



- (51) International Patent Classification:
G06F 9/30 (2006.01) G06F 9/305 (2006.01)
- (21) International Application Number:
PCT/US201 1/068230
- (22) International Filing Date:
30 December 201 1 (30. 12.201 1)
- (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, Santa Clara, California 95052 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): SAIR, Suleyman [TR/US]; 5000 West Chandler Boulevard, MS: CH7-401, Chandler, Arizona 85226 (US). OULD-AHMED-VALL, Elmoustapha [MR/US]; 5000 West Chandler Boulevard, MS: CH7-401, Chandler, Arizona 85226 (US). YOUNT, Charles R. [US/US]; 1373 East Desert Broom Way, Phoenix, Arizona 85048 (US). DOSHI, Kshitij A. [US/US]; 638 North Elm Street, Chandler, Arizona 85226 (US). TOLL, Bret L. [US/US]; 2868 NE Lorie Drive, Hillsboro, Oregon 97124 (US).

- (74) Agents: CRAWFORD, Jonathan P. 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).

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

(54) Title: VECTOR FREQUENCY COMPRESS INSTRUCTION




Figure 1

(57) Abstract: A processor core that includes a hardware decode unit to decode a vector frequency compress instruction that includes a source operand and a destination operand. The source operand specifying a source vector register that includes a plurality of source data elements including one or more runs of identical data elements that are each to be compressed in a destination vector register as a value and run length pair. The destination operand identifies the destination vector register. The processor core also includes an execution engine unit to execute the decoded vector frequency compress instruction which causes, for each source data element, a value to be copied into the destination vector register to indicate that source data element's value. One or more runs of the source data elements equal are encoded in the destination vector register as the predetermined compression value followed by a run length for that run.

WO 2013/101227 A1

VECTOR FREQUENCY COMPRESS INSTRUCTION

FIELD

[0001] The field of invention relates generally to computer processor architecture, and, more specifically, to a vector frequency compress instruction.

BACKGROUND

[0002] 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). It should be noted that the term instruction generally refers herein to a macro-instruction - that is instructions that are provided to the processor for execution - as opposed to micro-instructions or micro-ops - that result from a processor's decoder decoding macro-instructions). The instruction set architecture is distinguished from the microarchitecture, which is the internal design of the processor implementing the ISA.

Processors with different microarchitectures can share a common instruction set.

[0003] 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 and the operand(s) on which that operation is to be performed. A given instruction is expressed using a given instruction format and specifies the operation and the operands. An instruction stream is a specific sequence of instructions, where each instruction in the sequence is an occurrence of an instruction in an instruction format.

[0004] Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis)/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 the same 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 64-bit register may be specified as a source operand to be operated on as four separate 16-bit data elements, each of which represents a separate 16-bit value. As another 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 the 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).

[0005] 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 have only one or has more than two source vector operands; that operate in a horizontal fashion; that generate a result vector operand that is of a different size, that have a different size of data elements, and/or that have 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).

[0006] Certain instruction set architectures allow multiple vector and scalar operations to complete in parallel and update the instruction set architecture register set. These instruction set architectures can be leveraged to implement compression/decompression instructions and algorithms such as instructions based on run-length encoding (RLE).

5 [0007] RLE is a form of lossless data compressing where sequences of data in a stream of data are compressed when those sequences contain one or more sets of consecutive data values. Rather than storing each data element in the set of consecutive data values, a single element with the value is stored followed by an element with the count of consecutive elements. This form of compression is most useful on data that contains many such runs.

10 [0008] For example, zero-based compression/decompression takes advantage of frequently occurring zero elements in data streams. In some data types, particularly data associated with graphics processing, it is common to have a significant portion of data contain the value zero and, in turn, many runs of zeros. RLE based on zero is often referred to as zero-based compression. Although other RLE schemes may be based on a value other than zero if
15 compression would benefit from another RLE value being selected.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the
20 drawings:

[0010] Figure 1 illustrates an exemplary execution of a vector frequency compress instruction according to one embodiment;

[0011] Figure 2 illustrates an exemplary execution of a vector frequency compress instruction according to one embodiment;

25 [0012] Figure 3 is a flow diagram illustrating exemplary operations for compressing values from a source vector register to a destination vector register by executing a vector frequency compress instruction in a processor according to one embodiment;

[0013] Figure 4A illustrates an exemplary AVX instruction format including a VEX prefix, real opcode field, Mod R/M byte, SIB byte, displacement field, and IMM8 according to one
30 embodiment;

[0014] Figure 4B illustrates which fields from Figure 4A make up a full opcode field and a base operation field according to one embodiment;

[0015] Figure 4C illustrates which fields from Figure 4A make up a register index field according to one embodiment;

- [0016] Figure 5A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments of the invention;
- [0017] Figure 5B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments of the invention
- 5 [0018] Figure 6A is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention;
- [0019] Figure 6B is a block diagram illustrating the fields of the specific vector friendly instruction format of Figure 6a that make up the full opcode field according to one embodiment of the invention;
- 10 [0020] Figure 6C is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the register index field according to one embodiment of the invention;
- [0021] Figure 6D is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the augmentation operation field according to one embodiment
- 15 of the invention;
- [0022] Figure 7 is a block diagram of a register architecture according to one embodiment of the invention;
- [0023] Figure 8A 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
- 20 the invention;
- [0024] Figure 8B 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;
- [0025] Figure 9A is a block diagram of a single processor core, along with its connection to
- 25 the on-die interconnect network and with its local subset of the Level 2 (L2) cache, according to embodiments of the invention;
- [0026] Figure 9B is an expanded view of part of the processor core in Figure 9A according to embodiments of the invention;
- [0027] Figure 10 is a block diagram of a processor that may have more than one core, may
- 30 have an integrated memory controller, and may have integrated graphics according to embodiments of the invention;
- [0028] Figure 11 is a block diagram of a system in accordance with one embodiment of the present invention;

[0029] Figure 12 is a block diagram of a first more specific exemplary system in accordance with an embodiment of the present invention;

[0030] Figure 13 is a block diagram of a second more specific exemplary system in accordance with an embodiment of the present invention;

5 [0031] Figure 14 is a block diagram of a SoC in accordance with an embodiment of the present invention; and

[0032] Figure 15 is 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.

10

DESCRIPTION OF EMBODIMENTS

[0033] In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description.

15

[0034] References in the specification to "one embodiment," "an embodiment," "an example embodiment," etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

20

[0035] Figure 1 illustrates an exemplary execution of a vector frequency compress instruction according to one embodiment. The vector frequency compress instruction implements RLE such that a set of source data elements are compressed based on sequential appearance of a predetermined value occurring in the set of source data elements. For example, a run of seven zeros in the set of source data elements would be encoded as a zero followed by seven to indicate the number of zeros.

25

30 [0036] The vector frequency compress instruction 100 includes a destination operand 105 and a source operand 115. The vector frequency compress instruction 100 belongs to an instruction set architecture, and each "occurrence" of the instruction 100 within an instruction stream would include values within the destination operand 105 and the source operand 115. In this example, both the destination operand 105 and the source operand 115 are vector registers (such as 128-,

256-, 512-bit registers). The vector registers may be zmm registers with 16 32-bit data elements, however, other data element and register sizes may be used such as xmm or ymm registers and 16- or 64-bit data elements. Thus, the source operand 115 and destination operand 105 are illustrated with 16 data elements, labeled as ele[0] for the first element using the 0-index notation where the first element is in the 0 position. The last element for each operand is then labeled ele[15].

[0037] The contents of the source vector register specified by the source operand include multiple source data elements. As illustrated in Figure 1, the source data element at index 0 contains the value 54. The source data elements at indexes 1-7 contain the value 0. The source data elements at indexes 8-10 contain the value 35. The source data element at index 11 contains the value 12, the source data element at index 12 contains the value 0, the source data element at index 13 contains the value 15, and the source data elements at indexes 14 and 15 contain the value 0.

[0038] The vector frequency compress instruction 100 is illustrated for compressing occurrences of data elements containing the value 0. However, optionally, vector frequency compress instruction 100 can be implemented for compressing occurrences of data elements containing other values. Thus, there is an optional operand 120 illustrated as an immediate value that could be encoded with the value that should be compressed. Thus, 120 would represent the value that should be compressed. Furthermore, the vector frequency compress instruction 100 is encoded with a vector mask MASK 110 that specifies which data elements match the value to be compressed and which data elements do not match the value to be compressed. The MASK 110 contains either zeros or ones depending on the operation of the instruction. In the illustrated embodiment, MASK 110 contains a one in the elements corresponding to the source data elements that do not match the value to be compressed and a zero in those elements that do match the value (zero as the case may be) to be compressed. Thus, MASK 110 contains a zero in mask elements 1-7, 12, 14, and 15.

[0039] The operation of the vector frequency compress instruction 100 can be described in the following manner. When a sequence of zeros is encountered in the MASK 110, indicating that a run of matching values are present in the source data elements, the instruction counts the number of successive zeros in the MASK 110 and inserts the value zero in the current position of the destination 105 vector followed by the count in the next position. The values in the source 115 data elements that have a corresponding one in the MASK 110 are copied into the current position of the destination 105 vector.

[0040] As shown in Figure 1, if the source 115 comprises the data elements 15 to 0 containing values: 0, 0, 15, 0, 12, 35, 35, 35, 0, 0, 0, 0, 0, 0, 0, 54 then the destination 105 vector register would comprise the following values in data elements 11 to 0: 2, 0, 15, 1, 0, 12, 35, 35, 35, 7, 0, 54. Going from element 0 to element 15, the first element of the source 115 vector register would be encountered with a 1 in the MASK 110 and 54 would be copied into the current position of the destination 105 vector register. Then a run of seven zeros would be counted and the values 0 and 7 would be copied into the next two positions of the destination 105 vector register. After which the values 35, 35, 35, and 12 would be copied into the next positions of the destination 105 vector register. Now a run of a single 0 is encountered so compression of the 0 actually requires 2 elements to store just a single 0. So, a 0 and a count of 1 are stored in the next data elements of the destination 105 vector register. Lastly a 15 followed by a 0 and a 2 are stored in the next positions of the destination 105 vector register representing the 15 and last two zeros of the source 115 vector register. Using the exemplary source 115 vector register data elements, zero-based RLE compressed the vector from 16 data elements to 12 data elements yielding a savings of 25 percent.

[0041] Embodiments of the invention output an indication of which elements of the destination 105 vector register were used in the encoding of the source 115 vector register. In one embodiment, this includes setting values in a used element control mask 125. Each used element in the destination 105 vector register corresponds with a 1 in the used element control mask 125. The used element control mask 125 may be a repurposed mask register or a special included mask register. Further embodiments of the instruction 100 allow the used element control mask 125 to be selected and encoded into the occurrence of the instruction. In Figure 1, elements 11-0 are set with a 1 in the used element control mask 125 to indicate which elements in the destination 105 vector register were used. In another embodiment, a used element count 130 is set to indicate the number of used elements in the destination 105 vector register, in this case with the value 12. As with the used element control mask 125, the used element count 130 may be set in a special register or into a repurposed general register. Furthermore, embodiments of the instruction 100 allow a destination, memory or register, to be encoded for the used element count 130. Although illustrated as counting the used data elements, the opposite could be used in other embodiments of the instructions. Specifically, the used element control mask and used element count could be an unused element control mask and an unused element count.

