



- (51) **International Patent Classification:**
G06F 9/38 (2006.01) *G06F 9/30* (2006.01)
- (21) **International Application Number:**
PCT/US2017/018883
- (22) **International Filing Date:**
22 February 2017 (22.02.2017)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
15/087,854 31 March 2016 (31.03.2016) US
- (71) **Applicant:** INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) **Inventors:** MOHAPATRA, Debabrata; 3600 Juliette Lane SC-12, MailStop - 303, Santa Clara, California 95054 (US). WANG, Perry H.; 15798 Alta Vista Way, San Jose, California 95127 (US). ZOU, Xiang; 2111, N.E 25th Avenue, MS: JF2-04, Hillsboro, Oregon 97124 (US). KIM, Sang Kyun; 3600 Juliette Lane, Santa Clara, California 95134 (US). MATHAIKUTTY, Deepak A.; 2200 Mission college blvd, SC12-303, Santa Clara, California 95054 (US). CHINYA, Gautham N.; 125 SE 49th CT, Hillsboro, Oregon 97123 (US).
- (74) **Agents:** KOMENDA, J. Kyle et al.; Patent Capital Group, 2816 Lago Vista Lane, Rockwall, TX 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, 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).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- 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))

[Continued on next page]

(54) **Title:** INSTRUCTION AND LOGIC FOR CONFIGURABLE ARITHMETIC LOGIC UNIT PIPELINE




FIG. 20

(57) **Abstract:** A processor includes a front end including circuitry to decode a first instruction to set a performance register for an execution unit and a second instruction and an allocator including circuitry to assign the second instruction to the execution unit to execute the second instruction. The execution unit includes circuitry to select between a normal computation and an accelerated computation based on a mode field of the performance register, perform the selected computation, and select between a normal result associated with the normal computation and an accelerated result associated with the accelerated computation based on the mode field.



Published:

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

INSTRUCTION AND LOGIC FOR CONFIGURABLE ARITHMETIC LOGIC UNIT PIPELINE

5 FIELD OF THE INVENTION

[0001] The present disclosure pertains to the field of processing logic, microprocessors, and associated instruction set architecture that, when executed by the processor or other processing logic, perform logical, mathematical, or other functional operations.

10

DESCRIPTION OF RELATED ART

[0002] Multiprocessor systems are becoming more and more common. Applications of multiprocessor systems include performing arithmetic. Applications of multiprocessor systems include dynamic domain partitioning all the way down to
15 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. Instructions as they are received on a processor may be decoded into terms or instruction words that are native, or more native, for execution on the processor. Processors may be implemented
20 in a system on chip. Arithmetic may include shifting, adding, or multiplication. These applications may include image, audio, and video.

DESCRIPTION OF THE FIGURES

[0003] Embodiments are illustrated by way of example and not limitation in the
25 Figures of the accompanying drawings:

[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
30 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;

5 [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;

10 [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;

15 [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;

20 [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;

25 [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;

30 [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 a block diagram of a system with a configurable arithmetic logic unit pipeline, in accordance with embodiments of the present disclosure;

[0031] FIGURE 19 is a logical representation of elements of a system with multiple selectable arithmetic logic unit pipelines, in accordance with embodiments of the present disclosure;

[0032] FIGURE 20 is a block diagram of elements of a system with a configurable arithmetic logic unit pipeline, in accordance with embodiments of the present disclosure;

[0033] FIGURE 21 is a more detailed block diagram of elements of a system with a configurable arithmetic logic unit pipeline, in accordance with embodiments of the present disclosure; and

5 [0034] FIGURE 22 is a diagram of operation of a method for configuring an arithmetic logic unit pipeline, in accordance with embodiments of the present disclosure.

DETAILED DESCRIPTION

[0035] The following description describes an instruction and processing logic for a configurable arithmetic logic unit (ALU) pipeline. The instruction and processing
10 logic may be implemented on an out-of-order processor. In the following description, numerous specific details such as processing logic, processor types, micro-architectural conditions, events, enablement mechanisms, and the like are set forth in order to provide a more thorough understanding of embodiments of the present disclosure. It will be appreciated, however, by one skilled in the art that the embodiments may be
15 practiced without such specific details. Additionally, some well-known structures, circuits, and the like have not been shown in detail to avoid unnecessarily obscuring embodiments of the present disclosure.

[0036] Although the following embodiments are described with reference to a processor, other embodiments are applicable to other types of integrated circuits and
20 logic devices. Similar techniques and teachings of embodiments of the present disclosure may be applied to other types of circuits or semiconductor devices that may benefit from higher pipeline throughput and improved performance. The teachings of embodiments of the present disclosure are applicable to any processor or machine that performs data manipulations. However, the embodiments are not limited to processors
25 or machines that perform 512-bit, 256-bit, 128-bit, 64-bit, 32-bit, or 16-bit data operations and may be applied to any processor and machine in which manipulation or management of data may be performed. In addition, the following description provides examples, and the accompanying drawings show various examples for the purposes of illustration. However, these examples should not be construed in a limiting sense as
30 they are merely intended to provide examples of embodiments of the present disclosure rather than to provide an exhaustive list of all possible implementations of embodiments of the present disclosure.

[0037] 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.

[0038] 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).

[0039] 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
5 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
10 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
15 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
20 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.

[0040] 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
25 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
30 point instructions, load/store operations, data moves, etc.

[0041] 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).

5 **[0042]** 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
10 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
15 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
20 (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.

[0043] An instruction may include one or more instruction formats. In one embodiment, an instruction format may indicate various fields (number of bits, location
25 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
30 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.

[0044] 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.

[0045] 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.).

[0046] 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
5 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
10 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.

[0047] 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
15 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[®],
20 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,
25 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.

[0048] Embodiments are not limited to computer systems. Embodiments of the
30 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.

5 **[0049]** 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
10 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
15 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.

[0050] In one embodiment, processor 102 may include a Level 1 (L1) internal
20 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
25 various registers including integer registers, floating point registers, status registers, and instruction pointer register.

[0051] 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.
30 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
5 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.

[0052] 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
10 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.

[0053] A system logic chip 116 may be coupled to processor bus 110 and
15 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
20 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.

[0054] System 100 may use a proprietary hub interface bus 122 to couple MCH
25 116 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,
30 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.

[0055] 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.

[0056] 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.

[0057] 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.

[0058] 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
5 indicate what operation should be performed on the corresponding data indicated within the instruction.

[0059] 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
10 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
15 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.

[0060] One embodiment of data processing system 140 provides for mobile,
20 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
25 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).

[0061] FIGURE 1C illustrates other embodiments of a data processing system
30 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
5 sufficient detail, may be suitable to facilitate the manufacture of all or part of data processing system 160 including processing core 170.

[0062] 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
10 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.

[0063] 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
20 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
25 received SIMD coprocessor instructions intended for it.

[0064] 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
30 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
5 embodiment.

[0065] 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
10 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,
15 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
20 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
25 needed to complete the operation.

[0066] 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
30 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.

[0067] 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.

[0068] 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.

[0069] 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.

[0070] 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.

[0071] 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.

