



- (51) **International Patent Classification:**  
*G06F 9/06* (2006.01)      *G06F 9/30* (2006.01)
- (21) **International Application Number:**  
PCT/US201 1/067071
- (22) **International Filing Date:**  
23 December 201 1 (23. 12.201 1)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant (for all designated States except US):** **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, MS: RNB-4-150, Santa Clara, California 95052 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **OULD-AHMED-VALL, Elmoustapha** [MR/US]; 5000 W Chandler Blvd, M/S: CH7-401, Chandler, Arizona 85226 (US). **HAGOG, Mostafa** [IL/IL]; Kaukab 323, 20185 Kaukab, HA (IL). **VALENTINE, Robert** [US/IL]; Rechov Hadganiot 33-5, 36054 Kiryat Tivon, HA (IL). **GRADSTEIN, Amit** [IL/IL]; 16th Hadas St., 30500 Binyamina, HA (IL). **RUBANOVICH, Simon** [IL/IL]; Egoz str. 10/4, 34792 Haifa, HA (IL). **SPERBER, Zeev** [IL/IL]; 32nd Igal Alon St., 30900 Zichron Yackov (IL).
- (74) **Agents:** **NICHOLSON, David F.** et al; Blakely Sokoloff Taylor & Zaffian LLP, 1279 Oakmead Parkway, Sunnyvale, California 94085 (US).

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

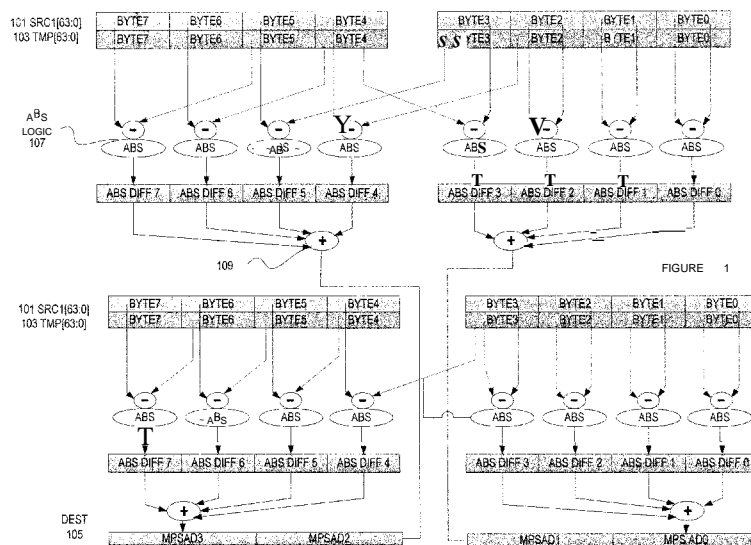
**Declarations under Rule 4.17 :**

— of inventorship (Rule 4.17(ivf))

**Published:**

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

(54) **Title:** SYSTEMS, APPARATUSES, AND METHODS FOR PERFORMING A DOUBLE BLOCKED SUM OF ABSOLUTE DIFFERENCES



(57) **Abstract:** Embodiments of systems, apparatuses, and methods for performing in a computer processor vector double block packed sum of absolute differences (SAD) in response to a single vector double block packed sum of absolute differences instruction that includes a destination vector register operand, first and second source operands, an immediate, and an opcode are described.

WO 2013/095599 A1

**SYSTEMS, APPARATUSES, AND METHODS FOR PERFORMING A DOUBLE  
BLOCKED SUM OF ABSOLUTE DIFFERENCES**

5

FIELD OF INVENTION

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

10 

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). The term instruction generally refers herein to macro-

15 instructions - that is instructions that are provided to the processor (or instruction converter that translates (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morphs, emulates, or otherwise converts an instruction to one or more other instructions to be processed by the processor) for execution - as opposed to micro-instructions or micro-operations (micro-ops) - that is the result of a processor's decoder decoding macro-  
20 instructions.

[0003] The ISA is distinguished from the microarchitecture, which is the internal design of the processor implementing the instruction set. Processors with different microarchitectures can share a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale CA implement  
25 nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. For example, the same register architecture of the ISA may be implemented in different ways in different microarchitectures using well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register  
30 Alias Table (RAT), a Reorder Buffer (ROB), and a retirement register file; the use of multiple maps and a pool of registers), etc. Unless otherwise specified, the phrases register architecture, register file, and register are used herein to refer to that which is visible to the software/programmer and the manner in which instructions specify registers. Where a

specificity is desired, the adjective logical, architectural, or software visible will be used to indicate registers/files in the register architecture, while different adjectives will be used to designation registers in a given microarchitecture (e.g., physical register, reorder buffer, retirement register, register pool).

5 **[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 (opcode) and the operand(s) on which that operation is to be performed. Some instruction formats are further broken down through the definition of instruction templates (or subformats). For example, the instruction templates of a given  
10 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  
15 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.

20 **[0005]** Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) often require the same operation to be performed on a large number of data items (referred to as "data parallelism"). Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes  
25 a processor to perform an operation on multiple data items. SIMD technology is especially suited to processors that can logically divide the bits in a register into a number of fixed-sized data elements, each of which represents a separate value. For example, the bits in a 256-bit register may be specified as a source operand to be operated on as four separate 64-bit packed data elements (quad-word (Q) size data elements), eight separate 32-bit packed data elements  
30 (double word (D) size data elements), sixteen separate 16-bit packed data elements (word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). This type of data is referred to as packed data type or vector data type, and operands of this data type are referred to as packed data operands or vector operands. In other words, a packed data item or

vector refers to a sequence of packed data elements, and a packed data operand or a vector operand is a source or destination operand of a SIMD instruction (also known as a packed data instruction or a vector instruction).

[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 has only one or has more than two source vector operands, that operate in a horizontal fashion, that generates a result vector operand that is of a different size, that has a different size data elements, and/or that has a different data element order). It should be understood that the term destination vector operand (or destination operand) is defined as the direct result of performing the operation specified by an instruction, including the storage of that destination operand at a location (be it a register or at a memory address specified by that instruction) so that it may be accessed as a source operand by another instruction (by specification of that same location by the another instruction).

[0007] The SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, has enabled a significant improvement in application

performance. An additional set of SIMD extensions, referred to the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme, has been , has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developers Manual, October 201 1; and see Intel® Advanced Vector Extensions Programming Reference, June 201 1).

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

[0008] 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:

10 [0009] **Figure 1** illustrates an exemplary illustration of part of an operation of an exemplary DBPSAD instruction operating on byte elements wherein the resulting SAD calculations are stored as word elements.

[0010] **Figure 2** shows the next 64-bit lane calculations. If source and destinations were 128-bit, then this would be all of the calculations performed.

15 [0011] **Figure 3** illustrates more detailed exemplary instruction formats.

[0012] **Figure 4** illustrates an embodiment of the use of a DBPSAD instruction in a processor.

[0013] **Figure 5** illustrates an embodiment of a method for processing a DBPSAD instruction.

[0014] **Figure 6** illustrates pseudo-code for DBPSAD for a byte source data element size and word destination data element size.

20 [0015] **Figure 7** illustrates a correlation between the number of one active bit vector writemask elements and the vector size and the data element size according to one embodiment of the invention.

[0016] **Figures 8A-8B** are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention.

25 [0017] **Figure 9** is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention

[0018] **Figure 10** is a block diagram of a register architecture according to one embodiment of the invention

30 [0019] **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.

[0020] **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.

[0021] **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

[0022] **Figure 13** is a block diagram of a processor that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention.

[0023] **Figure 14** is a block diagram of a system in accordance with one embodiment of the present invention.

[0024] **Figure 15**, is a block diagram of a first more specific exemplary system in accordance with an embodiment of the present invention.

[0025] **Figure 16** is a block diagram of a second more specific exemplary system in accordance with an embodiment of the present invention.

[0026] **Figure 17** is a block diagram of a SoC in accordance with an embodiment of the present invention.

[0027] **Figure 18** 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

[0028] 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.

[0029] 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.

**[0030]** Overview

**[0031]** In the description below, there are some items that may need explanation prior to describing the operations of this particular instruction in the instruction set architecture. One such item is called a "writemask register" which is generally used to predicate an operand to conditionally control per-element computational operation (below, the term mask register may also be used and it refers to a writemask register such as the "k" registers discussed below). As used below, a writemask register stores a plurality of bits (16, 32, 64, etc.) wherein each active bit of the writemask register governs the operation/update of a packed data element of a vector register during SIMD processing. Typically, there is more than one writemask register available for use by a processor core.

**[0032]** The instruction set architecture includes at least some SIMD instructions that specify vector operations and that have fields to select source registers and/or destination registers from these vector registers (an exemplary SIMD instruction may specify a vector operation to be performed on the contents of one or more of the vector registers, and the result of that vector operation to be stored in one of the vector registers). Different embodiments of the invention may have different sized vector registers and support more/less/different sized data elements.

**[0033]** The size of the multi-bit data elements specified by a SIMD instruction (e.g., byte, word, double word, quad word) determines the bit locations of the "data element positions" within a vector register, and the size of the vector operand determines the number of data elements. A packed data element refers to the data stored in a particular position. In other words, depending on the size of the data elements in the destination operand and the size of the destination operand (the total number of bits in the destination operand) (or put another way, depending on the size of the destination operand and the number of data elements within the destination operand), the bit locations of the multi-bit data element positions within the resulting vector operand change (e.g., if the destination for the resulting vector operand is a vector register, then the bit locations of the multi-bit data element positions within the destination vector register change). For example, the bit locations of the multi-bit data elements are different between a vector operation that operates on 32-bit data elements (data element position 0 occupies bit locations 31:0, data element position 1 occupies bit locations 63:32, and so on) and a vector operation that operates on 64-bit data elements (data element position 0 occupies bit locations 63:0, data element position 1 occupies bit locations 127:64, and so on).

