



(51) International Patent Classification:
G06F 9/345 (2006.01)

(21) International Application Number:
PCT/US2017/037553

(22) International Filing Date:
14 June 2017 (14.06.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/201,391 02 July 2016 (02.07.2016) US

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

(72) Inventors: **PLOTNIKOV, Mikhail**; Turgeneva Str 30, 603024 Nizhny Novgorod (RU). **OULD-AHMED-VALL, Elmoustapha**; 5000 W Chandler Blvd, M/S: CH7-401, Chandler, Arizona 85226 (US).

(74) Agent: **NICHOLSON, David F.**; Nicholson De Vos Webster & Elliott, LLP, 99 Almaden Blvd., Ste. 710, San Jose, California 95113 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

(54) Title: SYSTEMS, APPARATUSES, AND METHODS FOR STRIDED LOAD




FIG. 1

(57) Abstract: Systems, methods, and apparatuses for strided loads are described. In an embodiment, an instruction to include at least an opcode, a field for at least two packed data source operands, a field for a packed data destination operand, and an immediate is designated as a strided load instruction. This instruction is executed to load packed data elements from the at least two packed data source operands using a stride and storing results of the strided loads in the packed data destination operand starting from a defined position determined in part from the immediate.

MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

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

SYSTEMS, APPARATUSES, AND METHODS FOR STRIDED LOAD

FIELD OF INVENTION

[0001] The field of invention relates generally to computer processor architecture, and, more specifically, to an instruction which when executed causes a particular result.

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).

[0003] 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. For example, Intel Pentium 4 processors, Intel Core processors, and Advanced Micro Devices, Inc. of Sunnyvale CA processors implement nearly identical versions of the x86 instruction set (with some extensions having been added to newer versions), but have different internal designs. For example, the same register architecture of the ISA may be implemented in different ways in different micro-architectures using well known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB) and a retirement register file; the use of multiple maps and a pool of registers), etc. Unless otherwise specified, the phrases register architecture, register file, and register refer to that which is visible to the software/programmer and the manner in which instructions specify registers. Where specificity is desired, the adjective logical, architectural, or software visible will be used to indicate registers/files in the register architecture, while different adjectives will be used to designate registers in a given micro-architecture (e.g., physical register, reorder buffer, retirement register, register pool).

[0004] 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.

[0005] 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).

[0006] 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 as 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).

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0008] Figure 1 illustrates an embodiment of hardware to process a strideload instruction;

[0009] Figure 2 illustrates an example of an execution of a strideload instruction according to an embodiment;

[0010] Figure 3 illustrates an example of an execution of a strideload instruction according to an embodiment;

[0011] Figure 4 illustrates embodiments of the strideload instruction;

[0012] Figure 5 illustrates an embodiment of method performed by a processor to process a strideload instruction;

[0013] Figure 6 illustrates an embodiment of the execution portion of the method performed by a processor to process a strideload instruction;

[0014] Figure 7 illustrates an embodiment of a pseudocode representation of the execution of a strideload instruction;

[0015] Figures 8A-8B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention;

[0016] Figures 9A-D illustrate an exemplary specific vector friendly instruction format according to embodiments of the invention;

[0017] Figure 10 is a block diagram of a register architecture according to one embodiment of the invention;

[0018] Figure 11A 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;

[0019] Figure 11B 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;

[0020] Figures 12A-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;

[0021] Figure 13 is a block diagram of a processor 1300 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention;

[0022] Figures 14-16 are block diagrams of exemplary computer architectures; and

[0023] Figure 17 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.

DETAILED DESCRIPTION

[0024] 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.

[0025] 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.

[0026] Embodiments detailed herein may apply to the problem of vectorizing loops with loads from arrays of structures. Such pattern is also known as stride load. For example, consider an array of structures, which consists of 3 fields:

```
typedef struct {
    double x;
    double y;
    double z;
} xyz;
xyz *A;
```

and a with accesses to the fields of consecutive elements of the array:

```
for(i=0; i<N; i++){
    x_local = A[i].x;
    y_local = A[i].y;
    z_local = A[i].z;
    computation(x_local, y_local, z_local);
    ...
}
```

[0027] To vectorize the loop with vector registers of length VL bits having KL elements of a given data type (in the example KL=VL/64, as 64-bit floating point data is accessed, for VL=512, KL=8) need to gather elements to a vector register from memory. Elements reside in memory with a stride 3:

```

x_local_vec = [A[i+7].x:A[i+6].x:A[i+5].x:A[i+4].x:A[i+3].x:A[i+2].x:A[i+1].x:A[i+0].x]
y_local_vec = [A[i+7].y:A[i+6].y:A[i+5].y:A[i+4].y:A[i+3].y:A[i+2].y:A[i+1].y:A[i+0].y]
z_local_vec = [A[i+7].z:A[i+6].z:A[i+5].z:A[i+4].z:A[i+3].z:A[i+2].z:A[i+1].z:A[i+0].z]

```

[0028] In an instruction set which has a gather instruction, such loads can be done by 3 adjacent gathers with the same indexes:

```

x_local_vec = gatherqpd[&A[i+0] + v_index*8]
y_local_vec = gatherqpd[&A[i+0] + v_index*8 + 8]
z_local_vec = gatherqpd[&A[i+0] + v_index*8 + 16],

```

where $v_index = [21:18:15:12:9:6:3:0]$ – vector of indexes, which determine offsets of elements from the base address.

[0029] Unfortunately, using gathers may have some drawbacks. For example, the gather implementation itself is more complex than a regular vector load, because indexes are not known in general case and require processing of possible exceptions. Thus, performance is not optimal. Additionally, the locality of x_local , y_local , z_local is not taken into consideration (they may be in the same cache line).

[0030] To resolve these problems, a gather-to-shuffle (G2S) optimization may be enabled which loads several consecutive vectors to vector registers and then shuffles (or permutes) data accordingly to several destinations with a sequence of permute instructions. In an exemplary case, 3 vectors of 8 elements should be loaded and permuted to 3 destinations. A problem with G2S optimization is that permute controls for each destination are unique and need to be generated and preserved in separate vector registers for the time of the whole loop execution. A stride can have larger value (N) and the number of unique permute controls grows as $N*(N-1)$. As such, register pressure is increased dramatically. For stride 5 one would need 20 registers to hold all the permute controls. The number of permute instructions in sequences also grows.

[0031] A stride pattern also may appear in a loop going with increment value>1:

```

for(i=0; i<N; i+=5){
    x_local = B[i];
    y_local = B[i+1];
    z_local = B[i+2];
    u_local = B[i+3];
    v_local = B[i+4];
    computation(x_local, y_local, z_local, u_local, v_local);
...

```

}

[0032] Consider current compiler generated permute sequence for stride 3:

```
vmovdqu32 zmm5, ZMMWORD PTR .L_2il0floatpacket.2[rip] //load permute
vmovdqu32 zmm4, ZMMWORD PTR .L_2il0floatpacket.3[rip] //controls from
vmovdqu32 zmm3, ZMMWORD PTR .L_2il0floatpacket.4[rip] //memory before
vmovdqu32 zmm2, ZMMWORD PTR .L_2il0floatpacket.5[rip] //loop
vmovdqu32 zmm1, ZMMWORD PTR .L_2il0floatpacket.6[rip]
vmovdqu32 zmm0, ZMMWORD PTR .L_2il0floatpacket.7[rip]
```

..B1.12:

```
vmovdqu32 zmm8, ZMMWORD PTR [rcx] //load of 3 input vectors
vmovdqu32 zmm9, ZMMWORD PTR [64+rcx]
vmovdqu32 zmm11, ZMMWORD PTR [128+rcx]
vmovaps zmm6, zmm5 //restoring overwritten
vmovaps zmm7, zmm3 //permute controls
vmovaps zmm10, zmm1 //here
vperm2q zmm6, zmm8, zmm9
vperm2q zmm7, zmm8, zmm9
vperm2q zmm10, zmm8, zmm9
vmovaps zmm12, zmm4 //restoring overwritten
vmovaps zmm13, zmm2 //permute controls
vmovaps zmm14, zmm0 //here
vperm2q zmm12, zmm6, zmm11
vperm2q zmm13, zmm7, zmm11
vperm2q zmm14, zmm10, zmm11
cmp eax, r10d
jb ..B1.12
```

[0033] There are 6 permute controls and totally 12 instructions for permutation after 3 vectors are loaded (and plus 6 loads of permute controls from memory).