[0072] 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.

[0073] 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.

[0074] 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.

[0075] 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.

[0076] 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.

[0077] 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.

[0078] 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.

10 **[0079]** 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.

20 **[0080]** 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.

25 **[0081]** 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.

30 **[0082]** 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.

[0083] 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.

[0084] 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 162 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.

[0085] 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.

[0086] 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.

[0087] 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).

[0088] 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.

[0089] 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.

[0090] 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.

[0091] 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.

[0092] 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.

[0093] 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.

[0094] 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.

[0095] 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.

[0096] 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.

[0097] 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.

[0098] 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.

[0099] 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
5 500 and prepare the instructions to be used later in the processor pipeline as they are passed to out-of-order execution engine 580.

[00100] 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
10 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
15 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
20 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
25 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
30 588 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.

[00101] 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.

[00102] 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.

[00103] 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 5 6 with broken lines.

[00104] 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. 10

[00105] 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. 15

[00106] 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. 20

[00107] 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 25 30

asymmetry and heterogeneity amongst processors 610, 615. For at least one embodiment, various processors 610, 615 may reside in the same die package.

[00108] 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.

[00109] 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.

[00110] 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.

[00111] 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.

[00112] 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.

[00113] 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.

[00114] 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.

[00115] 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.

[00116] 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.

[00117] 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.

5 **[00118]** 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
10 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
15 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.

[00119] 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
20 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.

25 **[00120]** 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
30 and circuits may be included in the processor of FIGURE 10, including more CPUs or GPUs and other peripheral interface controllers.

[00121] 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.

[00122] 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.

[00123] 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.

[00124] 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.

15 **[00125]** 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.

[00126] 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.

[00127] 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.

[00128] 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.

[00129] 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.

[00130] 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.

[00131] In one embodiment, memory system 1540 may include an executed instruction pointer 1580. Executed instruction pointer 1580 may store a value

identifying the oldest, undispatched 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
5 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
10 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
15 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.

[00132] In another embodiment, memory system 1540 may include a
20 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.

[00133] Execution entities 1565 may include any suitable number and
25 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
30 collectively form an execution unit.

[00134] 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
5 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
10 the transfer of data between execution entities 1565 and the portions of a system external to instruction architecture 1500.

[00135] 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
15 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
20 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.

25 [00136] 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
30 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.

[00137] The operation of instruction architecture 1500 to execute an
5 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
10 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
15 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.

[00138] Dual instruction decode stage 1550 may translate a received
20 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
25 results may be input into branches 1557.

[00139] 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
30 issue stage 1560.

[00140] 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
5 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
10 execution.

[00141] 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
15 enable additional instructions to be executed. Performance of instruction set architecture 1500 may be monitored or debugged by trace unit 1575.

[00142] 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,
20 instruction architecture 1500 of FIGURE 15.

[00143] 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
25 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
30 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
5 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,
10 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
15 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.

[00144] FIGURE 17 is a block diagram of an electronic device 1700 for utilizing a processor 1710, in accordance with embodiments of the present disclosure.
20 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.

[00145] Electronic device 1700 may include processor 1710
25 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
30 Receiver/Transmitter (UART) bus.

[00146] 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
5 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.

[00147] Furthermore, in various embodiments other components may
10 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
15 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
20 next generation form factor (NGFF).

[00148] Embodiments of the present disclosure involve a configurable ALU pipeline. FIGURE 18 is an illustration of an example embodiment of a system
1800 for a configurable ALU pipeline. System 1800 may include a processor, SoC, integrated circuit, or other mechanism suitable for a configurable ALU pipeline. For
25 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. Processor 1802 may include any suitable mechanism for configuring the ALU pipeline. In one embodiment, such mechanisms may be implemented in hardware. In another embodiment, such mechanisms may include a memory-mapped address for
30 the configuration. In a further embodiment, such mechanisms may include one or more distinct entries in microcode to configure the ALU. In yet another embodiment, such mechanisms may include an instruction for a programmer, compiler, or firmware

to configure the ALU pipeline. In still yet another embodiment, such mechanisms may include distinct instructions for a given ALU configuration. Processor 1802 may be implemented fully or in part by the elements described in FIGURES 1–17.

5 **[00149]** In one embodiment, system 1800 may include an ALU 1826 to perform arithmetic operations or logic operations on input operands. The operands may specify immediate values, memory locations, or source registers. System 1800 may include an ALU 1826 in any suitable portion of system 1800. For example, ALU 1826 may be implemented as one of the execution units 1822 within an in-order or out-of-order execution pipeline 1816. ALU 1826 may be implemented by any
10 suitable combination of circuitry or hardware computational logic of a processor.

[00150] A configurable ALU pipeline for ALU 1826 may be used in a variety of applications requiring flexibility in the balance between performance and power efficiency, including hand-held and Internet of Things (IoT) devices. An ALU pipeline may execute appropriate instructions in a fixed number of cycles, which may
15 range based on an input characteristic. The input characteristic may include the data width of input operands, in addition to the number of 1's and 0's in the input operands. For example, smaller-sized inputs may execute in fewer fixed cycles than larger-sized inputs. ALU 1826 may execute a specific combination of a particular instruction and particular input operand in a fixed number of cycles. Some
20 applications may have additional needs to conserve battery life or reduce power consumption, while other applications may require superior ALU performance. A configurable ALU pipeline for ALU 1826 may enable a device to operate more optimally in both types of applications. In various embodiments, ALU 1826 may be configured for either performance or power efficiency.

25 **[00151]** Configuring an ALU pipeline may require additional silicon area. Hardware acceleration of ALU pipelines may enable increased performance at the expense of increased power and additional silicon. Moreover, hardware acceleration may require additional overhead to transfer data from a processor 1802 to the hardware accelerated unit, which may result in an execution time penalty.
30 Dynamic voltage and frequency scaling (DVFS) may also enable a configurable ALU pipeline. DVFS may also have a response time that is greater than the response time desired by a programmer. For example, a programmer may want ALU 1826 to

repeatedly switch between a performance mode and a power efficiency mode. An ISA may also enable a configurable ALU pipeline with specific instructions implementing existing functions with either a faster or a more efficient execution. The expansion of an ISA may require significant additions or changes to the ISA, and may rely on changes to compilers to support the new instructions. Such support may include the ability to recognize when a programmer intended to switch modes of ALU pipeline operation. Configuring the pipeline of an ALU 1826 using a specific instruction may enable a programmer or drafter of the code to configure the ALU pipeline with instruction-level control. Such an instruction may rely on the programmer alone, and not on an improved compiler. Moreover, such an instruction may not require new sets of arithmetic and logic instructions for an ALU 1826 with a configurable ALU pipeline. A software-controllable pipeline may also be implemented with minimal additional hardware overhead and minimal increases in complexity for ALU pipeline design.

15 **[00152]** 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 of 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, which may receive and decode instructions from instruction stream 1804 using a decode pipeline stage. The decoded instructions may be dispatched, allocated, and scheduled for execution by an allocation unit 1818 and scheduler 1820 of an execution pipeline 1816 and allocated to specific execution units 1822. After execution, instructions may be retired by a writeback stage or retirement stage in retirement unit 1824. If processor 1802 executes instructions out-of-order, allocation unit 1818 may rename instructions and the instructions may be input into a reorder buffer 1824 in association with the retirement unit. 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.

