



(11) **EP 4 439 285 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
02.10.2024 Bulletin 2024/40

(51) International Patent Classification (IPC):
G06F 9/30 (2018.01)

(21) Application number: **23214578.9**

(52) Cooperative Patent Classification (CPC):
G06F 9/30167; G06F 9/30036; G06F 9/30185

(22) Date of filing: **06.12.2023**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC ME MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA
Designated Validation States:
KH MA MD TN

(72) Inventors:
• **Subramaniam, Kameswar Austin, 78750 (US)**
• **Brandt, Jason Austin, 78703 (US)**

(30) Priority: **30.03.2023 US 202318193226**

(74) Representative: **Goddar, Heinz J. et al Boehmert & Boehmert Anwaltspartnerschaft mbB Pettenkofenstrasse 22 80336 München (DE)**

(71) Applicant: **Intel Corporation Santa Clara, CA 95054 (US)**

(54) **SUPPORT FOR A 32-BIT IMMEDIATE AS INDICATED BY AN INSTRUCTION PREFIX**

(57) Techniques for providing and using a 32-bit immediate in an instruction are described. In some examples, an instance of a single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a

field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction.

EP 4 439 285 A1

Description

BACKGROUND

5 **[0001]** Modern x86 processors define thousands of system registers, typically referred to as Model Specific Registers (MSRs) or control and/or status registers. Operating system or privileged software like BIOS, OS, drivers, and hypervisors use privileged instructions - read MSR (RDMSR) and write MSR (WRMSR) - to read/write these registers. However, application software cannot use these instructions, and thus needs to use system calls to the OS if it needs read/write access to any of this state. System calls add latency overhead during software execution and are thus have a detrimental effect on software performance.

BRIEF DESCRIPTION OF DRAWINGS

15 **[0002]** Various examples in accordance with the present disclosure will be described with reference to the drawings, in which:

FIG. 1 illustrates an example execution of a read control and/or status register (e.g., URDMSR) instruction.

20 FIG. 2 illustrates examples of descriptions of the URDMSR including an opcode, operands to use, a CPUID location to indicate support for the instruction, etc.

FIG. 3 illustrates an example execution of a write control and/or status register (e.g., UWRMSR) instruction.

25 FIG. 4 illustrates examples of descriptions of the UWRMSR including an opcode, operands to use, a CPUID location to indicate support for the instruction, etc.

FIG. 5 illustrates examples of a user-level access control bitmap to determine when a read and/or write of a control and/or status register is to be allowed.

30 FIG. 6 illustrates examples of a method to process a URDMSR instruction.

FIG. 7 illustrates examples of a method to process a URDMSR instruction.

35 FIG. 8 illustrates examples of a method to process a UWRMSR instruction.

FIG. 9 illustrates examples of a method to process a UWRMSR instruction.

40 FIG. 10 illustrates examples of an instruction format that support the UWRMSR or URDMSR instruction when a 32-bit immediate value is not used.

FIG. 11 illustrates examples of a prefix encoding to indicate a 32-bit immediate value usage.

FIG. 12 illustrates examples of a method to process an instruction.

45 FIG. 13 illustrates examples of computing hardware to process at least a UWRMSR and/or URDMSR instruction.

FIG. 14 illustrates an example computing system.

50 FIG. 15 illustrates a block diagram of an example processor and/or System on a Chip (SoC) that may have one or more cores and an integrated memory controller.

FIG. 16 is a block diagram illustrating a computing system 1600 configured to implement one or more aspects of the examples described herein.

55 FIG. 17A illustrates examples of a parallel processor.

FIG. 17B illustrates examples of a block diagram of a partition unit.

FIG. 17C illustrates examples of a block diagram of a processing cluster within a parallel processing unit.

FIG. 17D illustrates examples of a graphics multiprocessor in which the graphics multiprocessor couples with the pipeline manager of the processing cluster.

5

FIGS. 18A-18C illustrate additional graphics multiprocessors, according to examples.

FIG. 19 shows a parallel compute system 1900, according to some examples.

10

FIGS. 20A-20B illustrate a hybrid logical/physical view of a disaggregated parallel processor, according to examples described herein.

FIG. 21(A) is a block diagram illustrating both an example in-order pipeline and an example register renaming, out-of-order issue/execution pipeline according to examples.

15

FIG. 21(B) is a block diagram illustrating both an example in-order architecture core and an example register renaming, out-of-order issue/execution architecture core to be included in a processor according to examples.

FIG. 22 illustrates examples of execution unit(s) circuitry, such as execution unit(s) circuitry.

20

FIG. 23 is a block diagram of a register architecture according to some examples.

FIG. 24 illustrates examples of an instruction format.

25

FIG. 25 illustrates examples of an addressing information field.

FIG. 26 illustrates examples of a first prefix.

FIGS. 27(A)-(D) illustrate examples of how the R, X, and B fields of the first prefix are used.

30

FIGS. 28(A)-(B) illustrate examples of a second prefix.

FIG. 29 illustrates examples of a third prefix.

35

FIGS. 30A-30B illustrate thread execution logic including an array of processing elements employed in a graphics processor core according to examples described herein.

FIG. 31 illustrates an additional execution unit, according to an example.

40

FIG. 32 is a block diagram illustrating a graphics processor instruction formats 3200 according to some examples.

FIG. 33 is a block diagram of another example of a graphics processor.

FIG. 34A is a block diagram illustrating a graphics processor command format according to some examples.

45

FIG. 34B is a block diagram illustrating a graphics processor command sequence according to an example.

FIG. 35 is a block diagram illustrating the use of a software instruction converter to convert binary instructions in a source ISA to binary instructions in a target ISA according to examples.

50

FIG. 36 is a block diagram illustrating an IP core development system 3600 that may be used to manufacture an integrated circuit to perform operations according to some examples.

DETAILED DESCRIPTION

55

[0003] The present disclosure relates to methods, apparatus, systems, and non-transitory computer-readable storage media for reading and/or writing control and/or status registers. Every time there is a need for application software to read or write control (system) registers, the approach has been to add a new custom instruction set architecture (ISA)

extension that is specific to that state, or groups of states.

[0004] The current solutions of adding state specific ISA extensions are undesirable for a few different reasons. First, this solution only works for a few specified MSR's out of the several thousand currently defined MSR's. Software continues to rely on system calls if needs to access any state that does not have custom user ISA. Second, each state specific ISA that is defined adds custom hardware and microcode - logic cost to the processor to decode the new instructions and coordinating product intercepts for hybrid processors. This adds more and more logic cost to each core as the number of such MSR's grows. Third, each new ISA that is added needs to be enumerated to software and enabled by software. Because these features are almost always disabled by default, new enables are required and both the operating system (OS) and virtual machine monitor (VMM) need to be enabled to use these new features. There can also be overhead in that OS/VMM enabling (e.g. to change the settings for each application/guest). Fourth, a state specific ISA is inflexible and is not able to accommodate late additions / changes to system registers. It takes many years between seeing there is a software need and releasing a processor with the support.

[0005] Examples of instructions (URDMSR and UWRMSR) are detailed herein to read and write control and/or status registers (e.g., MSR's) regardless of a current privilege level (CPL). In some examples, these instructions are called "user MSR" instructions. In some examples, one or both instructions can be executed in CPL3 or CPL0. The kernel and/or VMM are in control of which control and/or status registers can be read and written. Each instruction supports, in some examples, providing the control and/or status register address as immediate value operand; or in a register.

[0006] These instructions should avoid the design and validation complexity due to the constant addition of dedicated instructions to read/write state in CPL3. They provide an OS controlled mechanism to enable high performance access to state in CPL3 as soon as the control and/or status register support is added to the processor. In some examples, the instructions are only supported in 64-bit mode.

[0007] In some examples, the OS controls exactly what MSR's can be read or written by this instruction by programming a bitmap in memory. This bitmap allows for OS opt-in and per-control and/or status register control of read and write access. The OS could choose to specify different bitmaps for different apps/guests. Allowing the OS the flexibility and power to control which control and/or status registers are readable / writeable by an application results in a cleaner and simpler ISA for the processor.

[0008] In some examples, the number of control and/or status registers that may be written to is less than the number of registers that may be read.

[0009] In some examples, just as with the OS, this is a one-time enabling for a VMM for all user control and/or status registers. Before user control and/or status registers, every dedicated state read/write instruction that was added to allow user access required new enabling and virtualization support. FIG. 1 illustrates an example execution of a read control and/or status register (e.g., URDMSR) instruction. In some examples, this instruction is executable at any privilege level (e.g., user (e.g., level 3) and supervisor/kernel (e.g., level 0)). In general, the execution of a single read control and/or status register instruction causes execution circuitry to read the content of a control and/or status register (e.g., a 64-bit model specific register (MSR)) 101 specified by an address in a source operand (either a register, memory, or immediate value) value 125 into destination operand 131 (e.g., register or memory) at any current privilege level (CPL). The destination operand is a general-purpose register in some examples.

[0010] When the address of the control and/or status register is stored in a register, in some examples the high-order 32 bits of the register specifying the control and/or status register address are ignored. Note the address may be a physical or linear address depending on the example. Specifying a reserved or unimplemented control and/or status register address causes a general protection exception in some examples. In some examples, if fewer than 64 bits are implemented in the control and/or status register being read, the values returned in unimplemented bit locations are undefined.

[0011] In some examples, most, if not all, control and/or status registers are potentially readable by URDMSR, but under an operating system (OS) control. The OS controls exactly what control and/or status registers can be read by programming a bitmap (e.g., user-level access control bitmap 123) or other data structure in memory. In particular, the OS programs a user-level access control and/or status register 121 (e.g., called IA32_USER_MSR_CTL[63:12] in some examples) to contain an aligned (linear) address of the bitmap. This control and/or status register also has an enable (e.g., bit 0), which the OS sets to enable URDMSR and/or UWRMSR instructions for that thread.

[0012] Access control logic 119 uses the user-level access control and/or status register data 115 to index information about a particular control and/or status register in the user-level access control bitmap data 117 based on a control and/or status register's address information provided by the control and/or status register address source 125. When the user-level access control bitmap data 117 indicates for that control and/or status register's address that access is allowed, then the register read and write circuitry 115 will read data out of the addressed control and/or status register (control and/or status register source 101) and store it in the general purpose register destination 131. While the access control logic 119 is shown as being a part of execution circuitry 111, in some examples, the access control logic 119 is a part of a decoder (for example, when an immediate value is used they access control aspects may happen before the execution circuitry 111 gets involved).

[0013] FIG. 2 illustrates examples of descriptions of the URDMSR including an opcode, operands to use, a CPUID location to indicate support for the instruction, etc.

[0014] An example of a format for a read control and/or status register instruction is URDMSR DST, SRC. In some examples, URDMSR is the opcode mnemonic of the instruction to read a control and/or status register and store the data from the control and/or status register into a destination register. DST is one or more fields for a location of destination operand such as a general-purpose register. SRC one or more fields for a location of a source operand such as general purpose register (that is to store a control and/or status register address). In some examples, the opcode is provided by field 2403, BPG12, or 3404. In some examples, aspects of the source and/or destination locations are provided by one or more of bits from a prefix 2401 (e.g., R-bit, VVVV, etc.) and addressing information 2405 (e.g., reg 2544 ("rrr" and, in some examples "R" from the prefix), R/M 2446 ("bbb" and, in some examples, "B" from the prefix), SIB byte 2504, etc.), 2018, 2020, 2022, 2024, 2026, etc. In some examples, the "rrr" information is used, at least in part, to identify the SRC and the "bbb" information is used, at least in part, to identify the DST. In some examples, the "rrr" information is used, at least in part, to identify the DST and the "bbb" information is used, at least in part, to identify the SRC.

[0015] An example of a format for a read control and/or status register instruction is URDMSR DST, IMM. In some examples, URDMSR is the opcode mnemonic of the instruction to read a control and/or status register and store the data from the control and/or status register into a destination register. DST is one or more fields for a location of destination operand such as a general-purpose register. IMM is a field (e.g., a 32-bit field) that is to encode a control and/or status register address. In some examples, the opcode is provided by field 2403, BPG12, or 3404. In some examples, aspects of the source and/or destination locations are provided by one or more of bits from a prefix 2401 (e.g., R-bit, VVVV, etc.) and addressing information 2405 (e.g., reg 2544 ("rrr" and, in some examples "R" from the prefix), R/M 2446 ("bbb" and, in some examples, "B" from the prefix), SIB byte 2504, etc.), 2018, 2020, 2022, 2024, 2026, etc., and 2409 for the immediate value. In some examples, the "rrr" information is used, at least in part, to identify the DST and in some examples the "bbb" information is used, at least in part, to identify the DST.

[0016] FIG. 3 illustrates an example execution of a write control and/or status register (e.g., UWRMSR) instruction. In some examples, this instruction is executable at any privilege level (e.g., user (e.g., level 3) and supervisor/kernel (e.g., level 0)). In general, the execution of a single write control and/or status register instruction causes execution circuitry to write content of a general purpose register (e.g., GPR 331) or memory into a control and/or status register (e.g., a 64-bit model specific register (MSR)) 301 specified by an address in a destination operand (either a register or immediate value) 325 at any current privilege level (CPL).

[0017] When the address of the control and/or status register is stored in a register, in some examples the high-order 32 bits of the register specifying the control and/or status register address are ignored. Note the address may be a physical or linear address depending on the example. Specifying a reserved or unimplemented control and/or status register address causes a general protection exception in some examples. In some examples, if fewer than 64 bits are implemented in the control and/or status register being read, the values returned in unimplemented bit locations are undefined.

[0018] In some examples, most, if not all, control and/or status registers are potentially readable by UWRMSR, but under an operating system (OS) control. The OS controls exactly what control and/or status registers can be read by programming a bitmap (e.g., user-level access control bitmap 323) or other data structure in memory. In particular, the OS programs a user-level access control and/or status register 321 (e.g., called IA32_USER_MSR_CTL[63:12] in some examples) to contain an aligned (linear) address of the bitmap. This control and/or status register also has an enable (e.g., bit 0), which the OS sets to enable URDMSR and/or UWRMSR instructions for that thread.

[0019] Access control logic 319 uses the user-level access control and/or status register data 315 to index information about a particular control and/or status register in the user-level access control bitmap data 317 based on a control and/or status register's address information provided by the control and/or status register address source 325. When the user-level access control bitmap data 317 indicates for that control and/or status register's address that access is allowed, then the register read and write circuitry 315 will not write data to the addressed control and/or status register (control and/or status register source 301). While the access control logic 319 is shown as being a part of execution circuitry 311, in some examples, the access control logic 319 is a part of a decoder (for example, when an immediate value is used they access control aspects may happen before the execution circuitry 311 gets involved).

[0020] An example of a format for a write control and/or status register instruction is UWRMSR DST, SRC. In some examples, UWRMSR is the opcode mnemonic of the instruction to write a control and/or status register with data from a general purpose register. SRCE is one or more fields for a location of source operand such as a general-purpose register. DST one or more fields for a location of a destination operand such as general purpose register (that is to store a control and/or status register address). In some examples, the opcode is provided by field 2403, 3212, or 3404. In some examples, aspects of the source and/or destination locations are provided by one or more of bits from a prefix 2401 (e.g., R-bit, VVVV, etc.) and addressing information 2405 (e.g., reg 2544 ("rrr" and, in some examples "R" from the prefix), R/M 2446 ("bbb" and, in some examples, "B" from the prefix), SIB byte 2504, etc.), 2018, 2020, 2022, 2024, 2026, etc. In some examples, the "rrr" information is used, at least in part, to identify the SRC and the "bbb" information

is used, at least in part, to identify the DST. In some examples, the "rrr" information is used, at least in part, to identify the DST and the "bbb" information is used, at least in part, to identify the SRC. These references are to subsequent figures.

[0021] An example of a format for a write control and/or status register instruction is UWRMSR DST, IMM. In some examples, UWRMSR is the opcode mnemonic of the instruction to write a control and/or status register with data from a general purpose register. DST one or more fields for a location of a destination operand such as general purpose register (that is to store a control and/or status register address). IMM is a field (e.g., a 32-bit field) that is to encode a control and/or status register address. In some examples, the opcode is provided by field 2403, BPG12, or 3404. In some examples, the opcode is provided by field 2403, 3212, or 3404. In some examples, aspects of the source and/or destination locations are provided by one or more of bits from a prefix 2401 (e.g., R-bit, VVVV, etc.) and addressing information 2405 (e.g., reg 2544 ("rrr" and, in some examples "R" from the prefix), R/M 2446 ("bbb" and, in some examples, "B" from the prefix), SIB byte 2504, etc.), 2018, 2020, 2022, 2024, 2026, etc. and 2409 for the immediate value. In some examples, the "rrr" information is used, at least in part, to identify the SRC and in some examples the "bbb" information is used, at least in part, to identify the SRC. These references are to subsequent figures.

[0022] FIG. 4 illustrates examples of descriptions of the UWRMSR including an opcode, operands to use, a CPUID location to indicate support for the instruction, etc.

[0023] FIG. 5 illustrates examples of a user-level access control bitmap to determine when a read and/or write of a control and/or status register is to be allowed. In this example, the bitmap is divided into two regions (e.g., two 2KB regions). Note that separate bitmaps or a single bit map may be used. The lower half of the bitmap is for URDMSR and the upper half for UWRMSR in some examples. Each region supports control and/or status registers with 14-bit addresses. These read and write bitmaps give the OS precise control over which MSRs can be read and/or written by an application and which cannot.

[0024] The bitmaps contain one bit for each control and/or status register address for a certain range (e.g., in the range 0H - 3FFFH). In some examples, a bitmap value of 0 indicates that access to the control and/or status register is to be prevented and a value of 1 indicates to allow access. Of course the opposite convention may be used. A request to read a register address outside of that range (e.g., beyond 3FFF) will cause an exception (e.g., a general protection fault).

[0025] Setting a bit allows URDMSR to the corresponding MSR (if the CPU supports the URDMSR). This bitmap thus allows for both OS opt-in and per-MSR control of read access. The OS could choose to specify different bitmaps for different apps by context switching the user-level access control and/or status register 121.

[0026] Accesses to these bitmaps for URDMSR and/or UWRMSR are implicit supervisor-mode accesses. Which means they use supervisor privilege regardless of CPL. The OS can create an alias to the bitmap in user address space if it wants the application to know which MSRs are permitted; but the alias should be mapped read-only to prevent the application from overwriting the bitmap.

[0027] In some examples, URDMSR and UWRMSR should not speculatively allow reads or writes to any state that is not allowed to be accessed. In some examples, virtual machine monitors (VMMs) expose the enumeration of these instructions to guests, allowing for guest access to the user-level access control and/or status register 121 and switching that MSR on context switches.

[0028] As shown in this figure, the user-level access control and/or status register 121 provides an indication of enablement for the URDMSR and UWRMSR instructions and includes a page aligned linear address of the user-level access control bitmap (in other words a pointer to this bitmap). In this example, the user-level access control bitmap data 117 is shown. Note that this also shows in some examples that control and/or status registers URDMSR and UWRMSR instructions index the same. Additionally, the indexes are the same for user and supervisor modes.

[0029] When control and/or status register address (from a control and/or status register address source 125) for a URDMSR or UWRMSR instruction is provided it is used to both index into the bitmap (using a byte offset) and, for a byte-sized entry of the bit map, to address a particular bit using a bit offset. Each of the bits of a byte-size entry correspond to a control and/or status register. When a particular bit is 1, then access is allowed. And when a particular bit is 0, then access is denied. Note that the opposite convention, different sized entries, offsets, etc. could be used.

[0030] FIG. 6 illustrates examples of a method to process a URDMSR instruction. In some examples, some aspects of the method includes using emulation or binary translation. In some examples, the aspects of the method are performed without emulation or binary translation. For example, a processor core as shown in FIG. 21(B), a pipeline and/or emulation/translation layer may perform aspects of this method.

[0031] At 601, an instance of single instruction is fetched. For example, an URDMSR instruction is fetched. The single instruction of a first instruction set architecture has fields for an opcode, an indication of a source register operand that is to store a control and/or status register address, and an indication of a destination register operand, wherein the opcode is to indicate that execution circuitry is to read data from the control and/or status register whose identity is stored in the source operand and write the data in the indicated destination register operand regardless of a current privilege level when access to the control and/or status register is allowed.

[0032] The fetched single instruction of the first instruction set architecture is translated into one or more instructions

of a second instruction set architecture at 602. This translation is performed by a translation and/or emulation layer of software in some examples. In some examples, this translation is performed by an instruction converter 2512 as shown in FIG. 25. In some examples, the translation is performed by hardware translation circuitry.

[0033] The one or more translated instructions of the second instruction set architecture are, or the instance of the instance of the single instruction of the first instruction set architecture is, decoded at 603. For example, the instruction(s) is/are decoded by decoder circuitry such as decoder circuitry 1305 or decode circuitry 2140 detailed herein. In some examples, the operations of translation and decoding at 602 and 603 are merged.

[0034] Data values associated with the source operand(s) of the decoded one or more instructions of the second instruction set architecture are retrieved and the one or more instructions are scheduled at 605. For example, when one or more of the source operands are memory operands, the data from the indicated memory location is retrieved.

[0035] At 607, the decoded instruction(s) of the second instruction set architecture is/are executed by execution circuitry (hardware) such as execution circuitry 111 shown in FIG. 1, execution circuitry 1309 shown in FIG. 13 (detailed below), or execution cluster(s) 2160 shown in FIG. 21(B), to perform the operation(s) indicated by the opcode of the single instruction of the first instruction set architecture. For the URDMSR instruction, the execution will cause execution circuitry to at least perform the operations described in connection with FIG. 1.

[0036] In some examples, the instruction is committed or retired at 609.

[0037] FIG. 7 illustrates examples of a method to process a URDMSR instruction. In some examples, some aspects of the method includes using emulation or binary translation. In some examples, the aspects of the method are performed without emulation or binary translation. For example, a processor core as shown in FIG. 21(B), a pipeline and/or emulation/translation layer may perform aspects of this method.

[0038] At 701, an instance of single instruction is fetched. For example, an URDMSR instruction is fetched. The instance of the single instruction of a first instruction set architecture has fields for an opcode, an immediate value to encode an indication of a source register operand that is to store a control and/or status register address, and an indication of a destination register operand, wherein the opcode is to indicate that execution circuitry is to read data from the control and/or status register whose identity is encoded in the immediate value and write the data in the indicated destination register operand regardless of a current privilege level when an access to control and/or status register is allowed.

[0039] The fetched single instruction of the first instruction set architecture is translated into one or more instructions of a second instruction set architecture at 702. This translation is performed by a translation and/or emulation layer of software in some examples. In some examples, this translation is performed by an instruction converter 2512 as shown in FIG. 25. In some examples, the translation is performed by hardware translation circuitry.

[0040] The one or more translated instructions of the second instruction set architecture are, or the instance of the instance of the single instruction of the first instruction set architecture is, decoded at 703. For example, the instruction(s) is/are decoded by decoder circuitry such as decoder circuitry 1305 or decode circuitry 2140 detailed herein. In some examples, the operations of translation and decoding at 702 and 703 are merged.

[0041] Data values associated with the source operand(s) of the decoded one or more instructions of the second instruction set architecture are retrieved and the one or more instructions are scheduled at 705. For example, when one or more of the source operands are memory operands, the data from the indicated memory location is retrieved.

[0042] At 707, the decoded instruction(s) of the second instruction set architecture is/are executed by execution circuitry (hardware) such as execution circuitry 111 shown in FIG. 1, execution circuitry 1309 shown in FIG. 13 (detailed below), or execution cluster(s) 2160 shown in FIG. 21(B), to perform the operation(s) indicated by the opcode of the single instruction of the first instruction set architecture. For the URDMSR instruction, the execution will cause execution circuitry to at least perform the operations described in connection with FIG. 1.

[0043] In some examples, the instruction is committed or retired at 709.

[0044] FIG. 8 illustrates examples of a method to process a UWRMSR instruction. In some examples, some aspects of the method includes using emulation or binary translation. In some examples, the aspects of the method are performed without emulation or binary translation. For example, a processor core as shown in FIG. 21(B), a pipeline and/or emulation/translation layer may perform aspects of this method.

[0045] At 801, an instance of single instruction is fetched. For example, an UWRMSR instruction is fetched. The instance the single instruction of a first instruction set architecture having fields for an opcode, an indication of a destination register operand that is to store a control and/or status register address, and an indication of a source register operand, wherein the opcode is to indicate that execution circuitry is to read data from the source register operand and store the data to the control and/or status register whose identity is stored in the destination operand regardless of a current privilege level when access to the control and/or status register is allowed.

[0046] The fetched single instruction of the first instruction set architecture is translated into one or more instructions of a second instruction set architecture at 802. This translation is performed by a translation and/or emulation layer of software in some examples. In some examples, this translation is performed by an instruction converter 2512 as shown in FIG. 25. In some examples, the translation is performed by hardware translation circuitry.

[0047] The one or more translated instructions of the second instruction set architecture are, or the instance of the

instance of the single instruction of the first instruction set architecture is, decoded at 803. For example, the instruction(s) are decoded by decoder circuitry such as decoder circuitry 1305 or decode circuitry 2140 detailed herein. In some examples, the operations of translation and decoding at 802 and 803 are merged.

[0048] Data values associated with the source operand(s) of the decoded one or more instructions of the second instruction set architecture are retrieved and the one or more instructions are scheduled at 805. For example, when one or more of the source operands are memory operands, the data from the indicated memory location is retrieved.

[0049] At 807, the decoded instruction(s) of the second instruction set architecture is/are executed by execution circuitry (hardware) such as execution circuitry 311 shown in FIG. 3, execution circuitry 1309 shown in FIG. 13 (detailed below), or execution cluster(s) 2160 shown in FIG. 21(B), to perform the operation(s) indicated by the opcode of the single instruction of the first instruction set architecture. For the UWRMSR instruction, the execution will cause execution circuitry to at least perform the operations described in connection with FIG. 3.

[0050] In some examples, the instruction is committed or retired at 809.

[0051] FIG. 9 illustrates examples of a method to process a UWRMSR instruction. In some examples, some aspects of the method includes using emulation or binary translation. In some examples, the aspects of the method are performed without emulation or binary translation. For example, a processor core as shown in FIG. 21(B), a pipeline and/or emulation/translation layer may perform aspects of this method.

[0052] At 901, an instance of single instruction is fetched. For example, an UWRMSR instruction is fetched. The instance the single instruction of a first instruction set architecture having fields for an opcode, an immediate value to encode an indication of a destination control and/or status register operand, and an indication of a source register operand, wherein the opcode is to indicate that execution circuitry is to read data from the source register and store the data to the destination control and/or status register whose identity is encoded in the immediate value regardless of a current privilege level when access to the control and/or status register is allowed.

[0053] The fetched single instruction of the first instruction set architecture is translated into one or more instructions of a second instruction set architecture at 902. This translation is performed by a translation and/or emulation layer of software in some examples. In some examples, this translation is performed by an instruction converter 2512 as shown in FIG. 25. In some examples, the translation is performed by hardware translation circuitry.

[0054] The one or more translated instructions of the second instruction set architecture are, or the instance of the single instruction of the first instruction set architecture is, decoded at 903. For example, the instruction(s) are decoded by decoder circuitry such as decoder circuitry 1305 or decode circuitry 2140 detailed herein. In some examples, the operations of translation and decoding at 902 and 903 are merged.

[0055] Data values associated with the source operand(s) of the decoded one or more instructions of the second instruction set architecture are retrieved and the one or more instructions are scheduled at 905. For example, when one or more of the source operands are memory operands, the data from the indicated memory location is retrieved.

[0056] At 907, the decoded instruction(s) of the second instruction set architecture is/are executed by execution circuitry (hardware) such as execution circuitry 311 shown in FIG. 3, execution circuitry 1309 shown in FIG. 13 (detailed below), or execution cluster(s) 2160 shown in FIG. 21(B), to perform the operation(s) indicated by the opcode of the single instruction of the first instruction set architecture. For the UWRMSR instruction, the execution will cause execution circuitry to at least perform the operations described in connection with FIG. 3.

[0057] In some examples, the instruction is committed or retired at 909.

[0058] In some examples, execution of the user MSR instructions checks a virtual machine extension MSR bitmap. In some examples, no new virtual machines controls required as legacy VMMs would not allow the guest to enable these instructions. VMMs simply need to expose the enumeration to guests, allow the guest access to IA32_USER_MSR_CTL (User MSR bitmap pointer) and switch that MSR on context switches.

[0059] In some prior architectures the immediate value was an 8-bit value. As noted in the discussion above, in some examples, the immediate value used with an UWRMSR or URDMSR instruction will require an immediate value that is larger than 8 bits. In some examples, the immediate value for these instructions is a 32-bit value.

[0060] Previously, to support 32-bit values, memory or a register had to be used which required that requires time to access. Allows for instructions w/ 32-bit immediate values that do not need to be loaded into a register prior to execution. The immediate value can be used at issue time to look up a PLA in the frontend of the machine, greatly improving performance. Frees register for other use.

[0061] FIG. 10 illustrates examples of an instruction format that support the UWRMSR or URDMSR instruction when a 32-bit immediate value is not used. As illustrated, an instruction may include multiple components including, but not limited to, one or more prefixes 1001, an opcode 1003, addressing information 1005 (e.g., register identifiers, memory addressing information, etc.), a displacement value 1007, and/or an immediate value 1009. Note that some instructions utilize some or all the fields of the format whereas others may only use the field for the opcode 1003. In some examples, the order illustrated is the order in which these fields are to be encoded, however, it should be appreciated that in other examples these fields may be encoded in a different order, combined, etc.

[0062] The prefix(es) field(s) 1001, when used, modifies an instruction. In some examples, one or more prefixes are

EP 4 439 285 A1

used to repeat string instructions (e.g., 0xF0, 0xF2, 0xF3, etc.), to provide section overrides (e.g., 0x2E, 0x36, 0x3E, 0x26, 0x64, 0x65, 0x2E, 0x3E, etc.), to perform bus lock operations, and/or to change operand (e.g., 0x66) and address sizes (e.g., 0x67). Certain instructions require a mandatory prefix (e.g., 0x66, 0xF2, 0xF3, etc.). Certain of these prefixes may be considered "legacy" prefixes.

5 **[0063]** In some examples, a prefix used to support the UWRMSR or URDMSR instructions is as shown. In some examples, the prefix is an example of a VEX prefix. The prefix, in general, allows instructions to be encoded to have more than two operands, and allows SIMD vector registers (e.g., vector/SIMD registers 2310) to be longer than 64-bits (e.g., 128-bit and 256-bit).

10 **[0064]** In some examples, a format field 1011 (byte 0 1013) contains the value C4H. Byte 1 1015 includes in bits[7:5] "R," "X," and "B" which are the complements of the same values of first prefix 2401(A) that is detailed later. Bits[4:0] of byte 1 1015 (shown as mmmmm) include content to encode, as needed, one or more implied leading opcode bytes. For example, 00001 implies a 0FH leading opcode, 00010 implies a 0F38H leading opcode, 00011 implies a 0F3AH leading opcode, etc. In some examples, 00111 is encoded by m-mmmm for instructions that utilize a 32-bit immediate value.

15 **[0065]** Bit[7] of byte 2 1017 is used similar to W of the first prefix 2401(A) including helping to determine promotable operand sizes. Bit[2] is used to dictate the length (L) of the vector (where a value of 0 is a scalar or 128-bit vector and a value of 1 is a 256-bit vector). Bits[1:0] provide opcode extensionality equivalent to some legacy prefixes (e.g., 00 = no prefix, 01 = 66H, 10 = F3H, and 11 = F2H). Bits[6:3], shown as vvvv, may be used to: 1) encode the first source register operand, specified in inverted (1s complement) form and valid for instructions with 2 or more source operands; 2) encode the destination register operand, specified in 1s complement form for certain vector shifts; or 3) not encode
20 any operand, the field is reserved and should contain a certain value, such as 1111b.

[0066] The opcode field 1003 is used to at least partially define the operation to be performed upon a decoding of the instruction. In some examples, a primary opcode encoded in the opcode field 1003 is one, two, or three bytes in length. In other examples, a primary opcode can be a different length. An additional 3-bit opcode field is sometimes encoded in another field.

25 **[0067]** The addressing information field 1005 is used to address one or more operands of the instruction, such as a location in memory or one or more registers. Instructions that use the illustrated prefix may use a MOD R/M field 1002. This field may include a MOD R/M R/M field 1046 to encode an instruction operand that references a memory address or encode either the destination register operand or a source register operand. Instructions that use this prefix may use the MOD R/M reg field 1044 to encode either the destination register operand or a source register operand, or to be
30 treated as an opcode extension and not used to encode any instruction operand. Note that bits from the prefix (namely B and R from byte 1 1015) may also be needed to encode a register identifier. For a more discussion of other aspects of the addressing information 1005 please see later descriptions.

[0068] Some addressing forms utilize a displacement value to generate a memory address. For example, a memory address may be generated according to $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$, $\text{index} * \text{scale} + \text{displacement}$, $r/m + \text{displacement}$, instruction pointer (RIP/EIP) + displacement, register + displacement, etc. The displacement may be a 1-byte, 2-byte, 4-byte, etc. value. In some examples, the displacement field 2407 provides this value. Additionally, in some examples, a displacement factor usage is encoded in the MOD field of the addressing information field 2405 that indicates a compressed displacement scheme for which a displacement value is calculated and stored in the displacement field
35 2407.

40 **[0069]** In some examples, the immediate value field 2409 specifies an immediate value for the instruction. Typically the immediate value is encoded as a 1-byte value.

[0070] As noted above, in some examples, when m-mmmm of the prefix has certain value it is an indication to a decoder to view everything after the MOD R/M field 1002 to be a part of a 32-bit immediate value. FIG. 11 illustrates examples of a prefix encoding to indicate a 32-bit immediate value usage. As shown, bits 0-4 of byte 1 1015 of the prefix
45 are set to 00111 which will indicate to the decoder the format at the bottom is correct (that is after MOD R/M 1002 is a 32-bit immediate value 1113).

[0071] In some examples, 00111 refers to an opcode map (map 7). Other opcode maps may be used. However, in some examples, the following m-mmmm values have a defined meaning:

50 00000: Reserved for future use (will 4UD)

00001: implied 0F leading opcode byte

00010: implied 0F 38 leading opcode bytes

55 00011: implied 0F 3A leading opcode bytes

[0072] FIG. 12 illustrates examples of a method to process an instruction. In some examples, some aspects of the

method includes using emulation or binary translation. In some examples, the aspects of the method are performed without emulation or binary translation. For example, a processor core as shown in FIG. 21(B), a pipeline and/or emulation/translation layer may perform aspects of this method.

[0073] At 1201, an instance of single instruction is fetched. For example, an UWRMSR instruction is fetched. The instance the single instruction of a first instruction set architecture having fields for a prefix, one or more fields for an opcode, one or more fields for addressing information for one or more operands, and an immediate value field, wherein the opcode is to indicate that execution circuitry is to perform one or more particular operations using the one or more operands, and the prefix is to indicate that the immediate value field is a 32-bit value.

[0074] The fetched single instruction of the first instruction set architecture is translated into one or more instructions of a second instruction set architecture at 1202. This translation is performed by a translation and/or emulation layer of software in some examples. In some examples, this translation is performed by an instruction converter 2512 as shown in FIG. 25. In some examples, the translation is performed by hardware translation circuitry.

[0075] The one or more translated instructions of the second instruction set architecture are, or the instance of the instance of the single instruction of the first instruction set architecture is, decoded at 1203. For example, the instruction(s) is/are decoded by decoder circuitry such as decoder circuitry 1305 or decode circuitry 2140 detailed herein. In some examples, the operations of translation and decoding at 1202 and 1203 are merged. In some examples, because of settings within the prefix the 32 bits after at least a proper subset of addressing information for the one or more operands is to be treated as an immediate value.

[0076] In some examples, the decoder circuitry includes a PLA. The immediate value is used to index the PLA to determine one or more operation to perform (e.g., micro-operations by execution circuitry or other circuitry, access control, etc.). For example, the immediate value may be handled such that a determination is made of if a control and/or status register can be accessed before execution circuitry starts the process of reading and/or writing to control and/or status register. As such, the immediate value is handled before reaching the execution circuitry. In some examples, the one or more operations are performed outside of the execution circuitry. In some examples, the one or more operations are performed by the execution circuitry. In some examples, the one or more operations are performed by a combination of circuitry inside and outside of the execution circuitry.

[0077] Data values associated with the source operand(s) of the decoded one or more instructions of the second instruction set architecture are retrieved and the one or more instructions are scheduled at 1205. For example, when one or more of the source operands are memory operands, the data from the indicated memory location is retrieved.

[0078] At 1207, the decoded instruction(s) of the second instruction set architecture is/are executed by execution circuitry (hardware) such as execution circuitry 311 shown in FIG. 3, execution circuitry 1309 shown in FIG. 13 (detailed below), or execution cluster(s) 2160 shown in FIG. 21(B), to perform the operation(s) indicated by the opcode of the single instruction of the first instruction set architecture. For the UWRMSR instruction, the execution will cause execution circuitry to at least perform the operations described in connection with FIG. 3.

[0079] In some examples, the instruction is committed or retired at 1209.

[0080] FIG. 13 illustrates examples of computing hardware to process at least a UWRMSR and/or URDMSR instruction. As illustrated, storage 1303 stores a UWRMSR and/or URDMSR instruction 1301 to be executed. One or more other instructions 1304 may also be stored.

[0081] One of the instructions is received by decoder circuitry 1305. For example, the decoder circuitry 1305 receives this instruction from fetch circuitry (not shown). The instruction may be in any suitable format, such as that described with reference to FIG. 24 below. In some examples, the decoder circuitry 1305 supports the use of a 32-bit immediate value as indicated by a prefix. In some examples, the decoder circuitry 1305 includes a programmable logic array (PLA) that is indexable by an immediate value. The PLA allows for some operations to be performed before the execution circuitry. For example, different microops may be selected to handle the operation(s) indicated by the opcode (such as accessing a control and/or status register, etc.) In these examples, the immediate value is not consumed by the execution unit.

[0082] More detailed examples of at least one instruction format for the instruction will be detailed later. The decoder circuitry 1305 decodes the instruction into one or more operations. In some examples, this decoding includes generating a plurality of micro-operations to be performed by execution circuitry (such as execution circuitry 1309). The decoder circuitry 1305 also decodes instruction prefixes.

[0083] In some examples, register renaming, register allocation, and/or scheduling circuitry 1307 provides functionality for one or more of: 1) renaming logical operand values to physical operand values (e.g., a register alias table in some examples), 2) allocating status bits and flags to the decoded instruction, and 3) scheduling the decoded instruction for execution by execution circuitry out of an instruction pool (e.g., using a reservation station in some examples).

[0084] Registers (register file) and/or memory 1308 store data as operands of the instruction to be operated by execution circuitry 1309. Example register types include packed data registers, general purpose registers (GPRs), and floating-point registers.

[0085] Execution circuitry 1309 executes the decoded instruction. Example detailed execution circuitry includes exe-

cution circuitry 109 shown in FIG. 1, and execution cluster(s) 2160 shown in FIG. 21(B), etc. The execution of the decoded instruction causes the execution circuitry to perform the operation(s) indicated by the opcode.

[0086] In some examples, retirement/write back circuitry 1311 architecturally commits the destination register into the registers or memory 1308 and retires the instruction.

[0087] Some examples utilize instruction formats described herein. Some examples are implemented in one or more computer architectures, cores, accelerators, etc. Some examples are generated or are IP cores. Some examples utilize emulation and/or translation.

Example Architectures

[0088] Detailed below are descriptions of example computer architectures. Other system designs and configurations known in the arts for laptop, desktop, and handheld personal computers (PC)s, personal digital assistants, engineering workstations, servers, disaggregated servers, network devices, network hubs, switches, routers, 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, are also suitable. In general, a variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable.

Example Systems

[0089] FIG. 14 illustrates an example computing system. Multiprocessor system 1400 is an interfaced system and includes a plurality of processors or cores including a first processor 1470 and a second processor 1480 coupled via an interface 1450 such as a point-to-point (P-P) interconnect, a fabric, and/or bus. In some examples, the first processor 1470 and the second processor 1480 are homogeneous. In some examples, first processor 1470 and the second processor 1480 are heterogenous. Though the example system 1400 is shown to have two processors, the system may have three or more processors, or may be a single processor system. In some examples, the computing system is a system on a chip (SoC).

[0090] Processors 1470 and 1480 are shown including integrated memory controller (IMC) circuitry 1472 and 1482, respectively. Processor 1470 also includes interface circuits 1476 and 1478; similarly, second processor 1480 includes interface circuits 1486 and 1488. Processors 1470, 1480 may exchange information via the interface 1450 using interface circuits 1478, 1488. IMCs 1472 and 1482 couple the processors 1470, 1480 to respective memories, namely a memory 1432 and a memory 1434, which may be portions of main memory locally attached to the respective processors.

[0091] Processors 1470, 1480 may each exchange information with a network interface (NW I/F) 1490 via individual interfaces 1452, 1454 using interface circuits 1476, 1494, 1486, 1498. The network interface 1490 (e.g., one or more of an interconnect, bus, and/or fabric, and in some examples is a chipset) may optionally exchange information with a coprocessor 1438 via an interface circuit 1492. In some examples, the coprocessor 1438 is a special-purpose processor, such as, for example, a high-throughput processor, a network or communication processor, compression engine, graphics processor, general purpose graphics processing unit (GPGPU), neural-network processing unit (NPU), embedded processor, or the like.

[0092] A shared cache (not shown) may be included in either processor 1470, 1480 or outside of both processors, yet connected with the processors via an interface such as 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.

[0093] Network interface 1490 may be coupled to a first interface 1416 via interface circuit 1496. In some examples, first interface 1416 may be an interface such as a Peripheral Component Interconnect (PCI) interconnect, a PCI Express interconnect or another I/O interconnect. In some examples, first interface 1416 is coupled to a power control unit (PCU) 1417, which may include circuitry, software, and/or firmware to perform power management operations with regard to the processors 1470, 1480 and/or co-processor 1438. PCU 1417 provides control information to a voltage regulator (not shown) to cause the voltage regulator to generate the appropriate regulated voltage. PCU 1417 also provides control information to control the operating voltage generated. In various examples, PCU 1417 may include a variety of power management logic units (circuitry) to perform hardware-based power management. Such power management may be wholly processor controlled (e.g., by various processor hardware, and which may be triggered by workload and/or power, thermal or other processor constraints) and/or the power management may be performed responsive to external sources (such as a platform or power management source or system software).

[0094] PCU 1417 is illustrated as being present as logic separate from the processor 1470 and/or processor 1480. In other cases, PCU 1417 may execute on a given one or more of cores (not shown) of processor 1470 or 1480. In some cases, PCU 1417 may be implemented as a microcontroller (dedicated or general-purpose) or other control logic configured to execute its own dedicated power management code, sometimes referred to as P-code. In yet other examples, power management operations to be performed by PCU 1417 may be implemented externally to a processor, such as

by way of a separate power management integrated circuit (PMIC) or another component external to the processor. In yet other examples, power management operations to be performed by PCU 1417 may be implemented within BIOS or other system software.

