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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 October 2017 (05.10.2017)

WIPO | PCT

(10) International Publication Number
WO 2017/172124 A1

(51) International Patent Classification:
G06F 13/24 (2006.01) *G06F 13/14* (2006.01)
G06F 15/16 (2006.01) *G06F 9/48* (2006.01)

(21) International Application Number:
PCT/US2017/018939

(22) International Filing Date:
22 February 2017 (22.02.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/085,734 30 March 2016 (30.03.2016) US

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

(72) Inventor: **JAYAKUMAR, Sarathy**; 5311 NW Skycrest Pkwy, Portland, Oregon 97229 (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).

[Continued on next page]

(54) Title: ARBITER BASED SERIALIZATION OF PROCESSOR SYSTEM MANAGEMENT INTERRUPT EVENTS

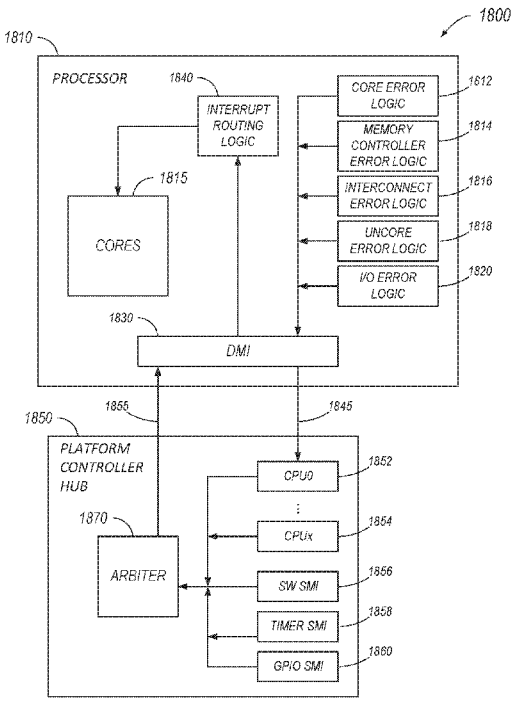


FIG. 18

(57) Abstract: A processor includes cores to execute instructions, and circuitry to detect a system management interrupt (SMI) event on the processor, direct an indication of the SMI event to an arbiter on a controller hub, and receive an interrupt signal from the arbiter. The processor also includes an SMI handler to take action in response to the interrupt, and circuitry to communicate the interrupt signal to the cores. The cores include circuitry to pause while the SMI handler responds to the interrupt. The interrupt handler includes circuitry to determine that a second SMI event detected on the processor or controller hub is pending, and to take action in response. The interrupt handler includes circuitry to set an end-of-SMI bit to indicate that the interrupt handler has completed its actions. The controller includes circuitry to prevent the arbiter from issuing another interrupt to the processor while this bit is false.

WO 2017/172124 A1

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))*

Published:

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

ARBITER BASED SERIALIZATION OF PROCESSOR SYSTEM MANAGEMENT INTERRUPT EVENTS

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.

DESCRIPTION OF RELATED ART

[0002] Multiprocessor systems are becoming more and more common. Applications of multiprocessor systems include dynamic domain partitioning all the way down to desktop computing. In order to take advantage of multiprocessor systems, code to be executed may be separated into multiple threads for execution by various processing entities. Each thread may be executed in parallel with one another. Pipelining of applications may be implemented in systems in order to more efficiently execute applications. System management interrupts and corresponding handlers may be used to manage the system in response to certain types of system errors and system events. Processors may be implemented in a system on chip.

DESCRIPTION OF THE FIGURES

[0003] Embodiments are illustrated by way of example and not limitation in the 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 embodiments of the present disclosure;

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[0030] FIGURE 18 is an illustration of an example system for arbitration based serialization of processor system management interrupt (SMI) events, according to embodiments of the present disclosure;

[0031] FIGURE 19 is an illustration of a portion of a platform controller hub (PCH) that includes SMI serialization logic, according to embodiments of the present disclosure;

[0032] FIGURE 20 is an illustration of a method for generating and responding to processor system management interrupt events, according to embodiments of the present disclosure;

[0033] FIGURE 21 is an illustration of a method for serialization of system management interrupt events by an arbiter, according to embodiments of the present disclosure; and

[0034] FIGURE 22 is an illustration of a method for handling serialized system management interrupt events, according to embodiments of the present disclosure.

DETAILED DESCRIPTION

[0035] The following description describes a processing apparatus and processing logic for arbitration based serialization of processor system management interrupt (SMI) events. Such a processing apparatus may include 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 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 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 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 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 manners. First, as may be useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, designs, at some stage, may reach a level of data representing the physical placement of various devices in the hardware model. In cases wherein some semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. A memory or a magnetic or optical storage such as a disc may be the machine-readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or retransmission of the electrical signal is performed, a new copy may be made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present disclosure.

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

[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).

[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 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 of a common instruction set, but may include different processor designs. For example, the same register architecture of the ISA may be implemented in different ways in different micro-architectures using new or well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB) and a retirement register file. In one embodiment, registers may include one or more registers, register architectures, register files, or other register sets that may or may not be addressable by a software programmer.

[0043] An instruction may include one or more instruction formats. In one embodiment, an instruction format may indicate various fields (number of bits, location of bits, etc.) to specify, among other things, the operation to be performed and the operands on which that operation will be performed. In a further embodiment, some instruction formats may be further defined by instruction templates (or sub-formats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields and/or defined to have a given field interpreted differently. In one embodiment, an instruction may be expressed using an instruction format (and, if defined, in a given one of the instruction templates of that

instruction format) and specifies or indicates the operation and the operands upon which the operation will operate.

[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 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 embodiment, "DEST1" may be a temporary storage register or other storage area, whereas "SRC1" and "SRC2" may be a first and second source storage register or other storage area, and so forth. In other embodiments, two or more of the SRC and DEST storage areas may correspond to different data storage elements within the same storage area (e.g., a SIMD register). In one embodiment, one of the source registers may also act as a destination register by, for example, writing back the result of an operation performed on the first and second source data to one of the two source registers serving as a destination registers.

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

[0048] Embodiments are not limited to computer systems. Embodiments of the present disclosure may be used in other devices such as handheld devices and embedded applications. Some examples of handheld devices include cellular phones, Internet Protocol devices, digital cameras, personal digital assistants (PDAs), and

handheld PCs. Embedded applications may include a micro controller, a digital signal processor (DSP), system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that may perform one or more instructions in accordance with at least one embodiment.

[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 example of a 'hub' system architecture. System 100 may include a processor 102 for processing data signals. Processor 102 may include a complex instruction set computer (CISC) microprocessor, a reduced instruction set computing (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, a processor implementing a combination of instruction sets, or any other processor device, such as a digital signal processor, for example. In one embodiment, processor 102 may be coupled to a processor bus 110 that may transmit data signals between processor 102 and other components in system 100. The elements of system 100 may perform conventional functions that are well known to those familiar with the art.

[0050] In one embodiment, processor 102 may include a Level 1 (L1) internal cache memory 104. Depending on the architecture, the processor 102 may have a single internal cache or multiple levels of internal cache. In another embodiment, the cache memory may reside external to processor 102. Other embodiments may also include a combination of both internal and external caches depending on the particular implementation and needs. Register file 106 may store different types of data in various registers including integer registers, floating point registers, status registers, and instruction pointer register.

[0051] Execution unit 108, including logic to perform integer and floating point operations, also resides in processor 102. Processor 102 may also include a microcode (ucode) ROM that stores microcode for certain macroinstructions. In one embodiment, execution unit 108 may include logic to handle a packed instruction set 109. By including the packed instruction set 109 in the instruction set of a general-purpose processor 102, along with associated circuitry to execute the instructions, the operations

used by many multimedia applications may be performed using packed data in a general-purpose processor 102. Thus, many multimedia applications may be accelerated and executed more efficiently by using the full width of a processor's data bus for performing operations on packed data. This may eliminate the need to transfer smaller units of data across the processor's data bus to perform one or more operations one data element at a time.

[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 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 119 and/or data 121 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 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 storage of instructions 119 and data 121 and for storage of graphics commands, data and textures. MCH 116 may direct data signals between processor 102, memory 120, and other components in system 100 and to bridge the data signals between processor bus 110, memory 120, and system I/O 122. In some embodiments, the system logic chip 116 may provide a graphics port for coupling to a graphics controller 112. MCH 116 may be coupled to memory 120 through a memory interface 118. Graphics card 112 may be coupled to MCH 116 through an Accelerated Graphics Port (AGP) interconnect 114.

[0054] System 100 may use a proprietary hub interface bus 122 to couple MCH 116 to I/O controller hub (ICH) 130. In one embodiment, ICH 130 may provide direct connections to some I/O devices via a local I/O bus. The local I/O bus may include a high-speed I/O bus for connecting peripherals to memory 120, chipset, and processor 102. Examples may include the audio controller 129, firmware hub (flash BIOS) 128, wireless transceiver 126, data storage 124, legacy I/O controller 123 containing user input interface 125 (which may include a keyboard interface), a serial expansion port 127 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 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 random access memory (SRAM) control 147, burst flash memory interface 148, personal computer memory card international association (PCMCIA)/compact flash (CF) card control 149, liquid crystal display (LCD) control 150, direct memory access (DMA) controller 151, and alternative bus master interface 152. In one embodiment, data processing system 140 may also comprise an I/O bridge 154 for communicating with various I/O devices via an I/O bus 153. Such I/O devices may include but are not limited to, for example, universal asynchronous receiver/transmitter (UART) 155, universal serial bus (USB) 156, Bluetooth wireless UART 157 and I/O expansion interface 158.

[0060] One embodiment of data processing system 140 provides for mobile, network and/or wireless communications and a processing core 159 that may perform SIMD operations including a text string comparison operation. Processing core 159 may be programmed with various audio, video, imaging and communications algorithms including discrete transformations such as a Walsh-Hadamard transform, a fast Fourier transform (FFT), a discrete cosine transform (DCT), and their respective inverse transforms; compression/decompression techniques such as color space transformation, video encode motion estimation or video decode motion compensation; and modulation/demodulation (MODEM) functions such as pulse coded modulation (PCM).

[0061] FIGURE 1C illustrates other embodiments of a data processing system that performs SIMD text string comparison operations. In one embodiment, data processing system 160 may include a main processor 166, a SIMD coprocessor 161, a cache memory 167, and an input/output system 168. Input/output system 168 may optionally be coupled to a wireless interface 169. SIMD coprocessor 161 may perform operations

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

[0062] In one embodiment, SIMD coprocessor 161 comprises an execution unit 162 and a set of register files 164. One embodiment of main processor 166 comprises a decoder 165 to recognize instructions of instruction set 163 including instructions in accordance with one embodiment for execution by execution unit 162. In other embodiments, SIMD coprocessor 161 also comprises at least part of decoder 165 (shown as 165B) 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 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 171, these instructions may be received by any attached SIMD coprocessors. In this case, SIMD coprocessor 161 may accept and execute any received SIMD coprocessor instructions intended for it.

[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 communications. For another example, compressed audio and/or video may be received in the form of a digital bit stream, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples and/or motion video frames. In one embodiment of processing core 170, main processor 166, and a SIMD

coprocessor 161 may be integrated into a single processing core 170 comprising an execution unit 162, a set of register files 164, and a decoder 165 to recognize instructions of instruction set 163 including instructions in accordance with one embodiment.

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

[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 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 in allocator/register renamer 215 allocates the machine buffers and resources that each uop needs in order to execute. The register renaming logic in allocator/register renamer 215 renames logic registers onto entries in a register file. The allocator 215 also allocates an entry for each uop in one of the two uop queues, one for memory operations (memory uop queue 207) and one for non-memory operations (integer/floating point uop queue 205), in front of the instruction schedulers: memory scheduler 209, 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 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 data 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 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.

[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.

[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.

[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.

[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 454 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using one or more reorder buffers and one or more retirement register files, using one or more future files, one or more history buffers, and one or more retirement register files; using register maps and a pool of registers; etc.). Generally, the architectural registers may be visible from the outside of the processor or from a programmer's perspective. The registers might not be limited to any known particular type of circuit. Various different types of registers may be suitable as long as they store and provide data as described herein. Examples of

suitable registers include, but might not be limited to, dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. Retirement unit 454 and physical register file units 458 may be coupled to execution clusters 460. Execution clusters 460 may include a set of one or more execution units 462 and a set of one or more memory access units 464. Execution units 462 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. Scheduler units 456, physical register file units 458, and execution clusters 460 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file unit, and/or execution cluster – and in the case of a separate memory access pipeline, certain embodiments may be implemented in which only the execution cluster of this pipeline has memory access units 464). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

[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 caches 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 514 for communications busses for graphics. In one embodiment, interface 514 may be implemented by PCI Express (PCIe). In a further embodiment, interface 514 may be implemented by PCI Express Graphics (PEG). 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 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 from front end 570 for execution. Out-of-order execution engine 580 may include an allocate module 582. In one embodiment, allocate module 582 may allocate resources of processor 500 or other resources, such as registers or buffers, to execute a given instruction. Allocate module 582 may make allocations in schedulers, such as a memory scheduler, fast scheduler, or floating point scheduler. Such schedulers may be represented in FIGURE 5B by resource schedulers 584. Allocate module 582 may be implemented fully or in part by the allocation logic described in conjunction with FIGURE 2. Resource schedulers 584 may determine when an instruction is ready to execute based on the readiness of a given resource's sources and the availability of execution resources needed to execute an instruction. Resource schedulers 584 may be implemented by, for example, schedulers 202, 204, 206 as discussed above. Resource schedulers 584 may schedule the execution of instructions upon one or more resources. In one embodiment, such resources may be internal to core 502, and may be illustrated, for example, as resources 586. In another embodiment, such resources may be external to core 502 and may be accessible by, for example, cache hierarchy 503. Resources may include, for example, memory, caches, register files, or registers. Resources internal to core 502 may be represented by resources 586 in FIGURE 5B. As necessary, values written to or read from resources 586 may be coordinated with other portions of processor 500 through, for example, cache hierarchy 503. As instructions are assigned resources, they may be placed into a reorder buffer 588. Reorder buffer 588 may track instructions as they are executed and may selectively reorder their execution based upon any suitable criteria of processor 500. In one embodiment, reorder buffer 588 may identify instructions or a series of instructions that may be executed independently. Such instructions or a series of instructions may be executed in parallel from other such instructions. Parallel execution in core 502 may be

performed by any suitable number of separate execution blocks or virtual processors. In one embodiment, shared resources—such as memory, registers, and caches—may be accessible to multiple virtual processors within a given core 502. In other embodiments, shared resources may be accessible to multiple processing entities within processor 500.

[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 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.

[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.

[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 a discrete graphics device coupled to ICH 650 along with another peripheral device 670.

[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 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 770, 780 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 732, 734 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 790.

[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 502A-N and shared cache units 506; a system agent unit 510; a bus controller units 916; an integrated memory controller units 914; a set of 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.

[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 performed by the GPU. In still another embodiment, the instruction may be performed through a combination of operations performed by the GPU and the CPU. For example, in one embodiment, an instruction in accordance with one embodiment may be received and decoded for execution on the GPU. However, one or more operations within the decoded instruction may be performed by a CPU and the result returned to the GPU for final retirement of the instruction. Conversely, in some embodiments, the CPU may act as the primary processor and the GPU as the co-processor.

[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 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.

[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 I^2S/I^2C controller 1070. Other logic and circuits may be included in the processor of FIGURE 10, including more CPUs or GPUs and other peripheral interface controllers.

[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 1100 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 1100 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 1165 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.

[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 1411. 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 1460. Flash controller 1445 may provide access to or from memory such as flash memory 1465 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 1565 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 1555, 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 ordering instructions to ensure correct execution semantics of code. A PO may be reconstructed by mechanisms such as evaluating increments to PO encoded in the instruction rather than an absolute value. Such a reconstructed PO may be known as an "RPO." Although a PO may be referenced herein, such a PO may be used

interchangeably with an RPO. A strand may include a sequence of instructions that are data dependent upon each other. The strand may be arranged by a binary translator at compilation time. Hardware executing a strand may execute the instructions of a given strand in order according to the PO of the various instructions. A thread may include multiple strands such that instructions of different strands may depend upon each other. A PO of a given strand may be the PO of the oldest instruction in the strand which has not yet been dispatched to execution from an issue stage. Accordingly, given a thread of multiple strands, each strand including instructions ordered by PO, executed instruction pointer 1580 may store the oldest—illustrated by the lowest number—PO in the thread.

[00132] In another embodiment, memory system 1540 may include a retirement pointer 1582. Retirement pointer 1582 may store a value identifying the PO of the last retired instruction. Retirement pointer 1582 may be set by, for example, retirement unit 454. If no instructions have yet been retired, retirement pointer 1582 may include a null value.

[00133] Execution entities 1565 may include any suitable number and kind of mechanisms by which a processor may execute instructions. In the example of FIGURE 15, execution entities 1565 may include ALU/multiplication units (MUL) 1566, ALUs 1567, and floating point units (FPU) 1568. In one embodiment, such entities may make use of information contained within a given address 1569. Execution entities 1565 in combination with stages 1530, 1550, 1555, 1560, 1570 may collectively form an execution unit.

[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 may perform bus interfacing to other portions of a processor or electronic device. In such an embodiment, unit 1510 may thus include a bus interface unit 1520 for communicating over an interconnect, intraprocessor bus, interprocessor bus, or other communication bus, port, or line. Bus interface unit 1520 may provide interfacing in

order to perform, for example, generation of the memory and input/output addresses for the transfer of data between execution entities 1565 and the portions of a system external to instruction architecture 1500.

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

[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 1540 may include a load store unit 1546 for storing information such as buffers written to or read back from memory or registers. In another embodiment, memory system 1540 may include a translation lookaside buffer (TLB) 1545 that provides look-up of address values between physical and virtual addresses. In yet another embodiment, memory system 1540 may include a memory management unit (MMU) 1544 for facilitating access to virtual memory. In still yet another embodiment, memory system 1540 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 instruction may be performed through different stages. For example, using unit 1510 instruction prefetch stage 1530 may access an instruction through prefetcher 1543. Instructions retrieved may be stored in instruction cache 1532. Prefetch stage 1530 may enable an option 1531 for fast-loop mode, wherein a series of instructions forming a loop that is

small enough to fit within a given cache are executed. In one embodiment, such an execution may be performed without needing to access additional instructions from, for example, instruction cache 1532. Determination of what instructions to prefetch may be made by, for example, branch prediction unit 1535, which may access indications of execution in global history 1536, indications of target addresses 1537, or contents of a return stack 1538 to determine which of branches 1557 of code will be executed next. Such branches may be possibly prefetched as a result. Branches 1557 may be produced through other stages of operation as described below. Instruction prefetch stage 1530 may provide instructions as well as any predictions about future instructions to dual instruction decode stage 1550.

[00138] Dual instruction decode stage 1550 may translate a received instruction into microcode-based instructions that may be executed. Dual instruction decode stage 1550 may simultaneously decode two instructions per clock cycle. Furthermore, dual instruction decode stage 1550 may pass its results to register rename stage 1555. In addition, dual instruction decode stage 1550 may determine any resulting branches from its decoding and eventual execution of the microcode. Such results may be input into branches 1557.

[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 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 suitability of resources for execution of a given instruction. In one embodiment, issue stage 1560 may reorder the instructions within instruction queue 1561 such that the first instructions received might not be the first instructions executed. Based upon the ordering of instruction queue 1561, additional branching information may be provided

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

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

within a single clock cycle. At 1665, floating point arithmetic may be performed by one or more FPUs. The floating point operation may require multiple clock cycles to execute, such as two to ten cycles. At 1670, multiplication and division operations may be performed. Such operations may be performed in four clock cycles. At 1675, loading and storing operations to registers or other portions of pipeline 1600 may be performed. The operations may include loading and storing addresses. Such operations may be performed in four clock cycles. At 1680, write-back operations may be performed as required by the resulting operations of 1655-1675.

[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. 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 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 I^2C bus, system management bus (SMBus), low pin count (LPC) bus, SPI, high definition audio (HDA) bus, Serial Advance Technology Attachment (SATA) bus, USB bus (versions 1, 2, 3), or Universal Asynchronous Receiver/Transmitter (UART) bus.

[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 wireless wide area network (WWAN) unit 1756, a global positioning system (GPS) 1775, 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 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 1736, and touch pad 1730 may be communicatively coupled to EC 1735. Speakers 1763, headphones 1764, and a microphone 1765 may be communicatively coupled to an audio unit 1762, which may in turn be communicatively coupled to DSP 1760. Audio unit 1762 may include, for example, an audio codec and a class D amplifier. A SIM card 1757 may be communicatively coupled to WWAN unit 1756. Components such as WLAN unit 1750 and Bluetooth unit 1752, as well as WWAN unit 1756 may be implemented in a next generation form factor (NGFF).

[00148] Embodiments of the present disclosure involve processing logic or circuitry for serializing system management interrupt (SMI) events, including serializing SMIs that originate in a processor. FIGURE 18 is an illustration of an example system 1800 for arbitration based serialization of processor system management interrupt events, according to embodiments of the present disclosure. In general, a system management interrupt may be an unmaskable interrupt to all of the processors or cores in a given system. Following the issuance of an SMI, all of the processors or cores in the system may stop what they are doing and enter a system management mode. While in this mode, an SMI handler may be invoked to handle the interrupt, after which the processors or cores will resume operation at the point at which they were interrupted.

[00149] System 1800 may include a processor, SoC, integrated circuit, or other mechanism. For example, system 1800 may include processor 1810. Although processor 1810 is shown and described as an example in FIGURE 18, any suitable mechanism may be used. For example, some or all of the functionality of processor 1804 described herein may be implemented by circuitry, instructions for reconfiguring circuitry, a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor having more, fewer, or different elements than those illustrated in FIGURE 18. Processor 1810 may include any suitable mechanisms for serializing system management interrupt (SMI) events, including serializing SMIs that originate in a processor. In at least some embodiments, such mechanisms may be implemented in hardware. For example, in some embodiments, some or all of the elements of processor 1804 illustrated in FIGURE 18

and/or described herein may be implemented fully or in part using hardware circuitry. In some embodiments, this circuitry may include static (fixed-function) logic devices that collectively implement some or all of the functionality of processor 1804. In other embodiments, this circuitry may include programmable logic devices, such as field programmable logic gates or arrays thereof, that collectively implement some or all of the functionality of processor 1804. In still other embodiments, this circuitry may include static, dynamic, and/or programmable memory devices that, when operating in conjunction with other hardware elements, implement some or all of the functionality of processor 1804. For example, processor 1804 may include a hardware memory having stored therein instructions which may be used to program system 1800 to perform one or more operations according to embodiments of the present disclosure. Embodiments of system 1800 and processor 1804 are not limited to any specific combination of hardware circuitry and software. Processor 1810 may be implemented fully or in part by the elements described in FIGURES 1-17. Processor 1810 may include one or more cores 1815. Processor 1810 may also include circuitry or logic to implement the functionality of interrupt routing logic 1840, as described herein.

[00150] System 1800 may include a platform controller hub (PCH) 1850, which may be used to couple various devices and components of system 1800 to processor 1810. In one embodiment, PCH 1850 may include an input/output (I/O) controller hub (ICH), which may be used to couple various peripheral devices to one or more processors 1810 in system 1800. For example, PCH 1850 may include circuitry or logic to implement the functionality of ICH 130 shown in FIGURE 1A or ICH 650 shown in FIGURE 6. In another embodiment, PCH 1850 may include a memory controller hub (MCH), which may provide a high bandwidth path to memory within system 1800. For example, PCH 1850 may include circuitry or logic to implement the functionality of MCH 116 shown in FIGURE 1A or GMCH 620 shown in FIGURE 6. In yet another embodiment, PCH 1850 may include circuitry or logic to implement functionality other than that provided by an I/O controller hub or memory controller hub within system 1800.

[00151] In some embodiments, there may be multiple sources for SMI events that originate on a processor 1810, each of which may be implemented by circuitry or logic. For example, in one embodiment, the SMI event sources on the processor may include

core error logic 1812. In one embodiment, the SMI event sources on the processor may include memory controller error logic 1814. In one embodiment, the SMI event sources on the processor may include interconnect error logic 1816. In one embodiment, the SMI event sources on the processor may include uncore error logic 1818. In one embodiment, the SMI event sources on the processor may include input/output (I/O) error logic 1820.

[00152] In embodiments of the present disclosure, a respective indication of each SMI event that originates on processor 1810 (shown as indication 1845 in FIGURE 18) may be directed to SMI processing circuitry or logic on PCH 1850 for arbitration and serialization. In one embodiment, processor 1810 may communicate with PCH 1850 over a direct media interface (DMI) 1830. For example, indications (1845) of SMI events that originate on processor 1810 may be communicated to PCH 1850 over DMI link 1830. In one embodiment, DMI link 1830 may be implemented within a system agent on processor 1810 (not shown). For example, DMI link 1830 may include circuitry or logic to implement the functionality of DMI 516 shown in FIGURE 5A within a system agent 510 on processor 1810.

[00153] In one embodiment, the SMI processing logic on PCH 1850 may include circuitry or logic to arbitrate between any pending SMIs that originate on processor and/or that originate from SMI event sources on PCH 1850 and to serialize SMI signals communicated to processor 1810. In one embodiment, only one system management interrupt signal 1855 may be sent to the cores 1815 of processor 1810 at a time. For example, PCH 1850 may include an arbiter 1870 to serialize the communication of SMI signals to the cores 1815 of processor 1810 for handling. Each SMI issued to processor 1810 may trigger the handling of one pending SMI event or multiple pending SMI events on processor 1810. For example, in one embodiment, once arbiter 1870 on PCH 1850 asserts an SMI signal 1855, no other SMIs will be issued to processor 1810 until all of the pending SMI events have been handled by processor 1810.

[00154] In some embodiments, when an SMI event that originates from an SMI event source on processor 1810 is detected, rather than routing an indication of the SMI event directly to interrupt routing logic 1840 on processor 1810 for dissemination to cores 1815, an indication of the SMI event (shown as 1845) may be sent to PCH 1850 over DMI 1830. Processor 1810 may be one of multiple processors in system 1800 that

are communicatively coupled to PCH 1850. In one embodiment, these processors may all be similar to processor 1810. In another embodiment, at least two of the processors may be different from each other. In some embodiments, each of multiple processors within system 1800 (including processor 1810) may communicate with PCH 1850 through a respective direct media interface (DMI) link 1830. In embodiments in which system 1800 includes multiple processors, the indication of the SMI event 1845 that is directed to PCH 1850 in response to detecting an SMI event on the processor may include an identifier of the processor on which the SMI event originated. In another embodiment, the indication of the SMI event 1845 that is directed to PCH 1850 in response to detecting an SMI event on the processor may include an identifier of a core or thread that is associated with, or affected by, the detected SMI event. In yet another embodiment, the indication of the SMI event 1845 that is directed to PCH 1850 in response to detecting an SMI event on the processor may include an identifier of a socket in which the processor resides.

[00155] As noted above, PCH 1850 may include an arbiter 1870 to arbitrate and serialize SMIs to be handled by processor 1810, as described in detail herein. In one embodiment, PCH 1850 may include an SMI status register that includes a respective field for each of multiple possible SMI event sources. The value of each status register field may include whether or not an SMI event originating from a particular SMI event source is pending. In another embodiment, PCH 1850 may include circuitry or logic representing multiple SMI status indicators that are individually accessible by other elements of PCH 1850 rather than being fields of a single SMI status register. In one embodiment, the SMI status bits or indicators may be addressable by a software programmer. In other embodiments, the SMI status bits or indicators may not be addressable by a software programmer.

[00156] In embodiments of the present disclosure, one or more SMI event sources may reside on processor 1810. For example, an SMI status bit 1852 in the SMI status register may indicate whether or not an SMI event that originated on a particular processor or processor core (shown as CPU0) is pending. In this example, SMI status bit 1852 may be set by SMI processing circuitry or logic within PCH 1850 in response to the receipt of an indication of an SMI event 1845 that originated on the particular processor or processor core associated with SMI status bit 1852, regardless of the

specific source of the SMI event on the processor or processor core. In some embodiments, if an SMI event is detected in or by core error logic 1812, memory controller error logic 1814, interconnect error logic 1816, or I/O error logic 1820 on processor 1810, an indication 1845 that an SMI event has originated on processor 1810 may be sent to PCH 1850. In one embodiment, the indication of the SMI event 1845 may include an identifier of processor 1810, but may not include an identifier of the specific source of the SMI event. This may allow the SMI processing circuitry or logic within PCH 1850 to identify the particular SMI status bit or indicator that is to be set in response to the SMI event. In another embodiment, the indication of the SMI event 1845 may include an identifier of a particular socket, core or thread that is associated with, or affected by, the SMI event. In yet another embodiment, the indication of the SMI event 1845 may include an identifier of the specific SMI event source on the processor or the type of the SMI event.

[00157] In the example embodiment illustrated in FIGURE 18, one or more other SMI status bits 1854 in the SMI status register may indicate whether or not an SMI event that originated on a respective different processor or processor core (shown as CPUx) is pending. While the example embodiment illustrated in FIGURE 18 includes one SMI status bit or indicator for all SMI events that originate on a given processor, regardless of their source within the processor, in other embodiments, PCH 1850 may include a respective SMI status bit or indicator for each SMI event source on the processor.

[00158] In some embodiments, one or more SMI event sources may reside outside of processor 1810 within system 1800. The value of a respective SMI status bit on PCH 1850 may indicate whether or not there is a pending SMI event associated with each such SMI event source. For example, in one embodiment, a bit 1856 in the SMI status register may indicate whether or not an SMI event that was triggered by software is pending. In one embodiment, a bit 1858 in the SMI status register may indicate whether or not an SMI event that was triggered on or by a timer is pending. In one embodiment, a bit 1860 in the SMI status register may indicate whether or not an SMI event that was triggered by general purpose input/output circuitry or logic on PCH 1850 is pending. In other embodiments, system 1800 may include more, fewer or

different SMI event sources and PCH 1850 may include more, fewer or different SMI status bits corresponding to those SMI event sources.

[00159] In some embodiments, arbiter 1870 on PCH 1850 may arbitrate and serialize any and all pending SMI events for handling by processor 1810, regardless of their sources. For example, arbiter 1870 arbitrate between any pending SMIs that originate on processor and/or on PCH 1850 and may issue a single interrupt signal to the processor at a time, regardless of the number of SMI events that are pending. In one embodiment, an SMI handler on the processor may then handle all of the pending SMI events before effecting an indication that the SMI handler has completed its actions in response to the interrupt signal. For example, an End-of-SMI (EOS) status bit or indicator on PCH 1850 may be cleared by the SMI processing circuitry or logic on PCH 1850 when it issues an interrupt to the processor. The SMI handler may then set the EOS status bit or indicator on PCH 1850 once it has finished handling all pending SMI events. In this way, the communication of SMI signals to the cores 1815 of processor 1810 may be serialized. In one embodiment, if there is at least one SMI event pending and the EOS status bit or indicator on PCH 1850 is set (true), PCH 1850 may issue a system management interrupt (SMI) signal 1855 to processor 1810 over DMI 1830. However, if the EOS status bit or indicator on PCH 1850 is clear (false), no interrupt signal will be issued to the processor 1810 even if one or more SMI events are pending. Upon its receipt by processor 1810, system management interrupt signal 1845 may be directed to interrupt routing logic 1840. Interrupt routing logic 1840 may then disseminate the interrupt signal to all of the cores 1815 of processor 1810.