[00153] Processor 1802 may recognize, either implicitly or through decoding and execution of specific instructions, that the ALU should be configured.

In such cases, execution pipeline 1816 may adjust how ALU 1826 computes results. In one embodiment, ALU 1826 may be targeted by specific instructions 1830 to be executed in instruction stream 1804. Such specific instructions may be generated by, for example, a compiler, or may be designated by a drafter of code resulting in instruction stream 1804. The instruction may be included in a library defined for execution by processor 1802 or ALU 1826. In another embodiment, ALU 1826 may be targeted by portions of processor 1802, wherein processor 1802 recognizes the need for reconfiguring ALU 1826.

[00154] Instructions 1830 may include SET_ALUPPCR(*Mode*, Reduced Size Operand, *Accuracy*) or SET_ALUPPCR(*Value*, [*Mask*]). *Mode* may correspond to the ALU mode of operation, which may include a normal and an accelerated mode. *Reduced Size Operand* may correspond to an arithmetic or logic operation where the data contained in the operands do not require the full width of the operand. For example, a 32-bit operand may only contain useful data in the lower 16-bits or lower term. The operand may be of any suitable width and the useful data may be of any suitable value less than the width of the operand. Such an operand may enable ALU 1826 to compute a result more quickly by ignoring the upper term of the operand or operands. *Accuracy* may correspond to the number of partial products to accumulate before outputting a result or may correspond to the number of cycles of execution by the ALU 1826 before outputting a result. *Value* may correspond to a value to set in the ALU performance register or ALUPPCR. This value may be masked by *Mask* to set a subset of the performance register. The mask may indicate which one or more bits to mask or modify the *Value* input. In one embodiment, the values stored may remain unchanged if the mask is set. In another embodiment, the values stored may be set to zero if the mask is set. A mask for a given bit may be represented with either a 0 or a 1.

[00155] Instructions 1830 may also include MULT(*D*, *A*, *B*), L/R_SHIFT(*D*, *A*, *Bits*), ADD(*D*, *A*, *B*). *MULT* may correspond to a multiply instruction, which may multiply *A* by *B* and store the result in *D*. *L/R_SHIFT* may correspond to a shift instruction, which may perform an arithmetic or logical shift on *A* to the left or right by *Bits* and store the result in *D*. *ADD* may correspond to an addition instruction, which may sum *A* and *B* and store the result in *D*. *D* may

correspond to a destination, which may be any suitable location to store data including memory and registers. *A* may correspond to a first operand, which may be any suitable location to source data including an immediate value, memory, or a register. *B* may correspond to a second operand, which may be any suitable location to source data including an immediate value, memory, or a register. *Bits* may represent the number of bits to shift an operand and may correspond to any suitable location to source data including an immediate value, memory, or a register. These parameters—Mode, Reduced Size Operand, Accuracy, Value, Mask, D, A, B, and Bits—may be in any suitable form, including parameter flags, explicit parameters, required parameters, optional parameters with an assumed default value, or inherent parameters stored in registers or other known locations that do not require the information to be passed as a parameter explicitly.

[00156] Configurable ALU pipeline for ALU 1826 may be implemented in any suitable manner. In one embodiment, configurable ALU pipeline for ALU 1826 may be implemented by separate pipelines, in which each pipeline corresponds to a specified mode. Moreover, the separate pipelines may be implemented by circuitry to compute the desired result from the input operands. The separate pipelines may be referenced with separate microcode entries. In another embodiment, configurable ALU pipeline for ALU 1826 may be implemented with a unified pipeline, in which the specified mode may activate additional logic for computing the desired result. The unified pipeline may be implemented by circuitry and may be referenced using one microcode entry.

[00157] In one embodiment, configurable ALU pipeline for ALU 1826 may be configured for execution in normal mode. Normal mode may require additional cycles to complete execution in a power efficient manner at the expense of performance. In another embodiment, configurable ALU pipeline for ALU 1826 may be configured for execution in accelerated mode. Accelerated mode may complete execution in one or a few cycles in a performance-oriented manner at the expense of power efficiency. Accelerated mode may complete execution faster than normal mode. In a further embodiment, configurable ALU pipeline for ALU 1826 may be configured for execution of a reduced size operand. A drafter of code or even a compiler may be aware that the most significant term in a number is unused. For

example, an application may only use the least significant term or the lower 16-bits of a 32-bit integer, leaving the most significant term or the upper 16-bits set to all 0's. In this instance, configurable ALU pipeline for ALU 1826 may complete execution in a power efficient manner using the normal mode that may also be as fast as an accelerated pipeline approach. This may be achieved by limiting the number of cycles for execution. Alternatively, configurable ALU pipeline for ALU 1826 may be configured for accelerated mode with a reduced size operand. For example, a 32-bit input with relevant data in the least significant term may be reduced to 16-bits before an accelerated mode ALU performs a 16-bit operation, and the result may be extended back to the expected number of bits after the ALU computes the result.

[00158] In yet another embodiment, configurable ALU pipeline for ALU 1826 may accelerate computation by reducing the accuracy of the result. For example, configurable ALU pipeline for ALU 1826 may be configured for normal mode with a reduced accuracy. A subsequent multiply instruction for the ALU may complete computation after calculating as few as one partial product. The accuracy field of the ALU performance register may be defined in any suitable manner to signal to the programmer a defined or relative level of accuracy. For example, the accuracy field may have four values corresponding to 25, 50, 75, and 100% accuracy. 25% accuracy may divide the expected number of cycles for execution by 4, the inverse of 25%. A drafter of code programming processor 1802 to execute highly fault tolerant applications may only need 25% accuracy to ensure proper operation. In such case, a 4-cycle multiplication, for instance, may end after the first cycle. Configurable ALU pipeline for ALU 1826 may also reorder the computation to ensure the most significant bits are computed first. This may ensure the reduced accuracy result is within a margin of error of the actual result. Alternatively, the accuracy field may have unique values corresponding to the number of cycles for execution. For example, the multiplication of two 32-bit input operands with full accuracy may require four execution cycles, one for each partial product of 16-bit operands. For this operation, the accuracy field of the performance register may define a unique value for two cycles of operation, in which the result is available after two cycles of execution. In yet a further embodiment, a reduced accuracy result, which may be obtained in relatively few cycles with less power consumed, may be stored in one

register, while a full accuracy result, obtained in relatively more cycles with more power consumed, may be stored in another register.

[00159] ALU 1826 may support any suitable data type, including but not limited to a byte, word, doubleword, quadword, single-precision floating point, or double-precision floating point. The memory addressing supported may include any suitable type, including but not limited to 32-bit and 64-bit addressing. The parameters—Mode, Reduced Size Operand, Accuracy, Value, Mask, D, A, B, and Bits—may be represented as either a parameter or operand of an instruction, or may be part of the name or designation of the instruction, and may be defined by any suitable data type, including an 8-bit immediate value, a general purpose register value represented as a byte, word, doubleword, or quadword, and may be from any suitable source.