[0042] Figure 2 illustrates an exemplary execution of a vector frequency compress instruction according to one embodiment. The vector frequency compress instruction 200 is shown compressing a value other than zero. As described above, one embodiment includes an

immediate value 220 encoded into the instruction 200 to indicate the value to be compressed. Another embodiment of the instruction uses the MASK 210 to determine a value that is to be compressed by reading the value stored in the source 215 vector register corresponding to one or more MASK 210 elements with a 0. By design, the values of the source 215 vector register can be determined and the MASK 210 computed to correspond with the most efficient encoding for a given input. In Figure 2, the value 35 is being compressed as it will create the longest run to optimize compression of the source 215 vector register. The data elements 15 to 0 of the source 215 vector register contain the values: 0, 0, 15, 0, 12, 35, 35, 35, 0, 7, 0, 98, 44, 0, 0, and 54. These values are compressed using RLE of the value 35 into the destination 205 vector register to data elements 14-0 as the value: 0, 0, 15, 0, 12, 3, 35, 0, 7, 0, 98, 44, 0, 0, and 54. As shown in Figure 2, the run of 3 35s are compressed to 35 followed by 3 in the destination 205 vector register. In Figure 2, the data elements 14-0 of the used element control mask 225 are set to 1 while the used element count 230 is set to 15 to indicate which elements of the destination 205 vector register comprise the compressed source 215 vector register.

[0043] There exist cases of input source vector register that may not be entirely handled by the instruction. For example, there may be cases where the source vector register contains values such that the compressed destination vector register would be larger than the source vector register. Such as when the source contains one or more single occurrences of the value to be compressed since each single value is translated into two data elements of the destination, value and count. This situation can be easily detected by software by counting the number of occurrences of the value to be compressed and counting the number of zeros in the computed control mask. Software can then break the compression into two cycles to account for this scenario or attempt to optimize for the input using a different RLE scheme. Another option for this input scenario is for the instruction to raise a special exception whenever this case occurs and allow software to be triggered to handle the case.

[0044] Figure 3 is a flow diagram illustrating exemplary operations for compressing values from a source vector register to a destination vector register by executing a vector frequency compress instruction in a processor according to one embodiment. At operation 310, a vector frequency compress instruction is fetched by the processor (e.g., by a fetch unit of the processor). The vector frequency compress instruction includes, at least, a source operand and a destination operand. The source operand specifies a set of data elements to be compressed into the destination operand (e.g., an xmm, ymm, or zmm register). In at least one occurrence of the instruction, the source operand comprises one or more runs of values that are to be RLE into a value and a count of said value.

[0045] Flow moves from operation 310 to operation 315 where the processor decodes the vector frequency compress instruction. For example, in some embodiments, the processor includes a hardware decode unit that is provided the instruction (e.g., by the fetch unit of the processor). A variety of different well known decode units could be used for the decode unit.

5 For example, the decode unit may decode the vector frequency compress instruction into a single wide micro instruction. As another example, the decode unit may decode the vector frequency compress instruction into multiple wide micro instructions. As another example particularly suited for out of order processor pipelines, the decode unit may decode the vector frequency compress instruction into one or more micro-ops, where each of the micro-ops may be issued
10 and executed out of order. Also, the decode unit may be implemented with one or more decoders and each decoder may be implemented as a programmable logic array (PLA), as is well known in the art. By way of example, a given decode unit may: 1) have steering logic to direct different macro instructions to different decoders; 2) a first decoder that may decode a subset of the instruction set (but more of it than the second, third, and fourth decoders) and generate two
15 micro-ops at a time; 3) a second, third, and fourth decoder that may each decode only a subset of the entire instruction set and generate only one micro-op at a time; 4) a micro-sequencer ROM that may decode only a subset of the entire instruction set and generate four micro-ops at a time; and 5) multiplexing logic feed by the decoders and the micro-sequencer ROM that determine whose output is provided to a micro-op queue. Other embodiments of the decode unit may have
20 more or less decoders that decode more or less instructions and instruction subsets. For example, one embodiment may have a second, third, and fourth decoder that may each generate two micro-ops at a time; and may include a micro-sequencer ROM that generates eight micro-ops at a time.

[0046] Flow then moves to operation 320 where the processor executes the vector frequency compress instruction causing, for each of the source data elements, a value to be copied into the destination register to indicate that source data element's value. At least a subset of the source data operand is encoded into the destination vector register as a value and count thereby compressing the source data elements into fewer destination data elements.

[0047] Exemplary Instruction Formats

30 **[0048]** 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.

[0049] *VEX Instruction Format*

[0050] VEX encoding allows instructions to have more than two operands, and allows SIMD vector registers to be longer than 128 bits. The use of a VEX prefix provides for three-operand (or more) syntax. For example, previous two-operand instructions performed operations such as $A = A + B$, which overwrites a source operand. The use of a VEX prefix enables operands to perform nondestructive operations such as $A = B + C$.

[0051] Figure 4A illustrates an exemplary AVX instruction format including a VEX prefix 402, real opcode field 430, Mod R/M byte 440, SIB byte 450, displacement field 462, and IMM8 472. Figure 4B illustrates which fields from Figure 4A make up a full opcode field 474 and a base operation field 442. Figure 4C illustrates which fields from Figure 4A make up a register index field 444.

[0052] VEX Prefix (Bytes 0-2) 402 is encoded in a three-byte form. The first byte is the Format Field 440 (VEX Byte 0, bits [7:0]), which contains an explicit C4 byte value (the unique value used for distinguishing the C4 instruction format). The second-third bytes (VEX Bytes 1-2) include a number of bit fields providing specific capability. Specifically, REX field 405 (VEX Byte 1, bits [7-5]) consists of a VEX.R bit field (VEX Byte 1, bit [7] - R), VEX.X bit field (VEX byte 1, bit [6] - X), and VEX.B bit field (VEX byte 1, bit[5] - B). 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 VEX.R, VEX.X, and VEX.B. Opcode map field 415 (VEX byte 1, bits [4:0] - mmmmm) includes content to encode an implied leading opcode byte. W Field 464 (VEX byte 2, bit [7] - W) - is represented by the notation VEX.W, and provides different functions depending on the instruction. The role of VEX.vvvv 420 (VEX Byte 2, bits [6:3]-vvvv) may include the following: 1) VEX.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) VEX.vvvv encodes the destination register operand, specified in Is complement form for certain vector shifts; or 3) VEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. If VEX.L 468 Size field (VEX byte 2, bit [2]-L) = 0, it indicates 128 bit vector; if VEX.L = 1, it indicates 256 bit vector. Prefix encoding field 425 (VEX byte 2, bits [1:0]-pp) provides additional bits for the base operation field.

[0053] Real Opcode Field 430 (Byte 3) is also known as the opcode byte. Part of the opcode is specified in this field.

[0054] MOD R/M Field 440 (Byte 4) includes MOD field 442 (bits [7-6]), Reg field 444 (bits [5-3]), and R/M field 446 (bits [2-0]). The role of Reg field 444 may include the following: encoding either the destination register operand or a source register operand (the rrr of Rrrr), or

be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 446 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.

[0055] Scale, Index, Base (SIB) - The content of Scale field 450 (Byte 5) includes SS452 (bits [7-6]), which is used for memory address generation. The contents of SIB.xxx 454 (bits [5-3]) and SIB.bbb 456 (bits [2-0]) have been previously referred to with regard to the register indexes Xxxx and Bbbb.

[0056] The Displacement Field 462 and the immediate field (IMM8) 472 contain address data.

[0057] *Exemplary Encoding into VEX*

10 **[0058]** *Generic Vector Friendly Instruction Format*

[0059] 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.

15 **[0060]** Figures 5A-5B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention. Figure 5A 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 5B 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 500 for which are defined class A and class B instruction templates, both of which include no memory access 505 instruction templates and memory access 520 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.

25 **[0061]** 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).

[0062] The class A instruction templates in Figure 5A include: 1) within the no memory access 505 instruction templates there is shown a no memory access, full round control type operation 510 instruction template and a no memory access, data transform type operation 515 instruction template; and 2) within the memory access 520 instruction templates there is shown a memory access, temporal 525 instruction template and a memory access, non-temporal 530 instruction template. The class B instruction templates in Figure 5B include: 1) within the no memory access 505 instruction templates there is shown a no memory access, write mask control, partial round control type operation 512 instruction template and a no memory access, write mask control, vsize type operation 517 instruction template; and 2) within the memory access 520 instruction templates there is shown a memory access, write mask control 527 instruction template.

[0063] The generic vector friendly instruction format 500 includes the following fields listed below in the order illustrated in Figures 5A-5B.

[0064] Format field 540 - 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.

[0065] Base operation field 542 - its content distinguishes different base operations.

[0066] Register index field 544 - 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).

[0067] Modifier field 546 - 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 505 instruction templates and memory access 520 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.

- 5 **[0068]** Augmentation operation field 550 - 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 568, an alpha field 552, and a beta field 554. The augmentation operation field 550 allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions.
- 10 **[0069]** Scale field 560 - its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base}$).
- [0070]** Displacement Field 562A- its content is used as part of memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$).
- [0071]** Displacement Factor Field 562B (note that the juxtaposition of displacement field 562A
15 directly over displacement factor field 562B 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
20 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 574 (described later herein) and the data manipulation field 554C. The displacement field 562A and the displacement factor field 562B are optional in the sense that they are not used for the no memory access 505 instruction templates and/or different
25 embodiments may implement only one or none of the two.
- [0072]** Data element width field 564 - 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
30 the opcodes.
- [0073]** Write mask field 570 - 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
5 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);
10 however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field 570 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 570 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 570 content indirectly identifies that masking to be performed),
15 alternative embodiments instead or additional allow the mask write field's 570 content to directly specify the masking to be performed.

[0074] Immediate field 572 - its content allows for the specification of an immediate. This field is optional in the sense that is it 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
20 use an immediate.

[0075] Class field 568 - its content distinguishes between different classes of instructions. With reference to Figures 5A-B, the contents of this field select between class A and class B instructions. In Figures 5A-B, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 568A and class B 568B for the class field 568 respectively in
25 Figures 5A-B).