[00160] In one embodiment, PCH 1850 may include SMI processing logic to serialize SMIs. This serialization logic may reside in whole or in part within arbiter 1870, in different embodiments. FIGURE 19 is an illustration of a portion of a platform controller hub (PCH) 1850 that includes SMI serialization logic 1900, according to embodiments of the present disclosure. In one embodiment, SMI serialization logic 1900 may include circuitry or logic 1910 to perform an OR function whose inputs include all of the fields of an SMI status register. For example, any or all of the SMI status register fields illustrated in FIGURE 18 as SMI status bits 1852, 1854, 1856, 1858 or 1860 may be inputs to 1910. In other embodiments, system 1800 may include more, fewer or different SMI event sources and corresponding SMI status

bits or indicators. In the example embodiment illustrated in FIGURE 19, if any one or more of the inputs to logic 1910 are true, the output of 1910 will be true. For example, if any of SMI status bits 1852, 1854, 1856, 1858 or 1860 has a value of "1", the output of 1910 will also be "1". This may indicate that at least one SMI event is pending.

[00161] In some embodiments, SMI serialization logic 1900 may include circuitry or logic 1920 to perform an AND function. The inputs to the AND function may include the output of 1910 and an input representing the value of an End-of-SMI (EOS) bit 1930. In one embodiment, EOS bit 1930 may be set (e.g., it may have a value of 1, or true) by default. The EOS bit 1930 may remain set until and unless something occurs to clear it. In one embodiment, EOS bit 1930 may be cleared (e.g., to a value of 0, or false) by hardware when an SMI is issued to processor 1810, which may prevent any other interrupts being issued to the processor based on SMI events. In one embodiment, the output of 1920 may be SMI signal 1855. In one embodiment, SMI signal 1855, which is also shown in FIGURE 18, may be communicated to processor 1810 over DMI 1830. In the example embodiment illustrated in FIGURE 19, if both the output of 1910 and the value of EOS bit 1930 are true, the output of 1920 (SMI signal 1855) will be true. In this case, the SMI signal 1855 will be issued to processor 1810 through DMI 1830, triggering an interrupt on processor 1810. However, if the value of EOS bit 1930 is not true, the output of 1920 (shown as SMI signal 1855) will be false. In this case, the SMI signal 1855 will not be asserted. Therefore, no interrupt will be issued to processor 1810, regardless of whether any SMI events are pending (as indicated by the output of 1910).

[00162] In one embodiment, circuitry/logic 1910, circuitry/logic 1920, and EOS bit 1930 may be implemented within arbiter 1870 shown in FIGURE 18. In other embodiments, at least a portion of circuitry/logic 1910, circuitry/logic 1920, and/or EOS bit 1930 may be implemented outside of arbiter 1870 on PCH 1850. In embodiments in which the output of circuitry/logic 1920 is generated outside of arbiter 1870, this output may be provided to arbiter 1870, triggering the assertion of SMI signal 1855 by arbiter 1870.

[00163] In some embodiments, PCH 1850 may include circuitry or logic to control the value of EOS bit 1930 (not shown). PCH 1850 may also include circuitry or logic to set and/or clear the respective SMI status bits for each of the SMI event sources (not

shown). In some embodiments, by controlling these inputs to SMI serialization logic 1900, PCH 1850 may serialize the system management interrupts that are communicated to interrupt routing logic 1840 and, in turn, to the cores 1815 of processor 1810.

[00164] In embodiments of the present disclosure, the mechanisms utilized in serializing processor SMI events may be described as follows:

1. The SMI arbiter in the PCH may include circuitry or logic to comprehend processor-generated SMIs in its arbitration along with PCH-generated SMIs.
2. The processor uncore may include circuitry or logic to send an SMI indication for a processor-generated SMI to the PCH. This indication may include the processor ID.
3. The PCH may include circuitry or logic to set the status bit in the SMI status register that corresponds to the processor ID.
4. The PCH may include circuitry or logic to clear the EOS bit and assert the SMI signal.
5. The processor may include circuitry or logic to deliver the SMI to all of the cores on all of the processors in the system.
6. The SMI handler may include circuitry or logic to set the EOS bit once the SMI event has been handled and the corresponding status bits have been cleared. This may indicate to the PCH that it can release further SMIs. For example, this may allow the SMI arbiter to reassert the SMI signal upon detection of another SMI event.

[00165] FIGURE 20 is an illustration of a method 2000 for generating and responding to system management interrupt events that originate on a processor, according to embodiments of the present disclosure. Method 2000 may be implemented by any of the elements shown in FIGURES 1-19. Method 2000 may be initiated by any suitable criteria and may initiate operation at any suitable point. In one embodiment, method 2000 may initiate operation at 2005. Method 2000 may include greater or fewer steps than those illustrated. Moreover, method 2000 may execute its steps in an order different than those illustrated below. Method 2000 may terminate at any suitable step. Moreover, method 2000 may repeat operation at any suitable step. Method 2000

may perform any of its steps in parallel with other steps of method 2000, or in parallel with steps of other methods.

[00166] At 2005, in one embodiment, a System Management Interrupt (SMI) event may be detected on a processor. At 2010, an indication of the SMI event may be directed to a Platform Controller Hub (PCH) for arbitration and serialization. The indication of the SMI event may include an identifier of the affected CPU. In one example, the indication of the SMI event may include an identifier of the processor on which the SMI event was detected. In another example, the indication of the SMI event may include an identifier of a particular core on the processor that is associated with the SMI event. In another example, the indication of the SMI event may include an identifier of particular thread of execution that is associated with the SMI event. In yet another example, the indication of the SMI event may include an identifier of the socket in which the processor on which the SMI event was detected resides.

[00167] In one embodiment, while no interrupt signal is received from the PCH (at 2015) and while no additional SMI events are detected (at 2020), the operation of the processor may continue without taking any action in response to the detected SMI event. In one embodiment, if, prior to an interrupt signal being received (at 2015), one or more other SMI events is detected (at 2020), an indication of each of the additional SMI events may be directed to a Platform Controller Hub (PCH) for arbitration and serialization. The indication of each additional SMI event may include an identifier of the CPU associated with the additional SMI event.

[00168] In one embodiment, if (at 2015) an interrupt signal is received from the PCH, then at 2025, the interrupt may be taken. In one embodiment, taking the interrupt may include delivering an SMI interrupt signal to all of the cores of the processor. Taking the interrupt may also include each of the cores entering System Management Mode (SMM) after completing a currently executing instruction (if execution of an instruction is in progress when the interrupt signal is received). Subsequently, at 2030, an SMI handler may deal with all currently pending SMI events including any that were detected since the interrupt signal was issued. For example, in one embodiment, when entering the SMM, each thread may go into the SMI handler and perform a rendezvous procedure, meaning that it waits for all other threads to come into the SMI handler. Once all of the threads have checked into the SMI handler, one of them may

handle the SMI events. Once the SMI handler has dealt with all currently pending SMI events, it may then set an End-of-SMI (EOS) bit in the PCH to indicate that it has finished taking its actions in response to the interrupt signal, after which the SMI handler may exit. At 2035, the cores may resume normal operation from the point at which they were interrupted. As illustrated in FIGURE 20, any or all of steps 2010 - 2035 of method 2000 may be repeated, as appropriate, if and when any additional SMI events are detected or any additional interrupt signals are received from the PCH.

[00169] In the example embodiment illustrated in FIGURE 20, the processor detected an SMI event that originated on the processor prior to receiving an SMI interrupt signal from the PCH. In other embodiments, the processor may receive an SMI interrupt signal (e.g., one that is triggered based on the detection of an SMI event originating on the PCH) prior to detecting an SMI event that originated on the processor.

[00170] In embodiments of the present disclosure, the handshaking protocol between the arbiter and each of the SMIs may be described as follows:

1. The source of the SMI (e.g., a timer, USB interface, general-purpose I/O component, processor, etc.) is logged in the SMI status register as a bit map.
2. SMI events are delivered to the arbiter on the PCH.
3. The arbiter checks the state of the EOS bit.
4. If the EOS bit is clear (a value of 0, or false), this blocks the delivery of the SMI, since it indicates that the handling of a previous SMI is in progress.
5. If the EOS bit is set (a value of 1, or true), this allows the PCH to deliver an SMI if there are any status bits set in the SMI status register, at which point the PCH clears the EOS bit.
6. The SMI handler queries the SMI status register to determine the SMI event source(s). If the SMI event source is a processor, the SMI handler may query the processor to determine the specific SMI event source on that processor. For example, the source of the SMI event may be the core error logic, memory controller error logic, interconnect error logic, uncore error logic, I/O error logic, or another component of the processor, in different embodiments.

7. The SMI handler handles the SMI events corresponding to bits that are set in the SMI status register. Once handling is complete for an SMI event, the SMI handler clears the corresponding bit in the SMI status register.
8. Once all of the SMI events have been handled, the processor sets the EOS bit on the PCH. For example, the SMI handler may write a value of 1 to the EOS bit. In another example, the BIOS may write a value of 1 to the EOS bit.
9. Once the EOS is set, this indicates to the arbiter that the previous SMIs have been handled. This will then unblock SMI delivery such that, if there are any set bits in the SMI status register, another SMI will be delivered. This mechanism may ensure that any new SMI event that occurs between steps 7 and 8 will not be lost.

[00171] In some embodiments, the SMI handler, when invoked, may handle all of the SMI events that were pending at the time the SMI was issued and any additional SMI events that are detected while the SMI handler is operating to handle the earlier SMI events, regardless of their sources. The SMI events may be handled in any order, in different embodiments. In some embodiments, multiple SMI events may be handled at the same time (e.g., substantially in parallel) by different circuitry or logic within one or more processors or cores. Once the SMI handler determines that all of the pending SMI events from all SMI sources have been handled, it may set the EOS bit and then exit the SMI handler. In some embodiments, if an SMI event was generated on a processor in a system that includes multiple processor sockets, an identifier of the socket or processor may be provided to the PCH. In some embodiments, the PCH may provide per-socket or per-processor SMI status bits and may use the identifier to populate the corresponding bit. This may ensure faster handling of the processor SMI events, in some embodiments. For example, instead of querying all of the processors to determine the specific source of an SMI event when a processor-generated SMI event is detected, the SMI handler may query only the identified processor to determine the specific SMI source.

[00172] FIGURE 21 is an illustration of a method 2100 for serialization of system management interrupt events by an arbiter, according to embodiments of the present disclosure. Method 2100 may be implemented by any of the elements shown in FIGURES 1-19. Method 2100 may be initiated by any suitable criteria and may initiate

operation at any suitable point. In one embodiment, method 2100 may initiate operation at 2105. Method 2100 may include greater or fewer steps than those illustrated. Moreover, method 2100 may execute its steps in an order different than those illustrated below. Method 2100 may terminate at any suitable step. Moreover, method 2100 may repeat operation at any suitable step. Method 2100 may perform any of its steps in parallel with other steps of method 2100, or in parallel with steps of other methods. In one embodiment, method 2100 may be invoked on a PCH 1850 in response to receipt, from a processor 1810, of an indication of a processor SMI event 1845 by PCH 1850.

[00173] At 2105, in one embodiment, an End-of-SMI (EOS) bit in a PCH may be initialized to a value of 1 (true). This may indicate that no SMI event is currently being handled by an interrupt handler (more specifically, an SMI handler) on a processor in the same system. If (at 2010) an SMI event indication is received from the processor, then at 2115, an SMI status bit in the PCH may be set. The SMI event indication may include an identifier of a CPU on which the SMI event was detected, and the SMI status bit that is set may be is associated with the identified CPU. Similarly, if (at 2120) an SMI event is detected on the PCH, then at 2125, a different SMI status bit in the PCH may be set. This SMI status bit may be associated with the source of the SMI event on the PCH. As illustrated in FIGURE 20, until and unless logic on the processor or on the PCH detects an SMI event, operation may continue without asserting an SMI signal to interrupt the processor.

[00174] If (at 2110) an SMI event indication is received from the processor or if (at 2120) an SMI event is detected on the PCH, and if (at 2130), it is determined that the EOS bit is set, an SMI interrupt may be issued to the processor. Issuing the interrupt may include asserting an interrupt signal and clearing the EOS bit to indicate that the handling of the interrupt is in progress. On the other hand, if (at 2110) an SMI event indication is received from the processor or if (at 2120) an SMI event is detected on the PCH, but if (at 2130), it is determined that the EOS bit is not set, operation may continue without issuing an interrupt to the processor. As illustrated in FIGURE 21, any or all of steps 2110 - 2135 of method 2100 may be repeated, as appropriate, if and when any additional SMI events are detected on the processor or on the PCH.

[00175] In embodiments of the present disclosure, the detection of SMI events on the processor and/or on the PCH may occur in any order. However, no interrupt will be issued to the processor for a pending SMI event that originated on the processor or on the PCH until and unless the value of the EOS bit is 1 (true).

[00176] FIGURE 22 is an illustration of a method 2200 for handling serialized system management interrupt events, according to embodiments of the present disclosure. Method 2200 may be implemented by any of the elements shown in FIGURES 1-19. 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 2205. 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. Moreover, 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 parallel with steps of other methods. In one embodiment, method 2200 may be invoked on a processor 1810 in response to receipt, from a PCH 1850, of an SMI signal 1855 by processor 1810.

[00177] At 2205, in one embodiment, all cores of a processor may enter System Management Mode (SMM), and an SMI handler may begin execution. In one embodiment, the cores may pause or halt normal operation while in the SMM. For example, they may refrain from the execution of any instructions in the instruction stream while in the SMM. In one embodiment, the SMI handler may be implemented by circuitry or logic in one of the cores of the processor while all of the cores are in the SMM. In another example, the SMI handler may be implemented by circuitry or logic outside of the cores while the cores are in the SMM. At 2210, the SMI handler may access a status register on the PCH to determine what SMI events are pending.

[00178] At 2215, in one embodiment, it may be determined whether or not there is a pending processor SMI event. If there is at least one pending SMI event that originated on the processor, as indicated by the value of an SMI status bit for the processor, then at 2220, the processor whose status bit is set may be polled for pending SMI events. Action may be taken by an SMI handler to deal with each pending SMI event that originated on the processor. For example, the processor may include multiple SMI event sources, each of which detects an SMI event of a different type. In one

embodiment, the actions taken by the SMI handler to deal with each of the pending SMI events that originated on the processor may be dependent on the source of the SMI event on the processor. In another embodiment, the actions taken by the SMI handler to deal with each of the pending SMI events that originated on the processor may be dependent on the type of the SMI event. Once all of the pending SMI events that originated on the processor have been dealt with, the SMI handler may clear the SMI status bit for the processor (on the PCH).

[00179] At 2225, in one embodiment, it may be determined whether or not there is a pending SMI event from an SMI event source on the PCH. If there is at least one pending SMI event that originated on the PCH, as indicated by a true value of one or more SMI status bits associated with SMI event sources on the PCH, then at 2230, the SMI handler may handle each SMI event from a PCH source whose SMI status bit is set. For example, the SMI handler may, for each pending SMI event, take one or more actions that are dependent on the source of the SMI event. Once all pending SMI events from a particular SMI event source on the PCH have been dealt with, the SMI handler may clear the SMI status bit associated with that SMI event source.

[00180] While (at 2235) there are more additional pending SMI events to handle, some or all of steps 2215 - 2230 may be repeated, as appropriate, to discover and handle each additional pending SMI event. If, or once (at 2235), there are no additional pending SMI events to handle, then at 2240, the SMI handler may set the EOS bit in PCH, after which the SMI handler may exit.

[00181] In some embodiments, methods 2100 and/or 2200 may be invoked based on events that occur and/or actions that are taken during execution of method 2000. For example, methods 2100 and/or 2200 may be invoked to arbitrate, serialize, and handle one or more SMI events, including SMI events that originate on the processor.

[00182] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may reduce or eliminate complications that occur in some existing systems due to SMIs. For example, in some existing server systems, SMIs are used extensively for firmware-based error handling and for handling various reliability, availability, and serializability events. In today's high core count server systems, such firmware-first error handling and reliance on SMIs may lead to complex corner cases and race conditions that require convoluted workarounds. In

some existing systems, these complex corner cases and race conditions are caused by the fact that the processor-generated SMIs are not serialized.

[00183] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may reduce or eliminate the SMI storms that occur in some existing systems. For example, in some existing systems, all processor-generated SMI events are delivered to the cores even if the core is already processing a previous SMI. Since there is no arbiter for these SMI events, multiple back-to-back SMIs can be delivered to the processor, causing an SMI storm with which the SMI handler cannot keep up. The systems described herein may serialize the SMI delivery, which may reduce or eliminate such SMI storms by allowing only a single SMI signal to be issued at a time, the handling of which may include handling multiple pending SMI events.

[00184] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may reduce or eliminate the types of complications that can be introduced by the merging of SMIs that is performed in some existing systems. For example, in some existing systems, there are situations in which there can be multiple back-to-back SMI events happening in rapid succession, and those events can get lost due to a mechanism called SMM merge. In these systems, each thread has one SMI pending bit. If more than one SMI is detected within an instruction boundary, the SMIs are merged into one SMI. In such systems, when two back-to-back SMIs occur within an instruction boundary in one thread, the two SMIs are merged, whereas if these same two SMI events happened to hit the thread across an instruction boundary, the thread will observe two separate SMI events. This situation can throw the SMI processing out of sync, leading to various corner cases. The systems described herein may serialize the SMI delivery, which may reduce or eliminate these types of SMM merging issues.

[00185] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may reduce the amount of time that the cores of the processor spend in system management mode in response to SMI events. For example, interrupting the processor fewer times and allowing the SMI handler to handle all currently pending SMIs each time all the cores are stopped in response to an interrupt, rather than repeatedly interrupting the processor and waiting for the cores to

rendezvous in the SMI handler in order to handle individual SMIs, may allow the processor to spend less time in the SMM and more time performing operating system and application workloads. This may reduce the number of "time-outs" observed by operating system processes and/or applications while the cores handle SMIs. In another example, by limiting the number of processors that are queried by the SMI handler to those that actually generated an SMI event (as indicated by corresponding SMI status bits on the PCH), the time spent by the SMI handler to identify the sources of all currently pending SMI events may be reduced.

[00186] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may provide consistent semantics for all of the SMI events in the system, regardless of their sources. The arbiter on the PCH may then be used to serialize processor-generated SMIs along with PCH-generated SMIs, such that an SMI is issued to the processor, and sent to its cores, only after the handling of any SMI currently being handled is complete. In some embodiments, these mechanisms may provide a cleaner solution for handling multiple SMI events than in existing systems, may prevent the storms, time-outs, and SMM merge issues than can occur in some existing system, and may avoid complex debug scenarios due to spurious SMI issues.

[00187] In some embodiments, the mechanisms described herein for serializing SMIs, including processor-generated SMIs, may handle multi-processor scenarios seamlessly by utilizing the arbiter in the PCH, which is a common resource shared by all the processor sockets in the system. In some embodiments, the PCH may reside on the same integrated circuit die as one or more processors in the system. In other embodiments, the PCH may reside in the same socket or package as one or more processors in the system. In still other embodiments, the PCH may reside in the same chipset as one or more processors in the system

[00188] Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the disclosure may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

[00189] Program code may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system may include any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

[00190] The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

[00191] One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine-readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

[00192] Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[00193] 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.

[00194] In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part-on and part-off processor.

[00195] Thus, techniques for performing one or more instructions according to at least one embodiment are disclosed. While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on other embodiments, and that such embodiments not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure. In an area of technology such as this, where growth is fast and further advancements are not easily foreseen, the disclosed embodiments may be readily modifiable in arrangement and detail as facilitated by enabling technological advancements without departing from the principles of the present disclosure or the scope of the accompanying claims.

[00196] Some embodiments of the present disclosure include a processor. In at least some of these embodiments, the processor may include a core to execute instructions, logic or circuitry to detect, on the processor, a first system management interrupt (SMI) event of a first SMI event type, logic or circuitry to direct an indication of the first SMI event to an arbiter, logic or circuitry to receive an interrupt signal from the arbiter, an interrupt handler to respond to the interrupt signal, including logic or circuitry to take action in response to detection of the first SMI event, the action to be dependent on the first SMI event type, and logic or circuitry to effect an indication that the interrupt handler has completed its actions in response to the interrupt signal. In combination with any of the above embodiments, the processor may include logic or circuitry to communicate the interrupt signal to the core, and the core may include logic or circuitry

to pause execution of instructions while the interrupt handler responds to the interrupt signal. In combination with any of the above embodiments, the processor may further include logic or circuitry to detect, on the processor and prior to receipt of the interrupt signal, a second SMI event of a second SMI event type different from the first SMI event type, and logic or circuitry to direct an indication of the second SMI event to the arbiter, and the interrupt handler may further include logic or circuitry to take action in response to detection of the second SMI event, the action to be dependent on the second SMI event type. In combination with any of the above embodiments, the interrupt handler may further include logic or circuitry to determine the first SMI event type. In combination with any of the above embodiments, the interrupt handler may further include logic or circuitry to determine that a second SMI event is pending, and logic or circuitry to take action in response to the second SMI event. In combination with any of the above embodiments, the arbiter may be a component of a controller hub, and the interrupt handler may further include logic or circuitry to determine that a second SMI event was detected on the controller hub, and logic or circuitry to take action in response to the second SMI event. In combination with any of the above embodiments, the arbiter may be a component of a controller hub, and the processor and the controller hub may reside in a single integrated circuit package. In combination with any of the above embodiments, the arbiter may be a component of a controller hub, and the logic or circuitry may set an end-of-SMI indicator on the controller hub. In combination with any of the above embodiments, the processor may be one of a plurality of processors communicatively coupled to the arbiter, and the indication of the first SMI event may include an identifier of the processor. In combination with any of the above embodiments, the processor may include a plurality of SMI event sources, and the interrupt handler may further include logic or circuitry to determine the source of the first SMI event on the processor. In combination with any of the above embodiments, the arbiter may be a component of a controller hub, the controller hub may include at least one SMI event source, and the interrupt handler may further include logic or circuitry to determine that a second SMI event was detected on the controller hub, and logic or circuitry to determine the source of the second SMI event on the controller hub. In combination with any of the above embodiments, processor may further include one or more SMI event sources, including core error logic or circuitry, memory controller

error logic or circuitry, interconnect error logic or circuitry, uncore error logic or circuitry, or input/output error logic or circuitry. In combination with any of the above embodiments, the processor may be a first one of a plurality of processors communicatively coupled to the arbiter, and the interrupt handler may further include logic or circuitry to determine that a second SMI event was detected on a second one of the plurality of processors, and logic or circuitry to determine the source of the second SMI event on the second processor.

[00197] Some embodiments of the present disclosure include a method. The method may include detecting, on a processor, a first system management interrupt (SMI) event of a first SMI event type, directing an indication of the first SMI event to a controller hub, receiving an interrupt signal from the controller hub, executing, in response to receiving the interrupt signal, an interrupt handler, including taking action in response to detecting the first SMI event, the action being dependent on the first SMI event type, and setting an indicator on the controller hub to indicate that the interrupt handler has completed its actions in response to the interrupt signal. In combination with any of the above embodiments, the processor may include a core for executing instructions, and the method may further include communicating the interrupt signal to the core, and pausing, by the core, execution of instructions while the interrupt handler responds to the interrupt signal. In combination with any of the above embodiments, the method may further include detecting, on the processor prior to receiving the interrupt signal, a second SMI event of a second SMI event type different from the first SMI event type, and directing an indication of the second SMI event to the controller hub, and executing the interrupt handler may further include taking action in response to detecting the second SMI event, the action being dependent on the second SMI event type. In combination with any of the above embodiments, executing the interrupt handler may further include determining the first SMI event type. In combination with any of the above embodiments, executing the interrupt handler may further include determining that a second SMI event is pending, and taking action in response to the second SMI event. In combination with any of the above embodiments, executing the interrupt handler may further include determining that a second SMI event was detected on the controller hub, and taking action in response to the second SMI event. In combination with any of the above embodiments, the processor may be one of a plurality of

processors communicatively coupled to the controller hub, and the indication of the first SMI event may include an identifier of the processor. In combination with any of the above embodiments, the processor may include a plurality of SMI event sources, and executing the interrupt handler may further include determining the source of the first SMI event on the processor. In combination with any of the above embodiments, the controller hub may include at least one SMI event source, and executing the interrupt handler may further include determining that a second SMI event was detected on the controller hub, and determining the source of the second SMI event on the controller hub. In combination with any of the above embodiments, the controller hub may include a respective source for each of one or more SMI event types, including a software SMI event type, a timer SMI event type, or an input/output SMI event type. In combination with any of the above embodiments, the processor may include one or more SMI event sources, including core error logic or circuitry, memory controller error logic or circuitry, interconnect error logic or circuitry, uncore error logic or circuitry, or input/output error logic or circuitry. In combination with any of the above embodiments, the processor may be a first one of a plurality of processors communicatively coupled to the controller hub, and executing the interrupt handler may further include determining that a second SMI event was detected on a second one of the plurality of processors, and determining the source of the second SMI event on the second processor. In combination with any of the above embodiments, executing the interrupt handler may further include taking, prior to setting the indicator on the controller hub, a respective action in response to at least one additional pending SMI event, the action being dependent on the type of the additional pending SMI event, and receiving, from the controller hub subsequent to setting the indicator on the controller hub, a second interrupt signal.

[00198] Some embodiments of the present disclosure include a system. In at least some of these embodiments, the system may include a controller hub, and a first processor. The first processor may include a core to execute instructions, logic or circuitry to detect a first system management interrupt (SMI) event from a first SMI event source on the first processor, logic or circuitry to direct an indication of the first SMI event to the controller hub, logic or circuitry to receive an interrupt signal from the controller hub, an interrupt handler to respond to the interrupt signal, including logic or

circuitry to take action in response to detection of the first SMI event, the action to be dependent on the source of the first SMI event, and logic or circuitry to set an indicator on the controller hub to indicate that the interrupt handler has completed its actions in response to the interrupt signal. In combination with any of the above embodiments, the first processor may further include logic or circuitry to communicate the interrupt signal to the core, the core to include logic or circuitry to pause execution of instructions while the interrupt handler responds to the interrupt signal. In combination with any of the above embodiments, the first processor may further include logic or circuitry to detect, prior to receipt of the interrupt signal, a second SMI event from a second SMI event source on the first processor different from the first SMI event source, and logic or circuitry to direct an indication of the second SMI event to the controller hub, and the interrupt handler may further include logic or circuitry to take action in response to detection of the second SMI event, the action to be dependent on the source of the second SMI event. In combination with any of the above embodiments, the interrupt handler may further include logic or circuitry to determine that a second SMI event is pending, and logic or circuitry to take action in response to the second SMI event. In combination with any of the above embodiments, the controller hub may include a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor, a second SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the controller hub, and the interrupt handler may further include logic or circuitry to determine, dependent on the second SMI status bit, that a second SMI event was detected on the controller hub, and logic or circuitry to take action in response to the second SMI event. In combination with any of the above embodiments, the controller hub may include an end-of-SMI indicator whose value indicates whether or not the interrupt handler has completed its actions in response to an interrupt issued to the first processor by the controller hub, logic or circuitry to prevent the controller hub from issuing an interrupt to the first processor while the end-of-SMI indicator is false, logic or circuitry to determine that at least one SMI event is pending, logic or circuitry to issue an interrupt to the first processor in response to determining that at least one SMI event is pending and that the end-of-SMI indicator is true, including logic or circuitry to clear the end-of-SMI indicator, and logic or circuitry

to assert the interrupt signal. In combination with any of the above embodiments, the system further may include a second processor, the indication of the first SMI event may include an identifier of the first processor, the controller hub may include a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor, and logic or circuitry to determine, dependent on receipt of the indication of the first SMI event and the identifier of the first processor, that the first SMI status bit is to be set. In combination with any of the above embodiments, the controller hub may include a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor, the interrupt handler may further include logic or circuitry to determine the source of the first SMI event on the first processor. In combination with any of the above embodiments, the controller hub may include a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor, a second SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the controller hub, logic or circuitry to determine that a second SMI event was detected on the controller hub, and logic or circuitry to set the second SMI status bit. In combination with any of the above embodiments, the system further may include a second processor, the controller hub may include a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor, a second SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the second processor, and logic or circuitry to determine, dependent on receipt of an indication of a second SMI event including an identifier of the second processor, that the second SMI status bit is to be set. In combination with any of the above embodiments, the interrupt handler may further include logic or circuitry to determine the first SMI event type. In combination with any of the above embodiments, the logic or circuitry may set an end-of-SMI indicator on the controller hub. In combination with any of the above embodiments, the first processor may be one of a plurality of processors in the system, and the indication of the first SMI event may include an identifier of the first processor. In combination with any of the above embodiments, the first processor may include a plurality of SMI event sources, and the interrupt handler may further include logic or circuitry to determine the source of the

first SMI event on the first processor. In combination with any of the above embodiments, the interrupt handler may further include logic or circuitry to determine that a second SMI event was detected on the controller hub, and logic or circuitry to determine the source of the second SMI event on the controller hub. In combination with any of the above embodiments, the controller hub may include a respective source for each of one or more SMI event types, including a software SMI event type, a timer SMI event type, or an input/output SMI event type. In combination with any of the above embodiments, the processor further may include one or more SMI event sources, including core error logic or circuitry, memory controller error logic or circuitry, interconnect error logic or circuitry, uncore error logic or circuitry, or input/output error logic or circuitry. In combination with any of the above embodiments, the controller hub may include an arbiter to assert the interrupt signal. In combination with any of the above embodiments, the first processor and the controller hub may reside in a single integrated circuit package. In combination with any of the above embodiments, the first processor and the controller hub may reside on a single integrated circuit die.