[0095] Various I/O devices 1414 may be coupled to first interface 1416, along with a bus bridge 1418 which couples first interface 1416 to a second interface 1420. In some examples, one or more additional processor(s) 1415, such as coprocessors, high throughput many integrated core (MIC) processors, GPGPUs, accelerators (such as graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays (FPGAs), or any other processor, are coupled to first interface 1416. In some examples, second interface 1420 may be a low pin count (LPC) interface. Various devices may be coupled to second interface 1420 including, for example, a keyboard and/or mouse 1422, communication devices 1427 and storage circuitry 1428. Storage circuitry 1428 may be one or more non-transitory machine-readable storage media as described below, such as a disk drive or other mass storage device which may include instructions/code and data 1430 and may implement the storage 1303 in some examples. Further, an audio I/O 1424 may be coupled to second interface 1420. Note that other architectures than the point-to-point architecture described above are possible. For example, instead of the point-to-point architecture, a system such as multiprocessor system 1400 may implement a multi-drop interface or other such architecture.

Example Core Architectures, Processors, and Computer Architectures.

[0096] Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high-performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput) computing. Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip (SoC) that may be included on the same die as the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and additional functionality. Example core architectures are described next, followed by descriptions of example processors and computer architectures.

[0097] FIG. 15 illustrates a block diagram of an example processor and/or SoC 1500 that may have one or more cores and an integrated memory controller. The solid lined boxes illustrate a processor 1500 with a single core 1502(A), system agent unit circuitry 1510, and a set of one or more interface controller unit(s) circuitry 1516, while the optional addition of the dashed lined boxes illustrates an alternative processor 1500 with multiple cores 1502(A)-(N), a set of one or more integrated memory controller unit(s) circuitry 1514 in the system agent unit circuitry 1510, and special purpose logic 1508, as well as a set of one or more interface controller units circuitry 1516. Note that the processor 1500 may be one of the processors 1470 or 1480, or co-processor 1438 or 1415 of FIG. 14.

[0098] Thus, different implementations of the processor 1500 may include: 1) a CPU with the special purpose logic 1508 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores, not shown), and the cores 1502(A)-(N) being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, or a combination of the two); 2) a coprocessor with the cores 1502(A)-(N) being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 1502(A)-(N) being a large number of general purpose in-order cores. Thus, the processor 1500 may be a general-purpose processor, coprocessor or special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, GPGPU (general purpose graphics processing unit), a high throughput many integrated core (MIC) coprocessor (including 30 or more cores), embedded processor, or the like. The processor may be implemented on one or more chips. The processor 1500 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, complementary metal oxide semiconductor (CMOS), bipolar CMOS (BiCMOS), P-type metal oxide semiconductor (PMOS), or N-type metal oxide semiconductor (NMOS).

[0099] A memory hierarchy includes one or more levels of cache unit(s) circuitry 1504(A)-(N) within the cores 1502(A)-(N), a set of one or more shared cache unit(s) circuitry 1506, and external memory (not shown) coupled to the set of integrated memory controller unit(s) circuitry 1514. The set of one or more shared cache unit(s) circuitry 1506 may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, such as a last level cache (LLC), and/or combinations thereof. While in some examples interface network circuitry 1512 (e.g., a ring interconnect) interfaces the special purpose logic 1508 (e.g., integrated graphics logic), the set of shared cache

unit(s) circuitry 1506, and the system agent unit circuitry 1510, alternative examples use any number of well-known techniques for interfacing such units. In some examples, coherency is maintained between one or more of the shared cache unit(s) circuitry 1506 and cores 1502(A)-(N). In some examples, interface controller units circuitry 1516 couple the cores 1502 to one or more other devices 1518 such as one or more I/O devices, storage, one or more communication devices (e.g., wireless networking, wired networking, etc.), etc.

[0100] In some examples, one or more of the cores 1502(A)-(N) are capable of multi-threading. The system agent unit circuitry 1510 includes those components coordinating and operating cores 1502(A)-(N). The system agent unit circuitry 1510 may include, for example, power control unit (PCU) circuitry and/or display unit circuitry (not shown). The PCU may be or may include logic and components needed for regulating the power state of the cores 1502(A)-(N) and/or the special purpose logic 1508 (e.g., integrated graphics logic). The display unit circuitry is for driving one or more externally connected displays.

[0101] The cores 1502(A)-(N) may be homogenous in terms of instruction set architecture (ISA). Alternatively, the cores 1502(A)-(N) may be heterogeneous in terms of ISA; that is, a subset of the cores 1502(A)-(N) may be capable of executing an ISA, while other cores may be capable of executing only a subset of that ISA or another ISA.

[0102] FIG. 16 is a block diagram illustrating a computing system 1600 configured to implement one or more aspects of the examples described herein. The computing system 1600 includes a processing subsystem 1601 having one or more processor(s) 1602 and a system memory 1604 communicating via an interconnection path that may include a memory hub 1605. The memory hub 1605 may be a separate component within a chipset component or may be integrated within the one or more processor(s) 1602. The memory hub 1605 couples with an I/O subsystem 1611 via a communication link 1606. The I/O subsystem 1611 includes an I/O hub 1607 that can enable the computing system 1600 to receive input from one or more input device(s) 1608. Additionally, the I/O hub 1607 can enable a display controller, which may be included in the one or more processor(s) 1602, to provide outputs to one or more display device(s) 1610A. In some examples the one or more display device(s) 1610A coupled with the I/O hub 1607 can include a local, internal, or embedded display device.

[0103] The processing subsystem 1601, for example, includes one or more parallel processor(s) 1612 coupled to memory hub 1605 via a bus or other communication link 1613. The communication link 1613 may be one of any number of standards-based communication link technologies or protocols, such as, but not limited to PCI Express, or may be a vendor specific communications interface or communications fabric. The one or more parallel processor(s) 1612 may form a computationally focused parallel or vector processing system that can include a large number of processing cores and/or processing clusters, such as a many integrated core (MIC) processor. For example, the one or more parallel processor(s) 1612 form a graphics processing subsystem that can output pixels to one of the one or more display device(s) 1610A coupled via the I/O hub 1607. The one or more parallel processor(s) 1612 can also include a display controller and display interface (not shown) to enable a direct connection to one or more display device(s) 1610B.

[0104] Within the I/O subsystem 1611, a system storage unit 1614 can connect to the I/O hub 1607 to provide a storage mechanism for the computing system 1600. An I/O switch 1616 can be used to provide an interface mechanism to enable connections between the I/O hub 1607 and other components, such as a network adapter 1618 and/or wireless network adapter 1619 that may be integrated into the platform, and various other devices that can be added via one or more add-in device(s) 1620. The add-in device(s) 1620 may also include, for example, one or more external graphics processor devices, graphics cards, and/or compute accelerators. The network adapter 1618 can be an Ethernet adapter or another wired network adapter. The wireless network adapter 1619 can include one or more of a Wi-Fi, Bluetooth, near field communication (NFC), or other network device that includes one or more wireless radios.

[0105] The computing system 1600 can include other components not explicitly shown, including USB or other port connections, optical storage drives, video capture devices, and the like, which may also be connected to the I/O hub 1607. Communication paths interconnecting the various components in FIG. 16 may be implemented using any suitable protocols, such as PCI (Peripheral Component Interconnect) based protocols (e.g., PCI-Express), or any other bus or point-to-point communication interfaces and/or protocol(s), such as the NVLink high-speed interconnect, Compute Express Link™ (CXL™) (e.g., CXL.mem), Infinity Fabric (IF), Ethernet (IEEE 802.3), remote direct memory access (RDMA), InfiniBand, Internet Wide Area RDMA Protocol (iWARP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), quick UDP Internet Connections (QUIC), RDMA over Converged Ethernet (RoCE), Intel QuickPath Interconnect (QPI), Intel Ultra Path Interconnect (UPI), Intel On-Chip System Fabric (IOSF), Omnipath, HyperTransport, Advanced Microcontroller Bus Architecture (AMBA) interconnect, OpenCAPI, Gen-Z, Cache Coherent Interconnect for Accelerators (CCIX), 3GPP Long Term Evolution (LTE) (4G), 3GPP 5G, and variations thereof, or wired or wireless interconnect protocols known in the art. In some examples, data can be copied or stored to virtualized storage nodes using a protocol such as non-volatile memory express (NVMe) over Fabrics (NVMe-oF) or NVMe.

[0106] The one or more parallel processor(s) 1612 may incorporate circuitry optimized for graphics and video processing, including, for example, video output circuitry, and constitutes a graphics processing unit (GPU). Alternatively or additionally, the one or more parallel processor(s) 1612 can incorporate circuitry optimized for general purpose processing, while preserving the underlying computational architecture, described in greater detail herein. Components of the

computing system 1600 may be integrated with one or more other system elements on a single integrated circuit. For example, the one or more parallel processor(s) 1612, memory hub 1605, processor(s) 1602, and I/O hub 1607 can be integrated into a system on chip (SoC) integrated circuit. Alternatively, the components of the computing system 1600 can be integrated into a single package to form a system in package (SIP) configuration. In some examples at least a portion of the components of the computing system 1600 can be integrated into a multi-chip module (MCM), which can be interconnected with other multi-chip modules into a modular computing system.

[0107] It will be appreciated that the computing system 1600 shown herein is illustrative and that variations and modifications are possible. The connection topology, including the number and arrangement of bridges, the number of processor(s) 1602, and the number of parallel processor(s) 1612, may be modified as desired. For instance, system memory 1604 can be connected to the processor(s) 1602 directly rather than through a bridge, while other devices communicate with system memory 1604 via the memory hub 1605 and the processor(s) 1602. In other alternative topologies, the parallel processor(s) 1612 are connected to the I/O hub 1607 or directly to one of the one or more processor(s) 1602, rather than to the memory hub 1605. In other examples, the I/O hub 1607 and memory hub 1605 may be integrated into a single chip. It is also possible that two or more sets of processor(s) 1602 are attached via multiple sockets, which can couple with two or more instances of the parallel processor(s) 1612.

[0108] Some of the particular components shown herein are optional and may not be included in all implementations of the computing system 1600. For example, any number of add-in cards or peripherals may be supported, or some components may be eliminated. Furthermore, some architectures may use different terminology for components similar to those illustrated in FIG. 16. For example, the memory hub 1605 may be referred to as a Northbridge in some architectures, while the I/O hub 1607 may be referred to as a Southbridge.

[0109] FIG. 17A illustrates examples of a parallel processor 1700. The parallel processor 1700 may be a GPU, GPGPU or the like as described herein. The various components of the parallel processor 1700 may be implemented using one or more integrated circuit devices, such as programmable processors, application specific integrated circuits (ASICs), or field programmable gate arrays (FPGA). The illustrated parallel processor 1700 may be one or more of the parallel processor(s) 1612 shown in FIG. 16.

[0110] The parallel processor 1700 includes a parallel processing unit 1702. The parallel processing unit includes an I/O unit 1704 that enables communication with other devices, including other instances of the parallel processing unit 1702. The I/O unit 1704 may be directly connected to other devices. For instance, the I/O unit 1704 connects with other devices via the use of a hub or switch interface, such as memory hub 1605. The connections between the memory hub 1605 and the I/O unit 1704 form a communication link 1613. Within the parallel processing unit 1702, the I/O unit 1704 connects with a host interface 1706 and a memory crossbar 1716, where the host interface 1706 receives commands directed to performing processing operations and the memory crossbar 1716 receives commands directed to performing memory operations.

[0111] When the host interface 1706 receives a command buffer via the I/O unit 1704, the host interface 1706 can direct work operations to perform those commands to a front end 1708. In some examples the front end 1708 couples with a scheduler 1710, which is configured to distribute commands or other work items to a processing cluster array 1712. The scheduler 1710 ensures that the processing cluster array 1712 is properly configured and in a valid state before tasks are distributed to the processing clusters of the processing cluster array 1712. The scheduler 1710 may be implemented via firmware logic executing on a microcontroller. The microcontroller implemented scheduler 1710 is configurable to perform complex scheduling and work distribution operations at coarse and fine granularity, enabling rapid preemption and context switching of threads executing on the processing cluster array 1712. Preferably, the host software can prove workloads for scheduling on the processing cluster array 1712 via one of multiple graphics processing doorbells. In other examples, polling for new workloads or interrupts can be used to identify or indicate availability of work to perform. The workloads can then be automatically distributed across the processing cluster array 1712 by the scheduler 1710 logic within the scheduler microcontroller.

[0112] The processing cluster array 1712 can include up to "N" processing clusters (e.g., cluster 1714A, cluster 1714B, through cluster 1714N). Each cluster 1714A-1714N of the processing cluster array 1712 can execute a large number of concurrent threads. The scheduler 1710 can allocate work to the clusters 1714A-1714N of the processing cluster array 1712 using various scheduling and/or work distribution algorithms, which may vary depending on the workload arising for each type of program or computation. The scheduling can be handled dynamically by the scheduler 1710 or can be assisted in part by compiler logic during compilation of program logic configured for execution by the processing cluster array 1712. Optionally, different clusters 1714A-1714N of the processing cluster array 1712 can be allocated for processing different types of programs or for performing different types of computations.

[0113] The processing cluster array 1712 can be configured to perform various types of parallel processing operations. For example, the processing cluster array 1712 is configured to perform general-purpose parallel compute operations. For example, the processing cluster array 1712 can include logic to execute processing tasks including filtering of video and/or audio data, performing modeling operations, including physics operations, and performing data transformations.

[0114] The processing cluster array 1712 is configured to perform parallel graphics processing operations. In such

examples in which the parallel processor 1700 is configured to perform graphics processing operations, the processing cluster array 1712 can include additional logic to support the execution of such graphics processing operations, including, but not limited to texture sampling logic to perform texture operations, as well as tessellation logic and other vertex processing logic. Additionally, the processing cluster array 1712 can be configured to execute graphics processing related shader programs such as, but not limited to vertex shaders, tessellation shaders, geometry shaders, and pixel shaders. The parallel processing unit 1702 can transfer data from system memory via the I/O unit 1704 for processing. During processing the transferred data can be stored to on-chip memory (e.g., parallel processor memory 1722) during processing, then written back to system memory.

[0115] In examples in which the parallel processing unit 1702 is used to perform graphics processing, the scheduler 1710 may be configured to divide the processing workload into approximately equal sized tasks, to better enable distribution of the graphics processing operations to multiple clusters 1714A-1714N of the processing cluster array 1712. In some of these examples, portions of the processing cluster array 1712 can be configured to perform different types of processing. For example, a first portion may be configured to perform vertex shading and topology generation, a second portion may be configured to perform tessellation and geometry shading, and a third portion may be configured to perform pixel shading or other screen space operations, to produce a rendered image for display. Intermediate data produced by one or more of the clusters 1714A-1714N may be stored in buffers to allow the intermediate data to be transmitted between clusters 1714A-1714N for further processing.

[0116] During operation, the processing cluster array 1712 can receive processing tasks to be executed via the scheduler 1710, which receives commands defining processing tasks from front end 1708. For graphics processing operations, processing tasks can include indices of data to be processed, e.g., surface (patch) data, primitive data, vertex data, and/or pixel data, as well as state parameters and commands defining how the data is to be processed (e.g., what program is to be executed). The scheduler 1710 may be configured to fetch the indices corresponding to the tasks or may receive the indices from the front end 1708. The front end 1708 can be configured to ensure the processing cluster array 1712 is configured to a valid state before the workload specified by incoming command buffers (e.g., batch-buffers, push buffers, etc.) is initiated.

[0117] Each of the one or more instances of the parallel processing unit 1702 can couple with parallel processor memory 1722. The parallel processor memory 1722 can be accessed via the memory crossbar 1716, which can receive memory requests from the processing cluster array 1712 as well as the I/O unit 1704. The memory crossbar 1716 can access the parallel processor memory 1722 via a memory interface 1718. The memory interface 1718 can include multiple partition units (e.g., partition unit 1720A, partition unit 1720B, through partition unit 1720N) that can each couple to a portion (e.g., memory unit) of parallel processor memory 1722. The number of partition units 1720A-1720N may be configured to be equal to the number of memory units, such that a first partition unit 1720A has a corresponding first memory unit 1724A, a second partition unit 1720B has a corresponding second memory unit 1724B, and an Nth partition unit 1720N has a corresponding Nth memory unit 1724N. In other examples, the number of partition units 1720A-1720N may not be equal to the number of memory devices.

[0118] The memory units 1724A-1724N can include various types of memory devices, including dynamic random-access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory. Optionally, the memory units 1724A-1724N may also include 3D stacked memory, including but not limited to high bandwidth memory (HBM). Persons skilled in the art will appreciate that the specific implementation of the memory units 1724A-1724N can vary and can be selected from one of various conventional designs. Render targets, such as frame buffers or texture maps may be stored across the memory units 1724A-1724N, allowing partition units 1720A-1720N to write portions of each render target in parallel to efficiently use the available bandwidth of parallel processor memory 1722. In some examples, a local instance of the parallel processor memory 1722 may be excluded in favor of a unified memory design that utilizes system memory in conjunction with local cache memory.

[0119] Optionally, any one of the clusters 1714A-1714N of the processing cluster array 1712 has the ability to process data that will be written to any of the memory units 1724A-1724N within parallel processor memory 1722. The memory crossbar 1716 can be configured to transfer the output of each cluster 1714A-1714N to any partition unit 1720A-1720N or to another cluster 1714A-1714N, which can perform additional processing operations on the output. Each cluster 1714A-1714N can communicate with the memory interface 1718 through the memory crossbar 1716 to read from or write to various external memory devices. In one of the examples with the memory crossbar 1716 the memory crossbar 1716 has a connection to the memory interface 1718 to communicate with the I/O unit 1704, as well as a connection to a local instance of the parallel processor memory 1722, enabling the processing units within the different processing clusters 1714A-1714N to communicate with system memory or other memory that is not local to the parallel processing unit 1702. Generally, the memory crossbar 1716 may, for example, be able to use virtual channels to separate traffic streams between the clusters 1714A-1714N and the partition units 1720A-1720N.

[0120] While a single instance of the parallel processing unit 1702 is illustrated within the parallel processor 1700, any number of instances of the parallel processing unit 1702 can be included. For example, multiple instances of the parallel

processing unit 1702 can be provided on a single add-in card, or multiple add-in cards can be interconnected. For example, the parallel processor 1700 can be an add-in device, such as add-in device 1620 of FIG. 16, which may be a graphics card such as a discrete graphics card that includes one or more GPUs, one or more memory devices, and device-to-device or network or fabric interfaces. The different instances of the parallel processing unit 1702 can be configured to inter-operate even if the different instances have different numbers of processing cores, different amounts of local parallel processor memory, and/or other configuration differences. Optionally, some instances of the parallel processing unit 1702 can include higher precision floating point units relative to other instances. Systems incorporating one or more instances of the parallel processing unit 1702 or the parallel processor 1700 can be implemented in a variety of configurations and form factors, including but not limited to desktop, laptop, or handheld personal computers, servers, workstations, game consoles, and/or embedded systems. An orchestrator can form composite nodes for workload performance using one or more of: disaggregated processor resources, cache resources, memory resources, storage resources, and networking resources.

[0121] In some examples, the parallel processing unit 1702 can be partitioned into multiple instances. Those multiple instances can be configured to execute workloads associated with different clients in an isolated manner, enabling a pre-determined quality of service to be provided for each client. For example, each cluster 1714A-1714N can be compartmentalized and isolated from other clusters, allowing the processing cluster array 1712 to be divided into multiple compute partitions or instances. In such configuration, workloads that execute on an isolated partition are protected from faults or errors associated with a different workload that executes on a different partition. The partition units 1720A-1720N can be configured to enable a dedicated and/or isolated path to memory for the clusters 1714A-1714N associated with the respective compute partitions. This datapath isolation enables the compute resources within a partition can communicate with one or more assigned memory units 1724A-1724N without being subjected to inference by the activities of other partitions.

[0122] FIG. 17B is a block diagram of a partition unit 1720. The partition unit 1720 may be an instance of one of the partition units 1720A-1720N of FIG. 17A. As illustrated, the partition unit 1720 includes an L2 cache 1721, a frame buffer interface 1725, and a ROP 1726 (raster operations unit). The L2 cache 1721 is a read/write cache that is configured to perform load and store operations received from the memory crossbar 1716 and ROP 1726. Read misses and urgent write-back requests are output by L2 cache 1721 to frame buffer interface 1725 for processing. Updates can also be sent to the frame buffer via the frame buffer interface 1725 for processing. In some examples the frame buffer interface 1725 interfaces with one of the memory units in parallel processor memory, such as the memory units 1724A-1724N of FIG. 17A (e.g., within parallel processor memory 1722). The partition unit 1720 may additionally or alternatively also interface with one of the memory units in parallel processor memory via a memory controller (not shown).

[0123] In graphics applications, the ROP 1726 is a processing unit that performs raster operations such as stencil, z test, blending, and the like. The ROP 1726 then outputs processed graphics data that is stored in graphics memory. In some examples the ROP 1726 includes or couples with a CODEC 1727 that includes compression logic to compress depth or color data that is written to memory or the L2 cache 1721 and decompress depth or color data that is read from memory or the L2 cache 1721. The compression logic can be lossless compression logic that makes use of one or more of multiple compression algorithms. The type of compression that is performed by the CODEC 1727 can vary based on the statistical characteristics of the data to be compressed. For example, in some examples, delta color compression is performed on depth and color data on a per-tile basis. In some examples the CODEC 1727 includes compression and decompression logic that can compress and decompress compute data associated with machine learning operations. The CODEC 1727 can, for example, compress sparse matrix data for sparse machine learning operations. The CODEC 1727 can also compress sparse matrix data that is encoded in a sparse matrix format (e.g., coordinate list encoding (COO), compressed sparse row (CSR), compressed sparse column (CSC), etc.) to generate compressed and encoded sparse matrix data. The compressed and encoded sparse matrix data can be decompressed and/or decoded before being processed by processing elements or the processing elements can be configured to consume compressed, encoded, or compressed and encoded data for processing.

[0124] The ROP 1726 may be included within each processing cluster (e.g., cluster 1714A-1714N of FIG. 17A) instead of within the partition unit 1720. In such example, read and write requests for pixel data are transmitted over the memory crossbar 1716 instead of pixel fragment data. The processed graphics data may be displayed on a display device, such as one of the one or more display device(s) 1610A-1610B of FIG. 16, routed for further processing by the processor(s) 1602, or routed for further processing by one of the processing entities within the parallel processor 1700 of FIG. 17A.

[0125] FIG. 17C is a block diagram of a processing cluster 1714 within a parallel processing unit. For example, the processing cluster is an instance of one of the processing clusters 1714A-1714N of FIG. 17A. The processing cluster 1714 can be configured to execute many threads in parallel, where the term "thread" refers to an instance of a particular program executing on a particular set of input data. Optionally, single-instruction, multiple-data (SIMD) instruction issue techniques may be used to support parallel execution of a large number of threads without providing multiple independent instruction units. Alternatively, single-instruction, multiple-thread (SIMT) techniques may be used to support parallel execution of a large number of generally synchronized threads, using a common instruction unit configured to issue

instructions to a set of processing engines within each one of the processing clusters. Unlike a SIMD execution regime, where all processing engines typically execute identical instructions, SIMT execution allows different threads to more readily follow divergent execution paths through a given thread program. Persons skilled in the art will understand that a SIMD processing regime represents a functional subset of a SIMT processing regime.

5 **[0126]** Operation of the processing cluster 1714 can be controlled via a pipeline manager 1732 that distributes processing tasks to SIMT parallel processors. The pipeline manager 1732 receives instructions from the scheduler 1710 of FIG. 17A and manages execution of those instructions via a graphics multiprocessor 1734 and/or a texture unit 1736. The illustrated graphics multiprocessor 1734 is an exemplary instance of a SIMT parallel processor. However, various types of SIMT parallel processors of differing architectures may be included within the processing cluster 1714. One or more instances of the graphics multiprocessor 1734 can be included within a processing cluster 1714. The graphics multiprocessor 1734 can process data and a data crossbar 1740 can be used to distribute the processed data to one of multiple possible destinations, including other shader units. The pipeline manager 1732 can facilitate the distribution of processed data by specifying destinations for processed data to be distributed via the data crossbar 1740.

10 **[0127]** Each graphics multiprocessor 1734 within the processing cluster 1714 can include an identical set of functional execution logic (e.g., arithmetic logic units, load-store units, etc.). The functional execution logic can be configured in a pipelined manner in which new instructions can be issued before previous instructions are complete. The functional execution logic supports a variety of operations including integer and floating-point arithmetic, comparison operations, Boolean operations, bit-shifting, and computation of various algebraic functions. The same functional-unit hardware could be leveraged to perform different operations and any combination of functional units may be present.

15 **[0128]** The instructions transmitted to the processing cluster 1714 constitute a thread. A set of threads executing across the set of parallel processing engines is a thread group. A thread group executes the same program on different input data. Each thread within a thread group can be assigned to a different processing engine within a graphics multiprocessor 1734. A thread group may include fewer threads than the number of processing engines within the graphics multiprocessor 1734. When a thread group includes fewer threads than the number of processing engines, one or more of the processing engines may be idle during cycles in which that thread group is being processed. A thread group may also include more threads than the number of processing engines within the graphics multiprocessor 1734. When the thread group includes more threads than the number of processing engines within the graphics multiprocessor 1734, processing can be performed over consecutive clock cycles. Optionally, multiple thread groups can be executed concurrently on the graphics multiprocessor 1734.

20 **[0129]** The graphics multiprocessor 1734 may include an internal cache memory to perform load and store operations. Optionally, the graphics multiprocessor 1734 can forego an internal cache and use a cache memory (e.g., level 1 (L1) cache 1748) within the processing cluster 1714. Each graphics multiprocessor 1734 also has access to level 2 (L2) caches within the partition units (e.g., partition units 1720A-1720N of FIG. 17A) that are shared among all processing clusters 1714 and may be used to transfer data between threads. The graphics multiprocessor 1734 may also access off-chip global memory, which can include one or more of local parallel processor memory and/or system memory. Any memory external to the parallel processing unit 1702 may be used as global memory. Embodiments in which the processing cluster 1714 includes multiple instances of the graphics multiprocessor 1734 can share common instructions and data, which may be stored in the L1 cache 1748.

25 **[0130]** Each processing cluster 1714 may include an MMU 1745 (memory management unit) that is configured to map virtual addresses into physical addresses. In other examples, one or more instances of the MMU 1745 may reside within the memory interface 1718 of FIG. 17A. The MMU 1745 includes a set of page table entries (PTEs) used to map a virtual address to a physical address of a tile and optionally a cache line index. The MMU 1745 may include address translation lookaside buffers (TLB) or caches that may reside within the graphics multiprocessor 1734 or the L1 cache 1748 of processing cluster 1714. The physical address is processed to distribute surface data access locality to allow efficient request interleaving among partition units. The cache line index may be used to determine whether a request for a cache line is a hit or miss.

30 **[0131]** In graphics and computing applications, a processing cluster 1714 may be configured such that each graphics multiprocessor 1734 is coupled to a texture unit 1736 for performing texture mapping operations, e.g., determining texture sample positions, reading texture data, and filtering the texture data. Texture data is read from an internal texture L1 cache (not shown) or in some examples from the L1 cache within graphics multiprocessor 1734 and is fetched from an L2 cache, local parallel processor memory, or system memory, as needed. Each graphics multiprocessor 1734 outputs processed tasks to the data crossbar 1740 to provide the processed task to another processing cluster 1714 for further processing or to store the processed task in an L2 cache, local parallel processor memory, or system memory via the memory crossbar 1716. A preROP 1742 (pre-raster operations unit) is configured to receive data from graphics multiprocessor 1734, direct data to ROP units, which may be located with partition units as described herein (e.g., partition units 1720A-1720N of FIG. 17A). The preROP 1742 unit can perform optimizations for color blending, organize pixel color data, and perform address translations.

35 **[0132]** It will be appreciated that the core architecture described herein is illustrative and that variations and modifi-

cations are possible. Any number of processing units, e.g., graphics multiprocessor 1734, texture units 1736, preROPs 1742, etc., may be included within a processing cluster 1714. Further, while only one processing cluster 1714 is shown, a parallel processing unit as described herein may include any number of instances of the processing cluster 1714. Optionally, each processing cluster 1714 can be configured to operate independently of other processing clusters 1714

5 using separate and distinct processing units, L1 caches, L2 caches, etc.

[0133] FIG. 17D shows an example of the graphics multiprocessor 1734 in which the graphics multiprocessor 1734 couples with the pipeline manager 1732 of the processing cluster 1714. The graphics multiprocessor 1734 has an execution pipeline including but not limited to an instruction cache 1752, an instruction unit 1754, an address mapping unit 1756, a register file 1758, one or more general purpose graphics processing unit (GPGPU) cores 1762, and one or more load/store units 1766. The GPGPU cores 1762 and load/store units 1766 are coupled with cache memory 1772 and shared memory 1770 via a memory and cache interconnect 1768. The graphics multiprocessor 1734 may additionally

10 include tensor and/or ray-tracing cores 1763 that include hardware logic to accelerate matrix and/or ray-tracing operations. **[0134]** The instruction cache 1752 may receive a stream of instructions to execute from the pipeline manager 1732. The instructions are cached in the instruction cache 1752 and dispatched for execution by the instruction unit 1754. The instruction unit 1754 can dispatch instructions as thread groups (e.g., warps), with each thread of the thread group assigned to a different execution unit within GPGPU core 1762. An instruction can access any of a local, shared, or global address space by specifying an address within a unified address space. The address mapping unit 1756 can be used to translate addresses in the unified address space into a distinct memory address that can be accessed by the load/store units 1766.

20 **[0135]** The register file 1758 provides a set of registers for the functional units of the graphics multiprocessor 1734. The register file 1758 provides temporary storage for operands connected to the data paths of the functional units (e.g., GPGPU cores 1762, load/store units 1766) of the graphics multiprocessor 1734. The register file 1758 may be divided between each of the functional units such that each functional unit is allocated a dedicated portion of the register file 1758. For example, the register file 1758 may be divided between the different warps being executed by the graphics multiprocessor 1734.

25 **[0136]** The GPGPU cores 1762 can each include floating point units (FPUs) and/or integer arithmetic logic units (ALUs) that are used to execute instructions of the graphics multiprocessor 1734. In some implementations, the GPGPU cores 1762 can include hardware logic that may otherwise reside within the tensor and/or ray-tracing cores 1763. The GPGPU cores 1762 can be similar in architecture or can differ in architecture. For example and in some examples, a first portion of the GPGPU cores 1762 include a single precision FPU and an integer ALU while a second portion of the GPGPU cores include a double precision FPU. Optionally, the FPUs can implement the IEEE 754-2008 standard for floating point arithmetic or enable variable precision floating point arithmetic. The graphics multiprocessor 1734 can additionally include one or more fixed function or special function units to perform specific functions such as copy rectangle or pixel blending operations. One or more of the GPGPU cores can also include fixed or special function logic.

30 **[0137]** The GPGPU cores 1762 may include SIMD logic capable of performing a single instruction on multiple sets of data. Optionally, GPGPU cores 1762 can physically execute SIMD4, SIMD8, and SIMD16 instructions and logically execute SIMD1, SIMD2, and SIMD32 instructions. The SIMD instructions for the GPGPU cores can be generated at compile time by a shader compiler or automatically generated when executing programs written and compiled for single program multiple data (SPMD) or SIMT architectures. Multiple threads of a program configured for the SIMT execution model can be executed via a single SIMD instruction. For example and in some examples, eight SIMT threads that perform the same or similar operations can be executed in parallel via a single SIMD8 logic unit.

35 **[0138]** The memory and cache interconnect 1768 is an interconnect network that connects each of the functional units of the graphics multiprocessor 1734 to the register file 1758 and to the shared memory 1770. For example, the memory and cache interconnect 1768 is a crossbar interconnect that allows the load/store unit 1766 to implement load and store operations between the shared memory 1770 and the register file 1758. The register file 1758 can operate at the same frequency as the GPGPU cores 1762, thus data transfer between the GPGPU cores 1762 and the register file 1758 is very low latency. The shared memory 1770 can be used to enable communication between threads that execute on the functional units within the graphics multiprocessor 1734. The cache memory 1772 can be used as a data cache for example, to cache texture data communicated between the functional units and the texture unit 1736. The shared memory 1770 can also be used as a program managed cached. The shared memory 1770 and the cache memory 1772 can couple with the data crossbar 1740 to enable communication with other components of the processing cluster. Threads executing on the GPGPU cores 1762 can programmatically store data within the shared memory in addition to the automatically cached data that is stored within the cache memory 1772.

40 **[0139]** FIGS. 18A-18C illustrate additional graphics multiprocessors, according to examples. FIG. 18A-18B illustrate graphics multiprocessors 1825, 1850, which are related to the graphics multiprocessor 1734 of FIG. 17C and may be used in place of one of those. Therefore, the disclosure of any features in combination with the graphics multiprocessor 1734 herein also discloses a corresponding combination with the graphics multiprocessor(s) 1825, 1850, but is not limited to such. FIG. 18C illustrates a graphics processing unit (GPU) 1880 which includes dedicated sets of graphics

processing resources arranged into multi-core groups 1865A-1865N, which correspond to the graphics multiprocessors 1825, 1850. The illustrated graphics multiprocessors 1825, 1850 and the multi-core groups 1865A-1865N can be streaming multiprocessors (SM) capable of simultaneous execution of a large number of execution threads.

5 [0140] The graphics multiprocessor 1825 of FIG. 18A includes multiple additional instances of execution resource units relative to the graphics multiprocessor 1734 of FIG. 17D. For example, the graphics multiprocessor 1825 can include multiple instances of the instruction unit 1832A-1832B, register file 1834A-1834B, and texture unit(s) 1844A-1844B. The graphics multiprocessor 1825 also includes multiple sets of graphics or compute execution units (e.g., GPGPU core 1836A-1836B, tensor core 1837A-1837B, ray-tracing core 1838A-1838B) and multiple sets of load/store units 1840A-1840B. The execution resource units have a common instruction cache 1830, texture and/or data cache memory 1842, and shared memory 1846.

10 [0141] The various components can communicate via an interconnect fabric 1827. The interconnect fabric 1827 may include one or more crossbar switches to enable communication between the various components of the graphics multiprocessor 1825. The interconnect fabric 1827 may be a separate, high-speed network fabric layer upon which each component of the graphics multiprocessor 1825 is stacked. The components of the graphics multiprocessor 1825 communicate with remote components via the interconnect fabric 1827. For example, the cores 1836A-1836B, 1837A-1837B, and 1838A-1838B can each communicate with shared memory 1846 via the interconnect fabric 1827. The interconnect fabric 1827 can arbitrate communication within the graphics multiprocessor 1825 to ensure a fair bandwidth allocation between components.

15 [0142] The graphics multiprocessor 1850 of FIG. 18B includes multiple sets of execution resources 1856A-1856D, where each set of execution resource includes multiple instruction units, register files, GPGPU cores, and load store units, as illustrated in FIG. 17D and FIG. 18A. The execution resources 1856A-1856D can work in concert with texture unit(s) 1860A-1860D for texture operations, while sharing an instruction cache 1854, and shared memory 1853. For example, the execution resources 1856A-1856D can share an instruction cache 1854 and shared memory 1853, as well as multiple instances of a texture and/or data cache memory 1858A-1858B. The various components can communicate via an interconnect fabric 1852 similar to the interconnect fabric 1827 of FIG. 18A.

20 [0143] Persons skilled in the art will understand that the architecture described in FIG. 1, 17A-17D, and 18A-18B are descriptive and not limiting as to the scope of the present examples. Thus, the techniques described herein may be implemented on any properly configured processing unit, including, without limitation, one or more mobile application processors, one or more desktop or server central processing units (CPUs) including multi-core CPUs, one or more parallel processing units, such as the parallel processing unit 1702 of FIG. 17A, as well as one or more graphics processors or special purpose processing units, without departure from the scope of the examples described herein.

25 [0144] The parallel processor or GPGPU as described herein may be communicatively coupled to host/processor cores to accelerate graphics operations, machine-learning operations, pattern analysis operations, and various general-purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host processor/cores over a bus or other interconnect (e.g., a high-speed interconnect such as PCIe, NVLink, or other known protocols, standardized protocols, or proprietary protocols). In other examples, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

30 [0145] FIG. 18C illustrates a graphics processing unit (GPU) 1880 which includes dedicated sets of graphics processing resources arranged into multi-core groups 1865A-1865N. While the details of only a single multi-core group 1865A are provided, it will be appreciated that the other multi-core groups 1865B-1865N may be equipped with the same or similar sets of graphics processing resources. Details described with respect to the multi-core groups 1865A-1865N may also apply to any graphics multiprocessor 1734, 1825, 1850 described herein.

35 [0146] As illustrated, a multi-core group 1865A may include a set of graphics cores 1870, a set of tensor cores 1871, and a set of ray tracing cores 1872. A scheduler/dispatcher 1868 schedules and dispatches the graphics threads for execution on the various cores 1870, 1871, 1872. A set of register files 1869 store operand values used by the cores 1870, 1871, 1872 when executing the graphics threads. These may include, for example, integer registers for storing integer values, floating point registers for storing floating point values, vector registers for storing packed data elements (integer and/or floating-point data elements) and tile registers for storing tensor/matrix values. The tile registers may be implemented as combined sets of vector registers.

40 [0147] One or more combined level 1 (L1) caches and shared memory units 1873 store graphics data such as texture data, vertex data, pixel data, ray data, bounding volume data, etc., locally within each multi-core group 1865A. One or more texture units 1874 can also be used to perform texturing operations, such as texture mapping and sampling. A Level 2 (L2) cache 1875 shared by all or a subset of the multi-core groups 1865A-1865N stores graphics data and/or instructions for multiple concurrent graphics threads. As illustrated, the L2 cache 1875 may be shared across a plurality of multi-core groups 1865A-1865N. One or more memory controllers 1867 couple the GPU 1880 to a memory 1866

which may be a system memory (e.g., DRAM) and/or a dedicated graphics memory (e.g., GDDR6 memory).

[0148] Input/output (I/O) circuitry 1863 couples the GPU 1880 to one or more I/O devices 1862 such as digital signal processors (DSPs), network controllers, or user input devices. An on-chip interconnect may be used to couple the I/O devices 1862 to the GPU 1880 and memory 1866. One or more I/O memory management units (IOMMUs) 1864 of the I/O circuitry 1863 couple the I/O devices 1862 directly to the system memory 1866. Optionally, the IOMMU 1864 manages multiple sets of page tables to map virtual addresses to physical addresses in system memory 1866. The I/O devices 1862, CPU(s) 1861, and GPU(s) 1880 may then share the same virtual address space.

[0149] In one implementation of the IOMMU 1864, the IOMMU 1864 supports virtualization. In this case, it may manage a first set of page tables to map guest/graphics virtual addresses to guest/graphics physical addresses and a second set of page tables to map the guest/graphics physical addresses to system/host physical addresses (e.g., within system memory 1866). The base addresses of each of the first and second sets of page tables may be stored in control and/or status registers and swapped out on a context switch (e.g., so that the new context is provided with access to the relevant set of page tables). While not illustrated in FIG. 18C, each of the cores 1870, 1871, 1872 and/or multi-core groups 1865A-1865N may include translation lookaside buffers (TLBs) to cache guest virtual to guest physical translations, guest physical to host physical translations, and guest virtual to host physical translations.

[0150] The CPU(s) 1861, GPUs 1880, and I/O devices 1862 may be integrated on a single semiconductor chip and/or chip package. The illustrated memory 1866 may be integrated on the same chip or may be coupled to the memory controllers 1867 via an off-chip interface. In one implementation, the memory 1866 comprises GDDR6 memory which shares the same virtual address space as other physical system-level memories, although the underlying principles described herein are not limited to this specific implementation.

[0151] The tensor cores 1871 may include a plurality of execution units specifically designed to perform matrix operations, which are the fundamental compute operation used to perform deep learning operations. For example, simultaneous matrix multiplication operations may be used for neural network training and inferencing. The tensor cores 1871 may perform matrix processing using a variety of operand precisions including single precision floating-point (e.g., 32 bits), half-precision floating point (e.g., 16 bits), integer words (16 bits), bytes (8 bits), and half-bytes (4 bits). For example, a neural network implementation extracts features of each rendered scene, potentially combining details from multiple frames, to construct a high-quality final image.

[0152] In deep learning implementations, parallel matrix multiplication work may be scheduled for execution on the tensor cores 1871. The training of neural networks, in particular, requires a significant number of matrix dot product operations. In order to process an inner-product formulation of an $N \times N \times N$ matrix multiply, the tensor cores 1871 may include at least N dot-product processing elements. Before the matrix multiply begins, one entire matrix is loaded into tile registers and at least one column of a second matrix is loaded each cycle for N cycles. Each cycle, there are N dot products that are processed.

[0153] Matrix elements may be stored at different precisions depending on the particular implementation, including 16-bit words, 8-bit bytes (e.g., INT8) and 4-bit half-bytes (e.g., INT4). Different precision modes may be specified for the tensor cores 1871 to ensure that the most efficient precision is used for different workloads (e.g., such as inferencing workloads which can tolerate quantization to bytes and half-bytes). Supported formats additionally include 64-bit floating point (FP64) and non-IEEE floating point formats such as the bfloat16 format (e.g., Brain floating point), a 16-bit floating point format with one sign bit, eight exponent bits, and eight significand bits, of which seven are explicitly stored. One example includes support for a reduced precision tensor-float (TF32) mode, which performs computations using the range of FP32 (8-bits) and the precision of FP16 (10-bits). Reduced precision TF32 operations can be performed on FP32 inputs and produce FP32 outputs at higher performance relative to FP32 and increased precision relative to FP16. In some examples, one or more 8-bit floating point formats (FP8) are supported.

[0154] In some examples the tensor cores 1871 support a sparse mode of operation for matrices in which the vast majority of values are zero. The tensor cores 1871 include support for sparse input matrices that are encoded in a sparse matrix representation (e.g., coordinate list encoding (COO), compressed sparse row (CSR), compressed sparse column (CSC), etc.). The tensor cores 1871 also include support for compressed sparse matrix representations in the event that the sparse matrix representation may be further compressed. Compressed, encoded, and/or compressed and encoded matrix data, along with associated compression and/or encoding metadata, can be read by the tensor cores 1871 and the non-zero values can be extracted. For example, for a given input matrix A, a non-zero value can be loaded from the compressed and/or encoded representation of at least a portion of matrix A. Based on the location in matrix A for the non-zero value, which may be determined from index or coordinate metadata associated with the non-zero value, a corresponding value in input matrix B may be loaded. Depending on the operation to be performed (e.g., multiply), the load of the value from input matrix B may be bypassed if the corresponding value is a zero value. In some examples, the pairings of values for certain operations, such as multiply operations, may be pre-scanned by scheduler logic and only operations between non-zero inputs are scheduled. Depending on the dimensions of matrix A and matrix B and the operation to be performed, output matrix C may be dense or sparse. Where output matrix C is sparse and depending on the configuration of the tensor cores 1871, output matrix C may be output in a compressed format, a sparse encoding,

or a compressed sparse encoding.