[0034] With an strideload instruction the sequence will look like:

```
vmovdqu32 zmm1, ZMMWORD PTR [rcx] //load of 3 input vectors
vmovdqu32 zmm2, ZMMWORD PTR [64+rcx]
vmovdqu32 zmm3, ZMMWORD PTR [128+rcx]
strideload zmm12, zmm1, zmm2, 010:000:00b //permutation itself
strideload zmm12, zmm3, zmm3, 010:010:01b
```

strideload zmm13, zmm1, zmm2, 010:001:00b

strideload zmm13, zmm3, zmm3, 010:000:01b

strideload zmm14, zmm1, zmm2, 010:010:00b

strideload zmm14, zmm3, zmm3, 010:001:01b

[0035] Here 6 instructions are for permutations after 3 loads, no register is preserved for permute controls and no memory references are for getting permute controls.

[0036] Detailed herein are embodiments apparatuses, systems, and methods using a strided load (“strideload”) instruction that when executed by a hardware processor causes a loading of packed data elements from at least two concatenated packed data source operands (srcM,...,src2, src1) using a stride and storing the results of the strided loads in a destination operand starting from a defined position. The defined position of the destination is defined by $(\text{round} * M * KL + \text{offset}) / \text{stride}$. The “round” value is provided by an immediate of the strideload instruction. KL is a number of packed data elements in a source (vector length (VL) divided by packed data element size). M is the number of concatenated source operands. The stride value is at least partially supplied by an immediate of the strideload instruction. A starting data element position of the concatenated sources is defined by an offset is supplied by an immediate of the strideload instruction. In an embodiment, the immediate is an 8-bit value and bits 7:5 are used to calculate the stride (the value of these bits plus 1), bits 4:2 provide the offset, and bits 1:0 provide the round. In some embodiments, the storing is under control of a writemask and a field for the writemask is included in the instruction.

[0037] Figure 1 illustrates an embodiment of hardware to process a strideload instruction. The illustrated hardware is typically a part of a hardware processor or core such as a part of a central processing unit, accelerator, etc.

[0038] A strideload instruction is received by decode circuitry 101. For example, the decode circuitry 101 receives this instruction from fetch logic/circuitry. The strideload instruction includes fields for two or more source operands (e.g., packed data registers (sometimes called vector registers)), a destination operand (e.g., a packed data register (sometimes called a vector register)), an opcode, and an immediate. In some embodiments, the strideload instruction includes a field for an immediate. More detailed embodiments of instruction formats will be detailed later.

[0039] The decode circuitry 101 decodes the strideload instruction into one or more operations. In some embodiments, this decoding includes generating a plurality of micro-

operations to be performed by execution circuitry (such as execution circuitry 109). The decode circuitry 101 also decodes instruction prefixes.

[0040] In some embodiments, register renaming, register allocation, and/or scheduling circuitry 103 provides functionality for one or more of: 1) renaming logical operand values to physical operand values (e.g., a register alias table in some embodiments), 2) allocating status bits and flags to the decoded instruction, and/or 3) scheduling the decoded instruction for execution on execution circuitry out of an instruction pool (e.g., using a reservation station in some embodiments).

[0041] Registers (register file) 105 and memory 107 store data as operands of the strideload instruction to be operated on by execution circuitry 109. Exemplary register types include packed data registers, general purpose registers, and floating point registers. The registers 105 may also include write mask registers as detailed herein.

[0042] Execution circuitry 109 executes the decoded strideload instruction to load packed data elements from at least two concatenated source operands (srcM, ..., src2, src1) using a stride and store the results of the strided loads in the destination operand starting from a defined position. The defined position of the destination is defined by $(\text{round} * M * KL + \text{offset}) / \text{stride}$. The “round” value is provided by an immediate of the strideload instruction. KL is a number of packed data elements in a source (vector length (VL) divided by packed data element size). M is the number of concatenated source operands. The stride value is at least partially supplied by an immediate of the strideload instruction. A starting data element position of the concatenated sources is defined by an offset is supplied by an immediate of the strideload instruction. In an embodiment, the immediate is an 8-bit value and bits 7:5 are used to calculate the stride (the value of these bits plus 1), bits 4:2 provide the offset, and bits 1:0 provide the round. In some embodiments, the storing is under control of a writemask and a field for the writemask is included in the instruction. In some embodiments, retirement circuitry 111 architecturally commits result (e.g., commits a destination register into the registers 105) and retires the instruction.

[0043] Figure 2 illustrates an example of an execution of a strideload instruction according to an embodiment. This example is not meant to be limiting. For example, while this example uses a little endian format, the teachings herein allow for a big endian format execution.

[0044] Typically, the number of packed data elements to extract and their sizes is dependent upon the instruction encoding (data element size). As such, a different number of

packed data elements such as 2, 4, 8, 16, 32, or 64 may be in a packed data source. Packed data destination register sizes include 64-bit, 128-bit, 256-bit, and 512-bit.

[0045] In this example, the sources 201 (e.g., packed data registers) include 8 packed data elements each. These sources 201 are concatenated for the purposes of the strided loads. The least significant data element positions store a “0”, “2”, and “5” respectively. These values are shown as decimal, however, they are typically stored as binary or hexadecimal values.

[0046] Load/store circuitry (not shown for ease of understanding) is used to extract data elements from the sources 201 according to a strided value and initial offset and store them consecutively in the destination 205.

[0047] In the round 0 strideload operation, the destination is REG0 and the sources are REG1 and REG2. The immediate is shown as 01000100 in binary. This equates to an offset of 1 (bits 4:2 of the immediate have a value of 1). This means that the first data element to be chosen is offset 1 from the least significant data element of the concatenated sources of REG1 and REG2. This data element is marked in REG1 and has a value of 0. The stride value to be used is the value of the immediate bits 7:5 (which is 010b or 2 in decimal) plus 1, which, in this example is a stride value of 3. The round is found by the two least significant bits 1:0 of the immediate and is 0.

[0048] The starting point for storing strided values into the destination 205 from the concatenated sources REG1 and REG2 is found using the following $(\text{round} * M * KL + \text{offset}) / \text{stride}$. As noted, the round is 0, the number of sources M is 2, the offset is 1, and the stride is 3. KL for this example is 8 (there are 8 data elements per source register). As such, the starting point for the destination is 0. In this example, only 5 data element locations are loaded into the destination 205 as shown for this round.

[0049] In the round 1 strideload operation, the destination is REG0 and the source is REG3. The immediate is shown as 01000000 in binary. This equates to an offset of 0 (bits 4:2 of the immediate have a value of 0). This means that the first data element to be chosen is offset 0 from the least significant data element of the concatenated source of REG3. This data element is marked in REG3 and has a value of 5. The stride value to be used is the value of the immediate bits 7:5 (which is 010b or 2 in decimal) plus 1, which, in this example is a stride value of 3. The round is found by the two least significant bits 1:0 of the immediate and is 1.

[0050] The starting point for storing strided values into the destination 205 from the concatenated source REG3 is found using the following $(\text{round} * M * KL + \text{offset}) / \text{stride}$. As noted, the round is 1, the number of sources M is 2, the offset is 0, and the stride is 3. KL for

this example is 8 (there are 8 data elements per source register). As such, the starting point for the destination is 5. In this example, only 3 data element locations are loaded into the destination 205 as shown for this round.

[0051] Figure 3 illustrates an example of an execution of a strideload instruction according to an embodiment. This example is not meant to be limiting. For example, while this example uses a little endian format, the teachings herein allow for a big endian format execution.