[00199] Some embodiments of the present disclosure include a system for executing instructions. In at least some of these embodiments, the system may include means for detecting, on a processor, a first system management interrupt (SMI) event of a first SMI event type, means for directing an indication of the first SMI event to controller hub, means for receiving an interrupt signal from the controller hub, means for executing, in response to receiving the interrupt signal, an interrupt handler, including means for taking action in response to detecting the first SMI event, the action being dependent on the first SMI event type, and means for setting an indicator on the controller hub to indicate that the interrupt handler has completed its actions in response to the interrupt signal. In combination with any of the above embodiments, the processor may include a core for executing instructions, and the apparatus may further include means for communicating the interrupt signal to the core, and means for pausing, by the core, execution of instructions while the interrupt handler responds to the interrupt signal. In combination with any of the above embodiments, the apparatus may further include means for detecting, on the processor prior to receiving the interrupt signal, a second SMI event of a second SMI event type different from the first SMI event type, and means for directing an indication of the second SMI event to the

controller hub, the means for executing the interrupt handler may further include means for taking action in response to detecting the second SMI event, the action being dependent on the second SMI event type. In combination with any of the above embodiments, the means for executing the interrupt handler may further include means for determining the first SMI event type. In combination with any of the above embodiments, the means for executing the interrupt handler may further include means for determining that a second SMI event is pending, and means for taking action in response to the second SMI event. In combination with any of the above embodiments, the means for executing the interrupt handler may further include means for determining that a second SMI event was detected on the controller hub, and means for taking action in response to the second SMI event. In combination with any of the above embodiments, the processor may be one of a plurality of processors communicatively coupled to the controller hub, and the indication of the first SMI event may include an identifier of the processor. In combination with any of the above embodiments, the processor may include a plurality of SMI event sources, and the means for executing the interrupt handler may further include means for determining the source of the first SMI event on the processor. In combination with any of the above embodiments, the controller hub may include at least one SMI event source, and the means for executing the interrupt handler may further include means for determining that a second SMI event was detected on the controller hub, and means for determining the source of the second SMI event on the controller hub. In combination with any of the above embodiments, the controller hub may include a respective source for each of one or more SMI event types, including a software SMI event type, a timer SMI event type, or an input/output SMI event type. In combination with any of the above embodiments, the processor may include one or more SMI event sources, including core error logic or circuitry, memory controller error logic or circuitry, interconnect error logic or circuitry, uncore error logic or circuitry, or input/output error logic or circuitry. In combination with any of the above embodiments, the processor may be a first one of a plurality of processors communicatively coupled to the controller hub, and the means for executing the interrupt handler may further include means for determining that a second SMI event was detected on a second one of the plurality of processors, and means for determining the source of the second SMI event

on the second processor. In combination with any of the above embodiments, the means for executing the interrupt handler may further include means for taking, prior to setting the indicator on the controller hub, a respective action in response to at least one additional pending SMI event, the action being dependent on the type of the additional pending SMI event, and means for receiving, from the controller hub subsequent to setting the indicator on the controller hub, a second interrupt signal.

CLAIMS

What is claimed is:

1. A processor, comprising:

a core including circuitry to execute instructions; and

circuitry to:

detect, on the processor, a first system management interrupt (SMI) event of a first SMI event type;

direct an indication of the first SMI event to an arbiter that is outside the processor;

receive an interrupt signal from the arbiter;

an interrupt handler to respond to the interrupt signal, including circuitry to:

take action in response to detection of the first SMI event, the action to be dependent on the first SMI event type; and

effect an indication that the interrupt handler has completed its actions in response to the interrupt signal.

2. The processor of Claim 1, wherein:

the processor further comprises circuitry to:

communicate the interrupt signal to the core; and

the core includes circuitry to:

pause execution of instructions while the interrupt handler responds to the interrupt signal.

3. The processor of Claim 1, wherein:

the processor further comprises circuitry to:

detect, on the processor and prior to receipt of the interrupt signal, a second SMI event of a second SMI event type different from the first SMI event type; and

direct an indication of the second SMI event to the arbiter;

the interrupt handler further includes circuitry to:

take action in response to detection of the second SMI event, the action to be dependent on the second SMI event type.

4. The processor of Claim 1, wherein:

the interrupt handler further comprises circuitry to:

determine the first SMI event type.

5. The processor of Claim 1, wherein:

the interrupt handler further comprises circuitry to:

determine that a second SMI event is pending; and

take action in response to the second SMI event.

6. The processor of Claim 1, wherein:

the arbiter is a component of a controller hub that arbitrates between SMI events that originate on both the processor and on the controller hub; and

the interrupt handler further comprises circuitry to:

determine that a second SMI event was detected on the controller hub; and

take action in response to the second SMI event.

7. The processor of Claim 1, wherein:

the arbiter is a component of a controller hub; and

the processor and the controller hub reside in a single integrated circuit package.

8. The processor of Claim 1, wherein:

the processor comprises a plurality of SMI event sources; and

the interrupt handler further comprises circuitry to:

determine the source of the first SMI event on the processor.

9. The processor of Claim 1, wherein:

the arbiter is a component of a controller hub; and

the interrupt handler further comprises circuitry to:

set an end-of-SMI indicator on the controller hub.

10. The processor of Claim 1, wherein:

the processor is one of a plurality of processors communicatively coupled to the arbiter; and

the indication of the first SMI event includes an identifier of the processor.

11. A method, comprising, in a processor:

detecting, on the processor, a first system management interrupt (SMI) event of a first SMI event type;

directing an indication of the first SMI event to a controller hub that is outside the processor;

receiving an interrupt signal from the controller hub;

executing, in response to receiving the interrupt signal, an interrupt handler, including:

taking action in response to detecting the first SMI event, the action being dependent on the first SMI event type; and

setting an indicator on the controller hub to indicate that the interrupt handler has completed its actions in response to the interrupt signal.

12. The method of Claim 11, wherein:

the processor includes a core for executing instructions; and

the method further comprises:

communicating the interrupt signal to the core; and

pausing, by the core, execution of instructions while the interrupt handler responds to the interrupt signal.

13. The method of Claim 11, wherein:

the method further comprises:

detecting, on the processor prior to receiving the interrupt signal, a second SMI event of a second SMI event type different from the first SMI event type; and

directing an indication of the second SMI event to the controller hub;

executing the interrupt handler further includes:

taking action in response to detecting the second SMI event, the action being dependent on the second SMI event type.

14. The method of Claim 11, wherein:

executing the interrupt handler further comprises:

determining the first SMI event type.

15. The method of Claim 11, wherein:

executing the interrupt handler further comprises:

determining that a second SMI event is pending; and

taking action in response to the second SMI event.

16. The method of Claim 11, wherein:

the controller hub arbitrates between SMI events that originate on both the processor and on the controller hub; and

executing the interrupt handler further comprises:

determining that a second SMI event was detected on the controller hub; and

taking action in response to the second SMI event.

17. The method of Claim 11, wherein:

the processor comprises a plurality of SMI event sources; and

executing the interrupt handler further comprises:

determining the source of the first SMI event on the processor.

18. A system, comprising:

a controller hub; and

a first processor, comprising:

a core including circuitry to execute instructions; and

circuitry to:

detect a first system management interrupt (SMI) event from a first SMI event source on the first processor;

direct an indication of the first SMI event to the controller hub;

receive an interrupt signal from the controller hub;

an interrupt handler to respond to the interrupt signal, including circuitry to:

take action in response to detection of the first SMI event, the action to be dependent on the source of the first SMI event; and

set an indicator on the controller hub to indicate that the interrupt handler has completed its actions in response to the interrupt signal.

19. The system of Claim 18, wherein:

the first processor further comprises circuitry to:

communicate the interrupt signal to the core;

the core includes circuitry to:

pause execution of instructions while the interrupt handler responds to the interrupt signal.

20. The system of Claim 18, wherein:

the first processor further comprises circuitry to:

detect, prior to receipt of the interrupt signal, a second SMI event from a second SMI event source on the first processor different from the first SMI event source; and

direct an indication of the second SMI event to the controller hub;

the interrupt handler further includes circuitry to:

take action in response to detection of the second SMI event, the action to be dependent on the source of the second SMI event.

21. The system of Claim 18, wherein:

the interrupt handler further comprises circuitry to:

determine that a second SMI event is pending; and

take action in response to the second SMI event.

22. The system of Claim 18, wherein:

the controller hub comprises:

a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor;

a second SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the controller hub;

the interrupt handler further comprises circuitry to:

determine, dependent on the second SMI status bit, that a second SMI event was detected on the controller hub; and

take action in response to the second SMI event.

23. The system of Claim 18, wherein:

the controller hub comprises:

an end-of-SMI indicator whose value indicates whether or not the interrupt handler has completed its actions in response to an interrupt issued to the first processor by the controller hub; and

circuitry to:

prevent the controller hub from issuing an interrupt to the first processor while the end-of-SMI indicator is false;

determine that at least one SMI event is pending; and

issue an interrupt to the first processor in response to determining that at least one SMI event is pending and that the end-of-SMI indicator is true, including circuitry to:

clear the end-of-SMI indicator; and

assert the interrupt signal.

24. The system of Claim 18, wherein:

the system further comprises a second processor;

the indication of the first SMI event includes an identifier of the first processor;

the controller hub comprises:

a first SMI status bit whose value indicates whether or not an SMI event has been detected for an SMI event source on the first processor; and

circuitry to determine, dependent on receipt of the indication of the first SMI event and the identifier of the first processor, that the first SMI status bit is to be set.

25. An apparatus, including means for performing any of the methods of Claims 11-17.

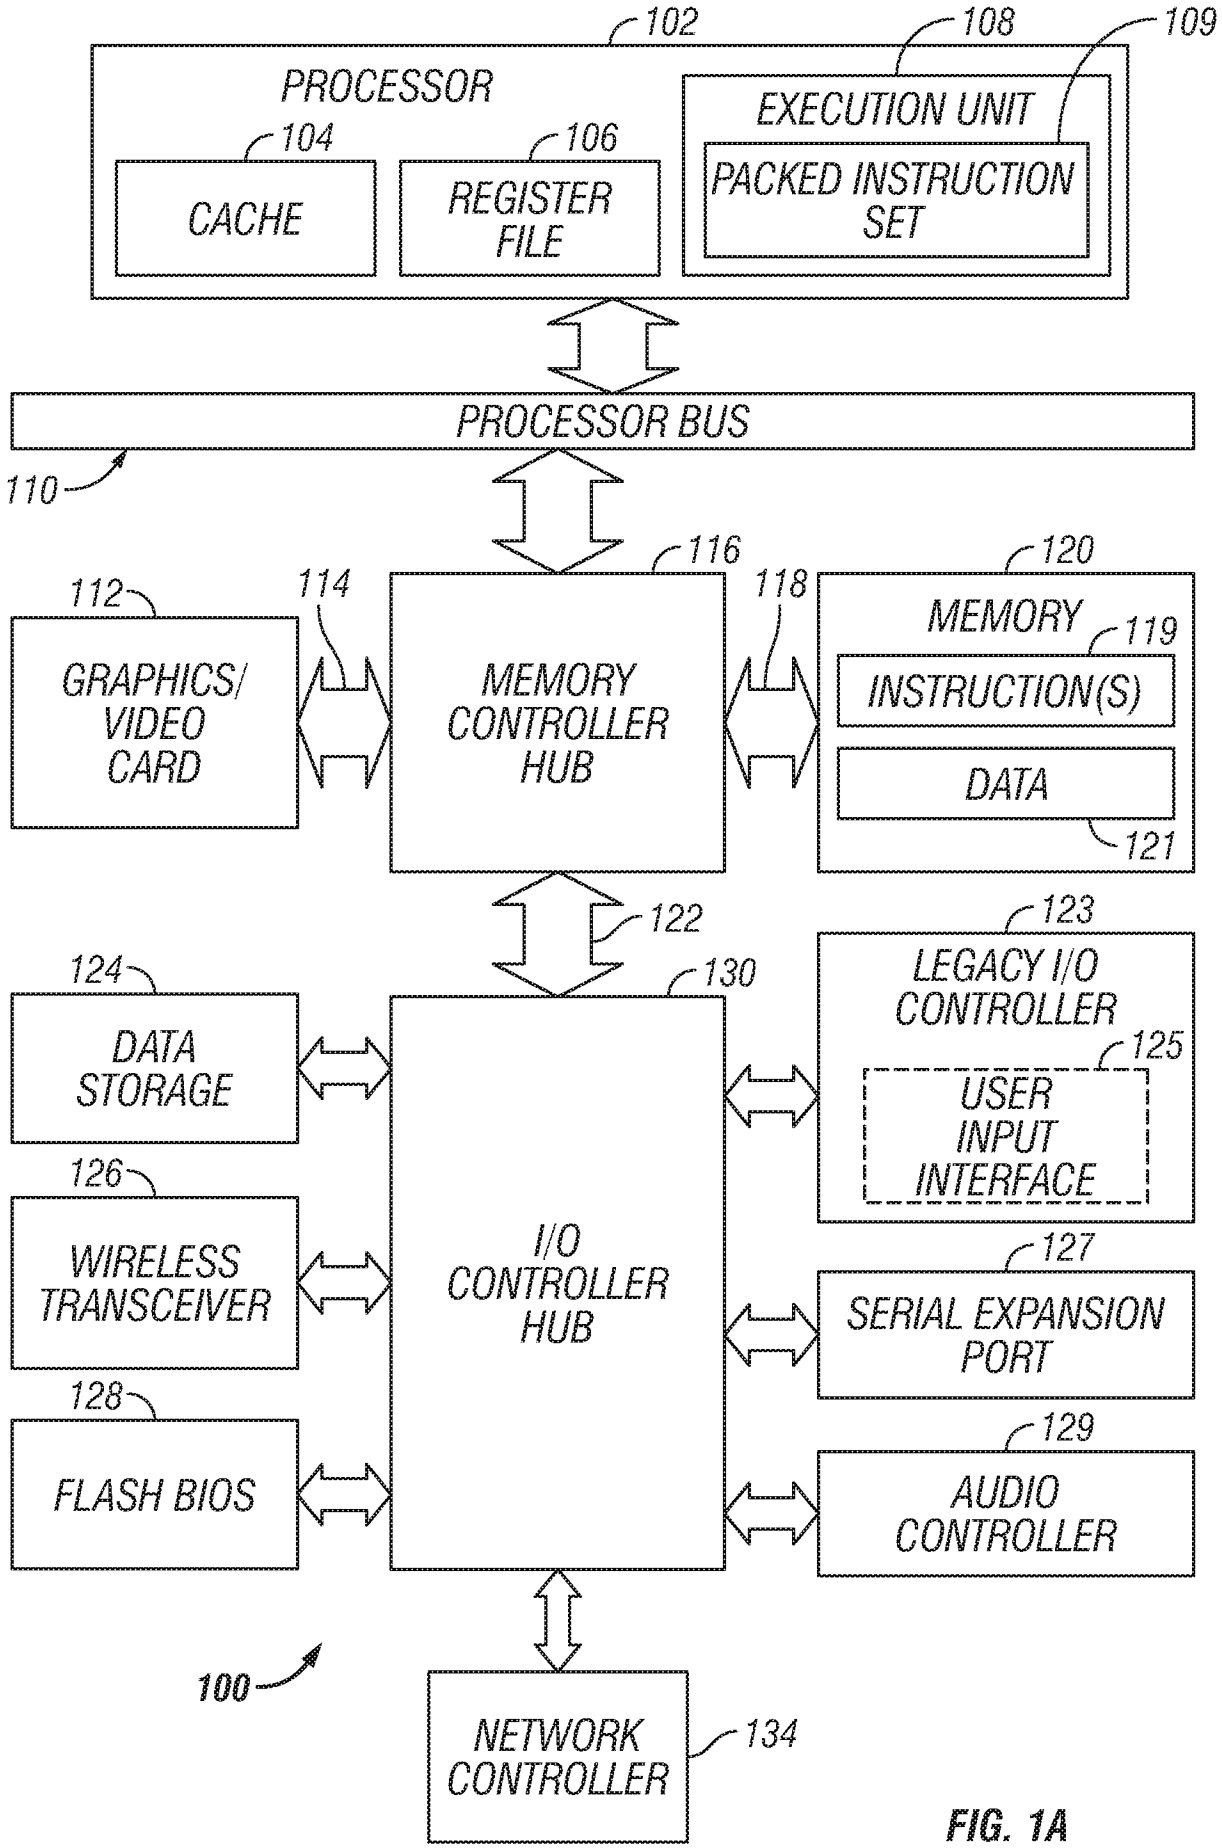

102
PROCESSOR
104
CACHE
106
REGISTER FILE
108
EXECUTION UNIT
109
PACKED INSTRUCTION SET
PROCESSOR BUS
110
112
GRAPHICS/ VIDEO CARD
114
116
MEMORY CONTROLLER HUB
118
120
MEMORY
119
INSTRUCTION(S)
DATA
121
122
123
LEGACY I/O CONTROLLER
125
USER INPUT INTERFACE
124
DATA STORAGE
126
WIRELESS TRANSCEIVER
128
FLASH BIOS
130
I/O CONTROLLER HUB
127
SERIAL EXPANSION PORT
129
AUDIO CONTROLLER
100
NETWORK CONTROLLER
134

FIG. 1A

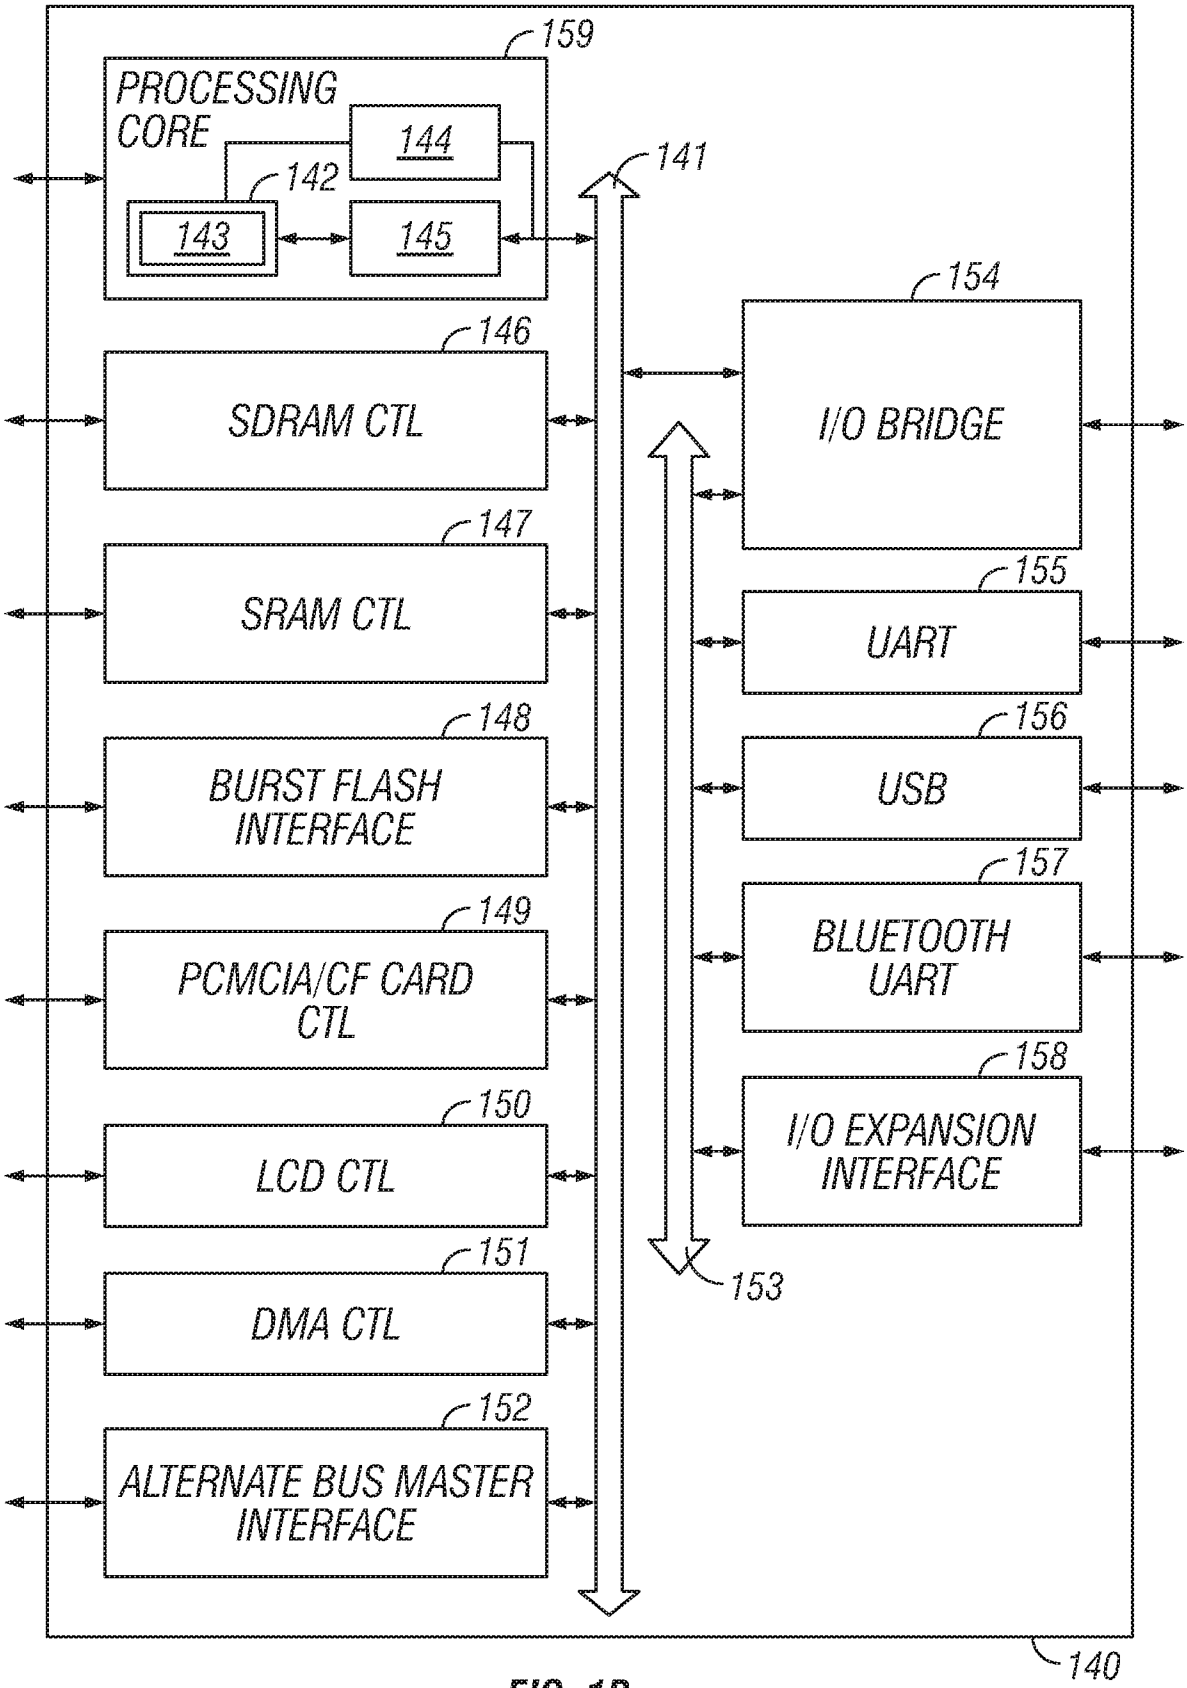

159
PROCESSING CORE
144
142
143
145
141
154
I/O BRIDGE
146
SDRAM CTL
147
SRAM CTL
155
UART
148
BURST FLASH INTERFACE
156
USB
149
PCMCIA/CF CARD CTL
157
BLUETOOTH UART
150
LCD CTL
158
I/O EXPANSION INTERFACE
151
DMA CTL
153
152
ALTERNATE BUS MASTER INTERFACE
140

