



- (51) **International Patent Classification:**  
*G06F 9/30* (2006.01)      *G06F 9/305* (2006.01)
- (21) **International Application Number:**  
PCT/US201 1/067087
- (22) **International Filing Date:**  
23 December 2011 (23.12.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant (for all designated States except US):** **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, MS: RNB-4-150, Sunnyvale, California 95052 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **VALENTINE, Robert** [US/IL]; Rechov Hadganiot 33-5, 36054 Kiryat Tivon (IL). **OULD-AHMED-VALL, Elmoustapha** [MR/US]; 5000 West Chandler Boulevard, MS: CH7-401, Chandler, Arizona 85226 (US). **CORBAL, Jesus** [ES/ES]; Jordi Girona 1-3 Intel Labs, E-08034 Barcelona (ES). **ULIEL, Tal** [IL/IL]; Imber 6/5, 64362 Tel Aviv (IL). **TOLL, Bret L.** [US/US]; 2868 NE Lorie Dr., Hillsboro, Oregon 97124 (US).
- (74) **Agents:** **WEBSTER, Thomas C.** et al; **BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP**, 1279 Oakmead Parkway, Sunnyvale, California 94085 (US).

- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

— of inventorship (Rule 4.17(ivf))

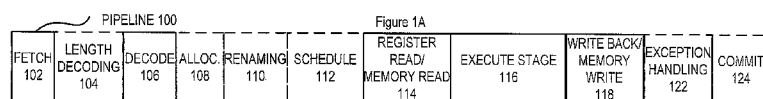
**Published:**

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



WO 2013/095610 A1

(54) **Title:** APPARATUS AND METHOD FOR SHUFFLING FLOATING POINT OR INTEGER VALUES



(57) **Abstract:** An apparatus and method are described for shuffling data elements from source registers to a destination register. For example, a method according to one embodiment includes the following operations: reading each mask bit stored in a mask data structure, the mask data structure containing mask bits associated with data elements of a destination register, the values usable for determining whether a masking operation or a shuffle operation should be performed on data elements stored within a first source register and a second source register; for each data element of the destination register, if a mask bit associated with the data element indicates that a shuffle operation should be performed, then shuffling data elements from the first source register and the second source register to the specified data element within the destination register; and if the mask bit indicates that a masking operation should be performed, then performing a specified masking operation with respect to the data element of the destination register.

## APPARATUS AND METHOD FOR SHUFFLING FLOATING POINT OR INTEGER VALUES

**Field of the Invention**

Embodiments of the invention relate generally to the field of computer systems. More particularly, the embodiments of the invention relate to an apparatus and method for shuffling floating point or integer values within a computer processor.

**Background**General Background

An instruction set, or instruction set architecture (ISA), is the part of the computer architecture related to programming, and may include the native data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O). The term instruction generally refers herein to macro-instructions - that is instructions that are provided to the processor (or instruction converter that translates (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morphs, emulates, or otherwise converts an instruction to one or more other instructions to be processed by the processor) for execution - as opposed to micro-instructions or micro-operations (micro-ops) - that is the result of a processor's decoder decoding macro-instructions.

The ISA is distinguished from the microarchitecture, which is the internal design of the processor implementing the instruction set. Processors with different microarchitectures can share a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale CA implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. For example, the same register architecture of the ISA may be implemented in different ways in different microarchitectures using well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB), and a retirement register file; the use of multiple maps and a pool of registers), etc. Unless otherwise specified, the phrases register architecture, register file, and register are used herein to refer to that which is visible to the software/programmer and the manner in which instructions specify registers. Where a specificity is desired, the adjective logical, architectural, or software visible will be used to indicate registers/files in the register architecture, while different adjectives will be used to

designation registers in a given microarchitecture (e.g., physical register, reorder buffer, retirement register, register pool).

An instruction set includes one or more instruction formats. A given instruction format defines various fields (number of bits, location of bits) to specify, among other things, the operation to be performed (opcode) and the operand(s) on which that operation is to be performed. Some instruction formats are further broken down through the definition of instruction templates (or subformats). 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 exemplary 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 (source 1/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.

Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) often require the same operation to be performed on a large number of data items (referred to as "data parallelism"). Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data items. SIMD technology is especially suited to processors that can logically divide the bits in a register into a number of fixed-sized data elements, each of which represents a separate value. For example, the bits in a 256-bit register may be specified as a source operand to be operated on as four separate 64-bit packed data elements (quad-word (Q) size data elements), eight separate 32-bit packed data elements (double word (D) 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). This type of data is referred to as packed data type or vector data type, and operands of this data type are referred to as packed data operands or vector operands. In other words, a packed data item or vector refers to a sequence of packed data elements, and a packed data operand or a vector operand is a source or destination operand of a SIMD instruction (also known as a packed data instruction or a vector instruction).

By way of example, one type of SIMD instruction specifies a single vector operation to be performed on two source vector operands in a vertical fashion to generate a destination vector operand (also referred to as a result vector operand) of the same size, with the same number of data elements, and in the same data element order. The data elements in the source vector operands are referred to as source data elements, while the data elements in the destination vector operand are referred to a destination or result data elements. These source vector operands are of the same size and contain data elements of the same width, and thus they contain the same number of data elements. The source data elements in the same bit positions in the two source vector operands form pairs of data elements (also referred to as corresponding data elements; that is, the data element in data element position 0 of each source operand correspond, the data element in data element position 1 of each source operand correspond, and so on). The operation specified by that SIMD instruction is performed separately on each of these pairs of source data elements to generate a matching number of result data elements, and thus each pair of source data elements has a corresponding result data element. Since the operation is vertical and since the result vector operand is the same size, has the same number of data elements, and the result data elements are stored in the same data element order as the source vector operands, the result data elements are in the same bit positions of the result vector operand as their corresponding pair of source data elements in the source vector operands. In addition to this exemplary type of SIMD instruction, there are a variety of other types of SIMD instructions (e.g., that has only one or has more than two source vector operands, that operate in a horizontal fashion, that generates a result vector operand that is of a different size, that has a different size data elements, and/or that has a different data element order). It should be understood that the term destination vector operand (or destination operand) is defined as the direct result of performing the operation specified by an instruction, including the storage of that destination operand at a location (be it a register or at a memory address specified by that instruction) so that it may be accessed as a source operand by another instruction (by specification of that same location by the another instruction).

The SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, has enabled a significant improvement in application performance. An additional set of SIMD extensions, referred to the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme, has been , has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developers Manual, October 2011; and see Intel® Advanced Vector Extensions Programming Reference, June 2011).

### Background Related to the Embodiments of the Invention

Shuffle instructions are used in current processor architectures to select data elements from two or more source registers and copy those data elements to different data element positions within a destination register. Current shuffle instructions, however, have not been implemented for use with conditional masking functionality as described herein and have not been implemented at a 256-bit granularity as described herein.

### **Brief Description of the Drawings**

10 **FIG. 1A** is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention;

**FIG. 1B** is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention;

**FIG. 2** is a block diagram of a single core processor and a multicore processor with integrated memory controller and graphics according to embodiments of the invention;

**FIG. 3** illustrates a block diagram of a system in accordance with one embodiment of the present invention;

20 **FIG. 4** illustrates a block diagram of a second system in accordance with an embodiment of the present invention;

**FIG. 5** illustrates a block diagram of a third system in accordance with an embodiment of the present invention;

**FIG. 6** illustrates a block diagram of a system on a chip (SoC) in accordance with an embodiment of the present invention;

**FIG. 7** illustrates a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention;

30 **FIG. 8** illustrates an apparatus for performing a shuffle operation according to one embodiment of the invention.

**FIG. 9** illustrates an apparatus for performing a shuffle operation according to another embodiment of the invention.

**FIG. 10** illustrates a method for performing a shuffle operation according to embodiments of the invention.

**FIGS. 11A and 11B** are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention;

**FIG. 12A-D** are block diagrams illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention; and

5 **FIG. 13** is a block diagram of a register architecture according to one embodiment of the invention;

**FIG. 14A** is a block diagram of a single processor core, along with its connection to the on-die interconnect network and with its local subset of the Level 2 (L2) cache, according to embodiments of the invention; and

10 **FIG. 14B** is an expanded view of part of the processor core in Figure 14A according to embodiments of the invention.

## Detailed Description

### Exemplary Processor Architectures and Data Types

15 **Figure 1A** is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention. **Figure 1B** is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention. The solid lined boxes in Figures 1A-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.

20 In Figure 1A, a processor pipeline 100 includes a fetch stage 102, a length decode stage 104, a decode stage 106, an allocation stage 108, a renaming stage 110, a scheduling (also known as a dispatch or issue) stage 112, a register read/memory read stage 114, an execute stage 116, a write back/memory write stage 118, an exception handling stage 122, and a commit stage 124.

30 Figure 1B shows processor core 190 including a front end unit 130 coupled to an execution engine unit 150, and both are coupled to a memory unit 170. The core 190 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 190 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.

The front end unit 130 includes a branch prediction unit 132 coupled to an instruction cache unit 134, which is coupled to an instruction translation lookaside buffer (TLB) 136, which is coupled to an instruction fetch unit 138, which is coupled to a decode unit 140. The decode unit 140 (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 unit 140 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. In one embodiment, the core 190 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 140 or otherwise within the front end unit 130). The decode unit 140 is coupled to a rename/allocator unit 152 in the execution engine unit 150.

The execution engine unit 150 includes the rename/allocator unit 152 coupled to a retirement unit 154 and a set of one or more scheduler unit(s) 156. The scheduler unit(s) 156 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 156 is coupled to the physical register file(s) unit(s) 158. Each of the physical register file(s) units 158 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 one embodiment, the physical register file(s) unit 158 comprises a vector registers unit, a write mask registers unit, and a scalar registers unit. These register units may provide architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) unit(s) 158 is overlapped by the retirement unit 154 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(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 154 and the physical register file(s) unit(s) 158 are coupled to the execution cluster(s) 160. The execution cluster(s) 160 includes a set of one or more execution units 162 and a set of one or more memory access units 164. The execution units 162 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all

perform all functions. The scheduler unit(s) 156, physical register file(s) unit(s) 158, and execution cluster(s) 160 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster - and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) 164). 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.

The set of memory access units 164 is coupled to the memory unit 170, which includes a data TLB unit 172 coupled to a data cache unit 174 coupled to a level 2 (L2) cache unit 176. In one exemplary embodiment, the memory access units 164 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 172 in the memory unit 170. The instruction cache unit 134 is further coupled to a level 2 (L2) cache unit 176 in the memory unit 170. The L2 cache unit 176 is coupled to one or more other levels of cache and eventually to a main memory.

By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 100 as follows: 1) the instruction fetch 138 performs the fetch and length decoding stages 102 and 104; 2) the decode unit 140 performs the decode stage 106; 3) the rename/allocator unit 152 performs the allocation stage 108 and renaming stage 110; 4) the scheduler unit(s) 156 performs the schedule stage 112; 5) the physical register file(s) unit(s) 158 and the memory unit 170 perform the register read/memory read stage 114; the execution cluster 160 perform the execute stage 116; 6) the memory unit 170 and the physical register file(s) unit(s) 158 perform the write back/memory write stage 118; 7) various units may be involved in the exception handling stage 122; and 8) the retirement unit 154 and the physical register file(s) unit(s) 158 perform the commit stage 124.

The core 190 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 190 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2, and/or some form of the generic vector friendly instruction format (U=0 and/or U=1), described below), thereby allowing the operations used by many multimedia applications to be performed using packed data.



It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor also includes separate instruction and data cache units 134/174 and a shared L2 cache unit 176, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

**Figure 2** is a block diagram of a processor 200 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention. The solid lined boxes in Figure 2 illustrate a processor 200 with a single core 202A, a system agent 210, a set of one or more bus controller units 216, while the optional addition of the dashed lined boxes illustrates an alternative processor 200 with multiple cores 202A-N, a set of one or more integrated memory controller unit(s) 214 in the system agent unit 210, and special purpose logic 208.

Thus, different implementations of the processor 200 may include: 1) a CPU with the special purpose logic 208 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 202A-N being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of the two); 2) a coprocessor with the cores 202A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 202A-N being a large number of general purpose in-order cores. Thus, the processor 200 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 200 may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units 206, and external memory (not shown) coupled to the set of integrated memory controller units 214. The set of shared cache units 206 may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof. While in one embodiment a ring based interconnect unit 212 interconnects the integrated graphics logic 208, the set of shared cache units 206, and the system agent unit 210/integrated memory controller unit(s) 214, alternative embodiments may use any number of well-known techniques for interconnecting such units. In one embodiment, coherency is maintained between one or more cache units 206 and cores 202-10 A-N.

In some embodiments, one or more of the cores 202A-N are capable of multi-threading. The system agent 210 includes those components coordinating and operating cores 202A-N. The system agent unit 210 may include for example a power control unit (PCU) and a display unit. The PCU may be or include logic and components needed for regulating the power state of the cores 202A-N and the integrated graphics logic 208. The display unit is for driving one or more externally connected displays. 15

The cores 202A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 202A-N may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set. 20

**Figures 3-6** are block diagrams of exemplary computer architectures. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable. 25

Referring now to **Figure 3**, shown is a block diagram of a system 300 in accordance with one embodiment of the present invention. The system 300 may include one or more processors 310, 315, which are coupled to a controller hub 320. In one embodiment the controller hub 320 includes a graphics memory controller hub (GMCH) 390 and an Input/Output Hub (IOH) 350 (which may be on separate chips); the GMCH 390 includes memory and graphics controllers to which are coupled memory 340 and a coprocessor 345; the IOH 350 is couples input/output (I/O) devices 360 to the GMCH 390. Alternatively, one or both of the memory and graphics 35

controllers are integrated within the processor (as described herein), the memory 340 and the coprocessor 345 are coupled directly to the processor 310, and the controller hub 320 in a single chip with the IOH 350.

The optional nature of additional processors 315 is denoted in Figure 3 with broken lines. Each processor 310, 315 may include one or more of the processing cores described herein and may be some version of the processor 200.

The memory 340 may be, for example, dynamic random access memory (DRAM), phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub 320 communicates with the processor(s) 310, 315 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection 395.

In one embodiment, the coprocessor 345 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like. In one embodiment, controller hub 320 may include an integrated graphics accelerator.

There can be a variety of differences between the physical resources 310, 315 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

In one embodiment, the processor 310 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 310 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 345. Accordingly, the processor 310 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 345. Coprocessor(s) 345 accept and execute the received coprocessor instructions.

Referring now to **Figure 4**, shown is a block diagram of a first more specific exemplary system 400 in accordance with an embodiment of the present invention. As shown in Figure 4, multiprocessor system 400 is a point-to-point interconnect system, and includes a first processor 470 and a second processor 480 coupled via a point-to-point interconnect 450. Each of processors 470 and 480 may be some version of the processor 200. In one embodiment of the invention, processors 470 and 480 are respectively processors 310 and 315, while coprocessor 438 is coprocessor 345. In another embodiment, processors 470 and 480 are respectively processor 310 coprocessor 345.

Processors 470 and 480 are shown including integrated memory controller (IMC) units 472 and 482, respectively. Processor 470 also includes as part of its bus controller units point-

to-point (P-P) interfaces 476 and 478; similarly, second processor 480 includes P-P interfaces 486 and 488. Processors 470, 480 may exchange information via a point-to-point (P-P) interface 450 using P-P interface circuits 478, 488. As shown in Figure 4, IMCs 472 and 482 couple the processors to respective memories, namely a memory 432 and a memory 434, which may be portions of main memory locally attached to the respective processors.

Processors 470, 480 may each exchange information with a chipset 490 via individual P-P interfaces 452, 454 using point to point interface circuits 476, 494, 486, 498. Chipset 490 may optionally exchange information with the coprocessor 438 via a high-performance interface 439. In one embodiment, the coprocessor 438 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

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

Chipset 490 may be coupled to a first bus 416 via an interface 496. In one embodiment, first bus 416 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

As shown in Figure 4, various I/O devices 414 may be coupled to first bus 416, along with a bus bridge 418 which couples first bus 416 to a second bus 420. In one embodiment, one or more additional processor(s) 415, such as coprocessors, high-throughput MIC processors, GPGPU's, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor, are coupled to first bus 416. In one embodiment, second bus 420 may be a low pin count (LPC) bus. Various devices may be coupled to a second bus 420 including, for example, a keyboard and/or mouse 422, communication devices 427 and a storage unit 428 such as a disk drive or other mass storage device which may include instructions/code and data 430, in one embodiment. Further, an audio I/O 424 may be coupled to the second bus 420. Note that other architectures are possible. For example, instead of the point-to-point architecture of Figure 4, a system may implement a multi-drop bus or other such architecture.

Referring now to **Figure 5**, shown is a block diagram of a second more specific exemplary system 500 in accordance with an embodiment of the present invention. Like elements in Figures 4 and 5 bear like reference numerals, and certain aspects of Figure 4 have been omitted from Figure 5 in order to avoid obscuring other aspects of Figure 5.

Figure 5 illustrates that the processors 470, 480 may include integrated memory and I/O control logic ("CL") 472 and 482, respectively. Thus, the CL 472, 482 include integrated memory controller units and include I/O control logic. Figure 5 illustrates that not only are the memories 432, 434 coupled to the CL 472, 482, but also that I/O devices 514 are also coupled to the control logic 472, 482. Legacy I/O devices 515 are coupled to the chipset 490.

Referring now to **Figure 6**, shown is a block diagram of a SoC 600 in accordance with an embodiment of the present invention. Similar elements in Figure 2 bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In Figure 6, an interconnect unit(s) 602 is coupled to: an application processor 610 which includes a set of one or more cores 202A-N and shared cache unit(s) 206; a system agent unit 210; a bus controller unit(s) 216; an integrated memory controller unit(s) 214; a set or one or more coprocessors 620 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 630; a direct memory access (DMA) unit 632; and a display unit 640 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 620 include a special-purpose processor, such as, for example, a network or communication processor, compression engine, GPGPU, a high-throughput MIC processor, embedded processor, or the like.

Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the invention 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.

Program code, such as code 430 illustrated in Figure 4, may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system includes any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

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.

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

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,  
10 compact disk read-only memories (CD-ROMs), compact disk rewritable's (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  
15 memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

Accordingly, embodiments of the invention 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  
20 system features described herein. Such embodiments may also be referred to as program products.

In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic  
25 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.

**Figure 7** is a block diagram contrasting the use of a software instruction converter to convert  
30 binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. Figure 7 shows a program in a high level language 702 may be compiled using an x86 compiler 704 to  
35 generate x86 binary code 706 that may be natively executed by a processor with at least one x86

instruction set core 716. The processor with at least one x86 instruction set core 716 represents any processor that can perform substantially the same functions as an Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. The x86 compiler 704 represents a compiler that is operable to generate x86 binary code 706 (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core 716.

