, a touch screen 1725, a touch pad 1730, a near field communications (NFC) unit 1745, a sensor hub 1740, a thermal sensor 1746, an express chipset (EC) 1735, a trusted platform module (TPM) 1738, BIOS/firmware/flash memory 1722, a digital signal processor 1760, a drive 1720 such as a solid state disk (SSD) or a hard disk drive (HDD), a wireless local area network (WLAN) unit 1750, a Bluetooth unit 1752, a wireless wide area network (WWAN) unit 1756, a global positioning system (GPS), a camera 1754 such as a USB 3.0 camera, or a low power double data rate (LPDDR) memory unit 1715 implemented in, for example, the LPDDR3 standard. These components may each be implemented in any suitable manner.

[0159] Furthermore, in various embodiments other components may be communicatively coupled to processor 1710 through the components discussed above. For example, an accelerometer 1741, ambient light sensor (ALS) 1742, compass 1743, and gyroscope 1744 may be communicatively coupled to sensor hub 1740. A thermal sensor 1739,

fan 1737, keyboard 1746, and touch pad 1730 may be communicatively coupled to EC 1735. Speaker 1763, headphones 1764, and a microphone 1765 may be communicatively coupled to an audio unit 1764, which may in turn be communicatively coupled to DSP 1760. Audio unit 1764 may include, for example, an audio codec and a class D amplifier. A SIM card 1757 may be communicatively coupled to WWAN unit 1756. Components such as WLAN unit 1750 and Bluetooth unit 1752, as well as WWAN unit 1756 may be implemented in a next generation form factor (NGFF).

[0160] Building on the disclosure of FIGURES 1 – 17, embodiments of the present specification may provide for register reclamation. In certain embodiments, early register reclamation depends on a binary translator, as described below.

[0161] FIGURE 18 is an illustration of a system 1800 with a binary translator (BT), in accordance with embodiments of the present disclosure. Embodiments of a BT may also provide additional features, such as total store elimination (TSE) as also described herein. However, these additional features should be understood as nonlimiting examples of additional features that may be selectively included in appropriate embodiments, and should not be understood as mandatory features.

[0162] System 1800 may include any suitable number and kind of elements to perform the operations described herein, including a processor, SoC, integrated circuit, and other mechanism suitable for total store elimination. Furthermore, although specific elements of system 1800 may be described herein as performing a specific function, any suitable portion of 1800 may perform the functionality described herein. For example, system 1800 may include processor 1802. Although processor 1802 is shown and described as an example in FIGURE 18, any suitable mechanism may be used. System 1800 may include any suitable mechanism for total store elimination. In one embodiment, such mechanisms may be implemented in hardware. In another embodiment, such mechanisms may include a memory-mapped address for the configuration. In a further embodiment, such mechanisms may include an instruction for a programmer, compiler, or firmware to configure the processor 1802 to enable total store elimination. In still yet another embodiment, such mechanisms may include distinct instruction flags or indicators for total store elimination.

[0163] Processor 1802 may be implemented fully or in part by the elements described in FIGURES 1–17. Instructions may be received from instruction stream 1804, which

may reside within a memory subsystem of system 1800. Instructions stream 1804 may be included in any suitable portion of processor 1802 or system 1800. In one embodiment, instruction stream 1804A may be included in an SoC, system, or other mechanism. In another embodiment, instruction stream 1804B may be included in a processor, integrated circuit, or other mechanism. Processor 1802 may include a front end 1806 to receive or retrieve instructions from any suitable location, including a cache or memory. Instructions may include instruction stream 1804. Front end 1806 may include a fetcher 1808 to fill the pipeline efficiently with possible instructions to execute. Front end 1806 may include an instruction decoder 1810 to decode an instruction into opcodes for execution, which may determine the meaning, side effects, data required, data consumed, and data to be produced for the instruction. A binary translator 1812 may be used to optimize code. Binary translator 1812 may be located in any suitable portion of processor 1802 or system 1800, and may be implemented using any suitable mechanism for total store elimination. In one embodiment, binary translator 1812A may be included in the front end 1806 of processor 1802. In another embodiment, binary translator 1812B may be included in an SoC, system, or other mechanism. In a further embodiment, binary translator 1812 may reside in programmable circuitry communicatively coupled to processor 1802 or system 1800. In yet another embodiment, binary translator 1812 may reside in memory in the form of instructions capable of execution by processor 1802.

[0164] Binary translator 1812 may optimize the instructions by eliminating redundant stores while maintaining total store order (TSO). TSO enables other threads or processors accessing data modified by a processor supporting TSO may expect a defined behavior. TSO guarantees that stores to memory will be committed and drained to the cache of the processor 1802 or system 1800 in the original program order. Instructions optimized by binary translator 1812 may be reside in any suitable portion of system 1800. In one embodiment, the optimized instructions may be stored back in the instruction stream 1804. In another embodiment, the optimized instructions may be stored in a data or instruction cache, including an L1, L2, L3, and L4 cache. In a further embodiment, the optimized instructions may be stored in memory. In still yet another embodiment, the optimized instructions may be stored in a dedicated storage device, including a dedicated memory, data array, or system cache. By keeping the optimized instructions separate from the non-

optimized instructions the binary translator, in some embodiments, may enable the processor to abandon all speculative execution of the optimized instructions and restart execution of the non-optimized instructions without any loss of data.

[0165] The decoded instructions may be passed to an out-of-order or in-order execution in an execution pipeline 1816. Execution pipeline 1816 may include a rename and allocate unit 1818 for renaming instructions for out-of-order execution, and a reorder buffer (ROB) coextensive with a retirement unit 1824 so that instructions may appear to be retired in the order that they were received. Rename and allocate unit 1818 may further rename or allocate resources for execution of instructions in parallel. Scheduler 1820 may schedule or allocate instructions to execute on execution units 1822 when inputs are available. Outputs of execution units 1822 may queue in the ROB 1824. Front end 1806 may attempt to anticipate any behaviors that will prevent instructions from executing in a sequential stream and may fetch streams of instructions that might execute. When there is, for example, a mis-prediction of a branch, the ROB may inform the front-end and a different set of instructions might be executed instead. Front end 1806 may store data such as metadata for branch prediction. The instructions may be retired as-if they were executed in-order. Various portions of such execution pipelining may be performed by one or more cores 1814. Each core 1814 may include one or more threads or logical cores for execution.

[0166] In one embodiment, core(s) 1814 may include a data cache unit (DCU). Even though a single DCU 1830 is shown in FIGURE 18, core(s) 1814 may include any suitable number and kind of DCUs. DCU 1830 may perform data caching within a core of processor 1802. DCU 1830 may be communicatively coupled with any suitable number or kind of other cache hierarchies or memory on system 1800 or processor 1802. For example, DCU 1830 may be communicatively coupled to an LLC, an L2 cache, an L3 cache, or a memory controller. DCU 1830 may include any suitable number and kind of elements. For example, DCU 1830 may include a data cache for storing data in lines of cache.