FIG. 1B

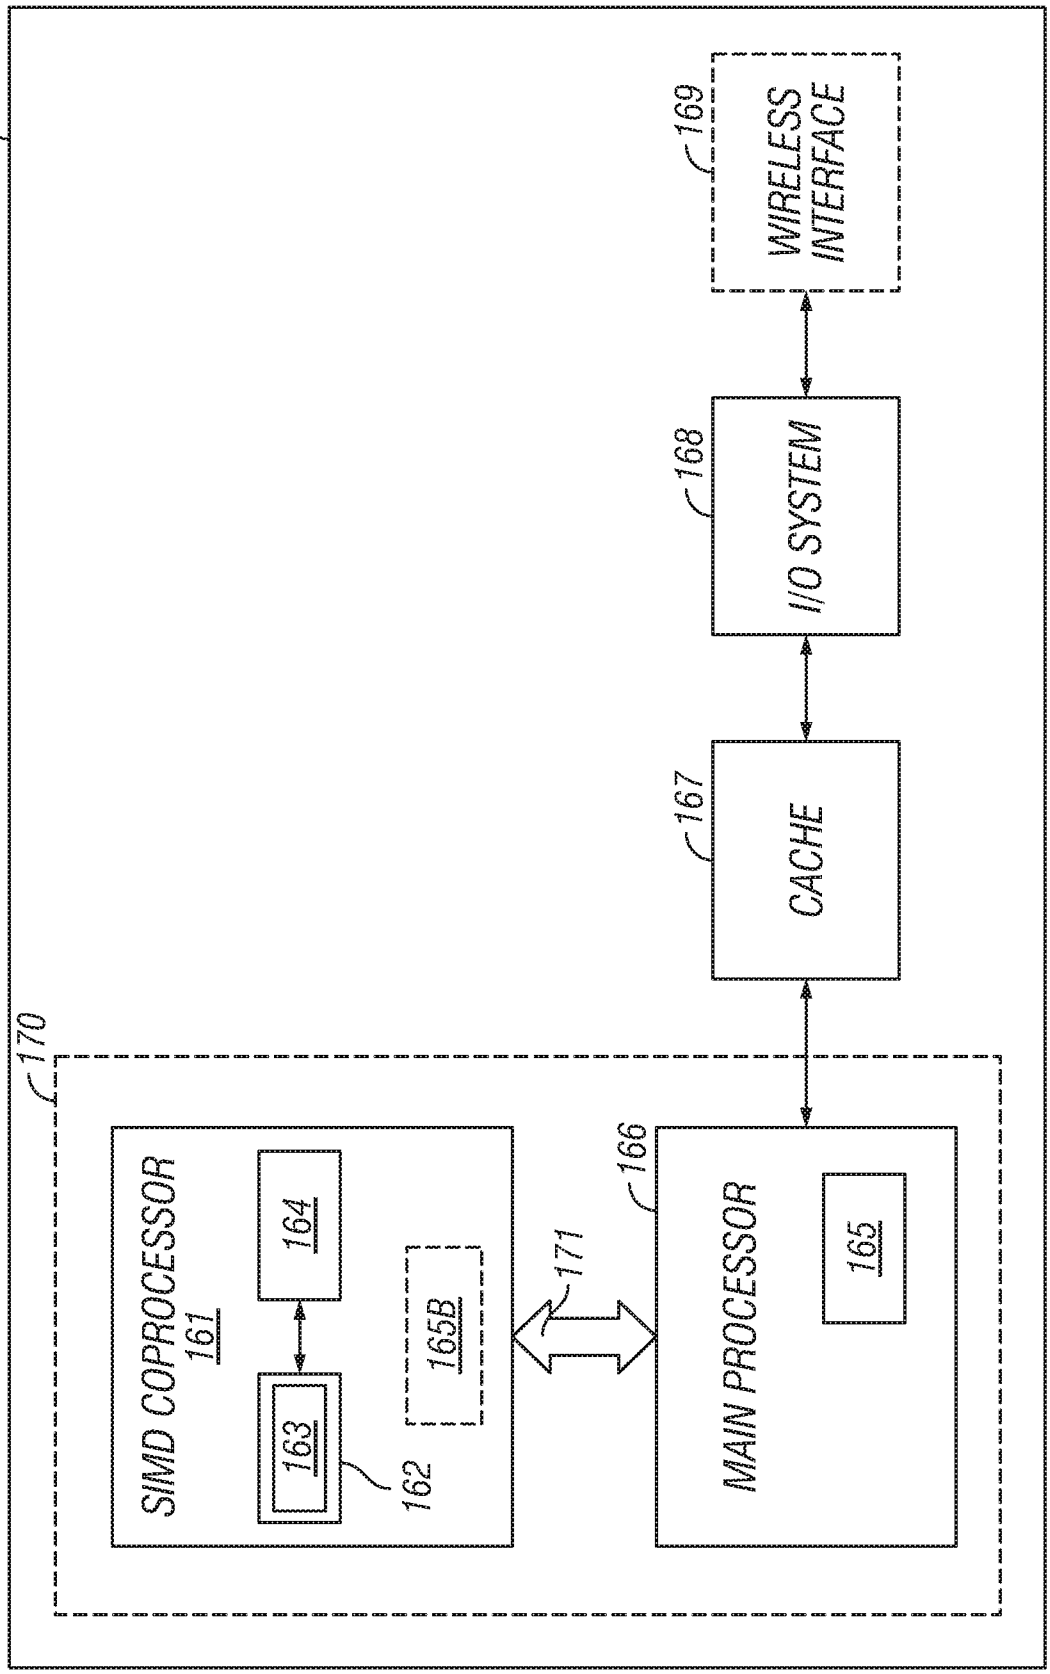

160
169
WIRELESS INTERFACE
168
I/O SYSTEM
167
CACHE
170
SIMD COPROCESSOR
161
164
163
162
165B
171
166
MAIN PROCESSOR
165

FIG. 1C

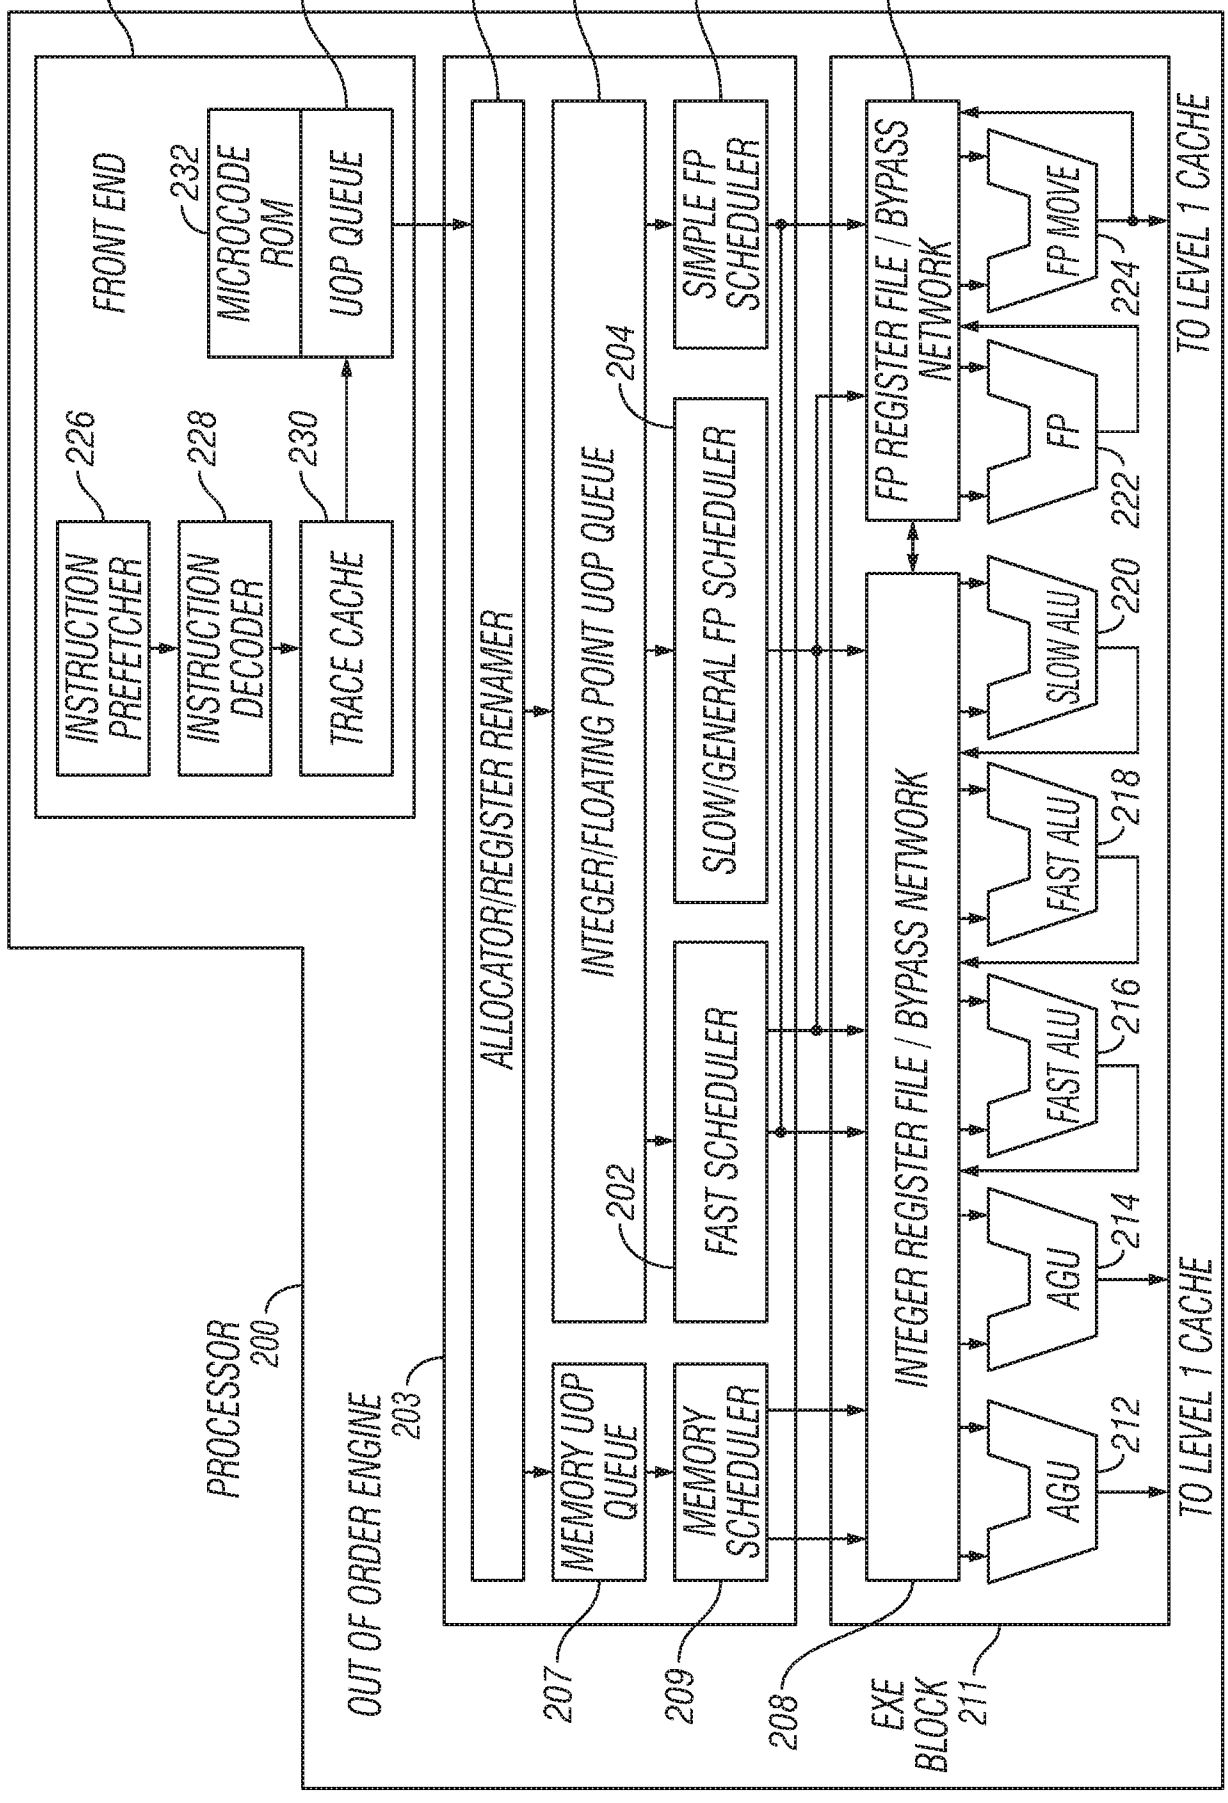

PROCESSOR
200
OUT OF ORDER ENGINE
203
FRONT END
201
INSTRUCTION PREFETCHER
226
INSTRUCTION DECODER
228
TRACE CACHE
230
MICROCODE ROM
232
UOP QUEUE
234
ALLOCATOR/REGISTER RENAMER
215
INTEGER/FLOATING POINT UOP QUEUE
205
202
204
SIMPLE FP SCHEDULER
206
FAST SCHEDULER
SLOW/GENERAL FP SCHEDULER
MEMORY UOP QUEUE
MEMORY SCHEDULER
207
209
208
EXE BLOCK
211
INTEGER REGISTER FILE / BYPASS NETWORK
FP REGISTER FILE / BYPASS NETWORK
210
AGU
212
AGU
214
FAST ALU
216
FAST ALU
218
SLOW ALU
220
FP
222
FP MOVE
224
TO LEVEL 1 CACHE
TO LEVEL 1 CACHE

FIG. 2

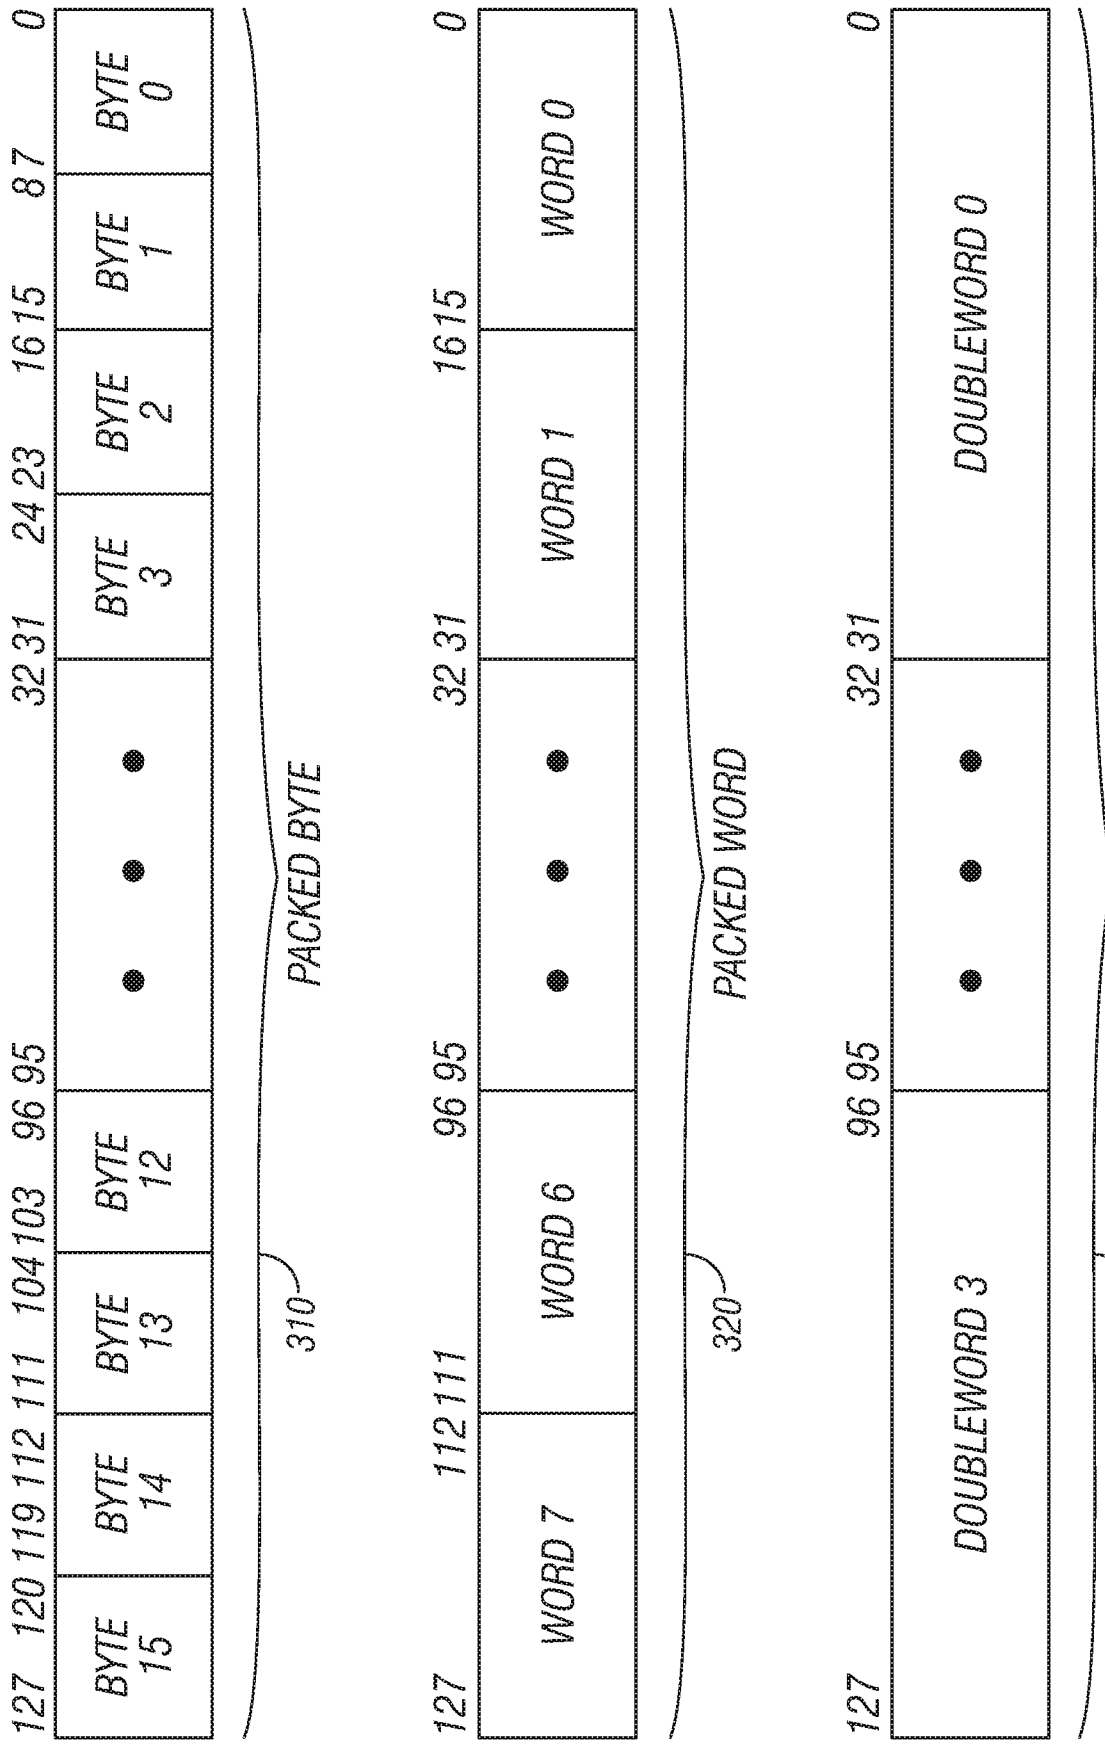

127 120 119 112 111 104 103 96 95 32 31 24 23 16 15 8 7 0
BYTE 15
BYTE 14
BYTE 13
BYTE 12
BYTE 3
BYTE 2
BYTE 1
BYTE 0
310
PACKED BYTE
127 112 111 96 95 32 31 16 15 0
WORD 7
WORD 6
WORD 1
WORD 0
320
PACKED WORD
127 96 95 32 31 0
DOUBLEWORD 3
DOUBLEWORD 0
330
PACKED DOUBLEWORD

FIG. 3A

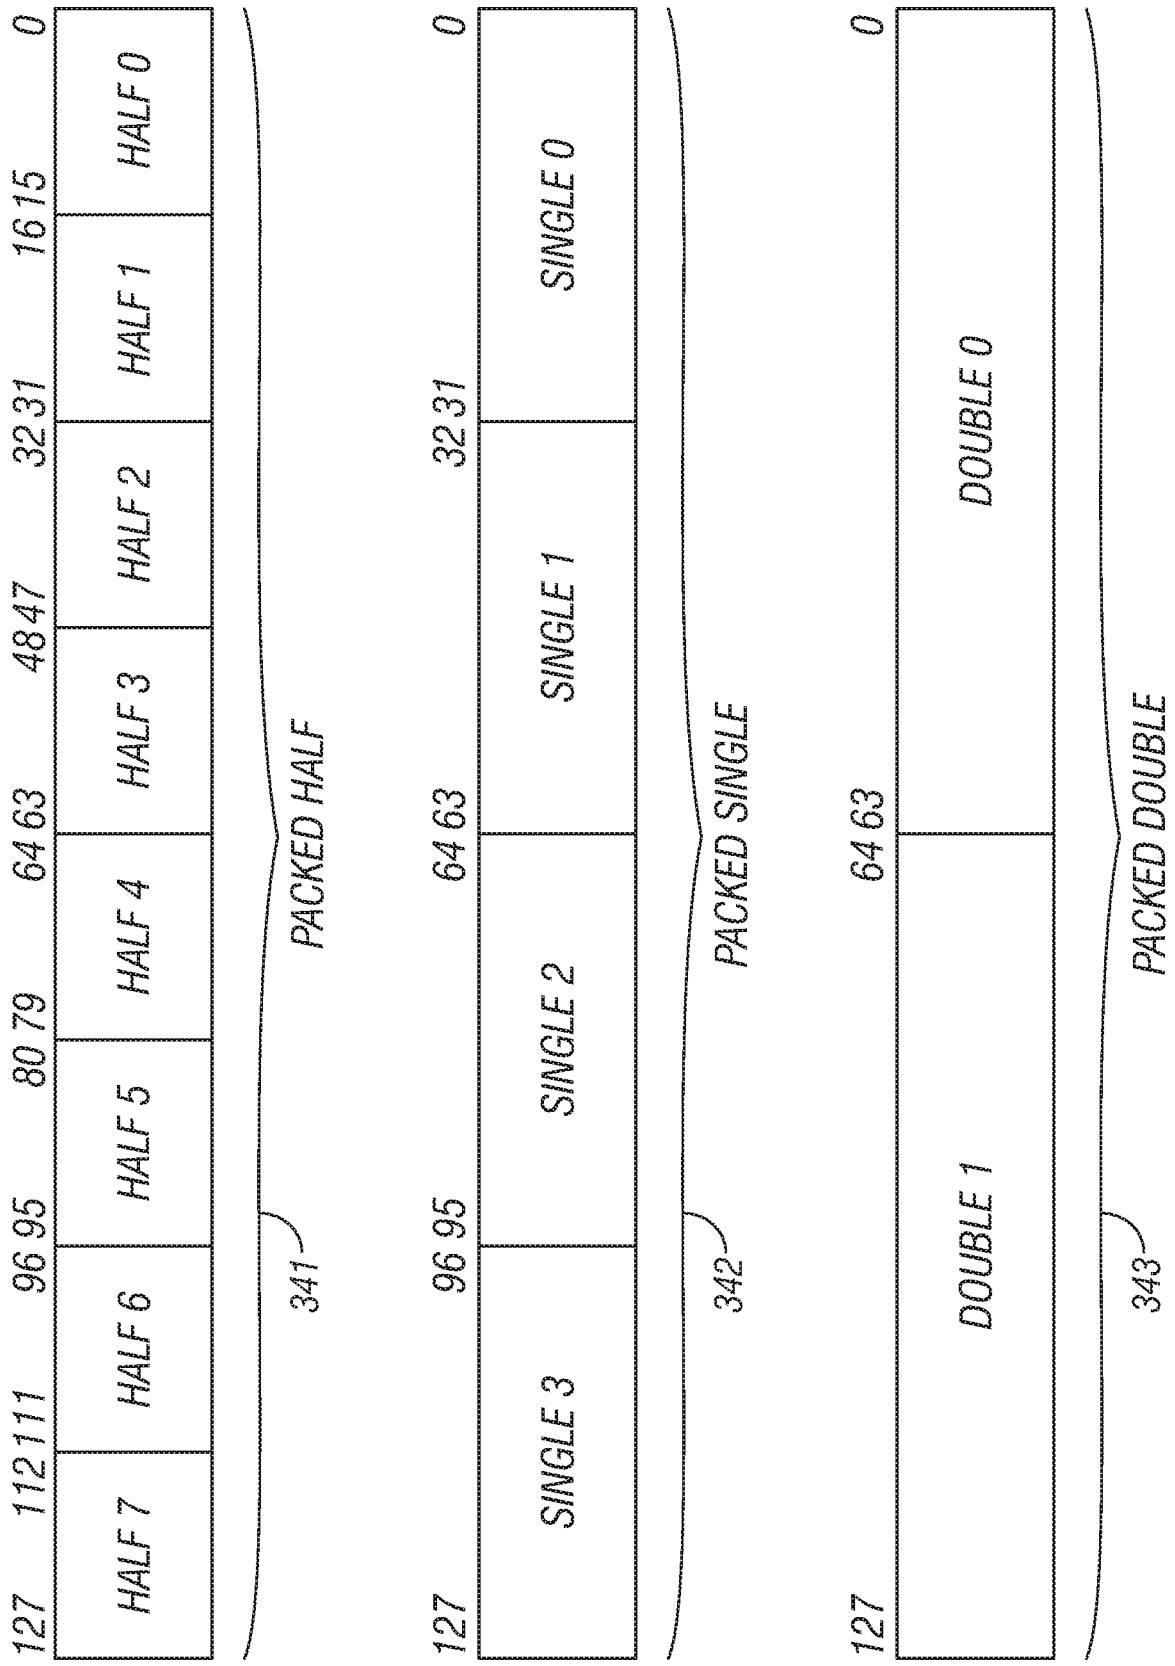
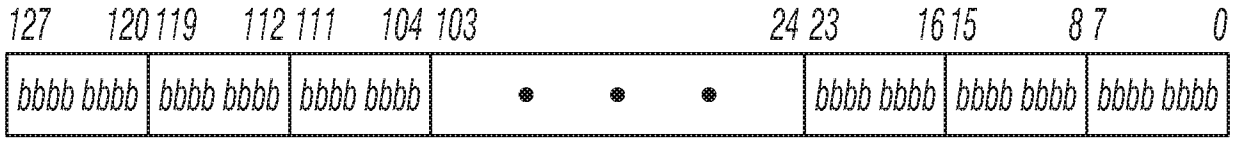
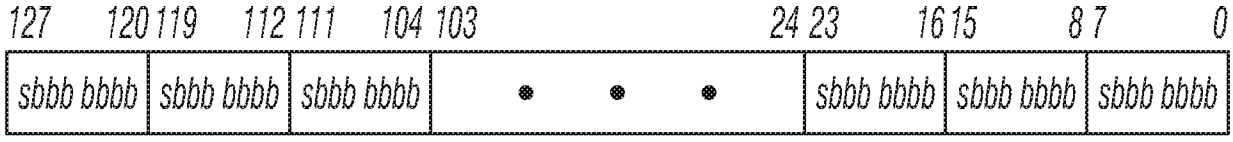

127 112 111 96 95 80 79 64 63 48 47 32 31 16 15 0
HALF 7
HALF 6
HALF 5
HALF 4
HALF 3
HALF 2
HALF 1
HALF 0
341
PACKED HALF
127 96 95 64 63 32 31 0
SINGLE 3
SINGLE 2
SINGLE 1
SINGLE 0
342
PACKED SINGLE
127 64 63 0
DOUBLE 1
DOUBLE 0
343
PACKED DOUBLE

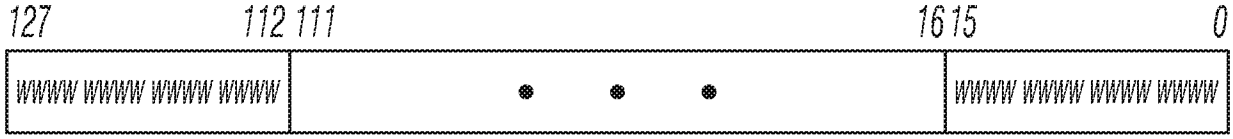
FIG. 3B



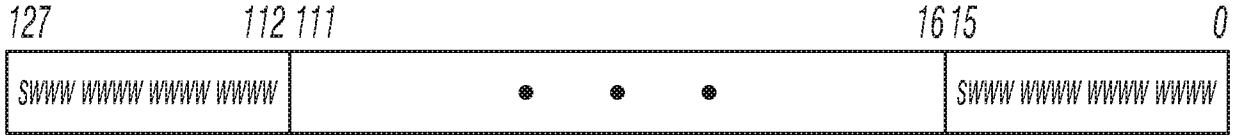
UNSIGNED PACKED BYTE REPRESENTATION 344



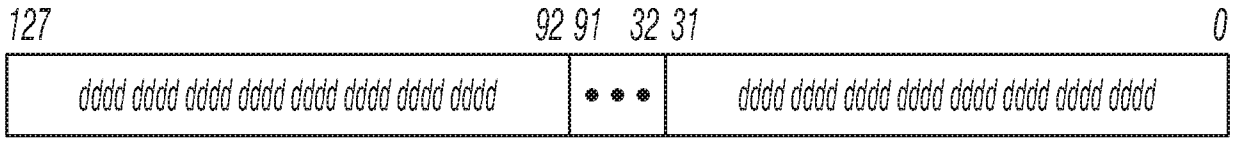
SIGNED PACKED BYTE REPRESENTATION 345



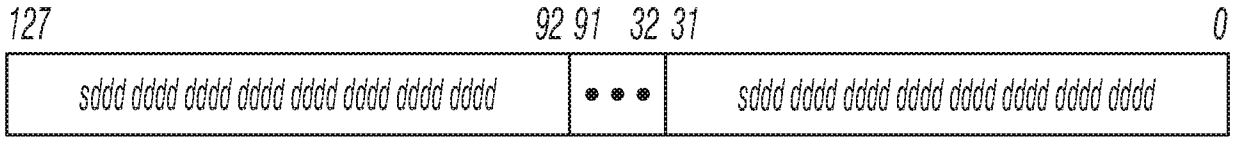
UNSIGNED PACKED WORD REPRESENTATION 346



SIGNED PACKED WORD REPRESENTATION 347



UNSIGNED PACKED DOUBLEWORD REPRESENTATION 348



SIGNED PACKED DOUBLEWORD REPRESENTATION 349

FIG. 3C

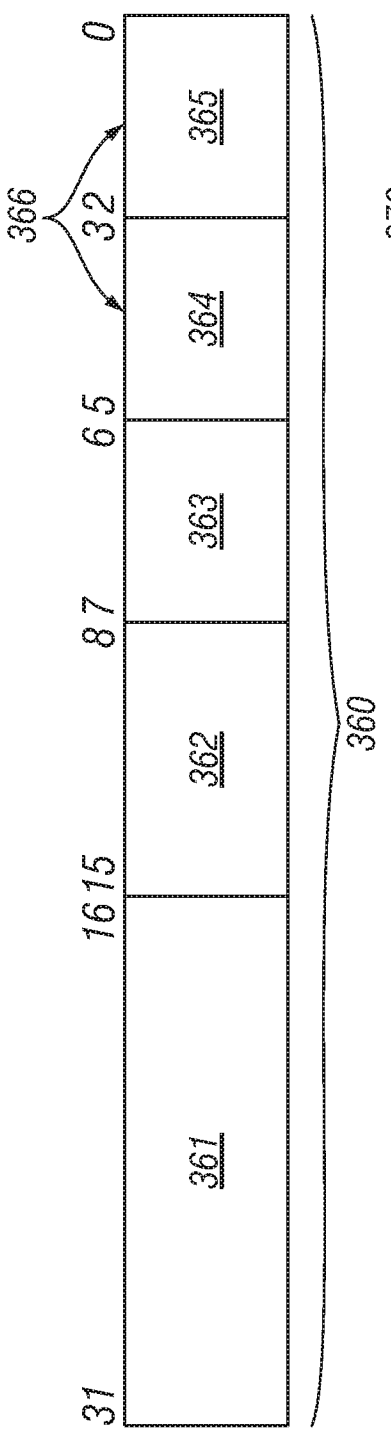


FIG. 3D

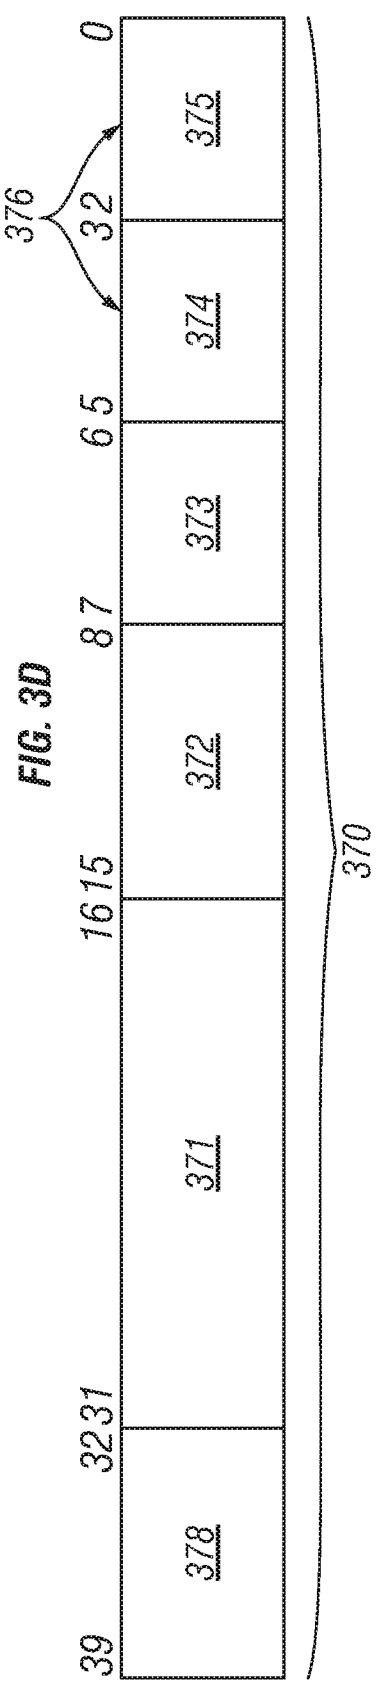


FIG. 3E

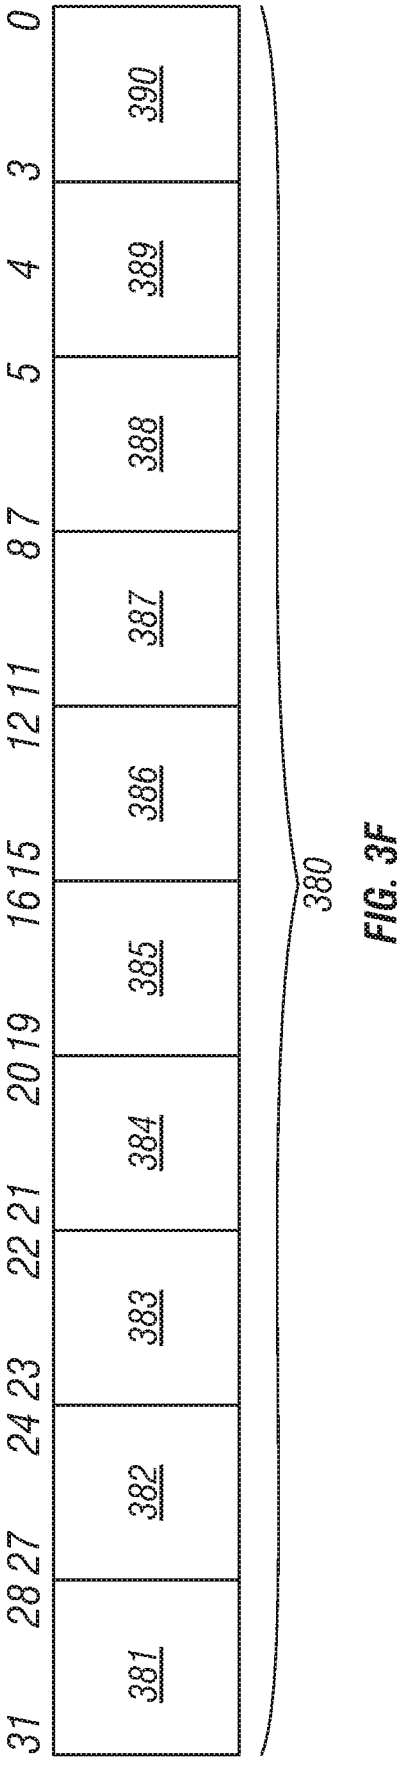


FIG. 3F

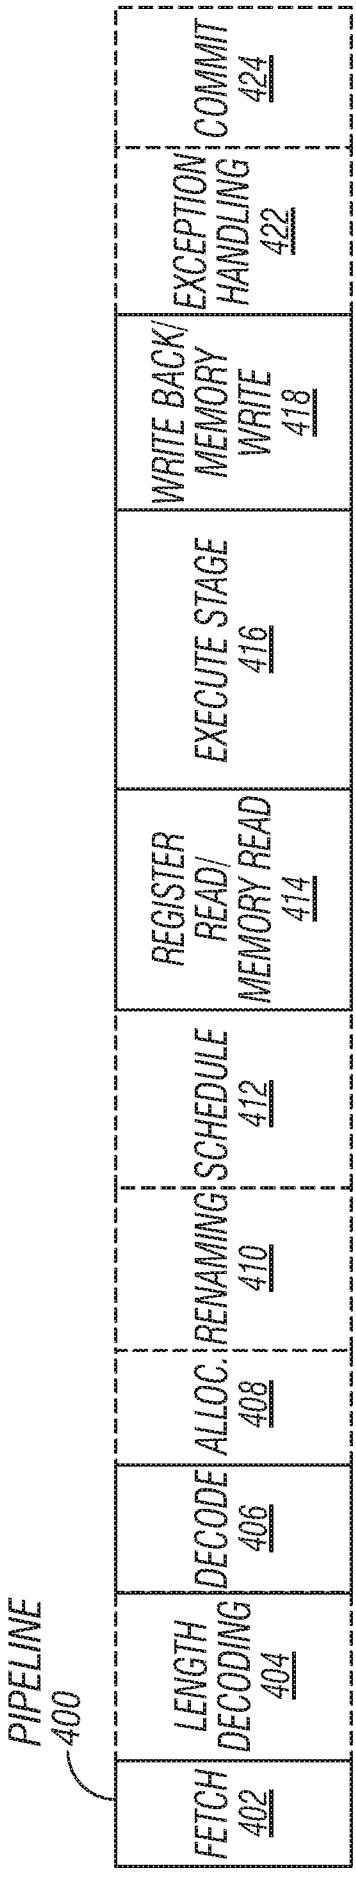

PIPELINE
400
FETCH
402
LENGTH DECODING
404
DECODE
406
ALLOC.
408
RENAMING
410
SCHEDULE
412
REGISTER READ/ MEMORY READ
414
EXECUTE STAGE
416
WRITE BACK/ MEMORY WRITE
418
EXCEPTION HANDLING
422
COMMIT
424