**[0034]** Additionally, there is a correlation between the number of one active bit vector writemask elements and the vector size and the data element size according to one embodiment

of the invention as shown in Figure 7. Vector sizes of 128-bits, 256-bits, and 512-bits are shown, although other widths are also possible. Data element sizes of 8-bit bytes (B), 16-bit words (W), 32-bit doublewords (D) or single precision floating point, and 64-bit quadwords (Q) or double precision floating point are considered, although other widths are also possible. As shown, when the vector size is 128-bits, 16-bits may be used for masking when the vector's data element size is 8-bits, 8-bits may be used for masking when the vector's data element size is 16-bits, 4-bits may be used for masking when the vector's data element size is 32-bits, and 2-bits may be used for masking when the vector's data element size is 64-bits. When the vector size is 256-bits, 32-bits may be used for masking when the packed data element width is 8-bits, 16-bits may be used for masking when the vector's data element size is 16-bits, 8-bits may be used for masking when the vector's data element size is 32-bits, and 4-bits may be used for masking when the vector's data element size is 64-bits. When the vector size is 512-bits, 64-bits may be used for masking when the vector's data element size is 8-bits, 32-bits may be used for masking when the vector's data element size is 16-bits, 16-bits may be used for masking when the vector's data element size is 32-bits, and 8-bits may be used for masking when the vector's data element size is 64-bits.

**[0035]** Depending upon the combination of the vector size and the data element size, either all 64-bits, or only a subset of the 64-bits, may be used as a write mask. Generally, when a single, per-element masking control bit is used, the number of bits in the vector writemask register used for masking (active bits) is equal to the vector size in bits divided by the vector's data element size in bits.

**[0036]** Described below is a vector instruction that computes multiple SAD (sum of absolute differences) of selected quadruplets of byte elements in a first operand vector compared to quadruplet of byte elements of a second operand vector. It also provides the flexibility to select which quadruplets to compare. This instruction provides the efficiency of comparing SAD of blocks on video images used for motion search.

**[0037]** Below are embodiments of an instruction generically called a double block packed sum of absolute differences (SAD) ("DBPSAD") instruction and embodiments of systems, architectures, instruction formats, etc. that may be used to execute such an instruction that is beneficial in several different areas. The execution of a DBPSAD instruction causes the storage of SAD on byte elements from selected quadruplets (four data element chunks) from a first and second source into word packed data elements of a destination register. Each quadruplet is the same size as a data element of the destination. More particularly, in some embodiments multiple



SAD calculations are performed on different 64-bit chunks of data from the two sources and stored as 4 word sized (64 bits in total) results in the destination register.

[0038] Figure 1 illustrates an exemplary illustration of part of an operation of an exemplary DBPSAD instruction operating on byte elements wherein the resulting SAD calculations are stored as word elements. From the first source operand the instruction takes two different quadruplets with offset of four bytes to participate in the four different SADs. In this illustration a 64-bit lane is having four SAD calculations performed between 64 selected bits of the first source and selected 64 bits from the second source. The process illustrate in this figure and described herein would be performed (with some slight modifications such as different quadruplet source selections and destination locations) for n number of times where n is the size of the vector (such as the first source) divided by 64.

[0039] In this particular example, the 64-bits of the first source 101 are the least significant 64 bits of the source. These 64-bits are broken down into 8 byte elements labeled byte0 to byte7.

[0040] A temporary 64-bit value 103 selected from the second source is also shown. How these bits are selected is detailed below. As such, there are two quadruplets (32 bit values broken comprising 4 bytes each) for both sources that have SAD calculations performed on them. The selection of quadruplets in the second source is done within 128-bit lanes of the vector. For a 128-bit second source there is obviously only one such lane, but there are two in a 256-bit vector and four in a 512-bit vector.

[0041] The illustration shows several independent absolute difference logic 107 and summation logic 109. This logic may be hardware such as an ALU or software routines running on an ALU. Additionally, while the figure illustrates separate absolute difference logic 107 and summation logic 109 to be used per SAD calculation, this logic may be combined into one unit per category (i.e., one absolute difference logic 107 and one summation logic 109) or into a single unit (i.e., a joint absolute difference logic 107 and summation logic 109).

[0042] The results from the 4 SAD operations are stored as word elements in the destination register 105. The specifics of each SAD calculation is detailed below. The calculations for the first 64-bit lane are as follows, where 1 = 0, SRC1 = the first 64 bits of the first source, and TMP1 is a selection of quadruplets from the second source that have been selected according to an immediate value (these are the lower two quadruplets).

$$\begin{aligned}
 \text{DEST}[\text{I}+15:\text{I}] <- \text{ABS}(\text{SRC1}[\text{I}+7: \text{I}] - \text{TMP1}[\text{I}+7: \text{I}]) + \text{ABS}(\text{SRC1}[\text{I}+15: \text{I}+8] - \\
 \text{TMP1}[\text{I}+15: \text{I}+8]) + \text{ABS}(\text{SRC1}[\text{I}+23: \text{I}+16] - \text{TMP1}[\text{I}+23: \text{I}+16]) + \text{ABS}(\text{SRC1}[\text{I}+31: \\
 \text{I}+24] - \text{TMP1}[\text{I}+31: \text{I}+24])
 \end{aligned}$$

- DEST[I+31: 1+16] <- (SRC 1[I+7: 1] - TMP1[I+15: 1+8]) + ABS(SRC1[I+15: 1+8]-  
TMP1[I+23: 1+16]) + ABS(SRC1[I+23: 1+16]- TMP1[I+31: 1+24]) + ABS(SRC1[I+31:  
1+24]- TMP1[I+39: 1+32])
- DEST[I+47: 1+32] <- ABS(SRC1[I+39: 1+32] - TMP1[I+23: 1+16]) +  
5 ABS(SRC1[I+47: 1+40]- TMP1[I+31: 1+24]) + ABS(SRC1[I+55: 1+48]- TMP1[I+39:  
1+32]) + ABS(SRC1[I+63: 1+56]- TMP 1[I+47: 1+40])
- DEST[I+63: 1+48] <- ABS(SRC1[I+39: 1+32] - TMP1[I+31: 1+24]) +  
ABS(SRC1[I+47: 1+40] - TMP 1[I+39: 1+32]) + ABS(SRC1[I+55: 1+48] - TMP1[I+47:  
1+40]) + ABS(SRC1[I+63: 1+56] - TMP 1[I+55: 1+48])

10 **[0043]** Generically, the what is stored into the destination data element positions is as follows wherein the value in [] is the data element position:

- DEST[0] <- ABS(SRC1[0] - TMP1[0]) + ABS(SRC1[1]- TMP1[1]) + ABS(SRC1[2]-  
TMP1[2]) + ABS(SRC1[3]- TMP1[3])
- DEST[1] <- (SRC1[0] - TMP1[1]) + ABS(SRC1[1]- TMP1[2]) + ABS(SRC1[2]-  
15 TMP1[3]) + ABS(SRC1[3]- TMP1[4])
- DEST[2] <- ABS(SRC1[4] - TMP1[2]) + ABS(SRC1[5]- TMP1[3]) + ABS(SRC1[6]-  
TMP1[4]) + ABS(SRC1[7]- TMP1[5])
- DEST[3] <- ABS(SRC1[4] - TMP1[3]) + ABS(SRC1[5] - TMP1[4]) + ABS(SRC1[6] -  
TMP1[5]) + ABS(SRC1[7] - TMP1[6])

20 **[0044]** Obvious variations to the above are made to do the next 64-bit lane. For example, the first source may be from bit 64 to bit 127 (the next 64-bit chunk) and the temporary value would be the upper two quadrants selected from the second source.

**[0045]** The selection of quadruplets from the second source is as follows:

- TMP1[I+31:I] <- select (SRC2[I+127: 1], imm8[1:0])
- 25 TMP1[I+63: 1+32] <- select (SRC2[I+127: 1], imm8[3:2])
- TMP1[I+95: 1+64] <- select (SRC2[I+127: 1], imm8[5:4])
- TMP1 [I+127: 1+96]<- select (SRC2[I+127: 1], imm8[7:6])

**[0046]** The two bits of the immediate allow for the selection of four 32-bit data elements from the second source. For example, for TMP1, if the immediate is "01" then bits 63:32 of the SRC2  
30 are selected.

**[0047]** The above is repeated, for example, 64-bits at a time, per data lane (128-bit) until the entire vector length has been processed.

[0048] Figure 2 shows the next 64-bit lane calculations. If source and destinations were 128-bit, then this would be all of the calculations performed.

[0049] Exemplary Format

[0050] An exemplary format of this instruction is "DBPSAD{B/W/D/Q} {BAV/D/Q}

5 XMM1/YMM1/ZMM1, XMM2/YMM2/ZMM2/m 128/m256/m512, imm8" where the operand XMM1/YMM1/ZMM1 is a source vector register (such as a 128-, 256-, or 512-bit register) and the destination XMM1/YMM1/ZMM1 which is a vector register (such as a 128-, 256-, or 512-bit register), or a memory location of 128-, 256-, or 512-bit size, imm8 is an 8-bit immediate (although other immediate sizes may be used), and DBPSAD{B/W/D/Q} is the instruction's  
 10 opcode. The size of the data elements in the source register may be defined in the "prefix" of the instruction such as through the use of an indication of data granularity bit. In most embodiments, this bit will indicate that each data element is either 32 or 64 bits, however, other variations may be used. In other embodiments, the size of the data elements is defined by the opcode itself. For example, a first {B/W/D/Q} identifier indicates a byte, word, double word, or  
 15 quadword source data element size respectively. Additionally, in some embodiments, a follow-on {B/W/D/Q} identifier indicates a byte, word, doubleword, or quadword destination data element size respectively. For example, DBPSADBW indicates byte source data elements and word destination data elements.

[0051] Figure 3 illustrates more detailed exemplary vector friendly instruction formats.

20 [0052] Exemplary Methods of Execution

[0053] Figure 4 illustrates an embodiment of the use of a DBPSAD instruction in a processor. A DBPSAD instruction with a first and second source operands, a destination operand, an immediate value, and an operand is fetched at 401. As noted above, the first source operand and destination operand are both vector registers. The second source operand may be either a vector  
 25 register or a memory location.

[0054] The DBPSAD instruction is decoded by decoding logic at 403. Depending on the instruction's format, a variety of data may be interpreted at this stage such as if there is to be a data transformation, which registers to write to and retrieve, what memory address to access, etc.

[0055] The source operand values are retrieved/read at 405. For example, the source register(s)  
 30 is/are read or the memory location of the second source operand is retrieved.