[00160] System 1800 may also have a power monitor 1828. Power monitor 1828 may measure the power or energy consumed by the various portions of system 1800, including processor 1802 and core(s) 1814. Power monitor 1828 may have sensors located across system 1800 to facilitate to measuring or monitoring power. Power monitor 1828 may be implemented in hardware, software, firmware, or a combination of such implementation approaches, including a microcontroller, an application specific integrated circuit (ASIC), a microprocessor, or a firmware unit. Power monitor 1828 may interface with instruction stream 1804 to insert instructions for execution by processor 1802. Alternatively, power monitor 1828 may interface directly with processor 1802 to issue commands necessary to control the power consumption of processor 1802. For example, power monitor 1828 may determine an application executing on processor 1802 is consuming an amount of power that exceeds a predetermined or programmed power threshold. In this case, power monitor 1828 may reduce the power consumed by system 1800 by any suitable type or combination of actions including reducing the voltage, reducing the frequency, partially or fully halting execution, or configuring ALU 1826 to consume less power at the expense of performance by setting the mode field of the performance register to normal mode. Power monitor 1828 may also determine an application executing on processor 1802 is not exceeding the power threshold and may increase the power consumed by system 1800 by any suitable type or combination of actions including

increasing the voltage, increasing the frequency, or configuring ALU 1826 to increase performance at the expense of consuming more power by setting the mode field of the performance register to accelerated mode. In one embodiment, power monitor 1828A resides external to processor 1802. In another embodiment, power monitor 1828B
5 may be internal to processor 1802.

[00161] 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.

[00162] FIGURE 19 illustrates a logical representation of elements of
10 system 1800 with separate selectable ALU pipelines, according to embodiments of the present disclosure. System 1900 may include an accelerated mode ALU pipeline and a normal mode ALU pipeline. Each pipeline may be targeted using unique microcode entries or by adding circuitry to execution pipeline 1816 to select the appropriate mode.

[00163] One or more instructions from an instruction stream may be
15 decoded at 1902. The decoding of an instruction may indicate the type of instruction. At 1904, the type of instruction may be evaluated to determine whether it corresponds to a multiply, add, or shift operation. Prior to the instruction being decoded, a previous instruction may have set the performance register of the ALU, including the
20 mode field of such a register. At 1906, the mode field of the ALU performance register may be evaluated to determine whether it corresponds to an accelerated mode. If an accelerated mode is specified, and the reduced size operand field of the performance register may indicate the presence of reduced size operands or the accuracy field of the performance register indicates the preference for a reduced
25 accuracy computation, the input operand width may be adjusted at 1910 before entering the accelerated pipeline 1912. For example, an addition instruction with 16-bit input operands may only need to compute the sum of the lower 8-bits of each operand. In this case, the 16-bit input operand width may be adjusted to 8-bits. The width of the output from the accelerated pipeline may be adjusted at 1914 back to an
30 expected width if necessary. 1902, 1904, 1906, and 1910 may be performed by a unit other than the execution units 1822. For example, front end 1806, execution pipeline 1816, or allocation unit 1818 may perform one or more of the operations. If an

accelerated mode is not specified, the normal mode ALU pipeline may be entered at 1908. System 1900 may select between the accelerated result from 1914 and the normal result from 1908. If the mode of the ALU corresponds to normal mode, the normal result from the normal mode ALU pipeline may be accepted at 1918. Otherwise, if the mode of the ALU corresponds to accelerated mode, the accelerated result from the accelerated mode ALU pipeline may be accepted at 1920.

[00164] FIGURE 20 is a block diagram of elements of a system 1800 with a shared or unified configurable ALU pipeline, according to embodiments of the present disclosure. System 2000 may share combinational logic and/or sequential elements between multiple ALU pipelines. For example, elements of a normal ALU pipeline may share elements with an accelerated ALU pipeline such that one or more stages at the front end and one or more stages on the backend of the ALU pipelines are common. A shared configurable ALU pipeline may reduce the amount of overhead logic and complexity to only an accelerated mode single-cycle multiplier, shifter, or adder 2012, or to only a normal mode multi-cycle multiplier, shifter, or adder 2010.

[00165] Inputs 2002 may be sent to 2004 ALU stage 1 where some initial processing may occur. Then, data may enter 2006 ALU stage 2 for additional processing that may be common between normal and accelerated modes. 2004 and 2006 may take one cycle total to complete execution in either the normal or accelerated modes. At 2008, the ALU mode may be evaluated. A normal ALU mode may proceed to 2010 for normal mode. An accelerated mode may proceed to 2012 for accelerated mode. Normal mode 2010 may contain several ALU stages 2014 and 2018. Data may initially be processed in 2014 ALU Stage 3. A multiple-cycle multiply, shift, or add may proceed, requiring 2018 ALU stage N. In total, normal mode execution in ALU 2000 may require N stages or N cycles if each stage is one cycle of execution. Accelerated mode 2012 may contain a single-cycle multiply, shift, or add 2020 and ALU stage M 2022. In total, accelerated mode execution in ALU 2000 may require M stages or M cycles if each stage is one cycle of execution.

[00166] At the expense of power, accelerated mode 2012 may provide additional performance for ALU 2000. By defining $M < N$, ALU 2000 may provide two options for ALU operations where one option, with M stages, is faster than the

other option, with N stages. By defining option M to be faster, requiring fewer stages, system 1800 may optimize or be optimized to balance the power and performance of ALU 2000. At 2024, the result selector uses the ALU mode 2026 to determine which output to select. 2024 may require one cycle to complete execution in either the normal or accelerated modes. If ALU mode 2026 is set to normal mode, result selector 2024 may select the output of normal mode 2010. If ALU mode 2026 is set to accelerated mode, result selector 2024 may select the output of accelerated mode 2012. The selected output may be set to the result 2028 for use by other logical units of processor 1802. Accelerated mode 2012 may complete execution in one cycle. However, sourcing operands to and from accelerated mode 2012 may take a total of two additional cycles via a bus of the ALU 2000. Thus, execution in accelerated mode may complete in three cycles. In contrast, normal mode 2010 may complete execution in four cycles, for a total of six cycles for execution in normal mode to complete. In such an example, execution in accelerated mode may be twice as fast as execution in normal mode.

[00167] FIGURE 21 is a more detailed block diagram of elements of a system with a configurable ALU pipeline, according to embodiments of the present disclosure. Configurable ALU pipeline 2100 may perform a variety of suitable arithmetic or logical operations, including multiplication as shown. 2102 describes an operation to multiply two 32-bit integers in which the 64-bit product D may be computed from input operands A and B. Input operands A and B may be of any suitable size and may be multiplied using partial products. A may have an upper or high term, AH 2104, and a lower term, AL 2106. 2104 may represent the most significant term and 2106 may represent the least significant term. Similarly, B may have an upper or high term 2108 and a lower term 2110. These values may be passed either to normal mode logic unit 2112 or accelerated logic unit 2114. The mode field of a performance register for ALU 2100 may control which logic unit is selected. This may be represented by ALU mode 2134.