[0155] The ray tracing cores 1872 may accelerate ray tracing operations for both real-time ray tracing and non-real-time ray tracing implementations. In particular, the ray tracing cores 1872 may include ray traversal/intersection circuitry for performing ray traversal using bounding volume hierarchies (BVHs) and identifying intersections between rays and primitives enclosed within the BVH volumes. The ray tracing cores 1872 may also include circuitry for performing depth testing and culling (e.g., using a Z buffer or similar arrangement). In one implementation, the ray tracing cores 1872 perform traversal and intersection operations in concert with the image denoising techniques described herein, at least a portion of which may be executed on the tensor cores 1871. For example, the tensor cores 1871 may implement a deep learning neural network to perform denoising of frames generated by the ray tracing cores 1872. However, the CPU(s) 1861, graphics cores 1870, and/or ray tracing cores 1872 may also implement all or a portion of the denoising and/or deep learning algorithms.

[0156] In addition, as described above, a distributed approach to denoising may be employed in which the GPU 1880 is in a computing device coupled to other computing devices over a network or high-speed interconnect. In this distributed approach, the interconnected computing devices may share neural network learning/training data to improve the speed with which the overall system learns to perform denoising for different types of image frames and/or different graphics applications.

[0157] The ray tracing cores 1872 may process all BVH traversal and/or ray-primitive intersections, saving the graphics cores 1870 from being overloaded with thousands of instructions per ray. For example, each ray tracing core 1872 includes a first set of specialized circuitry for performing bounding box tests (e.g., for traversal operations) and/or a second set of specialized circuitry for performing the ray-triangle intersection tests (e.g., intersecting rays which have been traversed). Thus, for example, the multi-core group 1865A can simply launch a ray probe, and the ray tracing cores 1872 independently perform ray traversal and intersection and return hit data (e.g., a hit, no hit, multiple hits, etc.) to the thread context. The other cores 1870, 1871 are freed to perform other graphics or compute work while the ray tracing cores 1872 perform the traversal and intersection operations.

[0158] Optionally, each ray tracing core 1872 may include a traversal unit to perform BVH testing operations and/or an intersection unit which performs ray-primitive intersection tests. The intersection unit generates a "hit", "no hit", or "multiple hit" response, which it provides to the appropriate thread. During the traversal and intersection operations, the execution resources of the other cores (e.g., graphics cores 1870 and tensor cores 1871) are freed to perform other forms of graphics work.

[0159] In some examples described below, a hybrid rasterization/ray tracing approach is used in which work is distributed between the graphics cores 1870 and ray tracing cores 1872.

[0160] The ray tracing cores 1872 (and/or other cores 1870, 1871) may include hardware support for a ray tracing instruction set such as Microsoft's DirectX Ray Tracing (DXR) which includes a DispatchRays command, as well as ray-generation, closest-hit, any-hit, and miss shaders, which enable the assignment of unique sets of shaders and textures for each object. Another ray tracing platform which may be supported by the ray tracing cores 1872, graphics cores 1870 and tensor cores 1871 is Vulkan API (e.g., Vulkan version 1.1.85 and later). Note, however, that the underlying principles described herein are not limited to any particular ray tracing ISA.

[0161] In general, the various cores 1872, 1871, 1870 may support a ray tracing instruction set that includes instructions/functions for one or more of ray generation, closest hit, any hit, ray-primitive intersection, per-primitive and hierarchical bounding box construction, miss, visit, and exceptions. More specifically, some examples includes ray tracing instructions to perform one or more of the following functions:

- Ray Generation - Ray generation instructions may be executed for each pixel, sample, or other user-defined work assignment.
- Closest Hit - A closest hit instruction may be executed to locate the closest intersection point of a ray with primitives within a scene.
- Any Hit - An any hit instruction identifies multiple intersections between a ray and primitives within a scene, potentially to identify a new closest intersection point.
- Intersection - An intersection instruction performs a ray-primitive intersection test and outputs a result.
- Per-primitive Bounding box Construction -This instruction builds a bounding box around a given primitive or group of primitives (e.g., when building a new BVH or other acceleration data structure).
- Miss - Indicates that a ray misses all geometry within a scene, or specified region of a scene.

- Visit - Indicates the child volumes a ray will traverse.
- Exceptions - Includes various types of exception handlers (e.g., invoked for various error conditions).

5 **[0162]** In some examples the ray tracing cores 1872 may be adapted to accelerate general-purpose compute operations that can be accelerated using computational techniques that are analogous to ray intersection tests. A compute framework can be provided that enables shader programs to be compiled into low level instructions and/or primitives that perform general-purpose compute operations via the ray tracing cores. Exemplary computational problems that can benefit from compute operations performed on the ray tracing cores 1872 include computations involving beam, wave, ray, or particle propagation within a coordinate space. Interactions associated with that propagation can be computed relative to a geometry or mesh within the coordinate space. For example, computations associated with electromagnetic signal propagation through an environment can be accelerated via the use of instructions or primitives that are executed via the ray tracing cores. Diffraction and reflection of the signals by objects in the environment can be computed as direct ray-tracing analogies.

15 **[0163]** Ray tracing cores 1872 can also be used to perform computations that are not directly analogous to ray tracing. For example, mesh projection, mesh refinement, and volume sampling computations can be accelerated using the ray tracing cores 1872. Generic coordinate space calculations, such as nearest neighbor calculations can also be performed. For example, the set of points near a given point can be discovered by defining a bounding box in the coordinate space around the point. BVH and ray probe logic within the ray tracing cores 1872 can then be used to determine the set of point intersections within the bounding box. The intersections constitute the origin point and the nearest neighbors to that origin point. Computations that are performed using the ray tracing cores 1872 can be performed in parallel with computations performed on the graphics cores 1872 and tensor cores 1871. A shader compiler can be configured to compile a compute shader or other general-purpose graphics processing program into low level primitives that can be parallelized across the graphics cores 1870, tensor cores 1871, and ray tracing cores 1872.

25 **[0164]** Building larger and larger silicon dies is challenging for a variety of reasons. As silicon dies become larger, manufacturing yields become smaller and process technology requirements for different components may diverge. On the other hand, in order to have a high-performance system, key components should be interconnected by high speed, high bandwidth, low latency interfaces. These contradicting needs pose a challenge to high performance chip development.

30 **[0165]** Embodiments described herein provide techniques to disaggregate an architecture of a system on a chip integrated circuit into multiple distinct chiplets that can be packaged onto a common chassis. In some examples, a graphics processing unit or parallel processor is composed from diverse silicon chiplets that are separately manufactured. A chiplet is an at least partially packaged integrated circuit that includes distinct units of logic that can be assembled with other chiplets into a larger package. A diverse set of chiplets with different IP core logic can be assembled into a single device. Additionally the chiplets can be integrated into a base die or base chiplet using active interposer technology. The concepts described herein enable the interconnection and communication between the different forms of IP within the GPU. The development of IPs on different process may be mixed. This avoids the complexity of converging multiple IPs, especially on a large SoC with several flavors IPs, to the same process.

35 **[0166]** Enabling the use of multiple process technologies improves the time to market and provides a cost-effective way to create multiple product SKUs. For customers, this means getting products that are more tailored to their requirements in a cost effective and timely manner. Additionally, the disaggregated IPs are more amenable to being power gated independently, components that are not in use on a given workload can be powered off, reducing overall power consumption.

40 **[0167]** FIG. 19 shows a parallel compute system 1900, according to some examples. In some examples the parallel compute system 1900 includes a parallel processor 1920, which can be a graphics processor or compute accelerator as described herein. The parallel processor 1920 includes a global logic unit 1901, an interface 1902, a thread dispatcher 1903, a media unit 1904, a set of compute units 1905A-1905H, and a cache/memory units 1906. The global logic unit 1901, in some examples, includes global functionality for the parallel processor 1920, including device configuration registers, global schedulers, power management logic, and the like. The interface 1902 can include a front-end interface for the parallel processor 1920. The thread dispatcher 1903 can receive workloads from the interface 1902 and dispatch threads for the workload to the compute units 1905A-1905H. If the workload includes any media operations, at least a portion of those operations can be performed by the media unit 1904. The media unit can also offload some operations to the compute units 1905A-1905H. The cache/memory units 1906 can include cache memory (e.g., L3 cache) and local memory (e.g., HBM, GDDR) for the parallel processor 1920.

45 **[0168]** FIGS. 20A-20B illustrate a hybrid logical/physical view of a disaggregated parallel processor, according to examples described herein. FIG. 20A illustrates a disaggregated parallel compute system 2000. FIG. 20B illustrates a chiplet 2030 of the disaggregated parallel compute system 2000.

[0169] As shown in FIG. 20A, a disaggregated compute system 2000 can include a parallel processor 2020 in which

the various components of the parallel processor SOC are distributed across multiple chiplets. Each chiplet can be a distinct IP core that is independently designed and configured to communicate with other chiplets via one or more common interfaces. The chiplets include but are not limited to compute chiplets 2005, a media chiplet 2004, and memory chiplets 2006. Each chiplet can be separately manufactured using different process technologies. For example, compute chiplets 2005 may be manufactured using the smallest or most advanced process technology available at the time of fabrication, while memory chiplets 2006 or other chiplets (e.g., I/O, networking, etc.) may be manufactured using a larger or less advanced process technologies.

[0170] The various chiplets can be bonded to a base die 2010 and configured to communicate with each other and logic within the base die 2010 via an interconnect layer 2012. In some examples, the base die 2010 can include global logic 2001, which can include scheduler 2011 and power management 2021 logic units, an interface 2002, a dispatch unit 2003, and an interconnect fabric module 2008 coupled with or integrated with one or more L3 cache banks 2009A-2009N. The interconnect fabric 2008 can be an inter-chiplet fabric that is integrated into the base die 2010. Logic chiplets can use the fabric 2008 to relay messages between the various chiplets. Additionally, L3 cache banks 2009A-2009N in the base die and/or L3 cache banks within the memory chiplets 2006 can cache data read from and transmitted to DRAM chiplets within the memory chiplets 2006 and to system memory of a host.

[0171] In some examples the global logic 2001 is a microcontroller that can execute firmware to perform scheduler 2011 and power management 2021 functionality for the parallel processor 2020. The microcontroller that executes the global logic can be tailored for the target use case of the parallel processor 2020. The scheduler 2011 can perform global scheduling operations for the parallel processor 2020. The power management 2021 functionality can be used to enable or disable individual chiplets within the parallel processor when those chiplets are not in use.

[0172] The various chiplets of the parallel processor 2020 can be designed to perform specific functionality that, in existing designs, would be integrated into a single die. A set of compute chiplets 2005 can include clusters of compute units (e.g., execution units, streaming multiprocessors, etc.) that include programmable logic to execute compute or graphics shader instructions. A media chiplet 2004 can include hardware logic to accelerate media encode and decode operations. Memory chiplets 2006 can include volatile memory (e.g., DRAM) and one or more SRAM cache memory banks (e.g., L3 banks).

[0173] As shown in FIG. 20B, each chiplet 2030 can include common components and application specific components. Chiplet logic 2036 within the chiplet 2030 can include the specific components of the chiplet, such as an array of streaming multiprocessors, compute units, or execution units described herein. The chiplet logic 2036 can couple with an optional cache or shared local memory 2038 or can include a cache or shared local memory within the chiplet logic 2036. The chiplet 2030 can include a fabric interconnect node 2042 that receives commands via the inter-chiplet fabric. Commands and data received via the fabric interconnect node 2042 can be stored temporarily within an interconnect buffer 2039. Data transmitted to and received from the fabric interconnect node 2042 can be stored in an interconnect cache 2040. Power control 2032 and clock control 2034 logic can also be included within the chiplet. The power control 2032 and clock control 2034 logic can receive configuration commands via the fabric can configure dynamic voltage and frequency scaling for the chiplet 2030. In some examples, each chiplet can have an independent clock domain and power domain and can be clock gated and power gated independently of other chiplets.

[0174] At least a portion of the components within the illustrated chiplet 2030 can also be included within logic embedded within the base die 2010 of FIG. 20A. For example, logic within the base die that communicates with the fabric can include a version of the fabric interconnect node 2042. Base die logic that can be independently clock or power gated can include a version of the power control 2032 and/or clock control 2034 logic.

[0175] Thus, while various examples described herein use the term SOC to describe a device or system having a processor and associated circuitry (e.g., Input/Output ("I/O") circuitry, power delivery circuitry, memory circuitry, etc.) integrated monolithically into a single Integrated Circuit ("IC") die, or chip, the present disclosure is not limited in that respect. For example, in various examples of the present disclosure, a device or system can have one or more processors (e.g., one or more processor cores) and associated circuitry (e.g., Input/Output ("I/O") circuitry, power delivery circuitry, etc.) arranged in a disaggregated collection of discrete dies, tiles and/or chiplets (e.g., one or more discrete processor core die arranged adjacent to one or more other die such as memory die, I/O die, etc.). In such disaggregated devices and systems the various dies, tiles and/or chiplets can be physically and electrically coupled together by a package structure including, for example, various packaging substrates, interposers, active interposers, photonic interposers, interconnect bridges and the like. The disaggregated collection of discrete dies, tiles, and/or chiplets can also be part of a System-on-Package ("SoP")."

Example Core Architectures - In-order and out-of-order core block diagram.

[0176] FIG. 21(A) is a block diagram illustrating both an example in-order pipeline and an example register renaming, out-of-order issue/execution pipeline according to examples. FIG. 21(B) is a block diagram illustrating both an example in-order architecture core and an example register renaming, out-of-order issue/execution architecture core to be included

in a processor according to examples. The solid lined boxes in FIGS. 21(A)-(B) illustrate the in-order pipeline and in-order core, while the optional addition of the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline and core. Given that the in-order aspect is a subset of the out-of-order aspect, the out-of-order aspect will be described.

[0177] In FIG. 21(A), a processor pipeline 2100 includes a fetch stage 2102, an optional length decoding stage 2104, a decode stage 2106, an optional allocation (Alloc) stage 2108, an optional renaming stage 2110, a schedule (also known as a dispatch or issue) stage 2112, an optional register read/memory read stage 2114, an execute stage 2116, a write back/memory write stage 2118, an optional exception handling stage 2122, and an optional commit stage 2124. One or more operations can be performed in each of these processor pipeline stages. For example, during the fetch stage 2102, one or more instructions are fetched from instruction memory, and during the decode stage 2106, the one or more fetched instructions may be decoded, addresses (e.g., load store unit (LSU) addresses) using forwarded register ports may be generated, and branch forwarding (e.g., immediate value offset or a link register (LR)) may be performed. In some examples, the decode stage 2106 and the register read/memory read stage 2114 may be combined into one pipeline stage. In some examples, during the execute stage 2116, the decoded instructions may be executed, LSU address/data pipelining to an Advanced Microcontroller Bus (AMB) interface may be performed, multiply and add operations may be performed, arithmetic operations with branch results may be performed, etc.

[0178] By way of example, the example register renaming, out-of-order issue/execution architecture core of FIG. 21(B) may implement the pipeline 2100 as follows: 1) the instruction fetch circuitry 2138 performs the fetch and length decoding stages 2102 and 2104; 2) the decode circuitry 2140 performs the decode stage 2106; 3) the rename/allocator unit circuitry 2152 performs the allocation stage 2108 and renaming stage 2110; 4) the scheduler(s) circuitry 2156 performs the schedule stage 2112; 5) the physical register file(s) circuitry 2158 and the memory unit circuitry 2170 perform the register read/memory read stage 2114; the execution cluster(s) 2160 perform the execute stage 2116; 6) the memory unit circuitry 2170 and the physical register file(s) circuitry 2158 perform the write back/memory write stage 2118; 7) various circuitry may be involved in the exception handling stage 2122; and 8) the retirement unit circuitry 2154 and the physical register file(s) circuitry 2158 perform the commit stage 2124.

[0179] FIG. 21(B) shows a processor core 2190 including front-end unit circuitry 2130 coupled to execution engine unit circuitry 2150, and both are coupled to memory unit circuitry 2170. The core 2190 may be a reduced instruction set architecture computing (RISC) core, a complex instruction set architecture computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 2190 may be a special-purpose core, such as, for example, a network or communication core, compression engine, coprocessor core, general purpose computing graphics processing unit (GPGPU) core, graphics core, or the like.

[0180] The front-end unit circuitry 2130 may include branch prediction circuitry 2132 coupled to instruction cache circuitry 2134, which is coupled to an instruction translation lookaside buffer (TLB) 2136, which is coupled to instruction fetch circuitry 2138, which is coupled to decode circuitry 2140. In some examples, the instruction cache circuitry 2134 is included in the memory unit circuitry 2170 rather than the front-end circuitry 2130. The decode circuitry 2140 (or decoder) 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 are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decode circuitry 2140 may further include address generation unit (AGU, not shown) circuitry. In some examples, the AGU generates an LSU address using forwarded register ports, and may further perform branch forwarding (e.g., immediate value offset branch forwarding, LR register branch forwarding, etc.). The decode circuitry 2140 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 some examples, the core 2190 includes a microcode ROM (not shown) or other medium that stores microcode for certain macroinstructions (e.g., in decode circuitry 2140 or otherwise within the front-end circuitry 2130). In some examples, the decode circuitry 2140 includes a micro-operation (micro-op) or operation cache (not shown) to hold/cache decoded operations, micro-tags, or micro-operations generated during the decode or other stages of the processor pipeline 2100. The decode circuitry 2140 may be coupled to rename/allocator unit circuitry 2152 in the execution engine circuitry 2150.

[0181] The execution engine circuitry 2150 includes the rename/allocator unit circuitry 2152 coupled to retirement unit circuitry 2154 and a set of one or more scheduler(s) circuitry 2156. The scheduler(s) circuitry 2156 represents any number of different schedulers, including reservations stations, central instruction window, etc. In some examples, the scheduler(s) circuitry 2156 can include arithmetic logic unit (ALU) scheduler/scheduling circuitry, ALU queues, address generation unit (AGU) scheduler/scheduling circuitry, AGU queues, etc. The scheduler(s) circuitry 2156 is coupled to the physical register file(s) circuitry 2158. Each of the physical register file(s) circuitry 2158 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, status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. In some examples, the physical register file(s) circuitry 2158 includes vector registers unit circuitry, writemask registers unit circuitry, and scalar register unit circuitry.

These register units may provide architectural vector registers, vector mask registers, general-purpose registers, etc. The physical register file(s) circuitry 2158 is coupled to the retirement unit circuitry 2154 (also known as a retire queue or a retirement queue) to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) (ROB(s)) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement unit circuitry 2154 and the physical register file(s) circuitry 2158 are coupled to the execution cluster(s) 2160. The execution cluster(s) 2160 includes a set of one or more execution unit(s) circuitry 2162 and a set of one or more memory access circuitry 2164. The execution unit(s) circuitry 2162 may perform various arithmetic, logic, floating-point or other types of operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar integer, scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point). While some examples may include a number of execution units or execution unit circuitry dedicated to specific functions or sets of functions, other examples may include only one execution unit circuitry or multiple execution units/execution unit circuitry that all perform all functions. The scheduler(s) circuitry 2156, physical register file(s) circuitry 2158, and execution cluster(s) 2160 are shown as being possibly plural because certain examples 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 circuitry, physical register file(s) circuitry, and/or execution cluster - and in the case of a separate memory access pipeline, certain examples are implemented in which only the execution cluster of this pipeline has the memory access unit(s) circuitry 2164). 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.

[0182] In some examples, the execution engine unit circuitry 2150 may perform load store unit (LSU) address/data pipelining to an Advanced Microcontroller Bus (AMB) interface (not shown), and address phase and writeback, data phase load, store, and branches.

[0183] The set of memory access circuitry 2164 is coupled to the memory unit circuitry 2170, which includes data TLB circuitry 2172 coupled to data cache circuitry 2174 coupled to level 2 (L2) cache circuitry 2176. In some examples, the memory access circuitry 2164 may include load unit circuitry, store address unit circuitry, and store data unit circuitry, each of which is coupled to the data TLB circuitry 2172 in the memory unit circuitry 2170. The instruction cache circuitry 2134 is further coupled to the level 2 (L2) cache circuitry 2176 in the memory unit circuitry 2170. In some examples, the instruction cache 2134 and the data cache 2174 are combined into a single instruction and data cache (not shown) in L2 cache circuitry 2176, level 3 (L3) cache circuitry (not shown), and/or main memory. The L2 cache circuitry 2176 is coupled to one or more other levels of cache and eventually to a main memory.

[0184] The core 2190 may support one or more instructions sets (e.g., the x86 instruction set architecture (optionally with some extensions that have been added with newer versions); the MIPS instruction set architecture; the ARM instruction set architecture (optionally with optional additional extensions such as NEON)), including the instruction(s) described herein. In some examples, the core 2190 includes logic to support a packed data instruction set architecture extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

Example Execution Unit(s) Circuitry.

[0185] FIG. 22 illustrates examples of execution unit(s) circuitry, such as execution unit(s) circuitry 2162 of FIG. 21(B). As illustrated, execution unit(s) circuitry 2162 may include one or more ALU circuits 2201, optional vector/single instruction multiple data (SIMD) circuits 2203, load/store circuits 2205, branch/jump circuits 2207, and/or Floating-point unit (FPU) circuits 2209. ALU circuits 2201 perform integer arithmetic and/or Boolean operations. Vector/SIMD circuits 2203 perform vector/SIMD operations on packed data (such as SIMD/vector registers). Load/store circuits 2205 execute load and store instructions to load data from memory into registers or store from registers to memory. Load/store circuits 2205 may also generate addresses. Branch/jump circuits 2207 cause a branch or jump to a memory address depending on the instruction. FPU circuits 2209 perform floating-point arithmetic. The width of the execution unit(s) circuitry 2162 varies depending upon the example and can range from 16-bit to 1,024-bit, for example. In some examples, two or more smaller execution units are logically combined to form a larger execution unit (e.g., two 128-bit execution units are logically combined to form a 256-bit execution unit).

Example Register Architecture.

[0186] FIG. 23 is a block diagram of a register architecture 2300 according to some examples. As illustrated, the register architecture 2300 includes vector/SIMD registers 2310 that vary from 128-bit to 1,024 bits width. In some examples, the vector/SIMD registers 2310 are physically 512-bits and, depending upon the mapping, only some of the lower bits are used. For example, in some examples, the vector/SIMD registers 2310 are ZMM registers which are 512

bits: the lower 256 bits are used for YMM registers and the lower 128 bits are used for XMM registers. As such, there is an overlay of registers. In some examples, a vector length field selects between a maximum length and one or more other shorter lengths, where each such shorter length is half the length of the preceding length. Scalar operations are operations performed on the lowest order data element position in a ZMM/YMM/XMM register; the higher order data element positions are either left the same as they were prior to the instruction or zeroed depending on the example.

[0187] In some examples, the register architecture 2300 includes writemask/predicate registers 2315. For example, in some examples, there are 8 writemask/predicate registers (sometimes called k0 through k7) that are each 16-bit, 32-bit, 64-bit, or 128-bit in size. Writemask/predicate registers 2315 may allow for merging (e.g., allowing any set of elements in the destination to be protected from updates during the execution of any operation) and/or zeroing (e.g., zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation). In some examples, each data element position in a given writemask/predicate register 2315 corresponds to a data element position of the destination. In other examples, the writemask/predicate registers 2315 are scalable and consists of a set number of enable bits for a given vector element (e.g., 8 enable bits per 64-bit vector element).

[0188] The register architecture 2300 includes a plurality of general-purpose registers 2325. These registers may be 16-bit, 32-bit, 64-bit, etc. and can be used for scalar operations. In some examples, these registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

[0189] In some examples, the register architecture 2300 includes scalar floating-point (FP) register file 2345 which is used for scalar floating-point operations on 32/64/80-bit floating-point data using the x87 instruction set architecture extension or as MMX registers to perform operations on 64-bit packed integer data, as well as to hold operands for some operations performed between the MMX and XMM registers.

[0190] One or more flag registers 2340 (e.g., EFLAGS, RFLAGS, etc.) store status and control information for arithmetic, compare, and system operations. For example, the one or more flag registers 2340 may store condition code information such as carry, parity, auxiliary carry, zero, sign, and overflow. In some examples, the one or more flag registers 2340 are called program status and control and/or status registers.

[0191] Segment registers 2320 contain segment points for use in accessing memory. In some examples, these registers are referenced by the names CS, DS, SS, ES, FS, and GS.

[0192] Model specific registers or machine specific registers (MSRs) 2335 control and report on processor performance. Most MSRs 2335 handle system-related functions and are not accessible to an application program. For example, MSRs may provide control for one or more of: performance-monitoring counters, debug extensions, memory type range registers, thermal and power management, instruction-specific support, and/or processor feature/mode support. Machine check registers 2360 consist of control, status, and error reporting MSRs that are used to detect and report on hardware errors. Control and/or status register(s) 2355 (e.g., CR0-CR4) determine the operating mode of a processor (e.g., processor 1470, 1480, 1438, 1415, and/or 1500) and the characteristics of a currently executing task. In some examples, MSRs 2335 are a subset of control and/or status registers 2355.

[0193] One or more instruction pointer register(s) 2330 store an instruction pointer value. Debug registers 2350 control and allow for the monitoring of a processor or core's debugging operations.

[0194] Memory (mem) management registers 2365 specify the locations of data structures used in protected mode memory management. These registers may include a global descriptor table register (GDTR), interrupt descriptor table register (IDTR), task register, and a local descriptor table register (LDTR) register.

[0195] Alternative examples may use wider or narrower registers. Additionally, alternative examples may use more, less, or different register files and registers. The register architecture 2300 may, for example, be used in register file / memory 1308, or physical register file(s) circuitry 2158.

Instruction set architectures.

[0196] An instruction set architecture (ISA) may include one or more instruction formats. A given instruction format may define various fields (e.g., number of bits, location of bits) to specify, among other things, the operation to be performed (e.g., opcode) and the operand(s) on which that operation is to be performed and/or other data field(s) (e.g., mask). Some instruction formats are further broken down through the definition of 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 (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an example ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands. In addition, though the description below is made in the context of x86 ISA, it is within the knowledge of one skilled in the art to apply the

teachings of the present disclosure in another ISA.

Example Instruction Formats.

5 **[0197]** Examples of the instruction(s) described herein may be embodied in different formats. Additionally, example systems, architectures, and pipelines are detailed below. Examples of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

10 **[0198]** FIG. 24 illustrates examples of an instruction format. As illustrated, an instruction may include multiple components including, but not limited to, one or more prefixes 2401, an opcode 2403, addressing information 2405 (e.g., register identifiers, memory addressing information, etc.), a displacement value 2407, and/or an immediate value 2409. Note that some instructions utilize some or all the fields of the format whereas others may only use the field for the opcode 2403. In some examples, the order illustrated is the order in which these fields are to be encoded, however, it should be appreciated that in other examples these fields may be encoded in a different order, combined, etc.

15 **[0199]** The prefix(es) field(s) 2401, when used, modifies an instruction. In some examples, one or more prefixes are used to repeat string instructions (e.g., 0xF0, 0xF2, 0xF3, etc.), to provide section overrides (e.g., 0x2E, 0x36, 0x3E, 0x26, 0x64, 0x65, 0x2E, 0x3E, etc.), to perform bus lock operations, and/or to change operand (e.g., 0x66) and address sizes (e.g., 0x67). Certain instructions require a mandatory prefix (e.g., 0x66, 0xF2, 0xF3, etc.). Certain of these prefixes may be considered "legacy" prefixes. Other prefixes, one or more examples of which are detailed herein, indicate, and/or provide further capability, such as specifying particular registers, etc. The other prefixes typically follow the "legacy" prefixes.

20 **[0200]** The opcode field 2403 is used to at least partially define the operation to be performed upon a decoding of the instruction. In some examples, a primary opcode encoded in the opcode field 2403 is one, two, or three bytes in length. In other examples, a primary opcode can be a different length. An additional 3-bit opcode field is sometimes encoded in another field.

25 **[0201]** The addressing information field 2405 is used to address one or more operands of the instruction, such as a location in memory or one or more registers. FIG. 25 illustrates examples of the addressing information field 2405. In this illustration, an optional MOD R/M byte 2502 and an optional Scale, Index, Base (SIB) byte 2504 are shown. The MOD R/M byte 2502 and the SIB byte 2504 are used to encode up to two operands of an instruction, each of which is a direct register or effective memory address. Note that both of these fields are optional in that not all instructions include one or more of these fields. The MOD R/M byte 2502 includes a MOD field 2542, a register (reg) field 2544, and R/M field 2546.

30 **[0202]** The content of the MOD field 2542 distinguishes between memory access and non-memory access modes. In some examples, when the MOD field 2542 has a binary value of 11 (11b), a register-direct addressing mode is utilized, and otherwise a register-indirect addressing mode is used.

35 **[0203]** The register field 2544 may encode either the destination register operand or a source register operand or may encode an opcode extension and not be used to encode any instruction operand. The content of register field 2544, directly or through address generation, specifies the locations of a source or destination operand (either in a register or in memory). In some examples, the register field 2544 is supplemented with an additional bit from a prefix (e.g., prefix 2401) to allow for greater addressing.

40 **[0204]** The R/M field 2546 may be used to encode an instruction operand that references a memory address or may be used to encode either the destination register operand or a source register operand. Note the R/M field 2546 may be combined with the MOD field 2542 to dictate an addressing mode in some examples.

45 **[0205]** The SIB byte 2504 includes a scale field 2552, an index field 2554, and a base field 2556 to be used in the generation of an address. The scale field 2552 indicates a scaling factor. The index field 2554 specifies an index register to use. In some examples, the index field 2554 is supplemented with an additional bit from a prefix (e.g., prefix 2401) to allow for greater addressing. The base field 2556 specifies a base register to use. In some examples, the base field 2556 is supplemented with an additional bit from a prefix (e.g., prefix 2401) to allow for greater addressing. In practice, the content of the scale field 2552 allows for the scaling of the content of the index field 2554 for memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base}$).

50 **[0206]** Some addressing forms utilize a displacement value to generate a memory address. For example, a memory address may be generated according to $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$, $\text{index} * \text{scale} + \text{displacement}$, $r/m + \text{displacement}$, instruction pointer (RIP/EIP) + displacement, register + displacement, etc. The displacement may be a 1-byte, 2-byte, 4-byte, etc. value. In some examples, the displacement field 2407 provides this value. Additionally, in some examples, a displacement factor usage is encoded in the MOD field of the addressing information field 2405 that indicates a compressed displacement scheme for which a displacement value is calculated and stored in the displacement field 2407.

55 **[0207]** In some examples, the immediate value field 2409 specifies an immediate value for the instruction. An immediate

value may be encoded as a 1-byte value, a 2-byte value, a 4-byte value, etc.

[0208] FIG. 26 illustrates examples of a first prefix 2401(A). In some examples, the first prefix 2401(A) is an example of a REX prefix. Instructions that use this prefix may specify general purpose registers, 64-bit packed data registers (e.g., single instruction, multiple data (SIMD) registers or vector registers), and/or control and/or status registers and debug registers (e.g., CR8-CR15 and DR8-DR15).

[0209] Instructions using the first prefix 2401(A) may specify up to three registers using 3-bit fields depending on the format: 1) using the reg field 2544 and the R/M field 2546 of the MOD R/M byte 2502; 2) using the MOD R/M byte 2502 with the SIB byte 2504 including using the reg field 2544 and the base field 2556 and index field 2554; or 3) using the register field of an opcode.

[0210] In the first prefix 2401(A), bit positions of the payload byte 7:4 are set as 0100. Bit position 3 (W) can be used to determine the operand size but may not solely determine operand width. As such, when W = 0, the operand size is determined by a code segment descriptor (CS.D) and when W = 1, the operand size is 64-bit.

[0211] Note that the addition of another bit allows for 16 (2⁴) registers to be addressed, whereas the MOD R/M reg field 2544 and MOD R/M R/M field 2546 alone can each only address 8 registers.

[0212] In the first prefix 2401(A), bit position 2 (R) may be an extension of the MOD R/M reg field 2544 and may be used to modify the MOD R/M reg field 2544 when that field encodes a general-purpose register, a 64-bit packed data register (e.g., a SSE register), or a control or debug register. R is ignored when MOD R/M byte 2502 specifies other registers or defines an extended opcode.

[0213] Bit position 1 (X) may modify the SIB byte index field 2554.

[0214] Bit position 0 (B) may modify the base in the MOD R/M R/M field 2546 or the SIB byte base field 2556; or it may modify the opcode register field used for accessing general purpose registers (e.g., general purpose registers 2325).

[0215] FIGS. 27(A)-(D) illustrate examples of how the R, X, and B fields of the first prefix 2401(A) are used. FIG. 27(A) illustrates R and B from the first prefix 2401(A) being used to extend the reg field 2544 and R/M field 2546 of the MOD R/M byte 2502 when the SIB byte 25 04 is not used for memory addressing. FIG. 27(B) illustrates R and B from the first prefix 2401(A) being used to extend the reg field 2544 and R/M field 2546 of the MOD R/M byte 2502 when the SIB byte 25 04 is not used (register-register addressing). FIG. 27(C) illustrates R, X, and B from the first prefix 2401(A) being used to extend the reg field 2544 of the MOD R/M byte 2502 and the index field 2554 and base field 2556 when the SIB byte 25 04 being used for memory addressing. FIG. 27(D) illustrates B from the first prefix 2401(A) being used to extend the reg field 2544 of the MOD R/M byte 2502 when a register is encoded in the opcode 2403.

[0216] FIGS. 28(A)-(B) illustrate examples of a second prefix 2401(B). In some examples, the second prefix 2401(B) is an example of a VEX prefix. The second prefix 2401(B) encoding allows instructions to have more than two operands, and allows SIMD vector registers (e.g., vector/SIMD registers 2310) to be longer than 64-bits (e.g., 128-bit and 256-bit). The use of the second prefix 2401(B) provides for three-operand (or more) syntax. For example, previous two-operand instructions performed operations such as A = A + B, which overwrites a source operand. The use of the second prefix 2401(B) enables operands to perform nondestructive operations such as A = B + C.

[0217] In some examples, the second prefix 2401(B) comes in two forms - a two-byte form and a three-byte form. The two-byte second prefix 2401(B) is used mainly for 128-bit, scalar, and some 256-bit instructions; while the three-byte second prefix 2401(B) provides a compact replacement of the first prefix 2401(A) and 3-byte opcode instructions.

[0218] FIG. 28(A) illustrates examples of a two-byte form of the second prefix 2401(B). In some examples, a format field 2801 (byte 0 2803) contains the value C5H. In some examples, byte 1 2805 includes an "R" value in bit[7]. This value is the complement of the "R" value of the first prefix 2401(A). Bit[2] is used to dictate the length (L) of the vector (where a value of 0 is a scalar or 128-bit vector and a value of 1 is a 256-bit vector). Bits[1:0] provide opcode extensionality equivalent to some legacy prefixes (e.g., 00 = no prefix, 01 = 66H, 10 = F3H, and 11 = F2H). Bits[6:3] shown as vvvv may be used to: 1) encode the first source register operand, specified in inverted (1s complement) form and valid for instructions with 2 or more source operands; 2) encode the destination register operand, specified in 1s complement form for certain vector shifts; or 3) not encode any operand, the field is reserved and should contain a certain value, such as 1111b.

[0219] Instructions that use this prefix may use the MOD R/M R/M field 2546 to encode the instruction operand that references a memory address or encode either the destination register operand or a source register operand.

[0220] Instructions that use this prefix may use the MOD R/M reg field 2544 to encode either the destination register operand or a source register operand, or to be treated as an opcode extension and not used to encode any instruction operand.

[0221] For instruction syntax that support four operands, vvvv, the MOD R/M R/M field 2546 and the MOD R/M reg field 2544 encode three of the four operands. Bits[7:4] of the immediate value field 2409 are then used to encode the third source register operand.

[0222] FIG. 28(B) illustrates examples of a three-byte form of the second prefix 2401(B). In some examples, a format field 2811 (byte 0 2813) contains the value C4H. Byte 1 2815 includes in bits[7:5] "R," "X," and "B" which are the complements of the same values of the first prefix 2401(A). Bits[4:0] of byte 1 2815 (shown as mmmmm) include content

to encode, as need, one or more implied leading opcode bytes. For example, 00001 implies a 0FH leading opcode, 00010 implies a 0F38H leading opcode, 00011 implies a 0F3AH leading opcode, etc.

[0223] Bit[7] of byte 2 2817 is used similar to W of the first prefix 2401(A) including helping to determine promotable operand sizes. Bit[2] is used to dictate the length (L) of the vector (where a value of 0 is a scalar or 128-bit vector and a value of 1 is a 256-bit vector). Bits[1:0] provide opcode extensionality equivalent to some legacy prefixes (e.g., 00 = no prefix, 01 = 66H, 10 = F3H, and 11 = F2H). Bits[6:3], shown as vvvv, may be used to: 1) encode the first source register operand, specified in inverted (1s complement) form and valid for instructions with 2 or more source operands; 2) encode the destination register operand, specified in 1s complement form for certain vector shifts; or 3) not encode any operand, the field is reserved and should contain a certain value, such as 1111b.

[0224] Instructions that use this prefix may use the MOD R/M R/M field 2546 to encode the instruction operand that references a memory address or encode either the destination register operand or a source register operand.

[0225] Instructions that use this prefix may use the MOD R/M reg field 2544 to encode either the destination register operand or a source register operand, or to be treated as an opcode extension and not used to encode any instruction operand.

[0226] For instruction syntax that support four operands, vvvv, the MOD R/M R/M field 2546, and the MOD R/M reg field 2544 encode three of the four operands. Bits[7:4] of the immediate value field 2409 are then used to encode the third source register operand.

[0227] FIG. 29 illustrates examples of a third prefix 2401(C). In some examples, the third prefix 2401(C) is an example of an EVEX prefix. The third prefix 2401(C) is a four-byte prefix.

[0228] The third prefix 2401(C) can encode 32 vector registers (e.g., 128-bit, 256-bit, and 512-bit registers) in 64-bit mode. In some examples, instructions that utilize a writemask/opmask (see discussion of registers in a previous figure, such as FIG. 23) or predication utilize this prefix. Opmask register allow for conditional processing or selection control. Opmask instructions, whose source/destination operands are opmask registers and treat the content of an opmask register as a single value, are encoded using the second prefix 2401(B).

[0229] The third prefix 2401(C) may encode functionality that is specific to instruction classes (e.g., a packed instruction with "load+op" semantic can support embedded broadcast functionality, a floating-point instruction with rounding semantic can support static rounding functionality, a floating-point instruction with non-rounding arithmetic semantic can support "suppress all exceptions" functionality, etc.).

[0230] The first byte of the third prefix 2401(C) is a format field 2911 that has a value, in some examples, of 62H. Subsequent bytes are referred to as payload bytes 2915-2919 and collectively form a 24-bit value of P[23:0] providing specific capability in the form of one or more fields (detailed herein).

[0231] In some examples, P[1:0] of payload byte 2919 are identical to the low two mm bits. P[3:2] are reserved in some examples. Bit P[4] (R') allows access to the high 16 vector register set when combined with P[7] and the MOD R/M reg field 2544. P[6] can also provide access to a high 16 vector register when SIB-type addressing is not needed. P[7:5] consist of R, X, and B which are operand specifier modifier bits for vector register, general purpose register, memory addressing and allow access to the next set of 8 registers beyond the low 8 registers when combined with the MOD R/M register field 2544 and MOD R/M R/M field 2546. P[9:8] provide opcode extensionality equivalent to some legacy prefixes (e.g., 00 = no prefix, 01 = 66H, 10 = F3H, and 11 = F2H). P[10] in some examples is a fixed value of 1. P[14:11], shown as vvvv, may be used to: 1) encode the first source register operand, specified in inverted (1s complement) form and valid for instructions with 2 or more source operands; 2) encode the destination register operand, specified in 1s complement form for certain vector shifts; or 3) not encode any operand, the field is reserved and should contain a certain value, such as 1111b.

[0232] P[15] is similar to W of the first prefix 2401(A) and second prefix 2411(B) and may serve as an opcode extension bit or operand size promotion.

[0233] P[18:16] specify the index of a register in the opmask (writemask) registers (e.g., writemask/predicate registers 2315). In some examples, the specific value aaa = 000 has a special behavior implying no opmask is used for the particular instruction (this may be implemented in a variety of ways including the use of a opmask hardwired to all ones or hardware that bypasses the masking hardware). When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other some examples, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in some examples, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the opmask field allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While examples are described in which the opmask field's content selects one of a number of opmask registers that contains the opmask to be used (and thus the opmask field's content indirectly identifies that masking to be performed),

alternative examples instead or additional allow the mask write field's content to directly specify the masking to be performed.

[0234] P[19] can be combined with P[14:11] to encode a second source vector register in a non-destructive source syntax which can access an upper 16 vector registers using P[19]. P[20] encodes multiple functionalities, which differs across different classes of instructions and can affect the meaning of the vector length/ rounding control specifier field (P[22:21]). P[23] indicates support for merging-writemasking (e.g., when set to 0) or support for zeroing and merging-writemasking (e.g., when set to 1).

[0235] Example examples of encoding of registers in instructions using the third prefix 2401(C) are detailed in the following tables.

Table 1: 32-Register Support in 64-bit Mode

	4	3	[2:0]	REG. TYPE	COMMON USAGES
REG	R'	R	MOD R/M reg	GPR, Vector	Destination or Source
VVV	V'	vvv		GPR, Vector	2nd Source or Destination
RM	X	B	MODR/M R/M	GPR, Vector	1st Source or Destination
BASE	0	B	MODR/M R/M	GPR	Memory addressing
INDEX	0	X	SIB.index	GPR	Memory addressing
VIDX	V'	X	SIB.index	Vector	VSIB memory addressing

Table 2: Encoding Register Specifiers in 32-bit Mode

	[2:0]	REG. TYPE	COMMON USAGES
REG	MOD R/M reg	GPR, Vector	Destination or Source
VVVV	vvv	GPR, Vector	2 nd Source or Destination
RM	MODR/M R/M	GPR, Vector	1 st Source or Destination
BASE	MODR/M R/M	GPR	Memory addressing
INDEX	SIB.index	GPR	Memory addressing
VIDX	SIB.index	Vector	VSIB memory addressing

Table 3: Opmask Register Specifier Encoding

	[2:0]	REG. TYPE	COMMON USAGES
REG	MOD R/M Reg	k0-k7	Source
VVVV	vvv	k0-k7	2 nd Source
RM	MOD R/M R/M	k0-k7	1 st Source
{k1}	aaa	k0-k7	Opmask

Graphics Execution Units

[0236] FIGS. 30A-30B illustrate thread execution logic 3000 including an array of processing elements employed in a graphics processor core according to examples described herein. Elements of Figures 30A-30B having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. Figure 30A is representative of an execution unit within a general-purpose graphics processor, while Figure 30B is representative of an execution unit that may be used within a compute accelerator.

[0237] As illustrated in Figure 30A, in some examples thread execution logic 3000 includes a shader processor 3002, a thread dispatcher 3004, instruction cache 3006, a scalable execution unit array including a plurality of execution units

3008A-3008N, a sampler 3010, shared local memory 3011, a data cache 3012, and a data port 3014. In some examples the scalable execution unit array can dynamically scale by enabling or disabling one or more execution units (e.g., any of execution units 3008A, 3008B, 3008C, 3008D, through 3008N-1 and 3008N) based on the computational requirements of a workload. In some examples the included components are interconnected via an interconnect fabric that links to each of the components. In some examples, thread execution logic 3000 includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache 3006, data port 3014, sampler 3010, and execution units 3008A-3008N. In some examples, each execution unit (e.g. 3008A) is a stand-alone programmable general-purpose computational unit that is capable of executing multiple simultaneous hardware threads while processing multiple data elements in parallel for each thread. In various examples, the array of execution units 3008A-3008N is scalable to include any number individual execution units.