[0076] Instruction Templates of Class A

[0077] In the case of the non-memory access 505 instruction templates of class A, the alpha field 552 is interpreted as an RS field 552A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 552A.1 and data
30 transform 552A.2 are respectively specified for the no memory access, round type operation 510 and the no memory access, data transform type operation 515 instruction templates), while the beta field 554 distinguishes which of the operations of the specified type is to be performed. In the no memory access 505 instruction templates, the scale field 560, the displacement field 562A, and the displacement scale filed 562B are not present.

[0078] No-Memory Access Instruction Templates - Full Round Control Type Operation

[0079] In the no memory access full round control type operation 510 instruction template, the beta field 554 is interpreted as a round control field 554A, whose content(s) provide static rounding. While in the described embodiments of the invention the round control field 554A includes a suppress all floating point exceptions (SAE) field 556 and a round operation control field 558, alternative embodiments may support may encode both these concepts into the same field or only have one or the other of these concepts/fields (e.g., may have only the round operation control field 558).

[0080] SAE field 556 - its content distinguishes whether or not to disable the exception event reporting; when the SAE field's 556 content indicates suppression is enabled, a given instruction does not report any kind of floating-point exception flag and does not raise any floating point exception handler.

[0081] Round operation control field 558 - 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 558 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 550 content overrides that register value.

[0082] No Memory Access Instruction Templates - Data Transform Type Operation

[0083] In the no memory access data transform type operation 515 instruction template, the beta field 554 is interpreted as a data transform field 554B, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

[0084] In the case of a memory access 520 instruction template of class A, the alpha field 552 is interpreted as an eviction hint field 552B, whose content distinguishes which one of the eviction hints is to be used (in Figure 5A, temporal 552B.1 and non-temporal 552B.2 are respectively specified for the memory access, temporal 525 instruction template and the memory access, non-temporal 530 instruction template), while the beta field 554 is interpreted as a data manipulation field 554C, whose content distinguishes which one of a number of data 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 520 instruction templates include the scale field 560, and optionally the displacement field 562A or the displacement scale field 562B.

[0085] Vector memory instructions perform vector loads from and vector stores to memory, with 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.

5 **[0086]** Memory Access Instruction Templates - Temporal

[0087] Temporal data is data likely to be reused soon enough to benefit from caching. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

[0088] Memory Access Instruction Templates - Non-Temporal

10 **[0089]** 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.

[0090] Instruction Templates of Class B

15 **[0091]** In the case of the instruction templates of class B, the alpha field 552 is interpreted as a write mask control (Z) field 552C, whose content distinguishes whether the write masking controlled by the write mask field 570 should be a merging or a zeroing.

20 **[0092]** In the case of the non-memory access 505 instruction templates of class B, part of the beta field 554 is interpreted as an RL field 557A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 557A.1 and vector length (VSIZE) 557A.2 are respectively specified for the no memory access, write mask control, partial round control type operation 512 instruction template and the no memory access, write mask control, VSIZE type operation 517 instruction template), while the rest of the beta field 554 distinguishes which of the operations of the specified type is to be performed. In the no memory access 505 instruction templates, the scale field 560, the displacement field 562A, and the
25 displacement scale field 562B are not present.

[0093] In the no memory access, write mask control, partial round control type operation 510 instruction template, the rest of the beta field 554 is interpreted as a round operation field 559A 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).

30 **[0094]** Round operation control field 559A - just as round operation control field 558, 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 559A 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 550 content overrides that register value.

5 [0095] In the no memory access, write mask control, VSIZE type operation 517 instruction template, the rest of the beta field 554 is interpreted as a vector length field 559B, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

10 [0096] In the case of a memory access 520 instruction template of class B, part of the beta field 554 is interpreted as a broadcast field 557B, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field 554 is interpreted the vector length field 559B. The memory access 520 instruction templates include the scale field 560, and optionally the displacement field 562A or the displacement scale field 562B.

15 [0097] With regard to the generic vector friendly instruction format 500, a full opcode field 574 is shown including the format field 540, the base operation field 542, and the data element width field 564. While one embodiment is shown where the full opcode field 574 includes all of these fields, the full opcode field 574 includes less than all of these fields in embodiments that do not support all of them. The full opcode field 574 provides the operation code (opcode).

20 [0098] The augmentation operation field 550, the data element width field 564, and the write mask field 570 allow these features to be specified on a per instruction basis in the generic vector friendly instruction format.

[0099] 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.

25 [00100] 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
30 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.

[00101] Exemplary Specific Vector Friendly Instruction Format

[00102] Figure 6A is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention. Figure 6A shows a specific vector friendly instruction format 600 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 600 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 5 into which the fields from Figure 6A map are illustrated.

[00103] It should be understood that, although embodiments of the invention are described with reference to the specific vector friendly instruction format 600 in the context of the generic vector friendly instruction format 500 for illustrative purposes, the invention is not limited to the specific vector friendly instruction format 600 except where claimed. For example, the generic vector friendly instruction format 500 contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format 600 is shown as having fields of specific sizes. By way of specific example, while the data element width field 564 is illustrated as a one bit field in the specific vector friendly instruction format 600, the invention is not so limited (that is, the generic vector friendly instruction format 500 contemplates other sizes of the data element width field 564).

[00104] The generic vector friendly instruction format 500 includes the following fields listed below in the order illustrated in Figure 6A.

[00105] EVEX Prefix (Bytes 0-3) 602 - is encoded in a four-byte form.

[00106] Format Field 540 (EVEX Byte 0, bits [7:0]) - the first byte (EVEX Byte 0) is the format field 540 and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment of the invention).

5 **[00107]** The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

[00108] REX field 605 (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 557BEX byte 1, bit[5] - B).

The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the
10 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.

[00109] REX' field 510 - this is the first part of the REX' field 510 and is the EVEX.R' bit
15 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
20 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.

[00110] Opcode map field 615 (EVEX byte 1, bits [3:0] - mmmm) - its content encodes an implied leading opcode byte (OF, OF 38, or OF 3).

25 **[00111]** Data element width field 564 (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).

[00112] EVEX.vvvv 620 (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
30 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 620 encodes the 4 low-order bits of the first source register

specifier stored in inverted (Is complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

[00113] EVEX.U 568 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.

5 **[00114]** Prefix encoding field 625 (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 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
 10 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 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
 15 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 support the 2 bit SIMD prefix encodings, and thus not require the expansion.

[00115] Alpha field 552 (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
 20 described, this field is context specific.

[00116] Beta field 554 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s_{2θ}, EVEX.r_{2θ}, EVEX.rrl, EVEX.LL0, EVEX.LLB; also illustrated with βββ) - as previously described, this field is context specific.

[00117] REX' field 510 - this is the remainder of the REX' field and is the EVEX.V bit field
 25 (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.

[00118] Write mask field 570 (EVEX byte 3, bits [2:0]-kkk) - its content specifies the index of
 30 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).

[00119] Real Opcode Field 630 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

[00120] MOD R/M Field 640 (Byte 5) includes MOD field 642, Reg field 644, and R/M field 646. As previously described, the MOD field's 642 content distinguishes between memory access and non-memory access operations. The role of Reg field 644 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 646 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.

[00121] Scale, Index, Base (SIB) Byte (Byte 6) - As previously described, the scale field's content is used for memory address generation. SIB.xxx and SIB.bbb - the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb.

[00122] Displacement field 562A (Bytes 7-10) - when MOD field 642 contains 10, bytes 7-10 are the displacement field 562A, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

[00123] Displacement factor field 562B (Byte 7) - when MOD field 642 contains 01, byte 7 is the displacement factor field 562B. 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 a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to disp8 and disp32, the displacement factor field 562B is a reinterpretation of disp8; when using displacement factor field 562B, 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 562B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 562B 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).

[00124] Immediate field 572 operates as previously described.

[00125] Full Opcode Field

5 **[00126]** Figure 6B is a block diagram illustrating the fields of the specific vector friendly instruction format 600 that make up the full opcode field 574 according to one embodiment of the invention. Specifically, the full opcode field 574 includes the format field 540, the base operation field 542, and the data element width (W) field 564. The base operation field 542 includes the prefix encoding field 625, the opcode map field 615, and the real opcode field 630.

10 **[00127] Register Index Field**

[00128] Figure 6C is a block diagram illustrating the fields of the specific vector friendly instruction format 600 that make up the register index field 544 according to one embodiment of the invention. Specifically, the register index field 544 includes the REX field 605, the REX' field 610, the MODR/M.reg field 644, the MODR/M.r/m field 646, the W W field 620, xxx field 654, and the bbb field 656.

[00129] Augmentation Operation Field

[00130] Figure 6D is a block diagram illustrating the fields of the specific vector friendly instruction format 600 that make up the augmentation operation field 550 according to one embodiment of the invention. When the class (U) field 568 contains 0, it signifies EVEX.U0 (class A 568A); when it contains 1, it signifies EVEX.U1 (class B 568B). When U=0 and the MOD field 642 contains 11 (signifying a no memory access operation), the alpha field 552 (EVEX byte 3, bit [7] - EH) is interpreted as the rs field 552A. When the rs field 552A contains a 1 (round 552A.1), the beta field 554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the round control field 554A. The round control field 554A includes a one bit SAE field 556 and a two bit round operation field 558. When the rs field 552A contains a 0 (data transform 552A.2), the beta field 554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data transform field 554B. When U=0 and the MOD field 642 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 552 (EVEX byte 3, bit [7] - EH) is interpreted as the eviction hint (EH) field 552B and the beta field 554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data manipulation field 554C.

[00131] When U=1, the alpha field 552 (EVEX byte 3, bit [7] - EH) is interpreted as the write mask control (Z) field 552C. When U=1 and the MOD field 642 contains 11 (signifying a no memory access operation), part of the beta field 554 (EVEX byte 3, bit [4]- So) is interpreted as the RL field 557A; when it contains a 1 (round 557A.1) the rest of the beta field 554 (EVEX

byte 3, bit [6-5]- S_{2-1}) is interpreted as the round operation field 559A, while when the RL field 557A contains a 0 (VSIZE 557.A2) the rest of the beta field 554 (EVEX byte 3, bit [6-5]- S_{2-i}) is interpreted as the vector length field 559B (EVEX byte 3, bit [6-5]- L_{1-0}). When U=1 and the MOD field 642 contains 00, 01, or 10 (signifying a memory access operation), the beta field 554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the vector length field 559B (EVEX byte 3, bit [6-5]- L_{i-0}) and the broadcast field 557B (EVEX byte 3, bit [4]- B).

[00132] Exemplary Register Architecture

[00133] Figure 7 is a block diagram of a register architecture 700 according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers 710 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 600 operates on these overlaid register file as illustrated in the below tables.

[00134] Adjustable Vector Length	[00135] Class	[00136] Operations	[00137] Registers
[00138] Instruction Templates that do not include the vector length field 559B	[00139] A (Figure 5A; U=0)	[00140] 510, 515, 525, 530	[00141] zmm registers (the vector length is 64 byte)
	[00142] B (Figure 5B; U=1)	[00143] 512	[00144] zmm registers (the vector length is 64 byte)
[00145] Instruction Templates that do include the vector length field 559B	[00146] B (Figure 5B; U=1)	[00147] 517, 527	[00148] zmm, ymm, or xmm registers (the vector length is 64 byte, 32 byte, or 16 byte) depending on the vector length field 559B

[00149] In other words, the vector length field 559B 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 559B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 600 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.