[00168] Normal mode logic unit 2112 may execute over one or more cycles, computing one partial product per cycle. Alternatively, normal mode logic unit 2112 may execute over one or more cycles, computing multiple partial products per cycle. Accumulated value 2140 may be initialized to 0 or some other suitable

value. Partial product 2116 may represent a series of partial products, where each partial product corresponds to a cycle. For example, to multiply 32-bit operands, four partial products may need to be computed across four cycles. 2138 shows the logical representation of the partial product computation where D may be equal to $AH \times BH \times 2^{32} + AL \times BH \times 2^{16} + AH \times BL \times 2^{16} + AL \times BL \times 2^0$. In the first cycle, AH 2104 may be multiplied by BH 2108 and adjusted by multiplying the result by the partial product magnitude, 2^{32} . Multiplication by the magnitude may also be performing by shifting the data. The resulting partial product may be added to the accumulated value 2140 at 2118. Sum 2118 may be output from normal mode logic unit 2112 to ALU mode selector 2136. ALU mode selector may determine the ALU mode is set to normal mode. Moreover, it may be determined that an additional partial product needs to be computed to achieve the desired result. Such determination may be made by normal mode logic unit or ALU mode selector 2136. If it is determined that an additional partial product must be computed, accumulation value 2140 may be updated with the sum. Alternatively, accumulation value 2140 may be updated with the sum regardless of the determination.

[00169] Returning to the example shown in 2138, in the second cycle, normal mode logic unit 2112 may multiply AL 2106 by BH 2108 and adjust the result by multiplying it by the partial product magnitude, 2^{16} . The resulting partial product may be added to the accumulated value 2140 at 2118. Sum 2118 may be output from normal mode logic unit 2112. In the third cycle, normal mode logic unit 2112 may multiply AH 2104 by BL 2110 and adjust the result by multiplying it by the partial product magnitude, 2^{16} . The resulting partial product may be added to the accumulated value 2140 at 2118. Sum 2118 may be output from normal mode logic unit 2112. In the fourth cycle, normal mode logic unit 2112 may multiply AL 2106 by BL 2110, and adjust the result by multiplying it by the partial product magnitude, 2^0 . The resulting partial product may again be added to the accumulated value 2140 at 2118. Sum 2118 may be output from normal mode logic unit 2112. It may then be determined that no additional partial product needs to be calculated.

[00170] Accelerated mode logic unit 2114 may execute over one or more cycles, but in fewer cycles than normal mode logic unit 2112. Accelerated mode logic unit 2114 may achieve this accelerated operation by computing all partial

products in the same one or more cycles. In one embodiment, the partial products may all be computed in parallel. For example, logic unit 2120 may calculate the partial product associated with AH 2104 multiplied by BH 2108 and by 2^{32} ; logic unit 2122 may compute the partial product associated with AL 2106 multiplied by BH 2108 and by 2^{16} ; logic unit 2124 may compute the partial product associated with AH 2104 multiplied by BL 2110 and by 2^{16} ; and logic unit 2126 may compute the partial product associated with AL 2106 multiplied by BL 2110 multiplied by 2^0 . The partial products from logic units 2120 and 2122 may be summed at 2128, and the partial products from logic unit 2124 and 2126 may be summed at 2130. In one embodiment, the partial products 2120, 2122, 2124, and 2126 may not use the exponents (e.g., 2^{32}) and instead the exponents may be used in 2128 and 2130 to shift the partial products before summation. The sums from 2128 and 2130 may then be summed at 2132. The output from 2132 may represent the accelerated result from accelerated mode logic unit 2114. ALU mode selector 2136 may select the output from normal mode logic unit 2112 or accelerated mode logic unit 2114 and output the result to D 2142.

[00171] In one embodiment, the input operands may have a reduced size. For example, AH 2104 and BH 2108 may be unused. In this case, output D[63:0] may represent the multiplication of AL 2106 and BL 2110 and 2^0 . This multiplication may be performed by either the normal mode logic unit 2112 or the accelerated mode logic unit 2114. The normal mode logic unit 2112 may compute the result by computing the least significant partial products first and limiting the number of partial products to calculate. In this case, only one partial product may be needed. The accelerated mode logic unit 2114 may compute the result by reducing the size of the input operands to 16-bits, representing the least significant terms, before entering the accelerated pipeline. This may enable faster computation of the result or computation of the result with greater efficiency. Once the result is available, it may be extended back to the expected size before outputting the value from the execution pipeline 1816.

[00172] In another embodiment, the accuracy of the result may be reduced. For example, an application may only need the most significant partial product. In this case, output D[63:0] may represent the multiplication of AH 2104

and BH 2108 and 2^{32} . This multiplication may be performed by either the normal mode logic unit 2112 or the accelerated mode logic unit 2114. The normal mode logic unit 2112 may compute the result by computing the most significant partial products first and limiting the number of partial products to calculate. In this case, only one partial product may be needed. The accelerated mode logic unit 2114 may compute the result by reducing the size of the input operands to 16-bits, representing the most significant terms, before entering the accelerated pipeline. This may enable faster computation of the result or computation of the result with greater efficiency. Once the result is available, it may be extended back to the expected size before outputting the value from the execution pipeline 1816. In a further embodiment, the reduced size operand execution may be combined with the reduced accuracy execution to further optimize either the performance or power efficiency of configurable ALU pipeline 2100.

[00173] FIGURE 22 is a diagram of operation of a method for configuring an ALU pipeline, according to embodiments of the present disclosure. Method 2200 may be implemented by any of the elements shown in FIGURES 1–21. Method 2200 may be initiated by any suitable criteria and may initiate operation at any suitable point. In one embodiment, method 2200 may initiate operation at 2202. Method 2200 may include greater or fewer steps than those illustrated. Moreover, method 2200 may execute its steps in an order different than those illustrated below. Method 2200 may terminate at any suitable step. Furthermore, method 2200 may repeat operation at any suitable step. Method 2200 may perform any of its steps in parallel with other steps of method 2200, or in other methods. Method 2200 may perform any of its steps on any element of data in parallel with other elements of data, such that method 2200 operates in a vectorized way.

[00174] At 2202, in one embodiment an instruction for setting an ALU performance register may be received, decoded, and allocated. The ALU performance register may reside within the ALU, another execution unit, or any place in system 1800 suitable for access by ALU 1826, including front end 1806, core 1814, and execution pipeline 1816. At 2204, in one embodiment the ALU performance register may be set based on operands or parameters of the instruction. The ALU performance register may have one or more fields each corresponding to one or more

bits in the register. The ALU performance register may have any number of fields necessary to configure the performance of the ALU. For example, the register may have a mode field for configuring the ALU to operate in normal or accelerated mode. Normal mode may be represented by a value of 0 or 1, and accelerated mode may be represented by an opposite value of 1 or 0.