[0238] In some examples, the execution units 3008A-3008N are primarily used to execute shader programs. A shader processor 3002 can process the various shader programs and dispatch execution threads associated with the shader programs via a thread dispatcher 3004. In some examples the thread dispatcher includes logic to arbitrate thread initiation requests from the graphics and media pipelines and instantiate the requested threads on one or more execution unit in the execution units 3008A-3008N. For example, a geometry pipeline can dispatch vertex, tessellation, or geometry shaders to the thread execution logic for processing. In some examples, thread dispatcher 3004 can also process runtime thread spawning requests from the executing shader programs.

[0239] In some examples, the execution units 3008A-3008N support an instruction set that includes native support for many standard 3D graphics shader instructions, such that shader programs from graphics libraries (e.g., Direct 3D and OpenGL) are executed with a minimal translation. The execution units support vertex and geometry processing (e.g., vertex programs, geometry programs, vertex shaders), pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders). Each of the execution units 3008A-3008N is capable of multi-issue single instruction multiple data (SIMD) execution and multi-threaded operation enables an efficient execution environment in the face of higher latency memory accesses. Each hardware thread within each execution unit has a dedicated high-bandwidth register file and associated independent thread-state. Execution is multi-issue per clock to pipelines capable of integer, single and double precision floating point operations, SIMD branch capability, logical operations, transcendental operations, and other miscellaneous operations. While waiting for data from memory or one of the shared functions, dependency logic within the execution units 3008A-3008N causes a waiting thread to sleep until the requested data has been returned. While the waiting thread is sleeping, hardware resources may be devoted to processing other threads. For example, during a delay associated with a vertex shader operation, an execution unit can perform operations for a pixel shader, fragment shader, or another type of shader program, including a different vertex shader. Various examples can apply to use execution by use of Single Instruction Multiple Thread (SIMT) as an alternate to use of SIMD or in addition to use of SIMD. Reference to a SIMD core or operation can apply also to SIMT or apply to SIMD in combination with SIMT.

[0240] Each execution unit in execution units 3008A-3008N operates on arrays of data elements. The number of data elements is the "execution size," or the number of channels for the instruction. An execution channel is a logical unit of execution for data element access, masking, and flow control within instructions. The number of channels may be independent of the number of physical Arithmetic Logic Units (ALUs) or Floating Point Units (FPUs) for a particular graphics processor. In some examples, execution units 3008A-3008N support integer and floating-point data types.

[0241] The execution unit instruction set includes SIMD instructions. The various data elements can be stored as a packed data type in a register and the execution unit will process the various elements based on the data size of the elements. For example, when operating on a 256-bit wide vector, the 256 bits of the vector are stored in a register and the execution unit operates on the vector as four separate 64-bit packed data elements (Quad-Word (QW) size data elements), eight separate 32-bit packed data elements (Double Word (DW) size data elements), sixteen separate 16-bit packed data elements (Word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). However, different vector widths and register sizes are possible.

[0242] In some examples one or more execution units can be combined into a fused execution unit 3009A-3009N having thread control logic (3007A-3007N) that is common to the fused EUs. Multiple EUs can be fused into an EU group. Each EU in the fused EU group can be configured to execute a separate SIMD hardware thread. The number of EUs in a fused EU group can vary according to examples. Additionally, various SIMD widths can be performed per-EU, including but not limited to SIMD8, SIMD16, and SIMD32. Each fused graphics execution unit 3009A-3009N includes at least two execution units. For example, fused execution unit 3009A includes a first EU 3008A, second EU 3008B, and thread control logic 3007A that is common to the first EU 3008A and the second EU 3008B. The thread control logic 3007A controls threads executed on the fused graphics execution unit 3009A, allowing each EU within the fused execution units 3009A-3009N to execute using a common instruction pointer register.

[0243] One or more internal instruction caches (e.g., 3006) are included in the thread execution logic 3000 to cache thread instructions for the execution units. In some examples, one or more data caches (e.g., 3012) are included to cache thread data during thread execution. Threads executing on the execution logic 3000 can also store explicitly

managed data in the shared local memory 3011. In some examples, a sampler 3010 is included to provide texture sampling for 3D operations and media sampling for media operations. In some examples, sampler 3010 includes specialized texture or media sampling functionality to process texture or media data during the sampling process before providing the sampled data to an execution unit.

[0244] During execution, the graphics and media pipelines send thread initiation requests to thread execution logic 3000 via thread spawning and dispatch logic. Once a group of geometric objects has been processed and rasterized into pixel data, pixel processor logic (e.g., pixel shader logic, fragment shader logic, etc.) within the shader processor 3002 is invoked to further compute output information and cause results to be written to output surfaces (e.g., color buffers, depth buffers, stencil buffers, etc.). In some examples, a pixel shader or fragment shader calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some examples, pixel processor logic within the shader processor 3002 then executes an application programming interface (API)-supplied pixel or fragment shader program. To execute the shader program, the shader processor 3002 dispatches threads to an execution unit (e.g., 3008A) via thread dispatcher 3004. In some examples, shader processor 3002 uses texture sampling logic in the sampler 3010 to access texture data in texture maps stored in memory. Arithmetic operations on the texture data and the input geometry data compute pixel color data for each geometric fragment, or discards one or more pixels from further processing.

[0245] In some examples, the data port 3014 provides a memory access mechanism for the thread execution logic 3000 to output processed data to memory for further processing on a graphics processor output pipeline. In some examples, the data port 3014 includes or couples to one or more cache memories (e.g., data cache 3012) to cache data for memory access via the data port.

[0246] In some examples, the execution logic 3000 can also include a ray tracer 3005 that can provide ray tracing acceleration functionality. The ray tracer 3005 can support a ray tracing instruction set that includes instructions/functions for ray generation.

[0247] Figure 30B illustrates exemplary internal details of an execution unit 3008, according to examples. A graphics execution unit 3008 can include an instruction fetch unit 3037, a general register file array (GRF) 3024, an architectural register file array (ARF) 3026, a thread arbiter 3022, a send unit 3030, a branch unit 3032, a set of SIMD floating point units (FPUs) 3034, and in some examples a set of dedicated integer SIMD ALUs 3035. The GRF 3024 and ARF 3026 includes the set of general register files and architecture register files associated with each simultaneous hardware thread that may be active in the graphics execution unit 3008. In some examples, per thread architectural state is maintained in the ARF 3026, while data used during thread execution is stored in the GRF 3024. The execution state of each thread, including the instruction pointers for each thread, can be held in thread-specific registers in the ARF 3026.

[0248] In some examples the graphics execution unit 3008 has an architecture that is a combination of Simultaneous Multi-Threading (SMT) and fine-grained Interleaved Multi-Threading (IMT). The architecture has a modular configuration that can be fine-tuned at design time based on a target number of simultaneous threads and number of registers per execution unit, where execution unit resources are divided across logic used to execute multiple simultaneous threads. The number of logical threads that may be executed by the graphics execution unit 3008 is not limited to the number of hardware threads, and multiple logical threads can be assigned to each hardware thread.

[0249] In some examples, the graphics execution unit 3008 can co-issue multiple instructions, which may each be different instructions. The thread arbiter 3022 of the graphics execution unit thread 3008 can dispatch the instructions to one of the send unit 3030, branch unit 3032, or SIMD FPU(s) 3034 for execution. Each execution thread can access 128 general-purpose registers within the GRF 3024, where each register can store 32 bytes, accessible as a SIMD 8-element vector of 32-bit data elements. In some examples, each execution unit thread has access to 4 Kbytes within the GRF 3024, although examples are not so limited, and greater or fewer register resources may be provided in other examples. In some examples the graphics execution unit 3008 is partitioned into seven hardware threads that can independently perform computational operations, although the number of threads per execution unit can also vary according to examples. For example, in some examples up to 16 hardware threads are supported. In an example in which seven threads may access 4 Kbytes, the GRF 3024 can store a total of 28 Kbytes. Where 16 threads may access 4Kbytes, the GRF 3024 can store a total of 64Kbytes. Flexible addressing modes can permit registers to be addressed together to build effectively wider registers or to represent strided rectangular block data structures.

[0250] In some examples, memory operations, sampler operations, and other longer-latency system communications are dispatched via "send" instructions that are executed by the message passing send unit 3030. In some examples, branch instructions are dispatched to a dedicated branch unit 3032 to facilitate SIMD divergence and eventual convergence.

[0251] In some examples the graphics execution unit 3008 includes one or more SIMD floating point units (FPU(s)) 3034 to perform floating-point operations. In some examples, the FPU(s) 3034 also support integer computation. In some examples the FPU(s) 3034 can SIMD execute up to M number of 32-bit floating-point (or integer) operations, or SIMD execute up to 2M 16-bit integer or 16-bit floating-point operations. In some examples, at least one of the FPU(s) provides extended math capability to support high-throughput transcendental math functions and double precision 64-

bit floating-point. In some examples, a set of 8-bit integer SIMD ALUs 3035 are also present, and may be specifically optimized to perform operations associated with machine learning computations.

[0252] In some examples, arrays of multiple instances of the graphics execution unit 3008 can be instantiated in a graphics sub-core grouping (e.g., a sub-slice). For scalability, product architects can choose the exact number of execution units per sub-core grouping. In some examples the execution unit 3008 can execute instructions across a plurality of execution channels. In a further example, each thread executed on the graphics execution unit 3008 is executed on a different channel.

[0253] Figure 31 illustrates an additional execution unit 3100, according to an example. In some examples, the execution unit 3100 includes a thread control unit 3101, a thread state unit 3102, an instruction fetch/prefetch unit 3103, and an instruction decode unit 3104. The execution unit 3100 additionally includes a register file 3106 that stores registers that can be assigned to hardware threads within the execution unit. The execution unit 3100 additionally includes a send unit 3107 and a branch unit 3108. In some examples, the send unit 3107 and branch unit 3108 can operate similarly as the send unit 3030 and a branch unit 3032 of the graphics execution unit 3008 of Figure 30B.

[0254] The execution unit 3100 also includes a compute unit 3110 that includes multiple different types of functional units. In some examples the compute unit 3110 includes an ALU unit 3111 that includes an array of arithmetic logic units. The ALU unit 3111 can be configured to perform 64-bit, 32-bit, and 16-bit integer and floating point operations. Integer and floating point operations may be performed simultaneously. The compute unit 3110 can also include a systolic array 3112, and a math unit 3113. The systolic array 3112 includes a W wide and D deep network of data processing units that can be used to perform vector or other data-parallel operations in a systolic manner. In some examples the systolic array 3112 can be configured to perform matrix operations, such as matrix dot product operations. In some examples the systolic array 3112 support 16-bit floating point operations, as well as 8-bit and 4-bit integer operations. In some examples the systolic array 3112 can be configured to accelerate machine learning operations. In such examples, the systolic array 3112 can be configured with support for the bfloat 16-bit floating point format. In some examples, a math unit 3113 can be included to perform a specific subset of mathematical operations in an efficient and lower-power manner than the ALU unit 3111. The math unit 3113 can include a variant of math logic that may be found in shared function logic of a graphics processing engine provided by other examples (e.g., math logic 422 of the shared function logic 420 of Figure 4). In some examples the math unit 3113 can be configured to perform 32-bit and 64-bit floating point operations.

[0255] The thread control unit 3101 includes logic to control the execution of threads within the execution unit. The thread control unit 3101 can include thread arbitration logic to start, stop, and preempt execution of threads within the execution unit 3100. The thread state unit 3102 can be used to store thread state for threads assigned to execute on the execution unit 3100. Storing the thread state within the execution unit 3100 enables the rapid pre-emption of threads when those threads become blocked or idle. The instruction fetch/prefetch unit 3103 can fetch instructions from an instruction cache of higher level execution logic (e.g., instruction cache 3006 as in Figure 30A). The instruction fetch/prefetch unit 3103 can also issue prefetch requests for instructions to be loaded into the instruction cache based on an analysis of currently executing threads. The instruction decode unit 3104 can be used to decode instructions to be executed by the compute units. In some examples, the instruction decode unit 3104 can be used as a secondary decoder to decode complex instructions into constituent micro-operations.

[0256] The execution unit 3100 additionally includes a register file 3106 that can be used by hardware threads executing on the execution unit 3100. Registers in the register file 3106 can be divided across the logic used to execute multiple simultaneous threads within the compute unit 3110 of the execution unit 3100. The number of logical threads that may be executed by the graphics execution unit 3100 is not limited to the number of hardware threads, and multiple logical threads can be assigned to each hardware thread. The size of the register file 3106 can vary across examples based on the number of supported hardware threads. In some examples, register renaming may be used to dynamically allocate registers to hardware threads.

[0257] Figure 32 is a block diagram illustrating a graphics processor instruction formats 3200 according to some examples. In one or more example, the graphics processor execution units support an instruction set having instructions in multiple formats. The solid lined boxes illustrate the components that are generally included in an execution unit instruction, while the dashed lines include components that are optional or that are only included in a sub-set of the instructions. In some examples, instruction format 3200 described and illustrated are macro-instructions, in that they are instructions supplied to the execution unit, as opposed to micro-operations resulting from instruction decode once the instruction is processed.

[0258] In some examples, the graphics processor execution units natively support instructions in a 128-bit instruction format 3210. A 64-bit compacted instruction format 3230 is available for some instructions based on the selected instruction, instruction options, and number of operands. The native 128-bit instruction format 3210 provides access to all instruction options, while some options and operations are restricted in the 64-bit format 3230. The native instructions available in the 64-bit format 3230 vary by example. In some examples, the instruction is compacted in part using a set of index values in an index field 3213. The execution unit hardware references a set of compaction tables based on the

index values and uses the compaction table outputs to reconstruct a native instruction in the 128-bit instruction format 3210. Other sizes and formats of instruction can be used.

[0259] For each format, instruction opcode 3212 defines the operation that the execution unit is to perform. The execution units execute each instruction in parallel across the multiple data elements of each operand. For example, in response to an add instruction the execution unit performs a simultaneous add operation across each color channel representing a texture element or picture element. By default, the execution unit performs each instruction across all data channels of the operands. In some examples, instruction control field 3214 enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For instructions in the 128-bit instruction format 3210 an exec-size field 3216 limits the number of data channels that will be executed in parallel. In some examples, exec-size field 3216 is not available for use in the 64-bit compact instruction format 3230.

[0260] Some execution unit instructions have up to three operands including two source operands, src0 3220, src1 3222, and one destination 3218. In some examples, the execution units support dual destination instructions, where one of the destinations is implied. Data manipulation instructions can have a third source operand (e.g., SRC2 3224), where the instruction opcode 3212 determines the number of source operands. An instruction's last source operand can be an immediate value (e.g., hard-coded) value passed with the instruction.

[0261] In some examples, the 128-bit instruction format 3210 includes an access/address mode field 3226 specifying, for example, whether direct register addressing mode or indirect register addressing mode is used. When direct register addressing mode is used, the register address of one or more operands is directly provided by bits in the instruction.

[0262] In some examples, the 128-bit instruction format 3210 includes an access/address mode field 3226, which specifies an address mode and/or an access mode for the instruction. In some examples the access mode is used to define a data access alignment for the instruction. Some examples support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in a first mode, the instruction may use byte-aligned addressing for source and destination operands and when in a second mode, the instruction may use 16-byte-aligned addressing for all source and destination operands.

[0263] In some examples, the address mode portion of the access/address mode field 3226 determines whether the instruction is to use direct or indirect addressing. When direct register addressing mode is used bits in the instruction directly provide the register address of one or more operands. When indirect register addressing mode is used, the register address of one or more operands may be computed based on an address register value and an address immediate value field in the instruction.

[0264] In some examples instructions are grouped based on opcode 3212 bit-fields to simplify Opcode decode 3240. For an 8-bit opcode, bits 4, 5, and 6 allow the execution unit to determine the type of opcode. The precise opcode grouping shown is merely an example. In some examples, a move and logic opcode group 3242 includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some examples, move and logic group 3242 shares the five most significant bits (MSB), where move (mov) instructions are in the form of 0000xxxxb and logic instructions are in the form of 0001xxxxb. A flow control instruction group 3244 (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group 3246 includes a mix of instructions, including synchronization instructions (e.g., wait, send) in the form of 0011xxxxb (e.g., 0x30). A parallel math instruction group 3248 includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group 3248 performs the arithmetic operations in parallel across data channels. The vector math group 3250 includes arithmetic instructions (e.g., dp4) in the form of 0101xxxxb (e.g., 0x50). The vector math group performs arithmetic such as dot product calculations on vector operands. The illustrated opcode decode 3240, in some examples, can be used to determine which portion of an execution unit will be used to execute a decoded instruction. For example, some instructions may be designated as systolic instructions that will be performed by a systolic array. Other instructions, such as ray-tracing instructions (not shown) can be routed to a ray-tracing core or ray-tracing logic within a slice or partition of execution logic.

Graphics Pipeline

[0265] Figure 33 is a block diagram of another example of a graphics processor 3300. Elements of Figure 33 having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0266] In some examples, graphics processor 3300 includes a geometry pipeline 3320, a media pipeline 3330, a display engine 3340, thread execution logic 3350, and a render output pipeline 3370. In some examples, graphics processor 3300 is a graphics processor within a multicore processing system that includes one or more general-purpose processing cores. The graphics processor is controlled by register writes to one or more control and/or status registers (not shown) or via commands issued to graphics processor 3300 via a ring interconnect 3302. In some examples, ring interconnect 3302 couples graphics processor 3300 to other processing components, such as other graphics processors

or general-purpose processors. Commands from ring interconnect 3302 are interpreted by a command streamer 3303, which supplies instructions to individual components of the geometry pipeline 3320 or the media pipeline 3330.

[0267] In some examples, command streamer 3303 directs the operation of a vertex fetcher 3305 that reads vertex data from memory and executes vertex-processing commands provided by command streamer 3303. In some examples, vertex fetcher 3305 provides vertex data to a vertex shader 3307, which performs coordinate space transformation and lighting operations to each vertex. In some examples, vertex fetcher 3305 and vertex shader 3307 execute vertex-processing instructions by dispatching execution threads to execution units 3352A-3352B via a thread dispatcher 3331.

[0268] In some examples, execution units 3352A-3352B are an array of vector processors having an instruction set for performing graphics and media operations. In some examples, execution units 3352A-3352B have an attached L1 cache 3351 that is specific for each array or shared between the arrays. The cache can be configured as a data cache, an instruction cache, or a single cache that is partitioned to contain data and instructions in different partitions.

[0269] In some examples, geometry pipeline 3320 includes tessellation components to perform hardware-accelerated tessellation of 3D objects. In some examples, a programmable hull shader 3311 configures the tessellation operations. A programmable domain shader 3317 provides back-end evaluation of tessellation output. A tessellator 3313 operates at the direction of hull shader 3311 and contains special purpose logic to generate a set of detailed geometric objects based on a coarse geometric model that is provided as input to geometry pipeline 3320. In some examples, if tessellation is not used, tessellation components (e.g., hull shader 3311, tessellator 3313, and domain shader 3317) can be bypassed.

[0270] In some examples, complete geometric objects can be processed by a geometry shader 3319 via one or more threads dispatched to execution units 3352A-3352B, or can proceed directly to the clipper 3329. In some examples, the geometry shader operates on entire geometric objects, rather than vertices or patches of vertices as in previous stages of the graphics pipeline. If the tessellation is disabled the geometry shader 3319 receives input from the vertex shader 3307. In some examples, geometry shader 3319 is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

[0271] Before rasterization, a clipper 3329 processes vertex data. The clipper 3329 may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some examples, a rasterizer and depth test component 3373 in the render output pipeline 3370 dispatches pixel shaders to convert the geometric objects into per pixel representations. In some examples, pixel shader logic is included in thread execution logic 3350. In some examples, an application can bypass the rasterizer and depth test component 3373 and access un-rasterized vertex data via a stream out unit 3323.

[0272] The graphics processor 3300 has an interconnect bus, interconnect fabric, or some other interconnect mechanism that allows data and message passing amongst the major components of the processor. In some examples, execution units 3352A-3352B and associated logic units (e.g., L1 cache 3351, sampler 3354, texture cache 3358, etc.) interconnect via a data port 3356 to perform memory access and communicate with render output pipeline components of the processor. In some examples, sampler 3354, caches 3351, 3358 and execution units 3352A-3352B each have separate memory access paths. In some examples the texture cache 3358 can also be configured as a sampler cache.

[0273] In some examples, render output pipeline 3370 contains a rasterizer and depth test component 3373 that converts vertex-based objects into an associated pixel-based representation. In some examples, the rasterizer logic includes a windower/masker unit to perform fixed function triangle and line rasterization. An associated render cache 3378 and depth cache 3379 are also available in some examples. A pixel operations component 3377 performs pixel-based operations on the data, though in some instances, pixel operations associated with 2D operations (e.g. bit block image transfers with blending) are performed by the 2D engine 3341, or substituted at display time by the display controller 3343 using overlay display planes. In some examples, a shared L3 cache 3375 is available to all graphics components, allowing the sharing of data without the use of main system memory.

[0274] In some examples, graphics processor media pipeline 3330 includes a media engine 3337 and a video front-end 3334. In some examples, video front-end 3334 receives pipeline commands from the command streamer 3303. In some examples, media pipeline 3330 includes a separate command streamer. In some examples, video front-end 3334 processes media commands before sending the command to the media engine 3337. In some examples, media engine 3337 includes thread spawning functionality to spawn threads for dispatch to thread execution logic 3350 via thread dispatcher 3331.

[0275] In some examples, graphics processor 3300 includes a display engine 3340. In some examples, display engine 3340 is external to processor 3300 and couples with the graphics processor via the ring interconnect 3302, or some other interconnect bus or fabric. In some examples, display engine 3340 includes a 2D engine 3341 and a display controller 3343. In some examples, display engine 3340 contains special purpose logic capable of operating independently of the 3D pipeline. In some examples, display controller 3343 couples with a display device (not shown), which may be a system integrated display device, as in a laptop computer, or an external display device attached via a display device connector.

[0276] In some examples, the geometry pipeline 3320 and media pipeline 3330 are configurable to perform operations based on multiple graphics and media programming interfaces and are not specific to any one application programming

interface (API). In some examples, driver software for the graphics processor translates API calls that are specific to a particular graphics or media library into commands that can be processed by the graphics processor. In some examples, support is provided for the Open Graphics Library (OpenGL), Open Computing Language (OpenCL), and/or Vulkan graphics and compute API, all from the Khronos Group. In some examples, support may also be provided for the Direct3D library from the Microsoft Corporation. In some examples, a combination of these libraries may be supported. Support may also be provided for the Open Source Computer Vision Library (OpenCV). A future API with a compatible 3D pipeline would also be supported if a mapping can be made from the pipeline of the future API to the pipeline of the graphics processor.

Graphics Pipeline Programming

[0277] Figure 34A is a block diagram illustrating a graphics processor command format 3400 according to some examples. Figure 34B is a block diagram illustrating a graphics processor command sequence 3410 according to an example. The solid lined boxes in Figure 34A illustrate the components that are generally included in a graphics command while the dashed lines include components that are optional or that are only included in a sub-set of the graphics commands. The exemplary graphics processor command format 3400 of Figure 34A includes data fields to identify a client 3402, a command operation code (opcode) 3404, and data 3406 for the command. A sub-opcode 3405 and a command size 3408 are also included in some commands.

[0278] In some examples, client 3402 specifies the client unit of the graphics device that processes the command data. In some examples, a graphics processor command parser examines the client field of each command to condition the further processing of the command and route the command data to the appropriate client unit. In some examples, the graphics processor client units include a memory interface unit, a render unit, a 2D unit, a 3D unit, and a media unit. Each client unit has a corresponding processing pipeline that processes the commands. Once the command is received by the client unit, the client unit reads the opcode 3404 and, if present, sub-opcode 3405 to determine the operation to perform. The client unit performs the command using information in data field 3406. For some commands an explicit command size 3408 is expected to specify the size of the command. In some examples, the command parser automatically determines the size of at least some of the commands based on the command opcode. In some examples commands are aligned via multiples of a double word. Other command formats can be used.

[0279] The flow diagram in Figure 34B illustrates an exemplary graphics processor command sequence 3410. In some examples, software or firmware of a data processing system that features an example of a graphics processor uses a version of the command sequence shown to set up, execute, and terminate a set of graphics operations. A sample command sequence is shown and described for purposes of example only as examples are not limited to these specific commands or to this command sequence. Moreover, the commands may be issued as batch of commands in a command sequence, such that the graphics processor will process the sequence of commands in at least partially concurrence.

[0280] In some examples, the graphics processor command sequence 3410 may begin with a pipeline flush command 3412 to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some examples, the 3D pipeline 3422 and the media pipeline 3424 do not operate concurrently. The pipeline flush is performed to cause the active graphics pipeline to complete any pending commands. In response to a pipeline flush, the command parser for the graphics processor will pause command processing until the active drawing engines complete pending operations and the relevant read caches are invalidated. Optionally, any data in the render cache that is marked 'dirty' can be flushed to memory. In some examples, pipeline flush command 3412 can be used for pipeline synchronization or before placing the graphics processor into a low power state.

[0281] In some examples, a pipeline select command 3413 is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some examples, a pipeline select command 3413 is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some examples, a pipeline flush command 3412 is required immediately before a pipeline switch via the pipeline select command 3413.

[0282] In some examples, a pipeline control command 3414 configures a graphics pipeline for operation and is used to program the 3D pipeline 3422 and the media pipeline 3424. In some examples, pipeline control command 3414 configures the pipeline state for the active pipeline. In some examples, the pipeline control command 3414 is used for pipeline synchronization and to clear data from one or more cache memories within the active pipeline before processing a batch of commands.

[0283] In some examples, return buffer state commands 3416 are used to configure a set of return buffers for the respective pipelines to write data. Some pipeline operations require the allocation, selection, or configuration of one or more return buffers into which the operations write intermediate data during processing. In some examples, the graphics processor also uses one or more return buffers to store output data and to perform cross thread communication. In some examples, the return buffer state 3416 includes selecting the size and number of return buffers to use for a set of pipeline operations.

[0284] The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination 3420, the command sequence is tailored to the 3D pipeline 3422 beginning with the 3D pipeline state 3430 or the media pipeline 3424 beginning at the media pipeline state 3440.

[0285] The commands to configure the 3D pipeline state 3430 include 3D state setting commands for vertex buffer state, vertex element state, constant color state, depth buffer state, and other state variables that are to be configured before 3D primitive commands are processed. The values of these commands are determined at least in part based on the particular 3D API in use. In some examples, 3D pipeline state 3430 commands are also able to selectively disable or bypass certain pipeline elements if those elements will not be used.

[0286] In some examples, 3D primitive 3432 command is used to submit 3D primitives to be processed by the 3D pipeline. Commands and associated parameters that are passed to the graphics processor via the 3D primitive 3432 command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D primitive 3432 command data to generate vertex data structures. The vertex data structures are stored in one or more return buffers. In some examples, 3D primitive 3432 command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline 3422 dispatches shader execution threads to graphics processor execution units.

[0287] In some examples, 3D pipeline 3422 is triggered via an execute 3434 command or event. In some examples, a register write triggers command execution. In some examples execution is triggered via a 'go' or 'kick' command in the command sequence. In some examples, command execution is triggered using a pipeline synchronization command to flush the command sequence through the graphics pipeline. The 3D pipeline will perform geometry processing for the 3D primitives. Once operations are complete, the resulting geometric objects are rasterized and the pixel engine colors the resulting pixels. Additional commands to control pixel shading and pixel back end operations may also be included for those operations.

[0288] In some examples, the graphics processor command sequence 3410 follows the media pipeline 3424 path when performing media operations. In general, the specific use and manner of programming for the media pipeline 3424 depends on the media or compute operations to be performed. Specific media decode operations may be offloaded to the media pipeline during media decode. In some examples, the media pipeline can also be bypassed and media decode can be performed in whole or in part using resources provided by one or more general-purpose processing cores. In some examples, the media pipeline also includes elements for general-purpose graphics processor unit (GPGPU) operations, where the graphics processor is used to perform SIMD vector operations using computational shader programs that are not explicitly related to the rendering of graphics primitives.

[0289] In some examples, media pipeline 3424 is configured in a similar manner as the 3D pipeline 3422. A set of commands to configure the media pipeline state 3440 are dispatched or placed into a command queue before the media object commands 3442. In some examples, commands for the media pipeline state 3440 include data to configure the media pipeline elements that will be used to process the media objects. This includes data to configure the video decode and video encode logic within the media pipeline, such as encode or decode format. In some examples, commands for the media pipeline state 3440 also support the use of one or more pointers to "indirect" state elements that contain a batch of state settings.

[0290] In some examples, media object commands 3442 supply pointers to media objects for processing by the media pipeline. The media objects include memory buffers containing video data to be processed. In some examples, all media pipeline states must be valid before issuing a media object command 3442. Once the pipeline state is configured and media object commands 3442 are queued, the media pipeline 3424 is triggered via an execute command 3444 or an equivalent execute event (e.g., register write). Output from media pipeline 3424 may then be post processed by operations provided by the 3D pipeline 3422 or the media pipeline 3424. In some examples, GPGPU operations are configured and executed in a similar manner as media operations.

[0291] Program code may be applied to input information 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 includes any system that has a processor, such as, for example, a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a microprocessor, or any combination thereof.

[0292] 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.

[0293] Examples of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Examples 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.

[0294] 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), phase change memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[0295] Accordingly, examples 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 examples may also be referred to as program products.

Emulation (including binary translation, code morphing, etc.).

[0296] In some cases, an instruction converter may be used to convert an instruction from a source instruction set architecture to a target instruction set architecture. 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.

[0297] FIG. 35 is a block diagram illustrating the use of a software instruction converter to convert binary instructions in a source ISA to binary instructions in a target ISA according to examples. In the illustrated example, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. FIG. 35 shows a program in a high-level language 3502 may be compiled using a first ISA compiler 3504 to generate first ISA binary code 3506 that may be natively executed by a processor with at least one first ISA core 3516. The processor with at least one first ISA core 3516 represents any processor that can perform substantially the same functions as an Intel® processor with at least one first ISA core by compatibly executing or otherwise processing (1) a substantial portion of the first ISA or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one first ISA core, in order to achieve substantially the same result as a processor with at least one first ISA core. The first ISA compiler 3504 represents a compiler that is operable to generate first ISA binary code 3506 (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one first ISA core 3516. Similarly, FIG. 35 shows the program in the high-level language 3502 may be compiled using an alternative ISA compiler 3508 to generate alternative ISA binary code 3510 that may be natively executed by a processor without a first ISA core 3514. The instruction converter 3512 is used to convert the first ISA binary code 3506 into code that may be natively executed by the processor without a first ISA core 3514. This converted code is not necessarily to be the same as the alternative ISA binary code 3510; however, the converted code will accomplish the general operation and be made up of instructions from the alternative ISA. Thus, the instruction converter 3512 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 a first ISA processor or core to execute the first ISA binary code 3506.

IP Core Implementations

[0298] One or more aspects of at least some examples may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as "IP cores," are reusable units of logic for an integrated circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated such that the circuit performs operations described in association with any of the examples described herein.

[0299] FIG. 36 is a block diagram illustrating an IP core development system 3600 that may be used to manufacture an integrated circuit to perform operations according to some examples. The IP core development system 3600 may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SOC integrated circuit). A design facility 3630 can generate a software simulation 3610 of an IP core design in a high-level programming language (e.g., C/C++). The software simulation 3610 can be used to design, test, and verify the behavior of the IP core using a simulation model 3612. The simulation model 3612 may include functional, behavioral, and/or timing simulations. A register transfer level (RTL) design 3615 can then be created

or synthesized from the simulation model 3612. The RTL design 3615 is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design 3615, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

[0300] The RTL design 3615 or equivalent may be further synthesized by the design facility into a hardware model 3620, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3rd party fabrication facility 3665 using non-volatile memory 3640 (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection 3650 or wireless connection 3660. The fabrication facility 3665 may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least some examples described herein.

[0301] References to "some examples," "an example," etc., indicate that the example described may include a particular feature, structure, or characteristic, but every example may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same example. Further, when a particular feature, structure, or characteristic is described in connection with an example, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other examples whether or not explicitly described.

[0302] Examples include, but are not limited to:

Example 1. An apparatus such as a processor, SoC, core, accelerator, etc. that at least includes decoder circuitry to decode an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction; and the execution circuitry to execute the decoded instance of the single instruction according to the opcode using the immediate value.

Example 2. The apparatus of example 1, wherein one of the one or more operands is a general purpose register.

Example 3. The apparatus of example 1, wherein one of the one or more operands is a vector register.

Example 4. The apparatus of example 1, wherein one of the one or more operands is memory.

Example 5. The apparatus of any of examples 1-4, wherein the prefix is a 3 byte prefix.

Example 6. The apparatus of example 5, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

Example 7. The apparatus of any of examples 1-6, wherein the decoder is to treat 32 bits following addressing information in a format of the instance of the single instruction as the immediate value.

Example 8. The apparatus of any of examples 1-7, wherein the immediate value is a 32-bit value.

Example 9. A method of processing using an immediate value comprising: decoding an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction; and executing the decoded instance of the single instruction according to the opcode using the immediate value.

Example 10. The method of example 9, wherein one of the one or more operands is a general purpose register.

Example 11. The method of example 9, wherein one of the one or more operands is a vector register.

Example 12. The method of example 9, wherein one of the one or more operands is memory.

Example 13. The method of any of examples 9-12, wherein the prefix is a 3 byte prefix.

Example 14. The method of example 13, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

5 Example 15. The method of any of examples 9-14, wherein the decoder is to treat 32 bits following addressing information in a format of the instance of the single instruction as the immediate value.

Example 16. The method of any of examples 9-15, wherein the immediate value is a 32-bit value

10 Example 17. A system such as a computer, SoC, hardware device, etc. that at least comprises memory to store an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction; decoder circuitry to decode the instance of the single instruction; and the execution circuitry to execute the decoded instance of the single instruction according to the opcode.

15 Example 18. The system of example 17, wherein the prefix is a 3 byte prefix.

20 Example 19. The system of example 18, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

Example 20. The system of any of examples 17-19, wherein the decoder is to treat 32 bits following addressing information in a format of the instance of the single instruction as the immediate value.

25 Example 21. An apparatus such as a processor, SoC, core, accelerator, etc. that at least includes decoder circuitry to decode an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value, wherein the immediate value is to be used to index a programmable logic array to select at least one operation to be performed; and the execution circuitry to execute the decoded instance of the single instruction according to the opcode.

30 Example 22. The apparatus of example 21, wherein the selected at least one operation is performed outside of the execution circuitry.

35 Example 23. The apparatus of any of examples 21-22, wherein one of the one or more operands is a vector register.

40 Example 24. The apparatus of any of examples 21-23, wherein the selected at least one operation is performed by the execution circuitry.

Example 25. The apparatus of any of examples 21-24, wherein the prefix is a 3 byte prefix.

45 Example 26. The apparatus of example 25, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

Example 27. The apparatus of any of examples 21-26, wherein the decoder is to treat 32 bits following addressing information as the immediate value.

50 Example 28. The apparatus of any of examples 21-27, wherein the immediate value is a 32-bit value

Example 29. A method comprising:

55 decoding an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction, wherein the immediate value is to be used to index a programmable logic array to select at least one operation to be performed; and

executing the decoded instance of the single instruction according to the opcode.

Example 30. The method of example 29, wherein the selected at least one operation is performed outside of execution circuitry.

Example 31. The method of any of examples 29-30, wherein one of the one or more operands is a vector register.

Example 32. The method of example 29, wherein the selected at least one operation is performed by execution circuitry.

Example 33. The method of any of examples 29-32, wherein the prefix is a 3 byte prefix.

Example 34. The method of example 33, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

Example 35. The method of any of examples 29-34, wherein the decoder is to treat 32 bits following addressing information in a format of the instance of the single instruction as the immediate value.

Example 36. The method of any of examples 29-35, wherein the immediate value is a 32-bit value.

Example 37. A system such as a computer, SoC, hardware device, etc. that at least comprises memory to store an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction, wherein the immediate value is to be used to index a programmable logic array to select at least one operation to be performed; decoder circuitry to decode the instance of the single instruction; and execution circuitry to execute the decoded instance of the single instruction according to the opcode.

Example 38. The system of example 37, wherein the prefix is a 3 byte prefix.

Example 39. The system of example 38, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.

Example 40. The system of any of examples 37-39, wherein the selected at least one operation is performed outside of the execution circuitry.

[0303] Moreover, in the various examples described above, unless specifically noted otherwise, disjunctive language such as the phrase "at least one of A, B, or C" or "A, B, and/or C" is intended to be understood to mean either A, B, or C, or any combination thereof (i.e. A and B, A and C, B and C, and A, B and C).

[0304] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the disclosure as set forth in the claims.

Claims

1. An apparatus comprising:

decoder means to decode an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed by execution means using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction; and
 the execution means to execute the decoded instance of the single instruction according to the opcode using the immediate value.

2. The apparatus of claim 1, wherein one of the one or more operands is a general purpose register.

3. The apparatus of claim 1, wherein one of the one or more operands is a vector register.
4. The apparatus of claim 1, wherein one of the one or more operands is memory.
- 5 5. The apparatus of any of claims 1-4, wherein the prefix is a 3 byte prefix.
6. The apparatus of claim 5, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.
- 10 7. The apparatus of any of claims 1-6, wherein the decoder means is to treat 32 bits following addressing information in a format of the instance of the single instruction as the immediate value.
8. The apparatus of any of claims 1-6, wherein the immediate value is a 32-bit value.
- 15 9. A method comprising:
- decoding an instance of a single instruction, the instance of the single instruction is to include one or more fields for an opcode, one or more fields for one or more operands, one or more fields for a prefix, and a field for an immediate value, wherein the opcode is to indicate one or more operations to be performed using the operands and the prefix is to indicate support for use of the immediate value during execution of the instance of the single instruction; and
- 20 executing the decoded instance of the single instruction according to the opcode using the immediate value.
10. The method of claim 9, wherein one of the one or more operands is a general purpose register.
- 25 11. The method of claim 9, wherein one of the one or more operands is a vector register.
12. The method of claim 9, wherein one of the one or more operands is memory.
- 30 13. The method of any of claims 9-12, wherein the prefix is a 3 byte prefix.
14. The method of claim 13, wherein the bits 0-4 of a second byte of the prefix are to indicate usage of the immediate value.
- 35 15. The method of any of claims 9-14, wherein the immediate value is a 32-bit value.