FIG. 4A

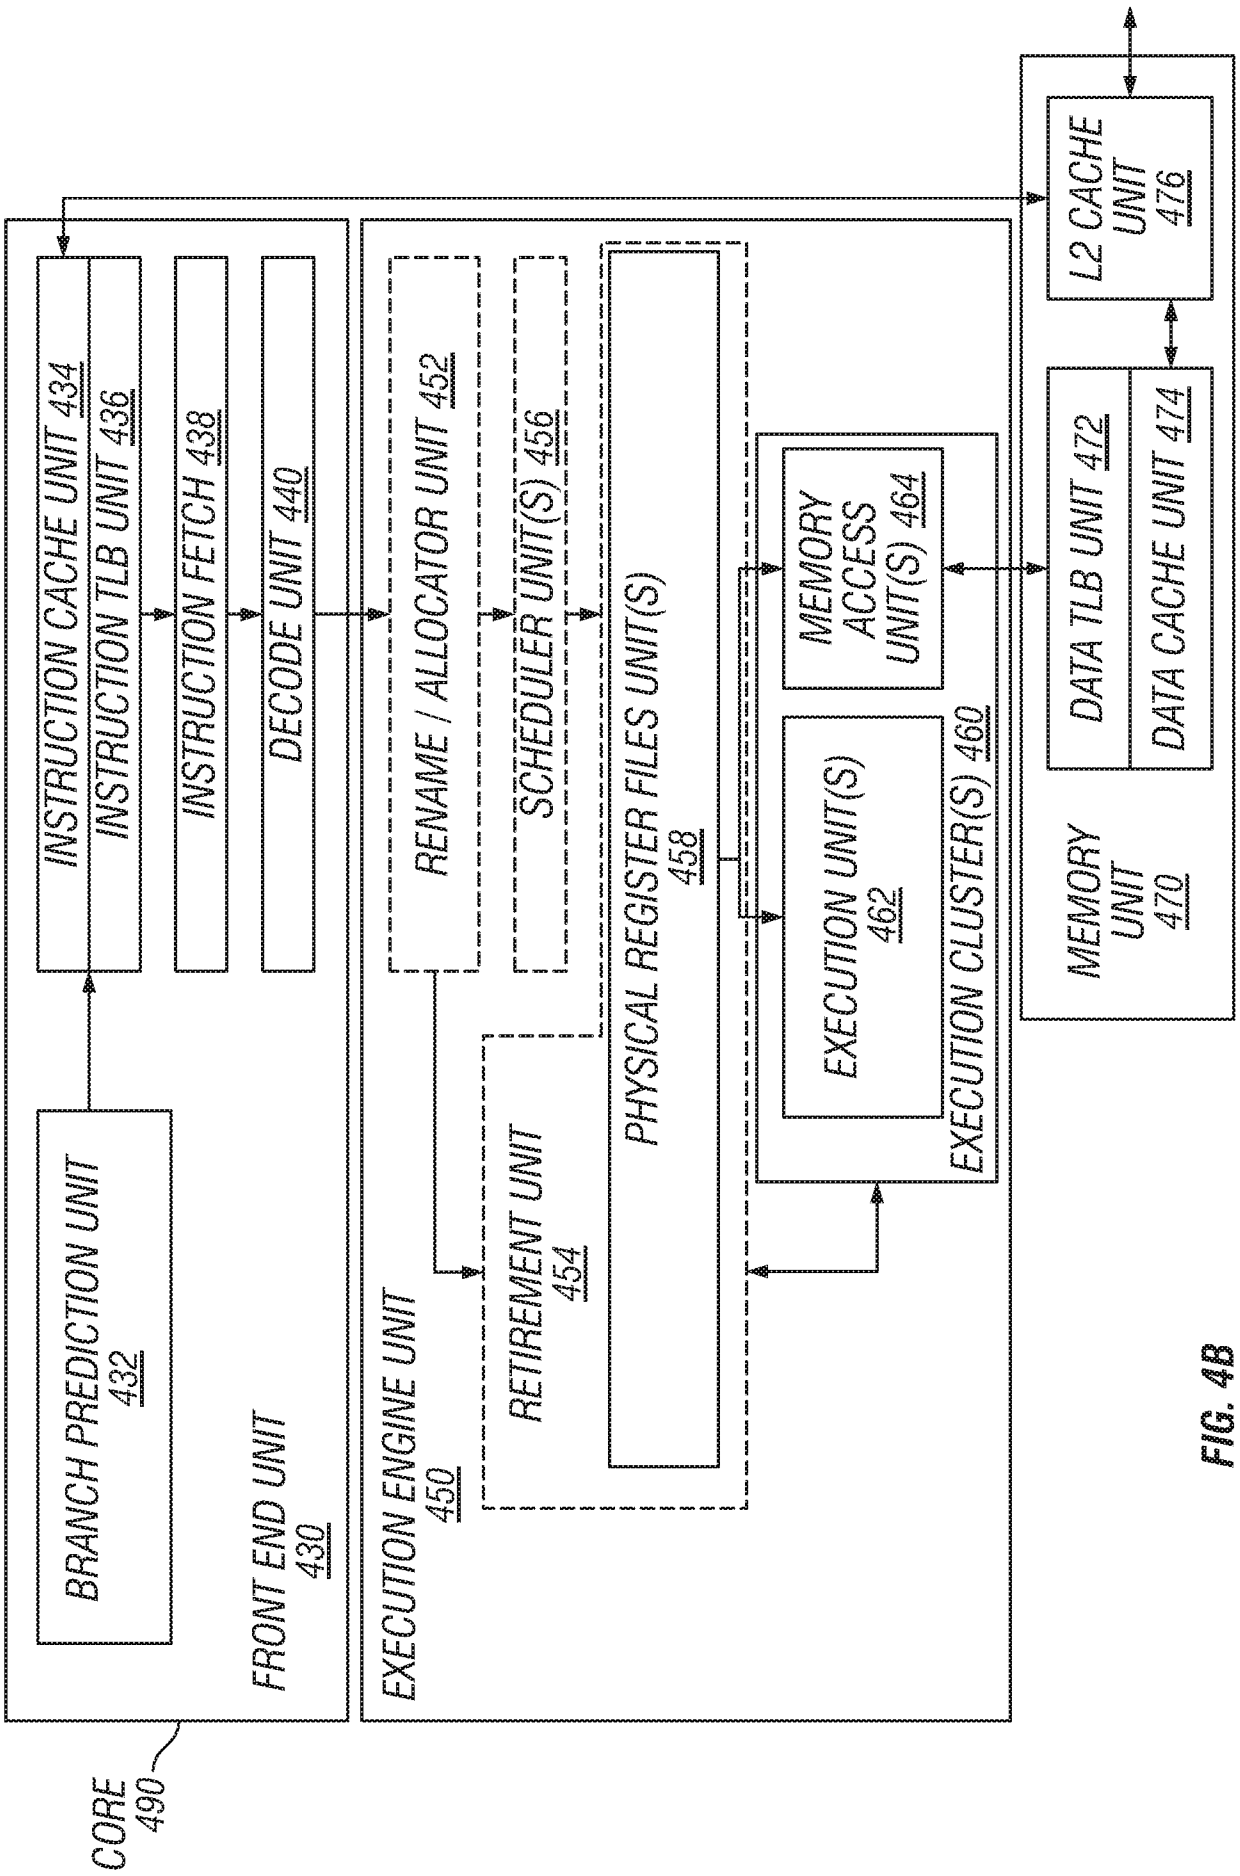

CORE 490
FRONT END UNIT 430
BRANCH PREDICTION UNIT 432
INSTRUCTION CACHE UNIT 434
INSTRUCTION TLB UNIT 436
INSTRUCTION FETCH 438
DECODE UNIT 440
EXECUTION ENGINE UNIT 450
RENAME / ALLOCATOR UNIT 452
RETIREMENT UNIT 454
SCHEDULER UNIT(S) 456
PHYSICAL REGISTER FILES UNIT(S) 458
EXECUTION CLUSTER(S) 460
EXECUTION UNIT(S) 462
MEMORY ACCESS UNIT(S) 464
MEMORY UNIT 470
DATA TLB UNIT 472
DATA CACHE UNIT 474
L2 CACHE UNIT 476

FIG. 4B

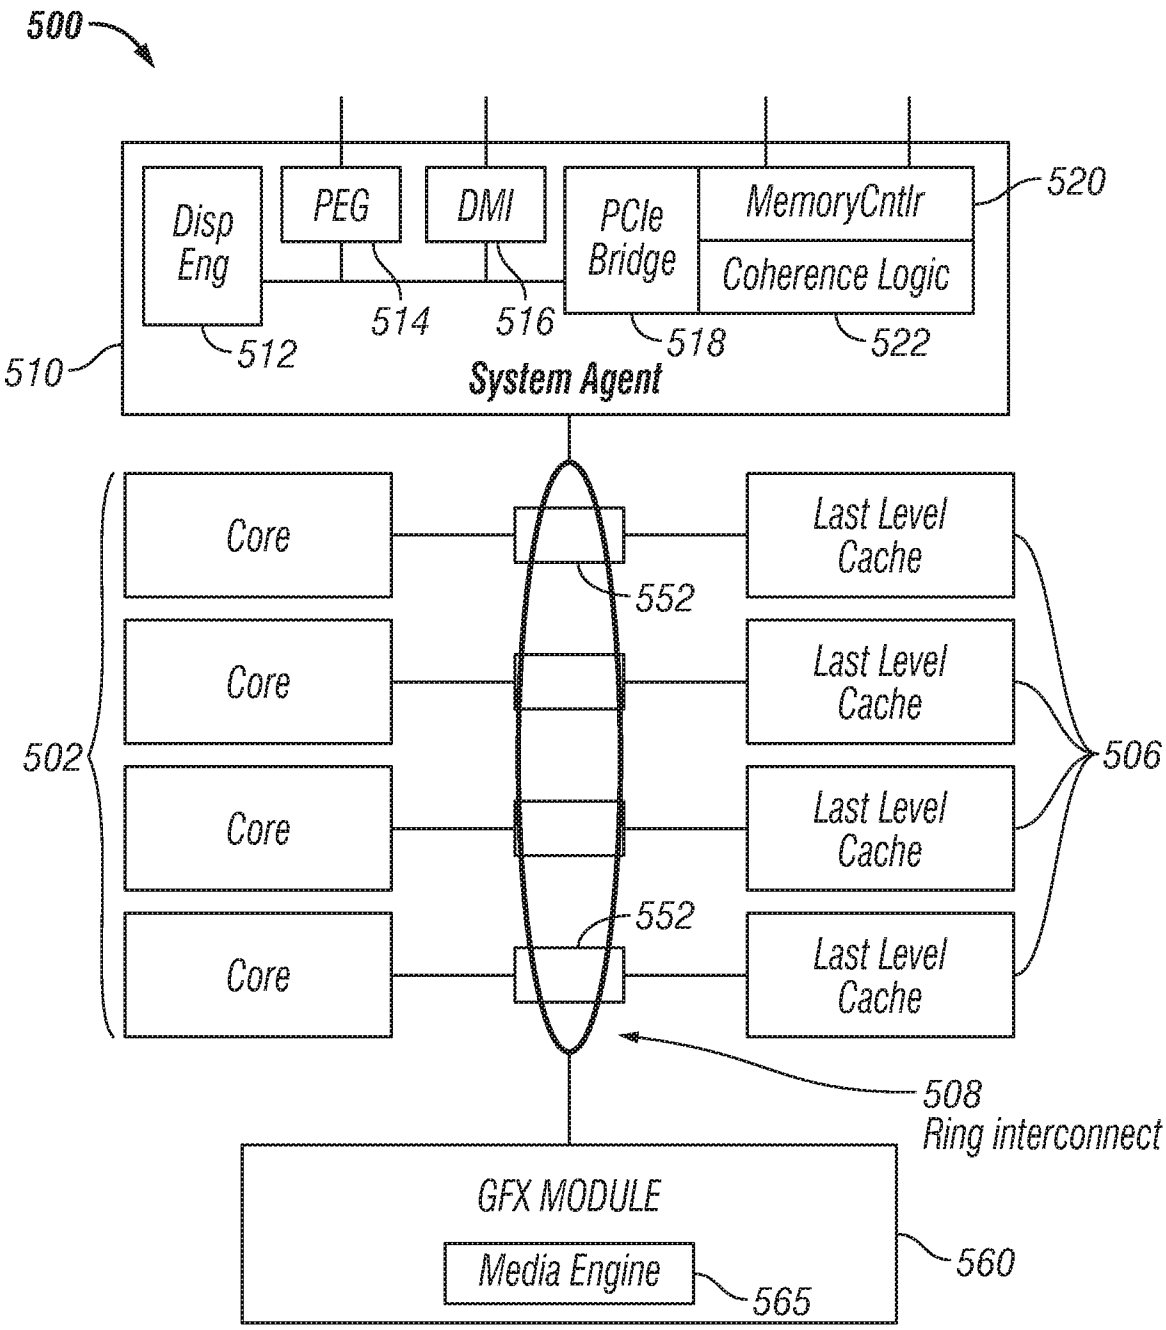

500
Disp Eng
PEG
DMI
PCIe Bridge
MemoryCntlr
Coherence Logic
520
514
516
518
522
512
510
System Agent
Core
Core
Core
Core
502
552
552
Last Level Cache
Last Level Cache
Last Level Cache
Last Level Cache
506
508
Ring interconnect
GFX MODULE
Media Engine
565
560

FIG. 5A

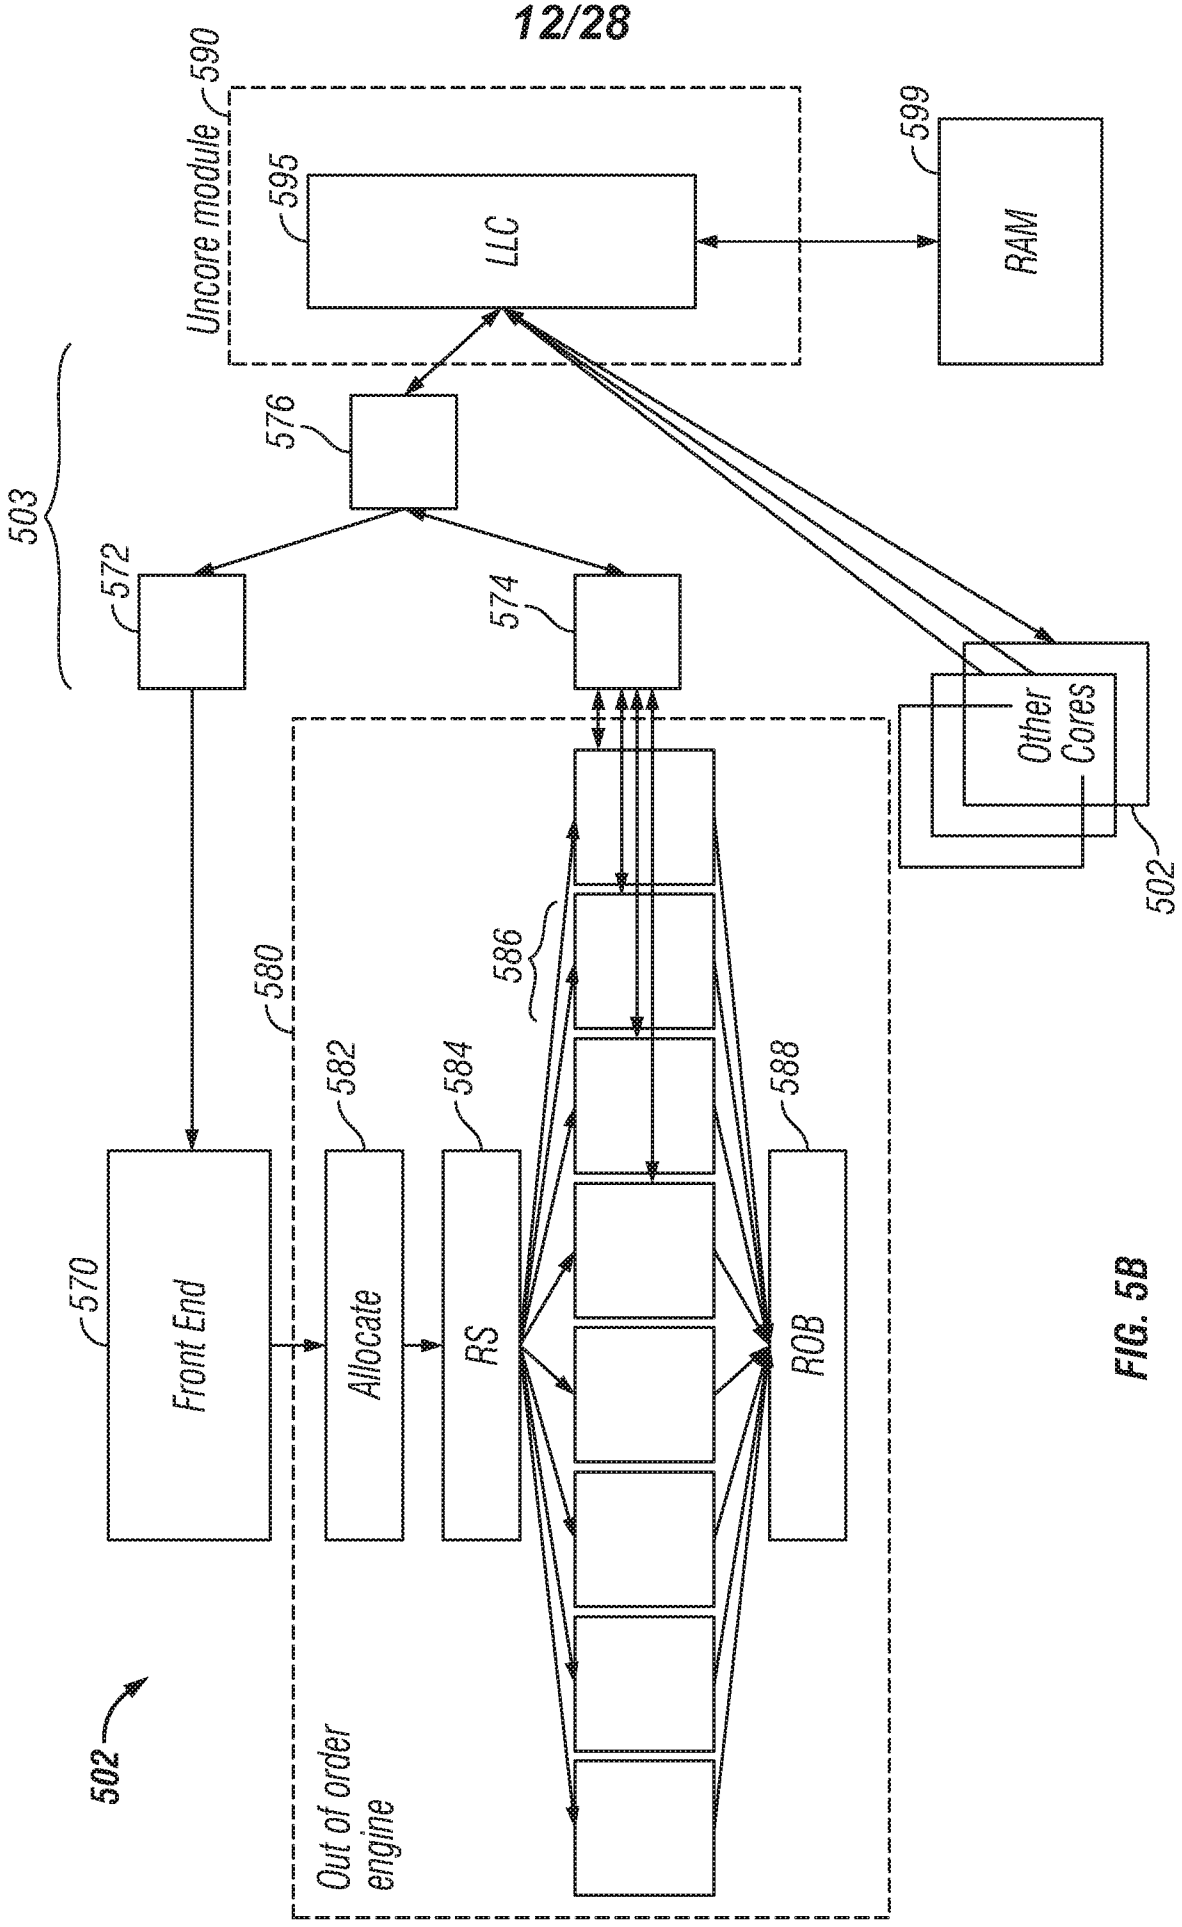

Uncore module
590
595
LLC
599
RAM
576
503
572
574
Other Cores
502
580
586
582
584
588
570
Front End
Allocate
RS
ROB
Out of order engine
502

FIG. 5B

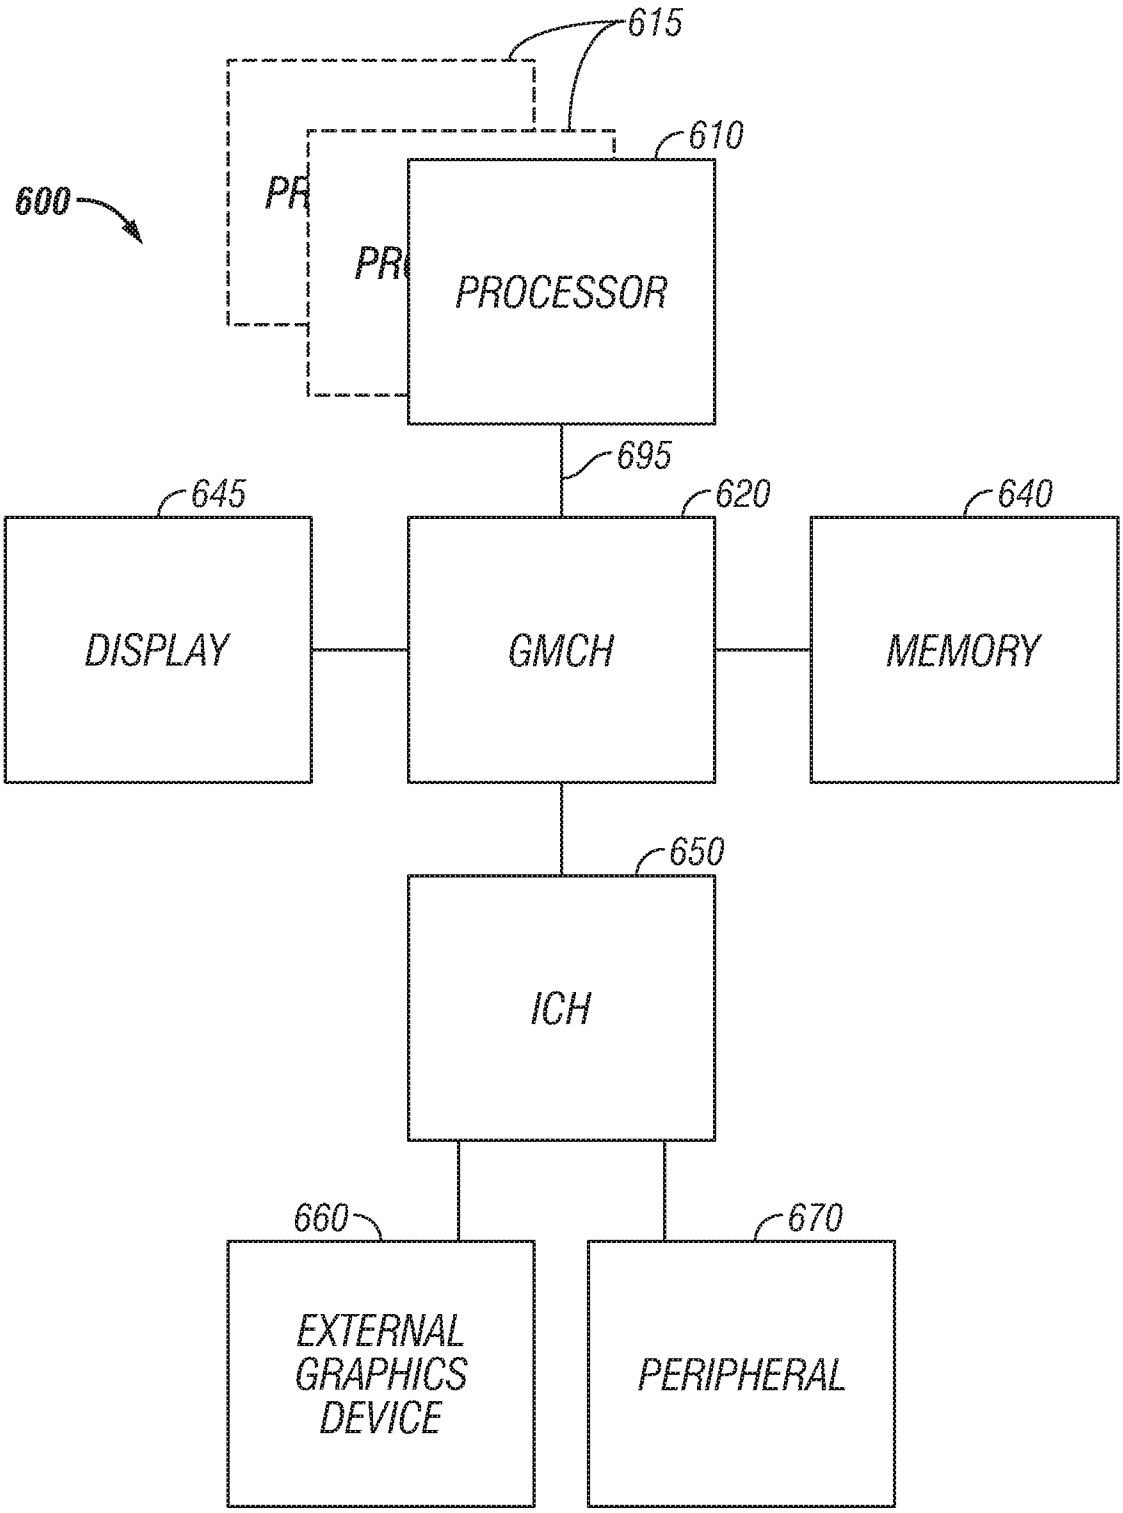

615
610
600
PR
PR
PROCESSOR
695
645
620
640
DISPLAY
GMCH
MEMORY
650
ICH
660
670
EXTERNAL GRAPHICS DEVICE
PERIPHERAL

FIG. 6

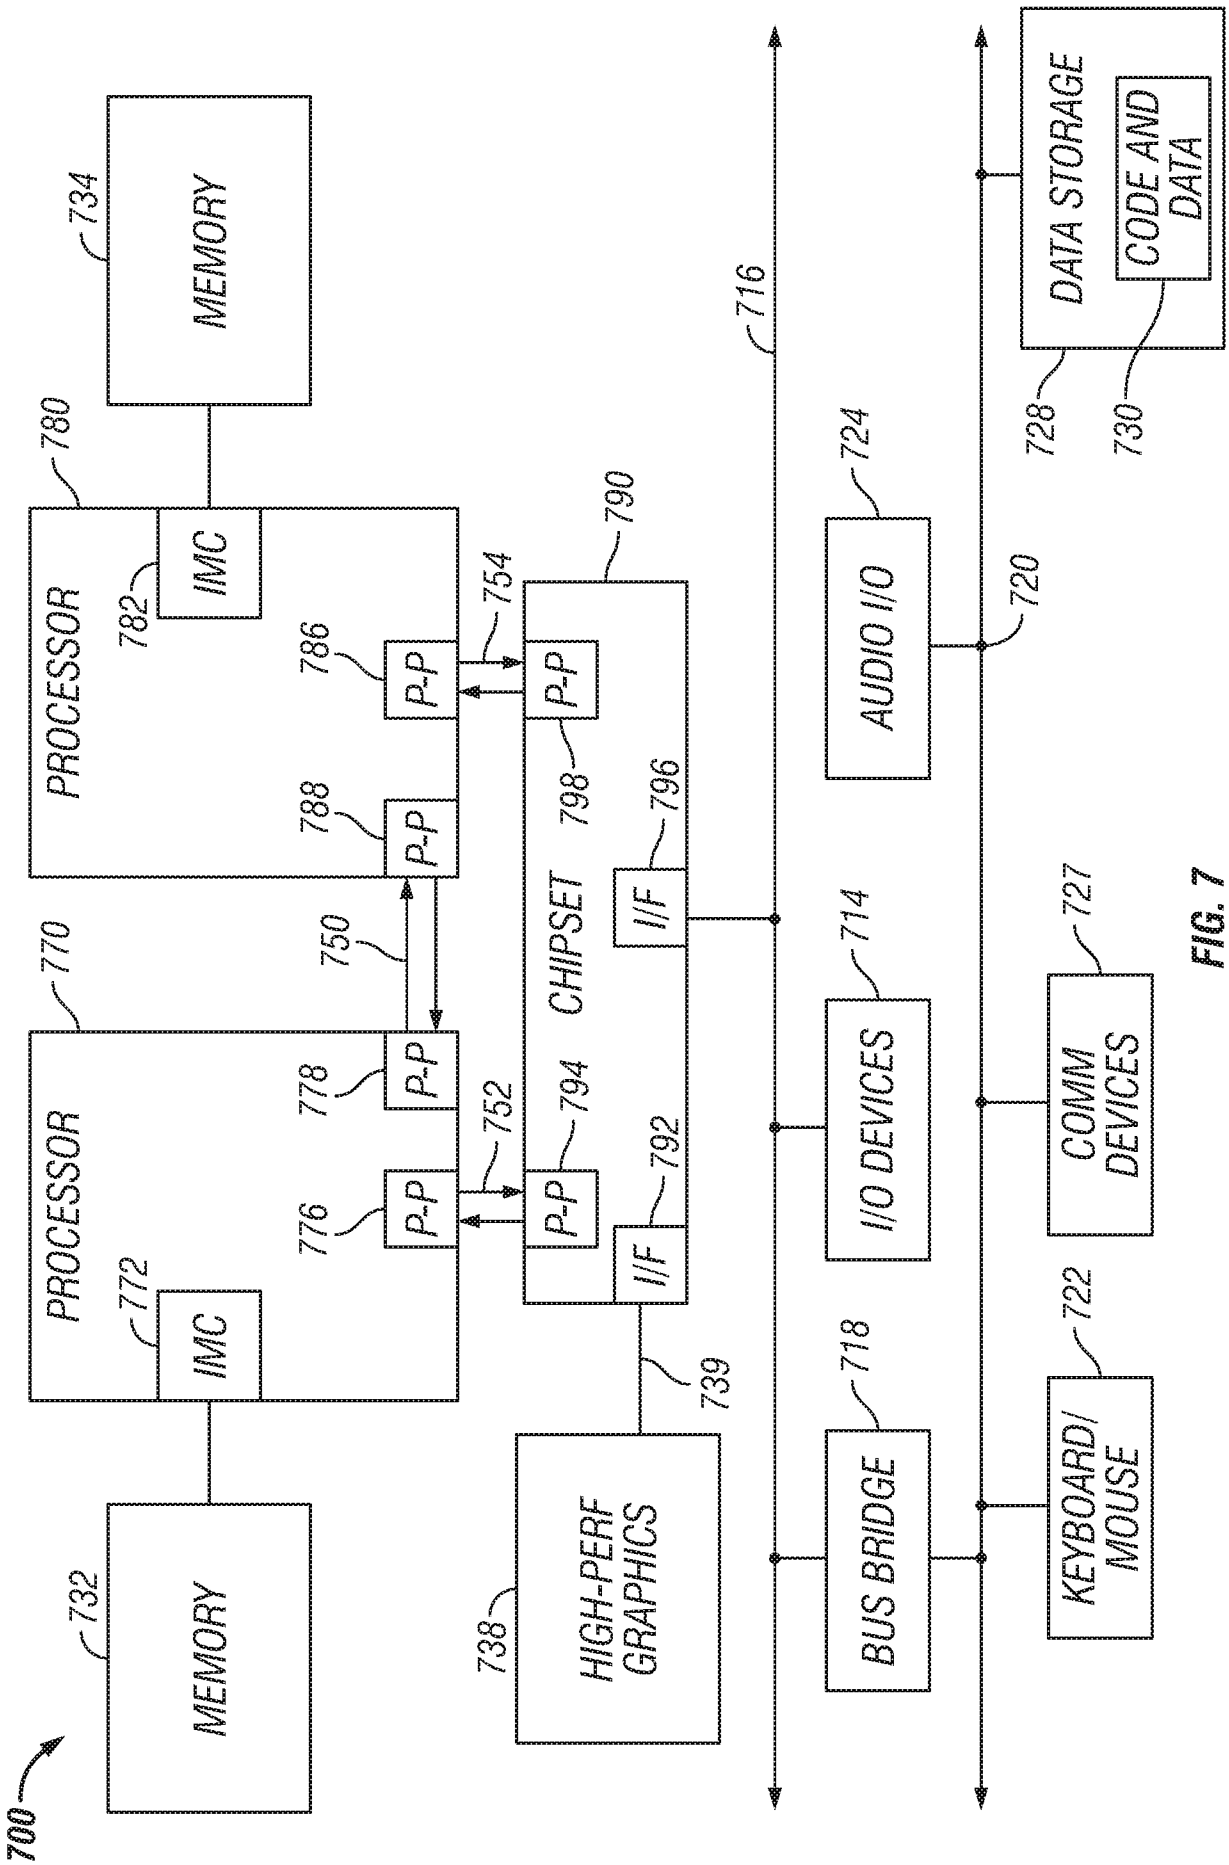

700
MEMORY
732
PROCESSOR
770
IMC
772
MEMORY
734
PROCESSOR
780
IMC
782
P-P
776
P-P
778
P-P
786
P-P
788
750
752
754
CHIPSET
790
P-P
794
P-P
798
I/F
792
I/F
796
739
HIGH-PERF GRAPHICS
738
BUS BRIDGE
718
I/O DEVICES
714
AUDIO I/O
724
716
720
DATA STORAGE
728
CODE AND DATA
730
KEYBOARD/ MOUSE
722
COMM DEVICES
727