5 [00150] Write mask registers 715 - 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 715 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.

[00151] General-purpose registers 725 - 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.

20 [00152] Scalar floating point stack register file (x87 stack) 745, on which is aliased the MMX packed integer flat register file 750 - 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.

[00153] 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.

[00154] Exemplary Core Architectures, Processors, and Computer Architectures

30 [00155] Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of

different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

[00156] Exemplary Core Architectures

15 [00157] In-order and out-of-order core block diagram

[00158] Figure 8A 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 8B 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 8A-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.

25 [00159] In Figure 8A, a processor pipeline 800 includes a fetch stage 802, a length decode stage 804, a decode stage 806, an allocation stage 808, a renaming stage 810, a scheduling (also known as a dispatch or issue) stage 812, a register read/memory read stage 814, an execute stage 816, a write back/memory write stage 818, an exception handling stage 822, and a commit stage 824.

30 [00160] Figure 8B shows processor core 890 including a front end unit 830 coupled to an execution engine unit 850, and both are coupled to a memory unit 870. The core 890 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 890 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.

[00161] The front end unit 830 includes a branch prediction unit 832 coupled to an instruction cache unit 834, which is coupled to an instruction translation lookaside buffer (TLB) 836, which is coupled to an instruction fetch unit 838, which is coupled to a decode unit 840. The decode unit 840 (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 840 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 890 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 840 or otherwise within the front end unit 830). The decode unit 840 is coupled to a rename/allocator unit 852 in the execution engine unit 850.

[00162] The execution engine unit 850 includes the rename/allocator unit 852 coupled to a retirement unit 854 and a set of one or more scheduler unit(s) 856. The scheduler unit(s) 856 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 856 is coupled to the physical register file(s) unit(s) 858. Each of the physical register file(s) units 858 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 858 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) 858 is overlapped by the retirement unit 854 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 854 and the physical register file(s) unit(s) 858 are coupled to the execution cluster(s) 860. The execution cluster(s) 860 includes a set of one or more execution units 862 and a set of one or more memory access units 864. The execution units 862 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) 856, physical register file(s) unit(s) 858, and execution cluster(s) 860 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) 864). 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.

[00163] The set of memory access units 864 is coupled to the memory unit 870, which includes a data TLB unit 872 coupled to a data cache unit 874 coupled to a level 2 (L2) cache unit 876. In one exemplary embodiment, the memory access units 864 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 872 in the memory unit 870. The instruction cache unit 834 is further coupled to a level 2 (L2) cache unit 876 in the memory unit 870. The L2 cache unit 876 is coupled to one or more other levels of cache and eventually to a main memory.

[00164] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 800 as follows: 1) the instruction fetch 838 performs the fetch and length decoding stages 802 and 804; 2) the decode unit 840 performs the decode stage 806; 3) the rename/allocator unit 852 performs the allocation stage 808 and renaming stage 810; 4) the scheduler unit(s) 856 performs the schedule stage 812; 5) the physical register file(s) unit(s) 858 and the memory unit 870 perform the register read/memory read stage 814; the execution cluster 860 perform the execute stage 816; 6) the memory unit 870 and the physical register file(s) unit(s) 858 perform the write back/memory write stage 818; 7) various units may be involved in the exception handling stage 822; and 8) the retirement unit 854 and the physical register file(s) unit(s) 858 perform the commit stage 824.

[00165] The core 890 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 890 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) previously described), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[00166] 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).

[00167] 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 834/874 and a shared L2 cache unit 876, 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.

[00168] Specific Exemplary In-Order Core Architecture

[00169] Figures 9A-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.

[00170] Figure 9A is a block diagram of a single processor core, along with its connection to the on-die interconnect network 902 and with its local subset of the Level 2 (L2) cache 904, according to embodiments of the invention. In one embodiment, an instruction decoder 900 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 906 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 908 and a vector unit 910 use separate register sets (respectively, scalar registers 912 and vector registers 914) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 906, 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).

[00171] The local subset of the L2 cache 904 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 904. Data read by a processor core is stored in its L2 cache subset 904 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 904 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.

10 [00172] Figure 9B is an expanded view of part of the processor core in Figure 9A according to embodiments of the invention. Figure 9B includes an LI data cache 906A part of the LI cache 904, as well as more detail regarding the vector unit 910 and the vector registers 914. Specifically, the vector unit 910 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 928), 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 920, numeric conversion with numeric convert units 922A-B, and replication with replication unit 924 on the memory input. Write mask registers 926 allow predicating resulting vector writes.

[00173] *Processor with integrated memory controller and graphics*

[00174] Figure 10 is a block diagram of a processor 1000 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to 20 embodiments of the invention. The solid lined boxes in Figure 10 illustrate a processor 1000 with a single core 1002A, a system agent 1010, a set of one or more bus controller units 1016, while the optional addition of the dashed lined boxes illustrates an alternative processor 1000 with multiple cores 1002A-N, a set of one or more integrated memory controller unit(s) 1014 in 25 the system agent unit 1010, and special purpose logic 1008.

[00175] Thus, different implementations of the processor 1000 may include: 1) a CPU with the special purpose logic 1008 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 1002A-N being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of 30 the two); 2) a coprocessor with the cores 1002A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 1002A-N being a large number of general purpose in-order cores. Thus, the processor 1000 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 1000 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.

[00176] The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units 1006, and external memory (not shown) coupled to the set of integrated memory controller units 1014. The set of shared cache units 1006 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 1012 interconnects the integrated graphics logic 1008, the set of shared cache units 1006, and the system agent unit 1010/integrated memory controller unit(s) 1014, 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 1006 and cores 1002-A-N.

[00177] In some embodiments, one or more of the cores 1002A-N are capable of multi-threading. The system agent 1010 includes those components coordinating and operating cores 1002A-N. The system agent unit 1010 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 1002A-N and the integrated graphics logic 1008. The display unit is for driving one or more externally connected displays.

[00178] The cores 1002A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 1002A-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.

[00179] *Exemplary Computer Architectures*

[00180] Figures 11-14 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.

[00181] Referring now to Figure 11, shown is a block diagram of a system 1100 in accordance with one embodiment of the present invention. The system 1100 may include one or more processors 1110, 1115, which are coupled to a controller hub 1120. In one embodiment the controller hub 1120 includes a graphics memory controller hub (GMCH) 1190 and an
5 Input/Output Hub (IOH) 1150 (which may be on separate chips); the GMCH 1190 includes memory and graphics controllers to which are coupled memory 1140 and a coprocessor 1145; the IOH 1150 is couples input/output (I/O) devices 1160 to the GMCH 1190. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory 1140 and the coprocessor 1145 are coupled directly to the processor 1110,
10 and the controller hub 1120 in a single chip with the IOH 1150.

[00182] The optional nature of additional processors 1115 is denoted in Figure 11 with broken lines. Each processor 1110, 1115 may include one or more of the processing cores described herein and may be some version of the processor 1000.

[00183] The memory 1140 may be, for example, dynamic random access memory (DRAM),
15 phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub 1120 communicates with the processor(s) 1110, 1115 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection 1195.

[00184] In one embodiment, the coprocessor 1145 is a special-purpose processor, such as, for
20 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 1120 may include an integrated graphics accelerator.

[00185] There can be a variety of differences between the physical resources 1110, 1115 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal,
25 power consumption characteristics, and the like.

[00186] In one embodiment, the processor 1110 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 1110 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 1145. Accordingly, the processor 1110
30 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 1145. Coprocessor(s) 1145 accept and execute the received coprocessor instructions.

[00187] Referring now to Figure 12, shown is a block diagram of a first more specific exemplary system 1200 in accordance with an embodiment of the present invention. As shown

in Figure 12, multiprocessor system 1200 is a point-to-point interconnect system, and includes a first processor 1270 and a second processor 1280 coupled via a point-to-point interconnect 1250. Each of processors 1270 and 1280 may be some version of the processor 1000. In one embodiment of the invention, processors 1270 and 1280 are respectively processors 1110 and
5 1115, while coprocessor 1238 is coprocessor 1145. In another embodiment, processors 1270 and 1280 are respectively processor 1110 coprocessor 1145.

[00188] Processors 1270 and 1280 are shown including integrated memory controller (IMC) units 1272 and 1282, respectively. Processor 1270 also includes as part of its bus controller units point-to-point (P-P) interfaces 1276 and 1278; similarly, second processor 1280 includes P-
10 P interfaces 1286 and 1288. Processors 1270, 1280 may exchange information via a point-to-point (P-P) interface 1250 using P-P interface circuits 1278, 1288. As shown in Figure 12, IMCs 1272 and 1282 couple the processors to respective memories, namely a memory 1232 and a memory 1234, which may be portions of main memory locally attached to the respective processors.

[00189] Processors 1270, 1280 may each exchange information with a chipset 1290 via individual P-P interfaces 1252, 1254 using point to point interface circuits 1276, 1294, 1286,
15 1298. Chipset 1290 may optionally exchange information with the coprocessor 1238 via a high-performance interface 1239. In one embodiment, the coprocessor 1238 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication
20 processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

[00190] 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.

[00191] Chipset 1290 may be coupled to a first bus 1216 via an interface 1296. In one embodiment, first bus 1216 may be a Peripheral Component Interconnect (PCI) bus, or a bus
25 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.

[00192] As shown in Figure 12, various I/O devices 1214 may be coupled to first bus 1216,
30 along with a bus bridge 1218 which couples first bus 1216 to a second bus 1220. In one embodiment, one or more additional processor(s) 1215, 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 1216. In one embodiment, second bus 1220 may be a low pin count (LPC) bus.

Various devices may be coupled to a second bus 1220 including, for example, a keyboard and/or mouse 1222, communication devices 1227 and a storage unit 1228 such as a disk drive or other mass storage device which may include instructions/code and data 1230, in one embodiment.

Further, an audio I/O 1224 may be coupled to the second bus 1220. Note that other architectures are possible. For example, instead of the point-to-point architecture of Figure 12, a system may implement a multi-drop bus or other such architecture.

[00193] Referring now to Figure 13, shown is a block diagram of a second more specific exemplary system 1300 in accordance with an embodiment of the present invention. Like elements in Figures 12 and 13 bear like reference numerals, and certain aspects of Figure 12 have been omitted from Figure 13 in order to avoid obscuring other aspects of Figure 13.