40

45

50

55

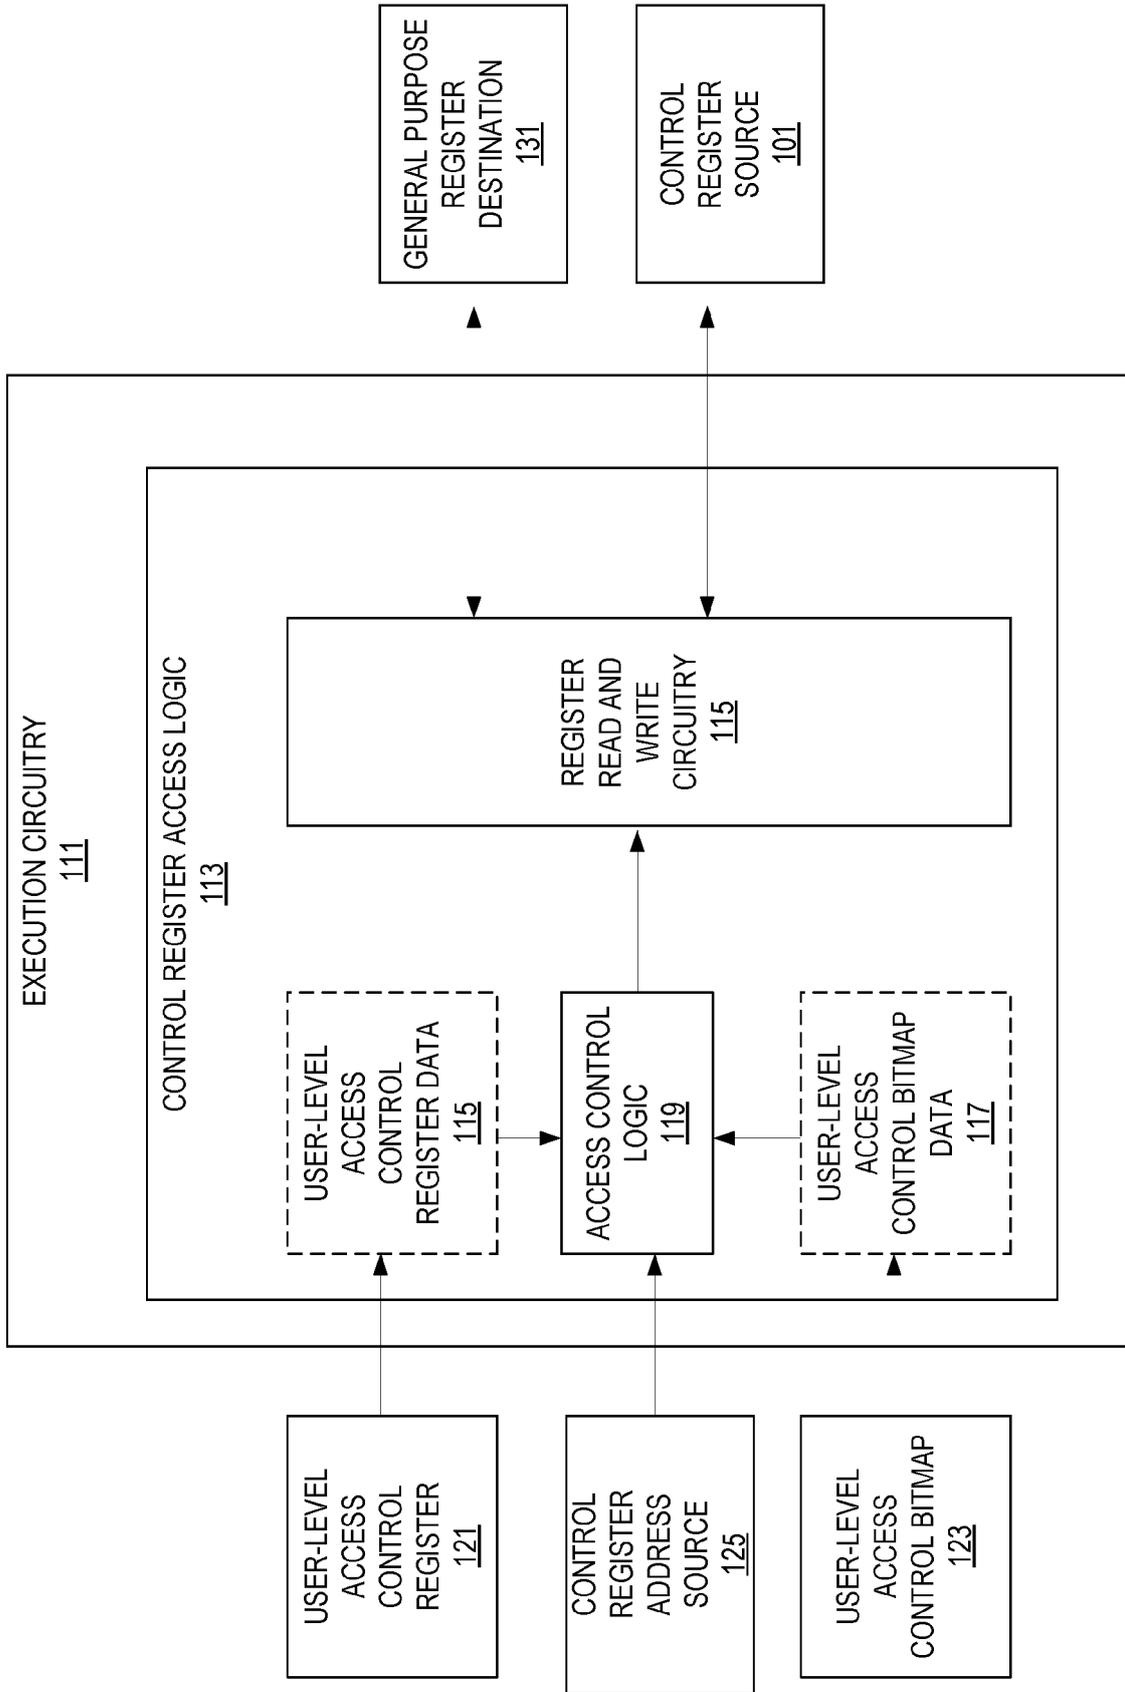


FIG. 1

<u>OPCODE</u>	<u>INSTRUCTION</u>	<u>OP/EN</u>	<u>64/32 BIT MODE</u>	<u>CPUID</u>	<u>DESCRIPTION</u>
F3 0F 38 F8 11:RRR:BBB	UWRMSR R64, R64	RN	V/N.E.	07.01H:EDX[D15] USER_MSR	WRITE THE VALUE FROM REGISTER BBB INTO THE MSR ADDRESS SPECIFIED IN RRR
VEX.128.F3.MAP7.W0.F8 11:000:BBB	UWRMSR IMM32, R64	IM	V/N.E.	07.01H:EDX[D15] USER_MSR	WRITE THE VALUE FROM REGISTER BBB INTO THE MSR ADDRESS SPECIFIED IN THE 32-BIT IMMEDIATE

<u>OP/EN</u>	<u>OPERAND 1</u>	<u>OPERAND 2</u>	<u>OPERAND 3</u>	<u>OPERAND 4</u>
RM	MODRM:REG (R)	MODRM:R/M (R)	NA	NA
IM	IM32	MODRM:R/M (R)	NA	NA

FIG. 2

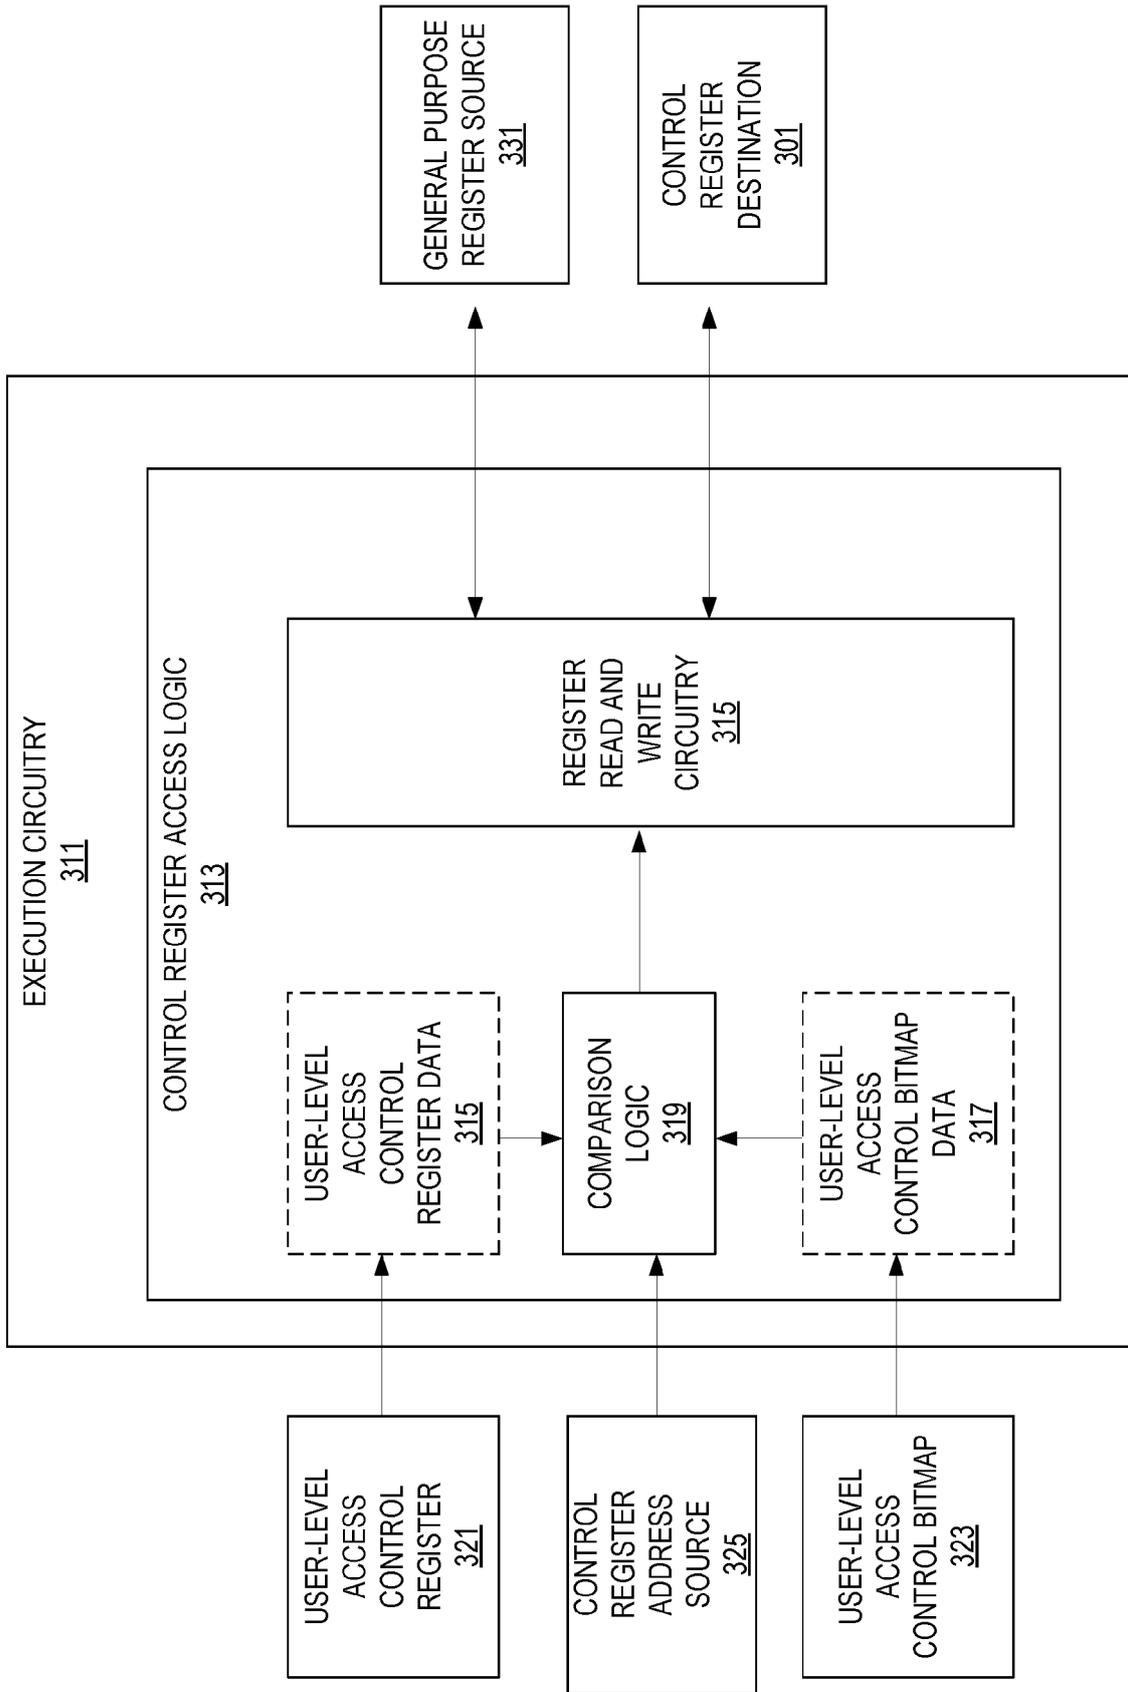


FIG. 3

<u>OPCODE</u>	<u>INSTRUCTION</u>	<u>OP/EN</u>	<u>64/32 BIT MODE</u>	<u>CPUID</u>	<u>DESCRIPTION</u>
F2 0F 38 F8 11:RRR:BBB	URDMSR R64, R64	MR	V/N.E.	07.01H:EDX[D15] USER_MSR	READ THE MSR ADDRESS SPECIFIED IN RRR AND WRITE THE VALUE INTO THE REGISTER SPECIFIED IN BBB
VEX.128.F2.MAP7.W0.F8 11:000:BBB	URDMSR R64, IMM32	MI	V/N.E.	07.01H:EDX[D15] USER_MSR	READ THE MSR ADDRESS SPECIFIED IN THE 32-BIT IMMEDIATE AND WRITE THE VALUE INTO THE REGISTER SPECIFIED IN BBB

<u>OP/EN</u>	<u>OPERAND 1</u>	<u>OPERAND 2</u>	<u>OPERAND 3</u>	<u>OPERAND 4</u>
MR	MODRM:R/M (R)	MODRM:REG (R)	NA	NA
MI	MODRM:R/M (R)	IMM32	NA	NA