[0056] The DBPSAD instruction (or operations comprising such an instruction such as microoperations) is executed by execution resources such as one or more functional units at 407 to, per data lanes (i.e., 64-bit lanes) of the sources, compute the SAD of selected quadruplets of

data elements (i.e., byte elements) of the first and second sources. The specifics of how this may be accomplished for byte elements was detailed above in the description of Figure 1.

[0057] At 409, the computed SAD values are stored in the destination vector register. For example, word results of the SAD calculations are stored in data element positions of the

5 destination register. Generically, the what is stored into the destination data element positions is as follows:

-  $DEST[0] <- ABS(SRC1[0] - TMP1[0]) + ABS(SRC1[1] - TMP1[1]) + ABS(SRC1[2] -$   
 $TMP1[2]) + ABS(SRC1[3] - TMP1[3])$

10 -  $DEST[1] <- (SRC1[0] - TMP1[1]) + ABS(SRC1[1] - TMP1[2]) + ABS(SRC1[2] -$   
 $TMP1[3]) + ABS(SRC1[3] - TMP1[4])$

-  $DEST[2] <- ABS(SRC1[4] - TMP1[2]) + ABS(SRC1[5] - TMP1[3]) + ABS(SRC1[6] -$   
 $TMP1[4]) + ABS(SRC1[7] - TMP1[5])$

-  $DEST[3] <- ABS(SRC1[4] - TMP1[3]) + ABS(SRC1[5] - TMP1[4]) + ABS(SRC1[6] -$   
 $TMP1[5]) + ABS(SRC1[7] - TMP1[6])$

15 [0058] While 407 and 409 have been illustrated separately, in some embodiments they are performed together as a part of the execution of the instruction.

[0059] Figure 5 illustrates an embodiment of a method for processing a DBPSAD instruction. Specifically, what is detailed below is for a data lane (such as a 64-bit data lane). This would be repeated as many times as necessary until all data lanes have been processed. In this

20 embodiment it is assumed that some, if not all, of the operations 401-405 have been performed earlier, however, they are not shown in order to not obscure the details presented below. For example, the fetching and decoding are not shown, nor is the operand retrieval shown.

[0060] At 501, quadruplets of data elements of the first source are selected. More particularly, two different quadruplets with an offset (such as four bytes) is selected from the first source. In

25 Figure 1, this is the least significant 64 bits of the source. If the source and destination registers were 128-bit, then in the subsequent iteration then the 64 most significant bits would be used.

[0061] Two quadruplets from the second source are selected according to control bits from the immediate value of the instruction at 503. Each quadruplet is the same size as a data element of the destination. The selection of quadruplets from the second source is done on a 128-bit step

30 size. In other words, the selection is made from a 128-bit chunk of the second source. If the second source is only 128-bits then only one selection of four quadruplets is made. If the second source is 256-bits, then two selections of four quadruplets is made - one on the 128 least significant bits and one on the 128 most significant bits.

[0062] Typically, the immediate is an 8-bit value and two bits are used per quadruplet selection. The two least significant bits are used to select the first quadruplet (least significant) and so on. For example, for TMP1 of Figure 1, if the immediate is "01" then bits 63:32 of the SRC2 are selected.

5 [0063] A SAD is calculated using the selected quadruplets at 505. Each SAD consists of multiple absolute difference calculations that have been summed together. Again, this is done on a data lane basis. In Figure 1, a 64-bit lane for the least significant 64-bit is illustrated. The specifics how the data elements are processed is discussed with respect to that figure.

[0064] At 507, the calculated SADs are stored into data element positions of the destination  
10 vector register. Again, exemplary specifics on what goes where has been discussed with respect to Figure 1.

[0065] Figure 6 illustrates pseudo-code for DBPSAD for a byte source data element size and word destination data element size.

[0066] Exemplary Instruction Formats

15 [0067] 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.

[0068] Generic Vector Friendly Instruction Format

20 [0069] 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.

25 [0070] **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  
30 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.

[0071] 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  
5 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  
10 (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).

[0072] The class A instruction templates in Figure 8A include: 1) within the no memory access  
15 805 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  
20 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.

25 [0073] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in Figures 8A-8B.

[0074] 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  
30 sense that it is not needed for an instruction set that has only the generic vector friendly instruction format.

[0075] Base operation field 842 - its content distinguishes different base operations.

[0076] 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).

[0077] 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 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.

[0078] 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.

[0079] 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  $2^{scale} * index + base$ ).

[0080] Displacement Field 862A- its content is used as part of memory address generation (e.g., for address generation that uses  $2^{scale} * index + base + displacement$ ).

[0081] 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  $2^{scale} * index + base + scaled\ displacement$ ). 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

5 embodiments may implement only one or none of the two.

**[0082]** 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

10 the opcodes.

**[0083]** 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-

15 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

20 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

25 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

30 specify the masking to be performed.

**[0084]** Immediate field 872 - its content allows for the specification of an immediate. This field is optional in the sense that is it not present in an implementation of the generic vector



friendly format that does not support immediate and it is not present in instructions that do not use an immediate.

**[0085]** 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).

**[0086]** Instruction Templates of Class A

**[0087]** 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.

**[0088]** No-Memory Access Instruction Templates - Full Round Control Type Operation

**[0089]** 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).

**[0090]** 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.

**[0091]** 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.

**[0092]** No Memory Access Instruction Templates - Data Transform Type Operation

5 **[0093]** 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).

**[0094]** 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  
10 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;  
15 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.

**[0095]** Vector memory instructions perform vector loads from and vector stores to memory, with conversion support. As with regular vector instructions, vector memory instructions  
20 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.

**[0096]** Memory Access Instruction Templates - Temporal

**[0097]** 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  
25 the hint entirely.

**[0098]** Memory Access Instruction Templates - Non-Temporal

**[0099]** 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.

30 **[00100]** Instruction Templates of Class B

**[00101]** 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.

[00102] 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.

10 [00103] 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).

[00104] 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.

20 [00105] 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).

[00106] 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.

30 [00107] 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).

[00108] 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.

[00109] 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.

[00110] 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.

[00111] Exemplary Specific Vector Friendly Instruction Format

[00112] **Figure 9** is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention. Figure 9 shows 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 Figure 9 map are illustrated.

[00113] 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).

[00114] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in Figure 9A.

[00115] EVEX Prefix (Bytes 0-3) 902 - is encoded in a four-byte form.

[00116] 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).

[00117] The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

[00118] 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.

[00119] 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.

10 [00120] Opcode map field 915 (EVEX byte 1, bits [3:0] - mmmm) - its content encodes an implied leading opcode byte (OF, OF 38, or OF 3).

[00121] 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).

15 [00122] 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 (Is complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in Is complement form for certain vector shifts; or 3) EVEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. Thus, EVEX.vvvv field 920 encodes the 4 low-order bits of the first source register specifier stored in inverted (Is complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

[00123] 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.

25 [00124] 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.

- 5 [00125] 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 a) - as previously described, this field is context specific.
- [00126] Beta field 854 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s<sub>20</sub>, EVEX.r<sub>20</sub>, EVEX.rri, EVEX.LL0, EVEX.LLB; also illustrated with βββ) - as previously described, this
- 10 field is context specific.
- [00127] 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,
- 15 EVEX.vvvv.
- [00128] 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
- 20 mask hardwired to all ones or hardware that bypasses the masking hardware).
- [00129] Real Opcode Field 930 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.
- [00130] 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
- 25 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.
- 30 [00131] Scale, Index, Base (SIB) Byte (Byte 6) - As previously described, the scale field's 850 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.

[00132] 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.

[00133] Displacement factor field 862B (Byte 7) - when MOD field 942 contains 01, byte 7 is  
5 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  
10 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 disp8\*N. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement is  
15 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  
20 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).

[00134] Immediate field 872 operates as previously described.

25 [00135] Full Opcode Field

[00136] 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  
30 includes the prefix encoding field 925, the opcode map field 915, and the real opcode field 930.

[00137] Register Index Field

[00138] 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 W W field 920, xxx field 954, and the bbb field 956.

[00139] Augmentation Operation Field

5 [00140] 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.

10

15 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.

[00141] 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]- So) 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<sub>2-i</sub>) 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<sub>2-i</sub>) is interpreted as the vector length field 859B (EVEX byte 3, bit [6-5]- L<sub>1,0</sub>). 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<sub>1,0</sub>) and the broadcast field 857B (EVEX byte 3, bit [4]- B).

20

25

[00142] *Exemplary Register Architecture*

30 [00143] **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

5 **[00144]** 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  
 10 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.

15 **[00145]** 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 kO is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

[00146] 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, 5 RSI, RDI, RSP, and R8 through R15.

[00147] 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 10 on 64-bit packed integer data, as well as to hold operands for some operations performed between the MMX and XMM registers.

[00148] 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. 15

[00149] Exemplary Core Architectures, Processors, and Computer Architectures

[00150] 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. 20 25 30

Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

[00151] Exemplary Core Architectures

[00152] In-order and out-of-order core block diagram

5 [00153] **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  
10 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.

[00154] In Figure 11A, a processor pipeline 1100 includes a fetch stage 1102, a length decode  
15 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.

[00155] Figure 11B shows processor core 1190 including a front end unit 1130 coupled to an  
20 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  
25 processing unit (GPGPU) core, graphics core, or the like.

[00156] 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  
30 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  
5 execution engine unit 1150.

**[00157]** 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.

10 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  
15 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  
20 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  
25 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  
30 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.

5 [00158] 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  
10 (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.

[00159] 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  
15 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  
20 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.

[00160] 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  
25 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, and/or some form of the generic vector friendly instruction format (U=0 and/or U=1) previously described), thereby allowing the operations used by many multimedia applications to be performed using packed data.

30 [00161] 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).

[00162] 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

5 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 (LI) 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.

10 Alternatively, all of the cache may be external to the core and/or the processor.

[00163] Specific Exemplary In-Order Core Architecture

[00164] **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

15 interfaces, and other necessary I/O logic, depending on the application.

[00165] 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

20 supports the x86 instruction set with a packed data instruction set extension. An LI 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 (LI) cache 1206,

25 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).

[00166] 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

30 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.

[00167] 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 LI data cache 1206A part of the LI 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.

[00168] Processor with integrated memory controller and graphics

[00169] **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.