[0167] Commit actions for stores or loads may be achieved by performing the underlying memory operations on memory, registers, or caches, such as DCU 1830. When operations within a memory ordered buffer (MOB) 1826 have been retired, committed, and are ready to send to cache, memory, or other portions of system 1800, they may be removed from MOB 1826. In the example of FIGURE 18, they may be drained to DCU 1830. MOB 1826

may interface with the execution pipeline 1816 and other portions of core(s) 1814 to provide buffering of data. MOB 1826 may be loaded from any suitable source, including DCU 1830, and may be drained 1828 to any suitable destination, including DCU 1830.

[0168] Depending upon the address of the load operation or store operation and the state of the contents of DCU 1830, DCU 1830 may need to interface with other caches. For example, if the address associated with a load operation is not present in DCU 1830, DCU 1830 may need to interface with another cache of a given core, another cache of another core within processor 1802, or another cache of another system 1800.

[0169] DCU 1830 may hold or block some or all snoops. The holding or blocking of snoops may delay the response to a snoop until the data in the DCU is updated and the speculative region of instructions with redundant stores has completed. In one embodiment, all snoops may be held during a speculative region. In another embodiment, snoops with addresses matching an address within the speculative region may be held. DCU 1830 may create track the addresses accessed in the speculative region. This set of addresses may be referred to as the write set. In some cases, the addresses may be known in advance because they are literals or variables that may be resolved in advance. In other cases, the addresses may not be known in advance and knowledge of the address or addresses may be determined after the first execution of the instruction. In a further embodiment, snoops with addresses matching an address similar to, but not exactly the same as an address in the write set may be held. The similar address may be on the same or similar cache line as the write set address.

[0170] System 1800 may support atomic or speculative regions of instructions, as discussed above. Furthermore, system 1800 may support out-of-order operation, as discussed above. The use of out-of-order operation may lead to potential errors. For example, other threads and/or cores may be deadlocked waiting for snoops from each other. The snoops may not process due to a speculative region. Thus, a speculative region may be abandoned and speculative results may be undone at a later time if an error is detected.

[0171] System 1800 may avoid deadlocks by measuring the snoop response. The measurement may be a raw measurement or an average to compensate for jitter in the response. The response may be represented by the number of snoops being held or the response time for a snoop to be handled by DU 1830. The measurement may be performed

using any suitable mechanism, including a performance counter for the snoop response count, the snoop response time, or the number of snoop requests being held in a snoop queue. The measurement may be compared to a response threshold to determine whether the measured response time exceeds the threshold. The threshold may be any suitable value and may be tuned to a particular system. The threshold may correspond to any suitable mechanism for the measurement. System 1800 may raise an interrupt based on the determination that the measured response time is greater than the threshold. Binary translator 1812 may process the interrupt and may restore the processor 1802 to an architectural state stored before the execution of the speculative region with “total store elimination” (TSE). Binary translator 1812 may then direct processor to execute the original instruction stream without redundant stores removed. The direction may be achieved by reverting to instruction stream 1805. Although various operations are described in this disclosure as performed by specific components of processor 1802, the functionality may be performed by any suitable portion of processor 1802.

[0172] FIGURE 19 illustrates a logical representation of elements of system 1900 with total store elimination (TSE), in accordance with embodiments of the present disclosure. System 1900 may execute and retire instructions 1902 into a store buffer 1904 without writing or committing the data to the DCU 1908. Instructions or instruction stream 1902 may be similar to optimized instruction stream 1813. However, optimized instruction stream 1813 may include any instructions suitable for processing in system 1800 and instruction stream 1902 may include any instructions suitable for processing in system 1900. Instructions 1902 may reflect two instructions which may be removed from the instruction stream and include seven instructions that may be executed. The two instructions that may be removed may be stores to memory address D and A (ST [D],0 and ST [A], 0). The seven instructions that may be executed may be stores to memory addresses B, C, E, D, F, A, and G (ST [D], 0; ST .BR {1} [B], 1; ST [A], 0; ST .BR {2} [C], 1; ST {1, 2} [E], 1; ST .ER {1} [D], 1; ST {2} [F], 1; ST .ER {2} [A], 1; and ST [G], 1).

[0173] Store buffer 1904 may include multiple entries. Each entry may include a TSE tag (TSE TAG), a write set identifier (Write Set ID), a retired flag (RETIRED?), a program order identifier (P.O.), and a speculative flag (Speculative?). The TSE tag field may indicate whether an entry is associated with the boundary of a speculative region with redundant stores

removed. BR may indicate the start of such a region and ER may indicate the end of such a region. The write set ID field may be used to associate a store with one or more write sets. An entry in the store buffer 1904, such as entry 1922 may not have a TSE tag but may include a write set ID, which may indicate that the store is within a speculative region rather than at the boundary of a speculative region. An entry may use the retired flag to determine whether the operation has been retired from execution. Store buffer 1904 may use the P.O. designator to determine the order in which data should be retired. In one embodiment, store buffer 1904 may drain senior stores to the DCU that have the lowest P.O. designator. In another embodiment, an instruction with a program order less than the program order of another instruction should be retired first. An entry in the store buffer may use the Speculative? flag to show whether the operation resides within a speculative region of instructions. An entry may have both the speculative flag set and have a defined write set ID. The store buffer may also be able to determine whether the speculative field is set dynamically by looking at the write set ID field. If the write set ID field is set to any value, then the speculative field may also be set. The write set ID field may include one bit for each write set to indicate whether the entry relates to the given write set. For example, entry 1918 may have a write set ID field of 1100, which may indicate that address of the store is in write sets 1 and 2, but not write sets 3 and 4.

[0174] Because TSE may remove redundant stores from the instruction stream, instructions 1902 and store buffer 1904 may not include the redundant operations. Instructions 1902 may be similar to any set of instructions with TSE. Instructions 1902 may be similar to instructions 1813. Although seven instructions are shown, any number of instructions suitable for TSE may be used.

[0175] As the instructions are retired from execution, the operations may be stored in the store buffer 1904. The first entry 1914 in the store buffer 1904 may be associated with a store to memory location B (ST [B]). This first entry may be the second operation in program order. Moreover, the first entry may have a BR TSE tag corresponding to the beginning of a speculative region with reordered stores, and may have a write set ID of {1}, which may correspond to the first write set. Because the first entry may be within a speculative region, the speculative field may be set.

[0176] The second entry 1916 in the store buffer may be associated with a store to memory location C, ST [C]. The second entry may be the fourth operation in program order. Moreover, the second entry may have a BR TSE tag corresponding to the beginning of a speculative region with reordered stores, and may have a write set ID of {2}, which may correspond to the second write set. Because the second entry may be within a speculative region, the speculative field may be set.