FIG. 4

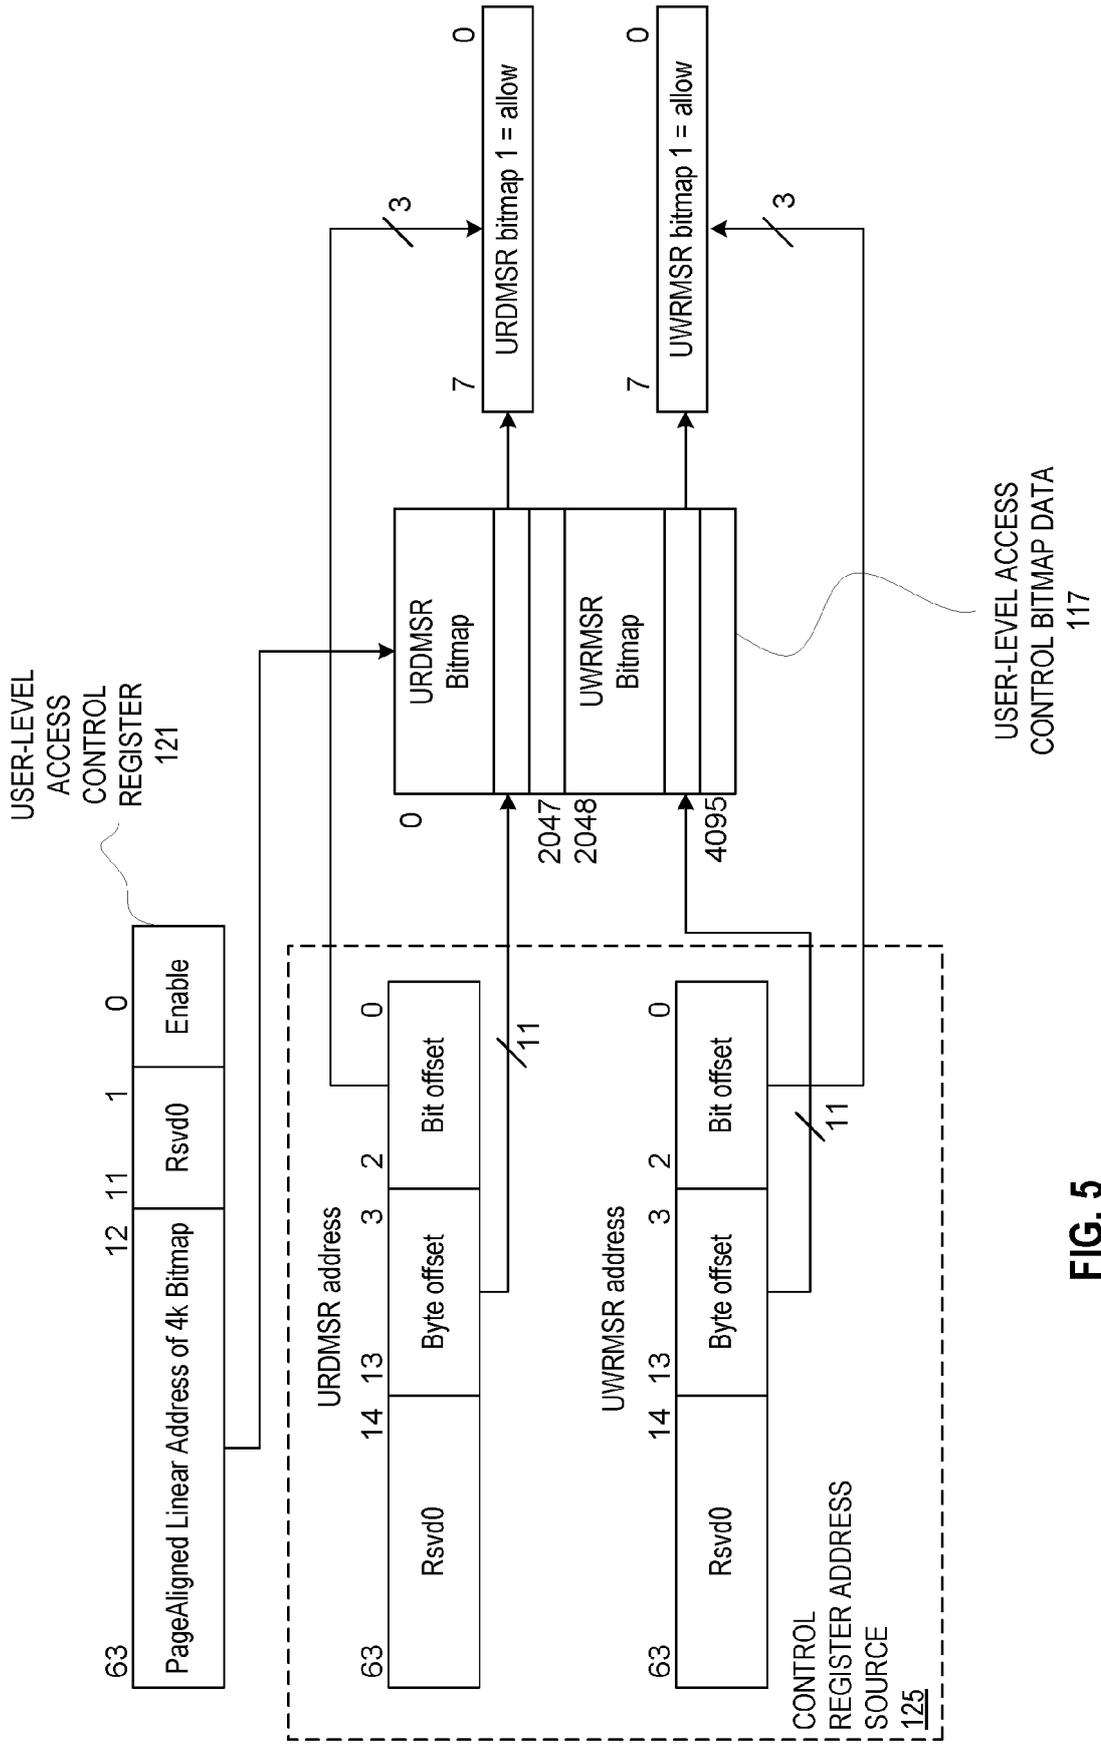


FIG. 5

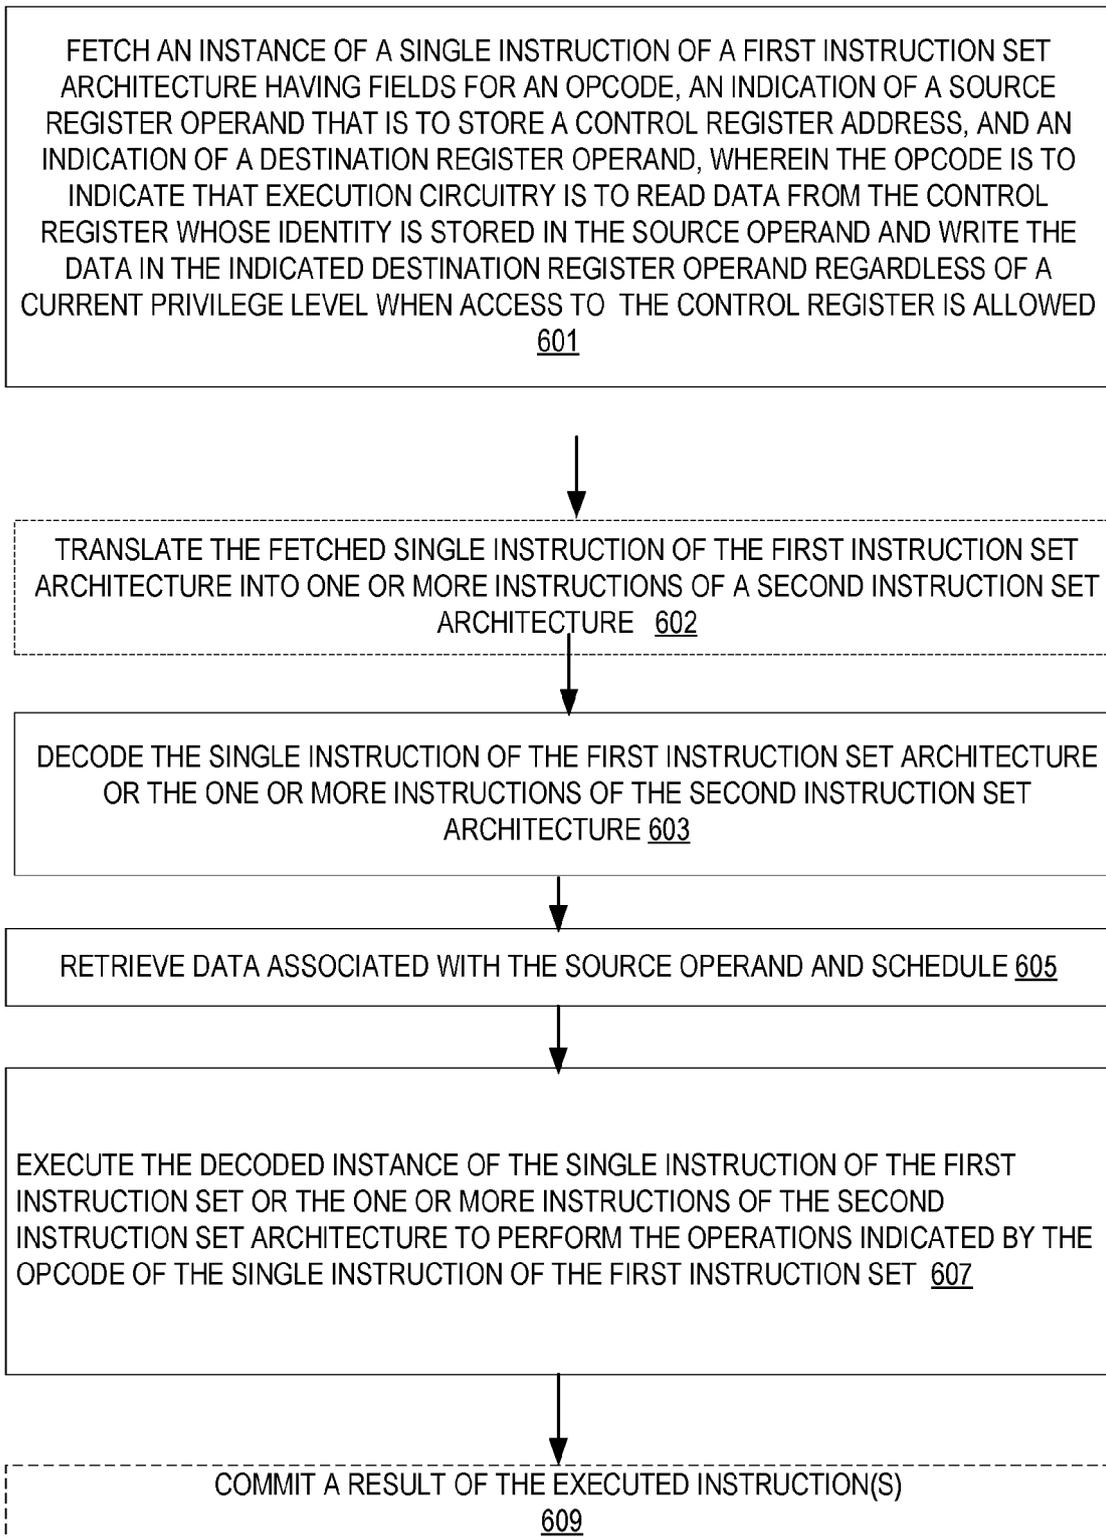


FIG. 6

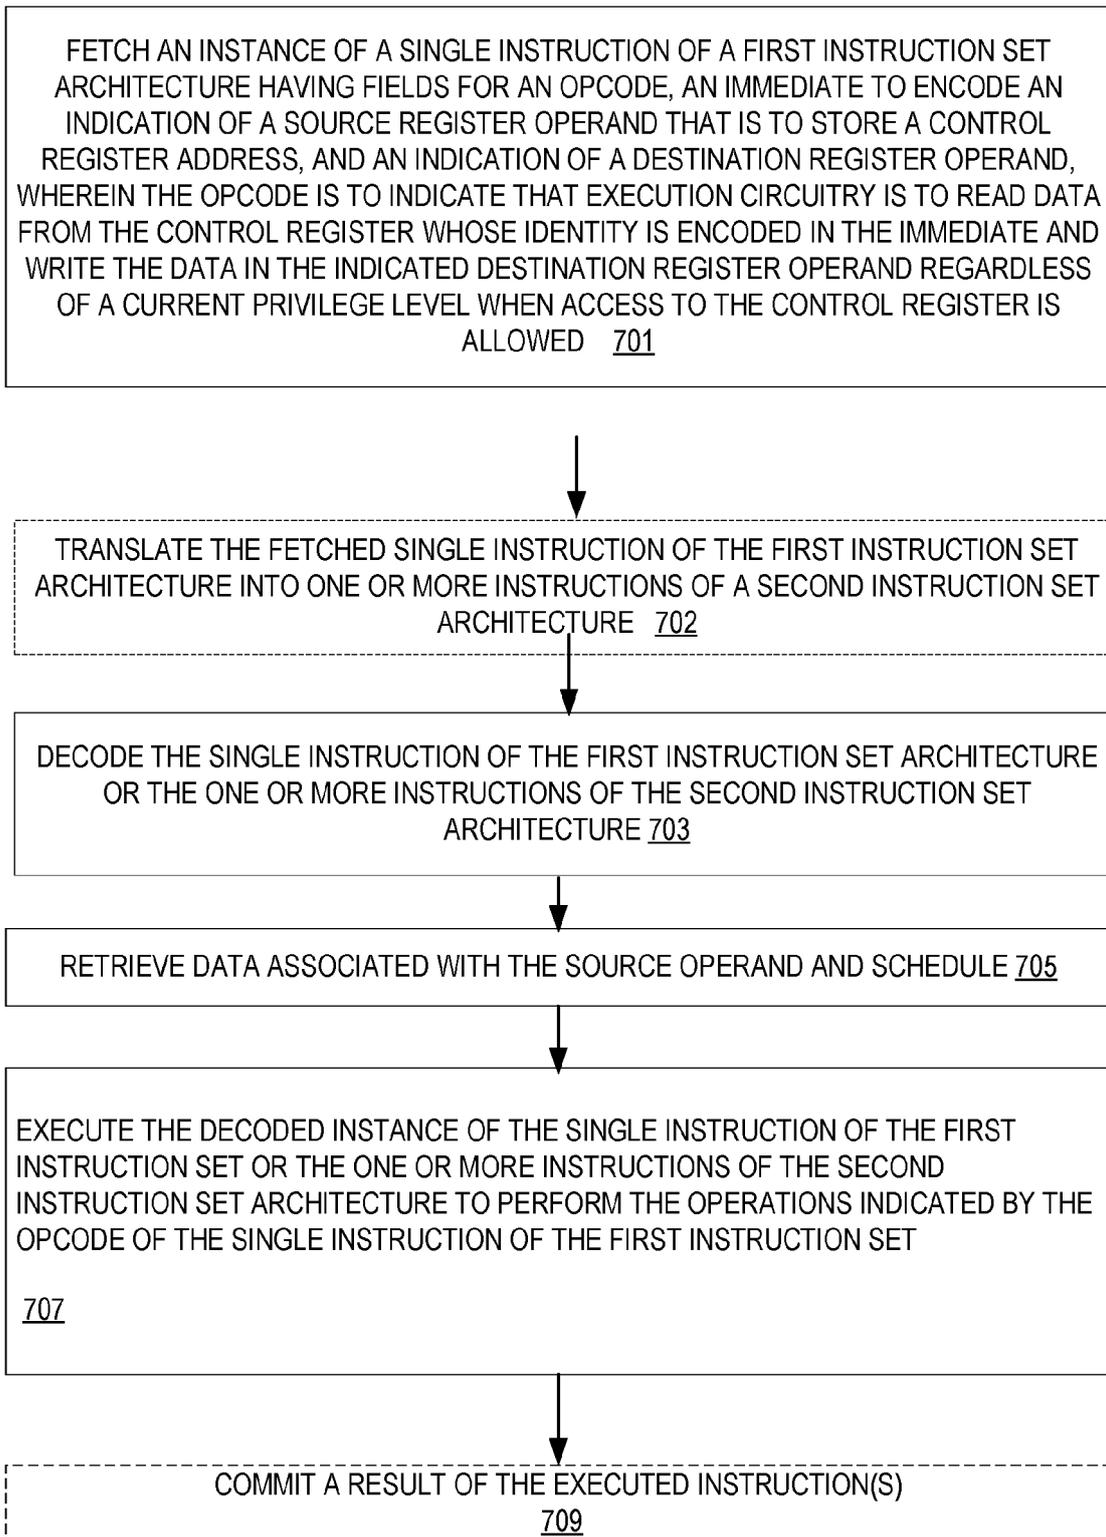


FIG. 7

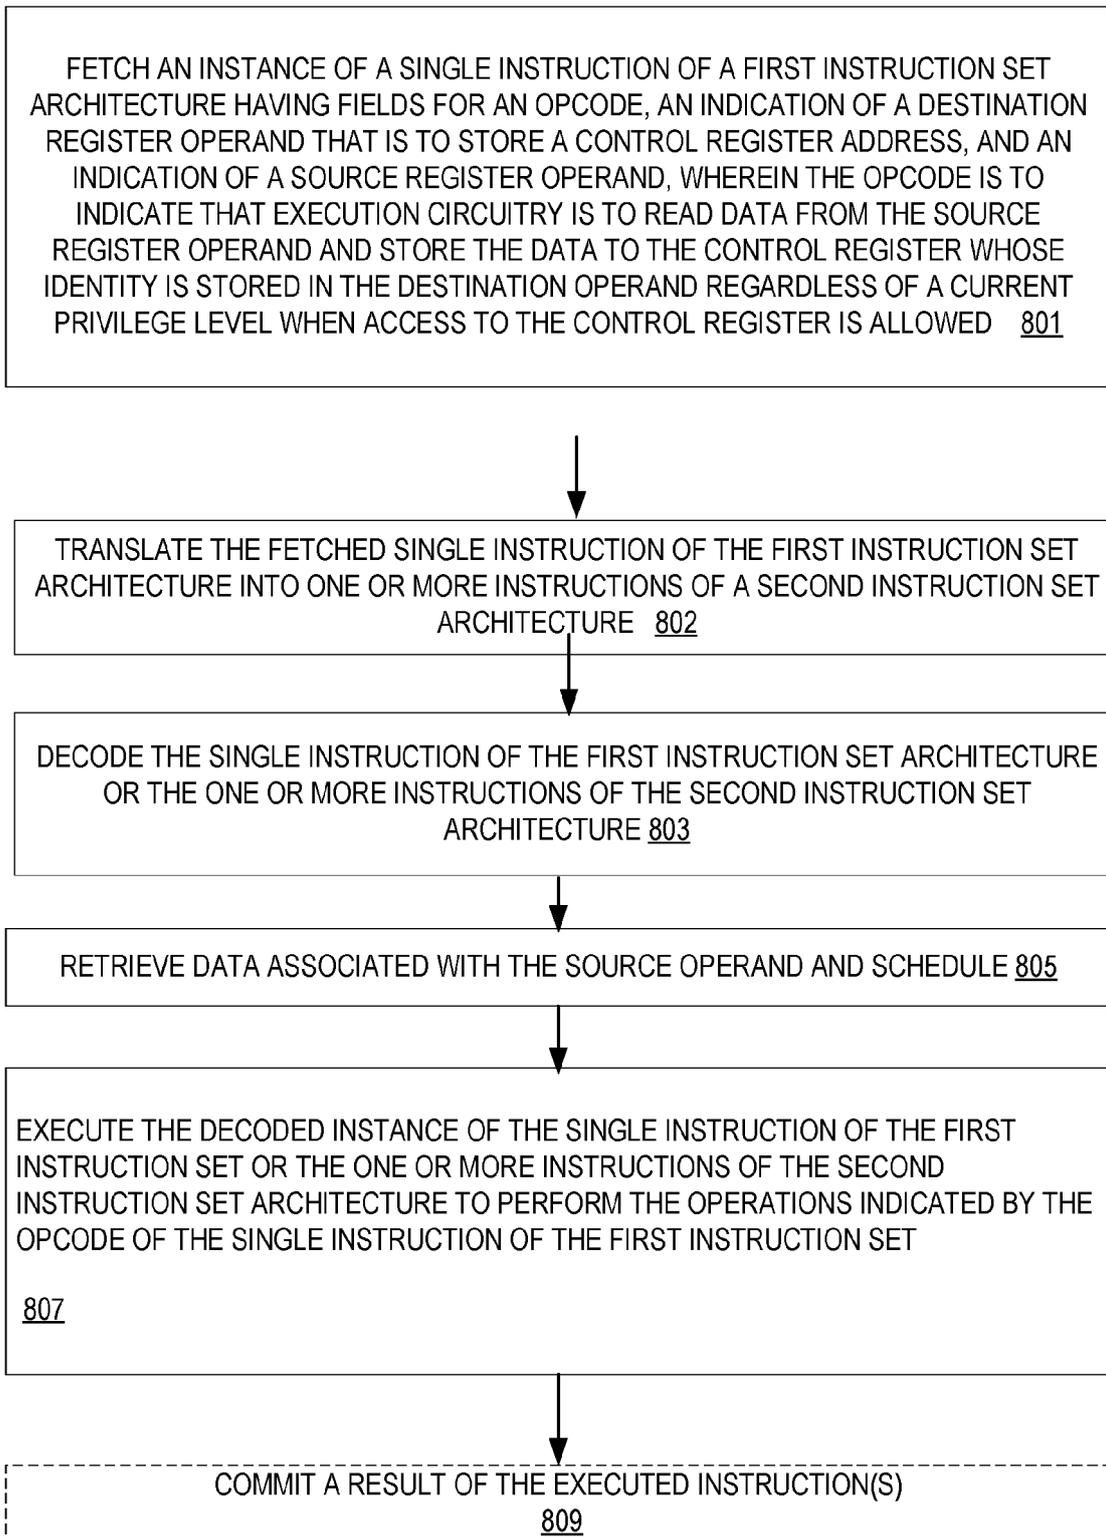


FIG. 8

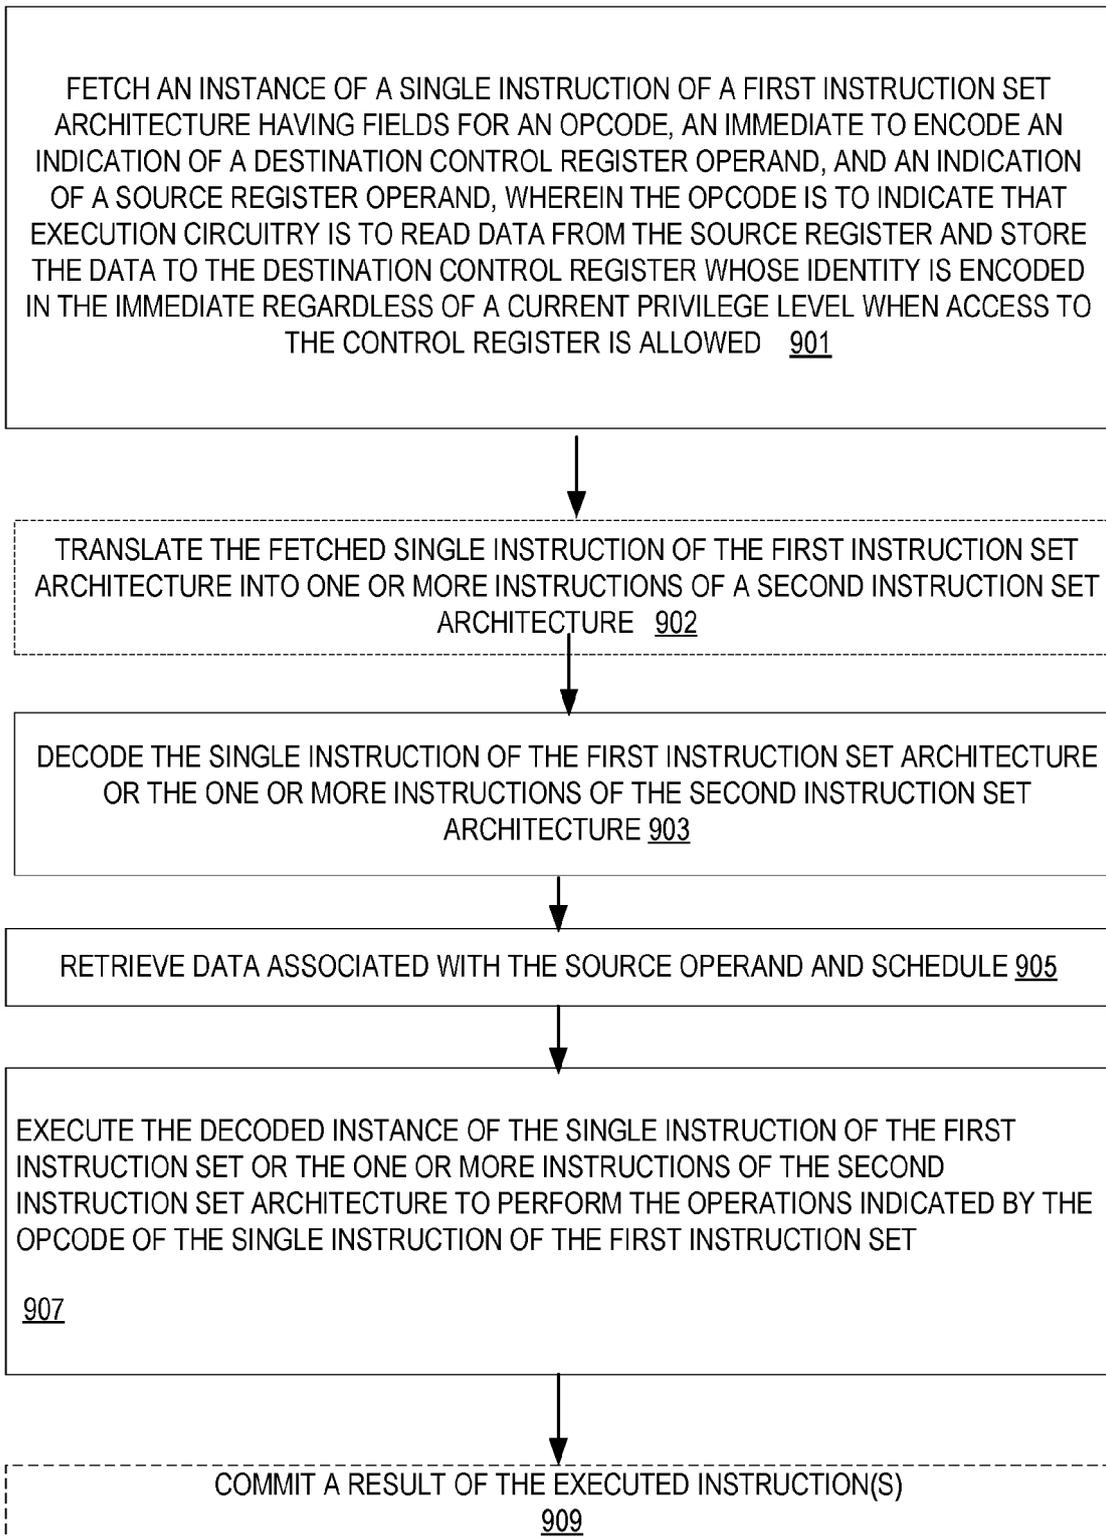


FIG. 9

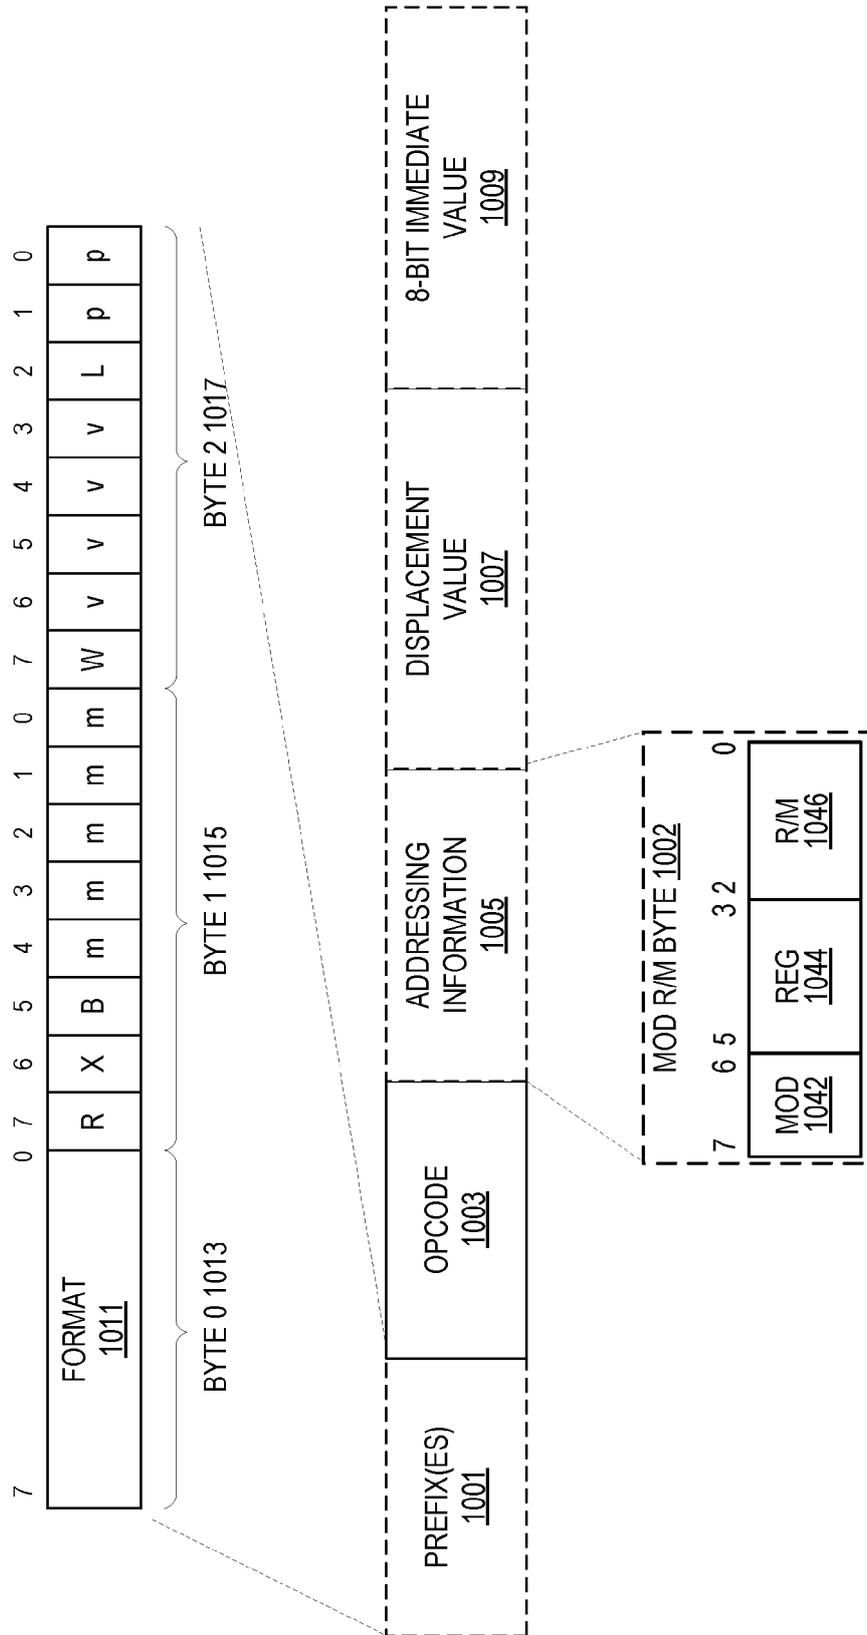


FIG. 10

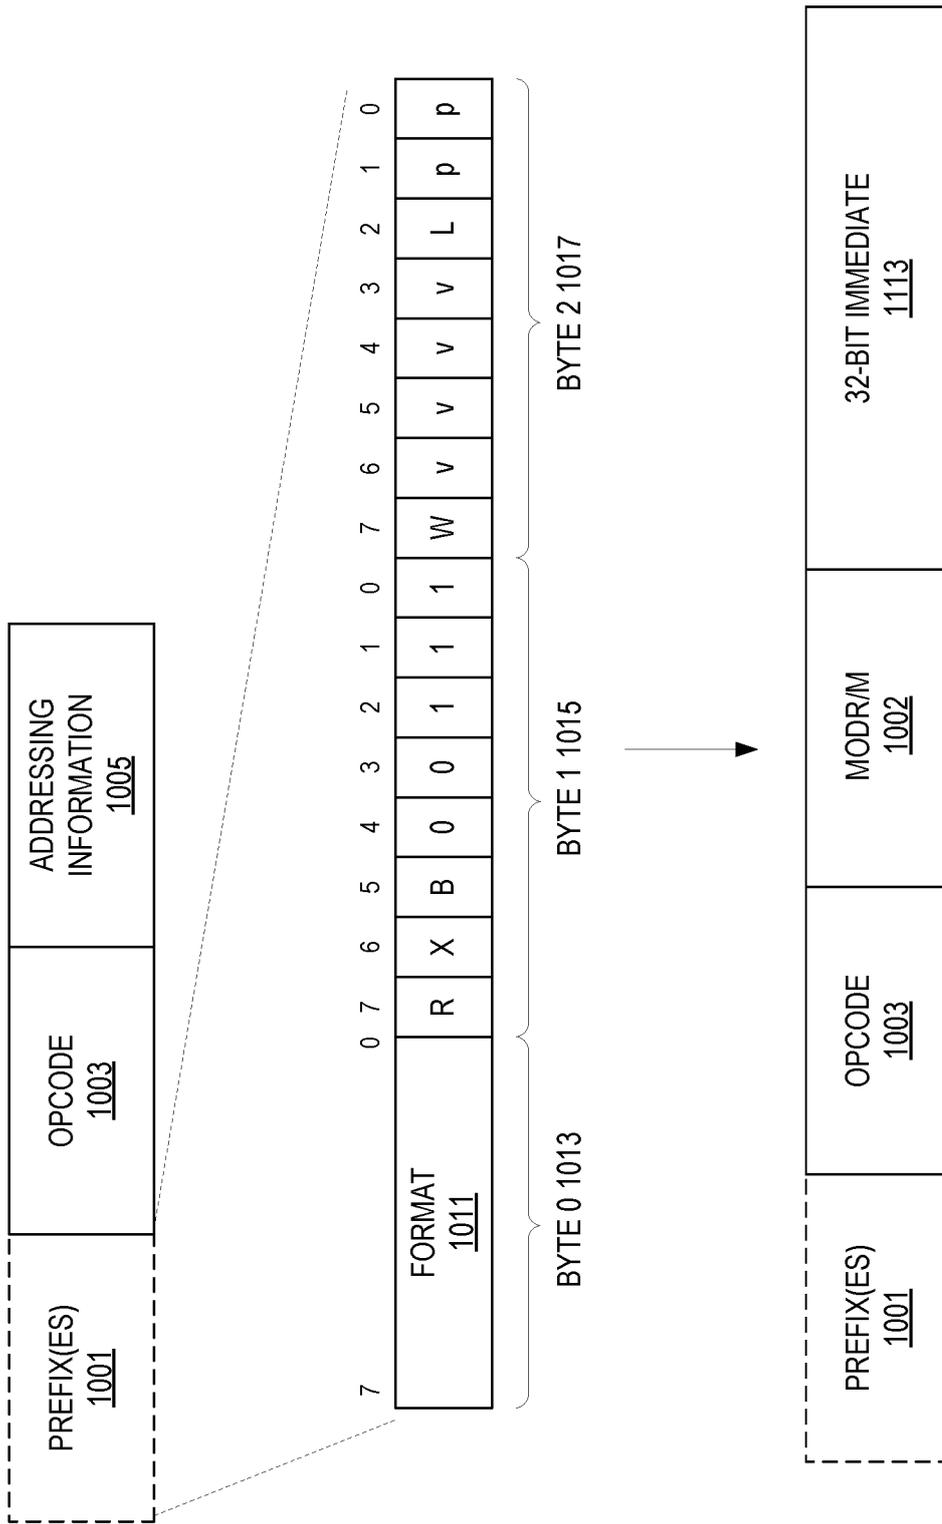


FIG. 11

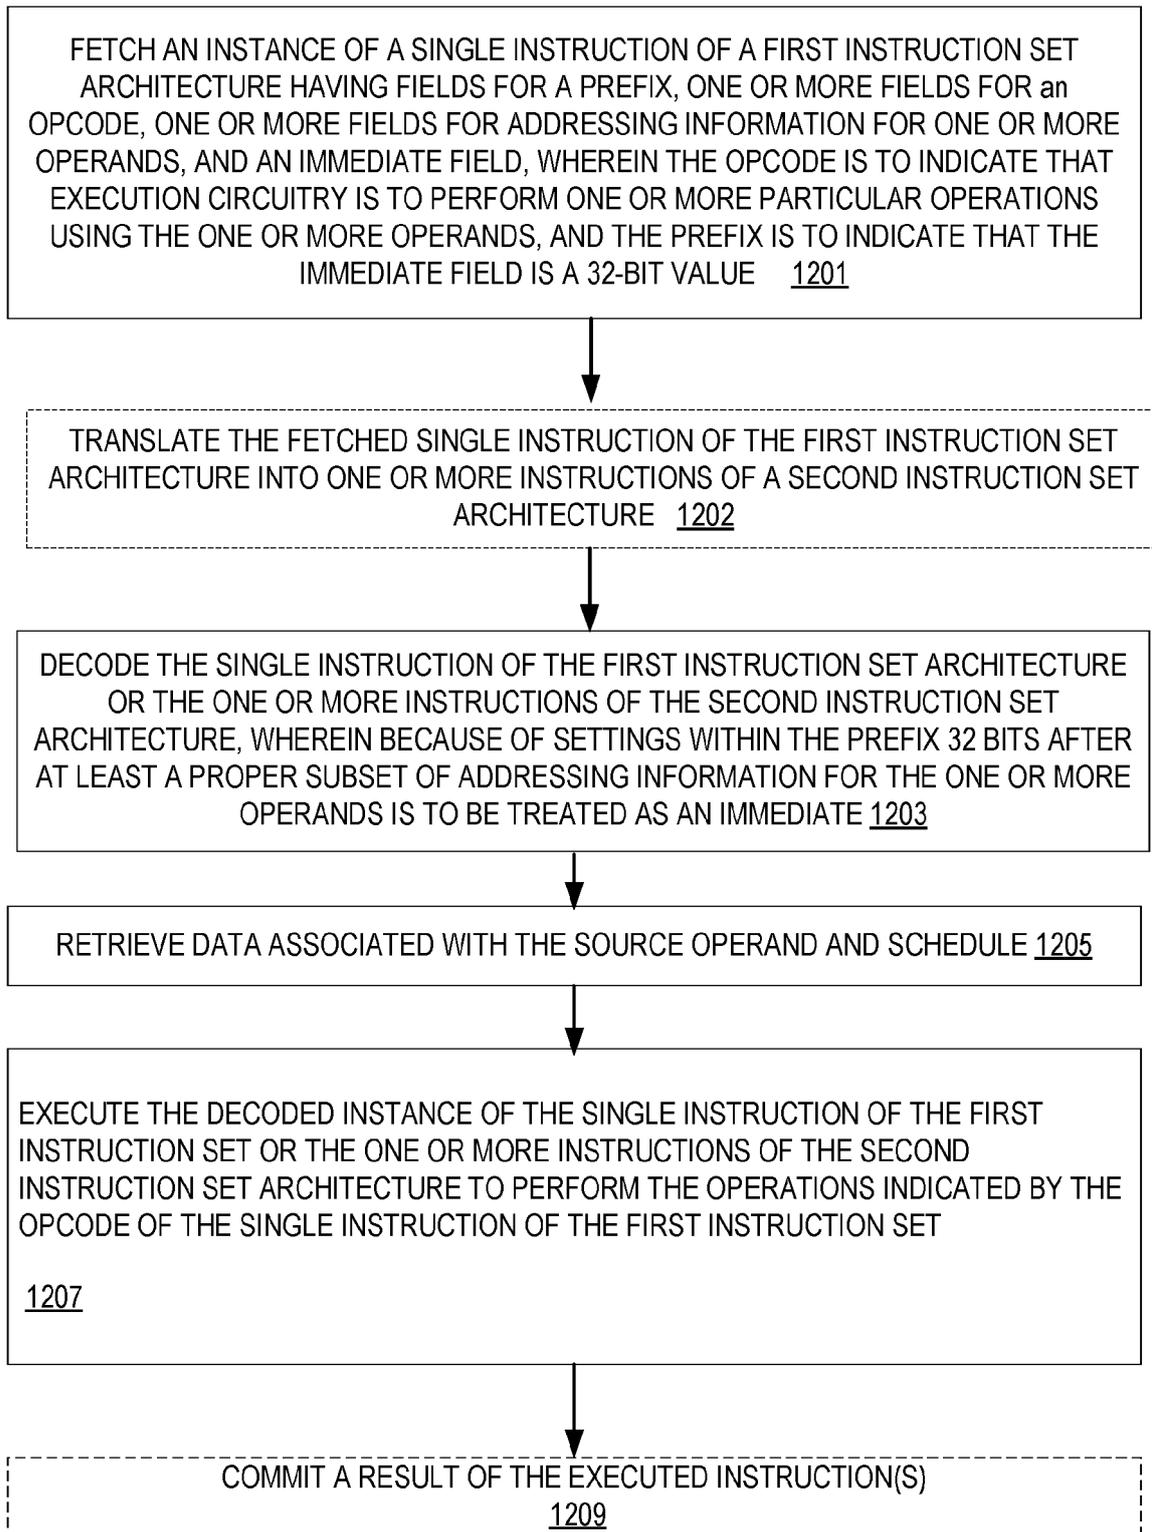


FIG. 12

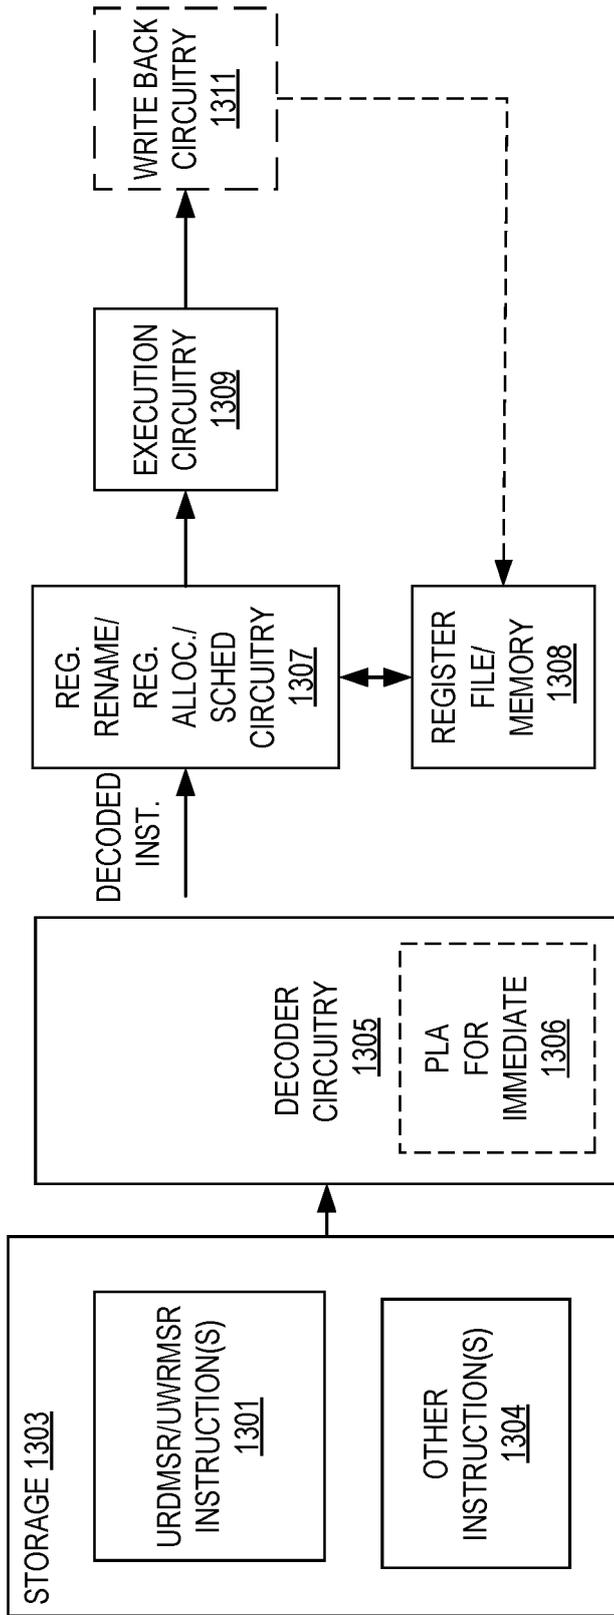


FIG. 13

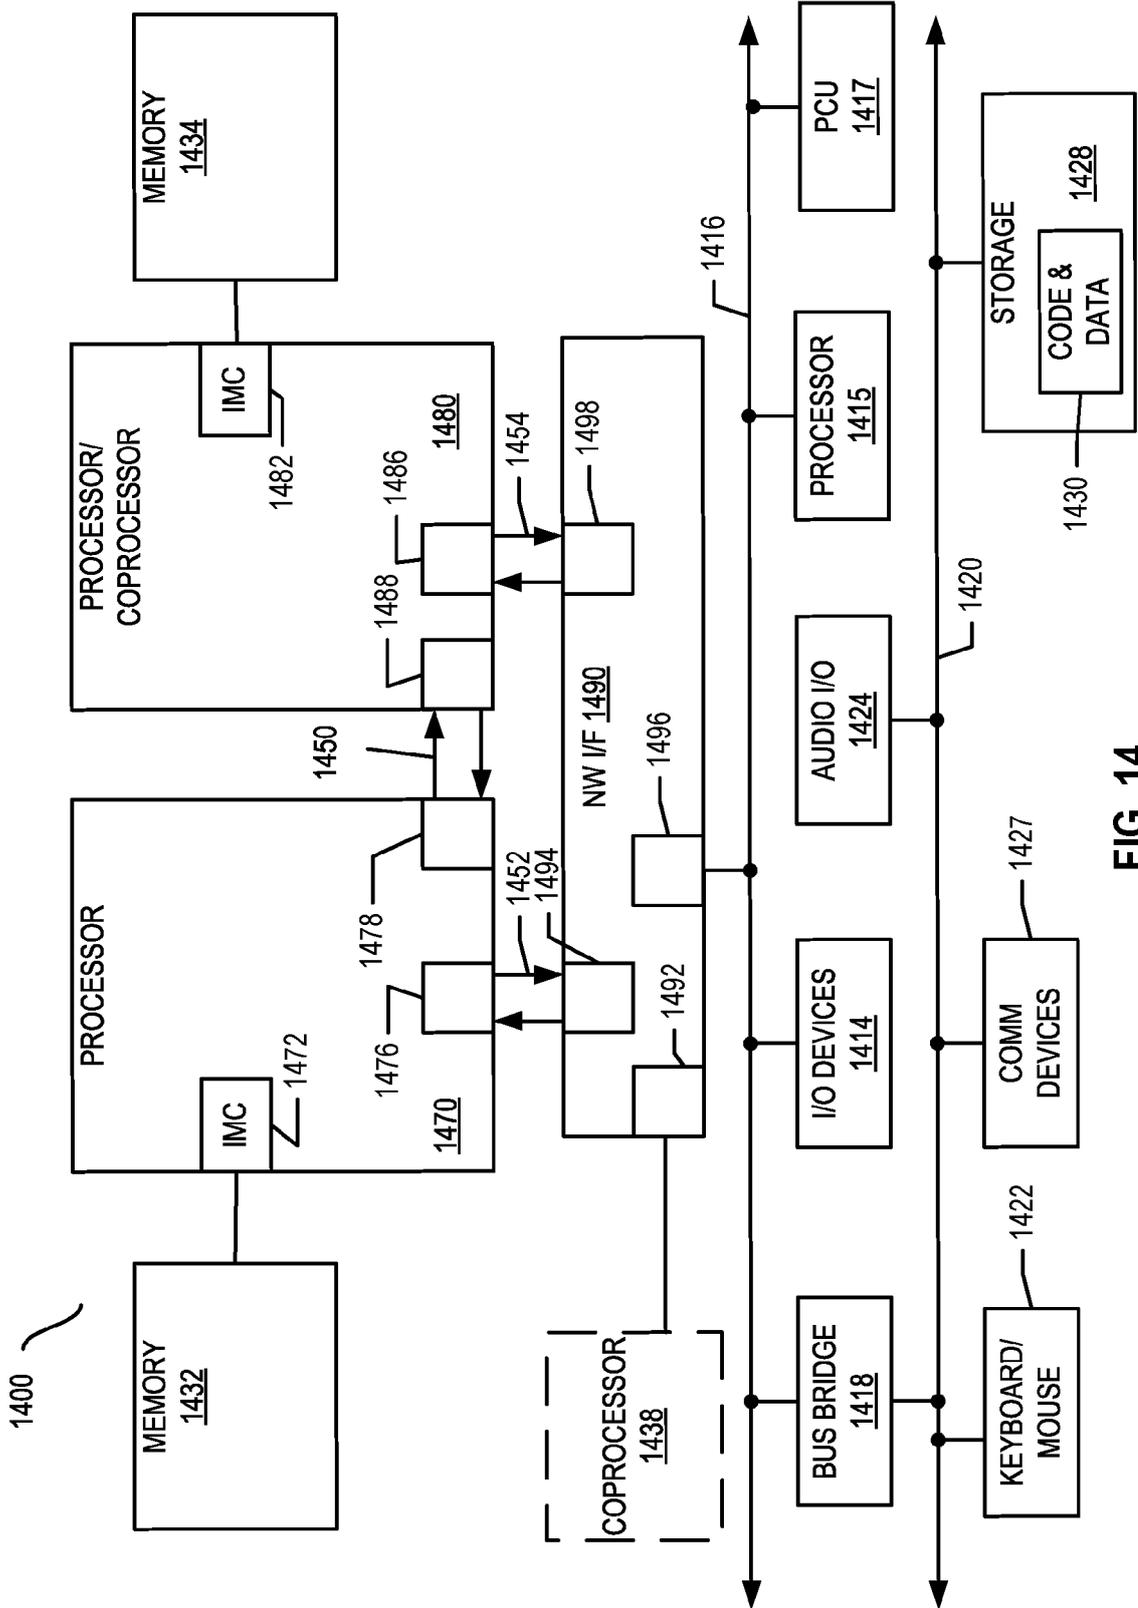


FIG. 14

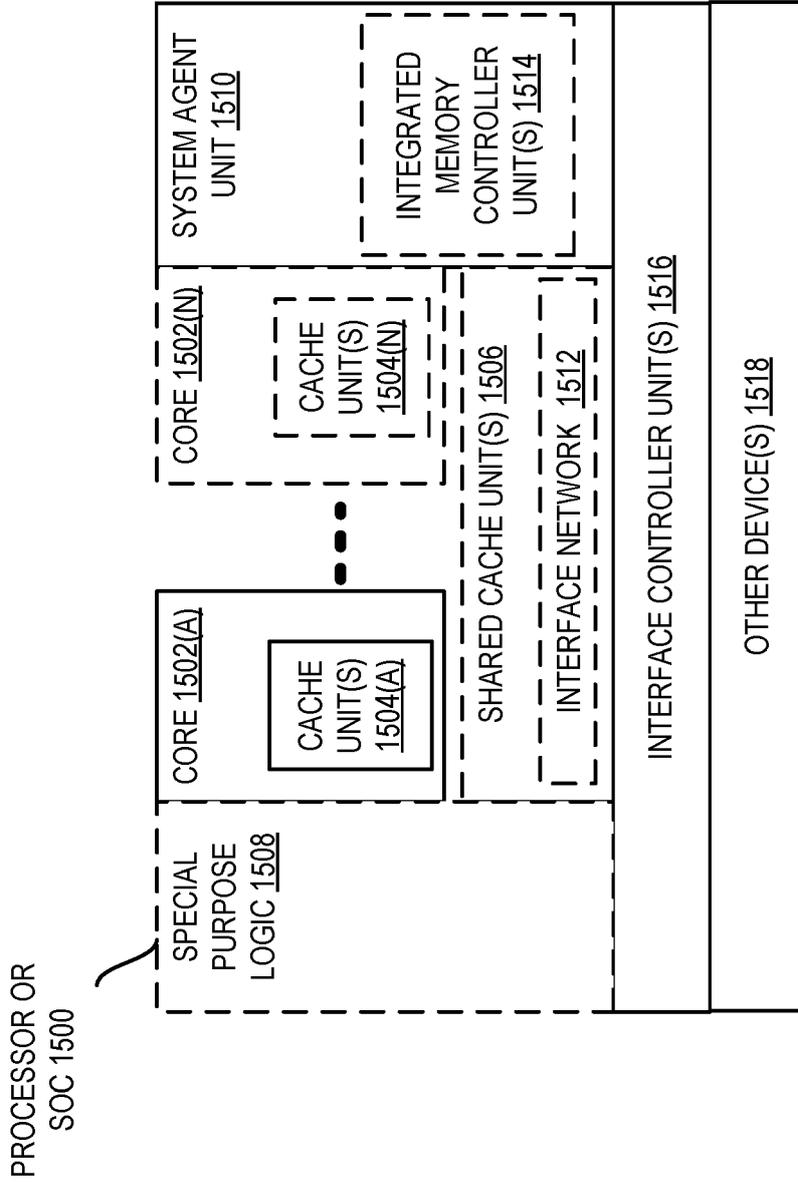


FIG. 15

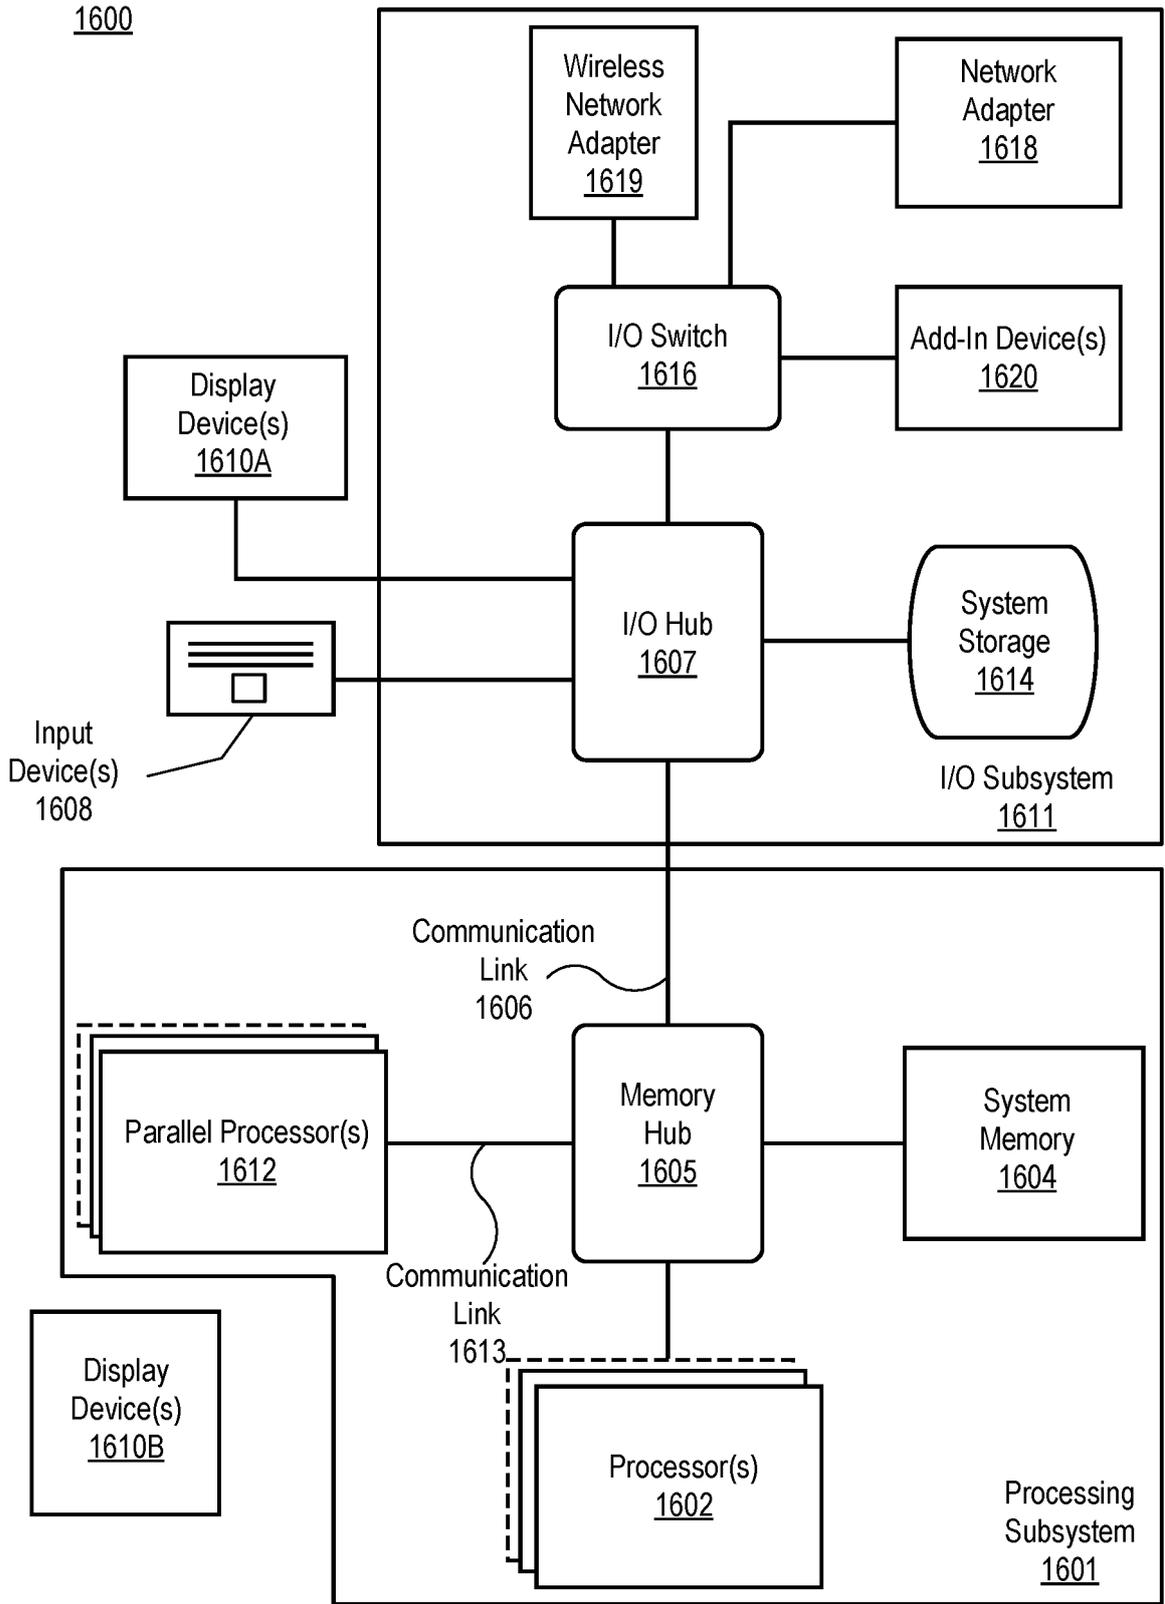