Similarly, Figure 7 shows the program in the high level language 702 may be compiled using an alternative instruction set compiler 708 to generate alternative instruction set binary code 710 that may be natively executed by a processor without at least one x86 instruction set core 714 (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, CA and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, CA). The instruction converter 712 is used to convert the x86 binary code 706 into code that may be natively executed by the processor without an x86 instruction set core 714. This converted code is not likely to be the same as the alternative instruction set binary code 710 because an instruction converter capable of this is difficult to make; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter 712 represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code 706.

#### EMBODIMENTS OF THE INVENTION FOR PERFORMING A SHUFFLE OPERATION

Embodiments of the invention described below provide for the shuffling of floating point or integer data elements from source registers to a destination register using conditional masking and with 128-bit or 256-bit granularity. While these embodiments are described within the context of specific mask register values and output vector register sizes, the underlying principles of the invention are not limited to these implementations.

**Figure 8** illustrates shuffle logic 805 for performing shuffle operations on 256-bit operands 802-803 with 128-bit packed data elements A-D in accordance with one embodiment of the invention. Specifically, in this embodiment, the shuffle logic 805 selects a 128-bit data element, A or B, to be copied to a first 128-bit position of a destination register (DST) 804 and

selects a second 128-bit data element, C or D, to be copied to a second 128-bit position of the destination register 804 based on a 2-bit immediate value 805. In one embodiment, the first bit of the immediate value specifies whether A or B is shuffled into the first half of the destination register (e.g., 0 = A, 1 = B) and the second bit of the immediate value specifies whether C or D is shuffled into the second half of the destination register (e.g., 0 = C, 1 = D).

In one embodiment, a mask bit may also be specified within a mask data structure 870 for each of the destination register data elements. If the mask bit associated with a particular data element in the destination register is set to true, then the shuffle logic 805 shuffles the data elements through to the destination register as described above. If the mask bit is set to false then, in one embodiment, the shuffle logic writes all zeroes to the associated destination register entry. For example, if a mask bit associated with destination register bits 0:127 is set to a false value, then in one embodiment, the shuffle logic 805 writes all zeros to bits 0:127. The foregoing technique of writing zeroes to a destination data element in response to a mask value is referred to herein as "zeroing masking." Alternatively, one embodiment of the invention uses "merging masking" in which the previous data element values stored in the destination register are maintained. Thus, returning to the above example, if "merging masking" is used, bits 0:127 would maintain their prior values. Of course, masking bits described above may be reversed while still complying with the underlying principles of the invention (e.g., true = masking; false = no masking).

As illustrated in **Figure 8**, in one embodiment, the shuffle logic 805 accesses the registers to perform the above operations by controlling multiplexers 810-811. The logic required for implementing a multiplexer is well understood by those of ordinary skill in the art and will not be described here in detail.

**Figure 9** illustrates another embodiment of shuffle logic 805 which shuffles four 128-bit data elements, A-D and E-F, stored in two 512-bit source registers 902-903. The shuffle logic 805 selects any two of data elements A-D based on the value of the first four bits of an immediate value imm8 905 and selects any two of data elements E-H based on the value of the last four bits of the immediate value imm8 905. More specifically, in one embodiment, the first two bits of the immediate value specify the selection for data element 0:127 of the destination register 904; the next two bits of the immediate value specify the selection for data element 128:255 of the destination register 904; the next two bits of the immediate value specify the selection for data element 256:383 of the destination register 904; and the final two bits of the immediate value specify the selection for data element 384:511 of the destination register 904.

In one embodiment, a mask data structure 870 may be specified to provide masking operations as described above. For example, if "zeroing masking" is used, then the shuffle logic



805 writes all zeroes for data elements in the destination 904 associated with a false mask bit and performs the shuffle operations as described above for data elements in the destination 904 associated with a true mask bit. Alternatively, if "merging masking" is used, then the shuffle logic 805 maintains prior values for data elements in the destination 904 associated with a false mask bit and performs the shuffle operations as described above for data elements in the destination 904 associated with a true mask bit.

A method according to one embodiment of the invention is illustrated in **Figure 10**. The method may be executed within the context of the architecture shown in **Figures 8-9** but is not limited to any particular architecture.

At 1002, control variable N is set to zero and, at 1002, the immediate value is read to determine how to shuffle the data elements. At 1004, data element N is selected for updating in the destination register. If a writemask condition is set to a false value, determined at 1005, then the data element N in the destination register is updated based on the particular masking type. For example, if zeroing masking is used, then all zeros are written to the destination data element N.

If the writemask condition is set to a true value then, at 1007, either SRC1 or SRC2 based on the current value of N. For example, if data element N is in the first half of the destination register, then SRC1 may be selected and if data element N is in the second half of the destination register, then SRC2 may be selected. At 1008, the data element is selected from the source register and stored in as data element N in the destination register based on the immediate value (e.g., as described above). If the final data element has been updated in the destination, determined at 1010, the process terminates. Otherwise, the next data element N in the destination register is selected for updating at 1011 and the process returns to 1003.

The pseudocode describing the implementation of one embodiment of the invention is set forth below. It will be understood, however, that the underlying principles of the invention are not limited to the specific implementation described in the pseudocode.

```
VSHUFF64x2 (EVEX 512-bit version)
(KL, VL) = (4, 256), (8, 512)
FOR j ← 0 TO KL-1
    i ← j *64
    IF (EVEX.b == 1) AND (SRC2 *is memory*)
        THEN TMP_SRC2[i+63:i] ← SRC2[63:0]
        ELSE TMP_SRC2[i+63:i] ← SRC2[i+63:i]
    FI;
ENDFOR;
```

```

IF VL = 256
    TMP_DEST[127:0] ← Select2(SRCI[255:0], imm8[0]);
    TMP_DEST[255:128] ← Select2(SRC[255:0], imm8[1]);
FI;
5 IF VL = 512
    TMP_DEST[127:0] ← Select4(SRCI[511:0], imm8[1:0]);
    TMP_DEST[255:128] ← Select4(SRCI[511:0], imm8[3:2]);
    TMP_DEST[383:256] ← Select4(SRCI[511:0], imm8[5:4]);
    TMP_DEST[511:384] ← Select4(SRCI[511:0], imm8[7:6]);
10 FI;
FOR j ← 0 TO KL-1
    i ← j * 64
    IF kl[j] OR *no writemask*
        THEN DEST[i+63:i] ← TMP_DEST[i+63:i]
15 ELSE
        IF *merging-masking* ;merging-masking
            THEN *DEST[i+63:i] remains unchanged*
            ELSE *zeroing-masking* ;zeroing-masking
                THEN DEST [i+63:i] ← 0
20 FI
FI;
ENDFOR
DEST[MAX-VL-1:VL] ← 0

```

### 25 *Exemplary Instruction Formats*

Embodiments of the instruction(s) described herein may be embodied in different formats. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

30 A vector friendly instruction format is an instruction format that is suited for vector instructions (e.g., there are certain fields specific to vector operations). While embodiments are described in which both vector and scalar operations are supported through the vector friendly instruction format, alternative embodiments use only vector operations the vector friendly instruction format.

**Figures 11A-11B** are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention. Figure 11A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments of the invention; while Figure 11B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments of the invention. Specifically, a generic vector friendly instruction format 1100 for which are defined class A and class B instruction templates, both of which include no memory access 1105 instruction templates and memory access 1120 instruction templates. The term generic in the context of the vector friendly instruction format refers to the instruction format not being tied to any specific instruction set.

While embodiments of the invention will be described in which the vector friendly instruction format supports the following: a 64 byte vector operand length (or size) with 32 bit (4 byte) or 64 bit (8 byte) data element widths (or sizes) (and thus, a 64 byte vector consists of either 16 doubleword-size elements or alternatively, 8 quadword-size elements); a 64 byte vector operand length (or size) with 16 bit (2 byte) or 8 bit (1 byte) data element widths (or sizes); a 32 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); and a 16 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); alternative embodiments may support more, less and/or different vector operand sizes (e.g., 256 byte vector operands) with more, less, or different data element widths (e.g., 128 bit (16 byte) data element widths).

The class A instruction templates in Figure 11A include: 1) within the no memory access 1105 instruction templates there is shown a no memory access, full round control type operation 1110 instruction template and a no memory access, data transform type operation 1115 instruction template; and 2) within the memory access 1120 instruction templates there is shown a memory access, temporal 1125 instruction template and a memory access, non-temporal 1130 instruction template. The class B instruction templates in Figure 11B include: 1) within the no memory access 1105 instruction templates there is shown a no memory access, write mask control, partial round control type operation 1112 instruction template and a no memory access, write mask control, vsize type operation 1117 instruction template; and 2) within the memory access 1120 instruction templates there is shown a memory access, write mask control 1127 instruction template.

The generic vector friendly instruction format 1100 includes the following fields listed below in the order illustrated in Figures 11A-11B.

Format field 1140 - a specific value (an instruction format identifier value) in this field uniquely identifies the vector friendly instruction format, and thus occurrences of instructions in the vector friendly instruction format in instruction streams. As such, this field is optional in the sense that it is not needed for an instruction set that has only the generic vector friendly instruction format.

Base operation field 1142 - its content distinguishes different base operations.

Register index field 1144 - its content, directly or through address generation, specifies the locations of the source and destination operands, be they in registers or in memory. These include a sufficient number of bits to select N registers from a PxQ (e.g. 32x512, 16x128, 32x1024, 64x1024) register file. While in one embodiment N may be up to three sources and one destination register, alternative embodiments may support more or less sources and destination registers (e.g., may support up to two sources where one of these sources also acts as the destination, may support up to three sources where one of these sources also acts as the destination, may support up to two sources and one destination).

Modifier field 1146 - its content distinguishes occurrences of instructions in the generic vector instruction format that specify memory access from those that do not; that is, between no memory access 1105 instruction templates and memory access 1120 instruction templates. Memory access operations read and/or write to the memory hierarchy (in some cases specifying the source and/or destination addresses using values in registers), while non-memory access operations do not (e.g., the source and destinations are registers). While in one embodiment this field also selects between three different ways to perform memory address calculations, alternative embodiments may support more, less, or different ways to perform memory address calculations.

Augmentation operation field 1150 - its content distinguishes which one of a variety of different operations to be performed in addition to the base operation. This field is context specific. In one embodiment of the invention, this field is divided into a class field 1168, an alpha field 1152, and a beta field 1154. The augmentation operation field 1150 allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions.

Scale field 1160 - its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses  $2^{\text{scale}} * \text{index} + \text{base}$ ).

Displacement Field 1162A- its content is used as part of memory address generation (e.g., for address generation that uses  $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$ ).

Displacement Factor Field 1162B (note that the juxtaposition of displacement field 1162A directly over displacement factor field 1162B indicates one or the other is used) - its content is used as part of address generation; it specifies a displacement factor that is to be scaled

by the size of a memory access (N) - where N is the number of bytes in the memory access (e.g., for address generation that uses  $2^{\text{scale}} * \text{index} + \text{base} + \text{scaled displacement}$ ). Redundant low-order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) in order to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field 1174 (described herein) and the data manipulation field 1154C. The displacement field 1162A and the displacement factor field 1162B are optional in the sense that they are not used for the no memory access 1105 instruction templates and/or different embodiments may implement only one or none of the two.

Data element width field 1164 - its content distinguishes which one of a number of data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

Write mask field 1170 - its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. 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 one embodiment, 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 one embodiment, 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 write mask field 1170 allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While embodiments of the invention are described in which the write mask field's 1170 content selects one of a number of write mask registers that contains the write mask to be used (and thus the write mask field's 1170 content indirectly identifies that masking to be performed), alternative embodiments instead or additional allow the mask write field's 1170 content to directly specify the masking to be performed.

Immediate field 1172 - its content allows for the specification of an immediate. This field is optional in the sense that it is not present in an implementation of the generic vector friendly format that does not support immediate and it is not present in instructions that do not use an immediate.

5           Class field 1168 - its content distinguishes between different classes of instructions. With reference to Figures 11A-B, the contents of this field select between class A and class B instructions. In Figures 11A-B, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 1168A and class B 1168B for the class field 1168 respectively in Figures 11A-B).

## 10 *Instruction Templates of Class A*

In the case of the non-memory access 1105 instruction templates of class A, the alpha field 1152 is interpreted as an RS field 1152A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 1152A.1 and data transform 1152A.2 are respectively specified for the no memory access, round type operation 15 1110 and the no memory access, data transform type operation 1115 instruction templates), while the beta field 1154 distinguishes which of the operations of the specified type is to be performed. In the no memory access 1105 instruction templates, the scale field 1160, the displacement field 1162A, and the displacement scale field 1162B are not present.

### No-Memory Access Instruction Templates - Full Round Control Type Operation

20           In the no memory access full round control type operation 1110 instruction template, the beta field 1154 is interpreted as a round control field 1154A, whose content(s) provide static rounding. While in the described embodiments of the invention the round control field 1154A includes a suppress all floating point exceptions (SAE) field 1156 and a round operation control field 1158, alternative embodiments may support may encode both these concepts into the same 25 field or only have one or the other of these concepts/fields (e.g., may have only the round operation control field 1158).

SAE field 1156 - its content distinguishes whether or not to disable the exception event reporting; when the SAE field's 1156 content indicates suppression is enabled, a given instruction does not report any kind of floating-point exception flag and does not raise any 30 floating point exception handler.

Round operation control field 1158 - its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 1158 allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes

a control register for specifying rounding modes, the round operation control field's 1150 content overrides that register value.

### ***No Memory Access Instruction Templates - Data Transform Type Operation***

In the no memory access data transform type operation 1115 instruction template, the  
5 beta field 1154 is interpreted as a data transform field 1154B, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

In the case of a memory access 1120 instruction template of class A, the alpha field 1152  
10 is interpreted as an eviction hint field 1152B, whose content distinguishes which one of the eviction hints is to be used (in Figure 11A, temporal 1152B.1 and non-temporal 1152B.2 are respectively specified for the memory access, temporal 1125 instruction template and the memory access, non-temporal 1130 instruction template), while the beta field 1154 is interpreted as a data manipulation field 1154C, whose content distinguishes which one of a number of data  
15 manipulation operations (also known as primitives) is to be performed (e.g., no manipulation; broadcast; up conversion of a source; and down conversion of a destination). The memory access 1120 instruction templates include the scale field 1160, and optionally the displacement field 1162A or the displacement scale field 1162B.

Vector memory instructions perform vector loads from and vector stores to memory, with  
20 conversion support. As with regular vector instructions, vector memory instructions transfer data from/to memory in a data element-wise fashion, with the elements that are actually transferred is dictated by the contents of the vector mask that is selected as the write mask.

### **Memory Access Instruction Templates - Temporal**

Temporal data is data likely to be reused soon enough to benefit from caching. This is,  
25 however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

### ***Memory Access Instruction Templates - Non-Temporal***

Non-temporal data is data unlikely to be reused soon enough to benefit from caching in  
the 1st-level cache and should be given priority for eviction. This is, however, a hint, and  
different processors may implement it in different ways, including ignoring the hint entirely.

### ***Instruction Templates of Class B***

In the case of the instruction templates of class B, the alpha field 1152 is interpreted as a  
write mask control (Z) field 1152C, whose content distinguishes whether the write masking  
controlled by the write mask field 1170 should be a merging or a zeroing.

In the case of the non-memory access 1105 instruction templates of class B, part of the  
35 beta field 1154 is interpreted as an RL field 1157A, whose content distinguishes which one of

the different augmentation operation types are to be performed (e.g., round 1157A.1 and vector length (VSIZE) 1157A.2 are respectively specified for the no memory access, write mask control, partial round control type operation 1112 instruction template and the no memory access, write mask control, VSIZE type operation 1117 instruction template), while the rest of the beta field 1154 distinguishes which of the operations of the specified type is to be performed. In the no memory access 1105 instruction templates, the scale field 1160, the displacement field 1162A, and the displacement scale field 1162B are not present.

In the no memory access, write mask control, partial round control type operation 1110 instruction template, the rest of the beta field 1154 is interpreted as a round operation field 1159A and exception event reporting is disabled (a given instruction does not report any kind of floating-point exception flag and does not raise any floating point exception handler).

Round operation control field 1159A - just as round operation control field 1158, its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 1159A allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes a control register for specifying rounding modes, the round operation control field's 1150 content overrides that register value.

In the no memory access, write mask control, VSIZE type operation 1117 instruction template, the rest of the beta field 1154 is interpreted as a vector length field 1159B, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

In the case of a memory access 1120 instruction template of class B, part of the beta field 1154 is interpreted as a broadcast field 1157B, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field 1154 is interpreted the vector length field 1159B. The memory access 1120 instruction templates include the scale field 1160, and optionally the displacement field 1162A or the displacement scale field 1162B.

With regard to the generic vector friendly instruction format 1100, a full opcode field 1174 is shown including the format field 1140, the base operation field 1142, and the data element width field 1164. While one embodiment is shown where the full opcode field 1174 includes all of these fields, the full opcode field 1174 includes less than all of these fields in embodiments that do not support all of them. The full opcode field 1174 provides the operation code (opcode).



The augmentation operation field 1150, the data element width field 1164, and the write mask field 1170 allow these features to be specified on a per instruction basis in the generic vector friendly instruction format.

The combination of write mask field and data element width field create typed instructions in that they allow the mask to be applied based on different data element widths.

The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments of the invention, different processors or different cores within a processor may support only class A, only class B, or both classes. For instance, a high performance general purpose out-of-order core intended for general-purpose computing may support only class B, a core intended primarily for graphics and/or scientific (throughput) computing may support only class A, and a core intended for both may support both (of course, a core that has some mix of templates and instructions from both classes but not all templates and instructions from both classes is within the purview of the invention). Also, a single processor may include multiple cores, all of which support the same class or in which different cores support different class. For instance, in a processor with separate graphics and general purpose cores, one of the graphics cores intended primarily for graphics and/or scientific computing may support only class A, while one or more of the general purpose cores may be high performance general purpose cores with out of order execution and register renaming intended for general-purpose computing that support only class B. Another processor that does not have a separate graphics core, may include one more general purpose in-order or out-of-order cores that support both class A and class B. Of course, features from one class may also be implement in the other class in different embodiments of the invention. Programs written in a high level language would be put (e.g., just in time compiled or statically compiled) into an variety of different executable forms, including: 1) a form having only instructions of the class(es) supported by the target processor for execution; or 2) a form having alternative routines written using different combinations of the instructions of all classes and having control flow code that selects the routines to execute based on the instructions supported by the processor which is currently executing the code.

**Figure 12A-D** are block diagrams illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention. Figure 12 shows a specific vector friendly instruction format 1200 that is specific in the sense that it specifies the location, size, interpretation, and order of the fields, as well as values for some of those fields. The specific vector friendly instruction format 1200 may be used to extend the x86 instruction set, and thus some of the fields are similar or the same as those used in the existing x86 instruction set and extension thereof (e.g., AVX). This format remains consistent with the prefix encoding field,

real opcode byte field, MOD R/M field, SIB field, displacement field, and immediate fields of the existing x86 instruction set with extensions. The fields from Figure 11 into which the fields from Figure 12 map are illustrated.