[00194] Figure 13 illustrates that the processors 1270, 1280 may include integrated memory and I/O control logic ("CL") 1272 and 1282, respectively. Thus, the CL 1272, 1282 include integrated memory controller units and include I/O control logic. Figure 13 illustrates that not only are the memories 1232, 1234 coupled to the CL 1272, 1282, but also that I/O devices 1314 are also coupled to the control logic 1272, 1282. Legacy I/O devices 1315 are coupled to the chipset 1290.

[00195] Referring now to Figure 14, shown is a block diagram of a SoC 1400 in accordance with an embodiment of the present invention. Similar elements in Figure 10 bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In Figure 14, an interconnect unit(s) 1402 is coupled to: an application processor 1410 which includes a set of one or more cores 202A-N and shared cache unit(s) 1006; a system agent unit 1010; a bus controller unit(s) 1016; an integrated memory controller unit(s) 1014; a set or one or more coprocessors 1420 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 1430; a direct memory access (DMA) unit 1432; and a display unit 1440 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 1420 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.

[00196] 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.

[00197] Program code, such as code 1230 illustrated in Figure 12, 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.

[00198] 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.

[00199] One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

[00200] Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk 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 memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[00201] 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 system features described herein. Such embodiments may also be referred to as program products.

[00202] *Emulation (including binary translation, code morphing, etc.)*

[00203] In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

[00204] Figure 15 is 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. 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 15 shows a program in a high level language 1502 may be compiled using an x86 compiler 1504 to generate x86 binary code 1506 that may be natively executed by a processor with at least one x86 instruction set core 1516. The processor with at least one x86 instruction set core 1516 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 1504 represents a compiler that is operable to generate x86 binary code 1506 (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 1516. Similarly, Figure 15 shows the program in the high level language 1502 may be compiled using an alternative instruction set compiler 1508 to generate alternative instruction set binary code 1510 that may be natively executed by a processor without at least one x86 instruction set core 1514 (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 1512 is used to convert the x86 binary code 1506 into code that may be natively executed by the processor without an x86 instruction set core 1514. This converted code is not likely to be the same as the alternative instruction set binary code 1510 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 1512 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 1506.

5 [00205] While the flow diagrams in the figures show a particular order of operations performed by certain embodiments of the invention, it should be understood that such order is exemplary (e.g., alternative embodiments may perform the operations in a different order, combine certain operations, overlap certain operations, etc.).

10 [00206] While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described, can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting.

CLAIMS

What is claimed is:

1. A method of performing a vector frequency compress instruction in a computer processor, comprising:
 - 5 fetching the vector frequency compress instruction that includes a source operand and a destination operand, wherein the source operand specifies a source vector register that includes a plurality of source data elements including one or more runs of identical data elements, wherein the destination operand identifies a destination vector register and wherein each of the one or more runs of identical values that
10 are to be compressed in the destination vector register as a value and run length pair;
decoding the fetched vector frequency compress instruction; and
executing the decoded vector frequency compress instruction causing, for each source data element, a value to be copied into the destination vector register to indicate
15 that source data element's value wherein one or more runs of one or more source data elements equal to a compression value are encoded in the destination vector register as the predetermined compression value followed by a run length for that run.
 2. The method of claim 1, wherein the fetched vector frequency compress instruction
20 further comprises the compression value that is to be encoded to a value and run length pair.
 3. The method of claim 1, wherein the executing the decoded vector frequency compress instruction further causes an exception be raised when the source data elements cannot be compressed into the destination vector register because the source data elements do not contain values that are optimized for run length encoding.
 - 25 4. The method of claim 1, wherein the executing the decoded vector frequency compress instruction further causes a value be written in a used element indicator to indicate which elements in the destination vector register were written during compression.
 5. The method of claim 4, wherein the fetched vector frequency compress instruction further comprises a used element indicator destination to indicate where the used element
30 indicator should be written.
 6. The method of claim 1, wherein the fetched vector frequency compress instruction further comprises a control mask that indicates one or more values from the source data elements to be copied to the destination vector register.

7. The method of claim 6, wherein the executing the decoded vector frequency compress instruction further causes determining the compression value by reading the control mask.
8. A processor core, comprising:
a hardware decode unit to decode a vector frequency compress instruction, wherein the
5 vector frequency compress instruction includes a source operand and a destination operand, wherein the source operand specifies a source vector register that includes a plurality of source data elements including one or more runs of identical data elements, wherein the destination operand identifies a destination vector register and wherein each of the one or more runs of identical values that
10 are to be compressed in the destination vector register as a value and run length pair; and
an execution engine unit to execute the decoded vector frequency compress instruction which causes, for each source data element, a value to be copied into the destination vector register to indicate that source data element's value wherein
15 one or more runs of one or more source data elements equal to a compression value are encoded in the destination vector register as the predetermined compression value followed by a run length for that run.
9. The processor core of claim 8, wherein the vector frequency compress instruction further comprises the compression value that is to be encoded to a value and run length pair.
- 20 10. The processor core of claim 8, the execution unit further causes an exception be raised when the source data elements cannot be compressed into the destination vector register because the source data elements do not contain values that are optimized for run length encoding.
11. The processor core of claim 8, the execution unit further causes a value be written in a used element indicator to indicate which elements in the destination vector register were written
25 during compression.
12. The processor core of claim 11, wherein the vector frequency compress instruction further comprises a used element indicator destination to indicate where the used element indicator should be written.
13. The processor core of claim 8, wherein the vector frequency compress instruction further
30 comprises a control mask that indicates one or more values from the source data elements to be copied to the destination vector register.
14. The processor core of claim 13, the execution unit further causes determining the compression value by reading the control mask.
15. An article of manufacture, comprising:

a tangible machine-readable storage medium having stored thereon a vector frequency compress instruction, wherein the vector frequency compress instruction includes a source operand and a destination operand, wherein the source operand specifies a source vector register that includes a plurality of source data elements including one or more runs of identical data elements, wherein the destination operand identifies a destination vector register and wherein each of the one or more runs of identical values that are to be compressed in the destination vector register as a value and run length pair; and

wherein the vector frequency compress instruction includes an opcode, which instructs a machine to execute the instruction that causes, for each source data element, a value to be copied into the destination vector register to indicate that source data element's value wherein one or more runs of one or more source data elements equal to a compression value are encoded in the destination vector register as the predetermined compression value followed by a run length for that run.

16. The article of manufacture of claim 15, wherein the vector frequency compress instruction further comprises the compression value that is to be encoded to a value and run length pair.

17. The article of manufacture of claim 15, wherein the vector frequency compress instruction further causes the machine to raise an exception when the source data elements cannot be compressed into the destination vector register because the source data elements do not contain values that are optimized for run length encoding.

18. The article of manufacture of claim 15, wherein the vector frequency compress instruction further causes the machine to write a value in a used element indicator to indicate which elements in the destination vector register were written during compression.

19. The article of manufacture of claim 18, wherein the vector frequency compress instruction further comprises a used element indicator destination to indicate where the used element indicator should be written.

20. The article of manufacture of claim 15, wherein the vector frequency compress instruction further comprises a control mask that indicates one or more values from the source data elements to be copied to the destination vector register.