FIG. 16

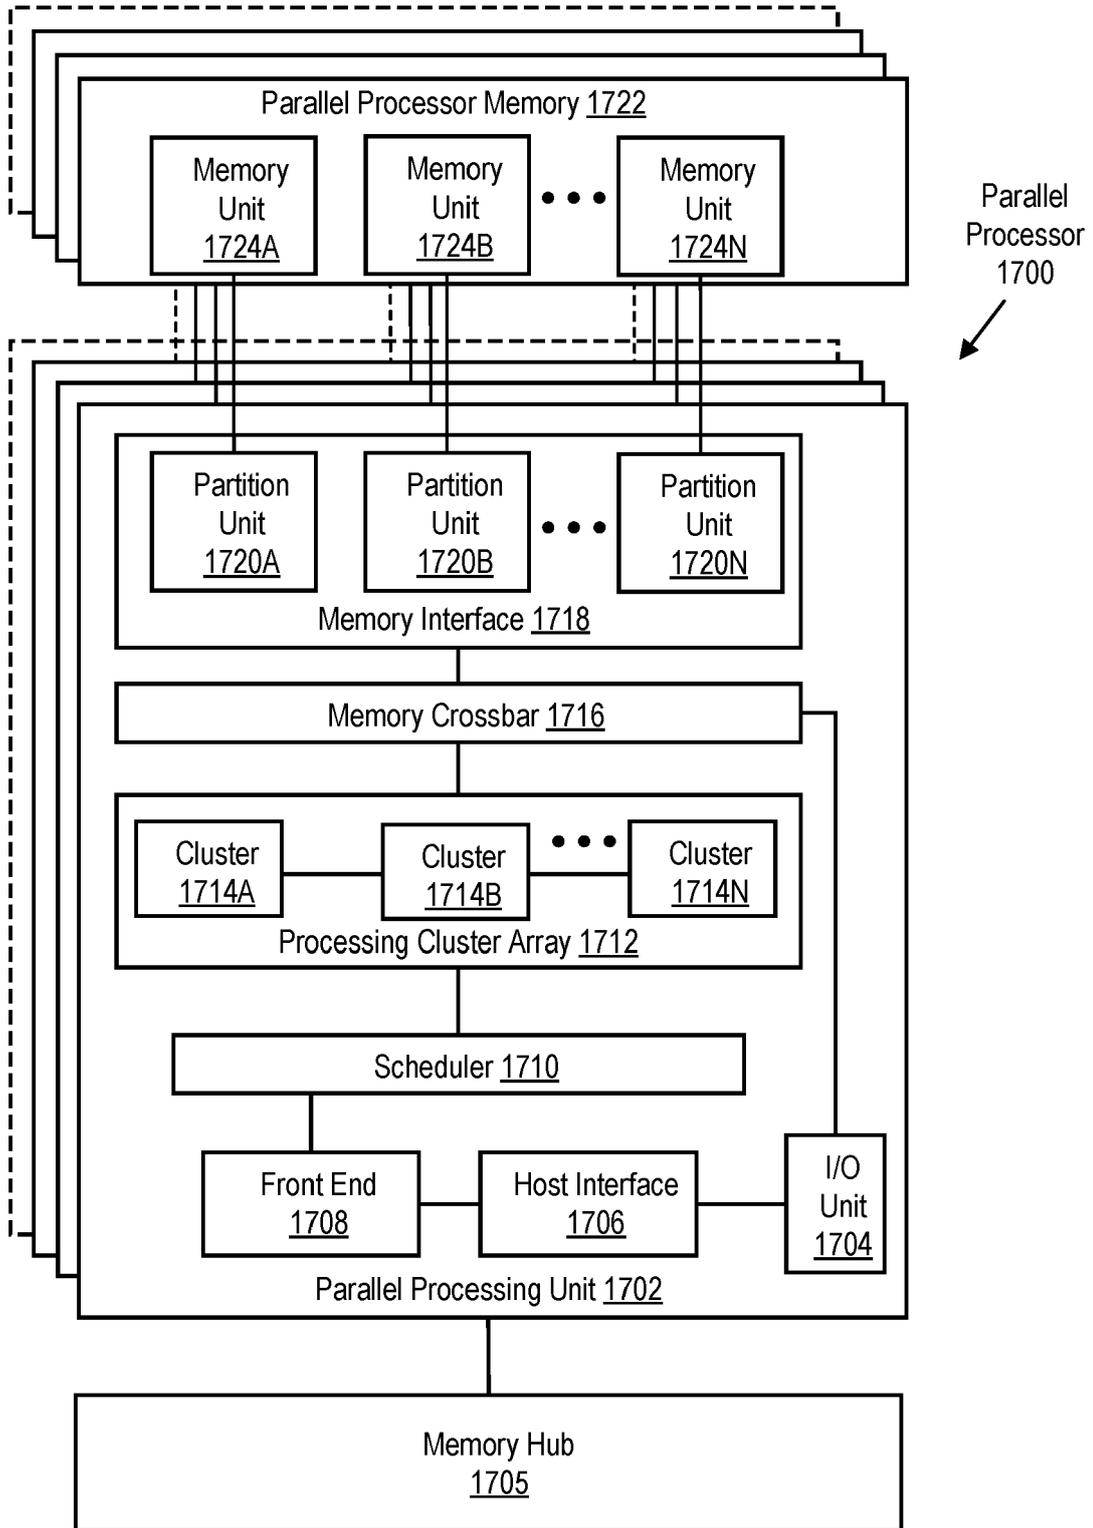


FIG. 17A

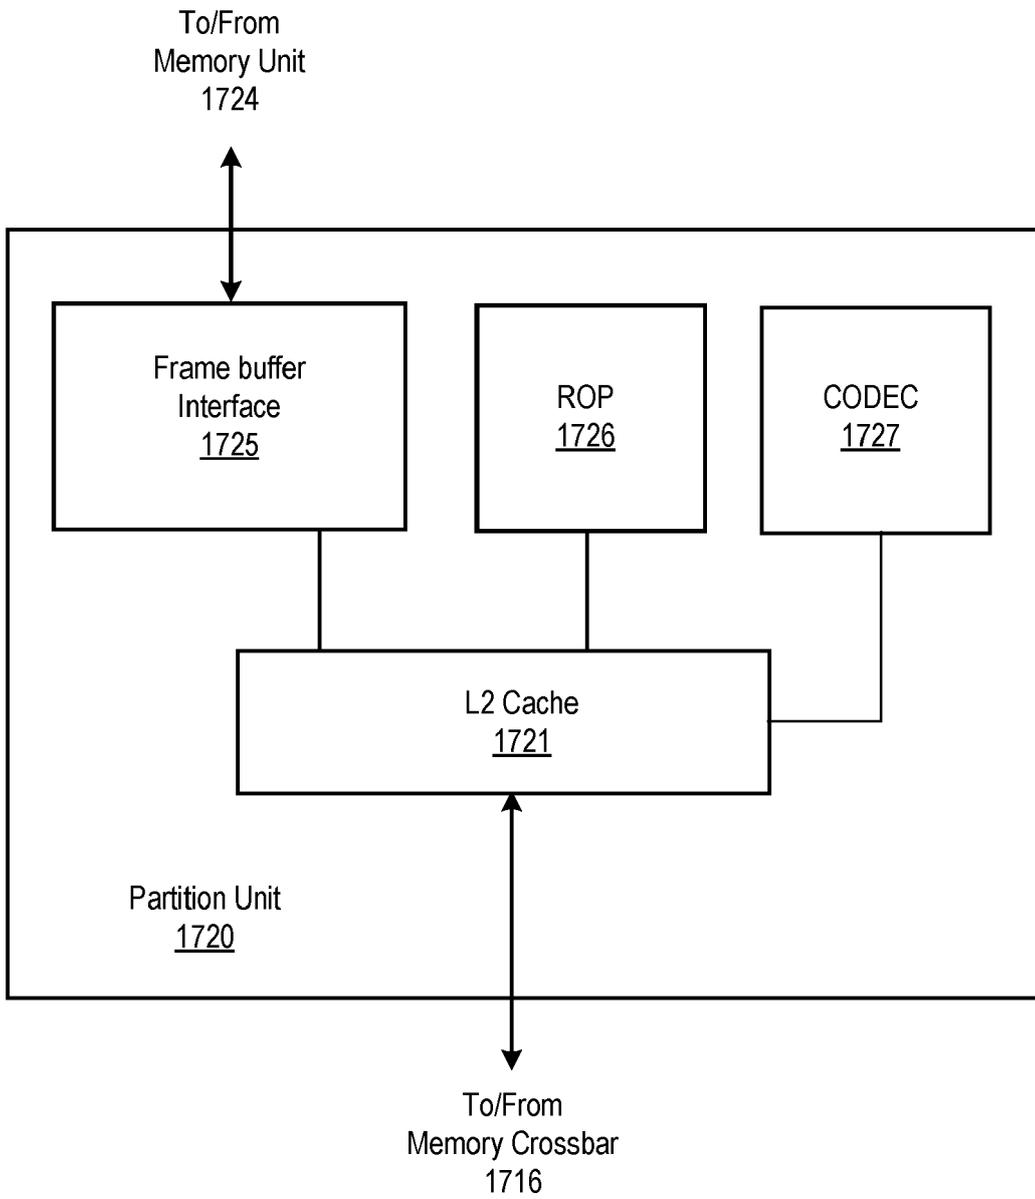


FIG. 17B

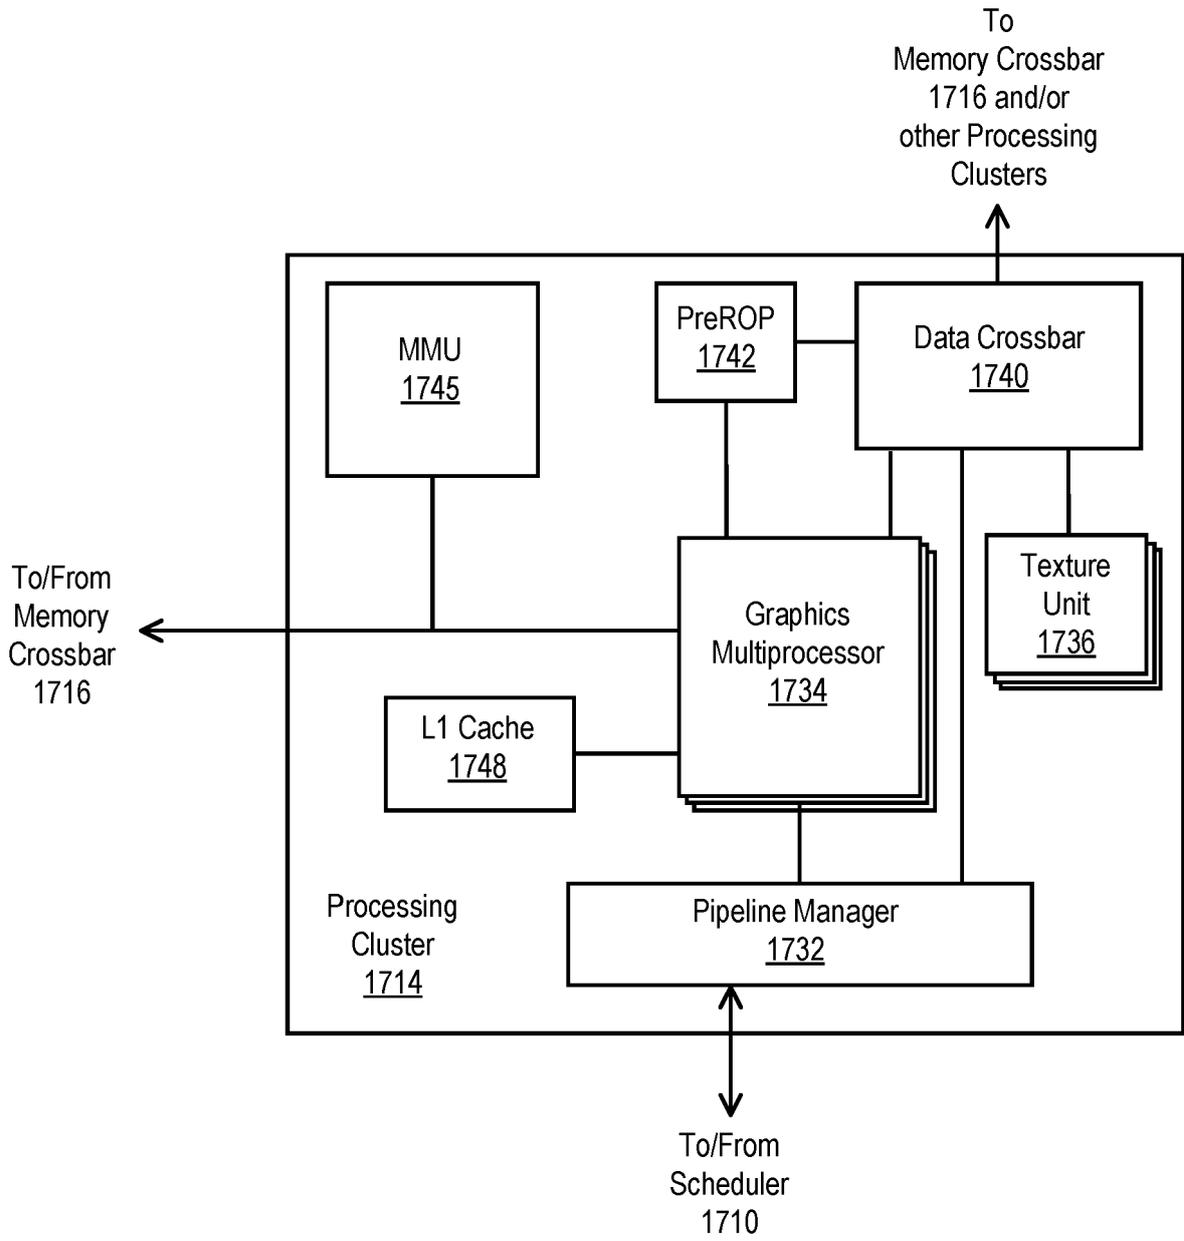


FIG. 17C

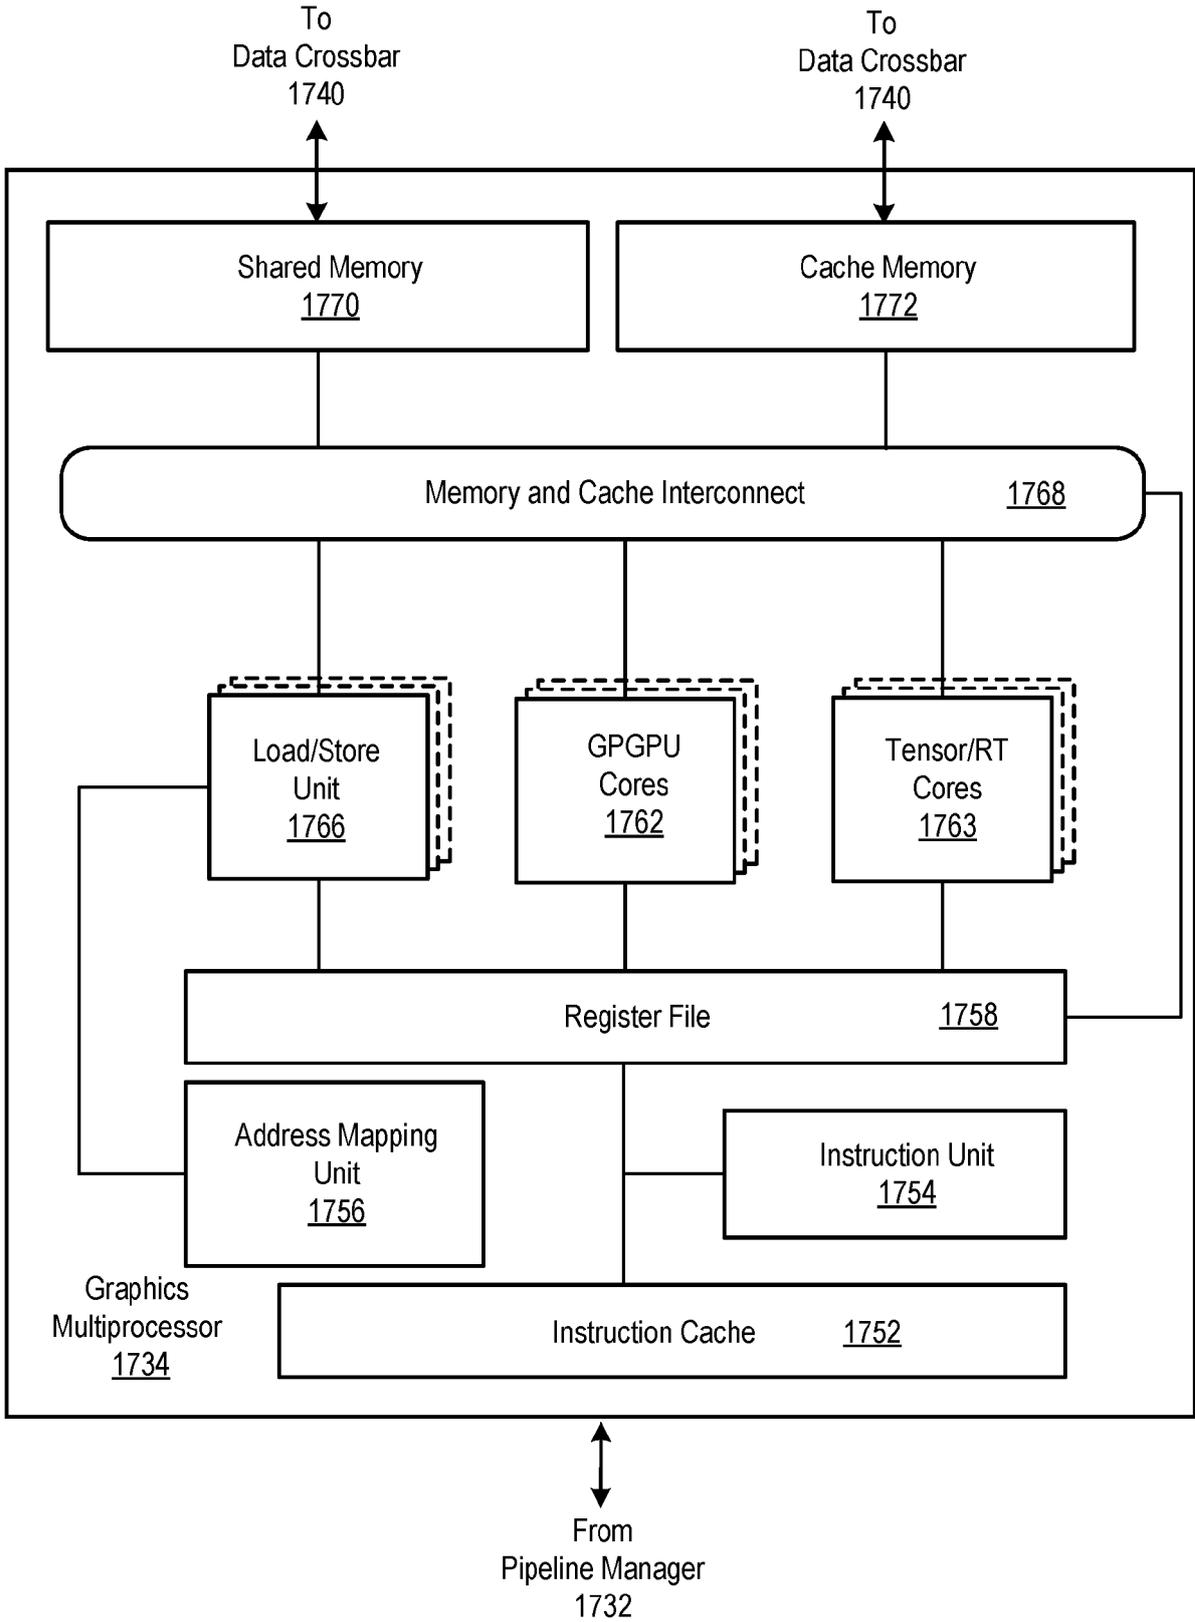


FIG. 17D

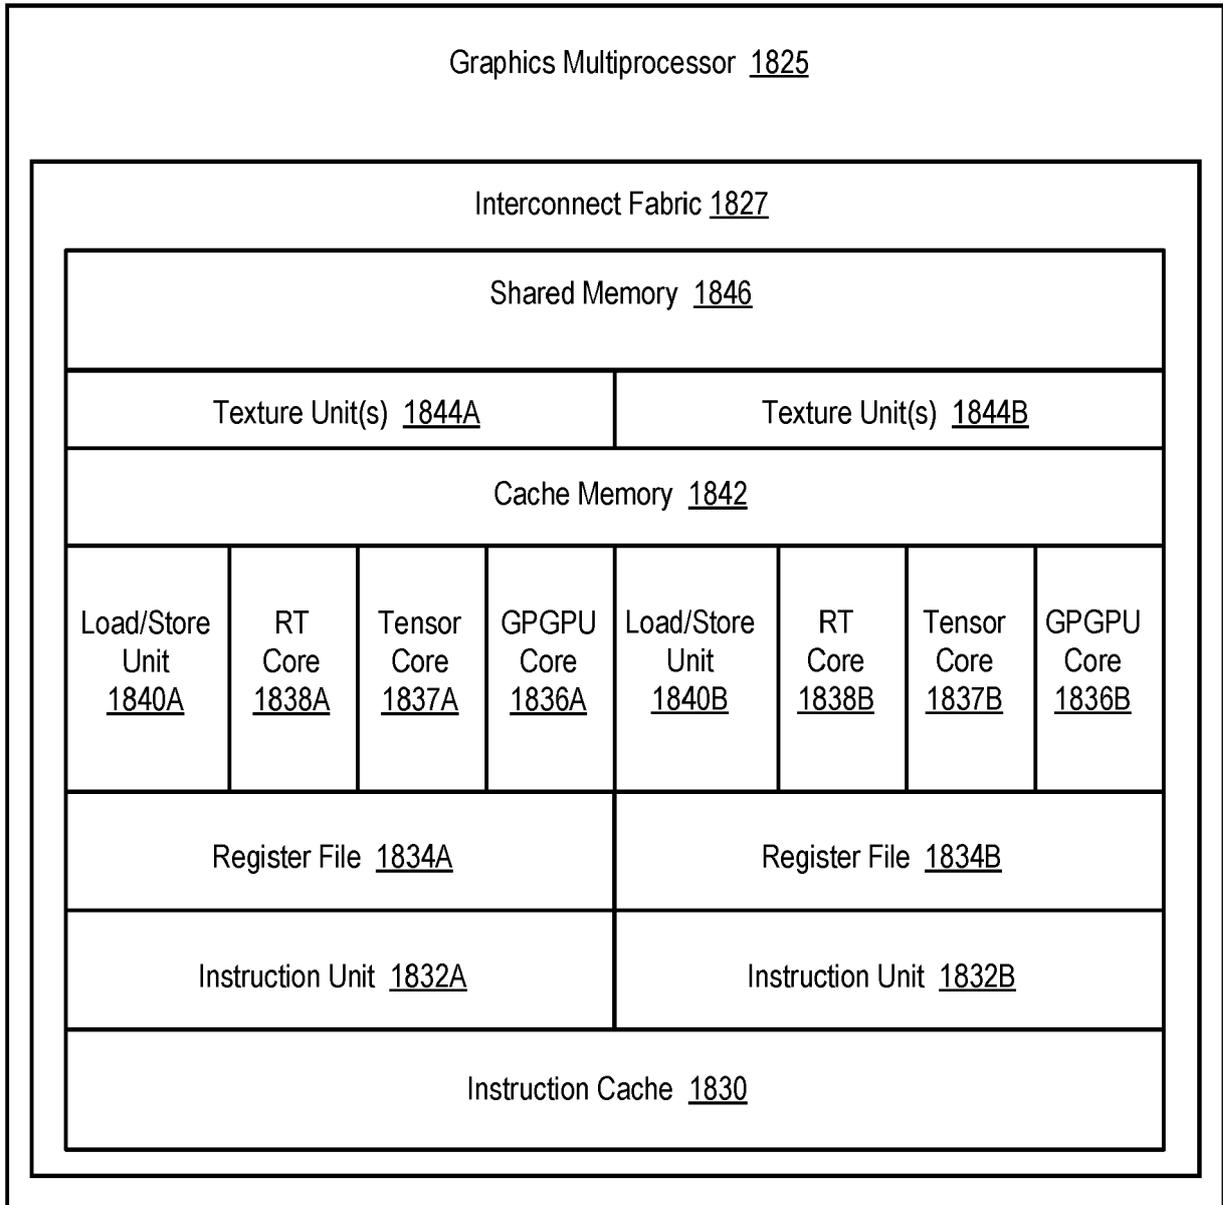


FIG. 18A

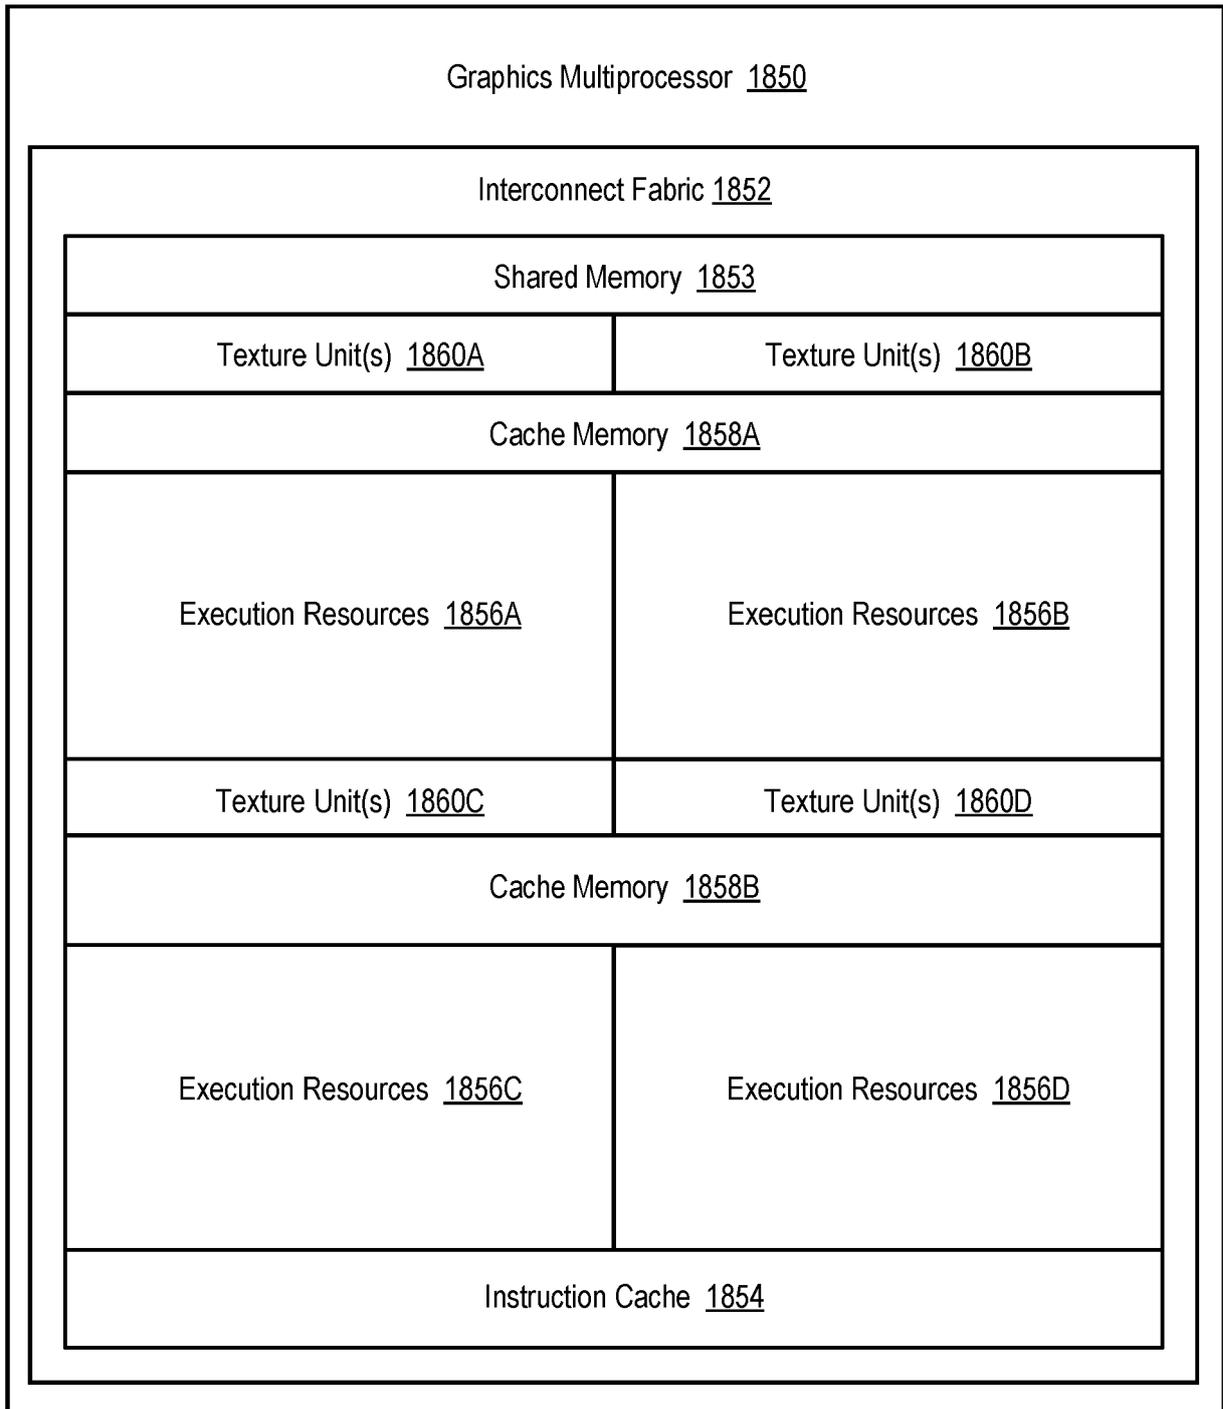


FIG. 18B

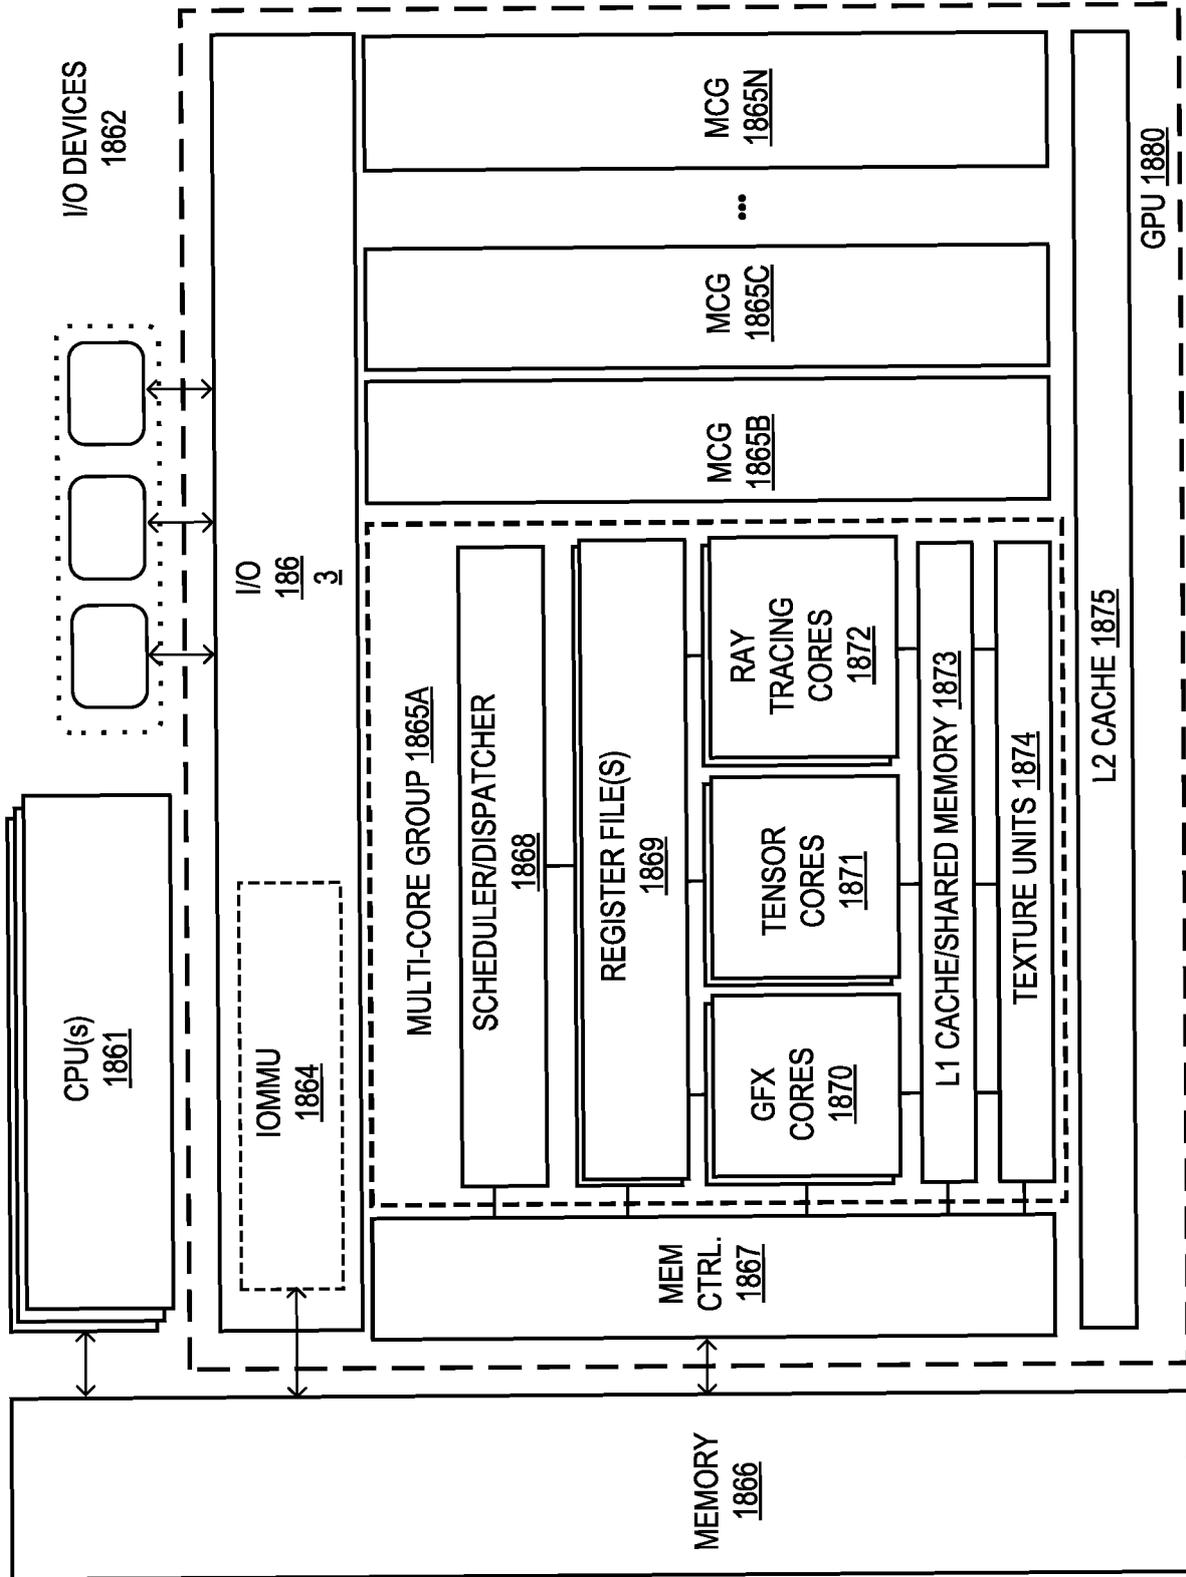


FIG. 18C

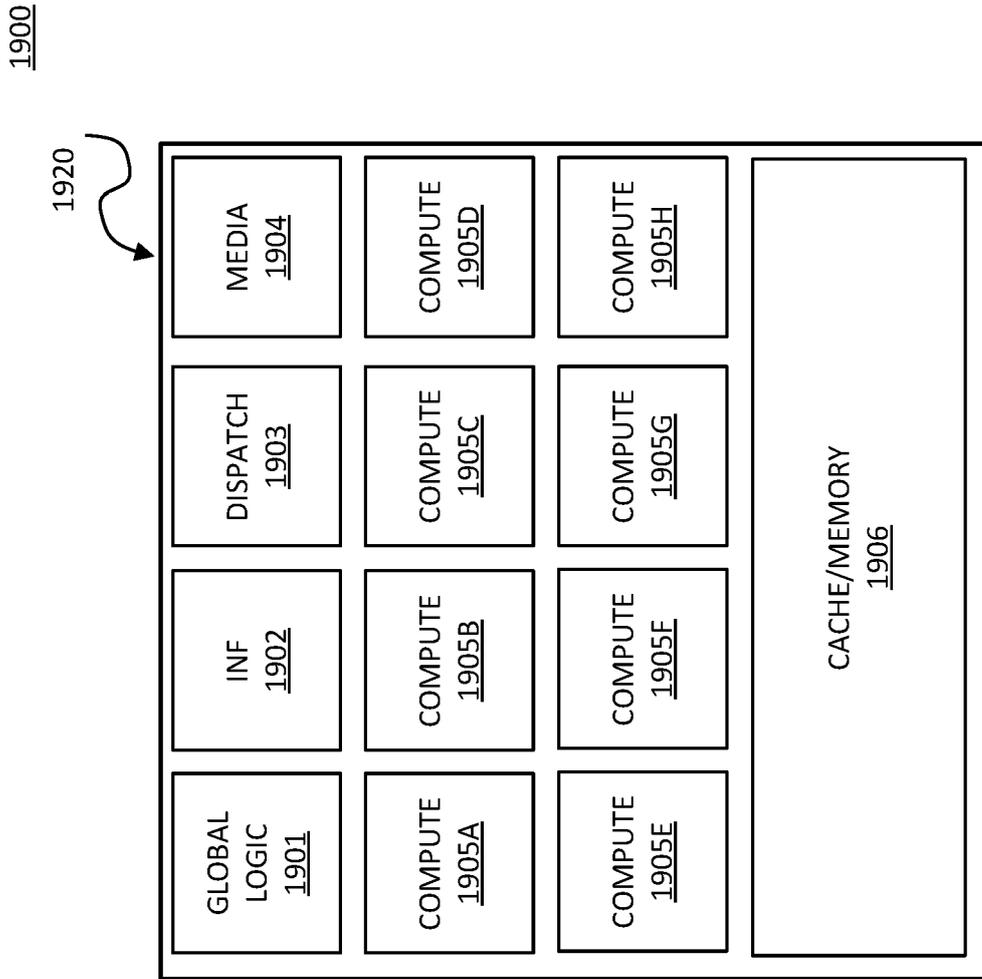


FIG. 19

2000

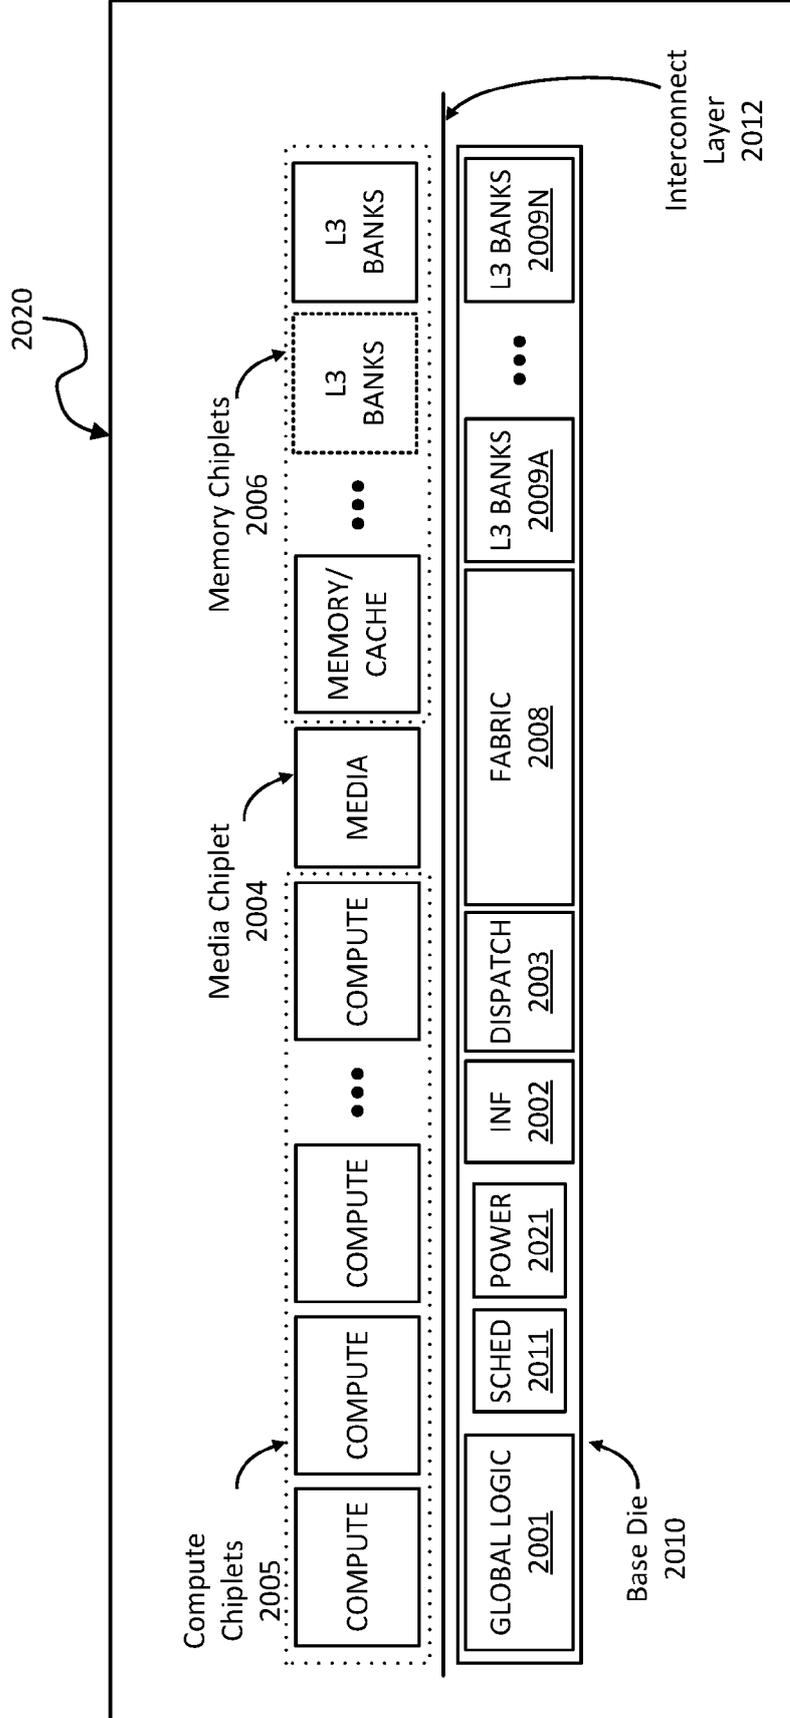


FIG. 20A

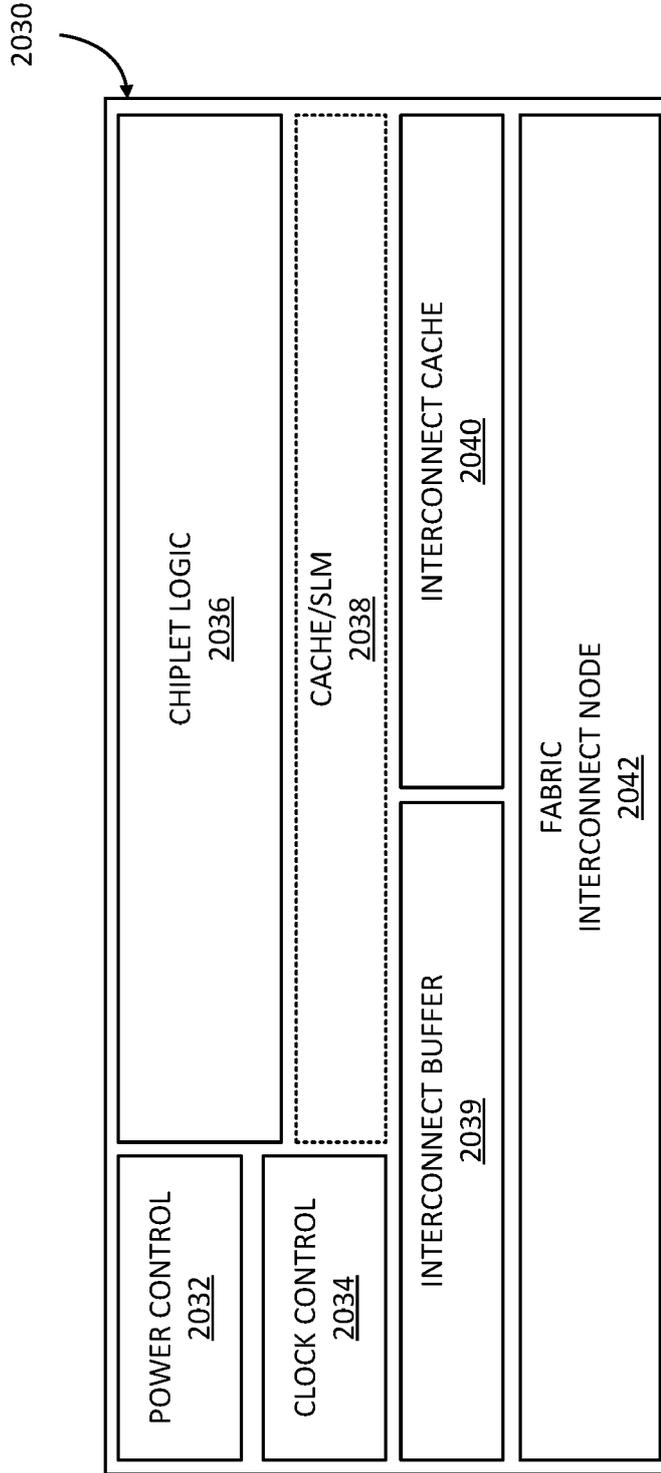


FIG. 20B



FIG. 21(A)

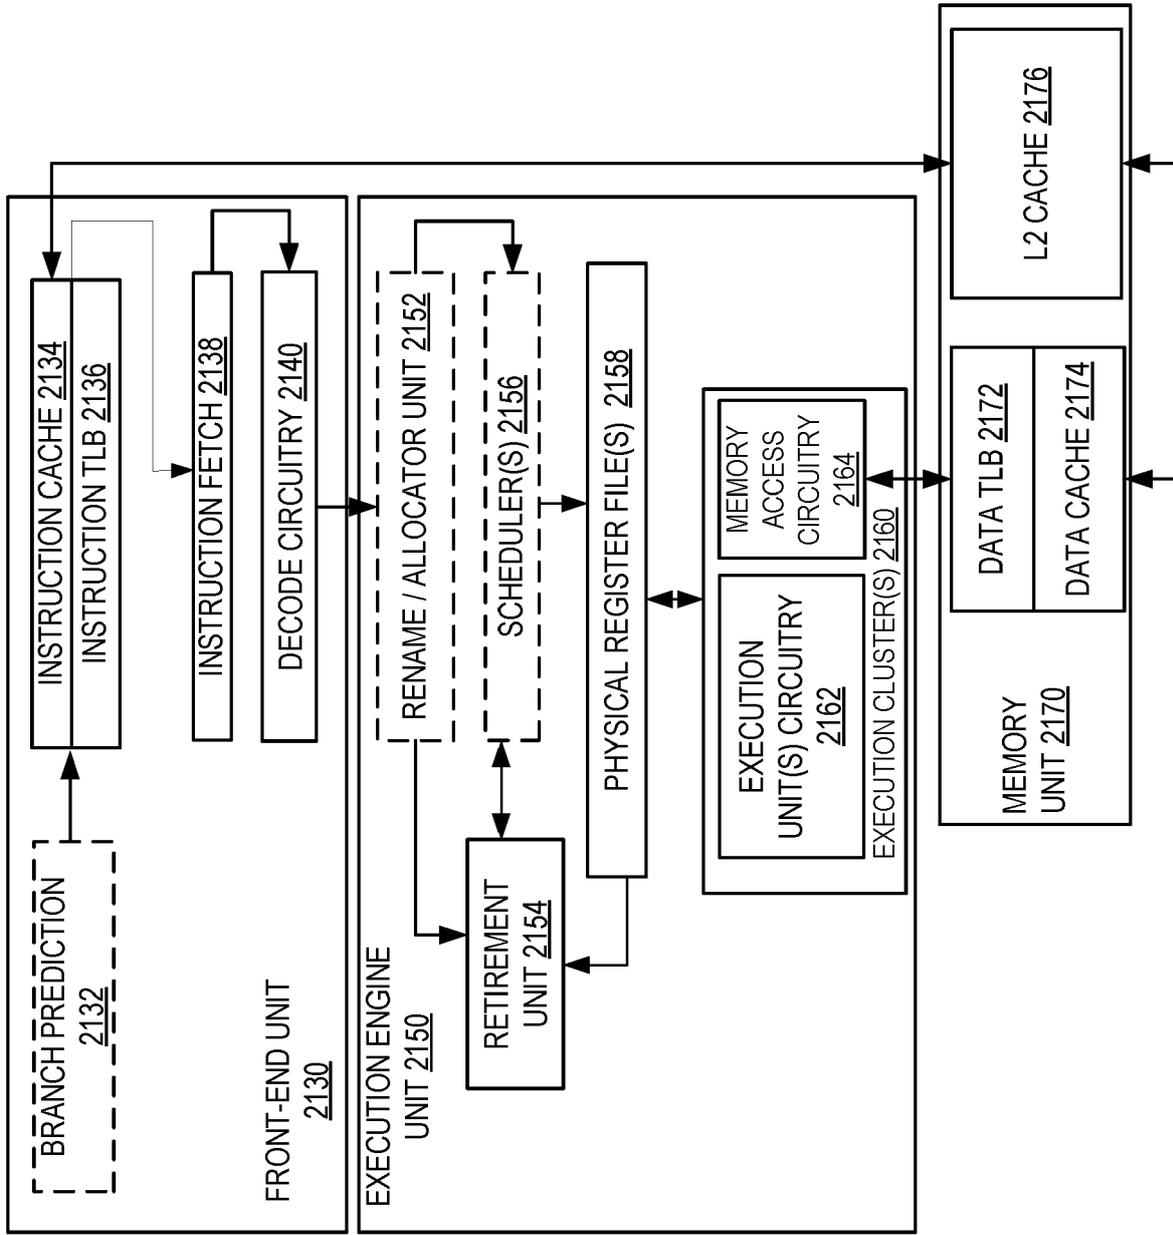


FIG. 21(B)