It should be understood that, although embodiments of the invention are described with reference to the specific vector friendly instruction format 1200 in the context of the generic vector friendly instruction format 1100 for illustrative purposes, the invention is not limited to the specific vector friendly instruction format 1200 except where claimed. For example, the generic vector friendly instruction format 1100 contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format 1200 is shown as having fields of specific sizes. By way of specific example, while the data element width field 1164 is illustrated as a one bit field in the specific vector friendly instruction format 1200, the invention is not so limited (that is, the generic vector friendly instruction format 1100 contemplates other sizes of the data element width field 1164).

The generic vector friendly instruction format 1100 includes the following fields listed below in the order illustrated in Figure 12A.

EVEX Prefix (Bytes 0-3) 1202 - is encoded in a four-byte form.

Format Field 1140 (EVEX Byte 0, bits [7:0]) - the first byte (EVEX Byte 0) is the format field 1140 and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment of the invention).

The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

REX field 1205 (EVEX Byte 1, bits [7-5]) - consists of a EVEX.R bit field (EVEX Byte 1, bit [7] - R), EVEX.X bit field (EVEX byte 1, bit [6] - X), and 1157BEX byte 1, bit[5] - B). The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the corresponding VEX bit fields, and are encoded using Is complement form, i.e. ZMM0 is encoded as 1111B, ZMM15 is encoded as 0000B. Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding EVEX.R, EVEX.X, and EVEX.B.

REX' field 1110 - this is the first part of the REX' field 1110 and is the EVEX.R' bit field (EVEX Byte 1, bit [4] - R') that is used to encode either the upper 16 or lower 16 of the extended 32 register set. In one embodiment of the invention, this bit, along with others as indicated below, is stored in bit inverted format to distinguish (in the well-known x86 32-bit mode) from the BOUND instruction, whose real opcode byte is 62, but does not accept in the MOD R/M field (described below) the value of 11 in the MOD field; alternative embodiments of the invention do not store this and the other indicated bits below in the inverted format. A value

of 1 is used to encode the lower 16 registers. In other words, R'Rrrr is formed by combining EVEX.R', EVEX.R, and the other RRR from other fields.

Opcode map field 1215 (EVEX byte 1, bits [3:0] - mmmm) - its content encodes an implied leading opcode byte (OF, OF 38, or OF 3).

5 Data element width field 1164 (EVEX byte 2, bit [7] - W) - is represented by the notation EVEX.W. EVEX.W is used to define the granularity (size) of the datatype (either 32-bit data elements or 64-bit data elements).

EVEX.vvvv 1220 (EVEX Byte 2, bits [6:3]-vvvv)- the role of EVEX.vvvv may include the following: 1) EVEX.vvvv encodes the first source register operand, specified in inverted (Is complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in Is complement form for certain vector shifts; or 3) EVEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. Thus, EVEX.vvvv field 1220 encodes the 4 low-order bits of the first source register specifier stored in inverted (Is complement) form. Depending on the instruction, an extra  
10  
15 different EVEX bit field is used to extend the specifier size to 32 registers.

EVEX.U 1168 Class field (EVEX byte 2, bit [2]-U) - If EVEX.U = 0, it indicates class A or EVEX.U0; if EVEX.U = 1, it indicates class B or EVEX.U1.

Prefix encoding field 1225 (EVEX byte 2, bits [1:0]-pp) - provides additional bits for the base operation field. In addition to providing support for the legacy SSE instructions in the  
20 EVEX prefix format, this also has the benefit of compacting the SIMD prefix (rather than requiring a byte to express the SIMD prefix, the EVEX prefix requires only 2 bits). In one embodiment, to support legacy SSE instructions that use a SIMD prefix (66H, F2H, F3H) in both the legacy format and in the EVEX prefix format, these legacy SIMD prefixes are encoded into the SIMD prefix encoding field; and at runtime are expanded into the legacy SIMD prefix prior  
25 to being provided to the decoder's PLA (so the PLA can execute both the legacy and EVEX format of these legacy instructions without modification). Although newer instructions could use the EVEX prefix encoding field's content directly as an opcode extension, certain embodiments expand in a similar fashion for consistency but allow for different meanings to be specified by these legacy SIMD prefixes. An alternative embodiment may redesign the PLA to  
30 support the 2 bit SIMD prefix encodings, and thus not require the expansion.

Alpha field 1152 (EVEX byte 3, bit [7] - EH; also known as EVEX.EH, EVEX.rs, EVEX.RL, EVEX.write mask control, and EVEX.N; also illustrated with a) - as previously described, this field is context specific.

Beta field 1154 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s<sub>2-0</sub>, EVEX.r<sub>2-0</sub>, EVEX.r<sub>rl</sub>, EVEX.LLO, EVEX.LLB; also illustrated with βββ) - as previously described, this field is context specific.

5 REX' field 1110 - this is the remainder of the REX' field and is the EVEX.V bit field (EVEX Byte 3, bit [3] - V ) that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V, EVEX.vvvv.

10 Write mask field 1170 (EVEX byte 3, bits [2:0]-kkk) - its content specifies the index of a register in the write mask registers as previously described. In one embodiment of the invention, the specific value EVEX.kkk=000 has a special behavior implying no write mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

15 Real Opcode Field 1230 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

20 MOD R/M Field 1240 (Byte 5) includes MOD field 1242, Reg field 1244, and R/M field 1246. As previously described, the MOD field's 1242 content distinguishes between memory access and non-memory access operations. The role of Reg field 1244 can be summarized to two situations: encoding either the destination register operand or a source register operand, or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 1246 may include the following: encoding the instruction operand that references a memory address, or encoding either the destination register operand or a source register operand.

25 Scale, Index, Base (SIB) Byte (Byte 6) - As previously described, the scale field's 1150 content is used for memory address generation. SIB.xxx 1254 and SIB.bbb 1256 - the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb.

Displacement field 1162A (Bytes 7-10) - when MOD field 1242 contains 10, bytes 7-10 are the displacement field 1162A, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

30 Displacement factor field 1162B (Byte 7) - when MOD field 1242 contains 01, byte 7 is the displacement factor field 1162B. The location of this field is that same as that of the legacy x86 instruction set 8-bit displacement (disp8), which works at byte granularity. Since disp8 is sign extended, it can only address between -128 and 127 bytes offsets; in terms of 64 byte cache lines, disp8 uses 8 bits that can be set to only four really useful values -128, -64, 0, and 64; since  
35 a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to

disp8 and disp32, the displacement factor field 1162B is a reinterpretation of disp8; when using displacement factor field 1162B, the actual displacement is determined by the content of the displacement factor field multiplied by the size of the memory operand access (N). This type of displacement is referred to as disp8\*N. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement is based on the assumption that the effective displacement is multiple of the granularity of the memory access, and hence, the redundant low-order bits of the address offset do not need to be encoded. In other words, the displacement factor field 1162B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 1162B is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to disp8\*N. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset).

Immediate field 1172 operates as previously described.

### ***Full Opcode Field***

Figure 12B is a block diagram illustrating the fields of the specific vector friendly instruction format 1200 that make up the full opcode field 1174 according to one embodiment of the invention. Specifically, the full opcode field 1174 includes the format field 1140, the base operation field 1142, and the data element width (W) field 1164. The base operation field 1142 includes the prefix encoding field 1225, the opcode map field 1215, and the real opcode field 1230.

### ***Register Index Field***

Figure 12C is a block diagram illustrating the fields of the specific vector friendly instruction format 1200 that make up the register index field 1144 according to one embodiment of the invention. Specifically, the register index field 1144 includes the REX field 1205, the REX' field 1210, the MODR/M.reg field 1244, the MODR/M.r/m field 1246, the W W field 1220, xxx field 1254, and the bbb field 1256.

### ***Augmentation Operation Field***

Figure 12D is a block diagram illustrating the fields of the specific vector friendly instruction format 1200 that make up the augmentation operation field 1150 according to one embodiment of the invention. When the class (U) field 1168 contains 0, it signifies EVEX.U0 (class A 1168A); when it contains 1, it signifies EVEX.U1 (class B 1168B). When U=0 and the MOD field 1242 contains 11 (signifying a no memory access operation), the alpha field 1152

(EVEX byte 3, bit [7] - EH) is interpreted as the rs field 1152A. When the rs field 1152A contains a 1 (round 1152A.1), the beta field 1154 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the round control field 1154A. The round control field 1154A includes a one bit SAE field 1156 and a two bit round operation field 1158. When the rs field 1152A contains a 0 (data transform 1152A.2), the beta field 1154 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data transform field 1154B. When U=0 and the MOD field 1242 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 1152 (EVEX byte 3, bit [7] - EH) is interpreted as the eviction hint (EH) field 1152B and the beta field 1154 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data manipulation field 1154C.

When U=1, the alpha field 1152 (EVEX byte 3, bit [7] - EH) is interpreted as the write mask control (Z) field 1152C. When U=1 and the MOD field 1242 contains 11 (signifying a no memory access operation), part of the beta field 1154 (EVEX byte 3, bit [4]- So) is interpreted as the RL field 1157A; when it contains a 1 (round 1157A.1) the rest of the beta field 1154 (EVEX byte 3, bit [6-5]- S<sub>2-i</sub>) is interpreted as the round operation field 1159A, while when the RL field 1157A contains a 0 (VSIZE 1157.A2) the rest of the beta field 1154 (EVEX byte 3, bit [6-5]- S<sub>2-1</sub>) is interpreted as the vector length field 1159B (EVEX byte 3, bit [6-5]- L<sub>1-0</sub>). When U=1 and the MOD field 1242 contains 00, 01, or 10 (signifying a memory access operation), the beta field 1154 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the vector length field 1159B (EVEX byte 3, bit [6-5]- L<sub>1-0</sub>) and the broadcast field 1157B (EVEX byte 3, bit [4]- B).

**Figure 13** is a block diagram of a register architecture 1300 according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers 1310 that are 512 bits wide; these registers are referenced as zmm0 through zmm31. The lower order 256 bits of the lower 16 zmm registers are overlaid on registers ymm0-16. The lower order 128 bits of the lower 16 zmm registers (the lower order 128 bits of the ymm registers) are overlaid on registers xmm0-15. The specific vector friendly instruction format 1200 operates on these overlaid register file as illustrated in the below tables.

Adjustable Vector Length	Class	Operations	Registers
Instruction Templates that do not include the vector length field 1159B	A (Figure 11A; U=0)	1110, 1115, 1125, 1130	zmm registers (the vector length is 64 byte)
	B (Figure 11B; U=1)	1112	zmm registers (the vector length is 64 byte)

Instruction Templates that do include the vector length field 1159B	B (Figure 11B; U=l)	1117, 1127	zmm, ymm, or xmm registers (the vector length is 64 byte, 32 byte, or 16 byte) depending on the vector length field 1159B
---	---------------------	------------	---

In other words, the vector length field 1159B 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; and instructions templates without the vector length field 1159B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 1200 operate on packed or scalar single/double-precision floating point data and packed or scalar integer data. Scalar operations are operations performed on the lowest order data element position in an 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 embodiment.

Write mask registers 1315 - in the embodiment illustrated, there are 8 write mask registers (k0 through k7), each 64 bits in size. In an alternate embodiment, the write mask registers 1315 are 16 bits in size. As previously described, in one embodiment of the invention, the vector mask register k0 cannot be used as a write mask; when the encoding that would normally indicate k0 is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

General-purpose registers 1325 - in the embodiment illustrated, there are sixteen 64-bit general-purpose registers that are used along with the existing x86 addressing modes to address memory operands. These registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

Scalar floating point stack register file (x87 stack) 1345, on which is aliased the MMX packed integer flat register file 1350 - in the embodiment illustrated, the x87 stack is an eight-element stack used to perform scalar floating-point operations on 32/64/80-bit floating point data using the x87 instruction set extension; while the MMX registers are used 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.

Alternative embodiments of the invention may use wider or narrower registers. Additionally, alternative embodiments of the invention may use more, less, or different register files and registers.

**Figures 14A-B** illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip. The logic blocks communicate through a high-bandwidth interconnect network (e.g., a ring network) with some fixed function logic, memory I/O interfaces, and other necessary I/O logic, depending on the application.

Figure 14A is a block diagram of a single processor core, along with its connection to the on-die interconnect network 1402 and with its local subset of the Level 2 (L2) cache 1404, according to embodiments of the invention. In one embodiment, an instruction decoder 1400 supports the x86 instruction set with a packed data instruction set extension. An LI cache 1406 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 1408 and a vector unit 1410 use separate register sets (respectively, scalar registers 1412 and vector registers 1414) and data transferred between them is written to memory and then read back in from a level 1 (LI) cache 1406, alternative embodiments of the invention may use a different approach (e.g., use a single register set or include a communication path that allow data to be transferred between the two register files without being written and read back).

The local subset of the L2 cache 1404 is part of a global L2 cache that is divided into separate local subsets, one per processor core. Each processor core has a direct access path to its own local subset of the L2 cache 1404. Data read by a processor core is stored in its L2 cache subset 1404 and can be accessed quickly, in parallel with other processor cores accessing their own local L2 cache subsets. Data written by a processor core is stored in its own L2 cache subset 1404 and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data. The ring network is bi-directional to allow agents such as processor cores, L2 caches and other logic blocks to communicate with each other within the chip. Each ring data-path is 1012-bits wide per direction.

Figure 14B is an expanded view of part of the processor core in Figure 14A according to embodiments of the invention. Figure 14B includes an LI data cache 1406A part of the LI cache 1404, as well as more detail regarding the vector unit 1410 and the vector registers 1414. Specifically, the vector unit 1410 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 1428), which executes one or more of integer, single-precision float, and double-precision



float instructions. The VPU supports swizzling the register inputs with swizzle unit 1420, numeric conversion with numeric convert units 1422A-B, and replication with replication unit 1424 on the memory input. Write mask registers 1426 allow predicating resulting vector writes.

Embodiments of the invention may include various steps, which have been described above. The steps may be embodied in machine-executable instructions which may be used to cause a general-purpose or special-purpose processor to perform the steps. Alternatively, these steps may be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

As described herein, instructions may refer to specific configurations of hardware such as application specific integrated circuits (ASICs) configured to perform certain operations or having a predetermined functionality or software instructions stored in memory embodied in a non-transitory computer readable medium. Thus, the techniques shown in the figures can be implemented using code and data stored and executed on one or more electronic devices (e.g., an end station, a network element, etc.). Such electronic devices store and communicate (internally and/or with other electronic devices over a network) code and data using computer machine-readable media, such as non-transitory computer machine-readable storage media (e.g., magnetic disks; optical disks; random access memory; read only memory; flash memory devices; phase-change memory) and transitory computer machine-readable communication media (e.g., electrical, optical, acoustical or other form of propagated signals - such as carrier waves, infrared signals, digital signals, etc.). In addition, such electronic devices typically include a set of one or more processors coupled to one or more other components, such as one or more storage devices (non-transitory machine-readable storage media), user input/output devices (e.g., a keyboard, a touchscreen, and/or a display), and network connections. The coupling of the set of processors and other components is typically through one or more busses and bridges (also termed as bus controllers). The storage device and signals carrying the network traffic respectively represent one or more machine-readable storage media and machine-readable communication media. Thus, the storage device of a given electronic device typically stores code and/or data for execution on the set of one or more processors of that electronic device. Of course, one or more parts of an embodiment of the invention may be implemented using different combinations of software, firmware, and/or hardware. Throughout this detailed description, for the purposes of explanation, numerous specific details were set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the invention may be practiced without some of these specific details. In certain instances, well known structures and functions were not described in elaborate detail in order to avoid obscuring

the subject matter of the present invention. Accordingly, the scope and spirit of the invention should be judged in terms of the claims which follow.

**CLAIMS**

We claim:

1. A processor to execute one or more instructions to perform the operations of:  
reading each mask bit stored in a mask data structure, the mask data structure containing  
5 mask bits associated with data elements of a destination register, the values usable for  
determining whether a masking operation or a shuffle operation should be performed on data  
elements stored within a first source register and a second source register;  
for each data element of the destination register, if a mask bit associated with the data  
element indicates that a shuffle operation should be performed, then shuffling data elements  
10 from the first source register and the second source register to the specified data element within  
the destination register; and  
if the mask bit indicates that a masking operation should be performed, then performing a  
specified masking operation with respect to the data element of the destination register.
- 15 2. The processor as in claim 1 wherein the specified masking operation comprises  
setting the bits of the data element in the destination register all equal to zero.
3. The processor as in claim 1 wherein the specified masking operation comprises  
maintaining existing values previously stored in the data element in the destination register.  
20
4. The processor as in claim 1 wherein each data element in the destination register  
comprises 128 bits of data.
5. The processor as in claim 4 wherein each data element in the first and second  
25 source registers comprises 128 bits of data.
6. A method comprising:  
reading each mask bit stored in a mask data structure, the mask data structure containing  
mask bits associated with data elements of a destination register, the values usable for  
30 determining whether a masking operation or a shuffle operation should be performed on data  
elements stored within a first source register and a second source register;  
for each data element of the destination register, if a mask bit associated with the data  
element indicates that a shuffle operation should be performed, then shuffling data elements  
from the first source register and the second source register to the specified data element within  
35 the destination register; and

if the mask bit indicates that a masking operation should be performed, then performing a specified masking operation with respect to the data element of the destination register.

7. The method as in claim 6 wherein the specified masking operation comprises setting the bits of the data element in the destination register all equal to zero.

8. The method as in claim 6 wherein the specified masking operation comprises maintaining existing values previously stored in the data element in the destination register.

9. The method as in claim 6 wherein each data element in the destination register comprises 128 bits of data.

10. The method as in claim 9 wherein each data element in the first and second source registers comprises 128 bits of data.

11. A processor comprising:

means for reading each mask bit stored in a mask data structure, the mask data structure containing mask bits associated with data elements of a destination register, the values usable for determining whether a masking operation or a shuffle operation should be performed on data elements stored within a first source register and a second source register;

means for shuffling data elements from the first source register and the second source register to a specified data element within the destination register, wherein, for each data element of the destination register, if a mask bit associated with the data element indicates that a shuffle operation should be performed, then shuffling data elements, wherein if the mask bit indicates that a masking operation should be performed, then performing a specified masking operation with respect to the data element of the destination register.

12. The processor as in claim 11 wherein the specified masking operation comprises setting the bits of the data element in the destination register all equal to zero.

13. The processor as in claim 11 wherein the specified masking operation comprises maintaining existing values previously stored in the data element in the destination register.

14. The processor as in claim 11 wherein each data element in the destination register comprises 128 bits of data.

35

15. The processor as in claim 14 wherein each data element in the first and second source registers comprises 128 bits of data.

5 16. A computer system comprising:  
a memory for storing program instructions and data;  
a processor to execute one or more of the program instructions to perform the operations of:

10 reading each mask bit stored in a mask data structure, the mask data structure containing mask bits associated with data elements of a destination register, the values usable for determining whether a masking operation or a shuffle operation should be performed on data elements stored within a first source register and a second source register;

15 for each data element of the destination register, if a mask bit associated with the data element indicates that a shuffle operation should be performed, then shuffling data elements from the first source register and the second source register to the specified data element within the destination register, and

if the mask bit indicates that a masking operation should be performed, then performing a specified masking operation with respect to the data element of the destination register.