[0052] Typically, the number of packed data elements to extract and their sizes is dependent upon the instruction encoding (data element size). As such, a different number of packed data elements such as 2, 4, 8, 16, 32, or 64 may be in a packed data source. Packed data destination register sizes include 64-bit, 128-bit, 256-bit, and 512-bit.

[0053] In this example, the sources 301 (e.g., packed data registers) include 8 packed data elements each. These sources 301 are concatenated for the purposes of the strided loads. The least significant data element positions store a “0”, “2”, and “5” respectively. These values are shown as decimal, however, they are typically stored as binary or hexadecimal values.

[0054] Load/store circuitry (not shown for ease of understanding) is used to extract data elements from the sources 301 according to a strided value and initial offset and store them consecutively in the destination 305.

[0055] In the round 0 strideload operation, the destination is REG0 and the sources are REG1 and REG2. The immediate is shown as 01000100 in binary. This equates to an offset of 1 (bits 4:2 of the immediate have a value of 1). This means that the first data element to be chosen is offset 1 from the least significant data element of the concatenated sources of REG1 and REG2. This data element is marked in REG1 and has a value of 0. The stride value to be used is the value of the immediate bits 7:5 (which is 010b or 2 in decimal) plus 1, which, in this example is a stride value of 3. The round is found by the two least significant bits 1:0 of the immediate and is 0.

[0056] The starting point for storing strided values into the destination 305 from the concatenated sources REG1 and REG2 is found using the following $(\text{round} * M * KL + \text{offset}) / \text{stride}$. As noted, the round is 0, the number of sources M is 2, the offset is 1, and the stride is 3. KL for this example is 8 (there are 8 data elements per source register). As such, the starting point for the destination is 0. However, in this example, a writemask 303 (e.g., a writemask register) is used. In data element (or bit) position 3 of the writemask 303 is a 0 indicating that the strided value will not be written into the destination

305. In this example, only 4 data element locations are loaded into the destination 305 as shown for this round.

[0057] In the round 1 strideload operation, the destination is REG0 and the source is REG3. The immediate is shown as 01000000 in binary. This equates to an offset of 0 (bits 4:2 of the immediate have a value of 0). This means that the first data element to be chosen is offset 0 from the least significant data element of the concatenated source of REG3. This data element is marked in REG3 and has a value of 5. The stride value to be used is the value of the immediate bits 7:5 (which is 010b or 2 in decimal) plus 1, which, in this example is a stride value of 3. The round is found by the two least significant bits 1:0 of the immediate and is 1.

[0058] The starting point for storing strided values into the destination 305 from the concatenated source REG3 is found using the following $(\text{round} * M * KL + \text{offset}) / \text{stride}$. As noted, the round is 1, the number of sources M is 2, the offset is 0, and the stride is 3. KL for this example is 8 (there are 8 data elements per source register). As such, the starting point for the destination is 5. Again, in this example, the same writemask (e.g., the previously detailed writemask register) is used. In data element (or bit) position 5 of the writemask is a 0 indicating that the strided value will not be written into the destination 305. In this example, only 2 data element locations are loaded into the destination 305 as shown for this round.

[0059] An embodiment of a format (including fields) for a strideload instruction is `strideload{B/W/D/Q}{k} DST, SRCM...SRC0`. In some embodiments, `strideload{B/W/D/Q}` is the opcode of the instruction. `B/W/D/Q` indicates the data element sizes of the sources/destination as byte, word, doubleword, and quadword. `DST` is a packed data destination register and `SRCM...SRC0` are packed data source registers. `K` is a writemask that is used in some embodiments as detailed herein.

[0060] In some embodiments, the strideload instruction includes a writemask register operand. A writemask is used to conditionally control per-element operations and updating of results. Depending upon the implementation, the writemask uses merging or zeroing masking. Instructions encoded with a predicate (writemask, write mask, or `k` register) operand use that operand to conditionally control per-element computational operation and updating of result to the destination operand. The predicate operand is known as the opmask (writemask) register. The opmask is a set of eight architectural registers of size `MAX_KL` (64-bit). Note that from this set of 8 architectural registers, only `k1` through `k7` can be addressed as predicate operand. `k0` can be used as a regular source or destination but cannot

be encoded as a predicate operand. Note also that a predicate operand can be used to enable memory fault-suppression for some instructions with a memory operand (source or destination). As a predicate operand, the opmask registers contain one bit to govern the operation/update to each data element of a vector register. In general, opmask registers can support instructions with element sizes: single-precision floating-point (float32), integer doubleword (int32), double-precision floating-point (float64), integer quadword (int64). The length of a opmask register, MAX_KL, is sufficient to handle up to 64 elements with one bit per element, i.e. 64 bits. For a given vector length, each instruction accesses only the number of least significant mask bits that are needed based on its data type. An opmask register affects an instruction at per-element granularity. So, any numeric or non-numeric operation of each data element and per-element updates of intermediate results to the destination operand are predicated on the corresponding bit of the opmask register. In most embodiments, an opmask serving as a predicate operand obeys the following properties: 1) the instruction's operation is not performed for an element if the corresponding opmask bit is not set (this implies that no exception or violation can be caused by an operation on a masked-off element, and consequently, no exception flag is updated as a result of a masked-off operation); 2). a destination element is not updated with the result of the operation if the corresponding writemask bit is not set. Instead, the destination element value must be preserved (merging-masking) or it must be zeroed out (zeroing-masking); 3) for some instructions with a memory operand, memory faults are suppressed for elements with a mask bit of 0. Note that this feature provides a versatile construct to implement control-flow predication as the mask in effect provides a merging behavior for vector register destinations. As an alternative the masking can be used for zeroing instead of merging, so that the masked out elements are updated with 0 instead of preserving the old value. The zeroing behavior is provided to remove the implicit dependency on the old value when it is not needed.

[0061] Figure 4 illustrates embodiments of the strideload instruction including fields for the opcode 401, packed data destination operand 403, packed data source operands 405, an immediate 409, and, in some embodiments, a writemask operand 407. In some embodiments, it is not possible to explicitly encode enough source operands for the instructions. In those embodiments, implicit registers may be used (e.g., SRC0 is always a particular register or specified in general purpose register) or by encoding one source that indicates this register and then next n registers (for example, when an instruction that takes four sources has bits indicating SRC0 is an operand that also means that SRC1, SRC2 and SRC3 are the other three sources).

[0062] Figure 5 illustrates an embodiment of method performed by a processor to process a strideload instruction.

[0063] At 501, an instruction is fetched. For example, a strideload instruction is fetched. The strideload instruction includes fields for an opcode, at least one packed data source operand, an immediate, and a packed data destination operand. In some embodiments, the strideload instruction includes field for a writemask operand. In some embodiments, the instruction is fetched from an instruction cache.

[0064] The fetched instruction is decoded at 503. For example, the fetched strideload instruction is decoded by decode circuitry such as that detailed herein. In some embodiments, the instruction is decoded into one or more micro-operations.

[0065] Data associated with the source operand of the decoded instruction is retrieved at 505. For example, contiguous elements from memory are accessed beginning at the source address or a source and/or destination packed data register is accessed.

[0066] At 507, the decoded instruction is executed by execution circuitry (hardware) such as that detailed herein. For the strideload instruction, the execution will, load packed data elements from at least two concatenated packed data source operands (srcM,...,src2, src1) using a stride and store the results of the strided loads in a destination operand starting from a defined position. The defined position of the destination is defined by $(\text{round} * M * KL + \text{offset}) / \text{stride}$. The “round” value is provided by an immediate of the strideload instruction. KL is a number of packed data elements in a source (vector length (VL) divided by packed data element size). M is the number of concatenated source operands. The stride value is at least partially supplied by an immediate of the strideload instruction. A starting data element position of the concatenated sources is defined by an offset is supplied by an immediate of the strideload instruction. In an embodiment, the immediate is an 8-bit value and bits 7:5 are used to calculate the stride (the value of these bits plus 1), bits 4:2 provide the offset, and bits 1:0 provide the round. In some embodiments, the storing is under control of a writemask and a field for the writemask is included in the instruction.