[00170] 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.



[00171] The memory hierarchy includes one or more levels of cache 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.

[00172] 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.

[00173] 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.

[00174] Exemplary Computer Architectures

[00175] **Figures 14-17** 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.

[00176] 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 coupled to 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  
5 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.

[00177] 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.

10 [00178] 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.

15 [00179] 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.

[00180] There can be a variety of differences between the physical resources 1410, 1415 in  
20 terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

[00181] 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  
25 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.

[00182] Referring now to **Figure 15**, shown is a block diagram of a first more specific  
30 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.

[00183] 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.

[00184] 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 1539. 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.

[00185] 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.

[00186] 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.

[00187] 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.

5 [00188] 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.

[00189] 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.

15 [00190] Referring now to **Figure 17**, shown is a block diagram of a SoC 1700 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 17, an interconnect unit(s) 1702 is coupled to: an application processor 1710 which includes a set of one or more cores 202A-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 1720 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 1730; a direct memory access (DMA) unit 1732; and a display unit 1740 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 1720 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.

[00191] 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.

[00192] 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.

[00193] 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.

[00194] 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.

[00195] 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.

[00196] 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.

[00197] Emulation (including binary translation, code morphing, etc.)

[00198] 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  
5 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.

[00199] **Figure 18** 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  
10 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 18 shows a program in a high level language 1802 may be compiled using an x86 compiler 1804 to generate x86 binary code 1806 that may be natively executed by a processor with at least one  
15 x86 instruction set core 1816. The processor with at least one x86 instruction set core 1816 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  
20 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 1804 represents a compiler that is operable to generate x86 binary code 1806 (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 1816. Similarly, Figure 18 shows the program in the high level language 1802 may be compiled using  
25 an alternative instruction set compiler 1808 to generate alternative instruction set binary code 1810 that may be natively executed by a processor without at least one x86 instruction set core 1814 (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 1812 is used to convert the x86 binary code 1806 into code that  
30 may be natively executed by the processor without an x86 instruction set core 1814. This converted code is not likely to be the same as the alternative instruction set binary code 1810 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 1812 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 1806.

5

**Claims**

What is claimed is:

1. A method of performing in a computer processor vector double block packed sum of absolute differences (SAD) in response to a single vector double block packed sum of absolute differences instruction that includes a destination vector register operand, first and second source operands, an immediate, and an opcode, the method comprising steps of:
  - executing the single vector double block packed sum of absolute differences instruction to, on a per data lane basis, compute a SAD of selected quadruplets of data elements of the first and second sources; and
  - storing each calculated SAD into the destination vector register.
2. The method of claim 1, wherein the data elements of the selected quadruplets of the first and second sources are byte sized.
3. The method of claim 2, wherein the data elements of the destination register are word sized.
4. The method of claim 1, wherein the first source operand is a vector register and the second source operand is a memory location.
5. The method of claim 1, wherein the first and second source operands are vector registers.
6. The method of claim 1, wherein the first and second sources and the destination vector register operand are all of the same size selected from the group consisting of 128-bit, 256-bit, and 512-bit.
7. The method of claim 1, wherein the stored SADs comprise:
  - in a least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a



third-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a fourth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source;

5           in a second-most least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute  
10 value of a third-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a fourth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source;

          in a third-most least significant position of the destination register, the stored SAD is  
15 an absolute value of a fifth-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a seventh-most least significant data element position of the first source minus  
20 a fifth-most least significant data element position of the second source that is added to an absolute value of a eighth-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source; and

          in a fourth-most least significant position of the destination register, the stored SAD is an absolute value of a fifth-most least significant data element position of the first source minus  
25 a fourth-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source that is added to an absolute value of a seventh-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source that is added to an absolute  
30 value of a eighth-most least significant data element position of the first source minus a seventh-most least significant data element position of the second source.

8.       An article of manufacture comprising:

a tangible machine-readable storage medium having stored thereon an occurrence of an instruction, wherein the instruction's format specifies as its source operands a first and second source and an immediate and specifies as its destination a single destination vector register, and wherein the instruction format includes an opcode which instructs a machine, responsive to the single occurrence of the single instruction, to cause on a per data lane basis, a computation of a SAD of selected quadruplets of data elements of the first and second sources and storage of each calculated SAD into the destination vector register.

9. The article of manufacture of claim 8, wherein the data elements of the selected quadruplets of the first and second sources are byte sized.

10. The article of manufacture of claim 9, wherein the data elements of the destination register are word sized.

11. The article of manufacture of claim 8, wherein the first source operand is a vector register and the second source operand is a memory location.

12. The article of manufacture of claim 8, wherein the first and second source operands are vector registers.

13. The article of manufacture of claim 8, wherein the first and second sources and the destination vector register operand are all of the same size selected from the group consisting of 128-bit, 256-bit, and 512-bit.

14. The article of manufacture of claim 8, wherein the stored SADs comprise:  
in a least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a third-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a

fourth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source;

in a second-most least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a third-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a fourth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source;

in a third-most least significant position of the destination register, the stored SAD is an absolute value of a fifth-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a seventh-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source that is added to an absolute value of an eighth-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source; and

in a fourth-most least significant position of the destination register, the stored SAD is an absolute value of a fifth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source that is added to an absolute value of a seventh-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source that is added to an absolute value of an eighth-most least significant data element position of the first source minus a seventh-most least significant data element position of the second source.

30

15. An apparatus comprising;

a hardware decoder to decode a single vector double block packed sum of absolute differences instruction that includes a destination vector register operand, first and second source operands, an immediate, and an opcode;

execution logic to, on a per data lane basis, compute a SAD of selected quadruplets of data elements of the first and second sources and store each calculated SAD into the destination vector register.

16. The apparatus of claim 15, wherein the data elements of the selected quadruplets of the first and second sources are byte sized.

17. The apparatus of claim 16, wherein the data elements of the destination register are word sized.

18. The apparatus of claim 15, wherein the first source operand is a vector register and the second source operand is a memory location.

19. The apparatus of claim 15, wherein the first and second source operands are vector registers.

20. The apparatus of claim 15, wherein the stored SADs comprise:

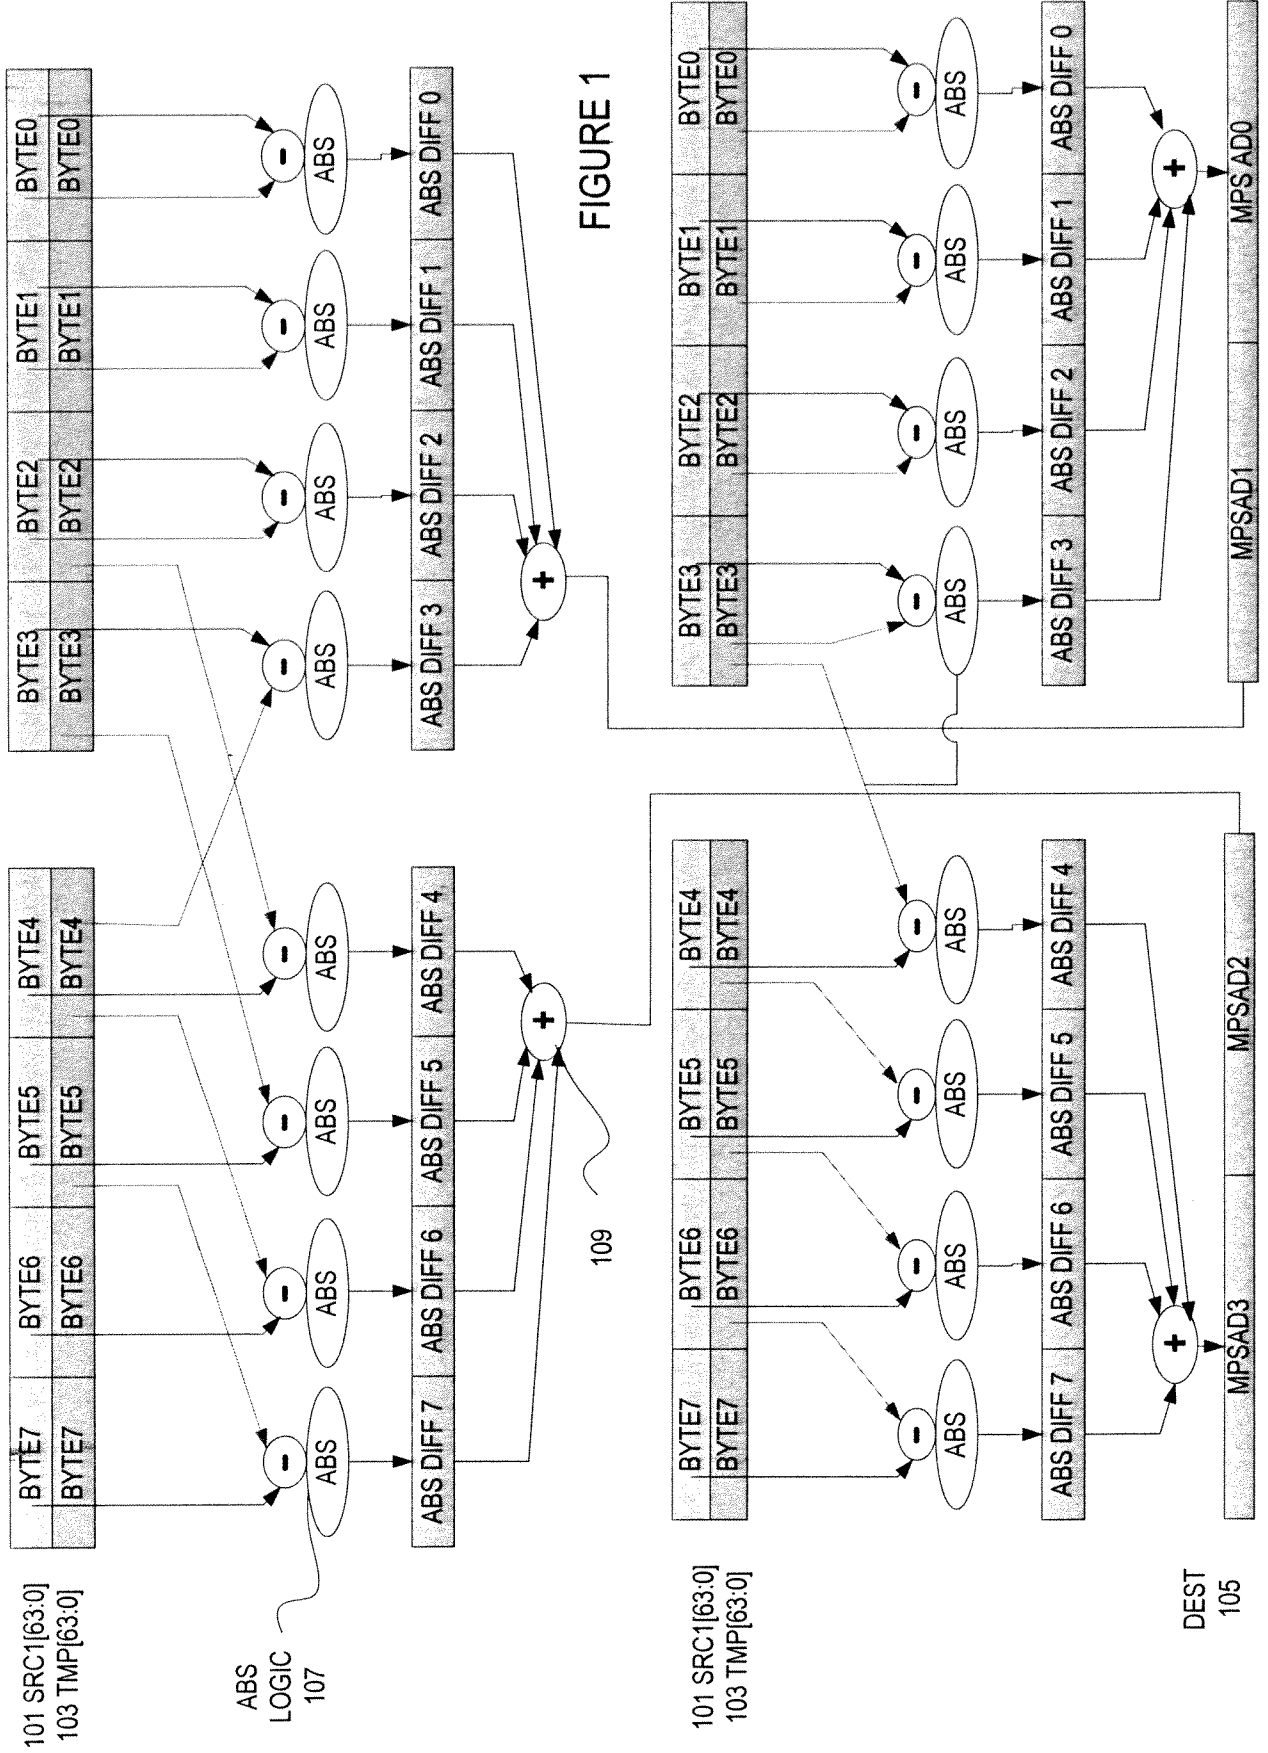
in a least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a third-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a fourth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source;