[0177] The third entry 1918 in the store buffer may be associated with a store to memory location E, ST [E]. The third entry may be the fifth operation in program order. The third entry may have no TSE tag because no speculative region begins or ends with the operation. The write set ID may be set to both the first write set, {1}, and the second write set, {2}. In one embodiment, the binary translator 1812 may provide these write set identifiers. In another embodiment, the front end 1806 of the processor 1802 may provide these identifiers based on the determination that the instruction is within one or more speculative regions. Because the third entry may be within a speculative region, the speculative field may be set.

[0178] The fourth entry 1920 in the store buffer may be associated with a store to memory location D, ST [D]. The fourth entry may be the sixth operation in program order. The fourth entry may have an ER TSE tag corresponding to the end of a speculative region with reordered stores, and may have a write set ID of {1}, which may correspond to the first write set. Because the fourth entry may be within a speculative region, the speculative field may be set.

[0179] The fifth entry 1922 in the store buffer may be associated with a store to memory location F, ST [F]. The fifth entry may be the seventh operation in program order. The fifth entry may have no TSE tag because no speculative region begins or ends with the operation. The write set ID may be set to the second write set, {2}. In one embodiment, the binary translator 1812 may provide the write set ID. In another embodiment, the front end 1806 of the processor 1802 may provide these identifiers based on the determination that the instruction is within a speculative region. Because the fifth entry may be within a speculative region, the speculative field may be set.

[0180] The sixth entry 1924 in the store buffer may be associated with a store to memory location A, ST [A]. The sixth entry may be the eighth operation in program order.

The sixth entry may have an ER TSE tag corresponding to the end of a speculative region with reordered stores, and may have a write set ID of {2}, which may correspond to the second write set. Because the sixth entry may be within a speculative region, the speculative field may be set.

[0181] The seventh entry 1926 in the store buffer may be associated with a store to memory location G, ST [G]. The seventh entry may be the ninth operation in program order. The seventh entry may have no TSE tag and may have no write set ID. Moreover, because the seventh entry may be outside of any speculative region, the speculative field may not be set.

[0182] Although eight entries in the store buffer are shown, the store buffer may include any number of entries suitable for TSE. Store buffer 1904 may contain similar entries for other instructions. Entries in the store buffer may be drained or written into the DCU 1908 at some time after execution.

[0183] DCU 1908 may include data drained from store buffer 1904 and may contain a snoop queue 1910 to handle snoops 1906. In some embodiments, DCU 1908 may include or have access to a write set validity register 1934, may include a write set ID field 1912A for each cache line, and/or may include a write set tracker 1912B. During the processing of instructions 1902, one or more write sets may be valid or active. DCU 1908 may hold snoops in any suitable manner. In one embodiment, DCU 1908 may hold all snoops to any address while any speculative region containing TSE is being processed. Write set validity register 1934 may be a single bit field, which may indicate whether any speculative region is active or valid. DCU 1908 may hold any snoop received by placing or pushing the snoop onto the snoop queue 1910 while the write set validity register 1934 is set.

[0184] In other embodiments, DCU 1908 may hold snoops that match an address in a write set and may process snoops that do not match. DCU 1908 may hold matched snoops in any suitable manner. In one embodiment, DCU 1908 may combine all write sets into one common write set. In this case, DCU 1908 may hold snoops corresponding to any address in the common write set while any region containing TSE is being processed. In another embodiment, DCU 1908 may distinguish between the addresses in each write set and may hold only those snoops corresponding to an address within an active write set where the speculative region containing TSE is still being processed. In this case, DCU 1908 may respond to snoops more quickly than a DCU that implements a common write set or universal snoop

blocking in speculative regions. DCU 1908 may determine the addresses in a write set in any suitable manner. In one embodiment, DCU 1908 may only handle memory locations or addresses that are resolved. Other structures, such as the store buffer 1904 or a translation table, may resolve the memory locations or address. Alternatively, all memory locations or addresses may be fixed before run-time of the instructions or may be defined as literals. In this case, no memory location or address will require resolution.

[0185] In another embodiment, DCU 1908 may determine one or more addresses in a write set by tracking the memory locations that are resolved after one iteration of execution. Until all the memory locations or addresses are known, DCU 1908 may hold all snoops to avoid violating TSO. After the memory locations or addresses are known, DCU 1908 may respond to snoops more quickly in subsequent iterations of a loop of instructions than in the first iteration. Moreover, DCU 1908 may associate more than one cache line with each resolved memory location or address to hold snoops to adjacent addresses, which may hold snoops to memory locations or addresses that are more likely to be used in subsequent iterations of a loop of instructions. For example, instructions for processing streaming media may iterate across a data structure spanning several cache lines. In this case, associating more than one cache line with an address in a write set may permit DCU 1908 to hold a snoop to an address that may be used in a subsequent iteration, even though that address may be different than the address in the write set.

[0186] DCU 1908 may receive snoops 1906 from other threads or cores in system 1900. For example, DCU 1908 may receive four snoops: LD [D], LD [C], LD [E], and LD [G]. Although four snoops are described, any number of snoops may be received and processed as is suitable for system 1900. After receiving a snoop, DCU 1908 may evaluate the snoop to determine whether the address of the snoop matches one or more valid or active write sets. Matching may occur in any suitable manner. In one embodiment, DCU 1908 may include additional information for each cache line corresponding to a write set ID field 1912A. The write set ID field may include one bit for each write set to indicate whether the cache line is used in more than one write set. For example, eight write sets may correspond to eight bits for the write set ID field. A value of zero for the write set ID field, in this case, may indicate that the cache line is not associated with any write set. Although seven entries in the cache are shown, any suitable number of entries may be used. Moreover, multiple stores may

correspond to the same cache line, separate cache lines, or any combination of cache lines. DCU 1908 may determine the one or more cache lines corresponding to the address of the snoop and may look at the write set ID field 1912A to determine whether to hold the snoop. If the matching cache line includes a write set ID, DCU 1908 will hold the snoop until the write set ID is no longer active or valid. DCU 1908 may track the validity of write set IDs in any suitable manner, including an additional field for each cache line, circuitry for adding or removing the write set ID values for cache lines when speculative regions begin or end, and a write set validity register 1934.

[0187] In another embodiment, DCU 1908 may include a separate structure 1912B, such as a look up table, to track the write sets. Write set tracker 1912B may include a write set ID, a valid field, and an address field. The write set tracker 1912B may be addressed using the address of the memory location. A matching address in the tracker may include one or more write set IDs and a valid field. A valid field may be set if any write set ID is still valid or active. In one embodiment, the write set tracker 1912B may have separate entries for each write set associated with an address. In another embodiment, the write set tracker 1912B may have one entry for each address with one or more write set IDs in the write set ID field. For example, eight write sets may correspond to eight bits in the write set ID field. A value of zero for the write set ID field, in this case, may indicate that the cache line is not associated with any write set.