[0067] In some embodiments, the instruction is committed or retired at 509.

[0068] Figure 6 illustrates an embodiment of the execution portion of the method performed by a processor to process a strideload instruction.

[0069] At 601, a determination of a number of data elements per packed data source register, an offset value, and a round value are made, along with a calculation of a stride value. The number of data elements per packed data source register (KL) is determined by

vector length (VL) divided by packed data element size (typically stated as a part of the opcode of the instruction). As detailed above, the instruction includes a field for an immediate which is used to get the other values. In an embodiment, the immediate is an 8-bit value and bits 7:5 are used to calculate the stride (the value of these bits plus 1), bits 4:2 provide the offset, and bits 1:0 provide the round. The offset is a starting data element position of the concatenated packed data source registers.

[0070] The source registers are concatenated at 603. In some embodiments, a temporary data structure is used for the concatenated result.

[0071] A determination of a starting data element position of the packed data destination register is made at 605. The starting position of the packed data destination register is defined by $(\text{round} * M * KL + \text{offset}) / \text{stride}$, where M is the number of concatenated packed data source registers.

[0072] At 607, a data element value from the concatenated packed data source registers at the offset data element position is loaded and stored into the determined starting data element position of the packed data destination register. In some embodiments, this storage is subject to a writemask.

[0073] At 609, a data element value from the concatenated packed data source registers at data element position stride values away from the immediately preceding loaded data element position is loaded and stored into the packed data destination register consecutively with the previous storage. In some embodiments, this storage is subject to a writemask.

[0074] A determination is made of if all of the strided data element positions of the concatenated packed data source registers have been exhausted is made at 611. If not, then the next strided data element position is loaded, etc. in 609.

[0075] If yes, in some embodiments, values in data element positions of the destination are zeroed at 613.

[0076] Figure 7 illustrates an embodiment of a pseudocode representation of the execution of a strideload instruction.

[0077] The figures below detail exemplary architectures and systems to implement embodiments of the above. In some embodiments, one or more hardware components and/or instructions described above are emulated as detailed below, or implemented as software modules.

[0078] Embodiments of the instruction(s) detailed above are embodied may be embodied in a “generic vector friendly instruction format” which is detailed below. In other embodiments, such a format is not utilized and another instruction format is used, however,

the description below of the writemask registers, various data transformations (swizzle, broadcast, etc.), addressing, etc. is generally applicable to the description of the embodiments of the instruction(s) above. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) above may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

[0079] An instruction set may include one or more instruction formats. A given instruction format may define various fields (e.g., number of bits, location of bits) to specify, among other things, the operation to be performed (e.g., opcode) and the operand(s) on which that operation is to be performed and/or other data field(s) (e.g., mask). Some instruction formats are further broken down through the definition of instruction templates (or subformats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands. A set of SIMD extensions referred to as the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developer's Manual, September 2014; and see Intel® Advanced Vector Extensions Programming Reference, October 2014).

Exemplary Instruction Formats

[0080] 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.

Generic Vector Friendly Instruction Format

[0081] 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.

[0082] **Figures 8A-8B** are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention. **Figure 8A** 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 8B** 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 800 for which are defined class A and class B instruction templates, both of which include no memory access 805 instruction templates and memory access 820 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.

[0083] 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).

[0084] The class A instruction templates in **Figure 8A** include: 1) within the no memory access 805 instruction templates there is shown a no memory access, full round control type operation 810 instruction template and a no memory access, data transform type operation 815 instruction template; and 2) within the memory access 820 instruction templates there is shown a memory access, temporal 825 instruction template and a memory access, non-temporal 830 instruction template. The class B instruction templates in **Figure 8B** include: 1) within the no memory access 805 instruction templates there is shown a no memory

access, write mask control, partial round control type operation 812 instruction template and a no memory access, write mask control, vsize type operation 817 instruction template; and 2) within the memory access 820 instruction templates there is shown a memory access, write mask control 827 instruction template.

[0085] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in **Figures 8A-8B**.

[0086] Format field 840 – 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.

[0087] Base operation field 842 – its content distinguishes different base operations.

[0088] Register index field 844 – 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).

[0089] Modifier field 846 – 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 805 instruction templates (846A) and memory access 820 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.

[0090] Augmentation operation field 850 – 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 868, an alpha field 852, and a beta field 854. The augmentation operation field 850 allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions.

[0091] Scale field 860 – its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base}$).

[0092] Displacement Field 862A– its content is used as part of memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$).

[0093] Displacement Factor Field 862B (note that the juxtaposition of displacement field 862A directly over displacement factor field 862B indicates one or the other is used) – its content is used as part of address generation; it specifies a displacement factor that is to be scaled by the size of a memory access (N) – where N is the number of bytes in the memory access (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{scaled displacement}$). Redundant low-order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) in order to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field 874 (described later herein) and the data manipulation field 854C. The displacement field 862A and the displacement factor field 862B are optional in the sense that they are not used for the no memory access 805 instruction templates and/or different embodiments may implement only one or none of the two.

[0094] Data element width field 864 – its content distinguishes which one of a number of data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

[0095] Write mask field 870 – its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other one embodiment, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in one embodiment, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control

the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field 870 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 870 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 870 content indirectly identifies that masking to be performed), alternative embodiments instead or additional allow the mask write field's 870 content to directly specify the masking to be performed.

[0096] Immediate field 872 – its content allows for the specification of an immediate. This field is optional in the sense that it is not present in an implementation of the generic vector friendly format that does not support immediate and it is not present in instructions that do not use an immediate.

[0097] Class field 868 – its content distinguishes between different classes of instructions. With reference to **Figures 8A-B**, the contents of this field select between class A and class B instructions. In **Figures 8A-B**, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 868A and class B 868B for the class field 868 respectively in **Figures 8A-B**).

Instruction Templates of Class A

[0098] In the case of the non-memory access 805 instruction templates of class A, the alpha field 852 is interpreted as an RS field 852A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 852A.1 and data transform 852A.2 are respectively specified for the no memory access, round type operation 810 and the no memory access, data transform type operation 815 instruction templates), while the beta field 854 distinguishes which of the operations of the specified type is to be performed. In the no memory access 805 instruction templates, the scale field 860, the displacement field 862A, and the displacement scale field 862B are not present.

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

[0099] In the no memory access full round control type operation 810 instruction template, the beta field 854 is interpreted as a round control field 854A, whose content(s) provide static rounding. While in the described embodiments of the invention the round control field 854A includes a suppress all floating point exceptions (SAE) field 856 and a

round operation control field 858, 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 858).

[0100] SAE field 856 – its content distinguishes whether or not to disable the exception event reporting; when the SAE field's 856 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.

[0101] Round operation control field 858 – 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 858 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 850 content overrides that register value.

No Memory Access Instruction Templates – Data Transform Type Operation

[0102] In the no memory access data transform type operation 815 instruction template, the beta field 854 is interpreted as a data transform field 854B, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

[0103] In the case of a memory access 820 instruction template of class A, the alpha field 852 is interpreted as an eviction hint field 852B, whose content distinguishes which one of the eviction hints is to be used (in **Figure 8A**, temporal 852B.1 and non-temporal 852B.2 are respectively specified for the memory access, temporal 825 instruction template and the memory access, non-temporal 830 instruction template), while the beta field 854 is interpreted as a data manipulation field 854C, 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 820 instruction templates include the scale field 860, and optionally the displacement field 862A or the displacement scale field 862B.

[0104] 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.

Memory Access Instruction Templates – Temporal

[0105] 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.

Memory Access Instruction Templates – Non-Temporal

[0106] Non-temporal data is data unlikely to be reused soon enough to benefit from caching in the 1st-level cache and should be given priority for eviction. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

Instruction Templates of Class B

[0107] In the case of the instruction templates of class B, the alpha field 852 is interpreted as a write mask control (Z) field 852C, whose content distinguishes whether the write masking controlled by the write mask field 870 should be a merging or a zeroing.