in a second-most least significant position of the destination register, the stored SAD is an absolute value of a least significant data element position of the first source minus a second-most least significant data element position of the second source that is added to an absolute value of a second-most least significant data element position of the first source minus a third-

most least significant data element position of the second source that is added to an absolute value of a third-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a fourth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source;

5 in a third-most least significant position of the destination register, the stored SAD is an absolute value of a fifth-most least significant data element position of the first source minus a third-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a  
10 fourth-most least significant data element position of the second source that is added to an absolute value of a seventh-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source that is added to an absolute value of a eighth-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source; and

15 in a fourth-most least significant position of the destination register, the stored SAD is an absolute value of a fifth-most least significant data element position of the first source minus a fourth-most least significant data element position of the second source that is added to an absolute value of a sixth-most least significant data element position of the first source minus a fifth-most least significant data element position of the second source that is added to an absolute  
20 value of a seventh-most least significant data element position of the first source minus a sixth-most least significant data element position of the second source that is added to an absolute value of a eighth-most least significant data element position of the first source minus a seventh-most least significant data element position of the second source.



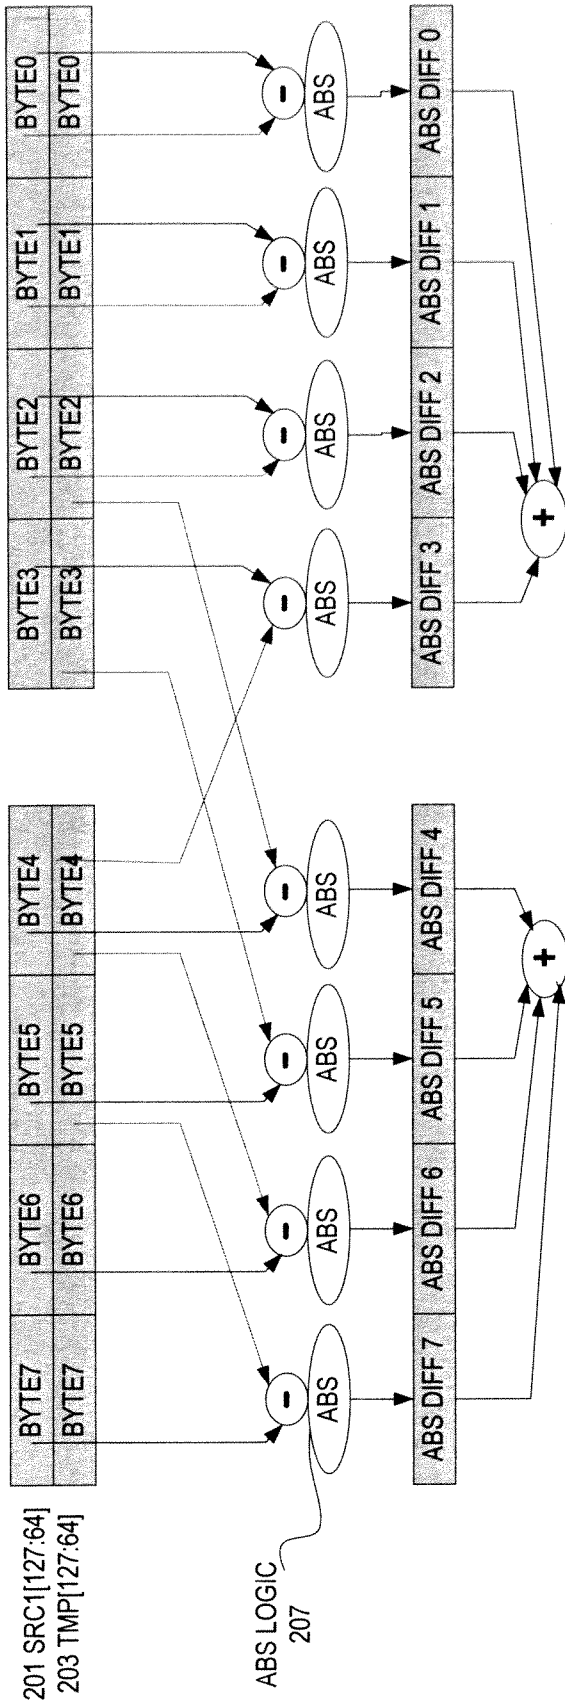



FIGURE 2



Opcode/ Instruction	Op En	Mode Suppo rt	CPUID Feature Flag	Description
EVEX.U1.NDS.128.66.0F3A.W 0 42 /ib VDBPSADBW xmm1, xmm2, xmm3/m128, imm8	A	V/V	AVX3.2	Compute the SAD (sum of absolute differences) of selected quadruplets of byte elements in xmm2 compared to quadruplet of byte elements in xmm3/m128. Select quadruplets from xmm3/m128 according to control in imm8. Word results of the SAD operations are stored in zmm1.
EVEX.U1.NDS.256.66.0F3A.W 0 42 VDBPSADBW ymm1, ymm2, ymm3/m256, imm8 /ib	A	V/V	AVX3.2	Compute the SAD (sum of absolute differences) of selected quadruplets of byte elements in ymm2 compared to quadruplet of byte elements in ymm3/m256. Select quadruplets from ymm3/m256 according to control in imm8. Word results of the SAD operations are stored in ymm1.
EVEX.U1.NDS.512.66.0F3A.W 0 42 VDBPSADBW zmm1, zmm2, zmm3/m512, imm8 /ib	A	V/V	AVX3.2	Compute the SAD (sum of absolute differences) of selected quadruplets of byte elements in zmm2 compared to quadruplet of byte elements in zmm3/m512. Select quadruplets from zmm3/m512 according to control in imm8. Word results of the SAD operations are stored in zmm1.