20 17. The system as in claim 16 wherein the specified masking operation comprises setting the bits of the data element in the destination register all equal to zero.

18. The system as in claim 16 wherein the specified masking operation comprises maintaining existing values previously stored in the data element in the destination register.

25 19. The system as in claim 16 wherein each data element in the destination register comprises 128 bits of data.

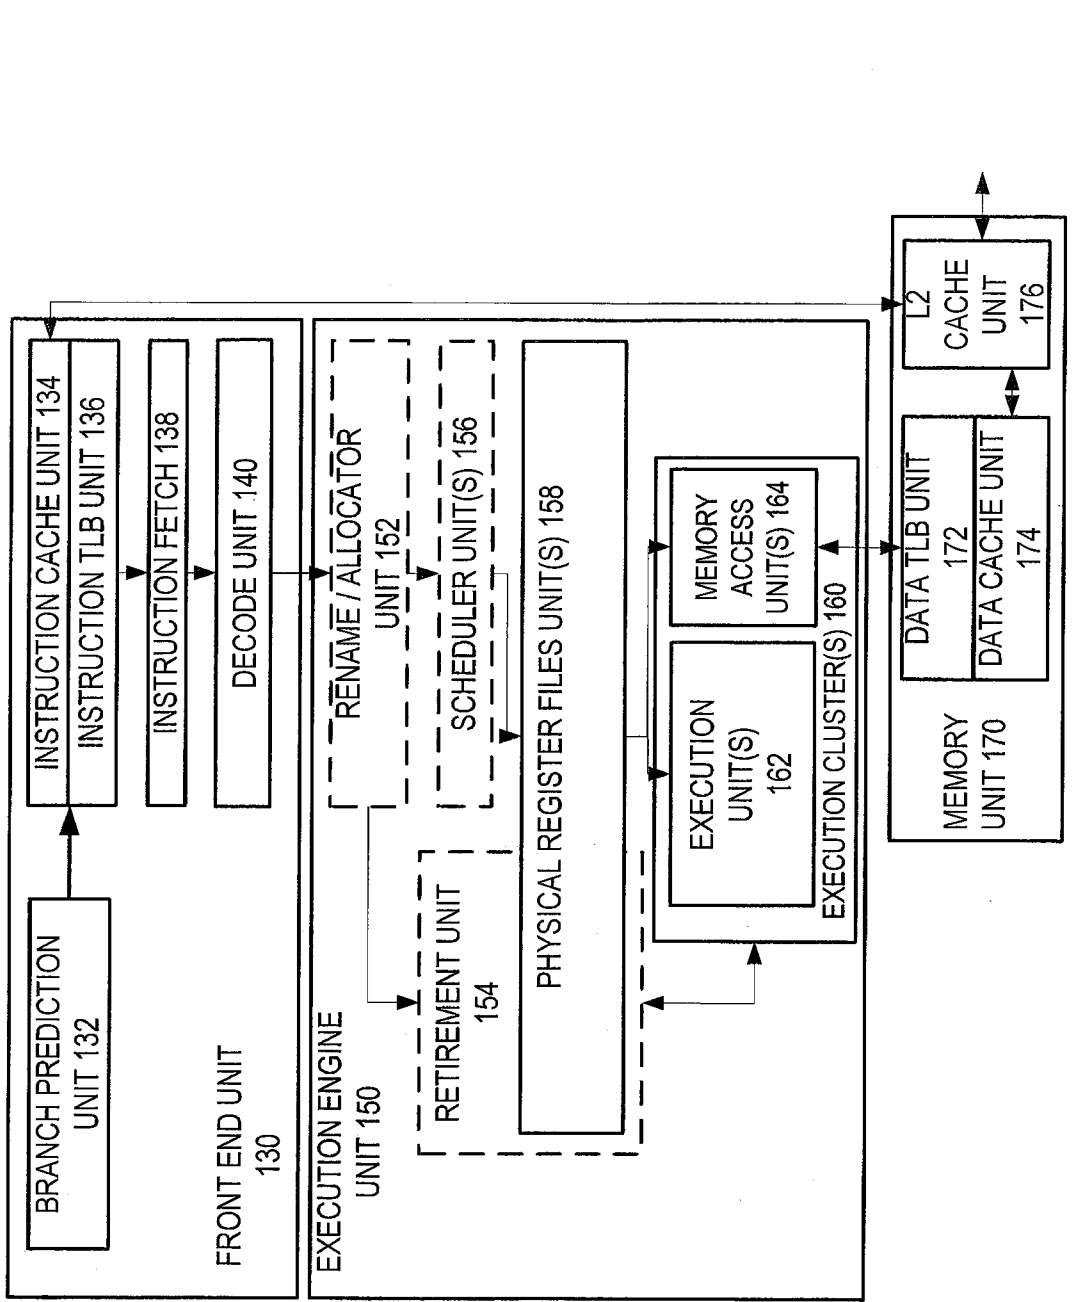
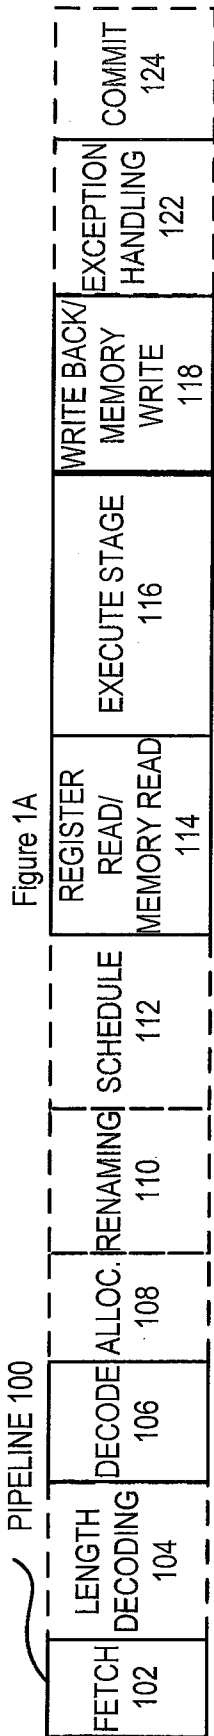
30 20. The system as in claim 19 wherein each data element in the first and second source registers comprises 128 bits of data.

21. The system as in claim 15 further comprising:  
a display adapter to render graphics images in response to execution of the program code by the processor.

35 22. The system as in claim 21 further comprising:

a user input interface to receive control signals **from** a user input device, the processor executing the program code in response to the control signals.