[0108] In the case of the non-memory access 805 instruction templates of class B, part of the beta field 854 is interpreted as an RL field 857A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 857A.1 and vector length (VSIZE) 857A.2 are respectively specified for the no memory access, write mask control, partial round control type operation 812 instruction template and the no memory access, write mask control, VSIZE type operation 817 instruction template), while the rest of the beta field 854 distinguishes which of the operations of the specified type is to be performed. In the no memory access 805 instruction templates, the scale field 860, the displacement field 862A, and the displacement scale field 862B are not present.

[0109] In the no memory access, write mask control, partial round control type operation 810 instruction template, the rest of the beta field 854 is interpreted as a round operation field 859A 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).

[0110] Round operation control field 859A – just as round operation control field 858, 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 859A 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 850 content overrides that register value.

[0111] In the no memory access, write mask control, VSIZE type operation 817 instruction template, the rest of the beta field 854 is interpreted as a vector length field 859B, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

[0112] In the case of a memory access 820 instruction template of class B, part of the beta field 854 is interpreted as a broadcast field 857B, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field 854 is interpreted the vector length field 859B. The memory access 820 instruction templates include the scale field 860, and optionally the displacement field 862A or the displacement scale field 862B.

[0113] With regard to the generic vector friendly instruction format 800, a full opcode field 874 is shown including the format field 840, the base operation field 842, and the data element width field 864. While one embodiment is shown where the full opcode field 874 includes all of these fields, the full opcode field 874 includes less than all of these fields in embodiments that do not support all of them. The full opcode field 874 provides the operation code (opcode).

[0114] The augmentation operation field 850, the data element width field 864, and the write mask field 870 allow these features to be specified on a per instruction basis in the generic vector friendly instruction format.

[0115] 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.

[0116] The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments of the invention, different processors or different cores within a processor may support only class A, only class B, or both classes. For instance, a high performance general purpose out-of-order core intended for general-purpose computing may support only class B, a core intended primarily for graphics and/or scientific (throughput) computing may support only class A, and a core intended for both may support both (of course, a core that has some mix of templates and instructions from both classes but not all templates and instructions from both classes is within the purview of the invention). Also, a single processor may include multiple cores, all of which support the same class or in which different cores support different class. For instance, in a processor with separate graphics and general purpose cores, one of the graphics cores intended primarily for graphics

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

Exemplary Specific Vector Friendly Instruction Format

[0117] **Figures 9A-D** illustrate an exemplary specific vector friendly instruction format according to embodiments of the invention. These show a specific vector friendly instruction format 900 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 900 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 8** into which the fields from **Figures 9A-D** map are illustrated.

[0118] It should be understood that, although embodiments of the invention are described with reference to the specific vector friendly instruction format 900 in the context of the generic vector friendly instruction format 800 for illustrative purposes, the invention is not limited to the specific vector friendly instruction format 900 except where claimed. For example, the generic vector friendly instruction format 800 contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format 900 is shown as having fields of specific sizes. By way of specific example, while the data element width field 864 is illustrated as a one bit field in the specific vector friendly instruction format 900, the invention is not so limited (that is, the generic vector friendly instruction format 800 contemplates other sizes of the data element width field 864).

[0119] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in **Figure 9A**.

[0120] EVEX Prefix (Bytes 0-3) 902 - is encoded in a four-byte form.

[0121] Format Field 840 (EVEX Byte 0, bits [7:0]) - the first byte (EVEX Byte 0) is the format field 840 and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment of the invention).

[0122] The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

[0123] REX field 905 (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 857BEX byte 1, bit[5] – B). The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the corresponding VEX bit fields, and are encoded using 1s 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.

[0124] REX' field 810 – this is the first part of the REX' field 810 and is the EVEX.R' bit field (EVEX Byte 1, bit [4] - R') that is used to encode either the upper 16 or lower 16 of the extended 32 register set. In one embodiment of the invention, this bit, along with others as indicated below, is stored in bit inverted format to distinguish (in the well-known x86 32-bit mode) from the BOUND instruction, whose real opcode byte is 62, but does not accept in the MOD R/M field (described below) the value of 11 in the MOD field; alternative embodiments of the invention do not store this and the other indicated bits below in the inverted format. A value of 1 is used to encode the lower 16 registers. In other words, R'Rrrr is formed by combining EVEX.R', EVEX.R, and the other RRR from other fields.

[0125] Opcode map field 915 (EVEX byte 1, bits [3:0] – mmmm) – its content encodes an implied leading opcode byte (0F, 0F 38, or 0F 3).

[0126] Data element width field 864 (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).

[0127] EVEX.vvvv 920 (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 (1s complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in 1s 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 920 encodes the 4 low-order bits of the first source register specifier stored in inverted (1s complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

[0128] EVEX.U 868 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.

[0129] Prefix encoding field 925 (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 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 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.

[0130] Alpha field 852 (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 α) – as previously described, this field is context specific.

[0131] Beta field 854 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s2-0, EVEX.r2-0, EVEX.rr1, EVEX.LL0, EVEX.LLB; also illustrated with $\beta\beta\beta$) – as previously described, this field is context specific.

[0132] REX' field 810 – this is the remainder of the REX' field and is the EVEX.V' bit field (EVEX Byte 3, bit [3] - V') that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V', EVEX.vvvv.

[0133] Write mask field 870 (EVEX byte 3, bits [2:0]-kkk) – its content specifies the index of a register in the write mask registers as previously described. In one embodiment of the invention, the specific value EVEX.kkk=000 has a special behavior implying no write

mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

[0134] Real Opcode Field 930 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

[0135] MOD R/M Field 940 (Byte 5) includes MOD field 942, Reg field 944, and R/M field 946. As previously described, the MOD field's 942 content distinguishes between memory access and non-memory access operations. The role of Reg field 944 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 946 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.

[0136] Scale, Index, Base (SIB) Byte (Byte 6) 950 - As previously described, the scale field's 852 content is used for memory address generation. SIB.xxx 954 and SIB.bbb 956 – the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb.

[0137] Displacement field 862A (Bytes 7-10) – when MOD field 942 contains 10, bytes 7-10 are the displacement field 862A, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

[0138] Displacement factor field 862B (Byte 7) – when MOD field 942 contains 01, byte 7 is the displacement factor field 862B. 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 862B is a reinterpretation of disp8; when using displacement factor field 862B, 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 $\text{disp8} \times N$. 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

862B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 862B is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to disp8*N. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset). Immediate field 872 operates as previously described.

Full Opcode Field

[0139] **Figure 9B** is a block diagram illustrating the fields of the specific vector friendly instruction format 900 that make up the full opcode field 874 according to one embodiment of the invention. Specifically, the full opcode field 874 includes the format field 840, the base operation field 842, and the data element width (W) field 864. The base operation field 842 includes the prefix encoding field 925, the opcode map field 915, and the real opcode field 930.

Register Index Field

[0140] **Figure 9C** is a block diagram illustrating the fields of the specific vector friendly instruction format 900 that make up the register index field 844 according to one embodiment of the invention. Specifically, the register index field 844 includes the REX field 905, the REX' field 910, the MODR/M.reg field 944, the MODR/M.r/m field 946, the VVVV field 920, xxx field 954, and the bbb field 956.