FIG. 7

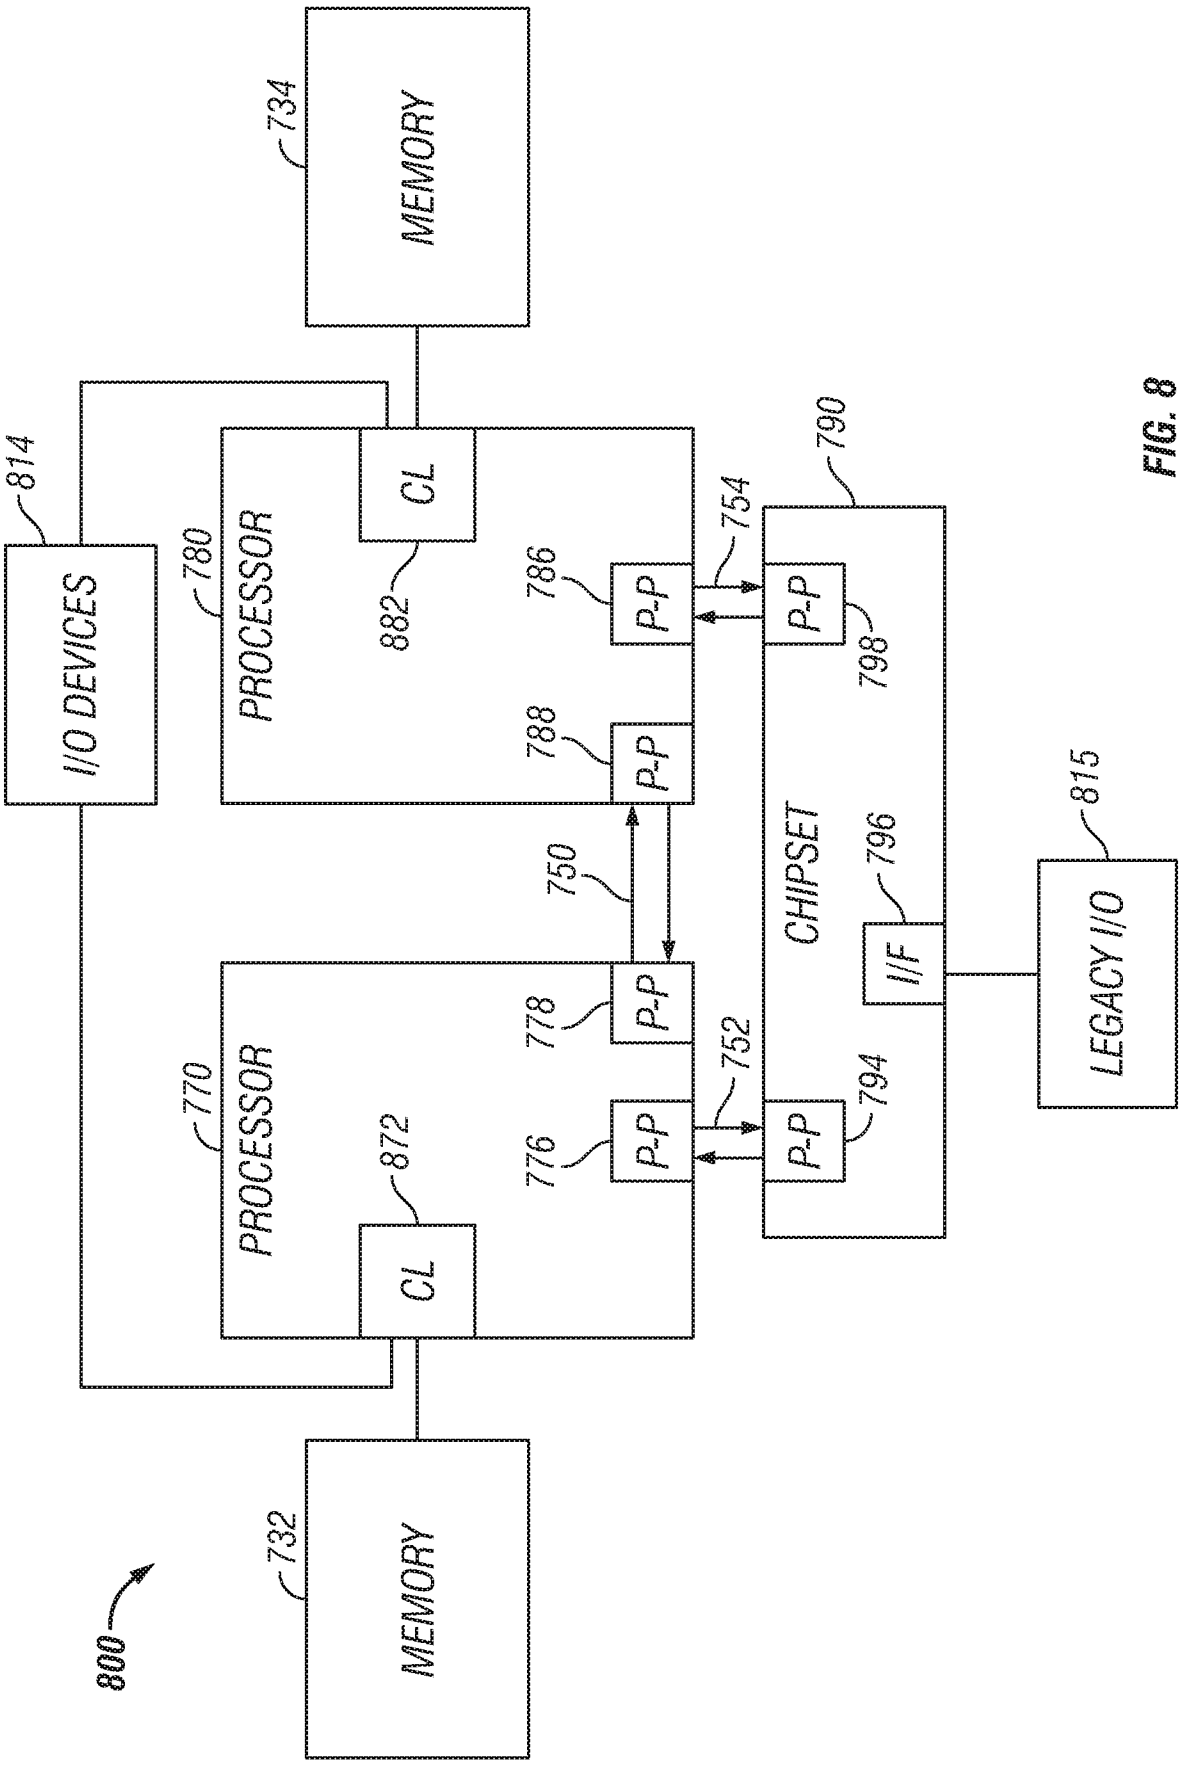

800
MEMORY
732
MEMORY
734
I/O DEVICES
814
PROCESSOR
770
CL
872
PROCESSOR
780
CL
882
P-P
776
P-P
778
P-P
786
P-P
788
750
752
754
CHIPSET
790
P-P
794
P-P
798
I/F
796
LEGACY I/O
815

FIG. 8

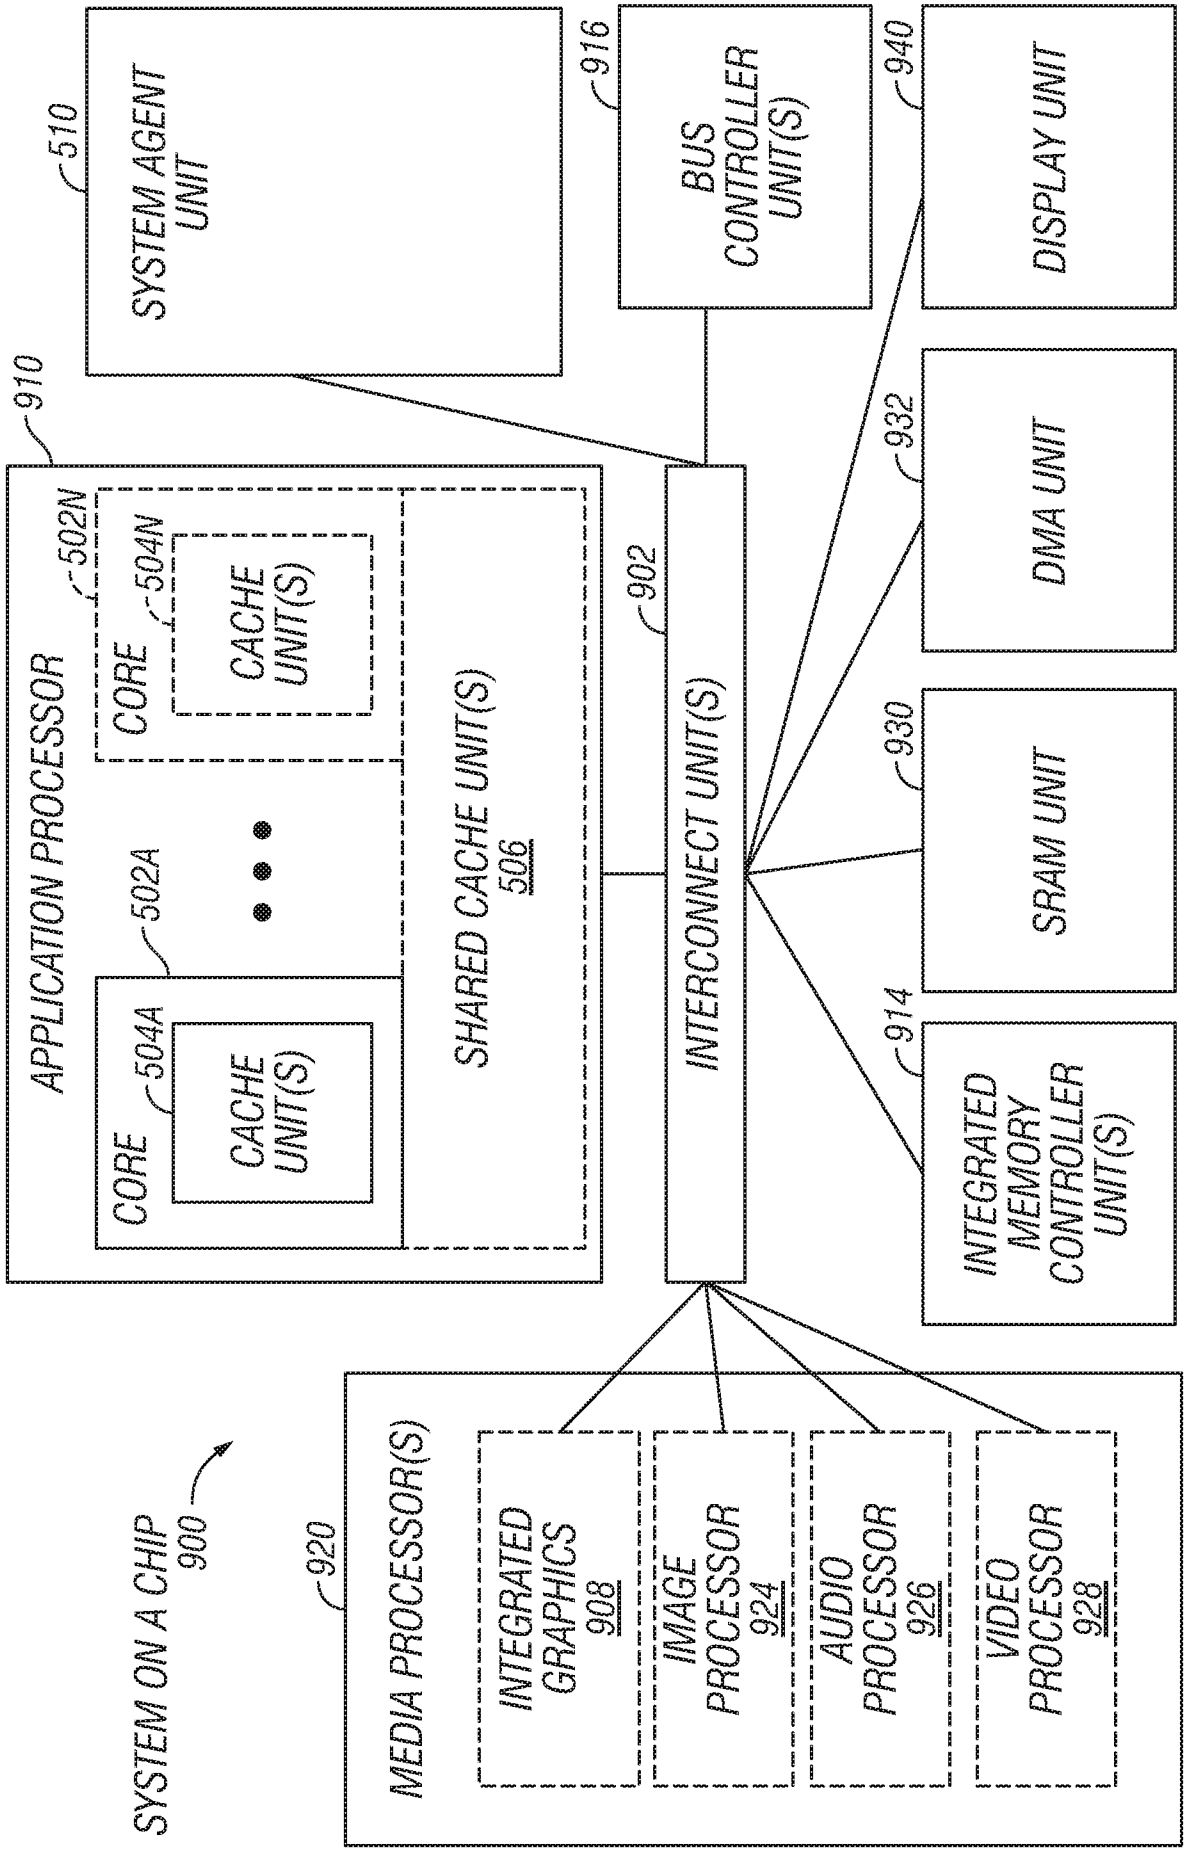

SYSTEM ON A CHIP
900
MEDIA PROCESSOR(S)
920
INTEGRATED GRAPHICS
908
IMAGE PROCESSOR
924
AUDIO PROCESSOR
926
VIDEO PROCESSOR
928
APPLICATION PROCESSOR
910
CORE
504A
CACHE UNIT(S)
502A
CORE
504N
CACHE UNIT(S)
502N
SHARED CACHE UNIT(S)
506
INTERCONNECT UNIT(S)
902
SYSTEM AGENT UNIT
510
BUS CONTROLLER UNIT(S)
916
DISPLAY UNIT
940
DMA UNIT
932
SRAM UNIT
930
INTEGRATED MEMORY CONTROLLER UNIT(S)
914

FIG. 9

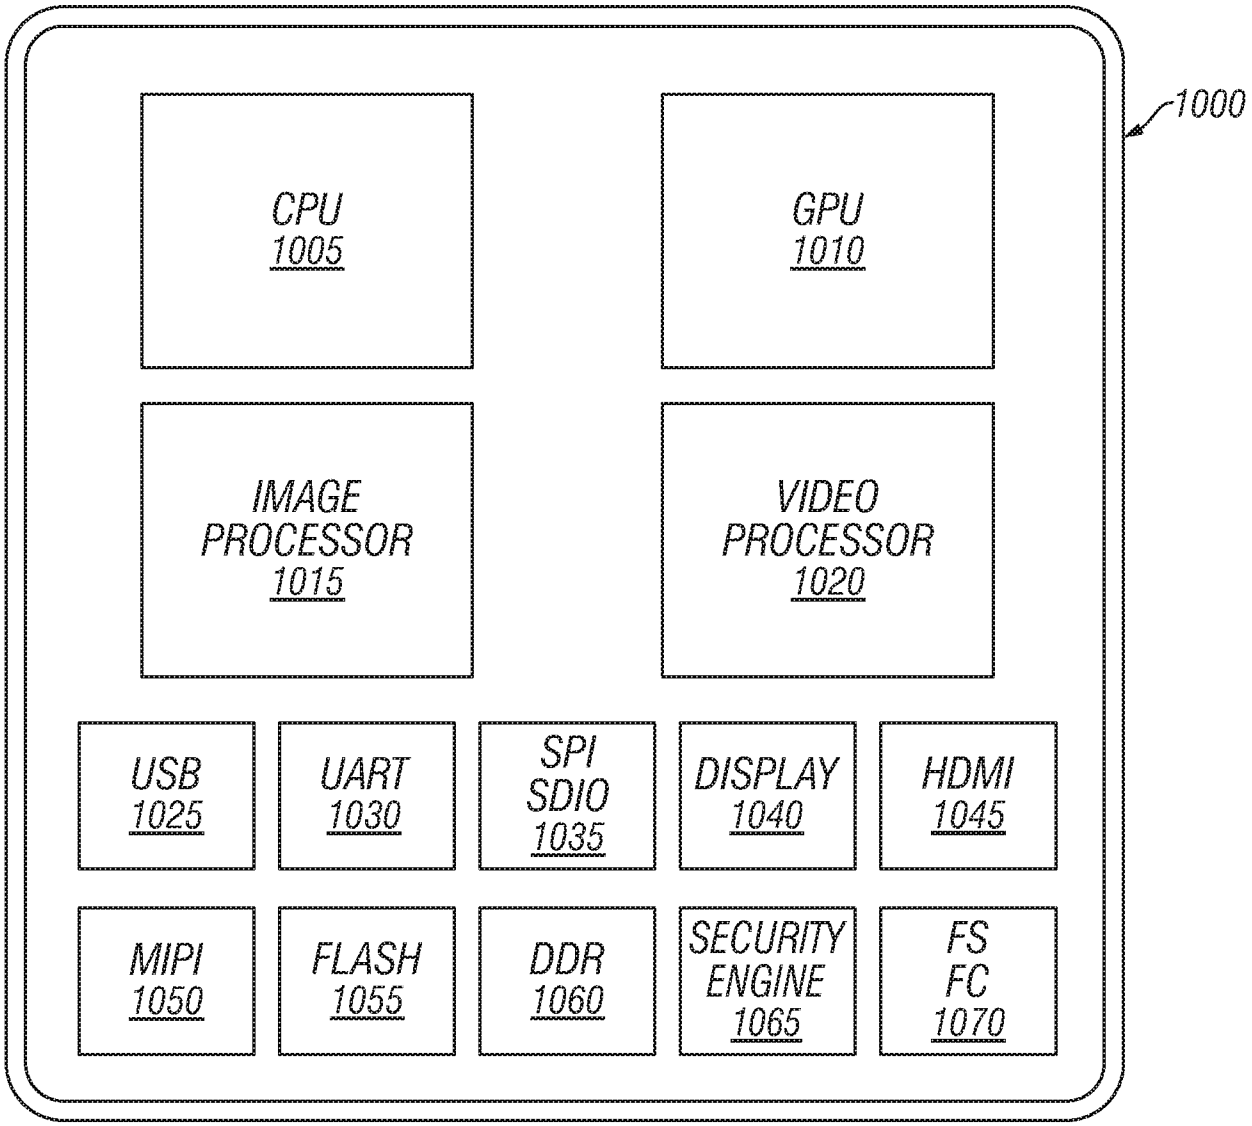

1000
CPU
1005
GPU
1010
IMAGE
PROCESSOR
1015
VIDEO
PROCESSOR
1020
USB
1025
UART
1030
SPI
SDIO
1035
DISPLAY
1040
HDMI
1045
MIPI
1050
FLASH
1055
DDR
1060
SECURITY
ENGINE
1065
FS
FC
1070

FIG. 10

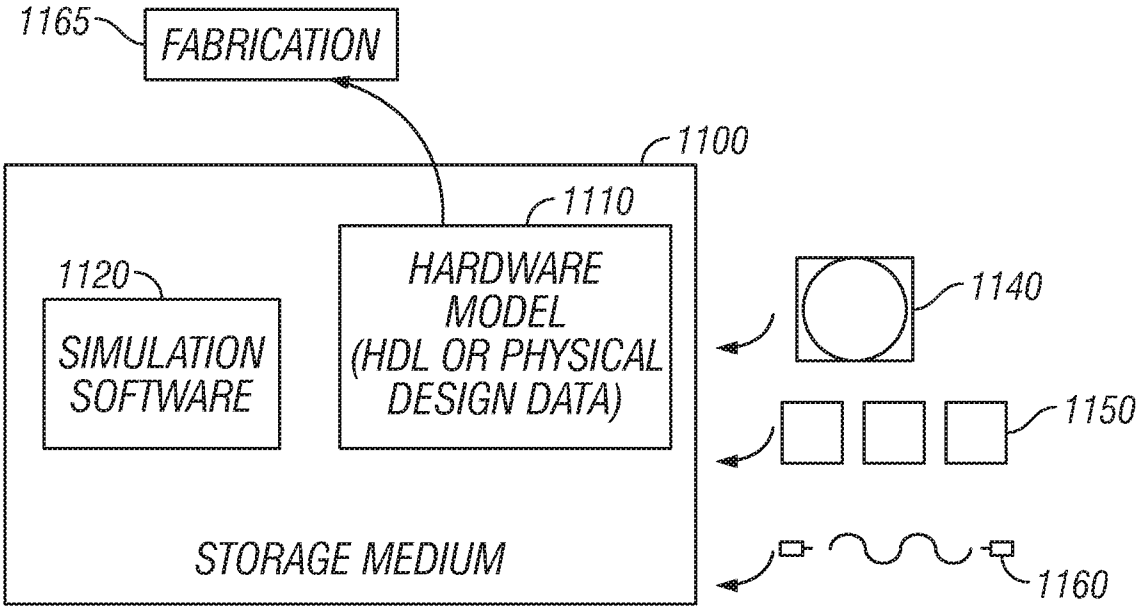

1165
FABRICATION
1100
1110
HARDWARE MODEL (HDL OR PHYSICAL DESIGN DATA)
1120
SIMULATION SOFTWARE
1140
1150
STORAGE MEDIUM
1160

FIG. 11

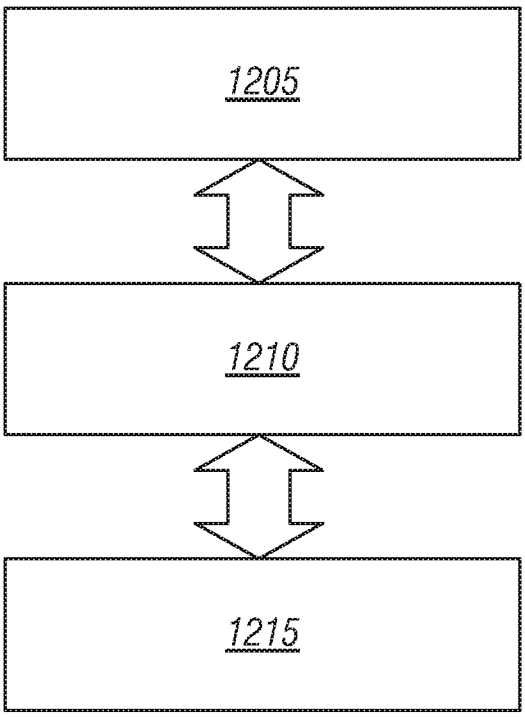

1205
1210
1215

FIG. 12

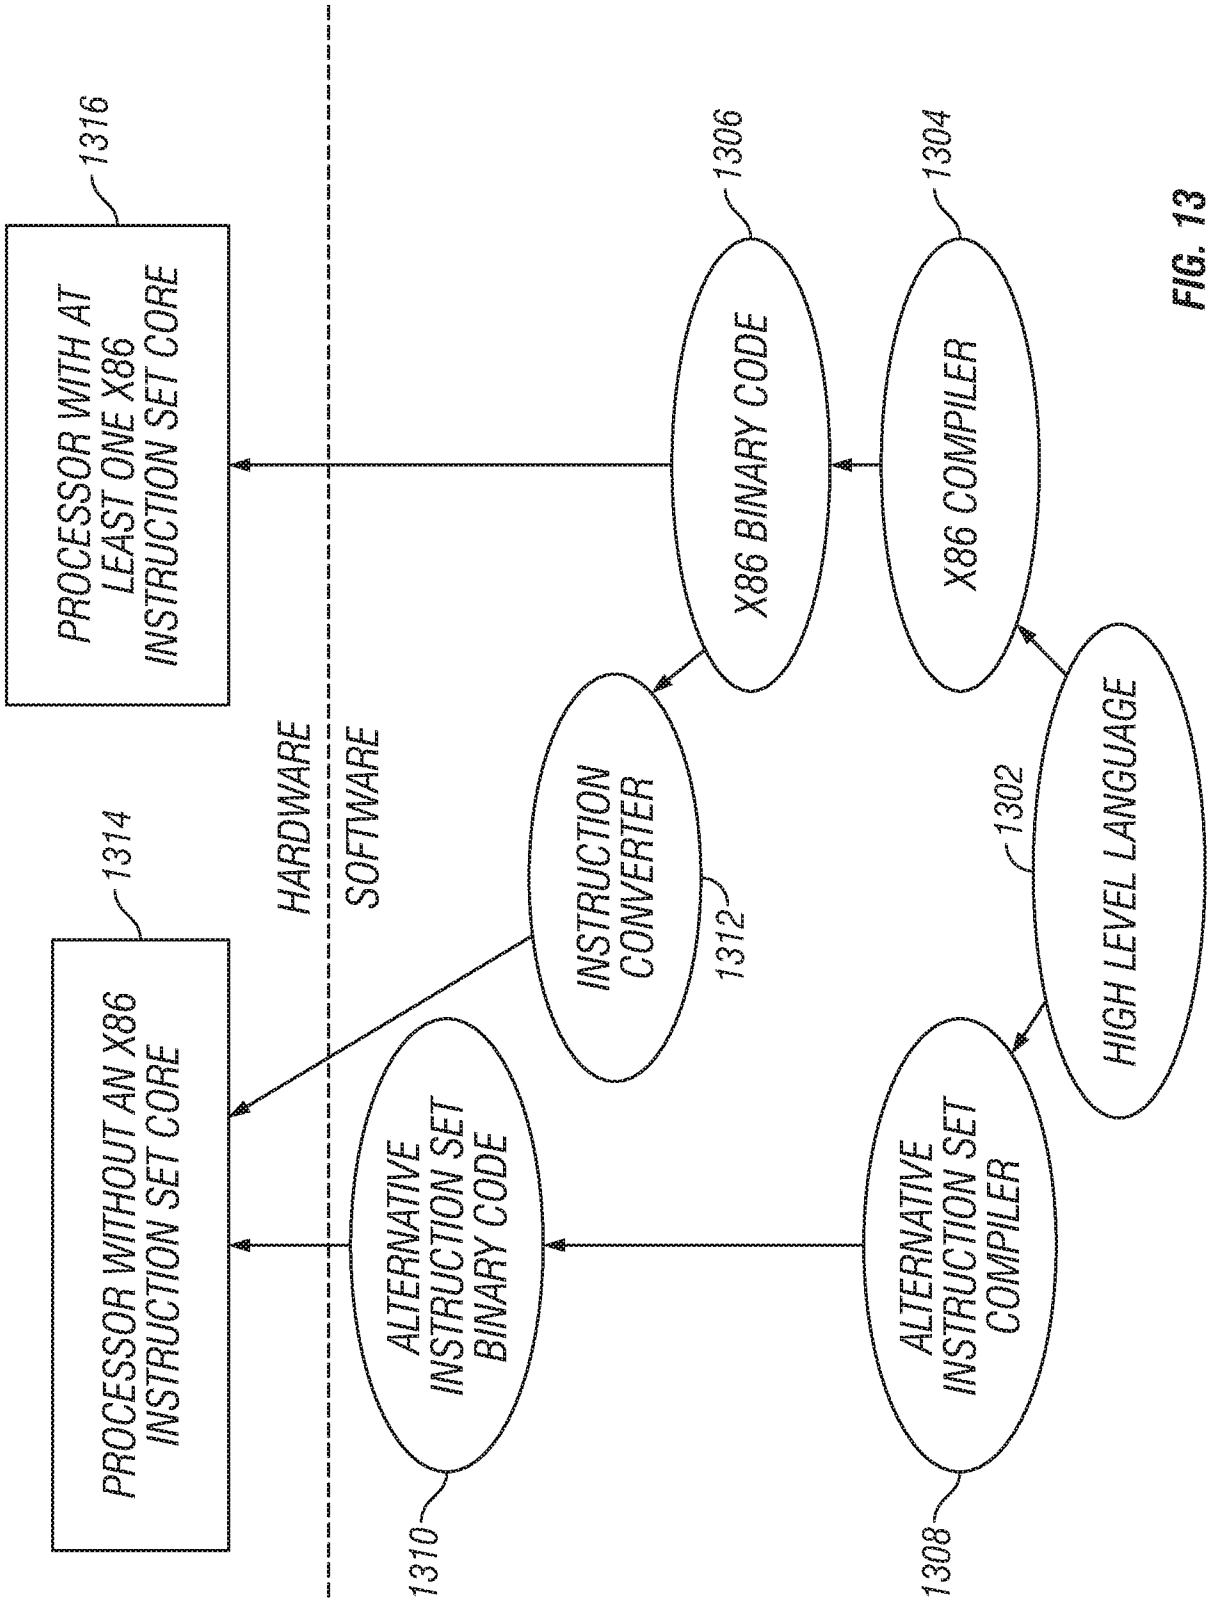

PROCESSOR WITH AT LEAST ONE X86 INSTRUCTION SET CORE
1316
PROCESSOR WITHOUT AN X86 INSTRUCTION SET CORE
1314
HARDWARE
SOFTWARE
X86 BINARY CODE
1306
X86 COMPILER
1304
INSTRUCTION CONVERTER
1312
HIGH LEVEL LANGUAGE
1302
ALTERNATIVE INSTRUCTION SET BINARY CODE
1310
ALTERNATIVE INSTRUCTION SET COMPILER
1308