5



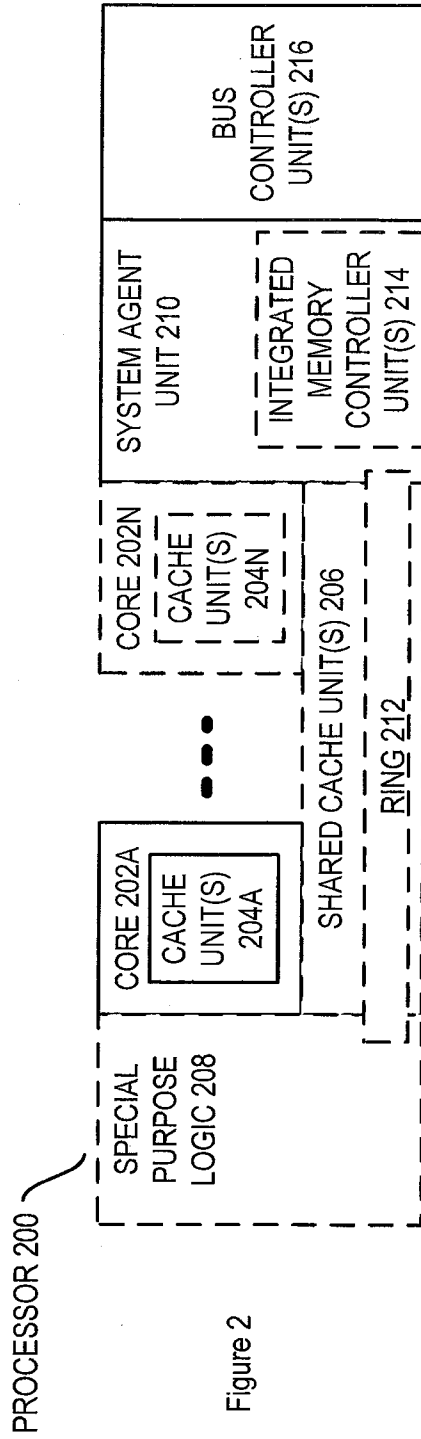


Figure 2



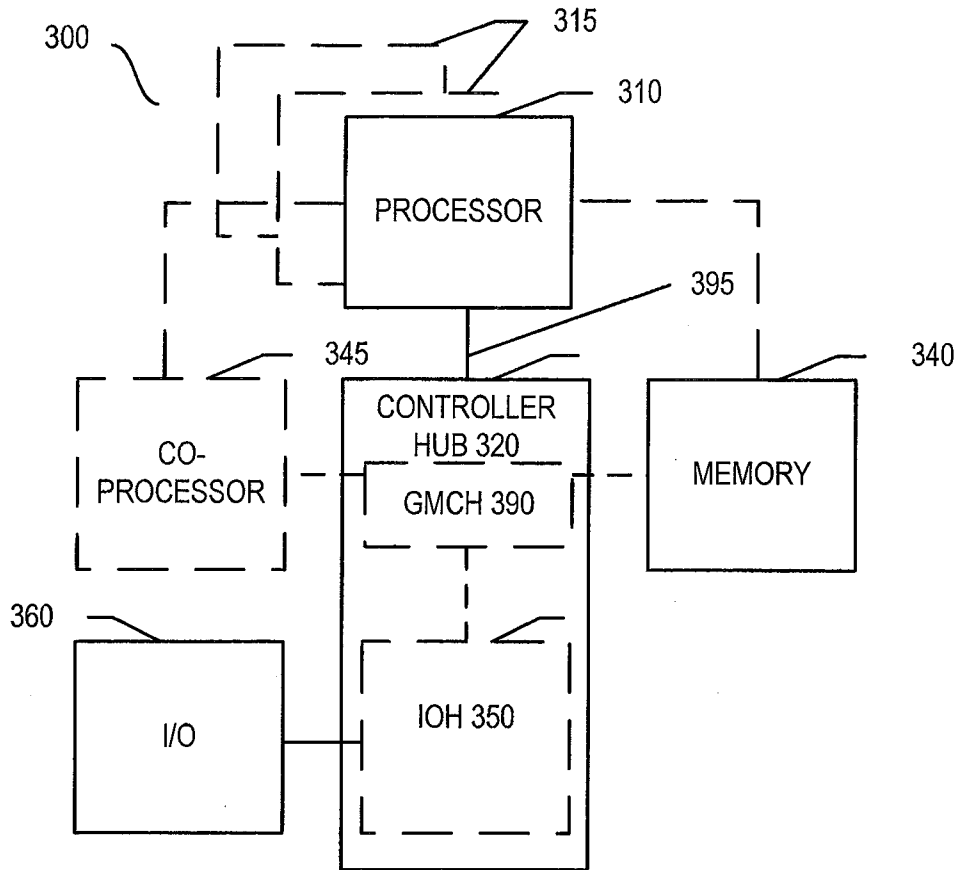


Figure 3

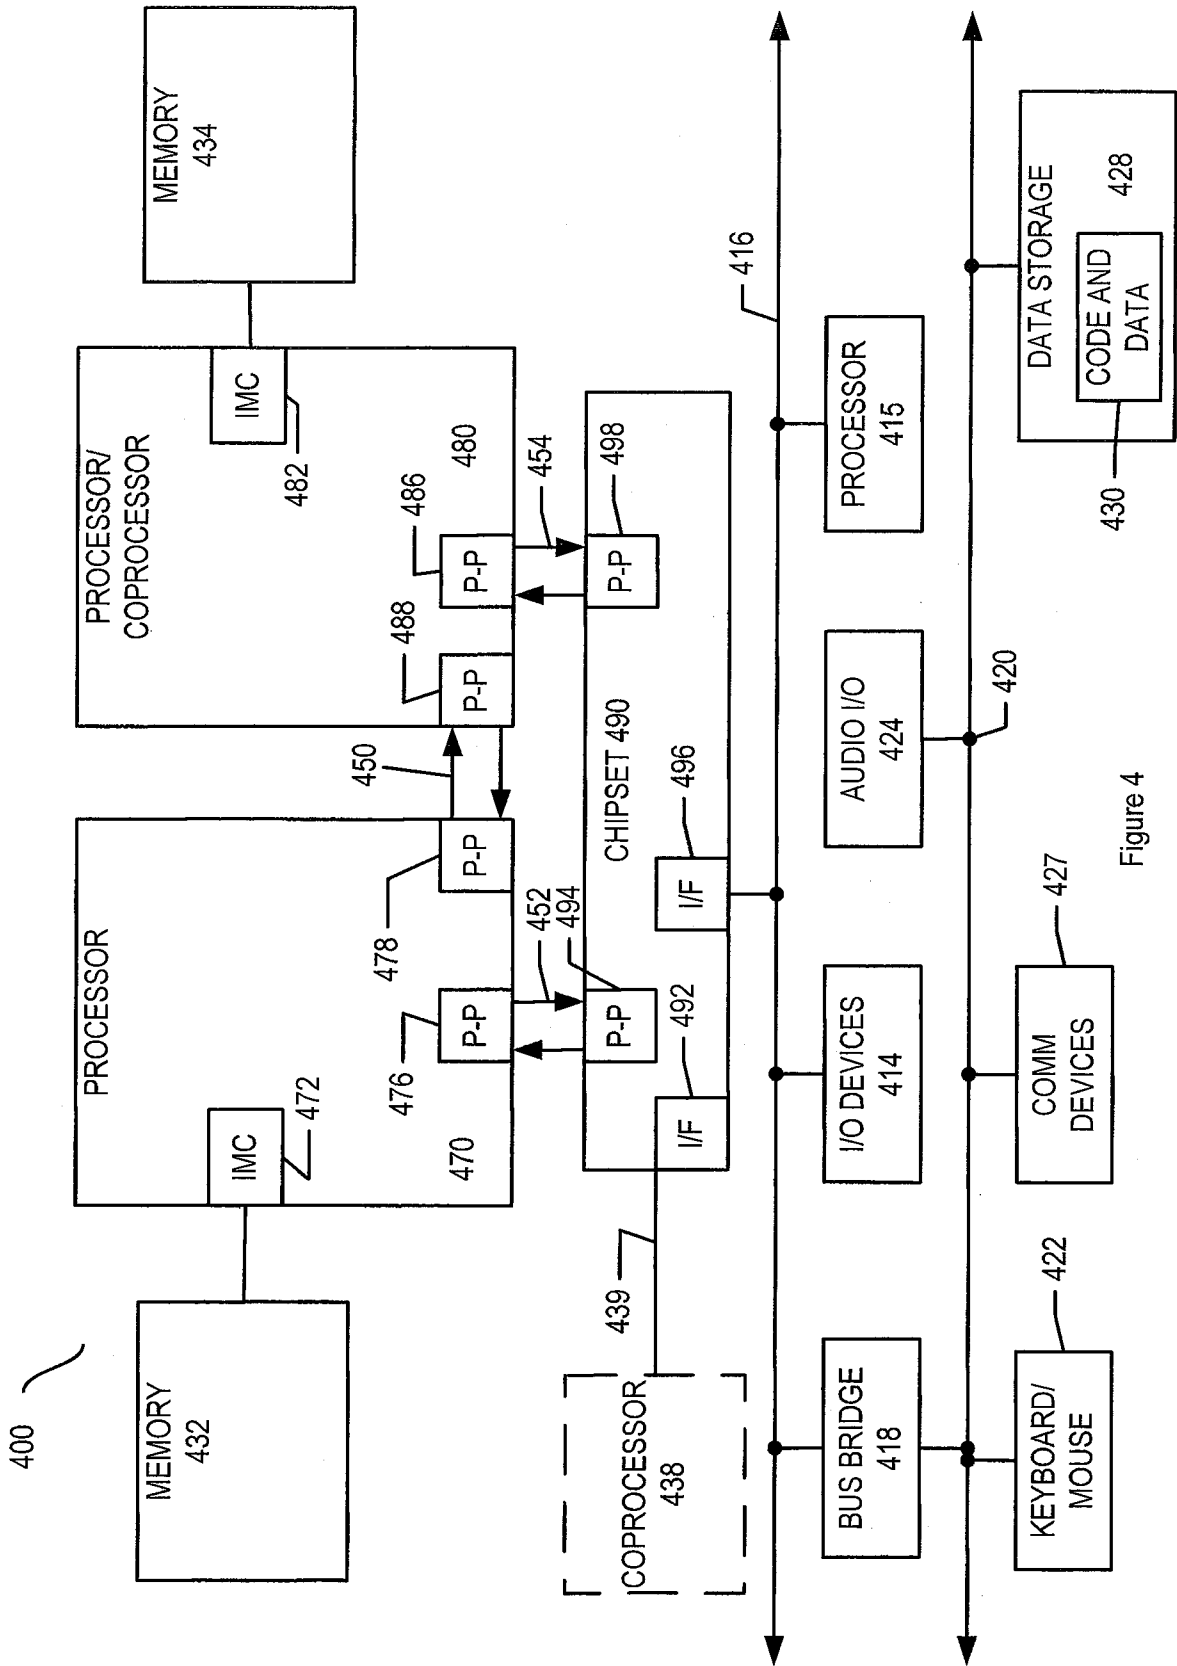


Figure 4

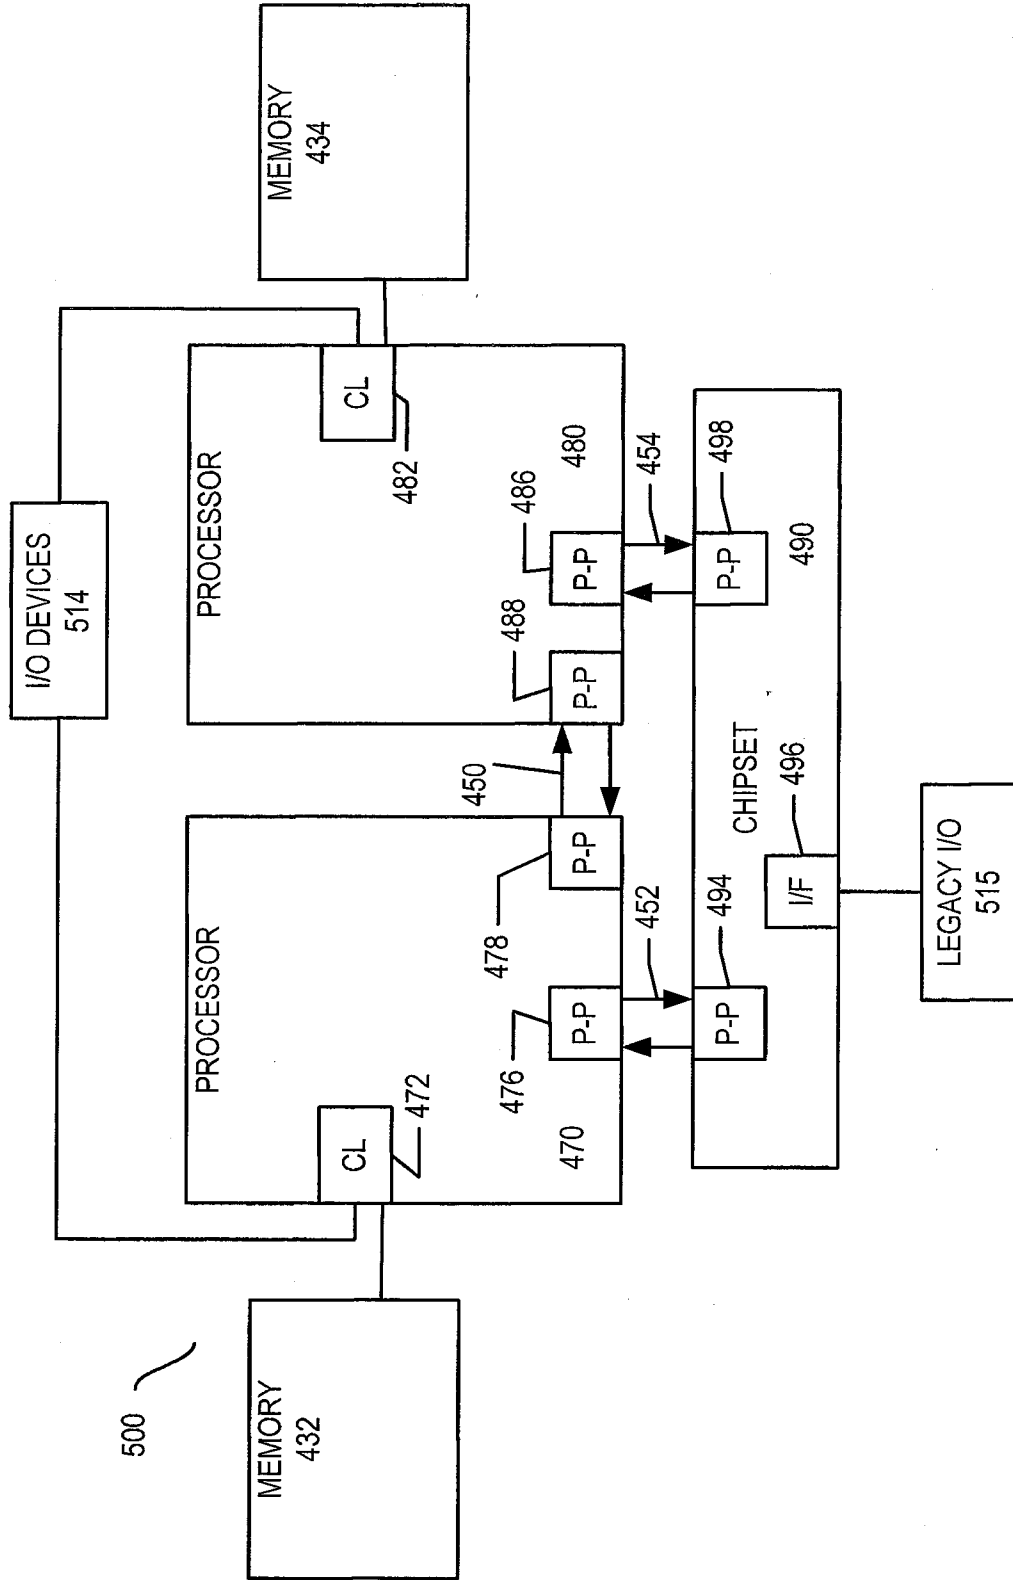


Figure 5

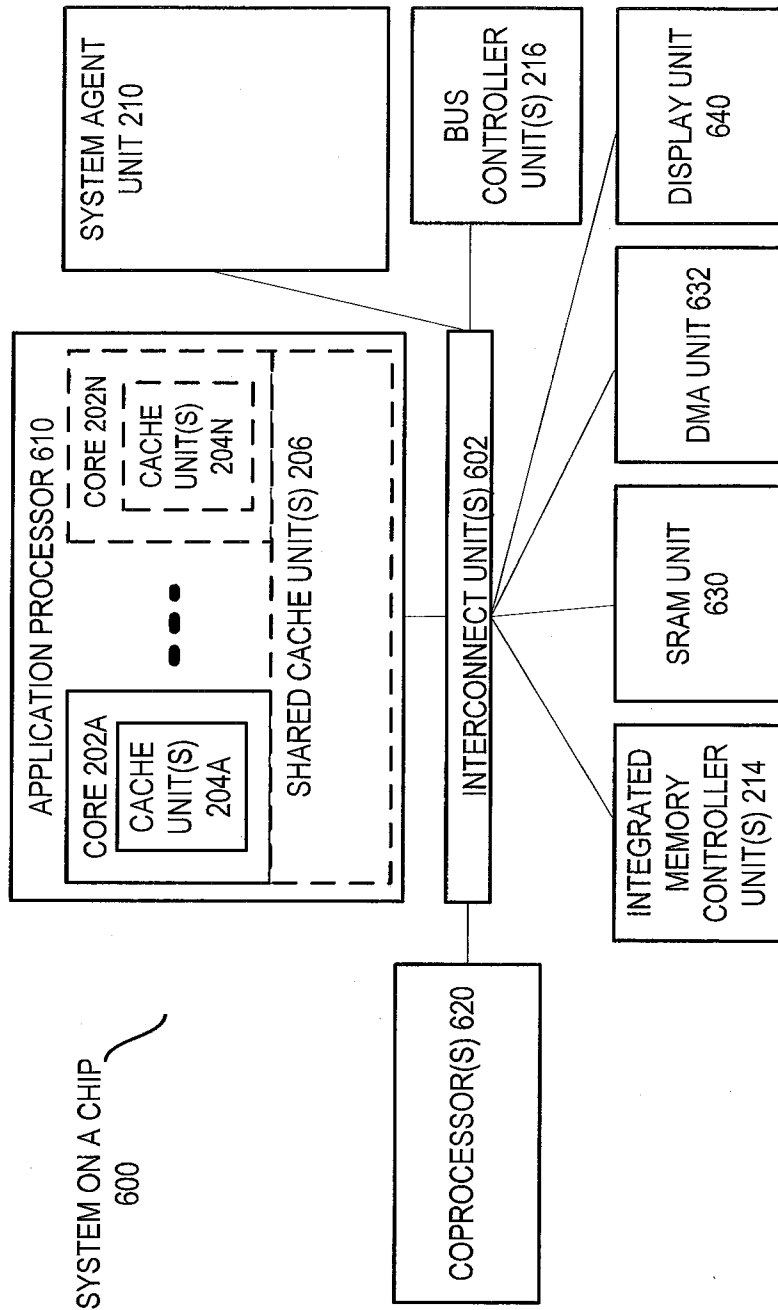


Figure 6

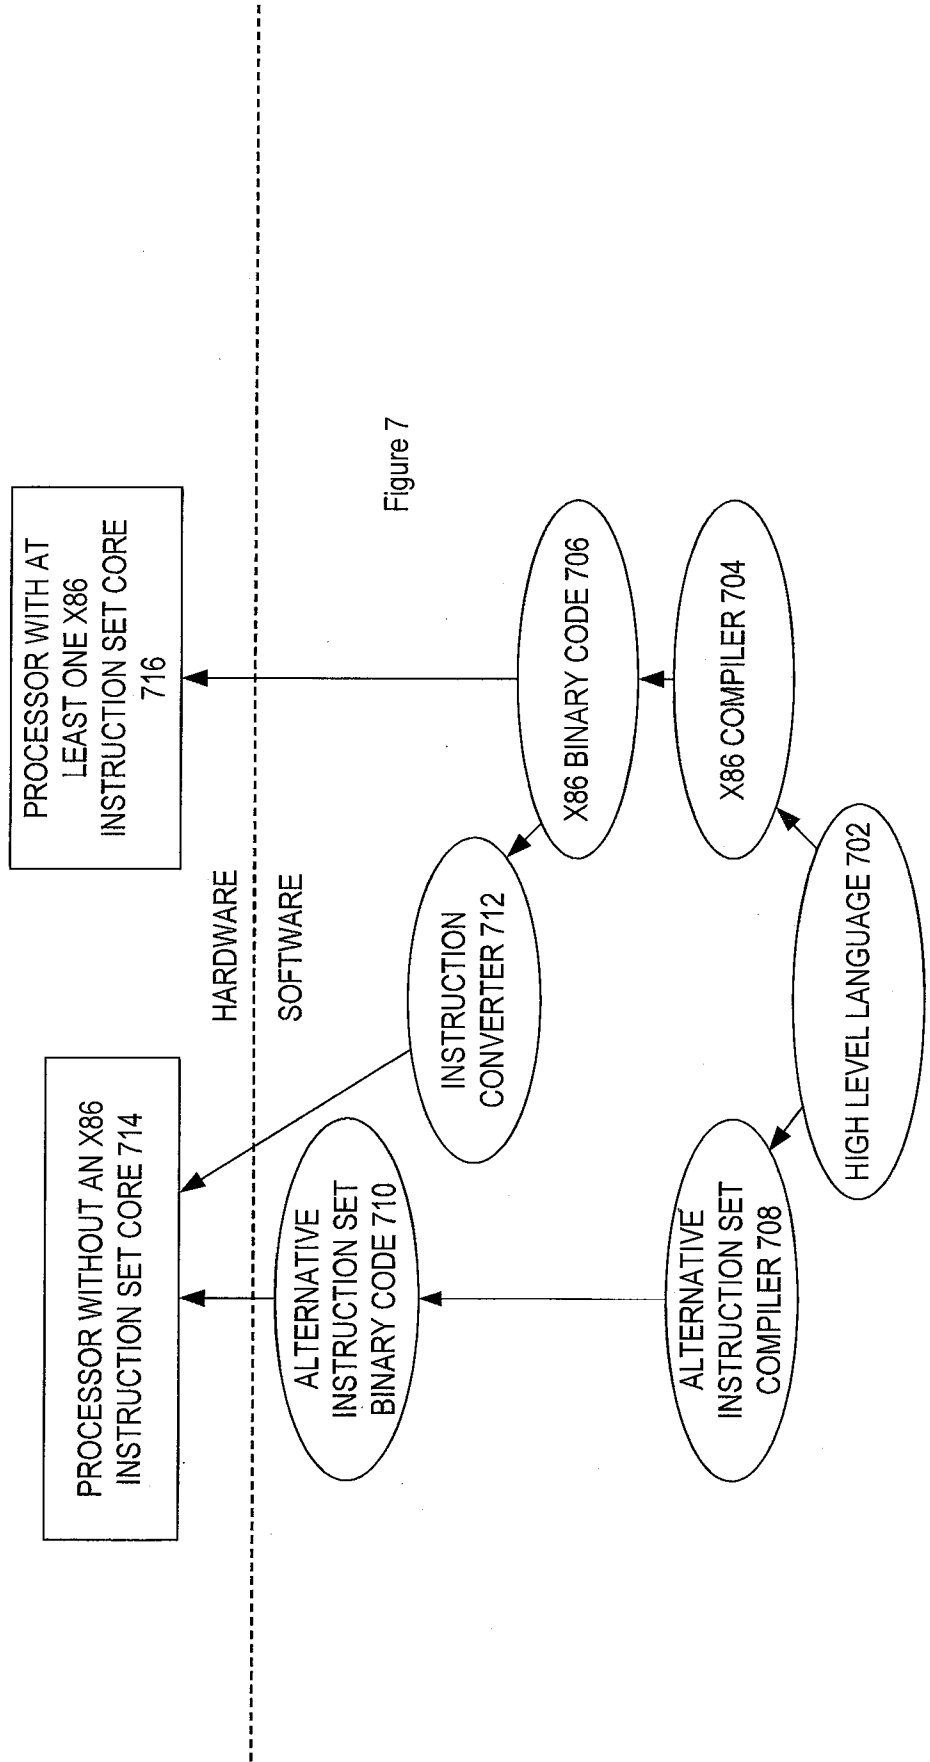
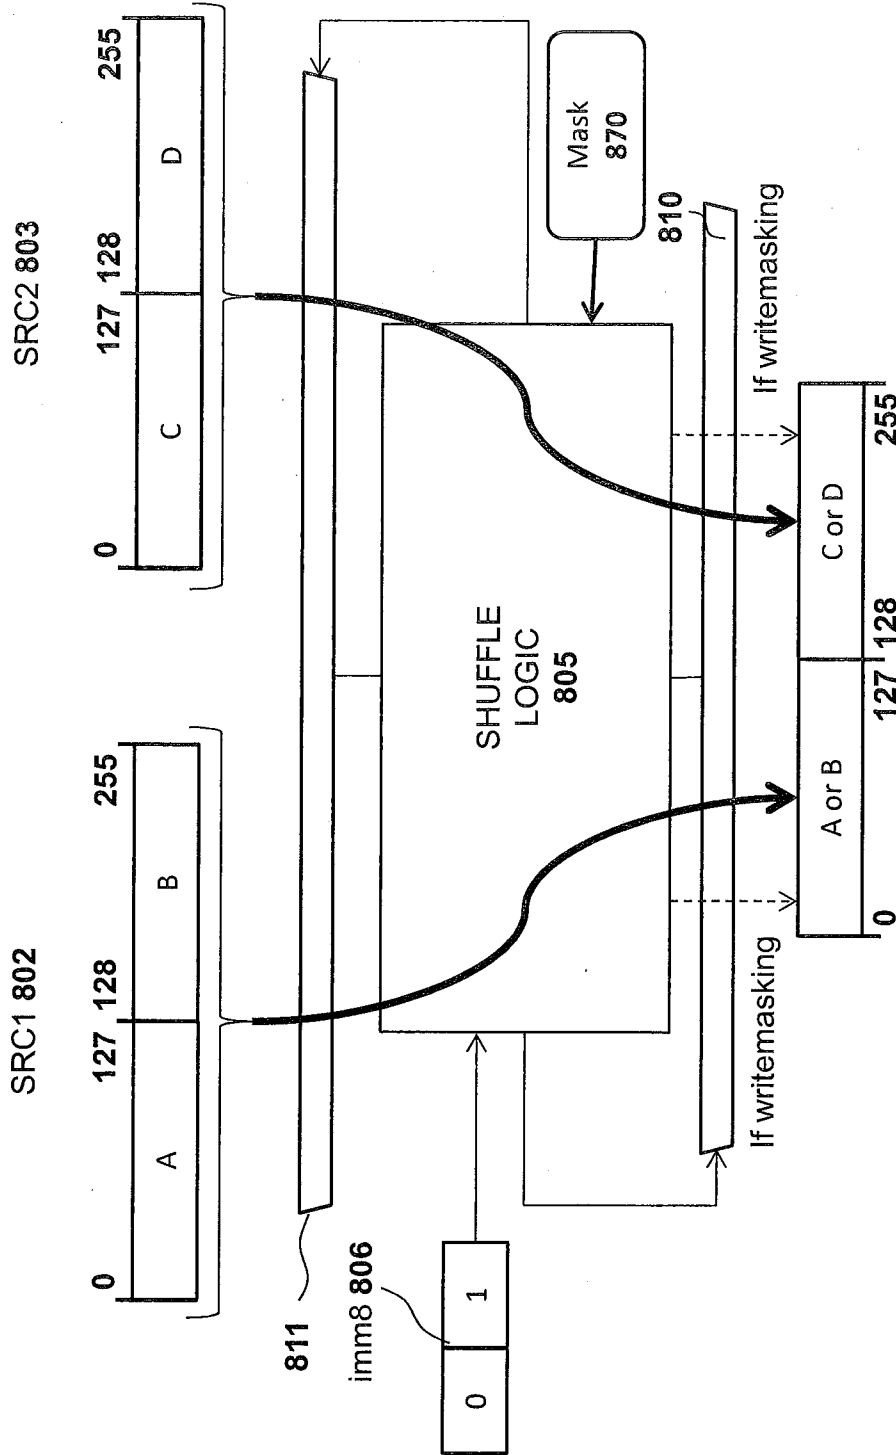