Augmentation Operation Field

[0141] **Figure 9D** is a block diagram illustrating the fields of the specific vector friendly instruction format 900 that make up the augmentation operation field 850 according to one embodiment of the invention. When the class (U) field 868 contains 0, it signifies EVEX.U0 (class A 868A); when it contains 1, it signifies EVEX.U1 (class B 868B). When U=0 and the MOD field 942 contains 11 (signifying a no memory access operation), the alpha field 852 (EVEX byte 3, bit [7] – EH) is interpreted as the rs field 852A. When the rs field 852A contains a 1 (round 852A.1), the beta field 854 (EVEX byte 3, bits [6:4]– SSS) is interpreted as the round control field 854A. The round control field 854A includes a one bit SAE field 856 and a two bit round operation field 858. When the rs field 852A contains a 0 (data transform 852A.2), the beta field 854 (EVEX byte 3, bits [6:4]– SSS) is interpreted as a three

bit data transform field 854B. When $U=0$ and the MOD field 942 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 852 (EVEX byte 3, bit [7] – EH) is interpreted as the eviction hint (EH) field 852B and the beta field 854 (EVEX byte 3, bits [6:4]– SSS) is interpreted as a three bit data manipulation field 854C.

[0142] When $U=1$, the alpha field 852 (EVEX byte 3, bit [7] – EH) is interpreted as the write mask control (Z) field 852C. When $U=1$ and the MOD field 942 contains 11 (signifying a no memory access operation), part of the beta field 854 (EVEX byte 3, bit [4]– S_0) is interpreted as the RL field 857A; when it contains a 1 (round 857A.1) the rest of the beta field 854 (EVEX byte 3, bit [6-5]– S_{2-1}) is interpreted as the round operation field 859A, while when the RL field 857A contains a 0 (VSIZE 857.A2) the rest of the beta field 854 (EVEX byte 3, bit [6-5]– S_{2-1}) is interpreted as the vector length field 859B (EVEX byte 3, bit [6-5]– L_{1-0}). When $U=1$ and the MOD field 942 contains 00, 01, or 10 (signifying a memory access operation), the beta field 854 (EVEX byte 3, bits [6:4]– SSS) is interpreted as the vector length field 859B (EVEX byte 3, bit [6-5]– L_{1-0}) and the broadcast field 857B (EVEX byte 3, bit [4]– B).

Exemplary Register Architecture

[0143] **Figure 10** is a block diagram of a register architecture 1000 according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers 1010 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 900 operates on these overlaid register file as illustrated in the below tables.

Adjustable Vector Length	Class	Operations	Registers
Instruction Templates that do not include the vector length field 859B	A (Figure 8A; $U=0$)	810, 815, 825, 830	zmm registers (the vector length is 64 byte)
	B (Figure 8B; $U=1$)	812	zmm registers (the vector length is 64 byte)
Instruction templates that do include the vector length field 859B	B (Figure 8B; $U=1$)	817, 827	zmm, ymm, or xmm registers (the vector length is 64 byte, 32 byte, or

			16 byte) depending on the vector length field 859B
--	--	--	--

[0144] In other words, the vector length field 859B 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 859B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 900 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.

[0145] Write mask registers 1015 - 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 1015 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.

[0146] General-purpose registers 1025 - 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.

[0147] Scalar floating point stack register file (x87 stack) 1045, on which is aliased the MMX packed integer flat register file 1050 - 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.

[0148] 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.

Exemplary Core Architectures, Processors, and Computer Architectures

[0149] 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.

Exemplary Core Architectures

In-order and out-of-order core block diagram

[0150] **Figure 11A** 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 11B** 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 11A-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.

[0151] In **Figure 11A**, a processor pipeline 1100 includes a fetch stage 1102, a length decode stage 1104, a decode stage 1106, an allocation stage 1108, a renaming stage 1110, a scheduling (also known as a dispatch or issue) stage 1112, a register read/memory read stage

1114, an execute stage 1116, a write back/memory write stage 1118, an exception handling stage 1122, and a commit stage 1124.

[0152] **Figure 11B** shows processor core 1190 including a front end unit 1130 coupled to an execution engine unit 1150, and both are coupled to a memory unit 1170. The core 1190 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 1190 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.

[0153] The front end unit 1130 includes a branch prediction unit 1132 coupled to an instruction cache unit 1134, which is coupled to an instruction translation lookaside buffer (TLB) 1136, which is coupled to an instruction fetch unit 1138, which is coupled to a decode unit 1140. The decode unit 1140 (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 1140 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 1190 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 1140 or otherwise within the front end unit 1130). The decode unit 1140 is coupled to a rename/allocator unit 1152 in the execution engine unit 1150.

[0154] The execution engine unit 1150 includes the rename/allocator unit 1152 coupled to a retirement unit 1154 and a set of one or more scheduler unit(s) 1156. The scheduler unit(s) 1156 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 1156 is coupled to the physical register file(s) unit(s) 1158. Each of the physical register file(s) units 1158 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 1158 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) 1158 is overlapped by the retirement unit 1154

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 1154 and the physical register file(s) unit(s) 1158 are coupled to the execution cluster(s) 1160. The execution cluster(s) 1160 includes a set of one or more execution units 1162 and a set of one or more memory access units 1164. The execution units 1162 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) 1156, physical register file(s) unit(s) 1158, and execution cluster(s) 1160 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) 1164). 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.

[0155] The set of memory access units 1164 is coupled to the memory unit 1170, which includes a data TLB unit 1172 coupled to a data cache unit 1174 coupled to a level 2 (L2) cache unit 1176. In one exemplary embodiment, the memory access units 1164 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 1172 in the memory unit 1170. The instruction cache unit 1134 is further coupled to a level 2 (L2) cache unit 1176 in the memory unit 1170. The L2 cache unit 1176 is coupled to one or more other levels of cache and eventually to a main memory.

[0156] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 1100 as follows: 1) the instruction fetch 1138 performs the fetch and length decoding stages 1102 and 1104; 2) the decode unit 1140 performs the decode stage 1106; 3) the rename/allocator unit 1152 performs the allocation stage 1108 and renaming stage 1110; 4) the scheduler unit(s) 1156 performs the schedule stage 1112; 5) the physical register file(s) unit(s) 1158 and the memory unit 1170 perform the

register read/memory read stage 1114; the execution cluster 1160 perform the execute stage 1116; 6) the memory unit 1170 and the physical register file(s) unit(s) 1158 perform the write back/memory write stage 1118; 7) various units may be involved in the exception handling stage 1122; and 8) the retirement unit 1154 and the physical register file(s) unit(s) 1158 perform the commit stage 1124.

[0157] The core 1190 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 1190 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[0158] 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).

[0159] 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 1134/1174 and a shared L2 cache unit 1176, 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.

Specific Exemplary In-Order Core Architecture

[0160] **Figures 12A-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.

[0161] **Figure 12A** is a block diagram of a single processor core, along with its connection to the on-die interconnect network 1202 and with its local subset of the Level 2 (L2) cache 1204, according to embodiments of the invention. In one embodiment, an instruction decoder 1200 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 1206 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 1208 and a vector unit 1210 use separate register sets (respectively, scalar registers 1212 and vector registers 1214) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 1206, 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).

[0162] The local subset of the L2 cache 1204 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 1204. Data read by a processor core is stored in its L2 cache subset 1204 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 1204 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.

[0163] **Figure 12B** is an expanded view of part of the processor core in **Figure 12A** according to embodiments of the invention. **Figure 12B** includes an L1 data cache 1206A part of the L1 cache 1204, as well as more detail regarding the vector unit 1210 and the vector registers 1214. Specifically, the vector unit 1210 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 1228), 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 1220, numeric conversion with numeric convert units 1222A-B, and replication with replication unit 1224 on the memory input. Write mask registers 1226 allow predicating resulting vector writes.

[0164] **Figure 13** is a block diagram of a processor 1300 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention. The solid lined boxes in **Figure 13** illustrate a processor 1300 with a single core 1302A, a system agent 1310, a set of one or more bus controller units

1316, while the optional addition of the dashed lined boxes illustrates an alternative processor 1300 with multiple cores 1302A-N, a set of one or more integrated memory controller unit(s) 1314 in the system agent unit 1310, and special purpose logic 1308.

[0165] Thus, different implementations of the processor 1300 may include: 1) a CPU with the special purpose logic 1308 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 1302A-N being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of the two); 2) a coprocessor with the cores 1302A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 1302A-N being a large number of general purpose in-order cores. Thus, the processor 1300 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 1300 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.