[0188] DCU 1908 may determine whether the address of a snoop matches a write set in any suitable manner. In one embodiment, an additional entry for the write set ID field 1912A may be implemented for each cache line. If the number of write sets is limited to a common write set, the write set ID field may be indicate only whether the cache line is used in a write set or not. A given write set may correspond to one or more cache lines. Write set {2}, for instance, may correspond to five cache lines, even though the write set only corresponds to four memory locations: [A], [B], [E], and [F]. Write set {1} may correspond to three cache lines, despite the write set containing the same number of memory locations: [C], [D], and [E]. In another embodiment, the write sets may be tracked with a write set tracker 1912B. Write set tracker 1912B may be indexed by the address and may indicate one or more write sets associated with the address and whether the one or more write sets is valid. DCU 1908 may hold snoops with an address that is present in the write set tracker 1912B with

a corresponding valid field set. In a further embodiment, the write sets may be tracked with a bloom filter (not shown). The bloom filter may be used to determine whether the address of a snoop is within the write set.

[0189] DCU 1908 may place a snoop to any write set in a snoop queue 1910 in any suitable manner. In one embodiment, DCU 1908 may determine whether the address of the snoop matches one or more write sets. In another embodiment, DCU 1908 may determine whether the snoop matches any write at least one write set is valid. Once a snoop is placed in the snoop queue 1910, it may remain in the queue until any suitable event occurs, including the ending of the speculative region corresponding to a particular write set or the ending of all speculative regions with TSE.

[0190] DCU 1908 may receive four snoops 1906: LD [D], LD [C], LD [E], and LD[G]. DCU 1908 may determine that the address of the snoop matches one or more write sets. For example, DCU 1908 may determine that address [D] is included in write set {1}. In one embodiment, DCU 1908 may locate the cache line associated with address [D] and then determine whether the cache line is associated with a write set by the write set ID field 1912A. In another embodiment, DCU 1908 may decode address [D] and determine whether the address has a valid write set in write set tracker 1912B. In the example, snoop queue 1910 may include two additional snoops: LD [C] and LD [E]. These snoops may correspond to write set {2} and {1,2}. Snoop queue 1910 may track whether to hold a given snoop based on whether the region with TSE corresponding to the write set is still in execution. For example, snoop LD [D] may remain in the snoop queue until the region with TSE corresponding to write set {2} has ended and snoop LD [E] may remain in the snoop queue until the two regions with TSE corresponding to write sets {1} and {2} have ended. When the regions end for each held snoop, the hold may be removed and DCU 1908 may respond to the held snoop with the dirty or modified status of the data. Thus, holding snoops may ensure that system 1900 maintains total store order (TSO) where no instructions ever appear to have executed out-of-order.

[0191] Some snoops may be held in a snoop queue 1910. For example, the snoop to load data from the address at G (LD [G]) may not match a write set because the write to address G (ST [G]) may be located later in program order than speculative regions with TSE. More specifically, ST [G] may be later in the instruction stream 1902 than ST .ER {2} [A], 1, and may not have any information related to TSE set in its entry 1926 the store buffer 1904. For

example, entry 1926 may not have the TSE tag or write set ID set, and may have the speculative flag unset (N). LD [D], however, may correspond to a memory location written to in write set {1}. Similarly, LD [C] may correspond to a memory location written to in write set {2}. Snoop queue 1910 may also specify that a memory location is included in more than one write set. For example, LD [E] may correspond to both write set {1} and write set {2}. The write set ID field may include eight bits, one bit for each write set. In this case, LD [E] may have a write set ID field value of 00000011, indicating that the snoop implicates both write set 1 and write set 2. The write set ID field, however, may use any mechanism suitable for tracking one or more write sets, including a list of write sets and a compressed representation of the write sets.

[0192] FIGURE 20 is a block diagram illustration of register mapping according to embodiments of this specification. In this illustration, architectural registers 2002 include a number of registers that a programmer or compiler may access by name (such as EAX (32-bit) or RAX (64-bit)). These are identified in this figure as $R_0 - R_n$. As explained above, to enable out-of-order execution, architectural registers may be mapped to a bank of physical registers at runtime. These physical registers are included in physical register file 2004, which here includes registers $P_0 - P_K$. In one common practice, $K > N$ so that multiple out-of-order instructions can use the same architectural register without conflict. By way of illustration, in this example, architectural register R_0 is mapped to physical register P_2 . Register R_1 is mapped to P_0 . Finally, register R_2 is mapped to a physical register P_M somewhere near the middle of physical register file 2004. Of interest to embodiments of this specification is determining when physical registers in physical register file 2004 may be safely reclaimed so that they can be used for later reassignment.

[0193] FIGURE 21 is a block diagram representation of code life cycle according to embodiments of the disclosure. The block diagram of FIGURE 21 includes only selected elements for the sake of simplicity and of focusing on certain novel features of the present specification, such as early register reclamation via reclamation hints. Note that in some embodiments, additional elements may be present, and in others, not all elements presented here are necessary. Rather, this illustration is presented to illustrate certain advantages achieved by early register reclamation.

[0194] In this example, a programmer 2102 writes source code 2104. In certain embodiments, programmer 2102 may place explicit reclamation hints in source code 2104. This is especially true if programmer 2102 is programming in a closer-to-hardware language such as C or x86 assembly.

[0195] Compiler 2106 compiles source code 2104 into object code, and ultimately into executable 2108. In certain embodiments, compiler 2106 could also place explicit reclamation hints into the object code, so that executable 2108 has pre-built reclamation hints. One advantage to this method is that compiler 2106 does not need to execute in real-time (like BT 2110 may need to do), so it can take the time to perform a more in-depth code analysis than BT 2110. There is a tradeoff however. Compiler 2108 may not have the same real-time view of actual hardware usage that BT 2110 has. Thus, in some embodiments, compiler 2106 and BT 2110 may work cooperatively, with compiler 2106 providing reclamation hints where possible, and BT 2110 filling in the gaps as it is able.

[0196] BT 2110 has a view of two versions of the code. Original executable 2112 is the original, non-optimized native code of executable 2108, and these two blocks may in some embodiments be a single block. Native code 2112 may be stored, for example, in transactional cache (“tcache”) of main memory. Translated executable 2114 is an optimized version of original executable 2112, and in some embodiments may be stored in faster on-chip “icache.” Icache may generally contain optimized pre-fetched code blocks that are selected according to recently executed instructions. BT 2110 may insert reclamation hints into translated executable 2114. When executing code, processor 2116 may first look to see whether the instructions are cached in icache. If the code can be successfully executed from icache, it is. Otherwise, processor 2116 may go out to main memory and retrieve the non-optimized original executable 2112. Note that this operation may be necessary not only in the case of icache misses, but also if the cached optimized instructions fail to execute properly, such as if a reclamation hint erroneously marks a register as “dead,” and the register turns out to be needed later. In that case, rather than failing catastrophically, processor 2116 may simply fetch the non-optimized code and try again.