Figure 7



DST 804

**FIG. 8**

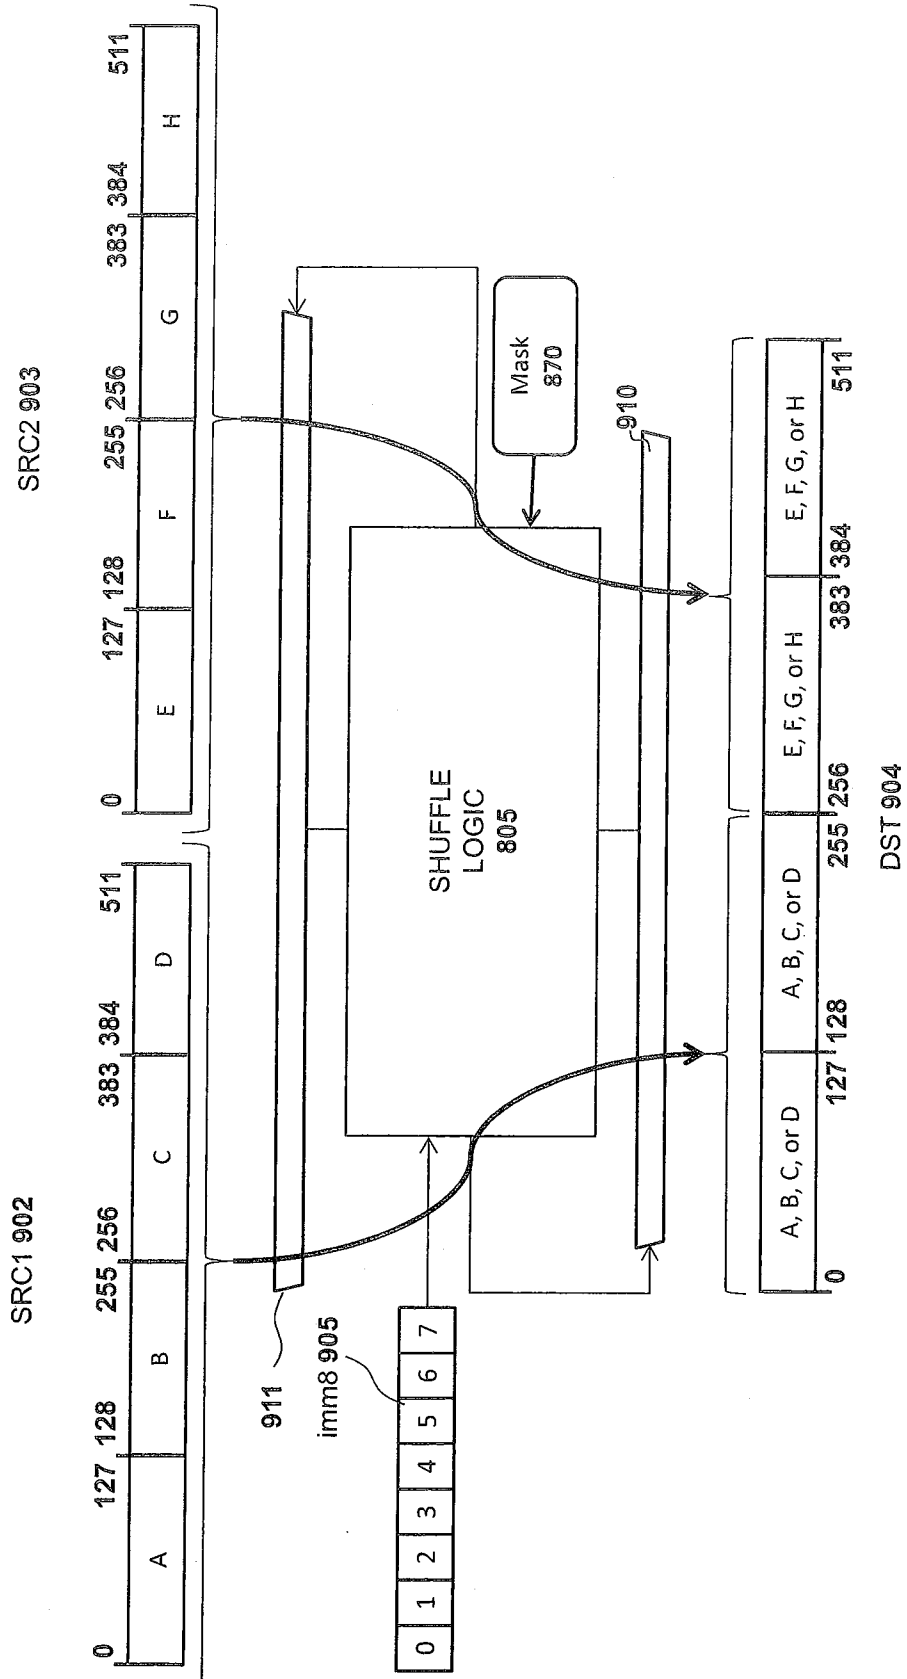


FIG. 9

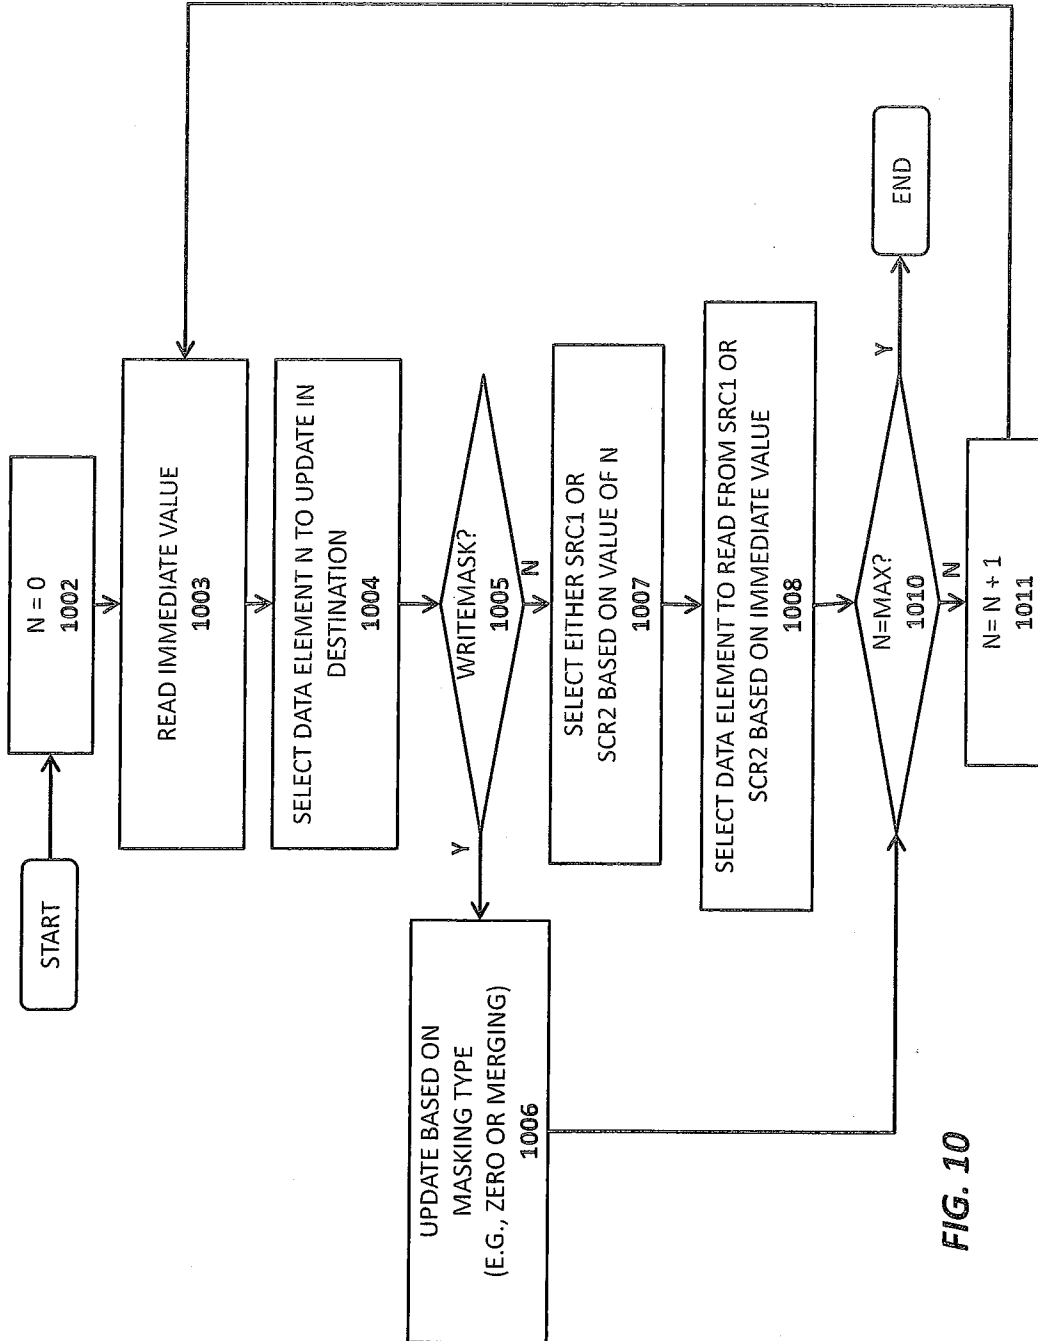
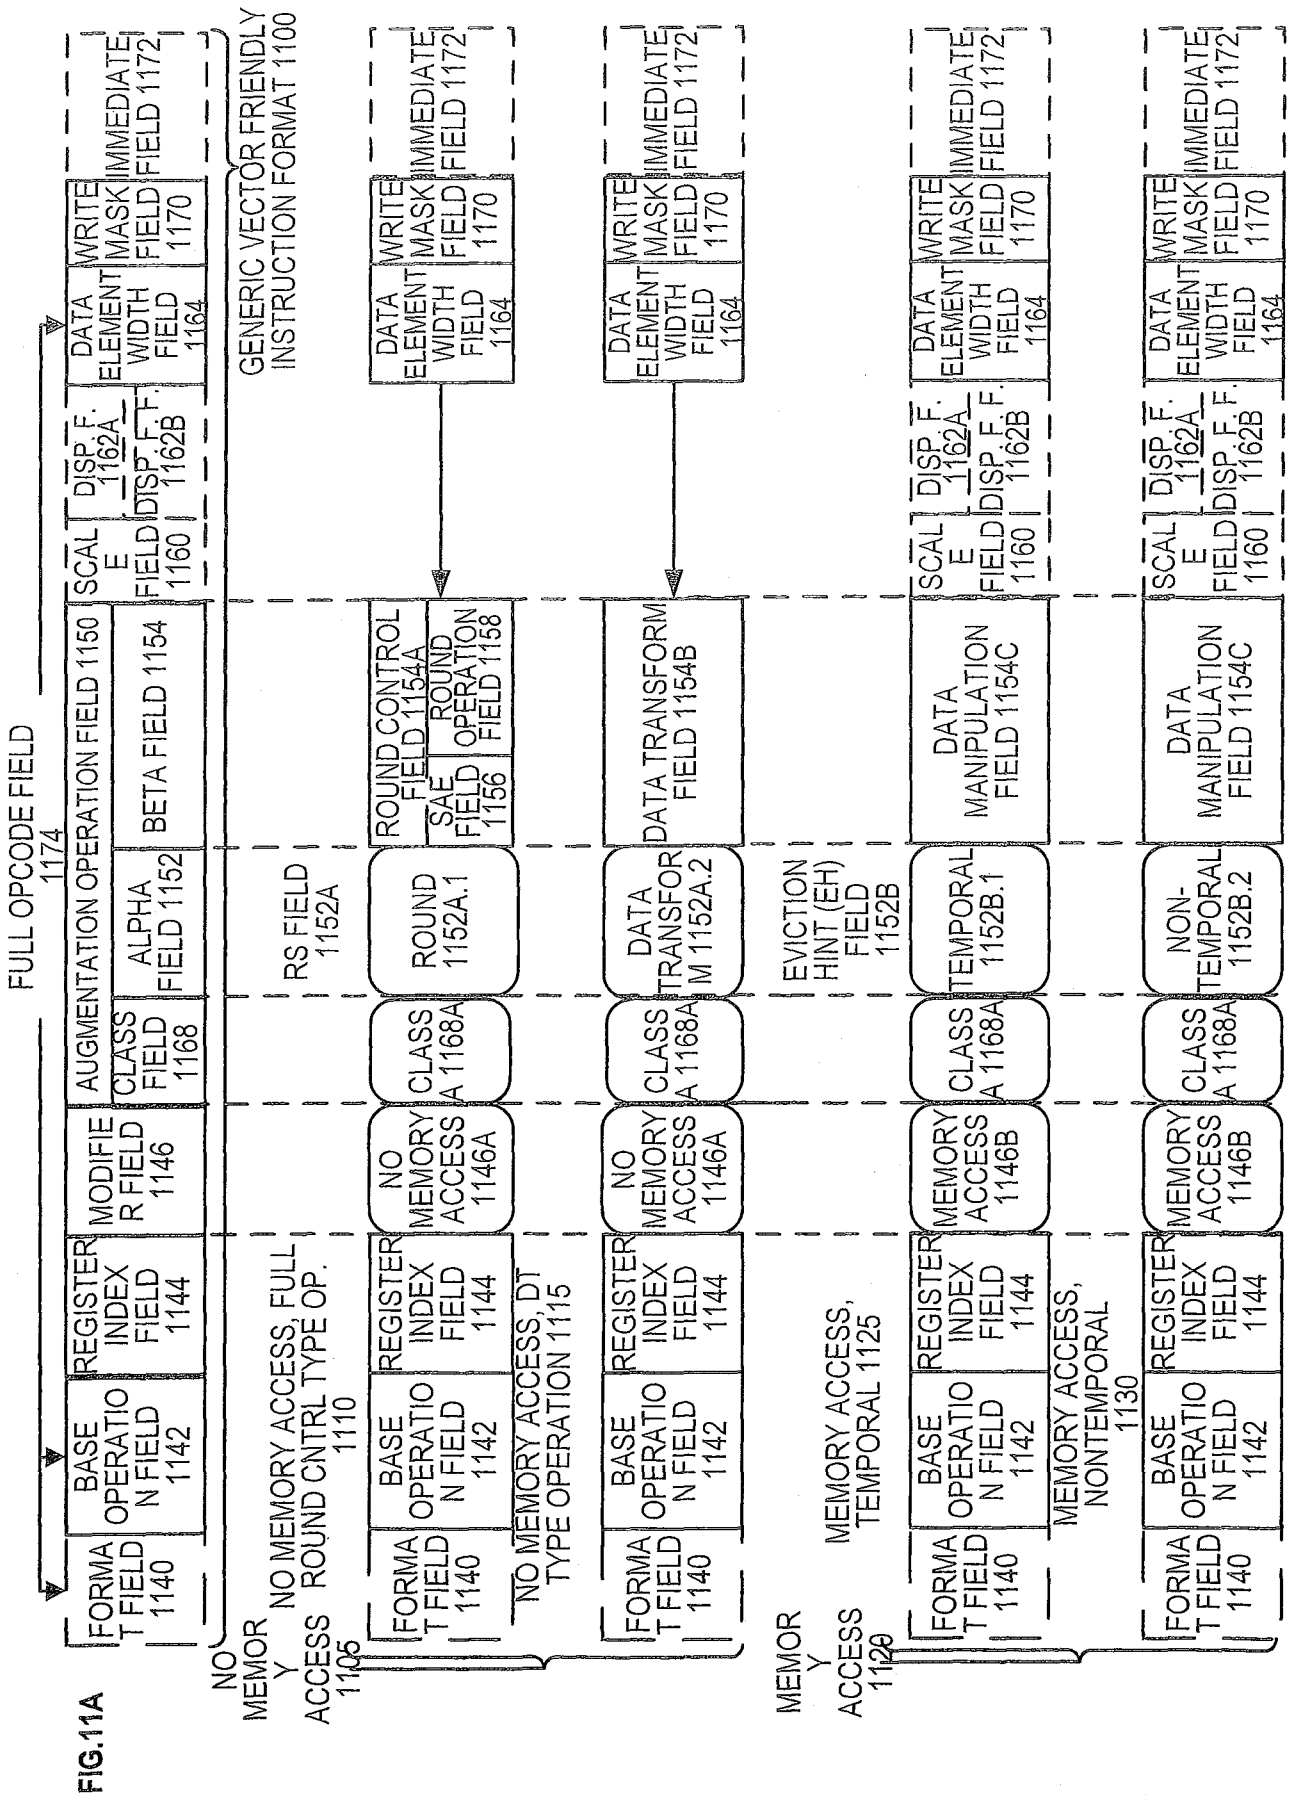
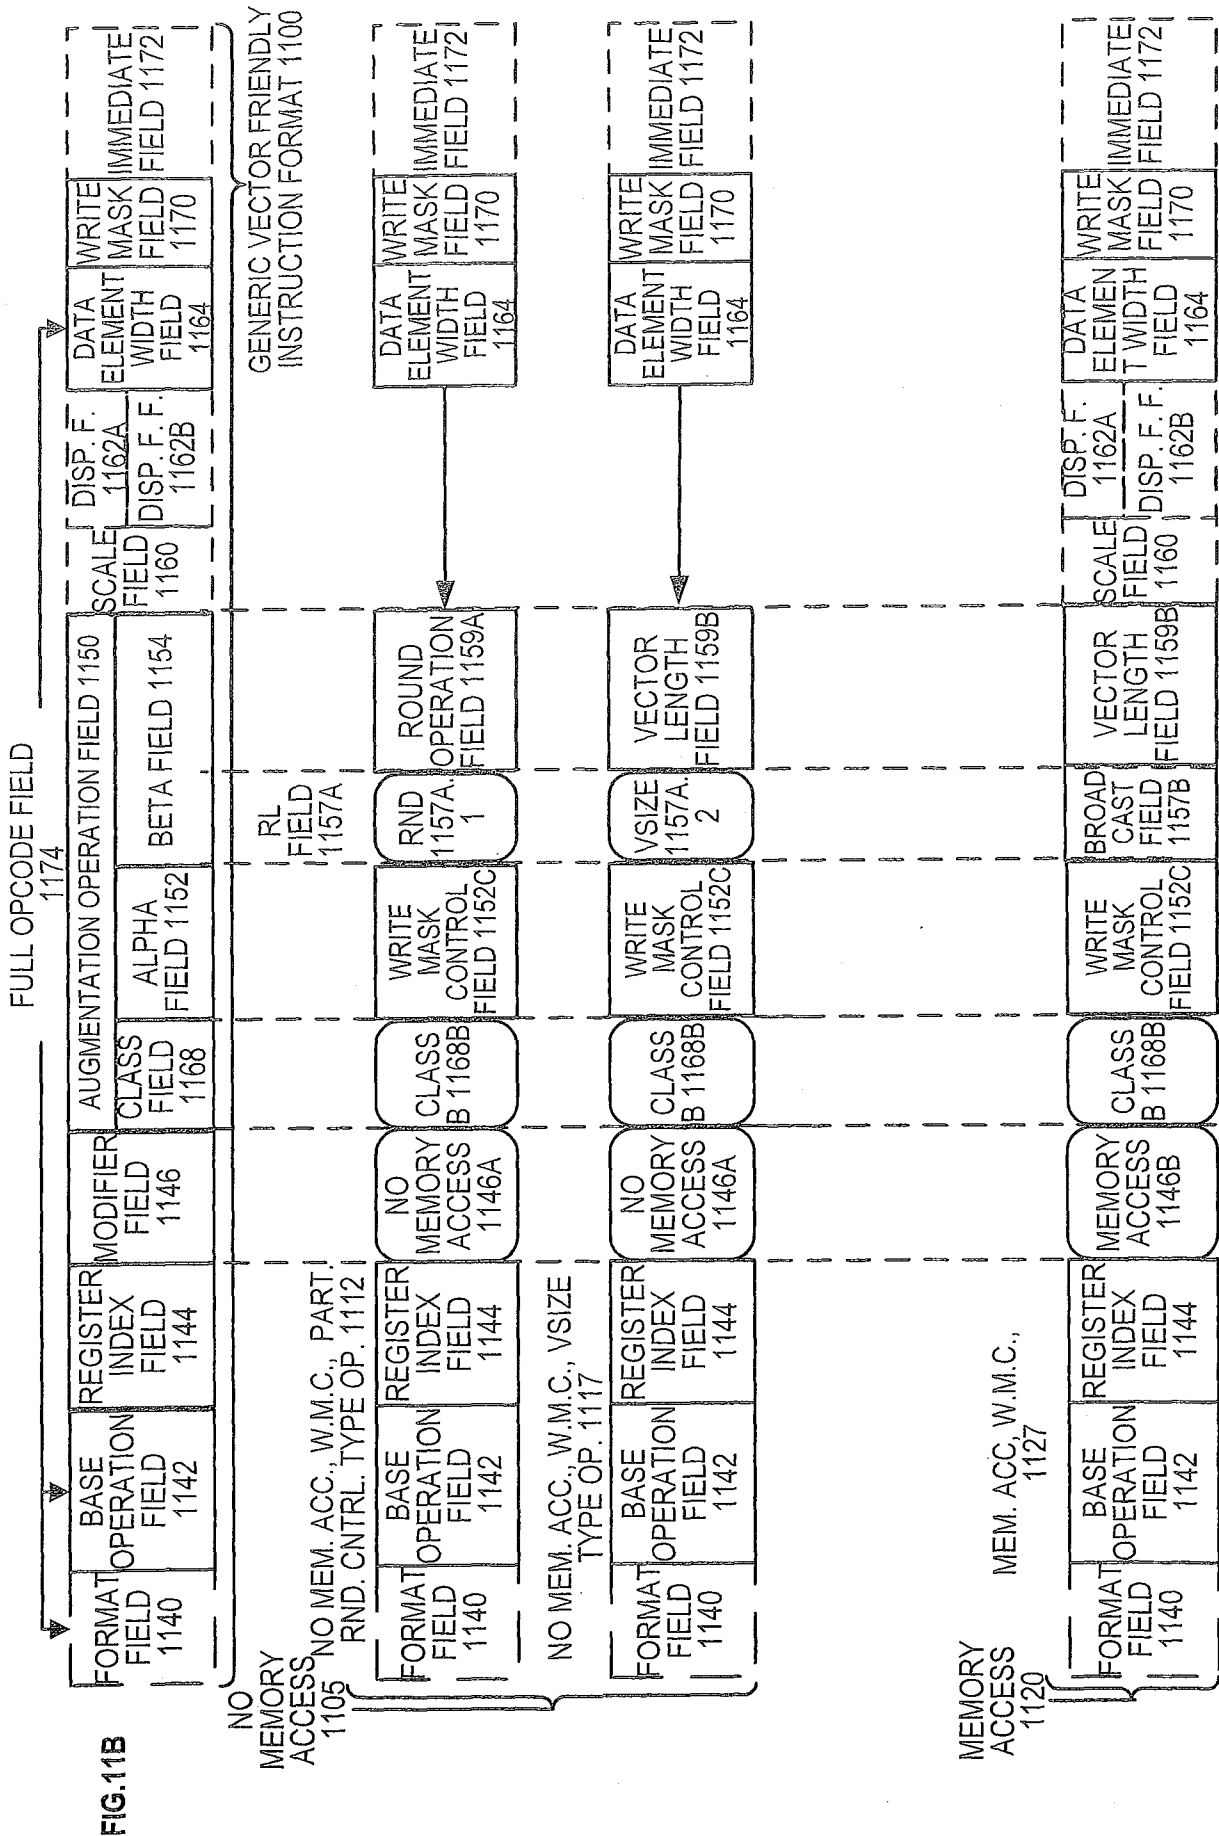