[0166] The memory hierarchy includes one or more levels of cache 1304A-N within the cores, a set or one or more shared cache units 1306, and external memory (not shown) coupled to the set of integrated memory controller units 1314. The set of shared cache units 1306 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 1312 interconnects the integrated graphics logic 1308, the set of shared cache units 1306, and the system agent unit 1310/integrated memory controller unit(s) 1314, 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 1306 and cores 1302-A-N.

[0167] In some embodiments, one or more of the cores 1302A-N are capable of multi-threading. The system agent 1310 includes those components coordinating and operating cores 1302A-N. The system agent unit 1310 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 1302A-N and the integrated graphics logic 1308. The display unit is for driving one or more externally connected displays.

[0168] The cores 1302A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 1302A-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.

Exemplary Computer Architectures

[0169] **Figures 14-16** 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.

[0170] Referring now to **Figure 14**, shown is a block diagram of a system 1400 in accordance with one embodiment of the present invention. The system 1400 may include one or more processors 1410, 1415, which are coupled to a controller hub 1420. In one embodiment the controller hub 1420 includes a graphics memory controller hub (GMCH) 1490 and an Input/Output Hub (IOH) 1450 (which may be on separate chips); the GMCH 1490 includes memory and graphics controllers to which are coupled memory 1440 and a coprocessor 1445; the IOH 1450 is couples input/output (I/O) devices 1460 to the GMCH 1490. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory 1440 and the coprocessor 1445 are coupled directly to the processor 1410, and the controller hub 1420 in a single chip with the IOH 1450.

[0171] The optional nature of additional processors 1415 is denoted in **Figure 14** with broken lines. Each processor 1410, 1415 may include one or more of the processing cores described herein and may be some version of the processor 1300.

[0172] The memory 1440 may be, for example, dynamic random access memory (DRAM), phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub 1420 communicates with the processor(s) 1410, 1415 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection 1495.

[0173] In one embodiment, the coprocessor 1445 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like. In one embodiment, controller hub 1420 may include an integrated graphics accelerator.

[0174] There can be a variety of differences between the physical resources 1410, 1415 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

[0175] In one embodiment, the processor 1410 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 1410 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 1445. Accordingly, the processor 1410 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 1445. Coprocessor(s) 1445 accept and execute the received coprocessor instructions.

[0176] Referring now to **Figure 15**, shown is a block diagram of a first more specific exemplary system 1500 in accordance with an embodiment of the present invention. As shown in **Figure 15**, multiprocessor system 1500 is a point-to-point interconnect system, and includes a first processor 1570 and a second processor 1580 coupled via a point-to-point interconnect 1550. Each of processors 1570 and 1580 may be some version of the processor 1300. In one embodiment of the invention, processors 1570 and 1580 are respectively processors 1410 and 1415, while coprocessor 1538 is coprocessor 1445. In another embodiment, processors 1570 and 1580 are respectively processor 1410 coprocessor 1445.

[0177] Processors 1570 and 1580 are shown including integrated memory controller (IMC) units 1572 and 1582, respectively. Processor 1570 also includes as part of its bus controller units point-to-point (P-P) interfaces 1576 and 1578; similarly, second processor 1580 includes P-P interfaces 1586 and 1588. Processors 1570, 1580 may exchange information via a point-to-point (P-P) interface 1550 using P-P interface circuits 1578, 1588. As shown in **Figure 15**, IMCs 1572 and 1582 couple the processors to respective memories, namely a memory 1532 and a memory 1534, which may be portions of main memory locally attached to the respective processors.

[0178] Processors 1570, 1580 may each exchange information with a chipset 1590 via individual P-P interfaces 1552, 1554 using point to point interface circuits 1576, 1594, 1586, 1598. Chipset 1590 may optionally exchange information with the coprocessor 1538 via a high-performance interface 1592. In one embodiment, the coprocessor 1538 is a special-

purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

[0179] 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.

[0180] Chipset 1590 may be coupled to a first bus 1516 via an interface 1596. In one embodiment, first bus 1516 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

[0181] As shown in **Figure 15**, various I/O devices 1514 may be coupled to first bus 1516, along with a bus bridge 1518 which couples first bus 1516 to a second bus 1520. In one embodiment, one or more additional processor(s) 1515, 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 1516. In one embodiment, second bus 1520 may be a low pin count (LPC) bus. Various devices may be coupled to a second bus 1520 including, for example, a keyboard and/or mouse 1522, communication devices 1527 and a storage unit 1528 such as a disk drive or other mass storage device which may include instructions/code and data 1530, in one embodiment. Further, an audio I/O 1524 may be coupled to the second bus 1520. Note that other architectures are possible. For example, instead of the point-to-point architecture of **Figure 15**, a system may implement a multi-drop bus or other such architecture.

[0182] Referring now to **Figure 16**, shown is a block diagram of a second more specific exemplary system 1600 in accordance with an embodiment of the present invention. Like elements in **Figures 15 and 16** bear like reference numerals, and certain aspects of **Figure 15** have been omitted from **Figure 16** in order to avoid obscuring other aspects of **Figure 16**.

[0183] **Figure 16** illustrates that the processors 1570, 1580 may include integrated memory and I/O control logic ("CL") 1572 and 1582, respectively. Thus, the CL 1572, 1582 include integrated memory controller units and include I/O control logic. **Figure 16** illustrates that not only are the memories 1532, 1534 coupled to the CL 1572, 1582, but also that I/O devices 1614 are also coupled to the control logic 1572, 1582. Legacy I/O devices 1615 are coupled to the chipset 1590.

[0184] Referring now to **Figure 16**, shown is a block diagram of a SoC 1600 in accordance with an embodiment of the present invention. Similar elements in **Figure 13** bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In **Figure 16**, an interconnect unit(s) 1602 is coupled to: an application processor 1610 which includes a set of one or more cores 1302A-N, cache 1304A-N, and shared cache unit(s) 1306; a system agent unit 1310; a bus controller unit(s) 1316; an integrated memory controller unit(s) 1314; a set or one or more coprocessors 1620 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 1630; a direct memory access (DMA) unit 1632; and a display unit 1640 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 1620 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.

[0185] 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.

[0186] Program code, such as code 1530 illustrated in **Figure 15**, 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.

[0187] 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.

[0188] 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.

[0189] 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.

[0190] 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.

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

[0191] 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.

[0192] **Figure 17** 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 17** shows a program in a high level language 1702 may be compiled using an x86 compiler 1704 to generate x86 binary code 1706 that may be

natively executed by a processor with at least one x86 instruction set core 1716. The processor with at least one x86 instruction set core 1716 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 1704 represents a compiler that is operable to generate x86 binary code 1706 (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 1716. Similarly, **Figure 17** shows the program in the high level language 1702 may be compiled using an alternative instruction set compiler 1708 to generate alternative instruction set binary code 1710 that may be natively executed by a processor without at least one x86 instruction set core 1714 (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 1712 is used to convert the x86 binary code 1706 into code that may be natively executed by the processor without an x86 instruction set core 1714. This converted code is not likely to be the same as the alternative instruction set binary code 1710 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 1712 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 1706.

CLAIMS

What is claimed is:

1. An apparatus comprising:
a decoder means to decode an instruction, wherein the instruction to include at least an opcode, a field for at least two packed data source operands, a field for a packed data destination operand, and an immediate; and
execution means to execute the decoded instruction to load packed data elements from the at least two packed data source operands using a stride and storing results of the strided loads in the packed data destination operand starting from a defined position determined in part from the immediate.
2. The apparatus of claim 1, wherein the packed data source operands are concatenated.
3. The apparatus of any of claims 1-2, wherein the defined position is determined by a round value provided by the immediate, a number of packed data elements in each of the packed data source operands, an offset provided by the immediate, the number of packed data source operands, and the stride determined from the immediate.
4. The apparatus of claim 3, wherein the stride is a value from the immediate plus 1.
5. The apparatus of claim 3, wherein the defined position is the round multiplied by the number of packed data elements in each of the packed data source operand multiplied by the number of packed data source operands plus an offset, wherein this value is divided by the stride.
6. The apparatus of any of claims 1-5, wherein the instruction to include a field for a writemask operand.
7. The apparatus of claim 6, wherein the execution means to store a result of the strided loads based on values of the writemask operand.
8. A method comprising:

decoding an instruction, wherein the instruction to include at least an opcode, a field for at least two packed data source operands, a field for a packed data destination operand, and an immediate; and

executing the decoded instruction to load packed data elements from the at least two packed data source operands using a stride and storing results of the strided loads in the packed data destination operand starting from a defined position determined in part from the immediate.