Figure 1




Figure 2

3/17

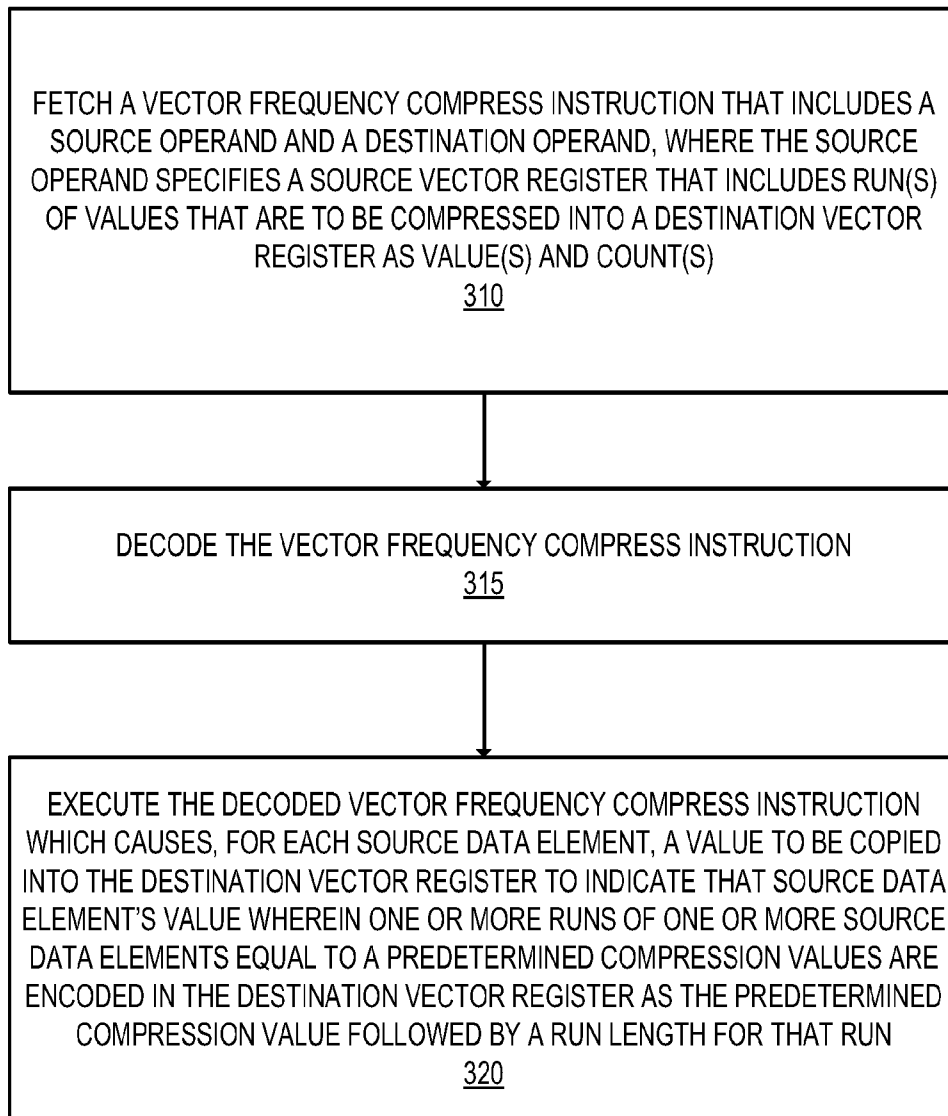
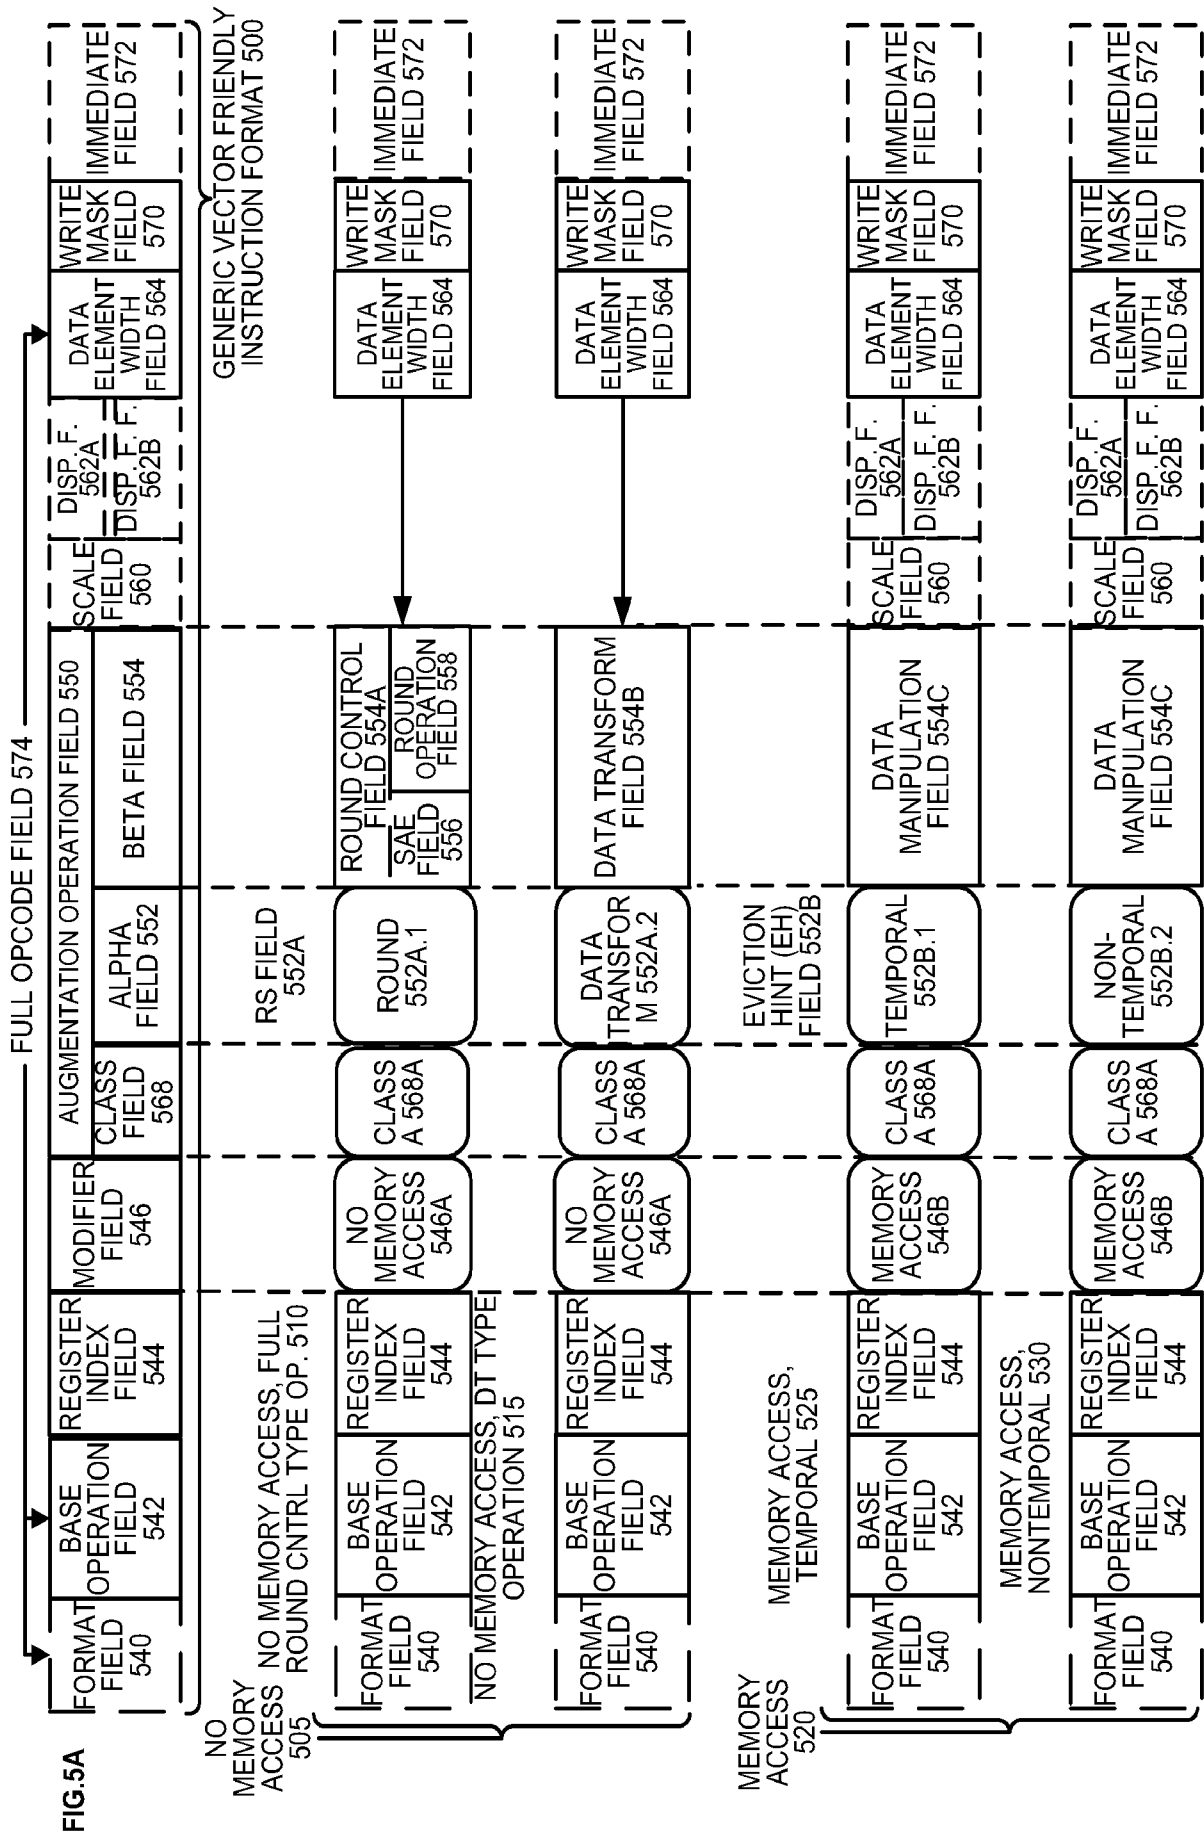
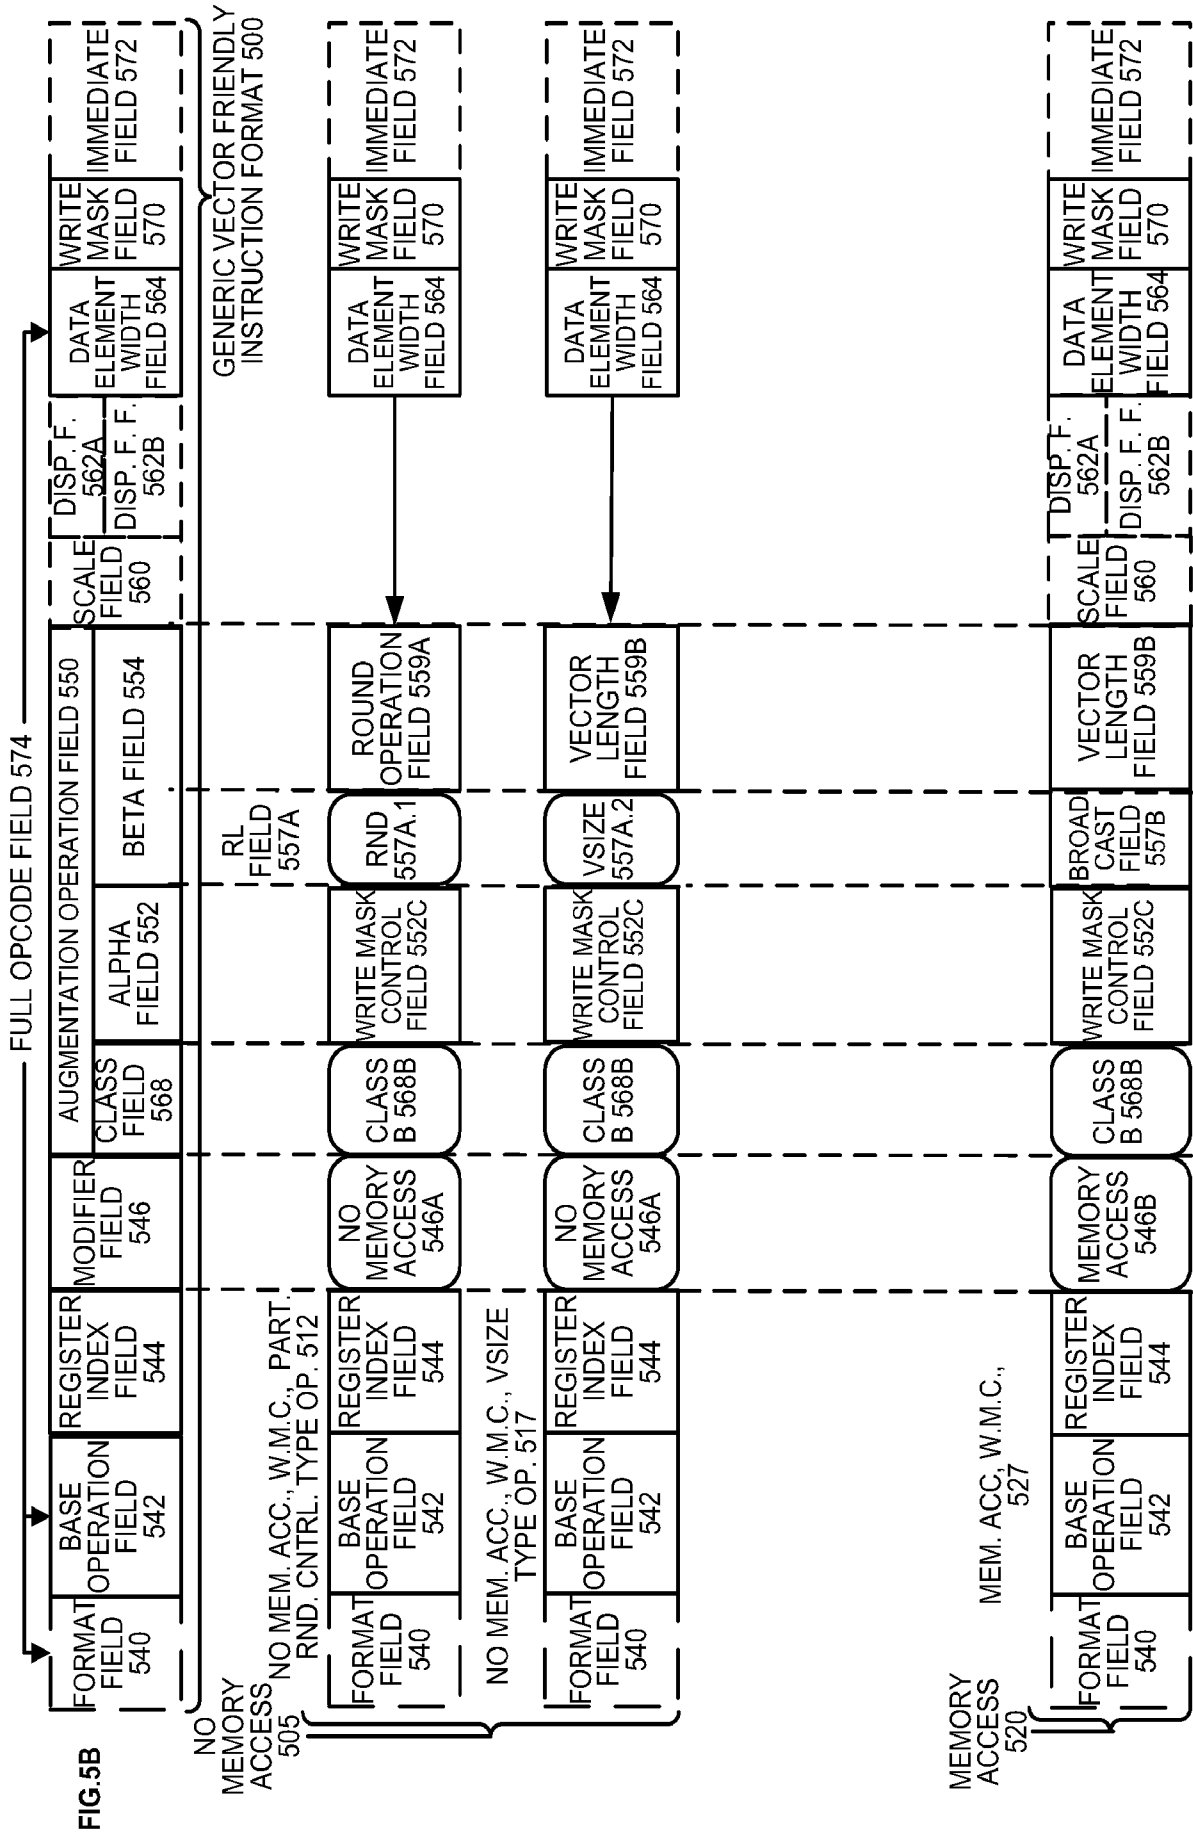