FIGURE 3



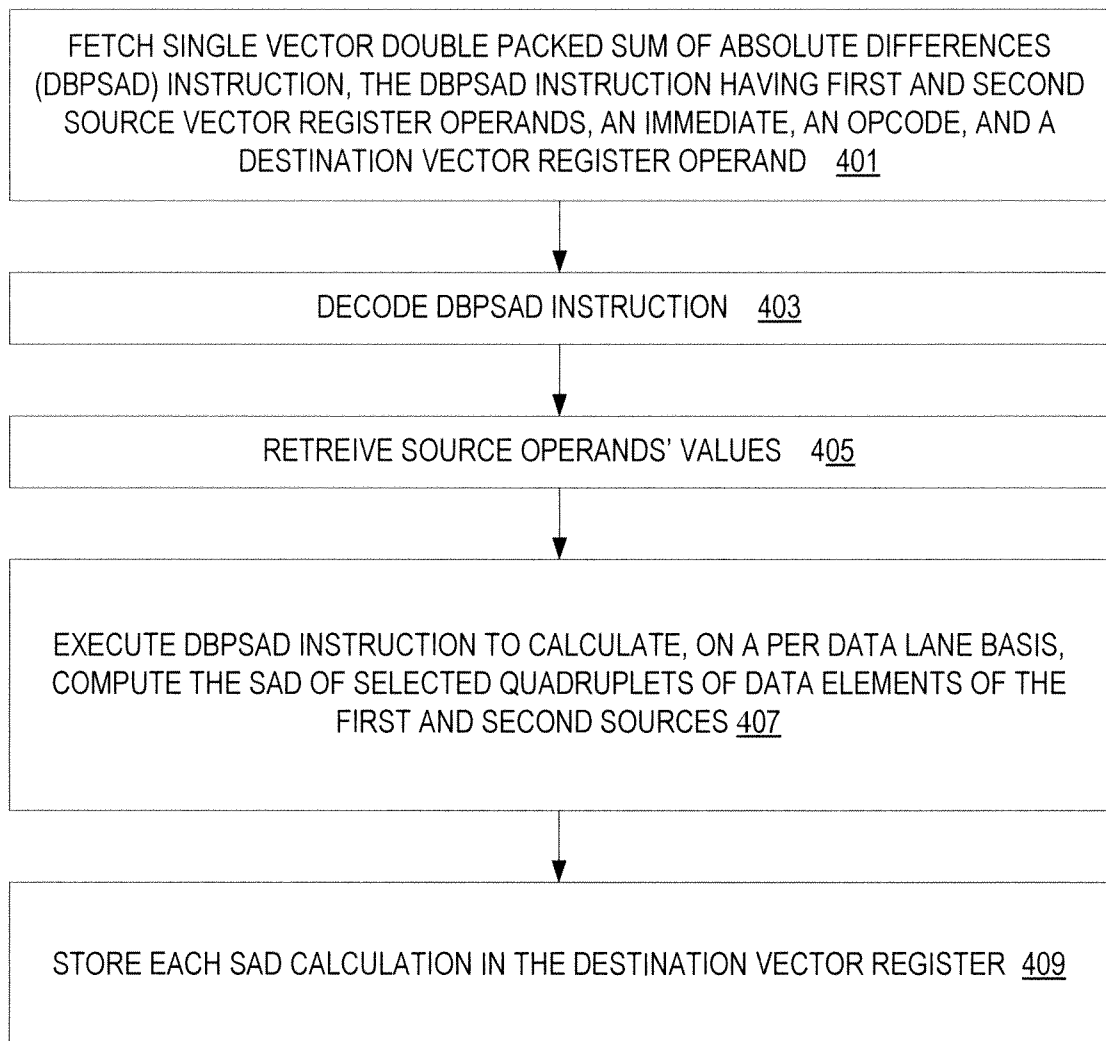


FIGURE 4

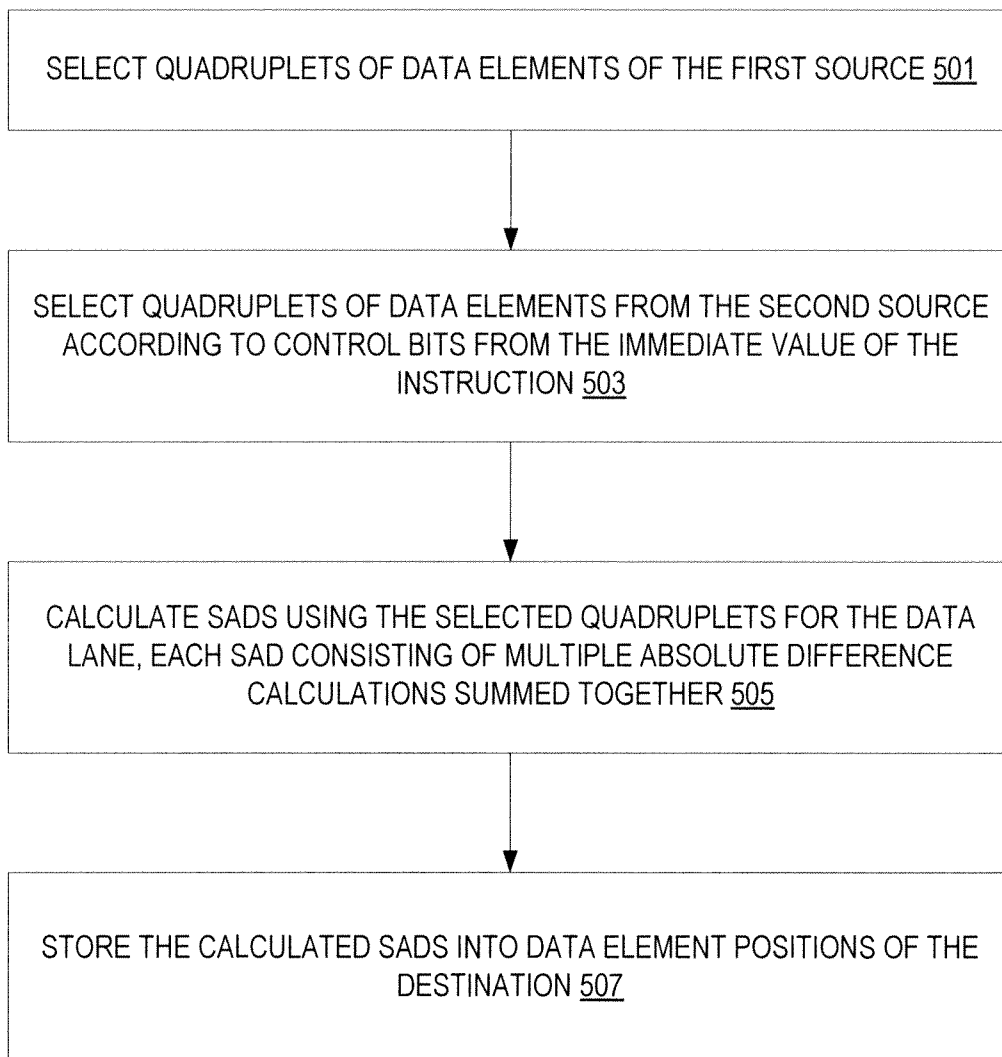


FIGURE 5

Pseudo-code for VDBPSADBW (Byte-sized Data Elements in the Source and Word-sized Data Elements in the Destination)

Per Data Lane Selection of Quadruplets

VL = 128, 256, 512 Selection of quadruplets:

FOR I = 0 to VL step 128

TMP1[I+31:I] <- select (SRC2[I+127: I], imm8[1:0])

TMP1[I+63: I+32] <- select (SRC2[I+127: I], imm8[3:2])

TMP1[I+95: I+64] <- select (SRC2[I+127: I], imm8[5:4])

TMP1[I+127: I+96] <- select (SRC2[I+127: I], imm8[7:6])

END FOR

SAD of Quadruplets

DEST[I+15:I] <- ABS(SRC1[I+7: I] - TMP1[I+7: I]) + ABS(SRC1[I+15: I+8] - TMP1[I+15: I+8]) + ABS(SRC1[I+23: I+16] - TMP1[I+23: I+16]) + ABS(SRC1[I+31: I+24] - TMP1[I+31: I+24])

DEST[I+31: I+16] <- (SRC1[I+7: I] - TMP1[I+15: I+8]) + ABS(SRC1[I+15: I+8] - TMP1[I+23: I+16]) + ABS(SRC1[I+23: I+16] - TMP1[I+31: I+24]) + ABS(SRC1[I+31: I+24] - TMP1[I+39: I+32])

DEST[I+47: I+32] <- ABS(SRC1[I+39: I+32] - TMP1[I+23: I+16]) + ABS(SRC1[I+47: I+40] - TMP1[I+31: I+24]) + ABS(SRC1[I+55: I+48] - TMP1[I+39: I+32]) + ABS(SRC1[I+63: I+56] - TMP1[I+47: I+40])

DEST[I+63: I+48] <- ABS(SRC1[I+39: I+32] - TMP1[I+31: I+24]) + ABS(SRC1[I+47: I+40] - TMP1[I+39: I+32]) + ABS(SRC1[I+55: I+48] - TMP1[I+47: I+40]) + ABS(SRC1[I+63: I+56] - TMP1[I+55: I+48])