9. The method of claim 8, wherein the packed data source operands are concatenated.
10. The method of any of claims 8-9, wherein the defined position is determined by a round value provided by the immediate, a number of packed data elements in each of the packed data source operands, an offset provided by the immediate, the number of packed data source operands, and the stride determined from the immediate.
11. The method of claim 9, wherein the stride is a value from the immediate plus 1.
12. The method of claim 9, wherein the defined position is the round multiplied by the number of packed data elements in each of the packed data source operand multiplied by the number of packed data source operands plus an offset, wherein this value is divided by the stride.
13. The method of any of claims 8-13, wherein the instruction to include a field for a writemask operand.
14. The method of claim 13, wherein the storing is based on values of the writemask operand.
15. A non-transitory machine-readable medium storing instructions which when executed to cause a processor to perform a method, the method comprising:

decoding an instruction, wherein the instruction to include at least an opcode, a field for at least two packed data source operands, a field for a packed data destination operand, and an immediate; and

executing the decoded instruction to load packed data elements from the at least two packed data source operands using a stride and storing results of the strided loads in the packed data destination operand starting from a defined position determined in part from the immediate.

16. The non-transitory machine-readable medium of claim 15, wherein the packed data source operands are concatenated.

17. The non-transitory machine-readable medium of claim 16, wherein the defined position is determined by a round value provided by the immediate, a number of packed data elements in each of the packed data source operands, an offset provided by the immediate, the number of packed data source operands, and the stride determined from the immediate.

18. The non-transitory machine-readable medium of claim 16, wherein the stride is a value from the immediate plus 1.

19. The non-transitory machine-readable medium of claim 16, wherein the defined position is the round multiplied by the number of packed data elements in each of the packed data source operand multiplied by the number of packed data source operands plus an offset, wherein this value is divided by the stride.

20. The non-transitory machine-readable medium of claim 15, wherein the instruction to include a field for a writemask operand and the storing is based on values of the writemask operand.

21. An apparatus comprising:

a decoder circuit to decode an instruction, wherein the instruction to include at least an opcode, a field for at least two packed data source operands, a field for a packed data destination operand, and an immediate; and

execution circuitry to execute the decoded instruction to load packed data elements from the at least two packed data source operands using a stride and storing results of the strided loads in the packed data destination operand starting from a defined position determined in part from the immediate.

22. The apparatus of claim 21, wherein the packed data source operands are concatenated.
23. The apparatus of claim 22, wherein the defined position is determined by a round value provided by the immediate, a number of packed data elements in each of the packed data source operands, an offset provided by the immediate, the number of packed data source operands, and the stride determined from the immediate.
24. The apparatus of claim 23, wherein the stride is a value from the immediate plus 1.
25. The apparatus of claim 23, wherein the defined position is the round multiplied by the number of packed data elements in each of the packed data source operand multiplied by the number of packed data source operands plus an offset, wherein this value is divided by the stride.




FIG. 1




FIG. 2




FIG. 3

OPCODE 401	DESTINATION 403	SOURCE(S) 405	IMMEDIATE 409	WITEMASK 40Z
STRIDELOAD{DE SIZE}	DST	SRC(S)	STRIDE[7:5] OFFSET[4:2] ROUND[1:0]	K

FIG. 4




FIG. 5




FIG. 6


```


STRIDELOAD[B,W,D,Q] dest{k1}, src1, src2, ..., srcM, imm8
    KL = VL/datasize //number of elements in a register
    stride = imm[7:5] + 1; //strides 1,2, ..., 8
    offset = imm[4:2]; //offsets 0, ..., 7
    round = imm[1:0]; //round 0, ..., 3
    src_tmp = concatenate(srcM,...,src2,src1);
    i0 = (round*M*KL + offset)/stride //i0 is a start position in a destination
    for(i=i0, j=offset; i< KL; i++, j+=stride){
        //do load of elements in the
        // concatenated source (M*KL) with a given stride
        if (j>= M*KL) break; //concatenated source finished earlier than dest.
        //Optionally raise CF flag about early out
        if(k1[i] || no_mask){
            dest[i] = src_tmp[j];
        }
    }



    //optionally zero upper elements to remove dependencies
    dest[KL-1: i] = 0;

```


FIG. 7







U.S. GPO







FIG. 10

REGISTER ARCHITECTURE 1000







FIG. 12A






FIG. 12B








FIG. 14




FIG. 15




FIG. 16




FIG. 17

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/037553**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/345(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 9/345; G06F 9/30; G06F 9/32

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: packed data, source, destination, immediate, stride, position

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2016-0179523 A1 (INTEL CORPORATION) 23 June 2016 See paragraphs [0062]-[0075]; claim 1; and figure 8.	1-25
Y	US 2013-0275727 A1 (SETH ABRAHAM et al.) 17 October 2013 See paragraphs [0086], [0089], [0123]-[0124]; claim 1; and figures 7B, 13.	1-25
Y	US 2015-0052333 A1 (CHRISTOPHER J. HUGHES et al.) 19 February 2015 See paragraphs [0049]-[0053]; claim 1; and figure 4.	6-7, 13-14, 20
A	US 2014-0195775 A1 (ELMOUSTAPHA OULD-AHMED-VALL et al.) 10 July 2014 See paragraphs [0080]-[0093]; claim 1; and figures 3D-3H.	1-25
A	US 5940876 A (JAMES K. PICKETT) 17 August 1999 See column 4, line 41 - column 11, line 47; column 13, lines 17-30; claim 1; and figures 2-3, 7.	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

12 September 2017 (12.09.2017)

Date of mailing of the international search report

12 September 2017 (12.09.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

Lee, Eun Kyu

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/037553

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016-0179523 A1	23/06/2016	TW 201636831 A WO 2016-105727 A1	16/10/2016 30/06/2016
US 2013-0275727 A1	17/10/2013	CN 104011645 A TW 201329861 A TW I511043 B WO 2013-095580 A1	27/08/2014 16/07/2013 01/12/2015 27/06/2013
US 2015-0052333 A1	19/02/2015	CN 103562856 A CN 103562856 B DE 112011105121 T5 GB 2503169 A JP 2014-513340 A JP 2016-040737 A JP 5844882 B2 KR 10-1607161 B1 KR 10-2013-0137702 A TW 201246065 A TW 201525856 A TW I476684 B TW I514273 B US 2012-0254591 A1 WO 2012-134555 A1	05/02/2014 16/11/2016 09/01/2014 18/12/2013 29/05/2014 24/03/2016 20/01/2016 29/03/2016 17/12/2013 16/11/2012 01/07/2015 11/03/2015 21/12/2015 04/10/2012 04/10/2012
US 2014-0195775 A1	10/07/2014	BR112014004603 A2 CN103827815 A DE 112011105665 T5 GB 2507018 A JP 2014-53042 6A JP 5930558 B2 KR 10-1714133 B1 KR 10-2014-0054371 A US 9672036 B2 WO 2013-048367 A1 WO 2013-048367 A9	13/06/2017 28/05/2014 21/08/2014 16/04/2014 17/11/2014 08/06/2016 08/03/2017 08/05/2014 06/06/2017 04/04/2013 04/04/2013
US 05940876 A	17/08/1999	US 6088781 A	11/07/2000