[0197] In an example code scenario, R1 is redefined in blocks B1 and B2. Thus, the physical register allocated to R1 can be reclaimed immediately after the execution of the load instruction in B0. However, certain architectures that lack early register reclamation may not

reclaim R1 until it is redefined (e.g., one of the two R1 redefinition instructions in B1 or B2). In the following table, the mapping to the specific physical register is omitted, to better focus on the work of the BT in marking when a register may be reclaimed. When a processor is provided with BT, the BT may be programmed to provide inline reclamation hints that indicate when a register is ready to be reclaimed. In an embodiment, these hints are evaluated at each branching operation or other change in flow control (e.g., branch, branch not equal, branch equal, jump, jump conditional, to name just a few). For ease of reference, all operations that result in a change in control flow are referred to herein as “branching operations.”

[0198] Consider the following code scenario:

Before	After
B0: R2 = LD [R1] ... JCC B2 B1: ... R3 = R2 + 1 R1 = R0 + 4 JZ B3 B2: ... R4 = LD [R5 + 4] R1 = R0 + 8 ... B3: ...	B0: R2 = LD [R1] ... JCC B2 B1: [R1] ... R3 = R2 + 1 R1 = R0 + 4 JZ B3 B2: [R1] ... R4 = LD [R5 + 4] R1 = R0 + 8 ... B3: ...

Table 2. Register Reclamation

[0199] Note that at the label of the target of each branching operation, there is a reclamation hint that indicates which registers can be reclaimed. The BT becomes aware of the redefinition of registers during its register allocation phase. BT can then annotate this information to blocks in the form of a new instruction, or as immediate values to BT-specific instructions. The information provided to hardware indicates, in the form of a bit-vector or similar, the registers that will be redefined before further reuse (such as R1). This is

annotated as “[R1]” for B1 and B2 in Table 2. This indicates to hardware that the physical register allocated to R1 can be reclaimed for power saving.

[0200] The BT-based system described here realizes advantage over some hardware based early register reclamation mechanisms. For example, a BT-based scheme may have a much broader view of the code than a hardware-only system has. In certain embodiments, BT can analyze liveness for the entire translation, from checkpoint to checkpoint. A BT-based mechanism may also be able to make more appropriate decisions when behavior varies across a branch. For example:

Before	After
B0: R2 = LD [R1] ... JCC B2 B1: R1 = R0 + 4 R3 = R2 + 1 JZ B3 B2: R4 = LD [R1 + 4] ... B3: ...	B0: R2 = LD [R1] ... JCC B2 B1: [R1] R1 = R0 + 4 R3 = R2 + 1 JZ B3 B2: R4 = LD [R1 + 4] ... B3 ...

[0201] In this example, R1 is redefined in B1, but then reused in B2. Thus, if the path through B1 is taken, the register can be reclaimed. But if the B2 path is taken, R1 should not be reclaimed yet. BT can provide this specific information to each block accurately.

[0202] In some embodiments, the analysis of whether a register is “dead” need not be wholly determinative. Note that the foregoing code blocks are trivially small, and so it is trivial to analyze them to determine whether R1 is definitively “dead.” But real-world code blocks may be very large. In some cases, analyzing a very large code block may be prohibitively expensive and may cancel out any advantages realized from early register reclamation. Thus, in some embodiments, rather than determinative register reclamation analysis, a more fuzzy analysis may be performed, such as a heuristic or other predictive but non-determinative analysis. In that case, it is anticipated that a certain number of mispredictions will occur. In the case of a mispredictions, it is not necessary for a code fault,

exception, or other error to occur. Rather, program control can revert back to the non-optimized codepage for a “do over.” This may represent, in certain embodiments, an acceptable tradeoff between code optimization and strict correctness of prediction.

[0203] After making a determination on register liveness, the BT may provide the reclamation hint to the processor using a special-purpose hardware instruction at the beginning of each block. In another embodiment, the BT may provide reclamation hints after the previous block using a bit-vector. If there is not enough space in the instruction to provide information about all registers (for example, modern processors may have 48 to 64 registers, requiring up to six bits to provide reclamation hints), as a compromise information on only the higher registers may be passed, on the assumption that the lower registers are used more actively, and thus more likely to be live at a given time. In cases where the blocks are small, the BT may aggregate reclamation hints for several blocks. Again, this may be provided in a special-purpose hardware instruction to avoid code bloat. The BT overhead of adding such an optimization may be minimal, as most of the information is already collected during register allocation.

[0204] FIGURE 22 is a flow chart of an example method 2200 of optimizing a program according to embodiments of this specification. In block 2202, BT 2110 begins its binary translation. This may be, for example, a method by which complex instruction set computing (CISC) are translated into reduced instruction set computing (RISC)-like microcode instructions. For additional discussion of the various functions of BT 2110, reference is made to FIGURE 18 herein, and the accompanying discussion beginning in paragraph [0161], and in other disclosures throughout this specification.

[0205] In block 2204, BT 2110 constructs a control flow graph (CFG). A CFG is a representation of the potential paths that the program flow may take according to various inputs. The CFG enables BT 2110 to analyze the branching instructions, and the resulting potentials for register reclamation.

[0206] In block 2206, BT 2110 identifies flow control changes. In one embodiment, these are candidates for analyzing early register reclamation. Note however that using flow control changes as “sign posts” for determining register reclamation is disclosed as one example only. In general, any suitable method of identifying reclaimable registers may be used.

[0207] In block 2208, BT 2110 identifies reclaimable registers. These are registers that BT 2110 has determined may be safely reclaimed without negatively affecting program operation. In one embodiment, BT 2110 performs a determinative analysis on registers, and inserts a reclamation hint only if it can definitively determine that the register is no longer needed. In another example, “fuzzy” logic is used, such as heuristic analysis, or partial analysis of a larger or more complex block. In that case, if BT 2110 determines within a threshold of confidence that the register can be reclaimed, it inserts a reclamation hint. If a failure is encountered because of an erroneous reclamation hint, the original, non-optimized code can be executed again to obtain a correct result.

[0208] In block 2210, BT 2110 inserts the reclamation hints for registers it has determined are reclaimable.

[0209] In block 2299, the method is done.

[0210] FIGURE 23 is a flow chart of an example method performed by a processor according to embodiments of this disclosure.

[0211] In block 2302, during program execution, processor 2116 assigns a physical register to an architectural register, as described in this specification.

[0212] In block 2304, processor 2116 begins to execute the code, assigning values to the register according to the instructions it encounters, and using those values as necessary.

[0213] In decision block 2306, processor 2116 determines whether there is a reclamation hint for the register. This may commonly be encountered at the retirement of a branching instruction, though in a general sense, reclamation hints may be inserted at any suitable point.

[0214] In block 2310, if there is no reclamation hint, then processor 2116 proceeds to reclaim the physical register according to its normal operating procedure, such as upon expiration of an instruction that writes a new value to the architectural register.

[0215] Returning to block 2306, if there is a reclamation hint, then in block 2308, processor 2116 reclaims the physical register early, such as upon encountering the reclamation hint.