[00175] Moreover, the register in the example above may have an accuracy field for specifying the accuracy of the ALU's operation. The accuracy field may correspond to an ALU's normal or accelerated mode. In one embodiment, the accuracy field may specify the number of cycles to execute before outputting a result. For example, the accuracy field value of 3 may correspond to 4 cycles of execution, which may represent full accuracy for 32-bit multiplication. Alternatively, the accuracy field value of 0 may correspond to 1 cycle of execution, which may represent minimum accuracy for 32-bit multiplication. In another embodiment, the accuracy field may specify a percentage of accuracy. For example, the accuracy field value of 2 may correspond to 50% accuracy, which may represent 8 partial products of an operation that requires 16 partial products for full accuracy. A programmer or compiler may understand that certain applications are fault tolerant and do not require full accuracy to achieve a desired purpose. In such circumstances, the programmer or compiler may set the accuracy field to a value corresponding to less than full accuracy, which may reduce power consumption while increasing performance of ALU 1826.

[00176] Furthermore, the performance register of the ALU may have a reduced size operand field for specifying the scope of the ALU's operation. The reduced size operand field may correspond to an ALU's normal or accelerated mode. In one embodiment, the reduced size operand field may specify the first and second operands, A and B, have a reduced useable width, which may be defined as one-half the operand's width. Accordingly, a 32-bit operand may have a reduced useable width of 16-bits, where the most-significant term or upper 16-bits are unused. A programmer or compiler may understand that certain uses of ALU instructions may specify an operand width wider than the actual data width. In such situations, the reduced useable width may be used to reduce power consumption or increase performance.

[00177] At 2206, in one embodiment an instruction for setting the ALU performance register may be retired by, for example, retirement unit 1824. At 2208, in one embodiment an instruction for an ALU may be received, decoded, and allocated. The instruction may be of any kind suitable for execution in an ALU, including multiply, add, and shift. At 2210, in one embodiment it may be determined whether the mode field of the performance register is set to accelerated mode. If the mode field is set to accelerated mode, method 2200 may proceed to 2212. Otherwise, the mode field may be set to normal mode, and method 2200 may proceed to 2214. At 2212, in one embodiment at least one input operand bit width is adjusted. The operand bit width may need to be adjusted to the bit width of the pipeline or to the reduced operand size. At 2214, in one embodiment the instruction may be received at the ALU. At 2216, in one embodiment it may be determined whether the mode field of the performance register is set to normal mode. If the mode field is set to accelerated mode, method 2200 may proceed to 2232. Otherwise, the mode field may be set to normal mode, and method 2200 may proceed to 2218.

[00178] At 2218, in one embodiment the normal accumulation and the partial product counter may be initialized. Initialization may include setting the values to zero or some other suitable value. At 2220, in one embodiment the total number of partial products may be determined based on the size of the ALU and the size of the operands. For example, if the instruction is a multiply operation and the ALU can only compute one 8-bit result at a time, the total number of partial products of two 16-bit operands may be four. At 2222, in one embodiment the number of partial product cycles may be determined based on at least the accuracy field in the ALU performance register. If the accuracy field is set to a reduced accuracy, the number of partial products may be reduced to increase performance at the expense of accuracy.

[00179] At 2224, in one embodiment the order and total number of partial products to compute may be adjusted or prioritized based on the reduced size field of the performance register. For example, if the most significant 8-bit terms of 16-bit operands are unused, the total number of partial products may be limited to one partial product, and the order of partial products to compute prioritize the least significant partial product first. The reduced size operand field may enable the ALU

to increase performance without reducing power efficiency or accuracy. At 2226, in one embodiment a partial product may be computed. The computation may include the exponent corresponding to the particular terms being multiplied. At 2228, in one embodiment the normal accumulation may be added to the partial product and the sum may be stored back in the normal accumulation. The computation at 2228 may include the exponent (e.g., 2^{32}) corresponding to the computed partial product. At 2230, in one embodiment, the partial product counter may be incremented. Method 2200 may then proceed to step 2234.

[00180] If the ALU mode is not normal mode at method step 2216, method 2200 may proceed to method step 2232. At 2232, in one embodiment all partial products may be computed and summed. 2232 may include adjusting the number of partial products or the operand size based on the size of the ALU, the size of the input operands, the accuracy field, or the reduced size field. At 2234, in one embodiment the normal accumulation and the accelerated accumulation may be received by the ALU mode selector. At 2236, in one embodiment it may be determined whether the mode field of the performance register is set to normal mode. If the field is set to normal mode, method 2200 may proceed to method step 2238. Otherwise if the field is set to accelerated mode, method 2200 may proceed to method step 2244.

[00181] At 2238, in one embodiment the partial product counter may be compared to the total number of partial products required. At 2240, in one embodiment it may be determined whether one or more additional partial products need to be computed based on the comparison. At 2242, it may be determined whether the normal mode computation is done. If so, method 2200 may proceed to method step 2246. Otherwise, method 2200 may proceed to method step 2226 to compute one or more additional partial products. If the partial product counter is less than the total number of partial products required, normal mode computation may not be done. At 2244, in one embodiment the output width from the accelerated mode may be adjusted. An output width may need to be adjusted if the accelerated mode pipeline bit width is less than the expected output bit width. This may occur, for example, if the reduced size operand field is set. For instance, a 32-bit result may be extended to 64-bits.

[00182] At 2246, in one embodiment the output result may be selected. The ALU mode selector may make such selection based on the mode field of the ALU performance register. At 2248, in one embodiment the instruction for the ALU may be retired by, for example, retirement unit 1824. Method 2200 may optionally
5 repeat or terminate.

[00183] 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
10 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.

[00184] 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
15 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.

[00185] The program code may be implemented in a high level procedural or object oriented programming language to communicate with a
20 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.

[00186] One or more aspects of at least one embodiment may be
25 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
30 load into the fabrication machines that actually make the logic or processor.

[00187] 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
5 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.

10 **[00188]** 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.

15 **[00189]** 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.
20 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.

[00190] Thus, techniques for performing one or more instructions according to at least one embodiment are disclosed. While certain exemplary
25 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
30 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.

[00191] 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.

[00192] 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.

5 **[00193]** 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
10 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
15 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
20 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
25 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.

30 **[00194]** 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.

[00195] 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.

[00196] 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
10 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
15 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
20 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
25 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
30 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.

5 **[00197]** 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
10 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
15 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.

[00198] 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
20 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.
25 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
30 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.

[00199] In combination with any of the above embodiments, in an embodiment the second instruction may correspond to a multiplication operation and
5 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
10 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
15 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
20 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
25 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.

[00200] In some embodiments of the present disclosure, an arithmetic logic unit may include a performance register including a mode field and four circuits.
30 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
5 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.

10 **[00201]** 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
15 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
20 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
25 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.

[00202] In combination with any of the above embodiments, in an embodiment the decoded instruction may correspond to a multiplication operation and
30 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.

[00203] 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.

[00204] 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
5 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
10 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
15 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
20 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.

[00205] In combination with any of the above embodiments, in an
25 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
30 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

5 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

10 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

15 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

20 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.