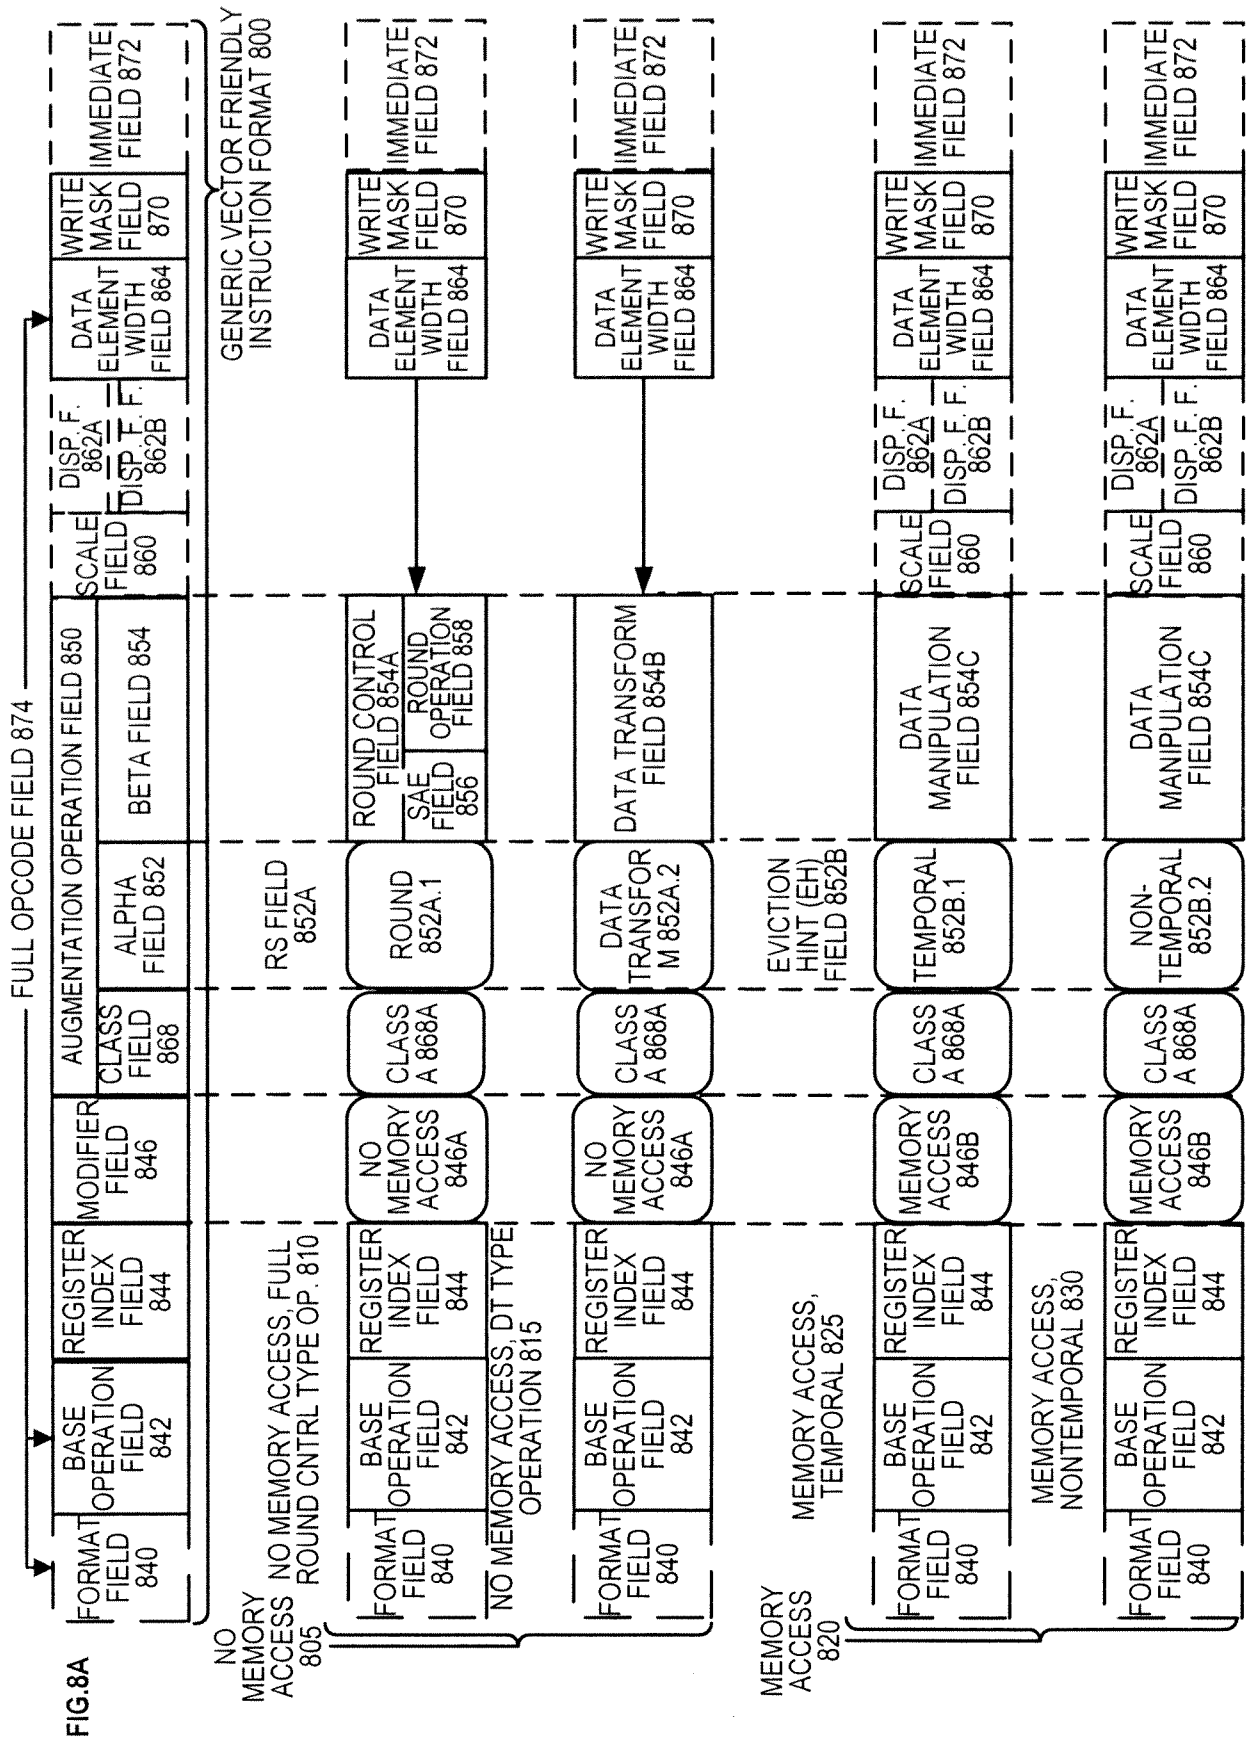
FIGURE 6

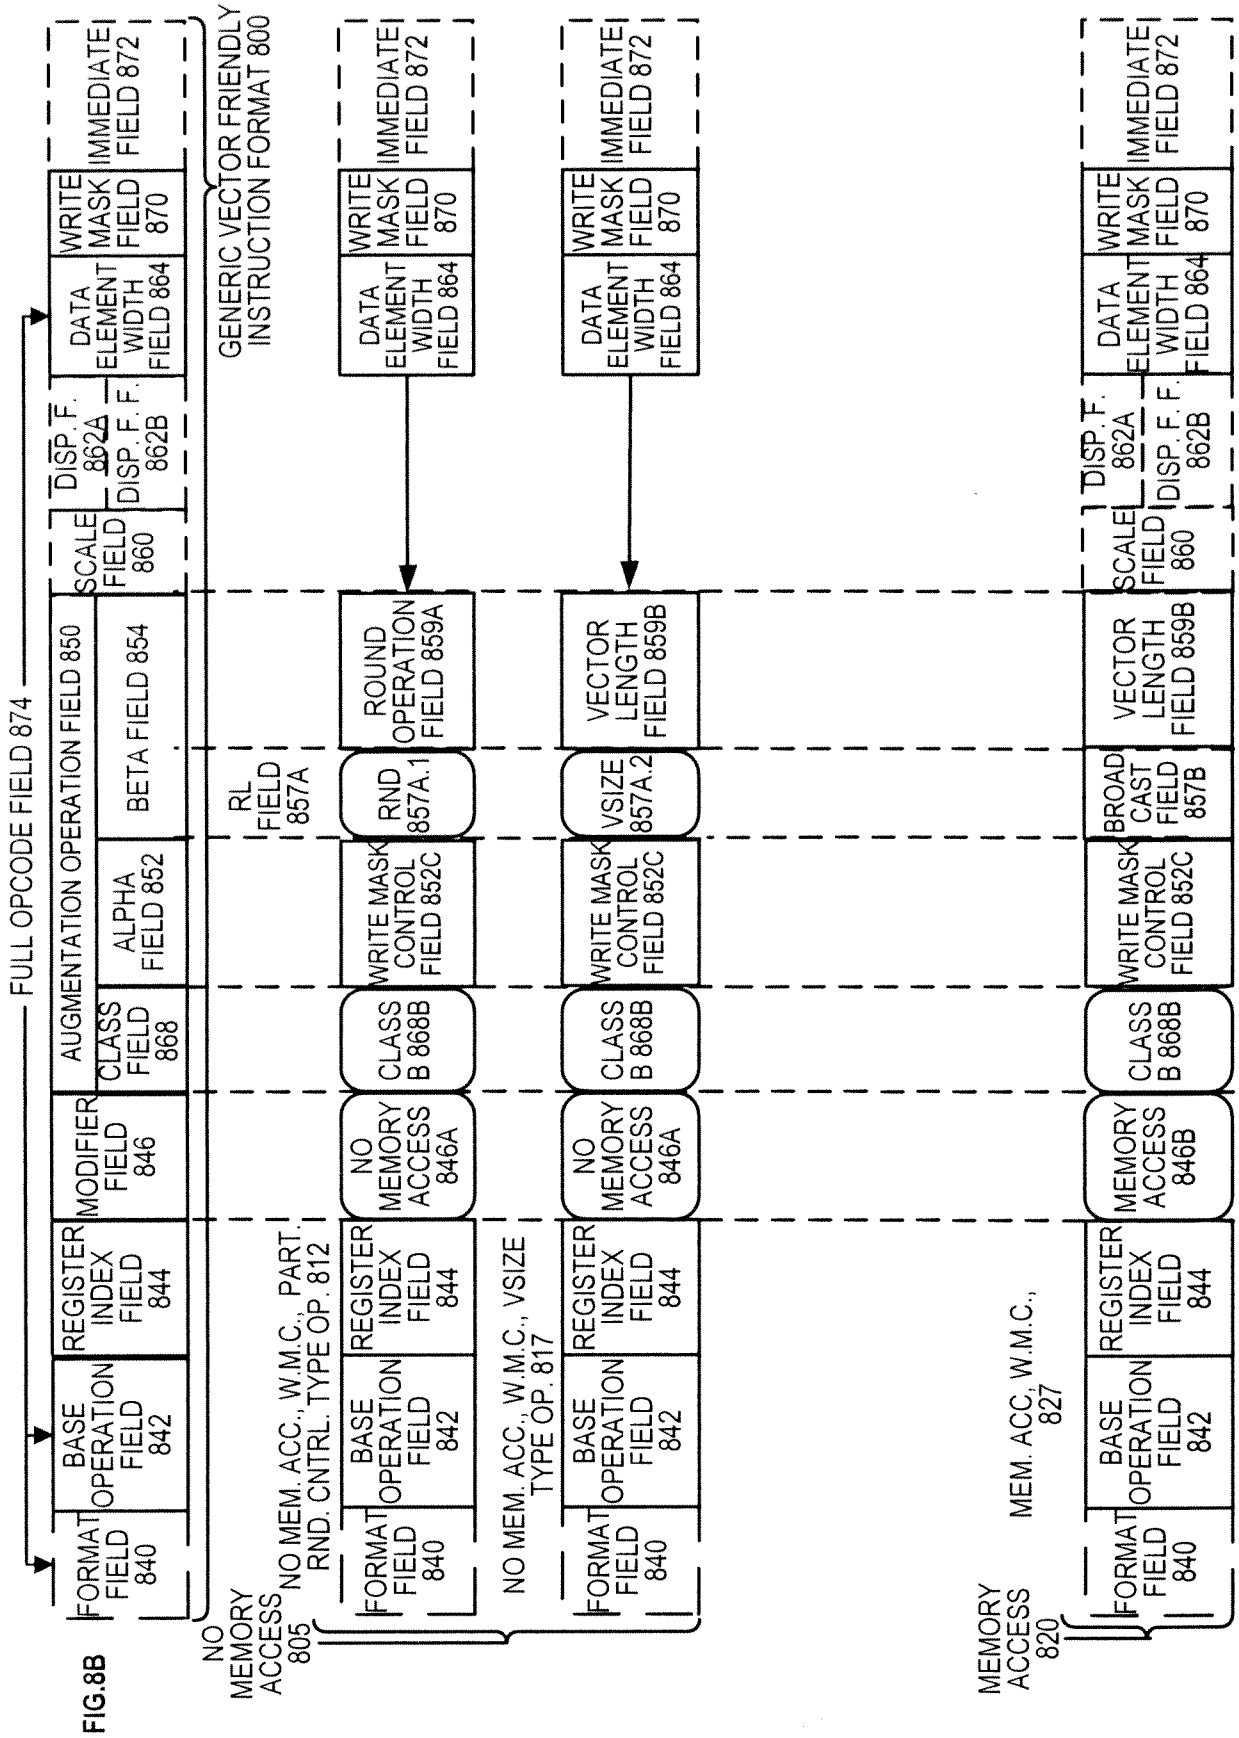
NUMBER OF ONE BIT VECTOR WRITE MASK ELEMENTS

DATA ELEMENT SIZE FOR VECTOR	VECTOR SIZE	
	128 BITS	256 BITS
8-BIT BYTES	16	32
16-BIT WORDS	8	16
32-BIT DWORDS/SP	4	8