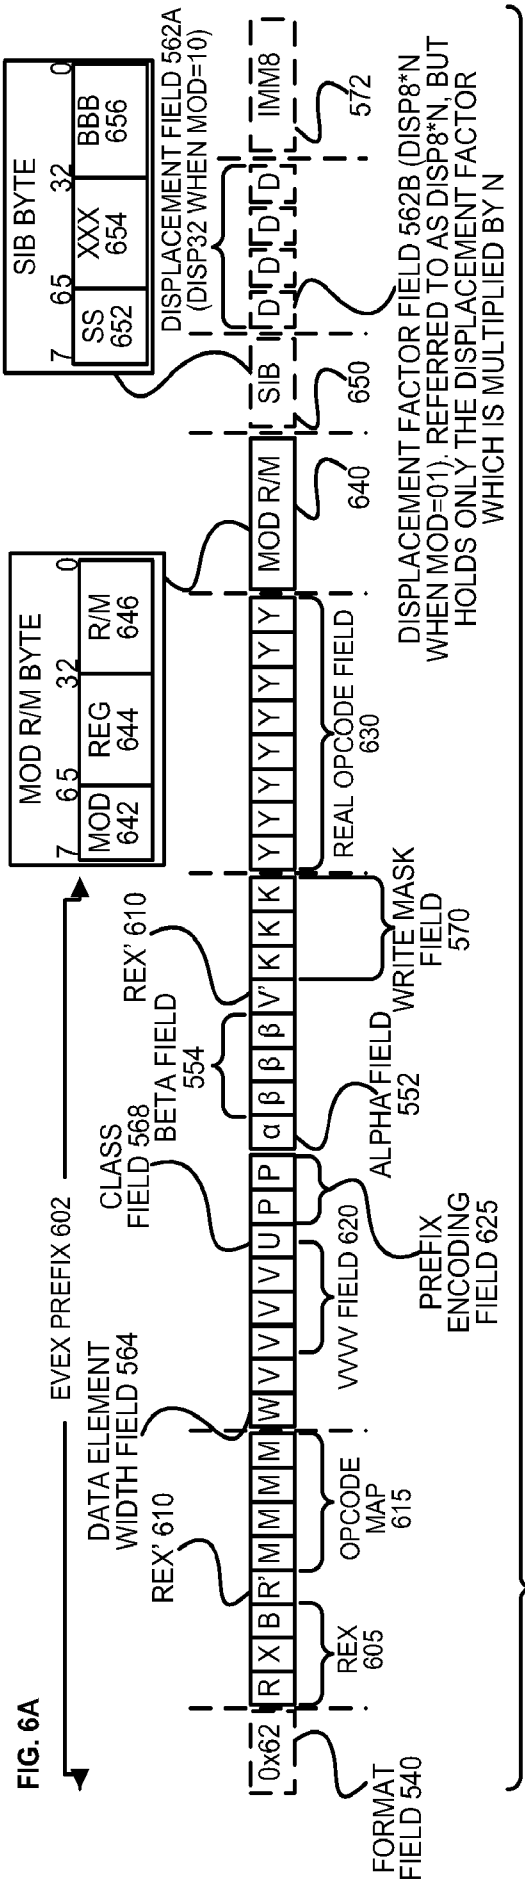
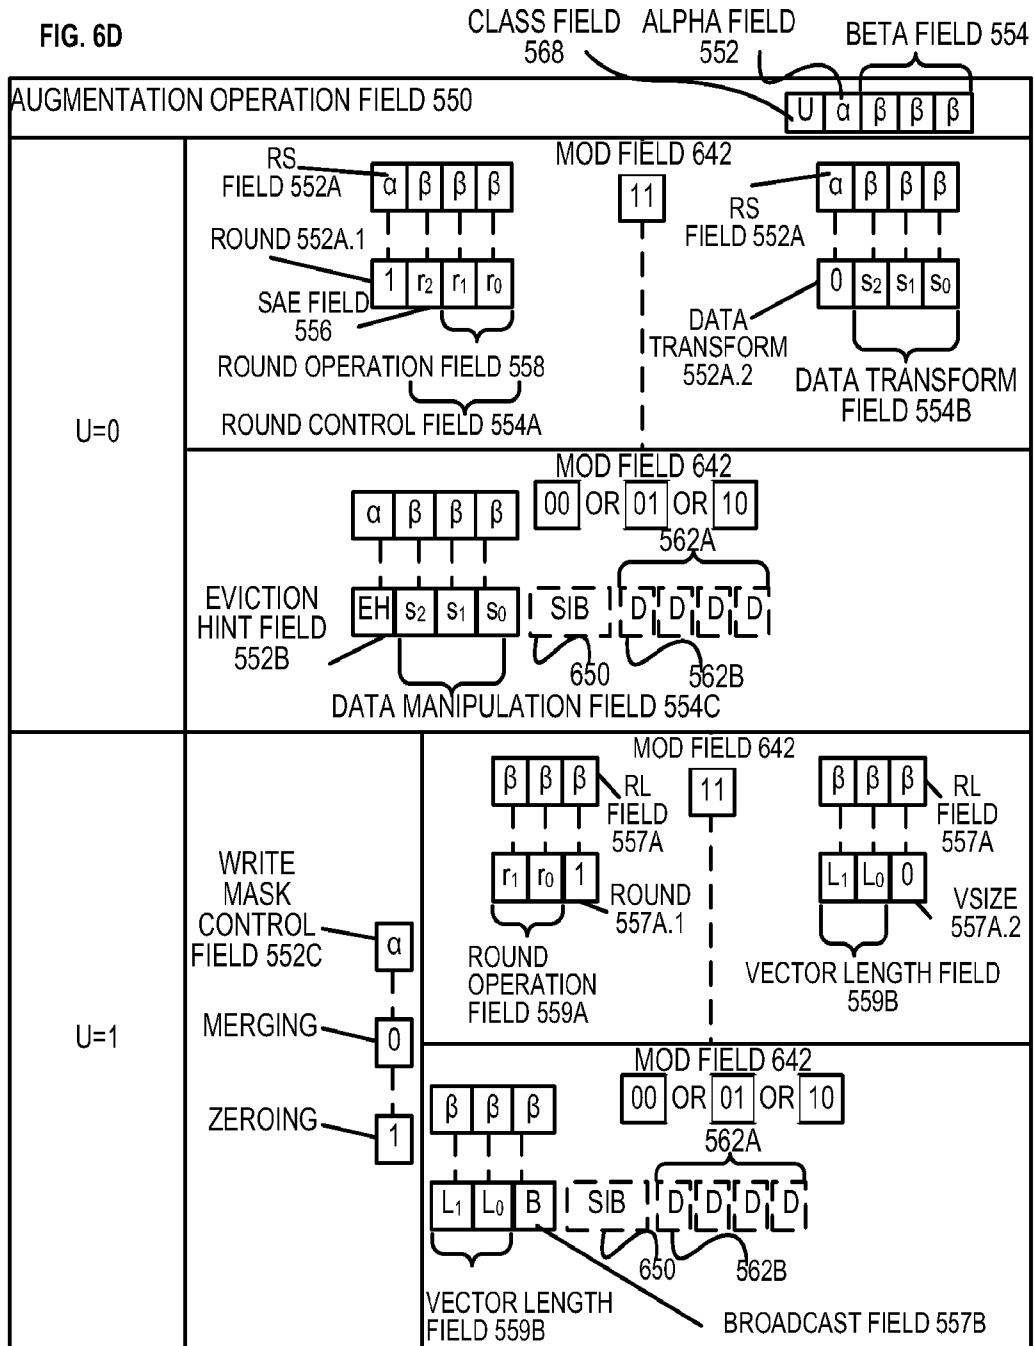