CLAIMS

1. A processor, comprising:
 - an execution unit;
 - a front end including circuitry to decode a first instruction to set a performance
 - 5 register for the execution unit, wherein the performance register includes a mode field to specify a mode of operation for the execution unit, and a second instruction;
 - an allocator including circuitry to assign the second instruction to the execution unit to execute the second instruction;
 - wherein the execution unit includes circuitry to:
 - 10 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, wherein the accelerated computation is faster than the normal computation;
 - perform the selected computation; and
 - 15 select between the normal result and the accelerated result based on the mode field of the performance register.
2. The processor of Claim 1, wherein the normal computation includes circuitry to:
 - 20 initialize a normal accumulation and a partial value counter;
 - determine a total number of partial values required based on input operands of the second instruction;
 - compute a partial values of the input operands;
 - increment the partial value counter;
 - 25 add the normal accumulation to the computed partial value and store the result of the addition back in the normal accumulation;
 - determine whether to compute an additional partial value based on the partial value counter and the total number of partial values required; and
 - store the normal accumulation in the normal result based on the determination
 - 30 that the additional partial value need not be computed.

3. The processor of Claim 1, wherein the second instruction corresponds to a multiplication operation and the normal computation includes circuitry to:
- initialize a normal accumulation and a partial product counter;
 - determine a total number of partial products required based on input operands
 - 5 of the second instruction;
 - compute a partial product of the input operands;
 - increment the partial product counter;
 - add the normal accumulation to the computed partial product and store the result of the addition back in the normal accumulation;
 - 10 determine whether to compute an additional partial product based on the partial product counter and the total number of partial products required; and
 - store the normal accumulation in the normal result based on the determination that the additional partial product need not be computed.
- 15 4. The processor of Claim 1, wherein the second instruction corresponds to an addition operation and the normal computation includes circuitry to:
- initialize a normal accumulation and a partial sum counter;
 - determine a total number of partial sums required based on input operands of the second instruction;
 - 20 compute a partial sum of the input operands;
 - increment the partial sum counter;
 - add the normal accumulation to the computed partial sum and store the result of the addition back in the normal accumulation;
 - determine whether to compute an additional partial sum based on the partial
 - 25 sum counter and the total number of partial sums required; and
 - store the normal accumulation in the normal result based on the determination that the additional partial sum need not be computed.
5. The processor of Claim 1, wherein the second instruction corresponds
- 30 to a shift operation and the normal computation includes circuitry to:
- initialize a normal accumulation and a partial shift counter;

determine a total number of partial shifts required based on input operands of the second instruction;
compute a partial shift of the input operands;
increment the partial shift counter;
5 add the normal accumulation to the computed partial shift and store the result of the addition back in the normal accumulation;
determine whether to compute an additional partial shift based on the partial sum counter and the total number of partial shifts required; and
store the normal accumulation in the normal result based on the determination
10 that the additional partial shift need not be computed.

6. The processor of Claim 1, wherein the accelerated computation includes circuitry to:
compute all partial values of input operands of the second
15 instruction; and
compute the accelerated result based on sum of all the computed partial values.

7. The processor of Claim 1, further comprising circuitry to:
20 monitor energy consumed;
compare the energy consumed to a power threshold; and
set the mode field of the performance register based on the comparison.

8. A method, comprising:
25 setting a mode field of a performance register associated with an arithmetic logic unit;
decoding an instruction for the arithmetic logic unit;
selecting between a normal computation corresponding to a normal result and an accelerated computation corresponding to an accelerated result based on the mode
30 field of the performance register, wherein the accelerated computation is faster than the normal computation;
performing the selected computation; and

selecting between the normal result and the accelerated result based on the mode field of the performance register.

9. The method of claim 8, wherein the normal computation includes:

5 initializing a normal accumulation and a partial value counter;

determining a total number of partial values required based on input operands of the decoded instruction;

computing a partial value of the input operands;

adding the normal accumulation to the computed partial value and storing the

10 result of the addition back in the normal accumulation;

incrementing the partial value counter;

comparing the partial value counter to the total number of partial values required;

determining whether to compute an additional partial value based on the

15 comparison; and

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

10. The method of Claim 8, wherein the decoded instruction corresponds

20 to a multiplication operation and the normal computation includes:

initializing a normal accumulation and a partial product counter;

determining a total number of partial products required based on input operands of the decoded instruction;

computing a partial product of the input operands;

25 adding the normal accumulation to the computed partial product and storing the result of the addition back in the normal accumulation;

incrementing the partial product counter;

comparing the partial product counter to the total number of partial products required;

30 determining whether to compute an additional partial product based on the comparison; and

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

11. The method of Claim 8, wherein the decoded instruction corresponds
5 to an addition operation and the normal computation includes:
initializing a normal accumulation and a partial sum counter;
determining a total number of partial sums required based on input operands
of the decoded instruction;
computing a partial sum of the input operands;
10 adding the normal accumulation to the computed partial sum and storing the
result of the addition back in the normal accumulation;
incrementing the partial sum counter;
comparing the partial sum counter to the total number of partial sums required;
determining whether to compute an additional partial sum based on the
15 comparison; and
storing the normal accumulation in the normal result based on the
determination not to compute the additional partial sum.

12. The method of Claim 8, wherein the decoded instruction corresponds
20 to a shift operation and the normal computation includes:
initializing a normal accumulation and a partial shift counter;
determining a total number of partial shifts required based on input operands
of the decoded instruction;
computing a partial shift of the input operands;
25 adding the normal accumulation to the computed partial shift and storing the
result of the addition back in the normal accumulation;
incrementing the partial shift counter;
comparing the partial shift counter to the total number of partial shifts
required;
30 determining whether to compute an additional partial shift based on the
comparison; and

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

13. The method of Claim 8, wherein the accelerated computation includes:
5 computing all partial values of input operands of the decoded instruction;
computing a sum of all partial values; and
storing the sum in the accelerated result.

14. The method of Claim 8, further comprising:
10 monitoring energy consumed by the arithmetic logic unit;
comparing the energy consumed to a power threshold; and
setting the mode field of the performance register based on the comparison.

15. The method of Claim 8, wherein the normal computation and the
15 accelerated computation are implemented with a shared computational pipeline.

16. The method of Claim 9, wherein:
the accuracy field of the performance register defines a total number of cycles
for the arithmetic unit to compute the normal result;
20 the step of determining the total number of partial values required is further
based on an accuracy field of the performance register, total number of cycles, and a
number of cycles for computing the partial value; and
the step of computing the partial value of the input operands further includes
prioritizing a most significant term of each input operand for the partial value.

25
17. An arithmetic logic unit, comprising circuitry to:
a performance register including a mode field;
receive a decoded instruction;
select between a normal computation corresponding to a normal result and an
30 accelerated computation corresponding to an accelerated result based on the mode
field of the performance register, wherein the accelerated computation is faster than
the normal computation;

perform the selected computation; and
select between the normal result and the accelerated result based on the mode field of the performance register.

- 5 18. The arithmetic logic unit of Claim 17, wherein the normal computation includes circuitry to:
- initialize a normal accumulation and a partial value counter;
 determine a total number of partial values required based on input operands of the instruction;
- 10 compute a partial value of the input operands;
 increment the partial value counter;
 add the normal accumulation to the computed partial value and store the result of the addition back in the normal accumulation;
 determine whether to compute an additional partial value based on the partial
- 15 value counter and the total number of partial values required; and
 store the normal accumulation in the normal result based on the determination that the additional partial value need not be computed.