FIG. 10







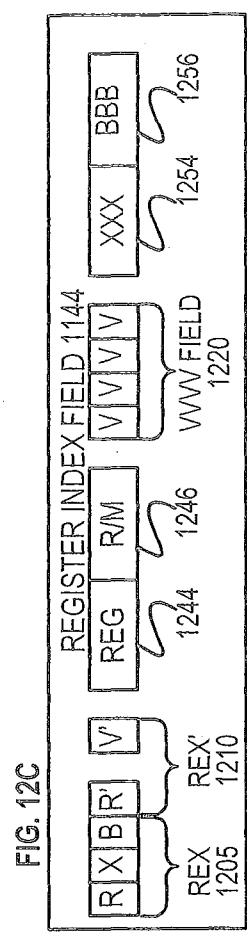
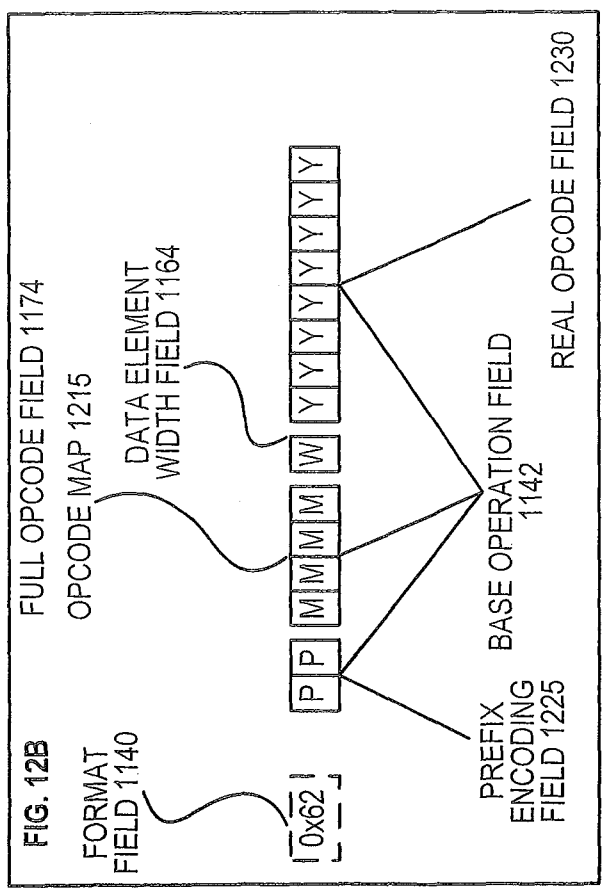
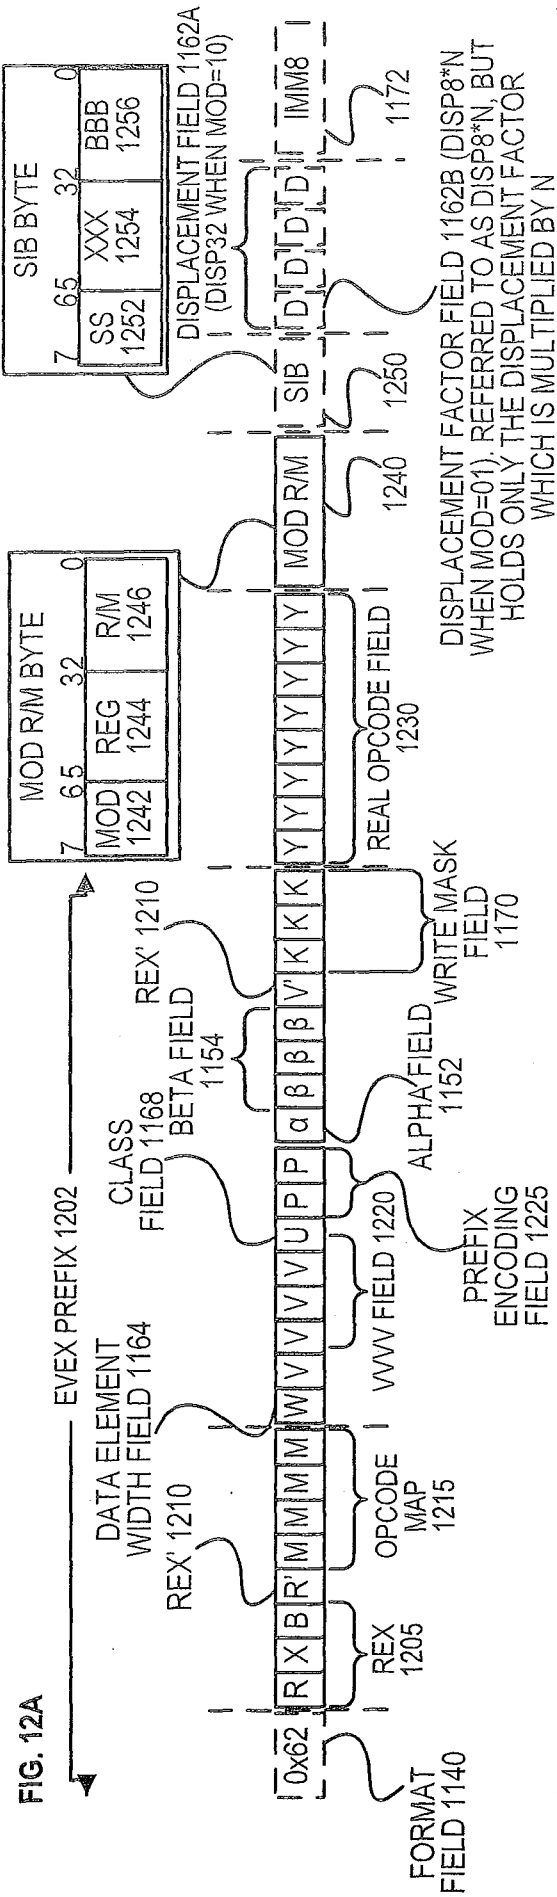
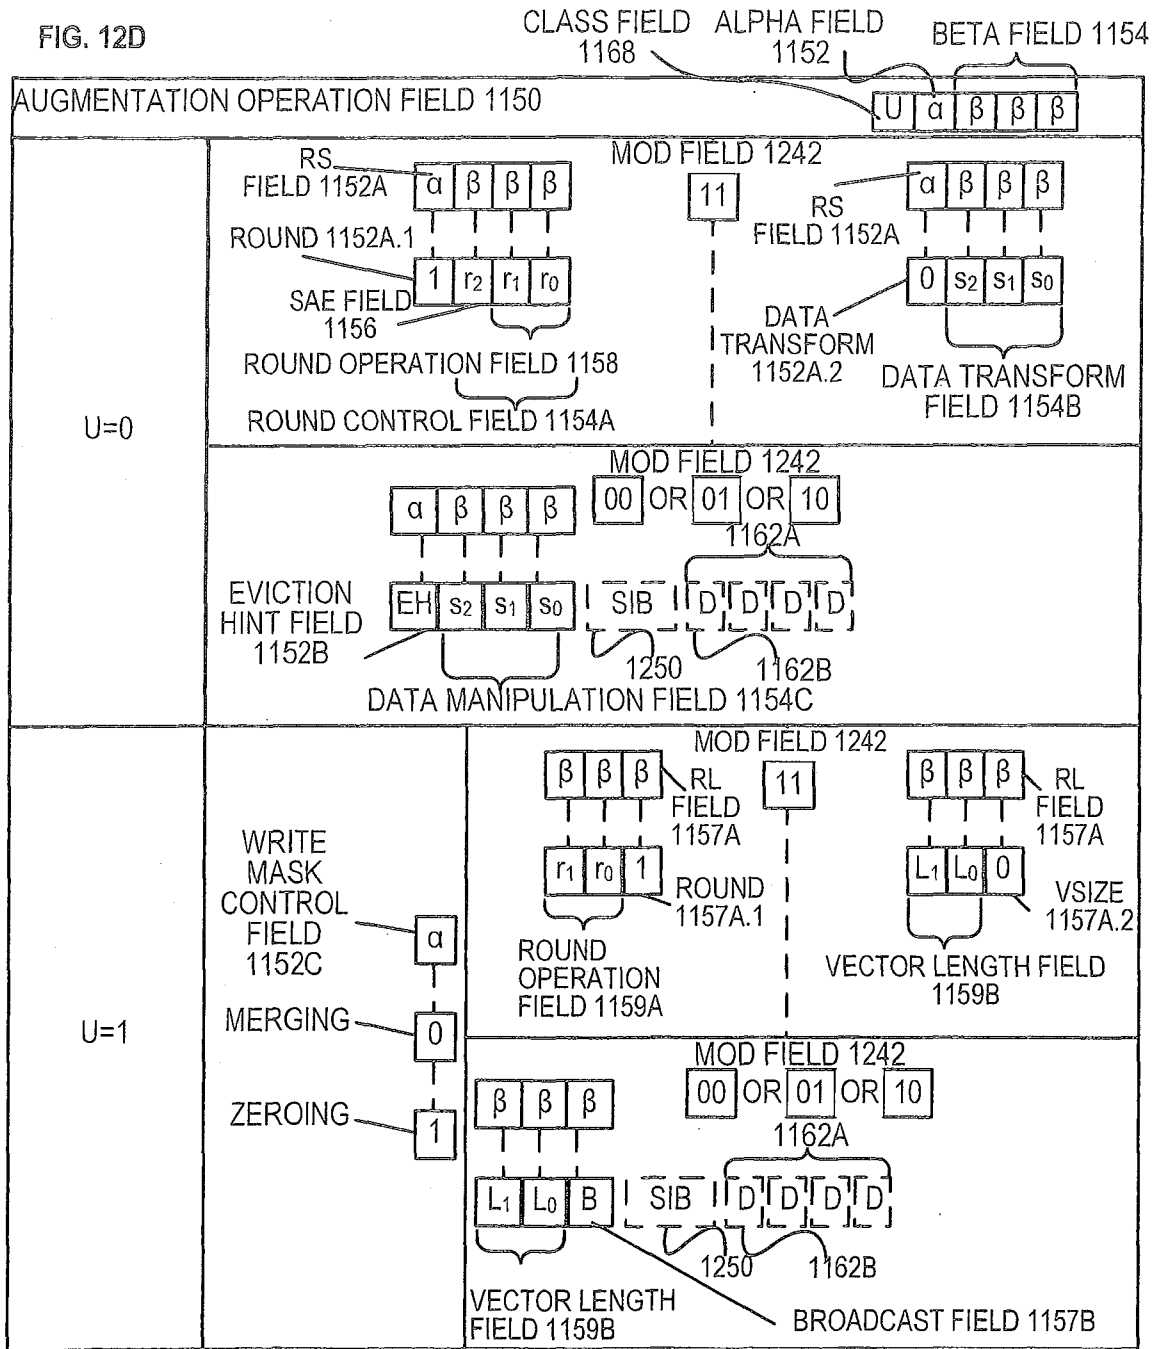