Figure 3

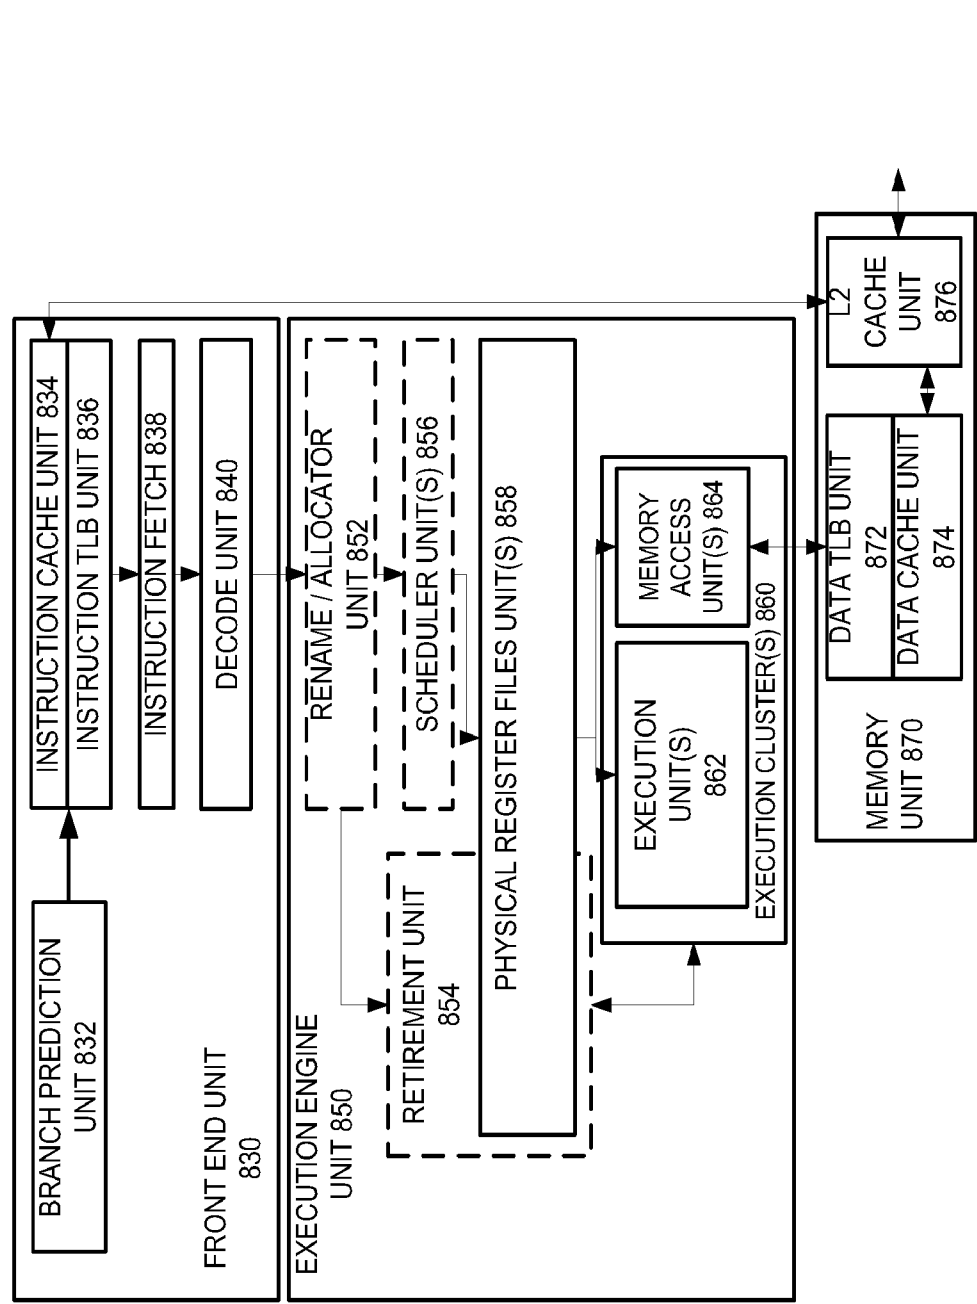




















FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14

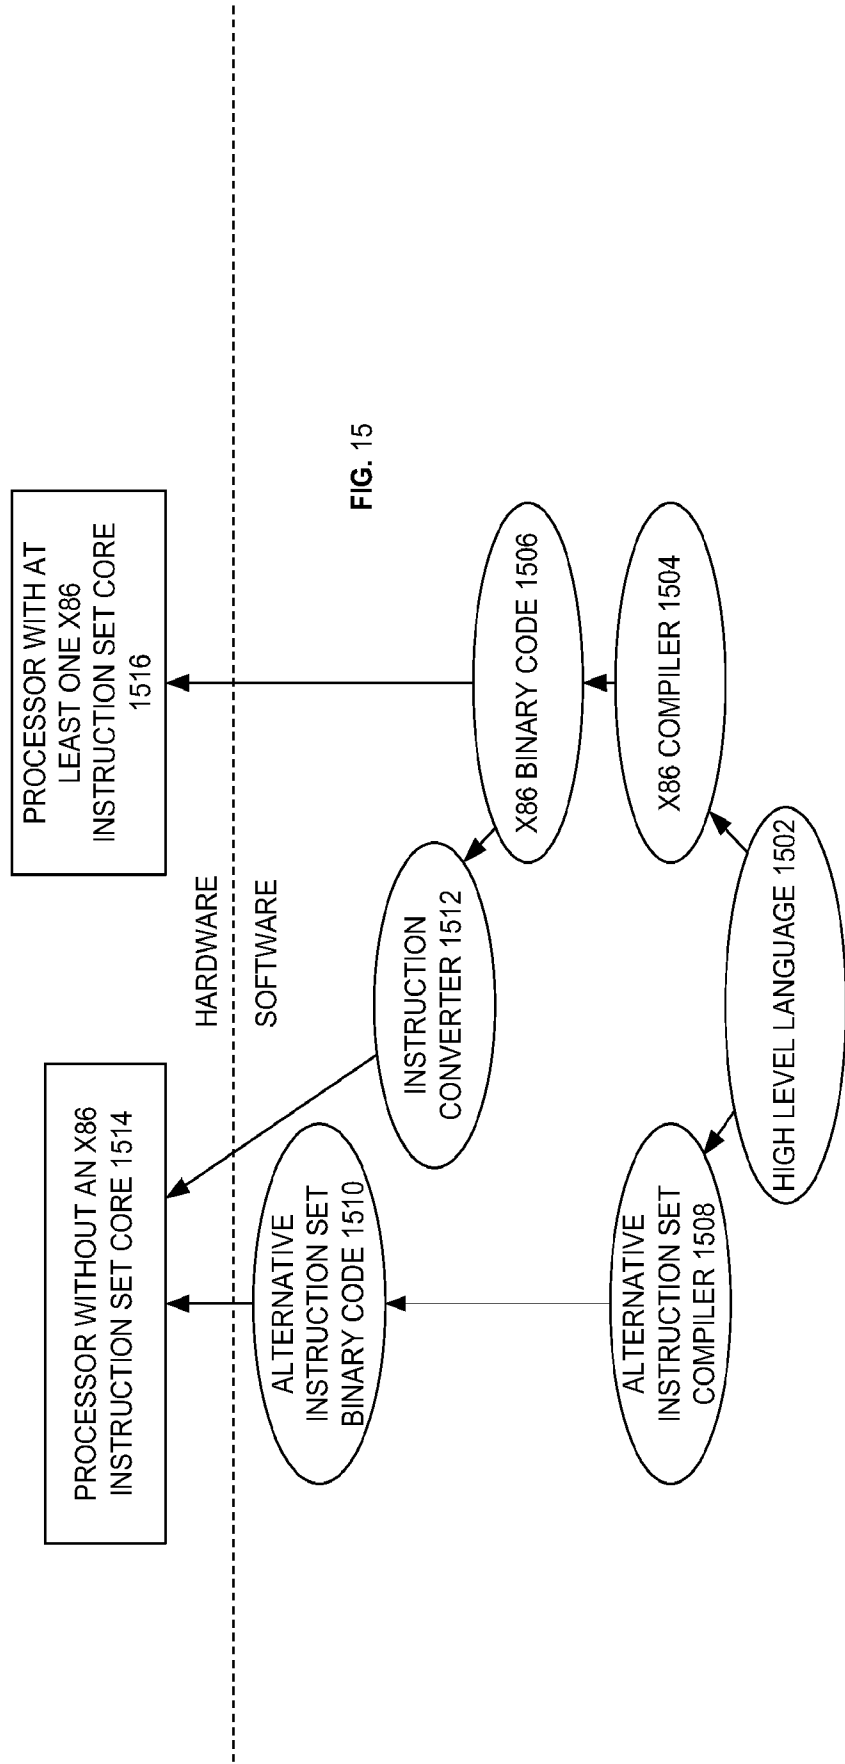


FIG. 15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2011/068230**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/30(2006.01)i, G06F 9/305(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/30; G10L 15/00; H03M 7/46; G06F 7/00; G06F 7/20

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: identical, value, compress, vector, frequency, encode;

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2007-0208792 A1 (ROBIN GERJON et al.) 6 September 2007 See abstract, paragraph [0011]-[0020], and claim 1.	1-20
A	US 2007-0016418 A1 (SANJEEV MEHROTRA et al.) 18 January 2007 See abstract, paragraphs [0017]-[0018], claims 1-3, and figures 1-4.	1-20
A	US 05109523 A (YASUSI KANDA et al.) 28 April 1992 See abstract, column 7, lines 30-48, claim 1, and figures 1-8.	1-20
A	US 6781528 B1 (CHIEN-HSIN LIN et al.) 24 August 2004 See abstract, column 8, lines 62-67, column 9, lines 1-14, claim 1, and figures 1-7.	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

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

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of 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

20 SEPTEMBER 2012 (20.09.2012)

Date of mailing of the international search report

21 SEPTEMBER 2012 (21.09.2012)

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

BOK, Jin Yo

Telephone No. 82-42-481-51 13



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US201 1/068230

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0208792 A1	06.09.2007	CA 2580 146 A1 CN 10 10 19325 A CN 10 10 19325 CO EP 1790080 A2 JP 2008-5129 17 A KR 10-2007-0083652 A Wo 2006-030288 A2 Wo 2006-030288 A3	23.03.2006 15.08.2007 15.08.2007 30.05.2007 24.04.2008 24.08.2007 23.03.2006 23.03.2006
US 2007-00 16418 A1	18.01.2007	AT 53600 1 T AU 2006-270259 A1 AU 2006-270259 B2 CA 26 12537 A1 CN 10 1223573 AO CN 10 1223573 B EP 1905000 A2 EP 1905000 A4 EP 1905000 B1 JP 2009-50 1943 A JP 2009-50 1943 T KR 10-2008-0025399 A MX 2008000524 A NO 2007626 1 A NO 2007626 1 B US 7599840 B2 wo 2007-01 1653 A2 wo 2007-01 1653 A3	15.12.2011 25.01.2007 03.03.2011 25.01.2007 16.07.2008 27.07.2011 02.04.2008 26.08.2009 30.11.2011 22.01.2009 22.01.2009 20.03.2008 06.03.2008 06.02.2008 06.02.2008 06.10.2009 25.01.2007 07.06.2007
US 05109523 A	28.04.1992	us 05 109523 A	28.04.1992
US 678 1528 B1	24.08.2004	None	