19. The arithmetic logic unit of Claim 17, wherein the decoded instruction
- 20 corresponds to a multiplication operation and the normal computation includes circuitry to:
- initialize a normal accumulation and a partial product counter;
 determine a total number of partial products required based on input operands of the instruction;
- 25 compute a partial product of the input operands;
 increment the partial product counter;
 add the normal accumulation to the computed partial product and store the result of the addition back in the normal accumulation;
 determine whether to compute an additional partial product based on the
- 30 partial product counter and the total number of partial products required; and
 store the normal accumulation in the normal result based on the determination that the additional partial product need not be computed.

20. The arithmetic logic unit of Claim 17, wherein the decoded instruction corresponds to an addition operation and the normal computation includes circuitry to:

- 5 initialize a normal accumulation and a partial sum counter;
- determine a total number of partial sums required based on input operands of the instruction;
- compute a partial sum of the input operands;
- increment the partial sum counter;
- 10 add the normal accumulation to the computed partial sum and store the result of the addition back in the normal accumulation;
- determine whether to compute an additional partial sum based on the partial sum counter and the total number of partial sums required; and
- store the normal accumulation in the normal result based on the determination
- 15 that the additional partial sum need not be computed.

21. The arithmetic logic unit of Claim 17, wherein the decoded instruction corresponds to a shift operation and the normal computation includes circuitry to:

- initialize a normal accumulation and a partial shift counter;
- 20 determine a total number of partial shifts required based on input operands of the instruction;
- compute a partial shift of the input operands;
- increment the partial shift counter;
- add the normal accumulation to the computed partial shift and store the result
- 25 of the addition back in the normal accumulation;
- determine whether to compute an additional partial shift based on the partial shift counter and the total number of partial shifts required; and
- store the normal accumulation in the normal result based on the determination that the additional partial shift need not be computed.


30

22. The arithmetic logic unit of Claim 17, wherein the accelerated computation includes circuitry to:

compute all partial values of input operands of the instruction; and
compute the accelerated result based on a sum of all the computed partial values.

- 5 23. The arithmetic logic unit of Claim 17, further comprising 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.

- 10 24. An apparatus comprising means for performing any of the methods of
Claims 8–23.

**FIG. 1A**

2/28




FIG. 1B

3/28




FIG. 1C




FIG. 2




FIG. 3A







FIG. 3B

7/28


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




FIG. 3D




FIG. 3E




FIG. 3F



FIG. 4A

10/28




FIG. 4B

11/28




FIG. 5A




FIG. 5B

13/28




FIG. 6

14/28




FIG. 7

15/28




FIG. 8

16/28




FIG. 9

17/28




FIG. 10

18/28




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15





FIG. 16



1800




FIG. 18




FIG. 19

2000




FIG. 20

2100




FIG. 21

28/28

2200




FIG. 22

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/018883**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/38(2006.01)i, 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 9/38; G06F 9/302; G06F 7/00; G06F 7/50; G06F 7/57

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: arithmetic logic unit, mode field, performance register, select, normal computation, accelerated computation, faster

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2005-0044124 A1 (SCOTT A. VANSTONE) 24 February 2005 See paragraphs [0013]-[0079]; claim 14; and figures 1, 4, 9.	1-24
Y	US 2003-0126178 A1 (DAVID ARNOLD LUICK) 03 July 2003 See paragraphs [0007]-[0018].	1-24
A	US 9081634 B1 (XILINX, INC.) 14 July 2015 See column 11, line 64 - column 16, line 2; and figures 2, 5-7.	1-24
A	US 2008-0189524 A1 (ELAINE POON et al.) 07 August 2008 See paragraphs [0013]-[0036]; and figures 1-2.	1-24
A	US 2006-0206557 A1 (ANNA WING WAH WONG et al.) 14 September 2006 See paragraphs [0078]-[0092]; and figures 7-12.	1-24



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

26 May 2017 (26.05.2017)

Date of mailing of the international search report

26 May 2017 (26.05.2017)

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/US2017/018883

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005-0044124 A1	24/02/2005	AU 7329198 A CA 2286647 A1 CA 2286647 C DE 69811877 D1 DE 69811877 T2 DE 69841492 D1 EP 0976027 A1 EP 0976027 B1 EP 1293891 A2 EP 1293891 A3 EP 1293891 B1 EP 1293891 B8 GB 9707861 D0 JP 04980834 B2 JP 05866128 B2 JP 2001-520775 A JP 2008-059595 A JP 2011-134346 A JP 2014-219994 A US 2002-0136402 A1 US 6230179 B1 US 6266717 B1 US 6349318 B1 US 6735611 B2 US 7424504 B2 WO 98-48345 A1	13/11/1998 29/10/1998 13/01/2015 10/04/2003 18/12/2003 25/03/2010 02/02/2000 05/03/2003 19/03/2003 08/09/2004 03/02/2010 14/04/2010 04/06/1997 18/07/2012 17/02/2016 30/10/2001 13/03/2008 07/07/2011 20/11/2014 26/09/2002 08/05/2001 24/07/2001 19/02/2002 11/05/2004 09/09/2008 29/10/1998
US 2003-0126178 A1	03/07/2003	US 7058678 B2	06/06/2006
US 9081634 B1	14/07/2015	None	
US 2008-0189524 A1	07/08/2008	US 9135017 B2	15/09/2015
US 2006-0206557 A1	14/09/2006	CA 2548327 A1 CA 2548327 C EP 1700231 A2 EP 1700231 B1 EP 2306331 A1 JP 04664311 B2 JP 2007-522699 A US 2005-0144212 A1 US 2006-0190516 A1 US 2006-0195496 A1 US 2006-0212499 A1 US 2006-0230092 A1 US 2006-0230093 A1 US 2006-0230094 A1 US 2006-0230095 A1 US 2006-0230096 A1	21/07/2005 20/10/2015 13/09/2006 17/10/2012 06/04/2011 06/04/2011 09/08/2007 30/06/2005 24/08/2006 31/08/2006 21/09/2006 12/10/2006 12/10/2006 12/10/2006 12/10/2006 12/10/2006

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/018883Patent document
cited in search reportPublication
datePatent family
member(s)Publication
date

US 2006-0288069 A1	21/12/2006
US 2006-0288070 A1	21/12/2006
US 7472155 B2	30/12/2008
US 7840627 B2	23/11/2010
US 7840630 B2	23/11/2010
US 7844653 B2	30/11/2010
US 7849119 B2	07/12/2010
US 7853632 B2	14/12/2010
US 7853634 B2	14/12/2010
US 7853636 B2	14/12/2010
US 7860915 B2	28/12/2010
US 7865542 B2	04/01/2011
US 7870182 B2	11/01/2011
US 7882165 B2	01/02/2011
WO 2005-066832 A2	21/07/2005
WO 2005-066832 A3	19/01/2006