FIG. 12D



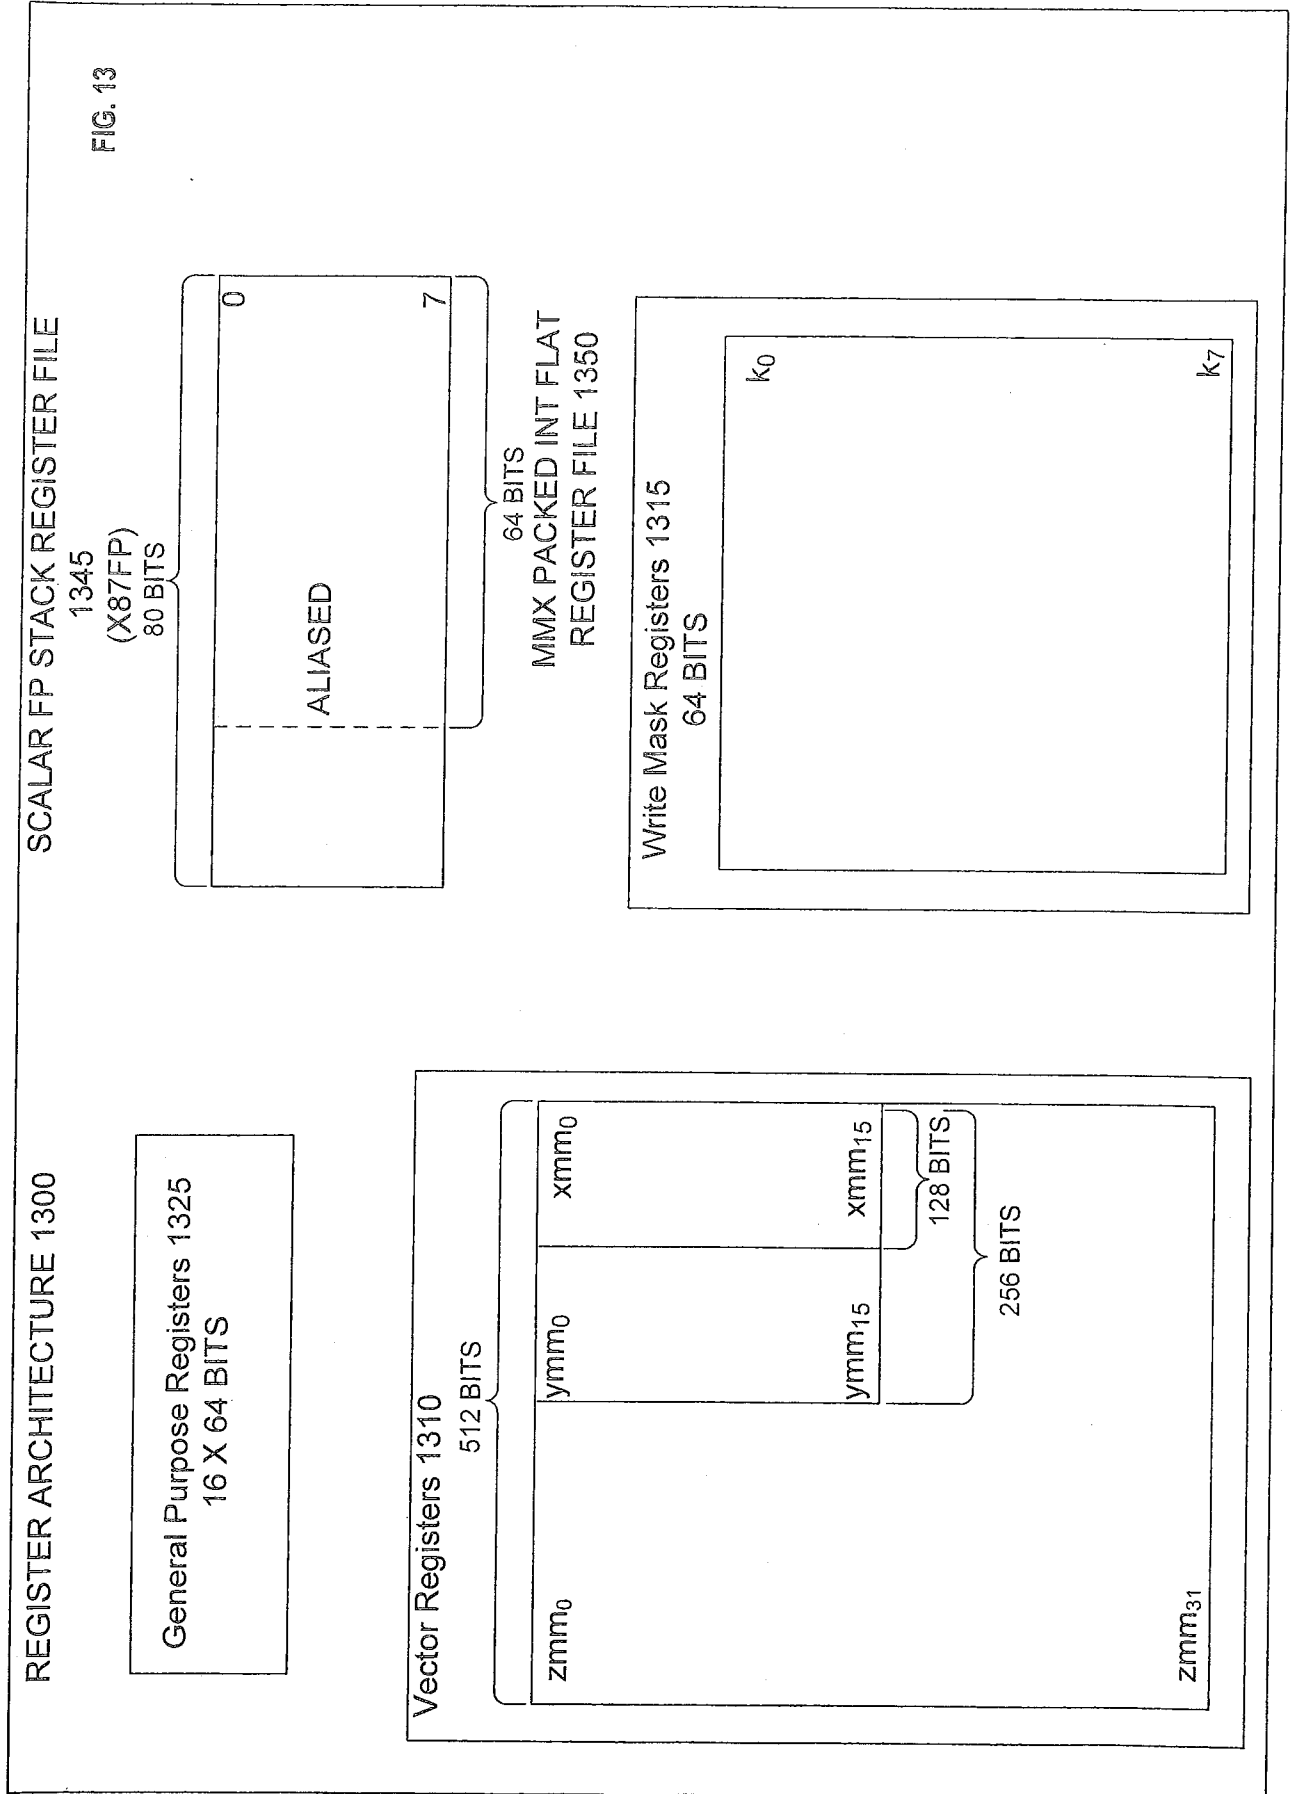


FIG. 14A

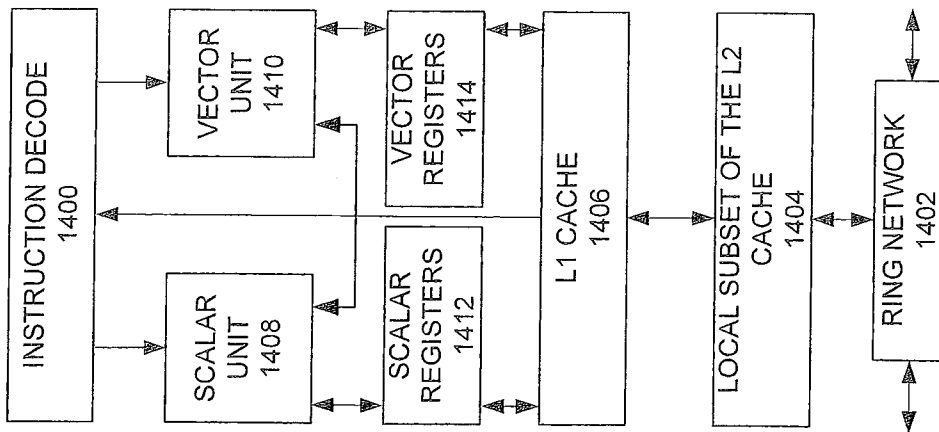
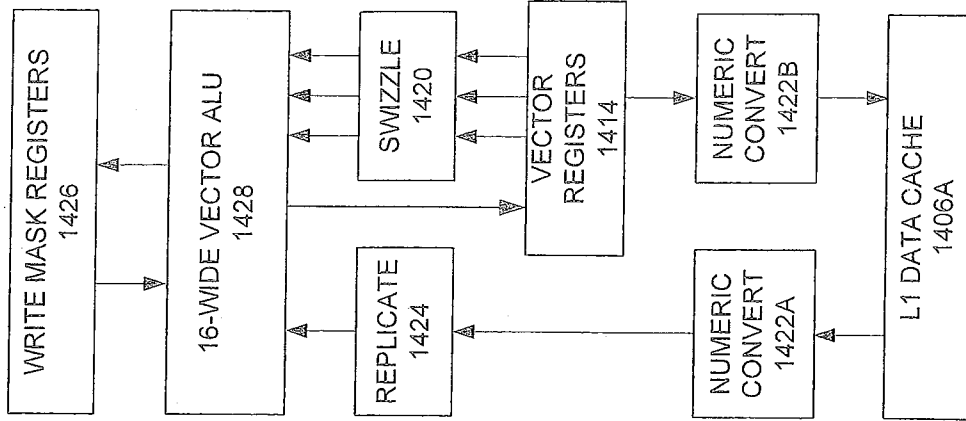




FIG. 14B



## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US201 1/067087**

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
<i>G06F 9/30(2006.01)i, G06F 9/305(2006.01)i</i>		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) G06F 9/30; G06F 5/01; G06F 9/315; G06F 9/00		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Korean utility models and applications for utility models Japanese utility models and applications for utility models		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) eKOMPASS(KIPO internal) & Keywords: shuffle, mask bit, register		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004-0054877 A1 (WILLIAM, W. MACY et al.) 18 March 2004 See the abstract, paras .[0101]-[0103] , and claim 1 .	1-22
A	US 7739319 B2 (MACY, JR. WILLIAM W. et al.) 15 June 2010 See the abstract, col.31 lines 44-67, col.32 lines 1-9, fig. 16, and claims 1 , 28.	1-22
A	US 7155601 B2 (SIRINIVAS CHENNUPATY) 26 December 2006 See the abstract, col.4 lines 53-67, col.5 lines 1-9, figs.3A, 3B, 3C, and claim 1 .	1-22
A	US 2011-0029759 A1 (MACY, JR. WILLIAM W. et al.) 03 February 2011 See the abstract, figs.9 , 11, and claim 1 .	1-22
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 28 AUGUST 2012 (28.08.2012)		Date of mailing of the international search report <b>29 AUGUST 2012 (29.08.2012)</b>
Name and mailing address of the ISA/KR  Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-70 1, Republic of Korea Facsimile No. 82-42-472-7140		Authorized officer Park Ji Eun Telephone No. 82-42-481-5696 

## INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

**PCT/US20 11/067087**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004-0054877 A1	18.03.2004	AU 1996-695 11 B2	23.03.2000
		AU 2003-30 17 18 A1	25.05.2004
		CA 2230108 A1	06.03.1997
		CA 2230108 C	12.12.2000
		CN 100338570 CO	19.09.2007
		CN 100461093 CO	11.02.2009
		CN 100492278 C	27.05.2009
		CN 101620525 A	06.01.2010
		CN 1107905 C	07.05.2003
		CN 1200821 A	02.12.1998
		CN 1200821 CO	02.12.1998
		CN 1506807 A	23.06.2004
		CN 1522401 A	18.08.2004
		CN 1522401 CO	09.08.2006
		CN 1801082 A	12.07.2006
		CN 1813241 A	02.08.2006
		CN 1813241 CO	02.08.2006
		EP 0847552 A1	30.01.2002
		EP 0847552 B1	30.10.2002
		EP 1639452 A2	29.03.2006
		EP 1639452 B1	09.09.2009
		JP 03-750820 B2	01.03.2006
		JP 04-064989 B2	19.03.2008
		JP 04-607105 B2	15.10.2010
		JP 04-623963 B2	12.11.2010
		JP 04-750157 B2	27.05.2011
		JP 11-51 1577 A	05.10.1999
		JP 2005-508043 A	24.03.2005
		JP 2006-107463 A	20.04.2006
		JP 2007-526536 A	13.09.2007
		JP 2009-009587 A	15.01.2009
		JP 2010-282649 A	16.12.2010
		JP 2011-138541 A	14.07.2011
		JP 4750157 B2	17.08.2011
		KR 10-0329339 B1	06.07.2002
		KR 10-0602532 B1	19.07.2006
		KR 10-0831472 B1	22.05.2008
		US 05721892A A	24.02.1998
		US 05859997A A	12.01.1999
		US 05983256A A	09.11.1999
		US 060353 16A A	07.03.2000
		US 2002-0059355 A1	16.05.2002
		US 2003-0050941 A1	13.03.2003
		US 2003-0084082 A1	01.05.2003
		US 2003-0123748 A1	03.07.2003
		us 2003-013 1030 A1	10.07.2003
		us 2004-0054878 A1	18.03.2004
		us 2004-0054879 A1	18.03.2004
		us 2004-0059889 A1	25.03.2004



## INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

**PCT/US20 11/067087**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2004-0073589 A1	15. 04. 2004
		US 2004-0078404 A1	22. 04. 2004
		US 2004-0098556 A1	20. 05. 2004
		US 2004-01 17422 A1	17. 06. 2004
		US 2004-01336 17 A1	08. 07. 2004
		US 2004-0139 138 A1	15. 07. 2004
		US 2004-02 106 16 A1	21. 10. 2004
		US 2004-0220992 A1	04. 11. 2004
		US 2005-01083 12 A1	19. 05. 2005
		US 2009-0265409 A1	22. 10. 2009
		US 2009-0265523 A1	22. 10. 2009
		US 2010-00 11042 A1	14. 01. 2010
		US 201 1-0029759 A1	03. 02. 2011
		US 201 1-0035426 A1	10. 02. 2011
		US 6385634 B1	07. 05. 2002
		US 6418529 B1	09. 07. 2002
		US 6961845 B2	01. 11. 2005
		US 7085795 B2	01. 08. 2006
		US 7272622 B2	18. 09. 2007
		US 7340495 B2	04. 03. 2008
		US 7392275 B2	24. 06. 2008
		US 7395298 B2	01. 07. 2008
		US 7395302 B2	01. 07. 2008
		US 7424505 B2	09. 09. 2008
		US 7430578 B2	30. 09. 2008
		US 7509367 B2	24. 03. 2009
		US 7624138 B2	24. 11. 2009
		US 7631025 B2	08. 12. 2009
		US 7685212 B2	23. 03. 2010
		US 7725521 B2	25. 05. 2010
		US 7739319 B2	15. 06. 2010
		US 7818356 B2	19. 10. 2010
		US 8185571 B2	22. 05. 2012
		US 8214626 B2	03. 07. 2012
		US 8225075 B2	17. 07. 2012
		Wo 03-03860 1 A1	08. 05. 2003
		Wo 2004-040439 A2	13. 05. 2004
		wo 2004-040439 A3	13. 05. 2004
		wo 2005-006 183 A2	20. 01. 2005
		wo 2005-006 183 A3	20. 01. 2005
		wo 97-08610 A1	06. 03. 1997
US 7739319 B2	15. 06. 2010	AU 1996-695 11 B2	23. 03. 2000
		AU 2003-30 17 18 A1	25. 05. 2004
		CA 2230108 A1	06. 03. 1997
		CA 2230108 C	12. 12. 2000
		CN 100338570 CO	19. 09. 2007
		CN 100461093 CO	11. 02. 2009
		CN 100492278 C	27. 05. 2009
		CN 101620525 A	06. 01. 2010

## INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

**PCT/US20 11/067087**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		CN 1107905 C	07. 05 ,2003
		CN 1200821 A	02. 12 , 1998
		CN 1200821 CO	02. 12 , 1998
		CN 1506807 A	23. 06 ,2004
		CN 1522401 A	18. 08 ,2004
		CN 1522401 CO	09. 08 ,2006
		CN 1801082 A	12. 07 .2006
		CN 1813241 A	02. 08 ,2006
		CN 1813241 CO	02. 08 ,2006
		EP 0847552 A1	30. 01 ,2002
		EP 0847552 B1	30. 10 ,2002
		EP 1639452 A2	29. 03 .2006
		EP 1639452 B1	09. 09 ,2009
		JP 03-750820 B2	01. 03 ,2006
		JP 04-064989 B2	19. 03 ,2008
		JP 04-607105 B2	15. 10 ,2010
		JP 04-623963 B2	12. 11 .2010
		JP 04-750157 B2	27. 05 ,2011
		JP 11-511577 A	05. 10 , 1999
		JP 2005-508043 A	24. 03 .2005
		JP 2006-107463 A	20. 04 ,2006
		JP 2007-526536 A	13. 09 ,2007
		JP 2009-009587 A	15. 01 ,2009
		JP 2010-282649 A	16. 12 ,2010
		JP 2011-138541 A	14. 07 .2011
		JP 4750157 B2	17. 08 ,2011
		KR 10-0329339 B1	06. 07 ,2002
		KR 10-0602532 B1	19. 07 ,2006
		KR 10-0831472 B1	22. 05 ,2008
		US 05721892A A	24. 02 . 1998
		US 05859997A A	12. 01 , 1999
		US 05983256A A	09. 11 , 1999
		US 06035316A A	07. 03 ,2000
		US 2002-0059355 A1	16. 05 ,2002
		US 2003-0050941 A1	13. 03 ,2003
		US 2003-0084082 A1	01. 05 ,2003
		US 2003-0123748 A1	03. 07 ,2003
		us 2003-0131030 A1	10. 07 .2003
		us 2004-0054877 A1	18. 03 ,2004
		us 2004-0054878 A1	18. 03 ,2004
		us 2004-0054879 A1	18. 03 ,2004
		us 2004-0059889 A1	25. 03 ,2004
		us 2004-0073589 A1	15. 04 .2004
		us 2004-0078404 A1	22. 04 ,2004
		us 2004-0098556 A1	20. 05 ,2004
		us 2004-0117422 A1	17. 06 .2004
		us 2004-0133617 A1	08. 07 ,2004
		us 2004-0139138 A1	15. 07 ,2004
		us 2004-0210616 A1	21. 10 ,2004
		us 2004-0220992 A1	04. 11 ,2004

## INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

**PCT /US2011/067087**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2005-01083 12 <b>A1</b>	19.05.2005
		US 2009-0265409 <b>A1</b>	22.10.2009
		us 2009-0265523 <b>A1</b>	22.10.2009
		us 2010-00 11042 <b>A1</b>	14.01.2010
		us 201 1-0029759 <b>A1</b>	03.02.2011
		us 201 1-0035426 <b>A1</b>	10.02.2011
		us 6385634 B1	07.05.2002
		us 6418529 B1	09.07.2002
		us 6961845 B2	01.11.2005
		us 7085795 B2	01.08.2006
		us 7272622 B2	18.09.2007
		us 7340495 B2	04.03.2008
		us 7392275 B2	24.06.2008
		us 7395298 B2	01.07.2008
		us 7395302 B2	01.07.2008
		us 7424505 B2	09.09.2008
		us 7430578 B2	30.09.2008
		us 7509367 B2	24.03.2009
		us 7624138 B2	24.11.2009
		us 7631025 B2	08.12.2009
		us 7685212 B2	23.03.2010
		us 7725521 B2	25.05.2010
		us 7818356 B2	19.10.2010
		us 8185571 B2	22.05.2012
		us 8214626 B2	03.07.2012
		us 8225075 B2	17.07.2012
		Wo 03-03860 1 A <b>1</b>	08.05.2003
		Wo 2004-040439 A2	13.05.2004
		wo 2004-040439 A3	13.05.2004
		wo 2005-006 183 A2	20.01.2005
		wo 2005-006 183 A3	20.01.2005
		wo 97-08610 A 1	06.03.1997
US 7155601 B2	26.12.2006	us 2002-1 12 147 <b>A1</b>	15.08.2002
US 201 1-0029759 <b>A1</b>	03.02.2011	AU 1996-695 11 B2	23.03.2000
		AU 2003-30 17 18 <b>A1</b>	25.05.2004
		CA 2230108 A 1	06.03.1997
		CA 2230108 C	12.12.2000
		CN 100338570 CO	19.09.2007
		CN 100461093 CO	11.02.2009
		CN 100492278 C	27.05.2009
		CN 101620525 A	06.01.2010
		CN 1107905 C	07.05.2003
		CN 1200821 A	02.12.1998
		CN 1200821 CO	02.12.1998
		CN 1506807 A	23.06.2004
		CN 1522401 A	18.08.2004
		CN 1522401 CO	09.08.2006
		CN 1801082 A	12.07.2006

## INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

**PCT/US20 11/067087**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		CN 1813241 A	02. 08 ,2006
		CN 1813241 co	02. 08 ,2006
		EP 0847552 A1	30. 01 ,2002
		EP 0847552 B1	30. 10 ,2002
		EP 1639452 A2	29. 03 ,2006
		EP 1639452 B1	09. 09 ,2009
		JP 03-750820 B2	01. 03 .2006
		JP 04-064989 B2	19. 03 ,2008
		JP 04-607105 B2	15. 10 ,2010
		JP 04-623963 B2	12. 11 ,2010
		JP 04-750157 B2	27. 05 ,2011
		JP 11-511577 A	05. 10 .1999
		JP 2005-508043 A	24. 03 ,2005
		JP 2006-107463 A	20. 04 ,2006
		JP 2007-526536 A	13. 09 ,2007
		JP 2009-009587 A	15. 01 ,2009
		JP 2010-282649 A	16. 12 .2010
		JP 2011-138541 A	14. 07 ,2011
		JP 4750157 B2	17. 08 ,2011
		KR 10-0329339 B1	06. 07 .2002
		KR 10-0602532 B1	19. 07 ,2006
		KR 10-0831472 B1	22. 05 ,2008
		US 05721892A A	24. 02 ,1998
		US 05859997A A	12. 01 ,1999
		US 05983256A A	09. 11 .1999
		US 06035316A A	07. 03 ,2000
		US 2002-0059355 A1	16. 05 ,2002
		US 2003-0050941 A1	13. 03 ,2003
		US 2003-0084082 A1	01. 05 ,2003
		US 2003-0123748 A1	03. 07 .2003
		us 2003-0131030 A1	10. 07 ,2003
		us 2004-0054877 A1	18. 03 ,2004
		us 2004-0054878 A1	18. 03 ,2004
		us 2004-0054879 A1	18. 03 ,2004
		us 2004-0059889 A1	25. 03 ,2004
		us 2004-0073589 A1	15. 04 ,2004
		us 2004-0078404 A1	22. 04 ,2004
		us 2004-0098556 A1	20. 05 .2004
		us 2004-0117422 A1	17. 06 ,2004
		us 2004-0133617 A1	08. 07 ,2004
		us 2004-0139138 A1	15. 07 ,2004
		us 2004-0210616 A1	21. 10 ,2004
		us 2004-0220992 A1	04. 11 .2004
		us 2005-0108312 A1	19. 05 ,2005
		us 2009-0265409 A1	22. 10 ,2009
		us 2009-0265523 A1	22. 10 .2009
		us 2010-0011042 A1	14. 01 ,2010
		us 2011-0035426 A1	10. 02 ,2011
		us 6385634 B1	07. 05 ,2002
		us 6418529 B1	09. 07 ,2002

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

PCT/US20 11/067087

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 6961845 B2	01.11.2005
		us 7085795 B2	01.08.2006
		us 7272622 B2	18.09.2007
		us 7340495 B2	04.03.2008
		us 7392275 B2	24.06.2008
		us 7395298 B2	01.07.2008
		us 7395302 B2	01.07.2008
		us 7424505 B2	09.09.2008
		us 7430578 B2	30.09.2008
		us 7509367 B2	24.03.2009
		us 7624138 B2	24.11.2009
		us 7631025 B2	08.12.2009
		us 7685212 B2	23.03.2010
		us 7725521 B2	25.05.2010
		us 7739319 B2	15.06.2010
		us 7818356 B2	19.10.2010
		us 8185571 B2	22.05.2012
		us 8214626 B2	03.07.2012
		us 8225075 B2	17.07.2012
		Wo 03-03860 1 A1	08.05.2003
		Wo 2004-040439 A2	13.05.2004
		wo 2004-040439 A3	13.05.2004
		wo 2005-006 183 A2	20.01.2005
		wo 2005-006 183 A3	20.01.2005
		wo 97-08610 A1	06.03.1997