FIG. 13

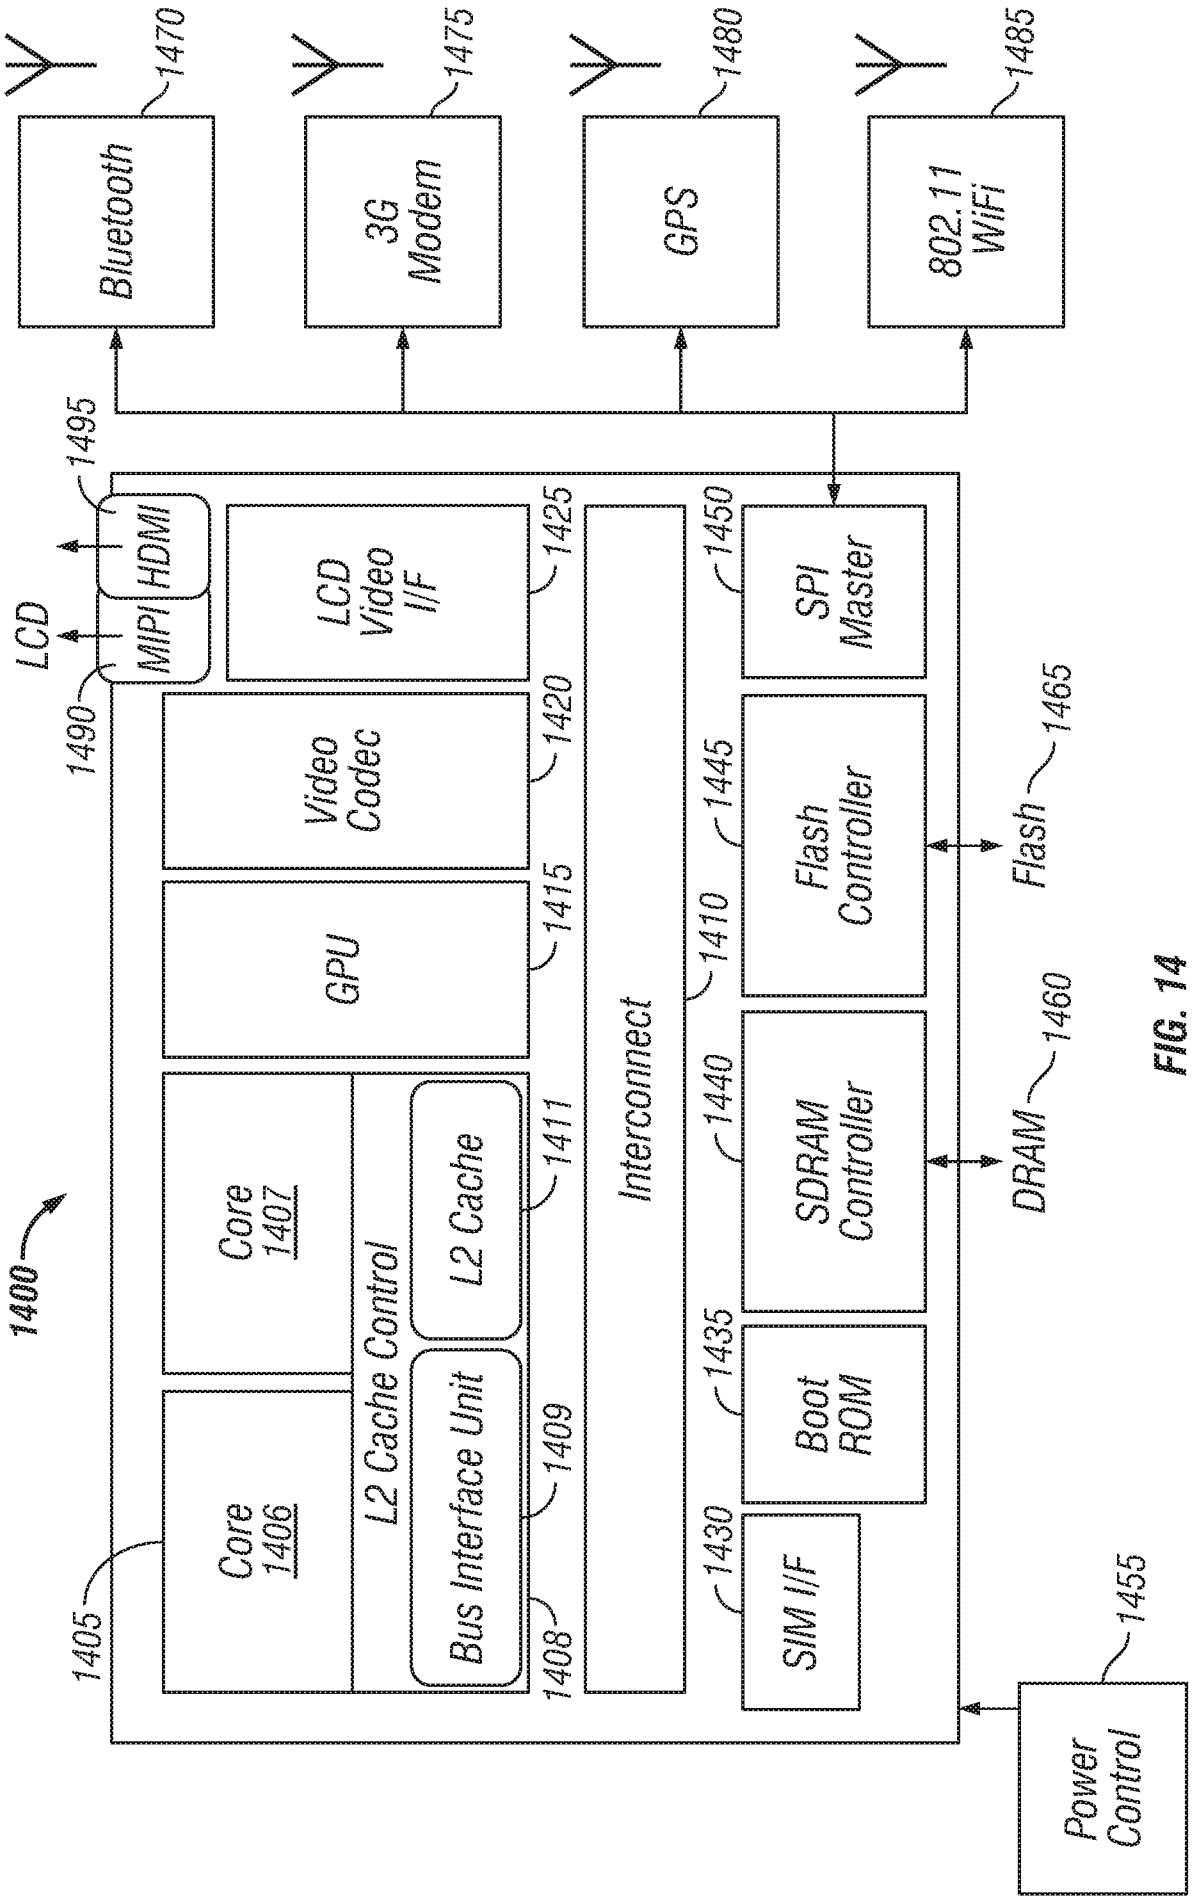

1400
1405
Core 1406
L2 Cache Control
Bus Interface Unit
1409
1408
L2 Cache
1411
Core 1407
GPU
1415
Video Codec
1420
LCD Video I/F
1425
Interconnect
1410
SIM I/F
1430
Boot ROM
1435
SDRAM Controller
1440
Flash Controller
1445
SPI Master
1450
MIPI
HDMI
1490
LCD
1495
Bluetooth
1470
3G Modem
1475
GPS
1480
802.11 WiFi
1485
DRAM
1460
Flash
1465
Power Control
1455

FIG. 14

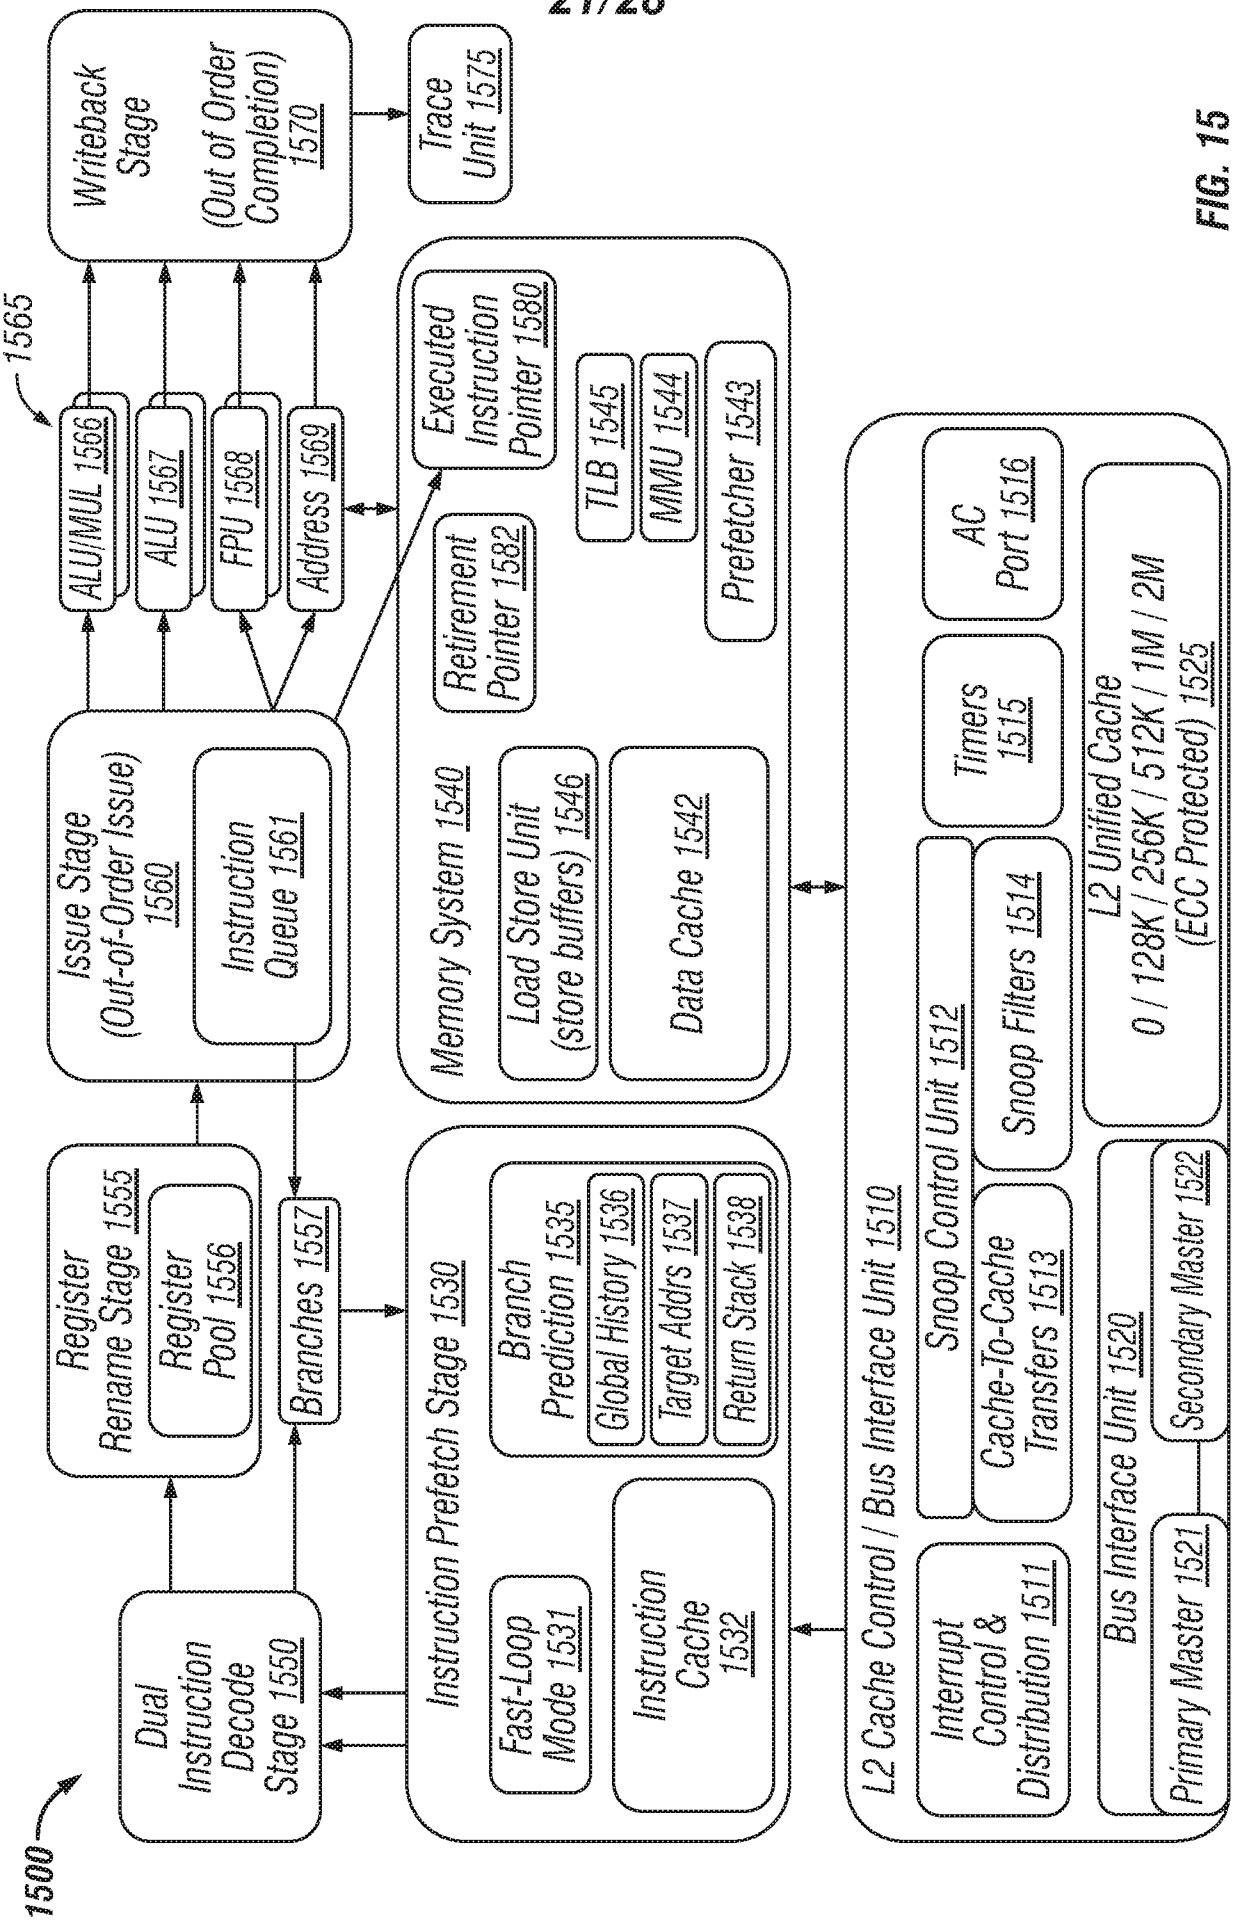

1500
1565
Dual Instruction Decode Stage 1550
Register Rename Stage 1555
Register Pool 1556
Branches 1557
Issue Stage (Out-of-Order Issue) 1560
Instruction Queue 1561
ALU/MUL 1566
ALU 1567
FPU 1568
Address 1569
Writeback Stage (Out of Order Completion) 1570
Trace Unit 1575
Instruction Prefetch Stage 1530
Fast-Loop Mode 1531
Instruction Cache 1532
Branch Prediction 1535
Global History 1536
Target Addrs 1537
Return Stack 1538
Memory System 1540
Load Store Unit (store buffers) 1546
Data Cache 1542
Retirement Pointer 1582
Executed Instruction Pointer 1580
TLB 1545
MMU 1544
Prefetcher 1543
L2 Cache Control / Bus Interface Unit 1510
Interrupt Control & Distribution 1511
Cache-To-Cache Transfers 1513
Snoop Control Unit 1512
Snoop Filters 1514
Timers 1515
AC Port 1516
Bus Interface Unit 1520
Primary Master 1521
Secondary Master 1522
L2 Unified Cache 0 / 128K / 256K / 512K / 1M / 2M (ECC Protected) 1525

FIG. 15

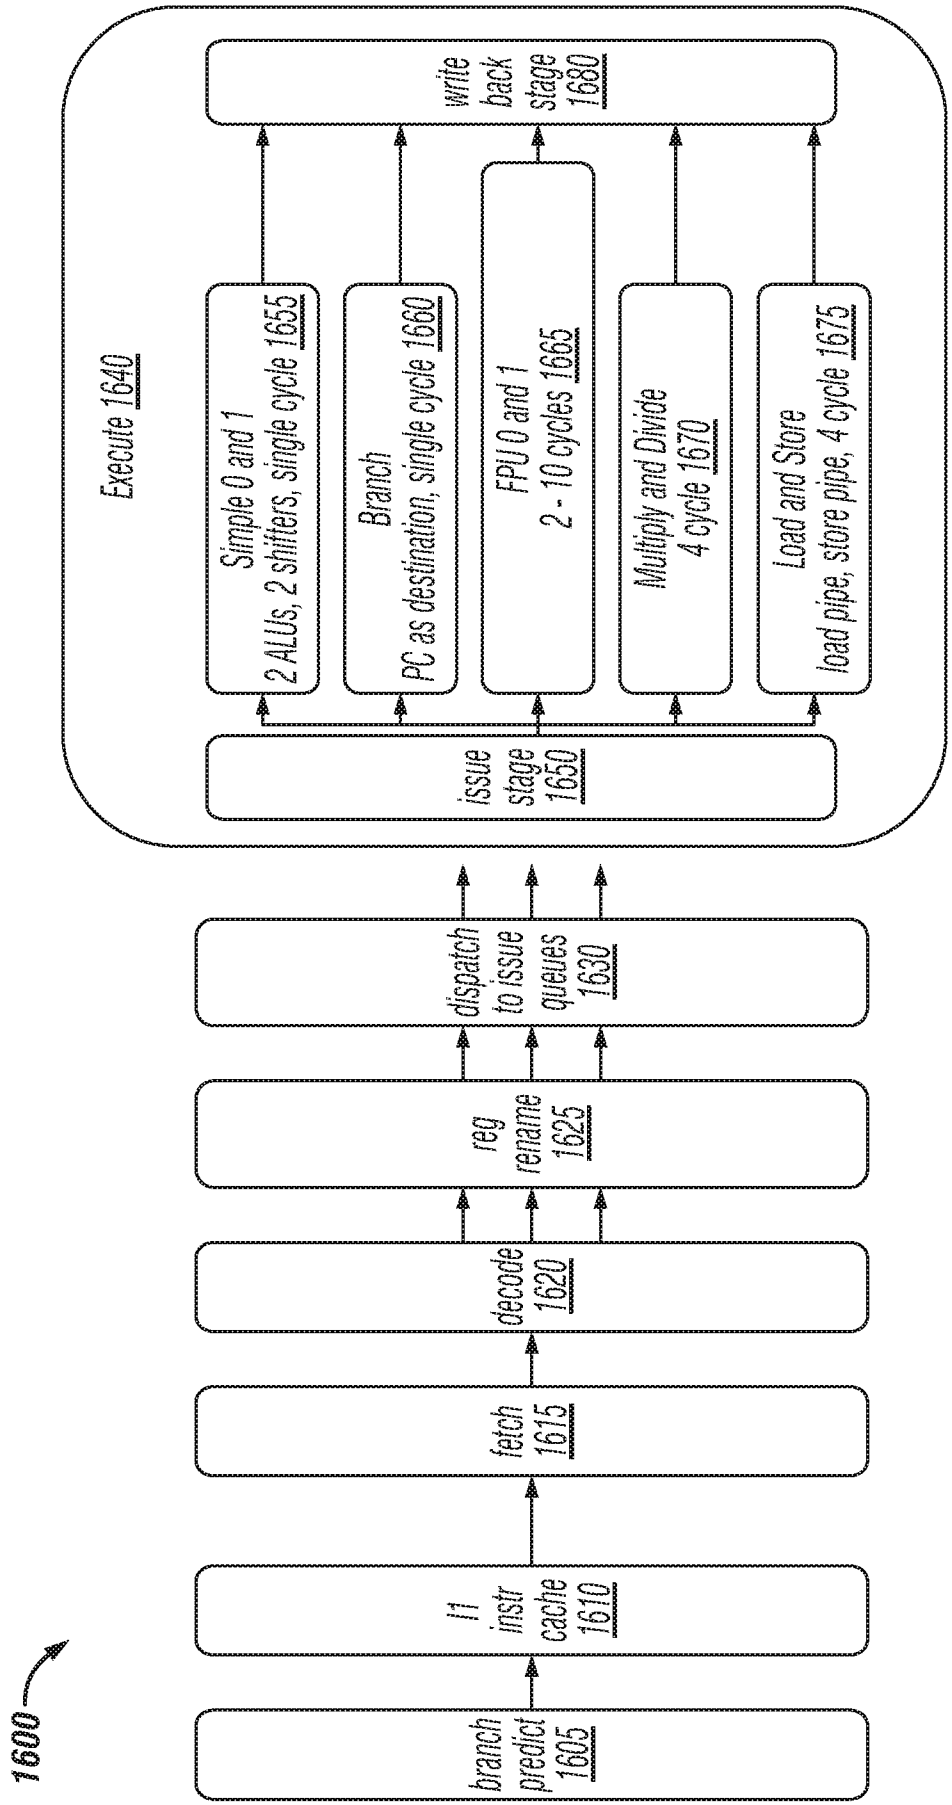

1600
branch predict 1605
I1 instr cache 1610
fetch 1615
decode 1620
reg rename 1625
dispatch to issue queues 1630
Execute 1640
issue stage 1650
Simple 0 and 1
2 ALUs, 2 shifters, single cycle 1655
Branch
PC as destination, single cycle 1660
FPU 0 and 1
2 - 10 cycles 1665
Multiply and Divide
4 cycle 1670
Load and Store
load pipe, store pipe, 4 cycle 1675
write back stage 1680

FIG. 16

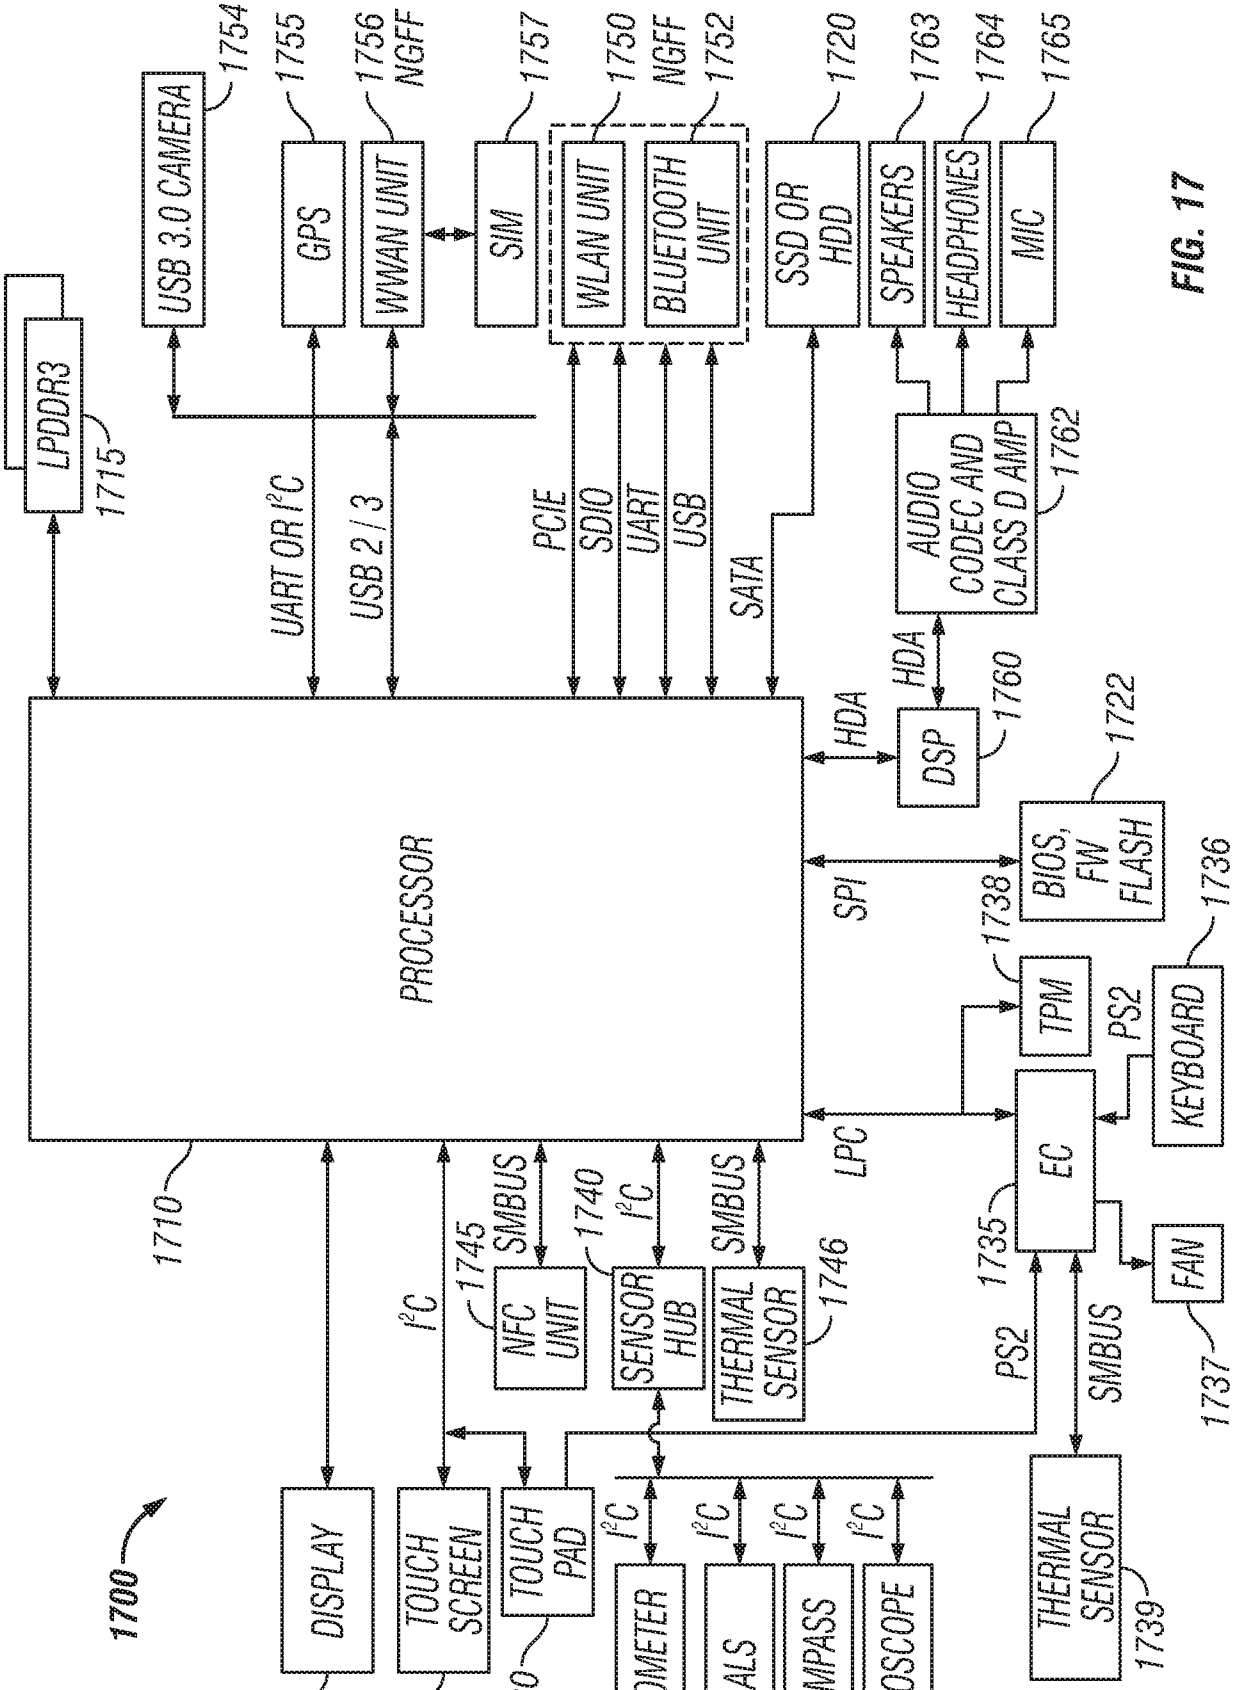

1700
PROCESSOR
1710
LPDDR3
1715
USB 3.0 CAMERA
1754
GPS
1755
WWAN UNIT
1756
NGFF
SIM
1757
WLAN UNIT
1750
NGFF
BLUETOOTH UNIT
1752
SSD OR HDD
1720
SPEAKERS
1763
HEADPHONES
1764
MIC
1765
AUDIO CODEC AND CLASS D AMP
1762
UART OR I2C
USB 2 / 3
PCIE
SDIO
UART
USB
SATA
HDA
DSP
1760
HDA
BIOS, FW FLASH
1722
SPI
1738
TPM
PS2
KEYBOARD
1736
LPC
EC
1735
FAN
1737
SMBUS
PS2
THERMAL SENSOR
1739
I2C
1745
SMBUS
NFC UNIT
1740
I2C
SENSOR HUB
SMBUS
THERMAL SENSOR
1746
DISPLAY
1724
TOUCH SCREEN
1725
TOUCH PAD
1730
I2C
ACCELEROMETER
1741
I2C
ALS
1742
I2C
COMPASS
1743
I2C
GYROSCOPE
1744