[0216] In decision block 2312, processor 2116 determines whether the code block successfully executed. This may be particularly important in embodiments where reclamation is not fully deterministic.

[0217] In block 2314, if the optimized code execution was not successful, then processor 2116 may go back and retrieve the original, non-optimized code and execute that. Note that in other embodiments, an exception or other error condition could also be thrown.

[0218] In block 2399, the method is done.

[0219] Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the disclosure may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

[0220] Program code may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system may include any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

[0221] The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

[0222] One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine-readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

[0223] Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks,

optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[0224] Accordingly, embodiments of the disclosure may also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

[0225] In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part-on and part-off processor.

[0226] Thus, techniques for performing one or more instructions according to at least one embodiment are disclosed. While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on other embodiments, and that such embodiments not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure. In an area of technology such as this, where growth is fast and further advancements are not easily foreseen, the disclosed embodiments may be readily modifiable in arrangement and detail as facilitated by enabling technological advancements without departing from the principles of the present disclosure or the scope of the accompanying claims.

[0227] In some embodiments of the present disclosure, a processor may include an execution unit, a front end, and an allocator. The front end may include circuitry to decode a first instruction to set a performance register for the execution unit and a second instruction. The performance register may include a mode field to specify a mode of operation for the execution unit. The allocator may include circuitry to assign the second instruction to the execution unit to execute the second instruction. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to select between a normal computation, which may correspond to a normal result, and an accelerated computation, which may correspond to an accelerated result, based on the mode field of the performance register. The accelerated computation may be faster than the normal computation. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to perform the selected computation. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to select between the normal result and the accelerated result based on the mode field of the performance register.

[0228] In combination with any of the above embodiments, in an embodiment the second instruction may correspond to a multiplication operation and the normal computation may include circuitry to initialize a normal accumulation and a partial product counter. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine a total number of partial products required based on input operands of the second instruction. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to compute a partial product of the input operands. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to increment the partial product counter, may add the normal accumulation to the computed partial product, and may store the result of the addition back in the normal accumulation. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine whether to compute an additional partial product based on the partial product counter and the total number of partial products required. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to store the normal

accumulation in the normal result based on the determination that the additional partial product need not be computed.

[0229] In combination with any of the above embodiments, in an embodiment the second instruction may correspond to a multiplication operation and the accelerated computation may include circuitry to compute all partial products of input operands of the second instruction. In combination with any of the above embodiments, in an embodiment the accelerated computation may include circuitry to compute the accelerated result based on a sum of all the computed partial products. In combination with any of the above embodiments, in an embodiment the circuitry to determine the total number of partial products required, may be further based on an accuracy field of the performance register, and the normal computation may further include circuitry to order terms of the input operands for the circuitry to compute the partial product of the input operands. The order may prioritize a most significant term of each input operand. In combination with any of the above embodiments, in an embodiment the circuitry to determine the total number of partial products required, may be further based on a reduced size field of the performance register, and the normal computation may further include circuitry to order terms of the input operands for the circuitry to compute the partial product of the input operands. The order may prioritize a least significant term of each input operand. In combination with any of the above embodiments, in an embodiment the processor may include circuitry to monitor energy consumed, and circuitry to compare the energy consumed to a power threshold and may set the mode field of the performance register based on the comparison. In combination with any of the above embodiments, in an embodiment the accuracy field of the performance register may define a relative percentage of accuracy for the execution unit and circuitry to determine the total number of partial products required may be further based on a product of the total number of partial products required and the relative percentage of accuracy.

[0230] In some of the present embodiments, a method may include, setting a mode field of a performance register associated with an arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the method may include decoding an instruction for the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the method may include selecting between a normal computation

corresponding to a normal result and an accelerated computation corresponding to an accelerated result based on the mode field of the performance register. The accelerated computation may be faster than the normal computation. In combination with any of the above embodiments, in an embodiment the method may include performing the selected computation. In combination with any of the above embodiments, in an embodiment the method may include selecting between the normal result and the accelerated result based on the mode field of the performance register.

[0231] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include initializing a normal accumulation and a partial product counter. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include determining a total number of partial products required based on input operands of the decoded instruction. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include computing a partial product of the input operands. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include adding the normal accumulation to the computed partial product and storing the result of the addition back in the normal accumulation. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include incrementing the partial product counter. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include comparing the partial product counter to the total number of partial products required. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include determining whether to compute an additional partial product based on the comparison. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include

storing the normal accumulation in the normal result based on the determination not to compute the additional partial product.

[0232] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the accelerated computation may include computing all partial products of input operands of the decoded instruction. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the accelerated computation may include computing a sum of all partial products. In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the accelerated computation may include storing the sum in the accelerated result. In combination with any of the above embodiments, in an embodiment determining the total number of partial products required may be further based on an accuracy field of the performance register. In combination with any of the above embodiments, in an embodiment computing the partial product of the input operands may further include prioritizing a most significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment determining the total number of partial products required may be further based on a reduced size field of the performance register. In combination with any of the above embodiments, in an embodiment computing the partial product of the input operands may further include prioritizing a least significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment the method may include monitoring energy consumed by the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the method may include comparing the energy consumed to a power threshold. In combination with any of the above embodiments, in an embodiment the method may include setting the mode field of the performance register based on the comparison. In combination with any of the above embodiments, in an embodiment the accuracy field of the performance register may define a relative percentage of accuracy for the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment determining the total number of partial products required may be further based on a product of the total number of partial products required and the relative percentage of accuracy.

[0233] In some embodiments of the present disclosure, a system may include an execution unit, a front end, and an allocator. The front end may decode a first instruction to set a performance register for the execution unit and a second instruction. The performance register may include circuitry to select between a normal computation, which may correspond to a normal result, and an accelerated computation, which may correspond to an accelerated result, based on the mode field of the performance register. The accelerated computation may be faster than the normal computation. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to perform the selected computation. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to select between the normal result and the accelerated result based on the mode field of the performance register.

[0234] In combination with any of the above embodiments, in an embodiment the second instruction may correspond to a multiplication operation and the normal computation may include circuitry to initialize a normal accumulation and a partial product counter. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine a total number of partial products required based on input operands of the second instruction. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to compute a partial product of the input operands. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to increment the partial product counter, may add the normal accumulation to the computed partial product, and may store the result of the addition back in the normal accumulation. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine whether to compute an additional partial product based on the partial product counter and the total number of partial products required. In combination with any of the above embodiments, in an embodiment the execution unit may include circuitry to store the normal accumulation in the normal result based on the determination that the additional partial product need not be computed.