CORE 2190

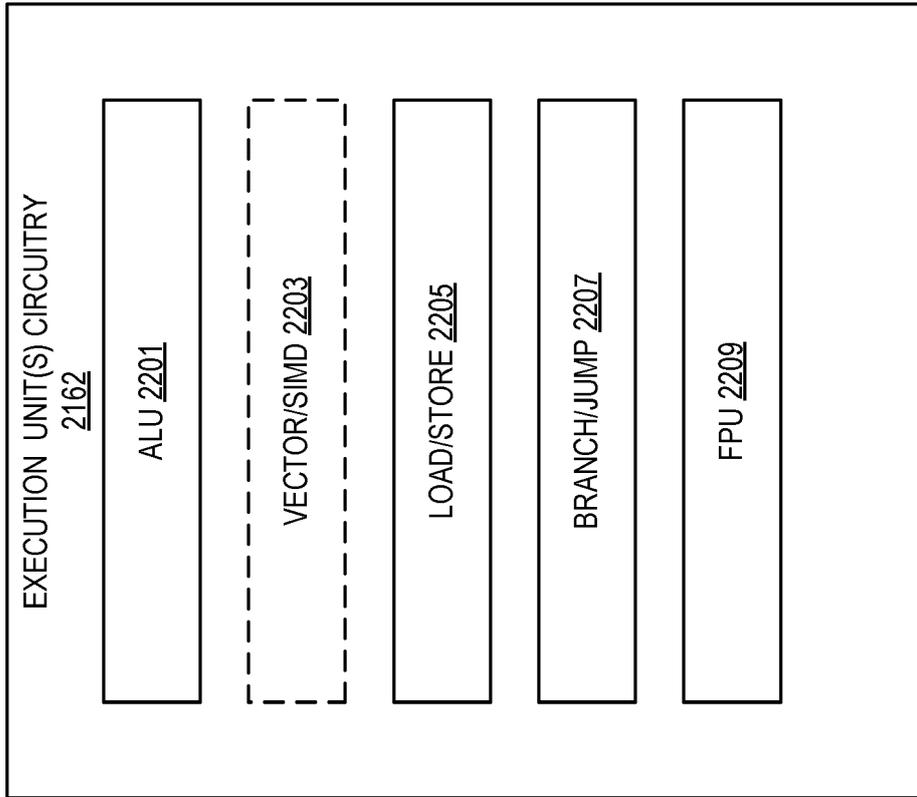


FIG. 22

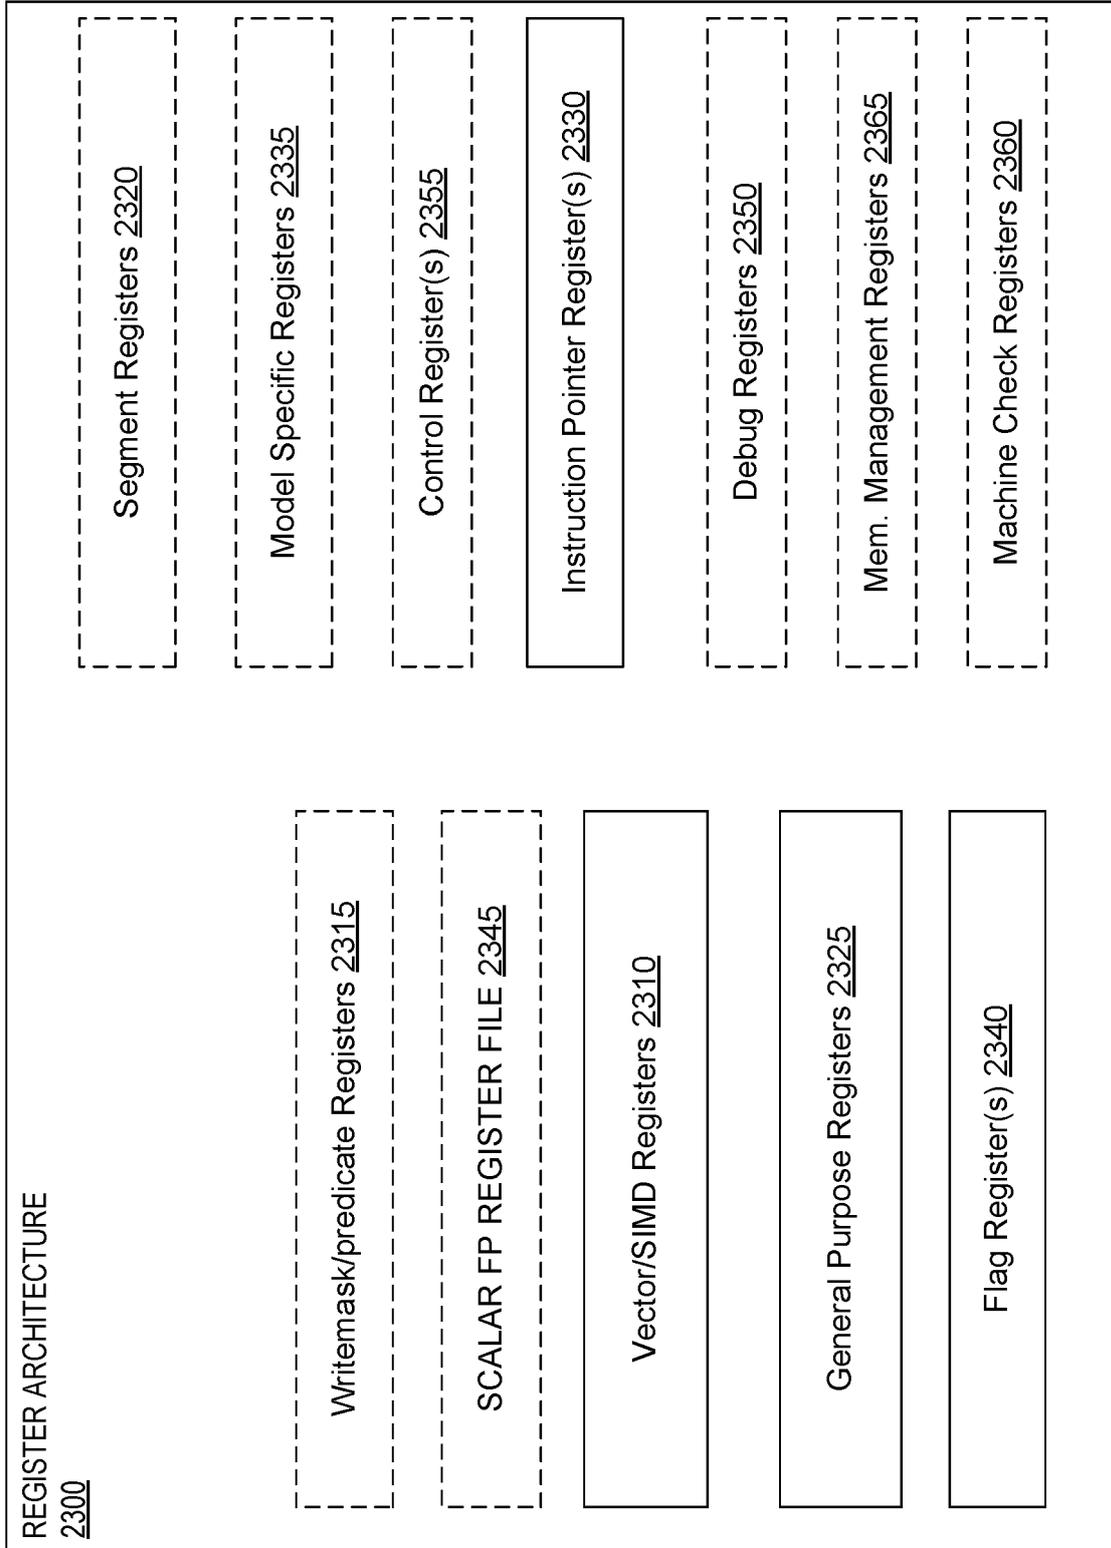


FIG. 23

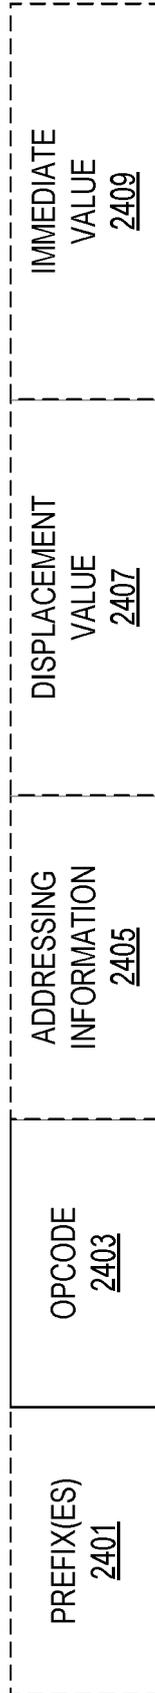


FIG. 24

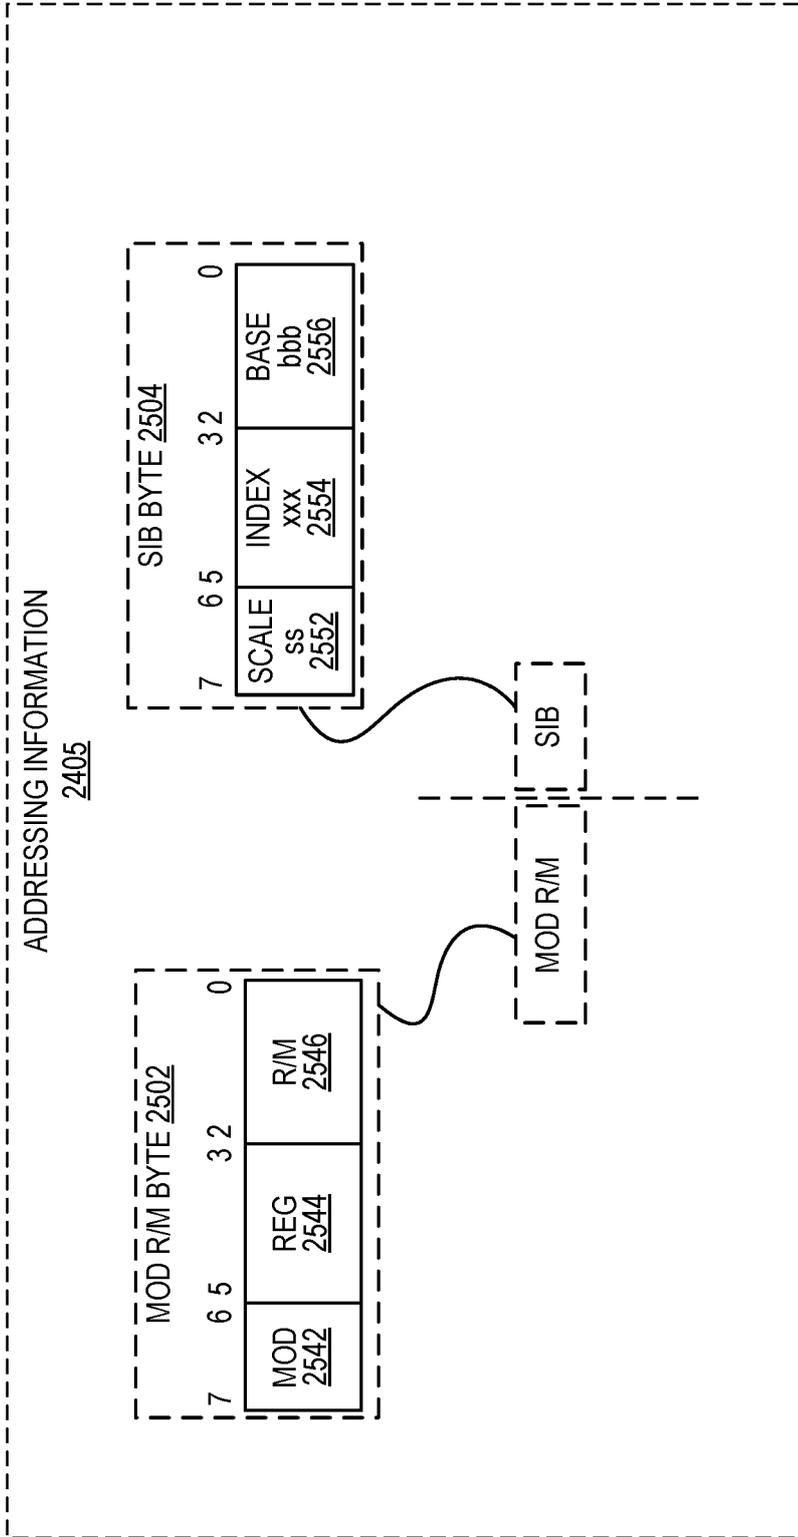
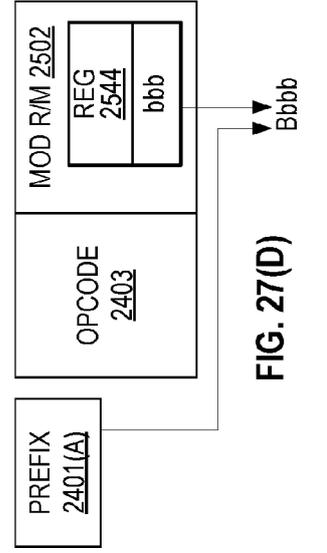
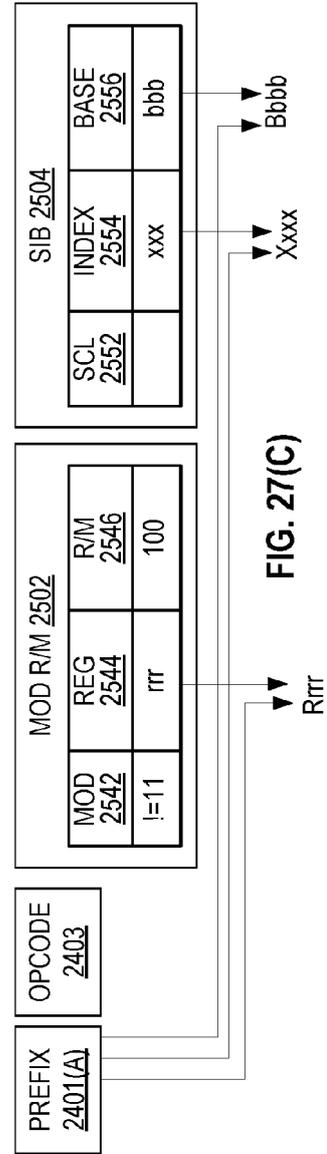
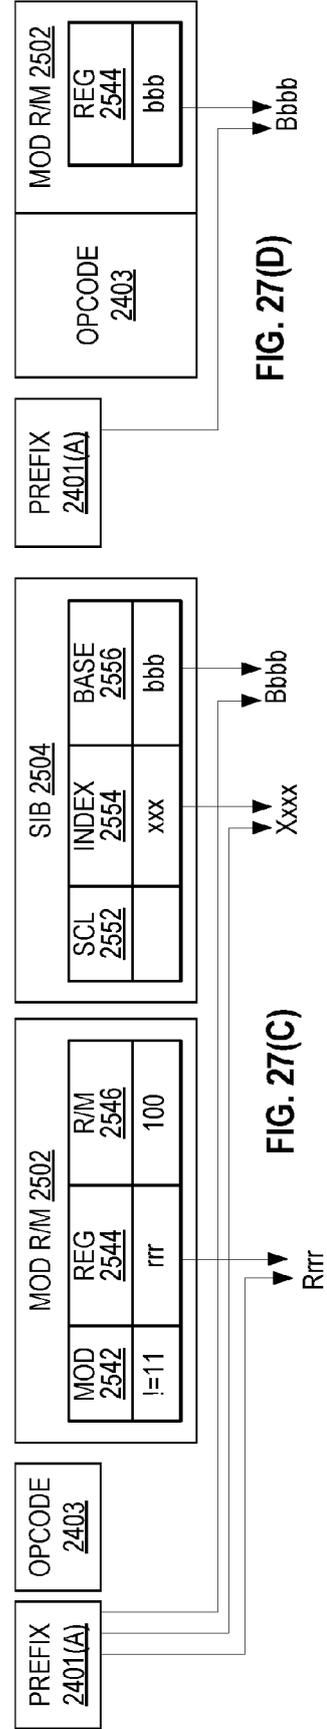
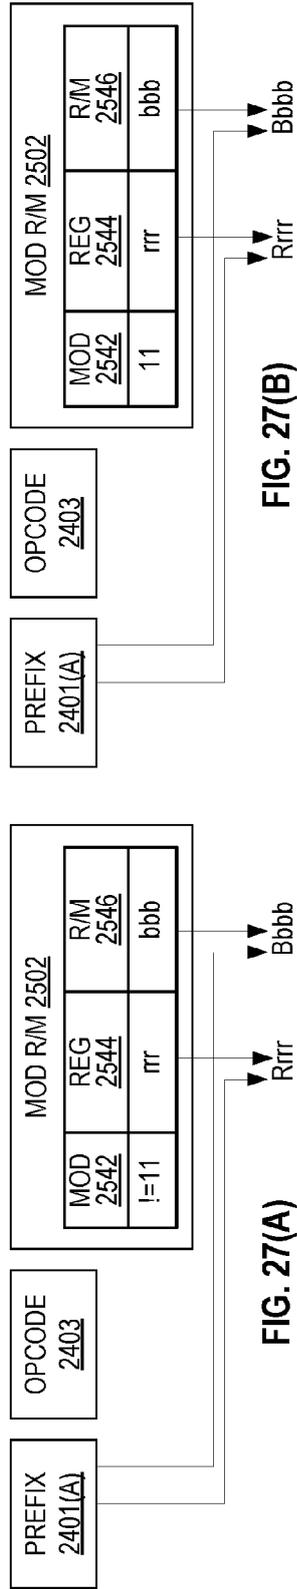
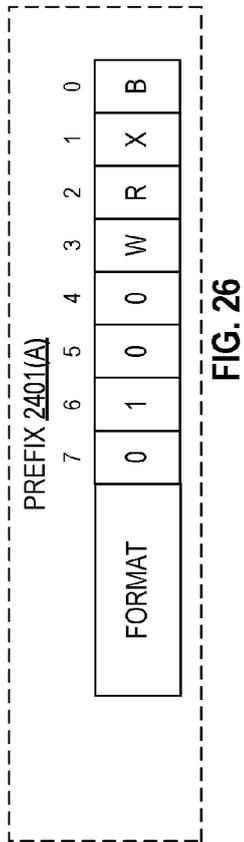


FIG. 25



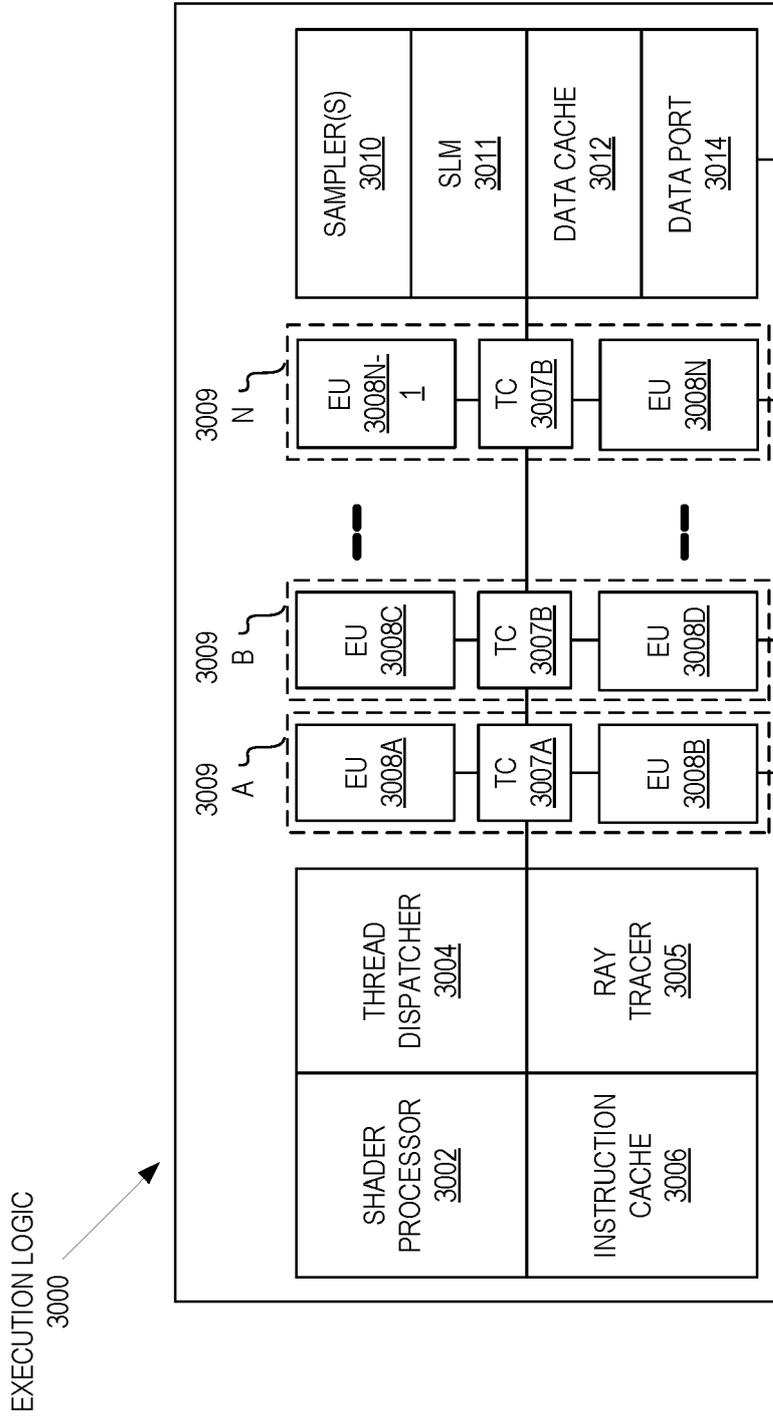


FIG. 30A

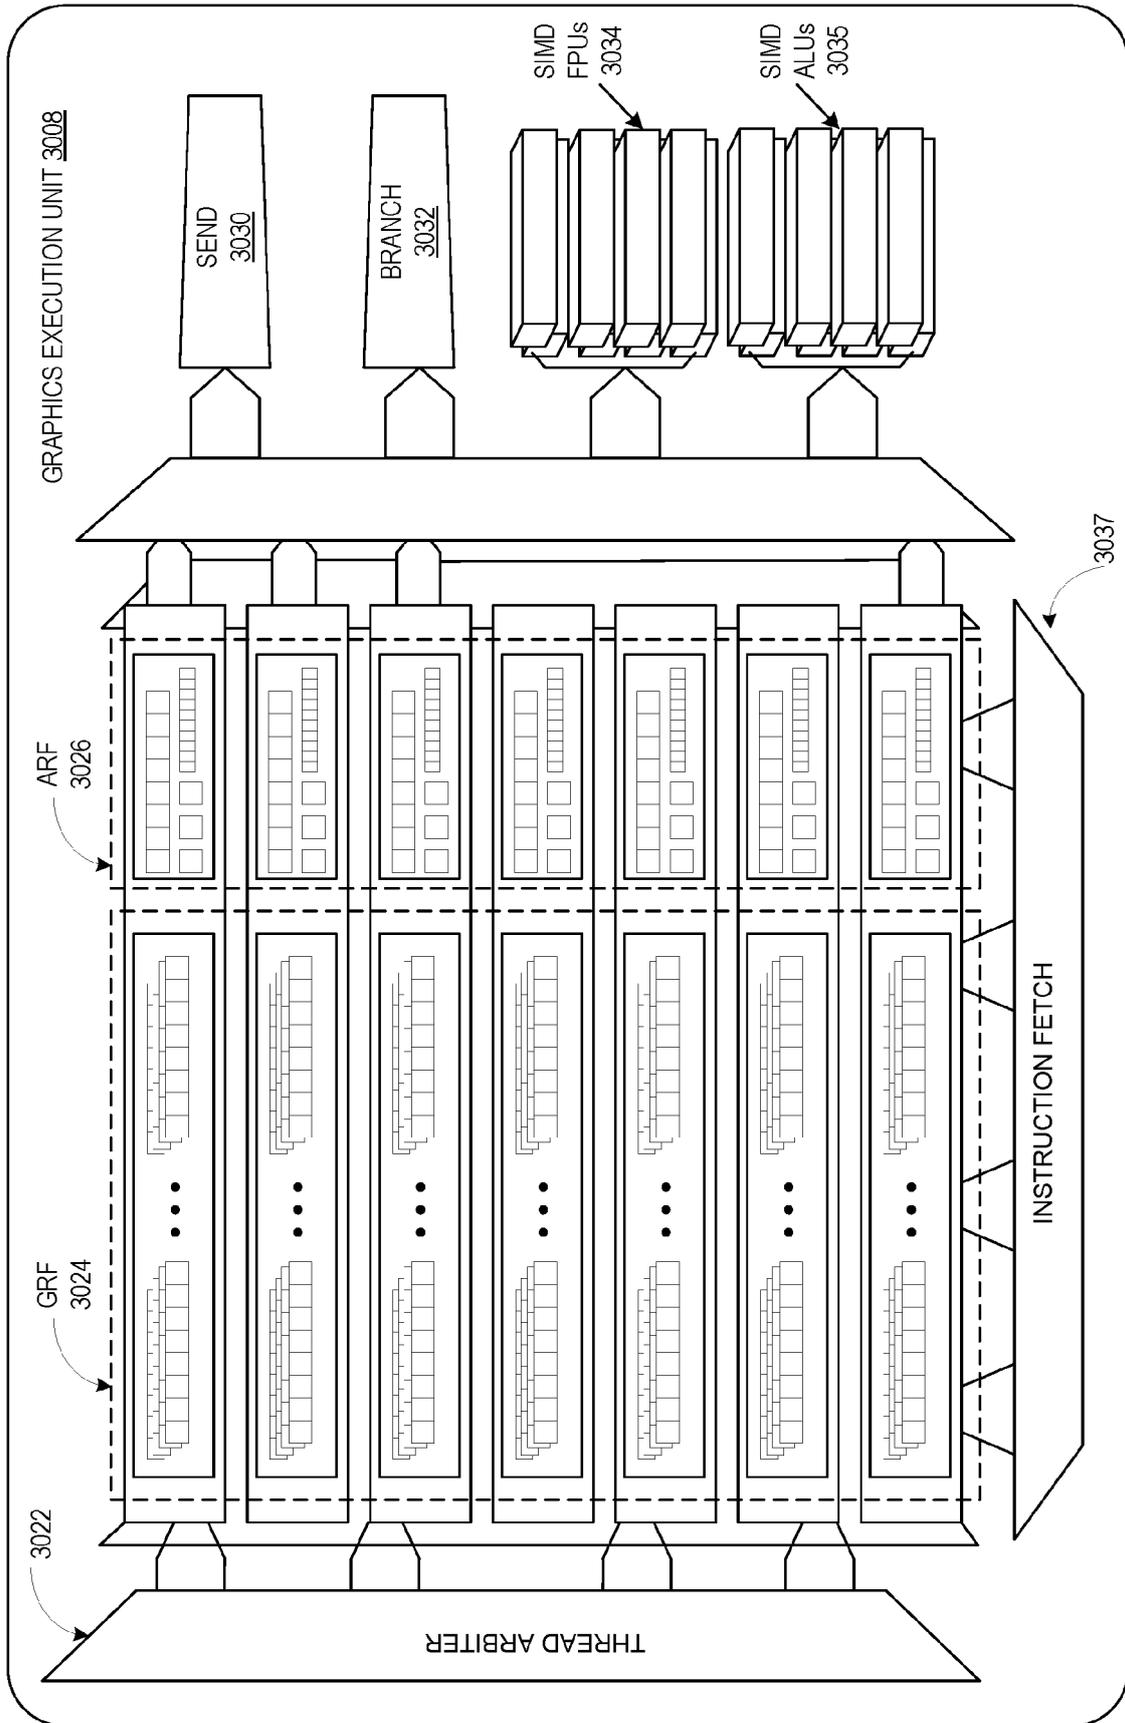


FIG. 30B

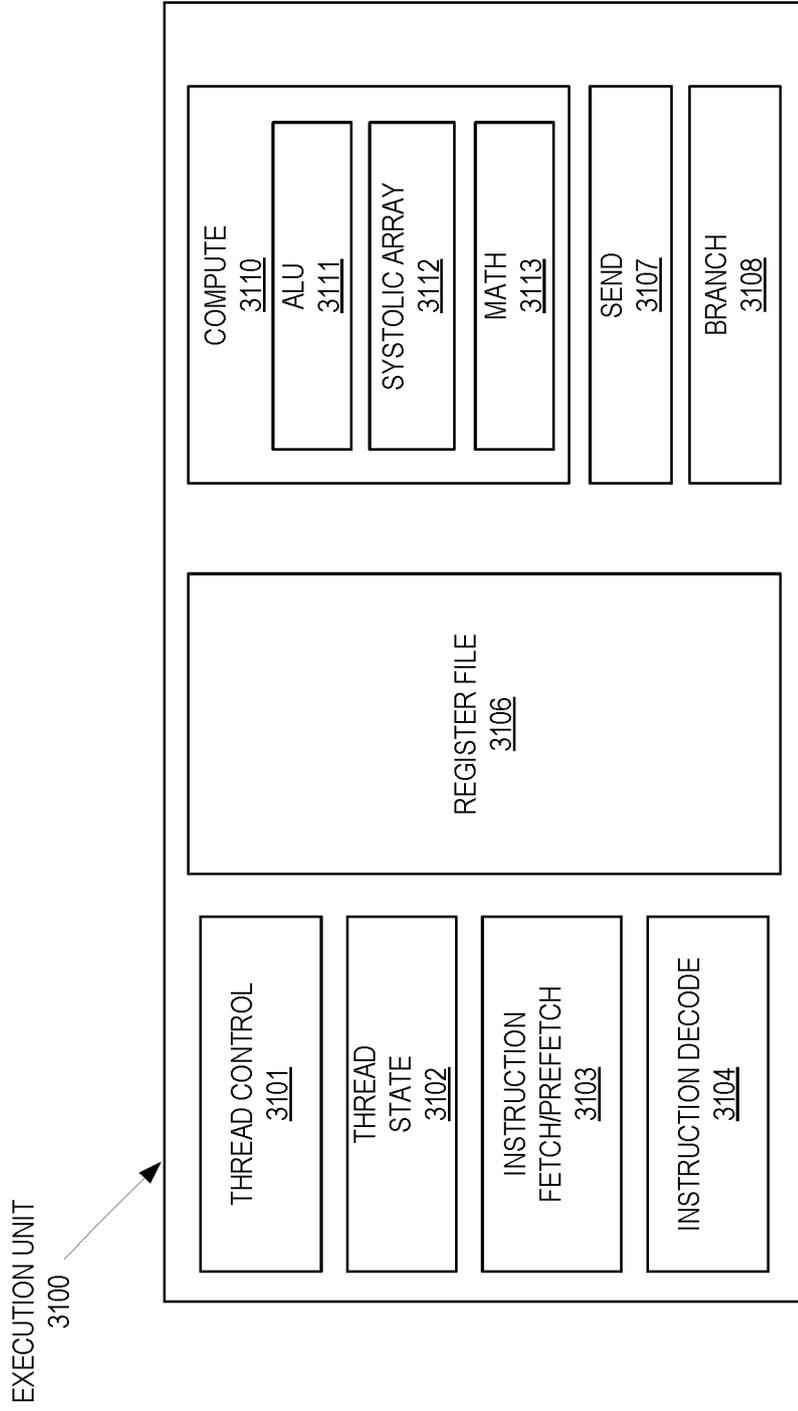


FIG. 31

GRAPHICS PROCESSOR INSTRUCTION FORMATS

3200



OPCODE DECODE
3240

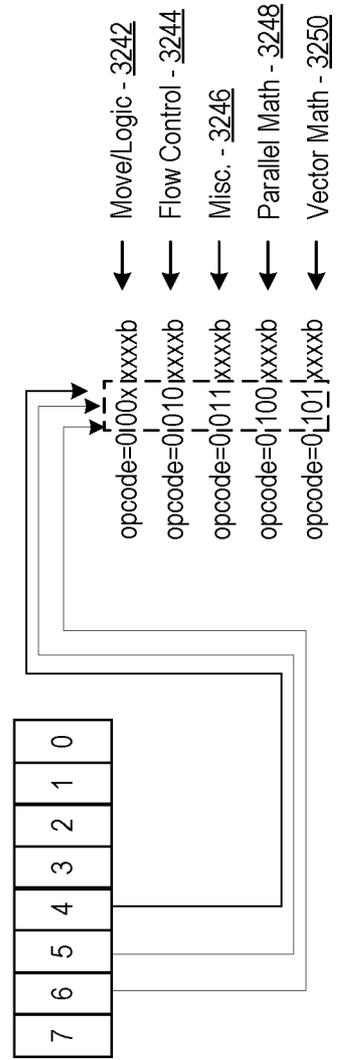


FIG. 32

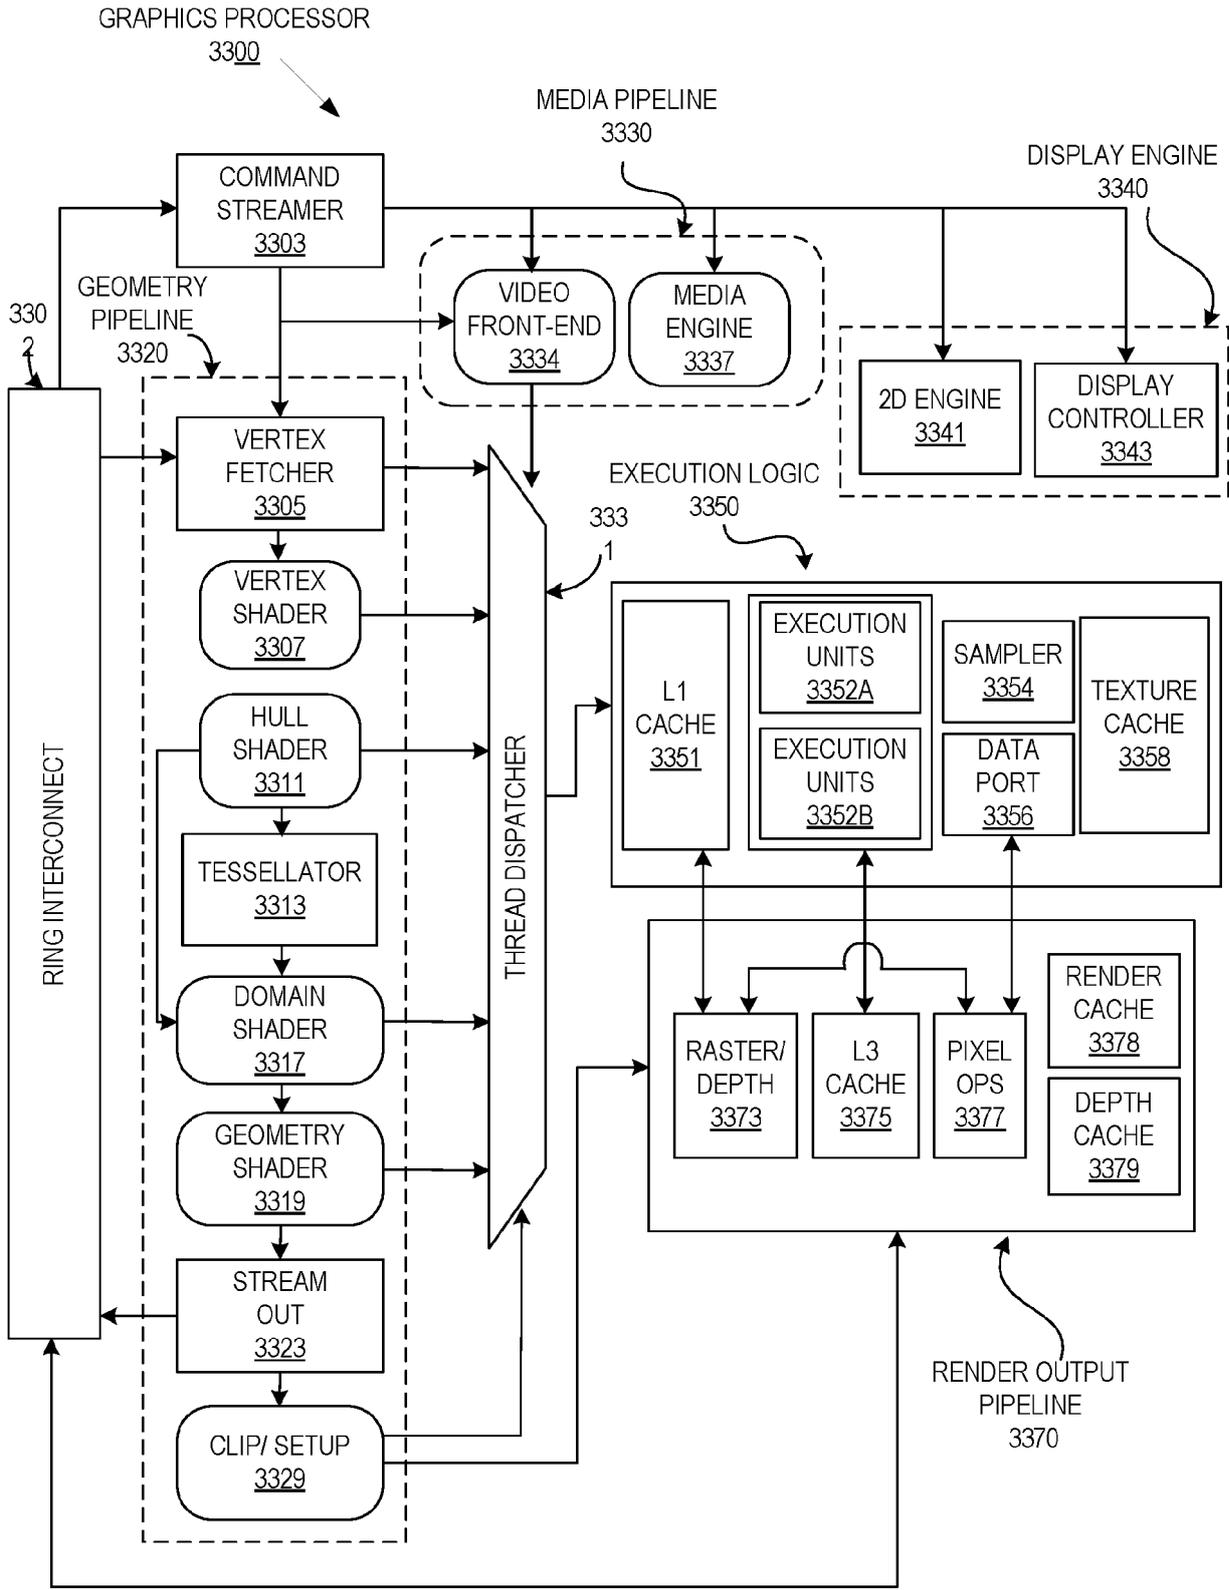


FIG. 33

FIG. 34A

GRAPHICS PROCESSOR COMMAND FORMAT

3400

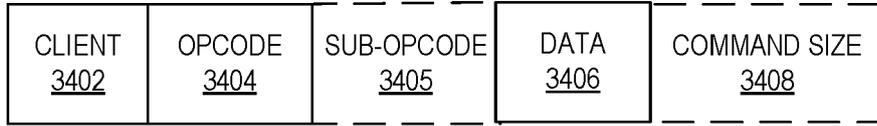
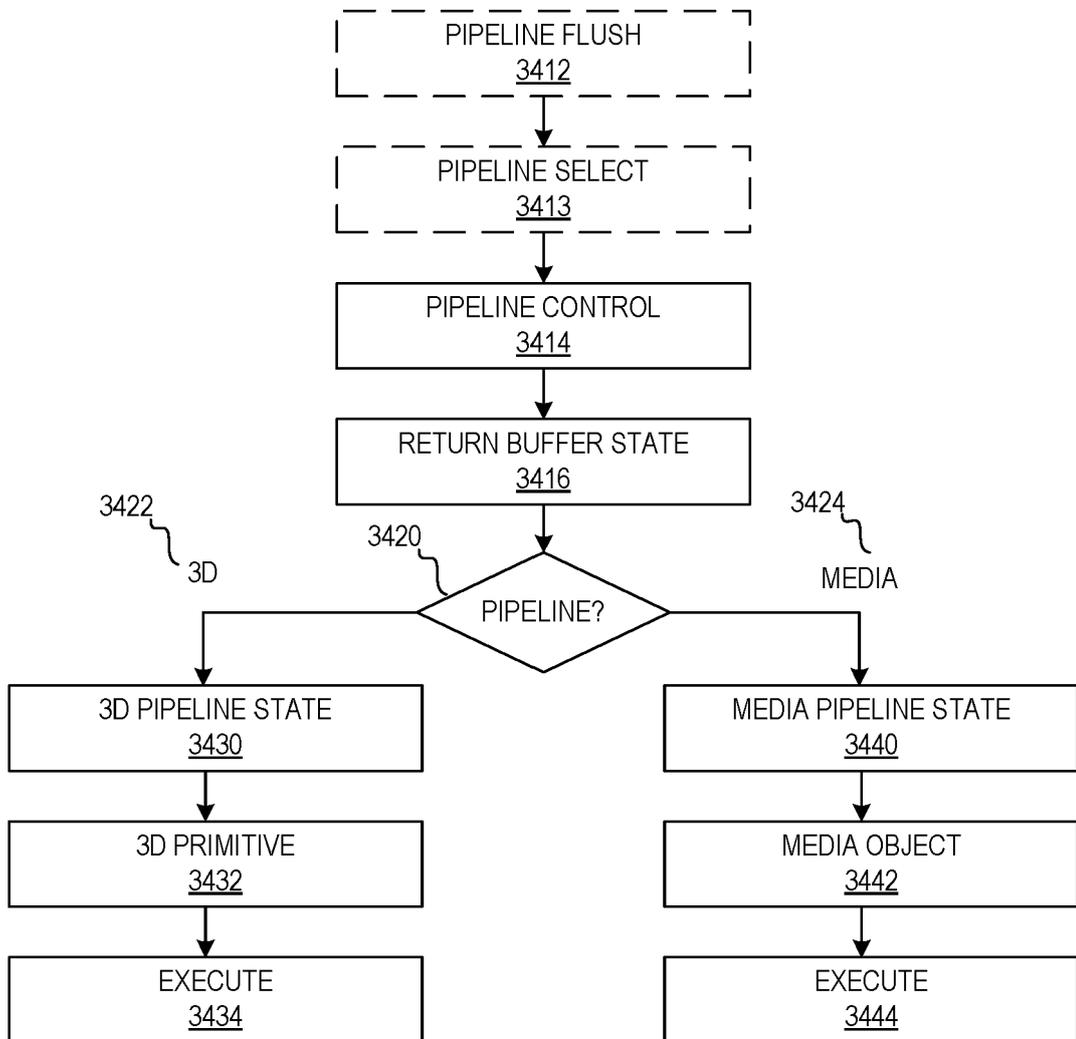


FIG. 34B

GRAPHICS PROCESSOR COMMAND SEQUENCE

3410



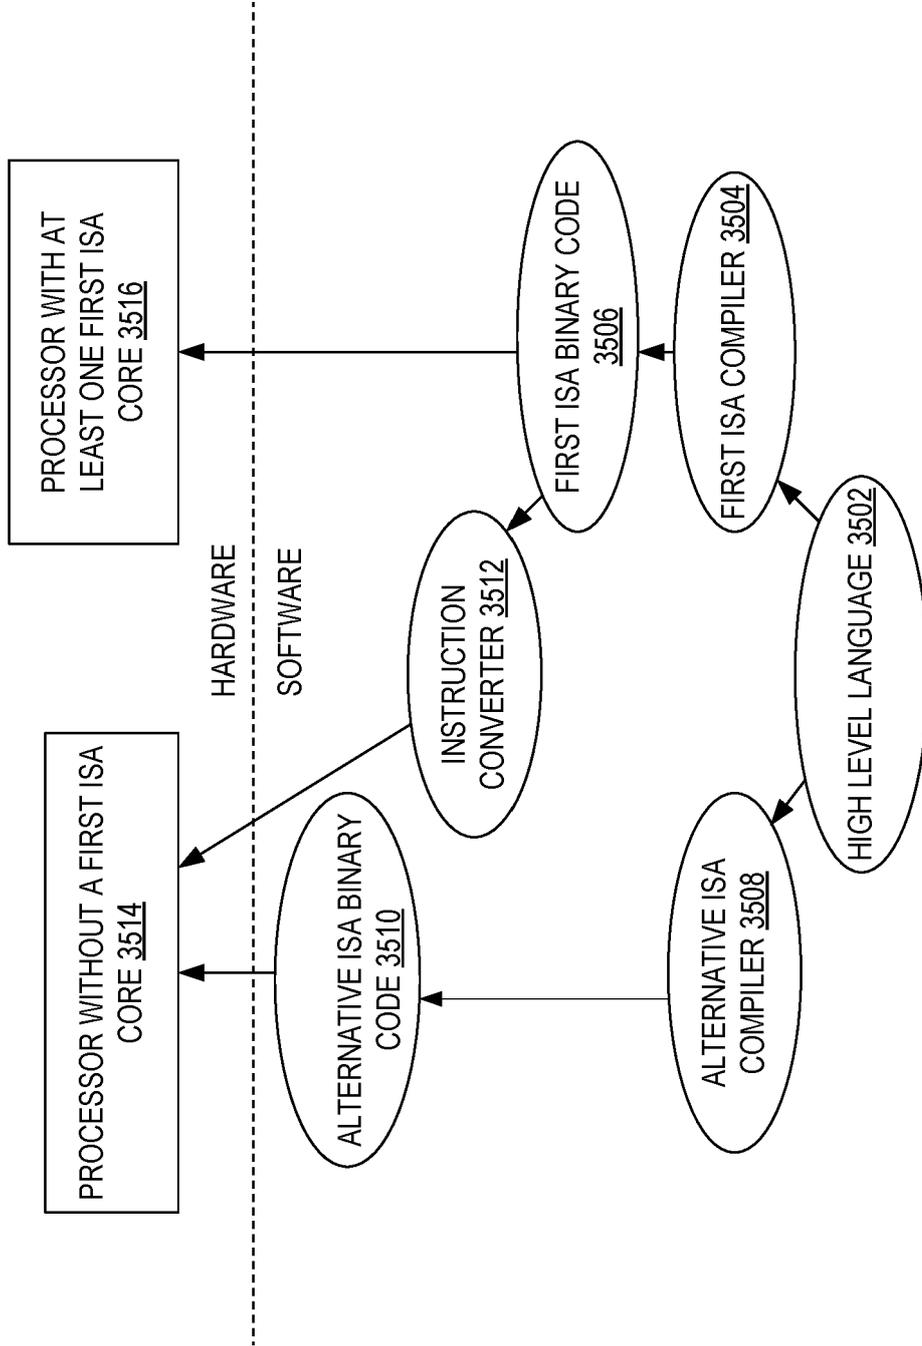


FIG. 35

IP CORE DEVELOPMENT - 3600

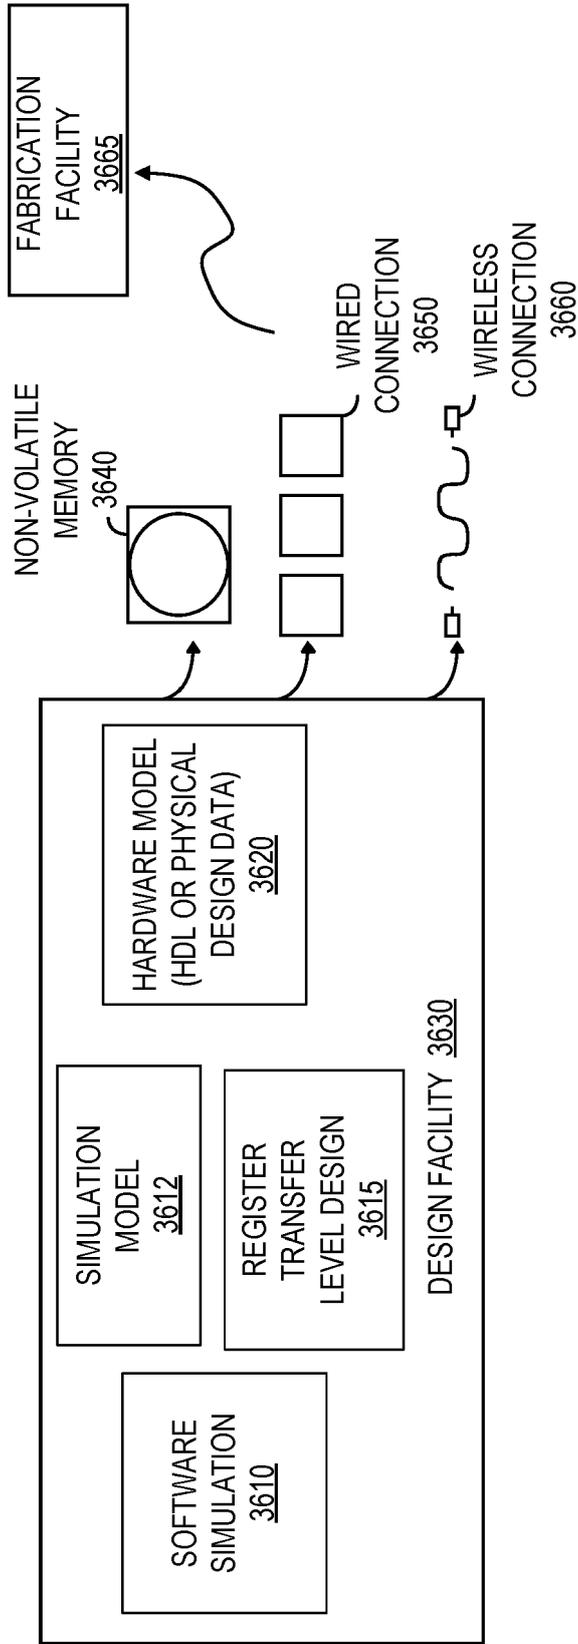


FIG. 36



EUROPEAN SEARCH REPORT

Application Number
EP 23 21 4578

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 7 941 651 B1 (TOLL BRET L [US] ET AL) 10 May 2011 (2011-05-10) * columns 1,3-12; figures 1-9b *	1-15	INV. G06F9/30
A	US 2011/208946 A1 (PROKOPENKO BORIS [US] ET AL) 25 August 2011 (2011-08-25) * paragraphs [0010], [0011], [0062] - [0233]; figures 1A - 31 *	1-15	
A	US 2010/299483 A1 (MCDONALD THOMAS C [US]) 25 November 2010 (2010-11-25) * the whole document *	1-15	
A	US 2004/039897 A1 (KUDO MAKOTO [JP]) 26 February 2004 (2004-02-26) * the whole document *	1-15	
			TECHNICAL FIELDS SEARCHED (IPC)
			G06F
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 17 May 2024	Examiner Freitas, Arthur
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03.82 (P04C01)

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 23 21 4578

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

17 - 05 - 2024

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7941651 B1	10-05-2011	NONE	
US 2011208946 A1	25-08-2011	CN 101359284 A CN 101615173 A TW I361379 B US 2007185953 A1 US 2007186082 A1 US 2011208946 A1	04-02-2009 30-12-2009 01-04-2012 09-08-2007 09-08-2007 25-08-2011
US 2010299483 A1	25-11-2010	TW 201042542 A US 2010299483 A1 US 2010299497 A1 US 2010299500 A1 US 2010299501 A1 US 2010299502 A1 US 2010299503 A1	01-12-2010 25-11-2010 25-11-2010 25-11-2010 25-11-2010 25-11-2010 25-11-2010
US 2004039897 A1	26-02-2004	CN 1490726 A JP 3627725 B2 JP 2004030015 A US 2004039897 A1	21-04-2004 09-03-2005 29-01-2004 26-02-2004