FIG. 17

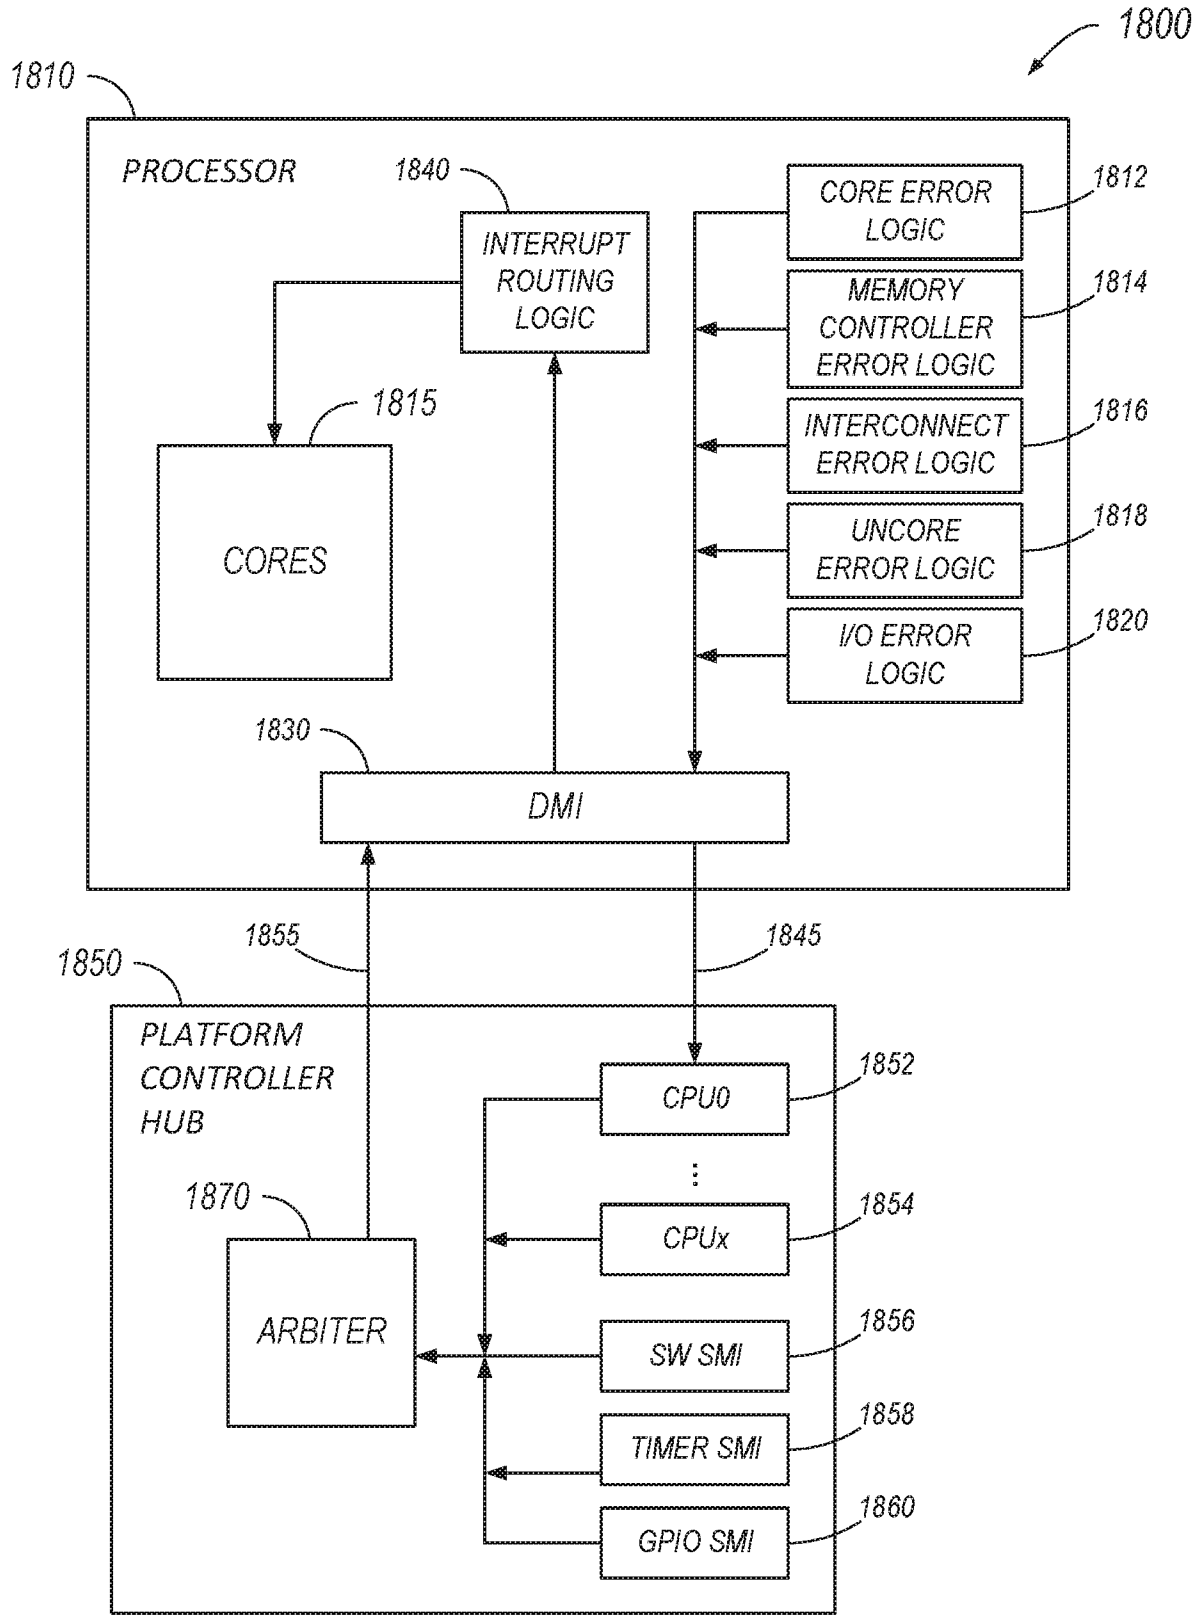

1800
1810
PROCESSOR
1840
INTERRUPT ROUTING LOGIC
1815
CORES
1830
DMI
CORE ERROR LOGIC
1812
MEMORY CONTROLLER ERROR LOGIC
1814
INTERCONNECT ERROR LOGIC
1816
UNCORE ERROR LOGIC
1818
I/O ERROR LOGIC
1820
1855
1845
1850
PLATFORM CONTROLLER HUB
1870
ARBITER
CPU0
1852
CPUx
1854
SW SMI
1856
TIMER SMI
1858
GPIO SMI
1860

FIG. 18

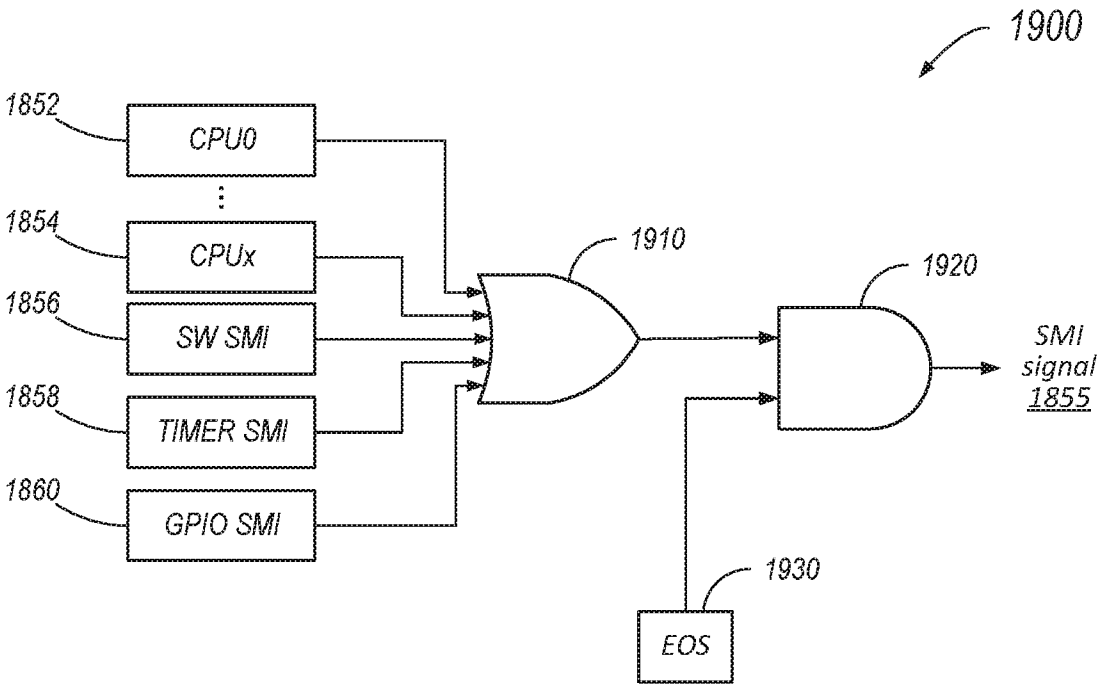

1900
1852
CPU0
1854
CPUx
1856
SW SMI
1858
TIMER SMI
1860
GPIO SMI
1910
1920
SMI signal 1855
1930
EOS

FIG. 19

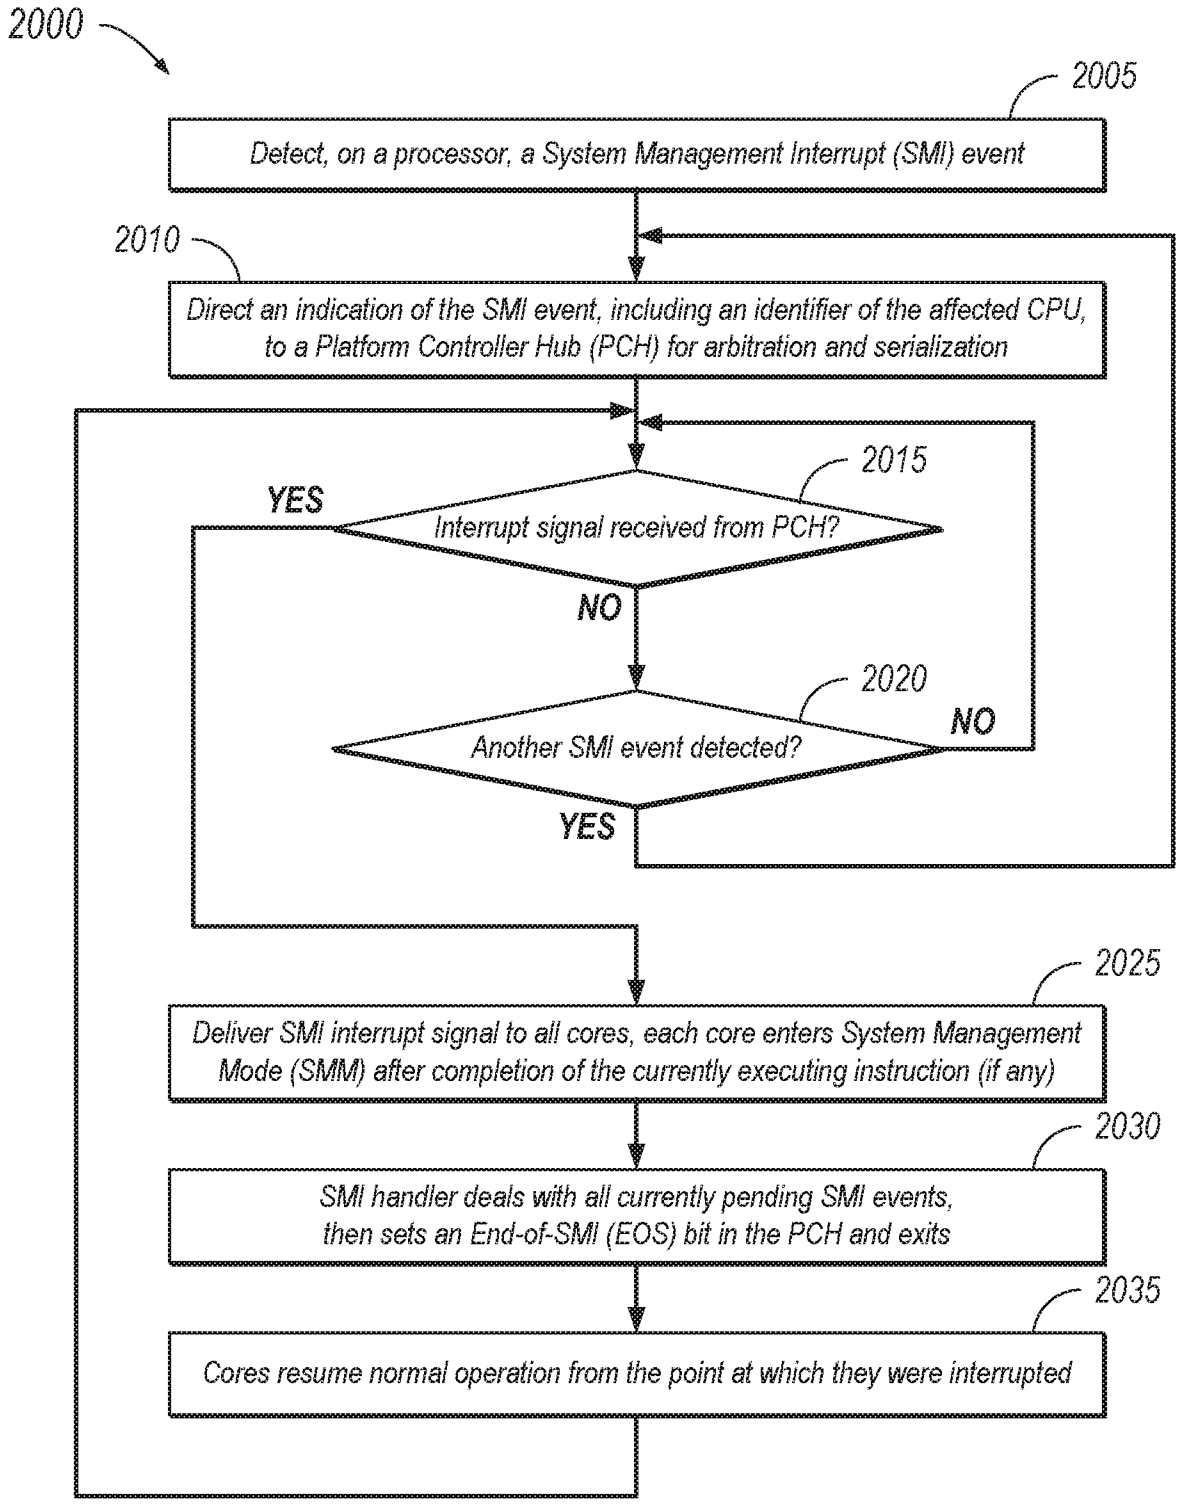

2000
2005
Detect, on a processor, a System Management Interrupt (SMI) event
2010
Direct an indication of the SMI event, including an identifier of the affected CPU, to a Platform Controller Hub (PCH) for arbitration and serialization
2015
YES
Interrupt signal received from PCH?
NO
2020
NO
Another SMI event detected?
YES
2025
Deliver SMI interrupt signal to all cores, each core enters System Management Mode (SMM) after completion of the currently executing instruction (if any)
2030
SMI handler deals with all currently pending SMI events, then sets an End-of-SMI (EOS) bit in the PCH and exits
2035
Cores resume normal operation from the point at which they were interrupted

FIG. 20

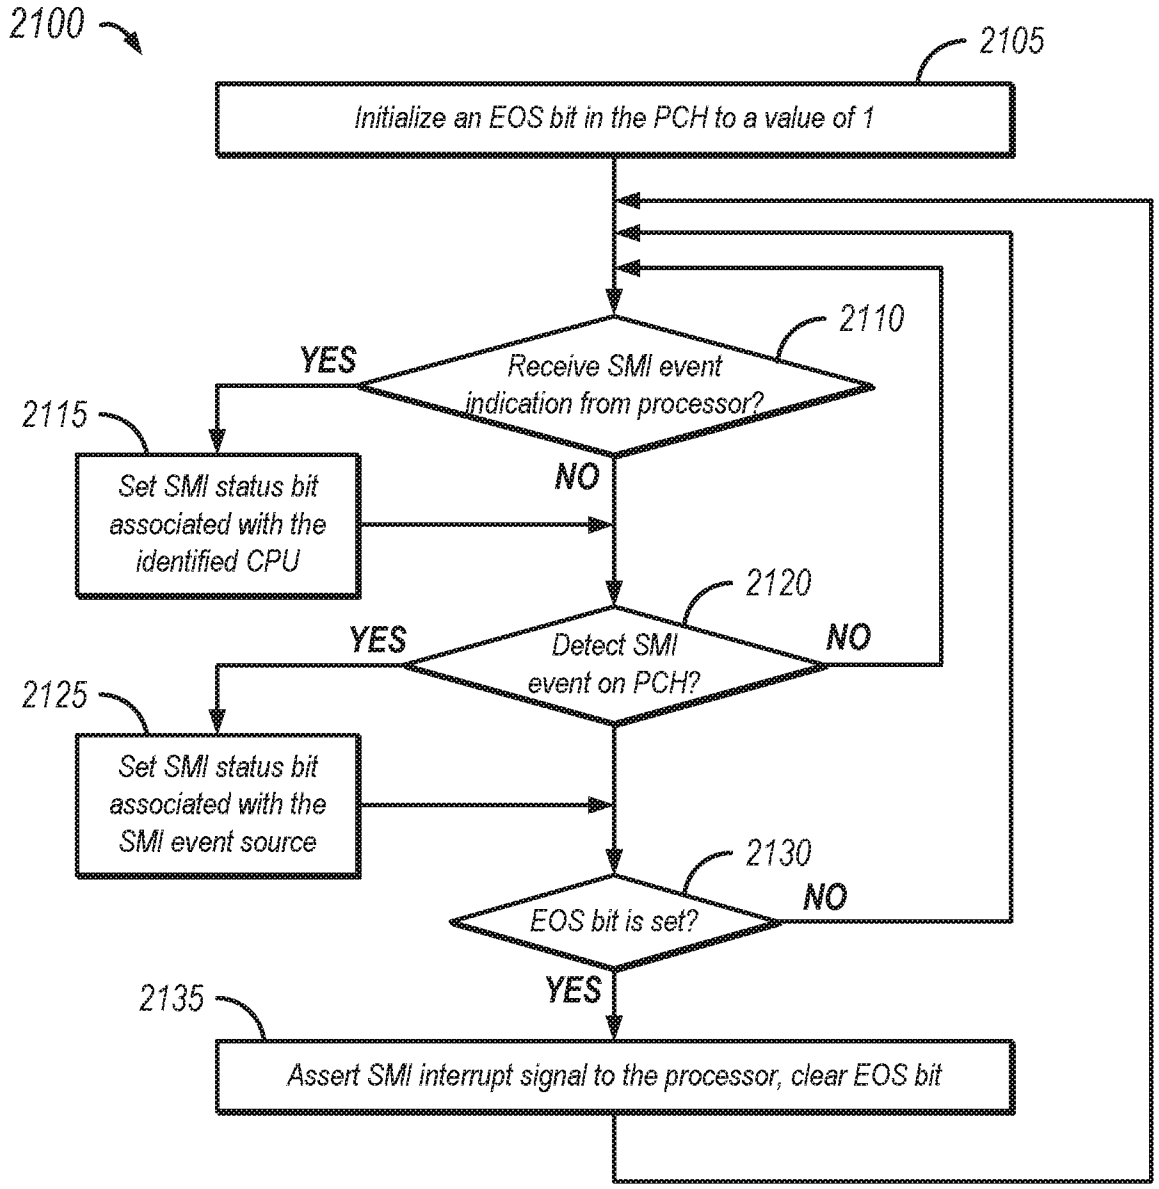

2100
2105
Initialize an EOS bit in the PCH to a value of 1
2110
Receive SMI event indication from processor?
YES
NO
2115
Set SMI status bit associated with the identified CPU
2120
Detect SMI event on PCH?
YES
NO
2125
Set SMI status bit associated with the SMI event source
2130
EOS bit is set?
NO
YES
2135
Assert SMI interrupt signal to the processor, clear EOS bit

FIG. 21

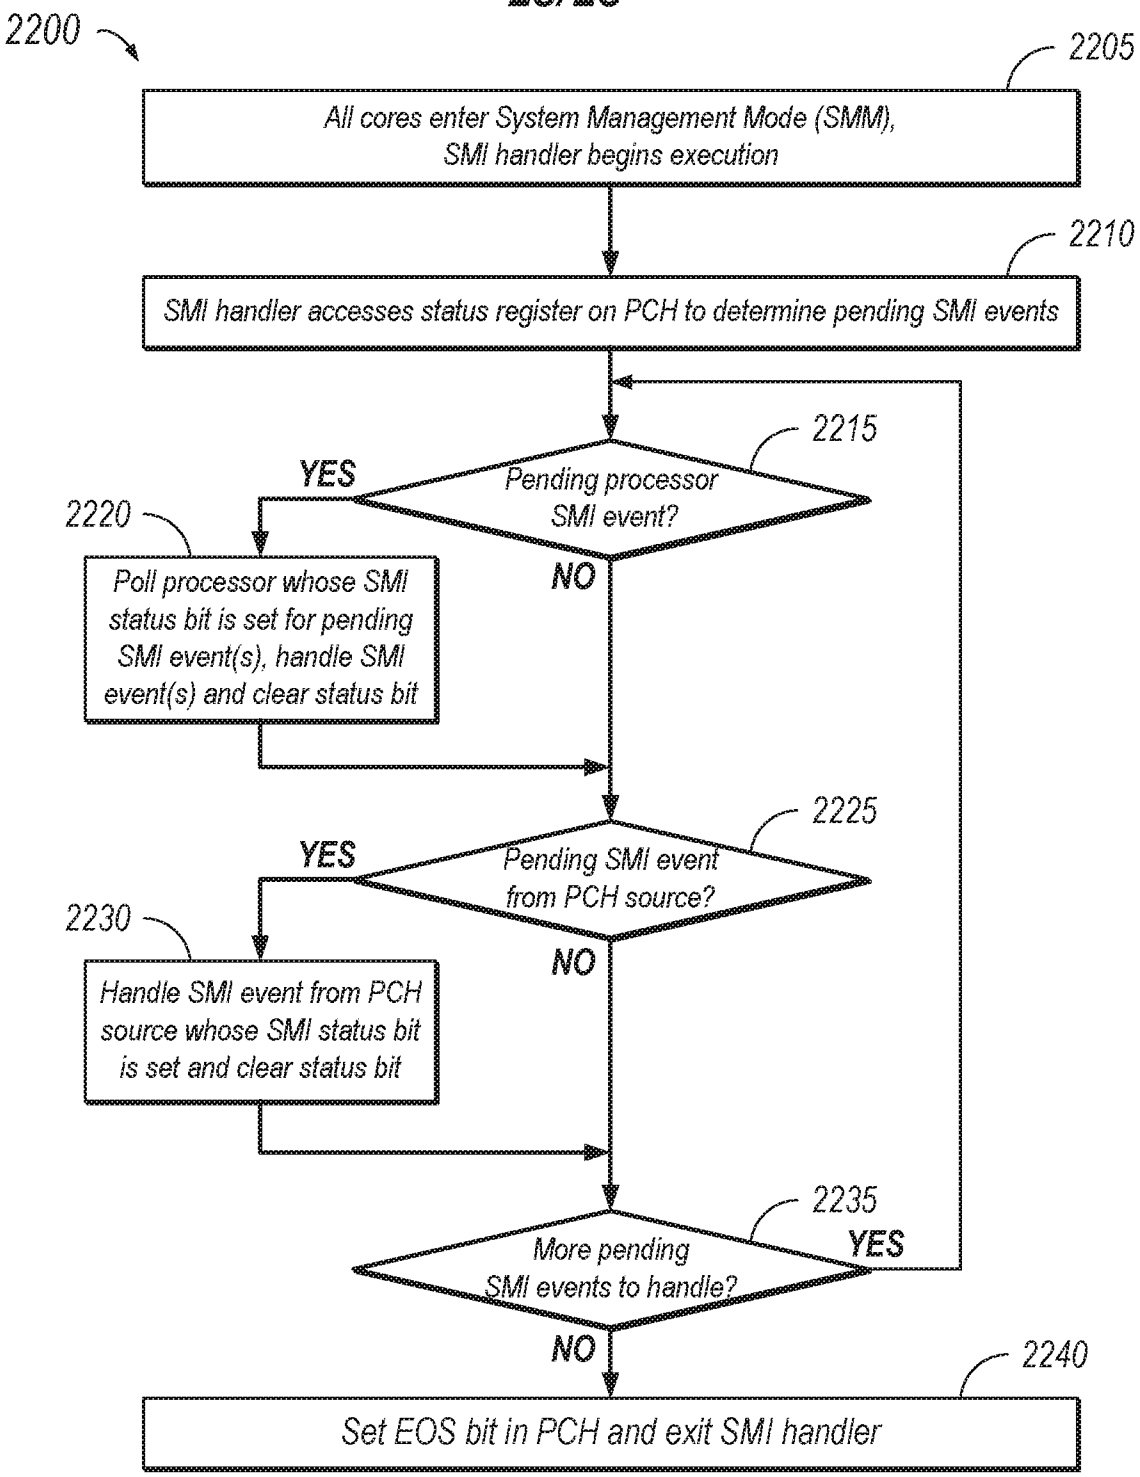

2200
2205
All cores enter System Management Mode (SMM), SMI handler begins execution
2210
SMI handler accesses status register on PCH to determine pending SMI events
2215
Pending processor SMI event?
YES
NO
2220
Poll processor whose SMI status bit is set for pending SMI event(s), handle SMI event(s) and clear status bit
2225
Pending SMI event from PCH source?
YES
NO
2230
Handle SMI event from PCH source whose SMI status bit is set and clear status bit
2235
More pending SMI events to handle?
YES
NO
2240
Set EOS bit in PCH and exit SMI handler

FIG. 22

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/018939

A. CLASSIFICATION OF SUBJECT MATTER

G06F 13/24(2006.01)i, G06F 15/16(2006.01)i, G06F 13/14(2006.01)i, G06F 9/48(2006.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F 13/24; G06F 13/26; G06F 13/36; G06F 13/00; G06F 15/16; G06F 13/14; G06F 9/48

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: multiprocessor, controller hub, arbiter, serialize, various sources, SMI event, interrupt handler, end-of-SMI, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006-0190668 A1 (GIORA BIRAN et al.) 24 August 2006 See paragraphs [0025] and [0052]-[0058]; and figures 1 and 9.	1-22,24-25
Y		23
Y	US 6,219,741 B1 (STEPHEN S. PAWLOWSKI et al.) 17 April 2001 See column 10, line 1 - column 11, line 17; and figures 8-9.	23
A	US 2008-0276027 A1 (YASUHARU HAGITA) 06 November 2008 See paragraphs [0033]-[0034], [0041], and [0048]-[0058]; and figures 2-6.	1-25
A	US 2001-0027502 A1 (TIMOTHY C. BRONSON et al.) 04 October 2001 See paragraphs [0063]-[0066] and figure 3.	1-25
A	US 5,560,018 A (JAMES F. MACON, JR. et al.) 24 September 1996 See column 7, line 16 - column 8, line 55; claim 1; and figures 3-5.	1-25

☐ 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	Date of mailing of the international search report
30 May 2017 (30.05.2017)	**30 May 2017 (30.05.2017)**

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
NHO, Ji Myong
Telephone No. +82-42-481-8528

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2017/018939

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006-0190668 A1	24/08/2006	US 7234017 B2	19/06/2007
US 6219741 B1	17/04/2001	AU 1309399 A	28/06/1999
		AU 1999-13093 A1	28/06/1999
		CN 1225708 C	02/11/2005
		CN 1290365 A	04/04/2001
		DE 69822221 T2	29/07/2004
		EP 1051678 A1	15/11/2000
		EP 1051678 A4	17/01/2001
		EP 1051678 B1	03/03/2004
		KR 10-0380197 B1	11/04/2003
		TW 504612 B	01/10/2002
		WO 99-30244 A1	17/06/1999
US 2008-0276027 A1	06/11/2008	JP 2009-169928 A	30/07/2009
		JP 5200646 B2	05/06/2013
		US 7805556 B2	28/09/2010
US 2001-0027502 A1	04/10/2001	US 6065088 A	16/05/2000
		US 6279064 B1	21/08/2001
		US 6442634 B2	27/08/2002
US 5560018 A	24/09/1996	None	