[0235] In combination with any of the above embodiments, in an embodiment the second instruction may correspond to a multiplication operation and the accelerated computation may include circuitry to compute all partial products of input operands of the

second instruction. In combination with any of the above embodiments, in an embodiment the accelerated computation may include circuitry to compute the accelerated result based on a sum of all the computed partial products. In combination with any of the above embodiments, in an embodiment circuitry to determine the total number of partial products required may be further based on an accuracy field of the performance register, and the normal computation may further include circuitry to order terms of the input operands for the circuitry to compute the partial product of the input operands. The order may prioritize a most significant term of each input operand. In combination with any of the above embodiments, in an embodiment circuitry to determine the total number of partial products required may be further based on a reduced size field of the performance register, and the normal computation may further include circuitry to order terms of the input operands for the circuitry to compute the partial product of the input operands. The order may prioritize a least significant term of each input operand. In combination with any of the above embodiments, in an embodiment the system may include circuitry to monitor energy consumed, compare the energy consumed to a power threshold, and set the mode field of the performance register based on the comparison. In combination with any of the above embodiments, in an embodiment the accuracy field of the performance register may define a relative percentage of accuracy for the execution unit and circuitry to determine the total number of partial products required may be further based on a product of the total number of partial products required and the relative percentage of accuracy.

[0236] In some embodiments of the present disclosure, an arithmetic logic unit may include a performance register including a mode field and four circuits. In combination with any of the above embodiments, in an embodiment the arithmetic logic unit may include circuitry to receive a decoded instruction. In combination with any of the above embodiments, in an embodiment the arithmetic logic unit may include circuitry to select between a normal computation corresponding to a normal result and an accelerated computation corresponding to an accelerated result based on the mode field of the performance register. The accelerated computation may be faster than the normal computation. In combination with any of the above embodiments, in an embodiment the arithmetic logic unit may include circuitry to perform the selected computation. In combination with any of the above embodiments, in an embodiment the arithmetic logic unit

may include circuitry to select between the normal result and the accelerated result based on the mode field of the performance register.

[0237] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the normal computation may include circuitry to initialize a normal accumulation and a partial product counter. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine a total number of partial products required based on input operands of the instruction. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to compute a partial product of the input operands. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to increment the partial product counter, may add the normal accumulation to the computed partial product, and may store the result of the addition back in the normal accumulation. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to determine whether to compute an additional partial product based on the partial product counter and the total number of partial products required. In combination with any of the above embodiments, in an embodiment the normal computation may include circuitry to store the normal accumulation in the normal result based on the determination that the additional partial product need not be computed.

[0238] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and the accelerated computation may include circuitry to compute all partial products of input operands of the instruction. In combination with any of the above embodiments, in an embodiment the accelerated computation may include circuitry to compute the accelerated result based on a sum of all the computed partial products. In combination with any of the above embodiments, in an embodiment the circuitry to determine the total number of partial products required, may be further based on an accuracy field of the performance register, and the circuitry to compute the partial product of the input operands, may further include circuitry to prioritize a most significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment the circuitry to determine the total number of partial products required, may be further based on a reduced

size field of the performance register, and circuitry to compute the partial product of the input operands, may further include circuitry to prioritize a least significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment the accuracy field of the performance register may define a relative percentage of accuracy for the arithmetic logic unit, and the circuitry to determine the total number of partial products required, may be further based on a product of the total number of partial products required and the relative percentage of accuracy.

[0239] In some of the present embodiments, an apparatus may include a means for setting a mode field of a performance register associated with an arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for decoding an instruction for the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for selecting between a normal computation means, which may correspond to a normal result, and an accelerated computation means, which may correspond to an accelerated result, based on the mode field of the performance register. The accelerated computation means may be faster than the normal computation means. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for performing the selected computation means. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for selecting between the normal result and the accelerated result based on the mode field of the performance register.

[0240] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication means and the normal computation means may include a means for initializing a normal accumulation and a partial product counter. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for determining a total number of partial products required based on input operands of the decoded instruction. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for computing a partial product of the input operands. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for adding the normal accumulation to the computed partial product and a means for storing the result of the addition back in the normal accumulation. In combination with any of the above

embodiments, in an embodiment the normal computation means may include a means for incrementing the partial product counter. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for comparing the partial product counter to the total number of partial products required. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for determining whether to compute an additional partial product based on the means for comparison. In combination with any of the above embodiments, in an embodiment the normal computation means may include a means for storing the normal accumulation in the normal result based on the determination not to compute the additional partial product.

[0241] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication means and the accelerated computation means may include a means for computing all partial products of input operands of the decoded instruction. In combination with any of the above embodiments, in an embodiment the accelerated computation means may include a means for computing a sum of all partial products. In combination with any of the above embodiments, in an embodiment the accelerated computation means may include a means for storing the sum in the accelerated result. In combination with any of the above embodiments, in an embodiment the means for determining the total number of partial products required may be further based on an accuracy field of the performance register. In combination with any of the above embodiments, in an embodiment the means for computing the partial product of the input operands may further include a means for prioritizing a most significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment the means for determining the total number of partial products required may be further based on a reduced size field of the performance register. In combination with any of the above embodiments, in an embodiment the means for computing the partial product of the input operands may further include a means for prioritizing a least significant term of each input operand for the partial product. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for monitoring energy consumed by the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for comparing the energy consumed to a

power threshold. In combination with any of the above embodiments, in an embodiment the apparatus may include a means for setting the mode field of the performance register based on the means for comparison. In combination with any of the above embodiments, in an embodiment the accuracy field of the performance register may define a relative percentage of accuracy for the arithmetic logic unit. In combination with any of the above embodiments, in an embodiment the means for determining the total number of partial products required may be further based on a product of the total number of partial products required and the relative percentage of accuracy.

[0242] By way of example, there is disclosed in one example, an apparatus, comprising: a front end including circuitry to decode instructions from an instruction stream; a physical register file; a register mapper to map an architectural register to a physical register in the physical register file; and circuitry and logic to: identify a reclamation hint in an instruction stream; and based at least in part on the reclamation hint, reclaim the physical register.

[0243] There is further disclosed an example, wherein the circuitry and logic are to reclaim the physical register at a specified operation.

[0244] There is further disclosed an example, wherein the reclamation hint is associated with a branching operation.

[0245] There is further disclosed an example, wherein the reclamation hint is at a target of the branching operation.

[0246] There is further disclosed an example, further comprising a binary translator to insert the reclamation hint.

[0247] There is further disclosed an example, wherein the circuitry and logic is also to determine that the reclamation has caused a program failure.

[0248] There is further disclosed an example, wherein the circuitry and logic is to raise an error condition upon the determining.

[0249] There is further disclosed an example, wherein the circuitry and logic is to revert to a non-optimized code block upon the determining.

[0250] There is further disclosed an example, wherein the circuitry and logic is to determine that the reclamation hint has introduced unacceptable overhead, and to revert to a non-optimized code block.

[0251] There is further disclosed an example of an apparatus, comprising: a binary translator (BT) including circuitry to: analyze a code block; determine that an architectural register mapped to a physical register in the physical register file is available for early reclamation; and insert a reclamation hint into the code block.

[0252] There is further disclosed an example, wherein the BT is to insert the reclamation hint in connection with a branching operation

[0253] There is further disclosed an example, wherein the BT is to insert the reclamation hint at a target of a branching operation.

[0254] There is further disclosed an example of a method of early reclamation of a physical, comprising: mapping an architectural register to a physical register in a physical register file; and identifying a reclamation hint in an instruction stream; and based at least in part on the reclamation hint, reclaiming the physical register.

[0255] There is further disclosed an example, further comprising the physical register at a specified operation.

[0256] There is further disclosed an example, wherein the reclamation hint is associated with a branching operation.

[0257] There is further disclosed an example, wherein the reclamation hint is at a target of the branching operation.

[0258] There is further disclosed an example, further comprising operating a binary translator to insert the reclamation hint.

[0259] There is further disclosed an example, further comprising determining that the reclamation has caused a program failure.

[0260] There is further disclosed an example, further comprising an error condition upon the determining.

[0261] There is further disclosed an example, further comprising reverting to a non-optimized code block upon the determining.

[0262] There is further disclosed an example, further comprising determining that the reclamation hint has introduced unacceptable overhead, and reverting to a non-optimized code block.

[0263] There is further disclosed an example of one or more tangible, non-transitory computer-readable storage mediums having stored thereon executable instructions for

instructing one or more processors for providing early register reclamation, comprising performing any or all of the operations of the preceding examples.

[0264] There is further disclosed a method of providing earlier register reclamation, comprising performing any or all of the operations described above.

[0265] There is further disclosed an example of an apparatus comprising means for performing the method.

[0266] There is further disclosed an example wherein the means comprise a processor and a memory.

[0267] There is further disclosed an example wherein the means comprise one or more tangible, non-transitory computer-readable storage mediums.

[0268] There is further disclosed an example wherein the apparatus is a computing device.

Claims

What is claimed is:

1. An apparatus, comprising:
a front end including circuitry to decode instructions from an instruction stream;
a physical register file;
a register mapper to map an architectural register to a physical register in the physical register file; and
circuitry and logic to:
identify a reclamation hint in an instruction stream; and
based at least in part on the reclamation hint, reclaim the physical register.
2. The apparatus of claim 1, wherein the circuitry and logic are to reclaim the physical register at a specified operation.
3. The apparatus of claim 1, wherein the reclamation hint is associated with a branching operation.
4. The apparatus of claim 3, wherein the reclamation hint is at a target of the branching operation.
5. The apparatus of claim 1, further comprising a binary translator to insert the reclamation hint.
6. The apparatus of claim 1, wherein the circuitry and logic is also to determine that the reclamation has caused a program failure.
7. The apparatus of claim 6, wherein the circuitry and logic is to raise an error condition upon the determining.
8. The apparatus of claim 6, wherein the circuitry and logic is to revert to a non-optimized code block upon the determining.

9. The apparatus of any of claims 1 – 8, wherein the circuitry and logic is to determine that the reclamation hint has introduced unacceptable overhead, and to revert to a non-optimized code block.
10. An apparatus, comprising:
 - a binary translator (BT) including circuitry to:
 - analyze a code block;
 - determine that an architectural register mapped to a physical register in the physical register file is available for early reclamation; and
 - insert a reclamation hint into the code block.
11. The apparatus of claim 10, wherein the BT is to insert the reclamation hint in connection with a branching operation
12. The apparatus of claim 11, wherein the BT is to insert the reclamation hint at a target of a branching operation.
13. A method of early reclamation of a physical, comprising:
 - mapping an architectural register to a physical register in a physical register file; and
 - identifying a reclamation hint in an instruction stream; and
 - based at least in part on the reclamation hint, reclaiming the physical register.
14. The method of claim 13, further comprising the physical register at a specified operation.
15. The method of claim 13, wherein the reclamation hint is associated with a branching operation.
16. The method of claim 15, wherein the reclamation hint is at a target of the branching operation.
17. The method of claim 13, further comprising operating a binary translator to insert the reclamation hint.

18. The method of claim 13, further comprising determining that the reclamation has caused a program failure.
19. The method of claim 18, further comprising an error condition upon the determining.
20. The method of claim 18, further comprising reverting to a non-optimized code block upon the determining.
21. The method of any of claims 13 – 20, further comprising determining that the reclamation hint has introduced unacceptable overhead, and reverting to a non-optimized code block.

1/29




FIG. 1A

2/29




FIG. 1B




FIG. 1C

4/29




FIG. 2







FIG. 3A




7/29




UNSIGNED PACKED BYTE REPRESENTATION 344




SIGNED PACKED BYTE REPRESENTATION 345




UNSIGNED PACKED WORD REPRESENTATION 346



SIGNED PACKED WORD REPRESENTATION 347




UNSIGNED PACKED DOUBLEWORD REPRESENTATION 348




SIGNED PACKED DOUBLEWORD REPRESENTATION 349

FIG. 3C




360

FIG. 3D



370

FIG. 3E



380

FIG. 3F




FIG. 4A

10/29





FIG. 4B

11/29

**FIG. 5A**

12/29





FIG. 5B

**FIG. 6**




FIG. 7




FIG. 8







FIG. 9

**FIG. 10**

**FIG. 11****FIG. 12**




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17

1800




FIG. 18

1900






FIG. 19


26/29

*Fig. 20*


27/29

*Fig. 21*

28/29

*Fig. 22*

29/29

*Fig. 23*

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/039577**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/30(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/30; G06F 12/06; G06F 9/38; G06F 9/312; G06F 9/305; G06F 12/08

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: reclamation hint, register, branching operation, renaming, determining, program failure, error, reverting

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012-0005444 A1 (JEFFREY P. RUPLEY et al.) 05 January 2012 See paragraphs [0018]-[0019], [0022]-[0023], [0025]-[0026], [0032]; claims 1, 15; and figures 1-5.	1-5, 10-17
Y		6-9, 18-21
Y	US 2013-0151819 A1 (FREDERIC CLAUDE MARIE PIRY et al.) 13 June 2013 See paragraph [0036]; and claims 8-11.	6-9, 18-21
A	US 2005-0120191 A1 (HAITHAM H. AKKARY et al.) 02 June 2005 See paragraphs [0022]-[0031]; and figures 2-3.	1-21
A	US 2014-0032884 A1 (ANIL KRISHNA et al.) 30 January 2014 See paragraphs [0042]-[0052]; and figures 3-4.	1-21
A	US 2009-0327661 A1 (ZEEV SPERBER et al.) 31 December 2009 See paragraphs [0027]-[0038]; and figure 3.	1-21



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

27 September 2017 (27.09.2017)

Date of mailing of the international search report

28 September 2017 (28.09.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office


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

Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/039577

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0005444 A1	05/01/2012	US 8707015 B2	22/04/2014
US 2013-0151819 A1	13/06/2013	US 9052909 B2	09/06/2015
US 2005-0120191 A1	02/06/2005	None	
US 2014-0032884 A1	30/01/2014	US 9262170 B2	16/02/2016
US 2009-0327661 A1	31/12/2009	None	