64-BIT QWORDS/DP	2	4
		512 BITS
		64
		32
		16
		8



Figure 7






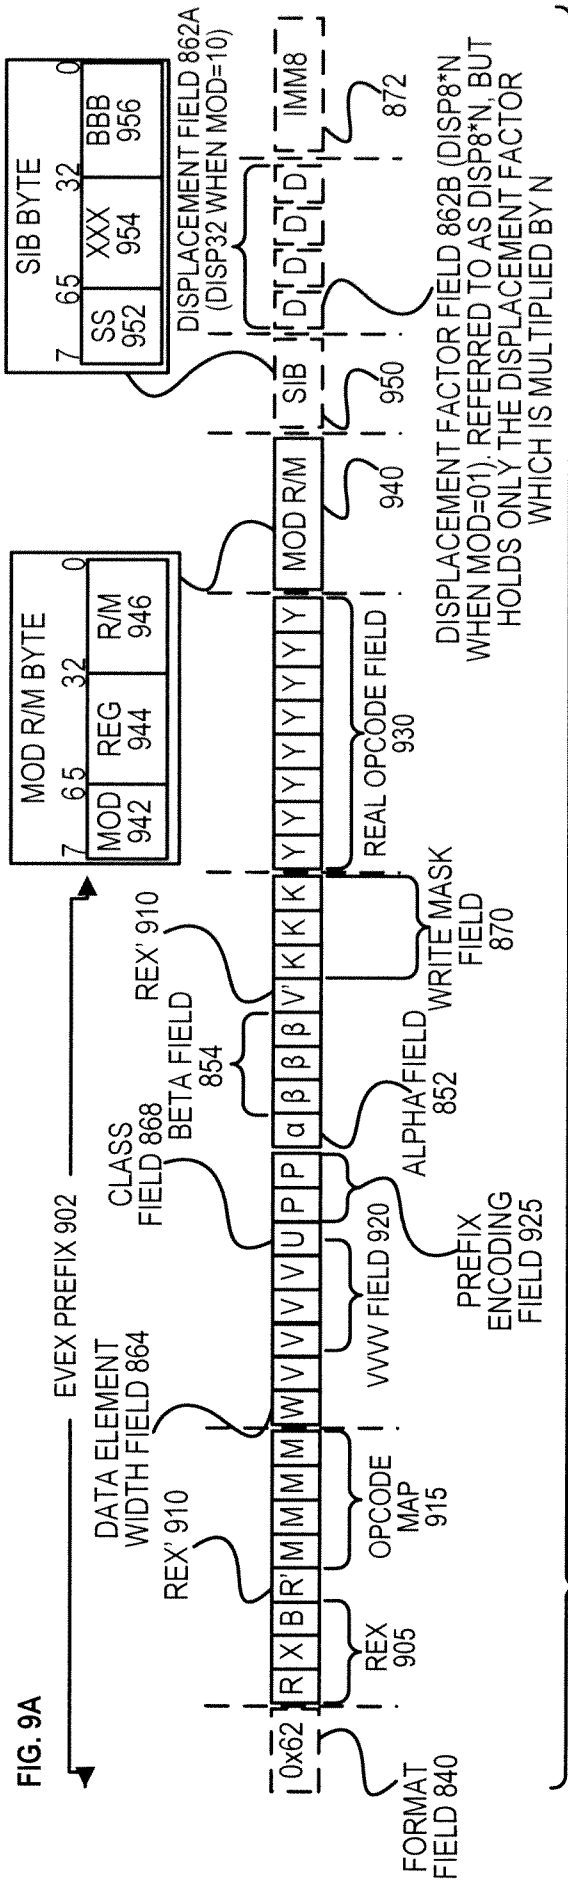


FIG. 9C


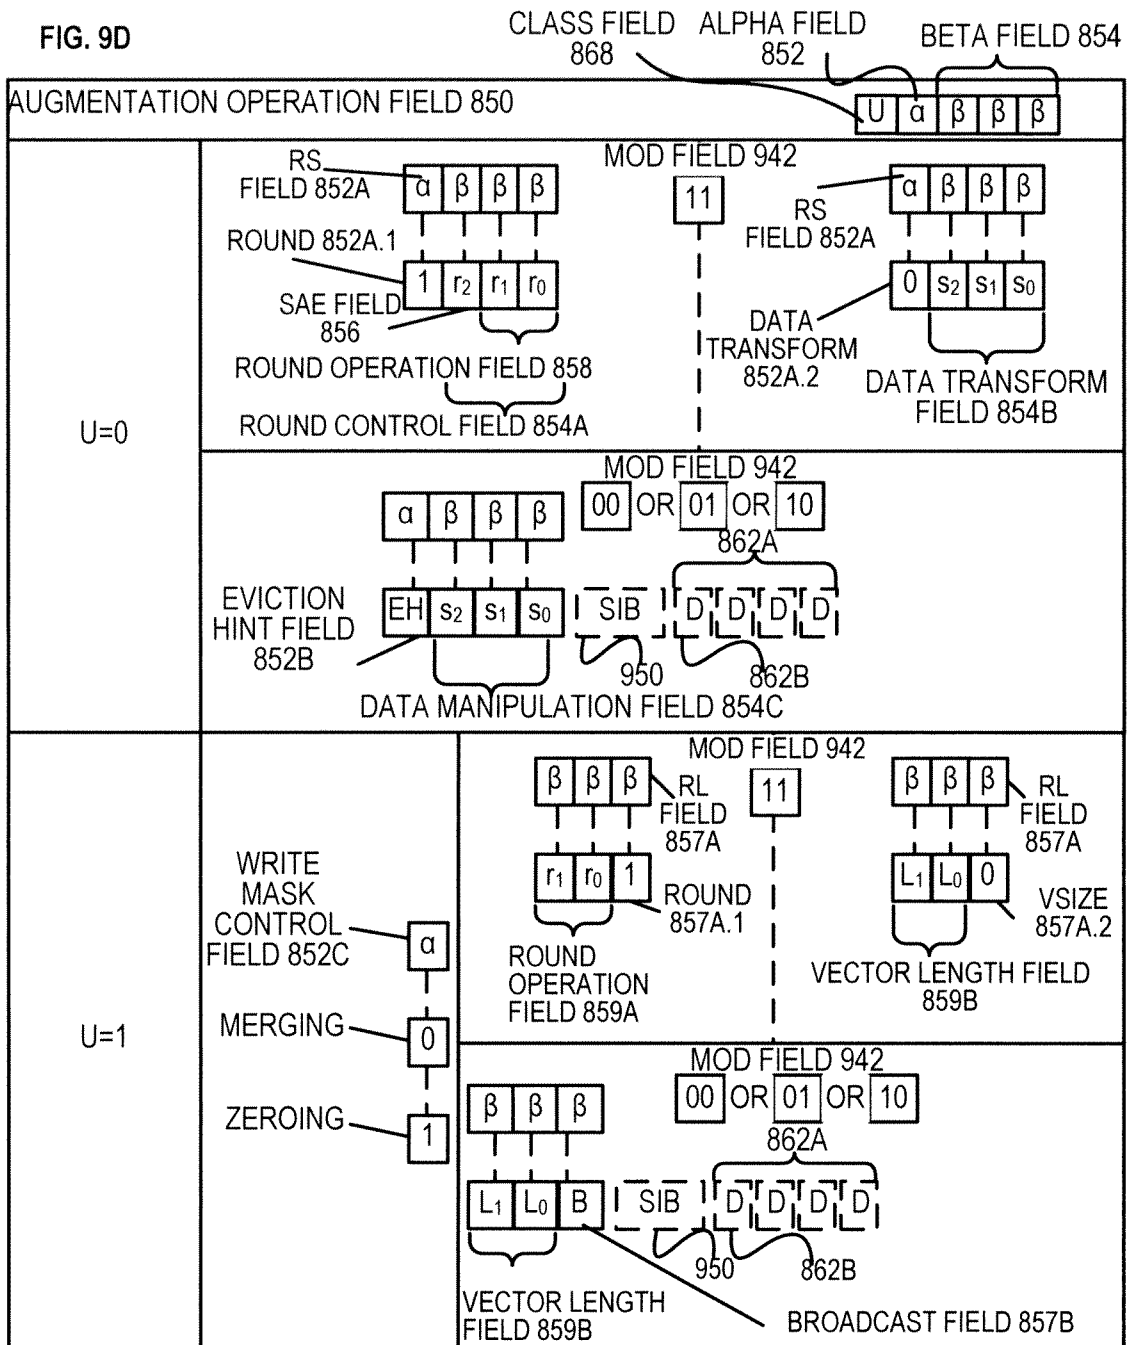







FIG. 9D

















FIG. 14




FIG. 15




FIG. 16




FIG. 17



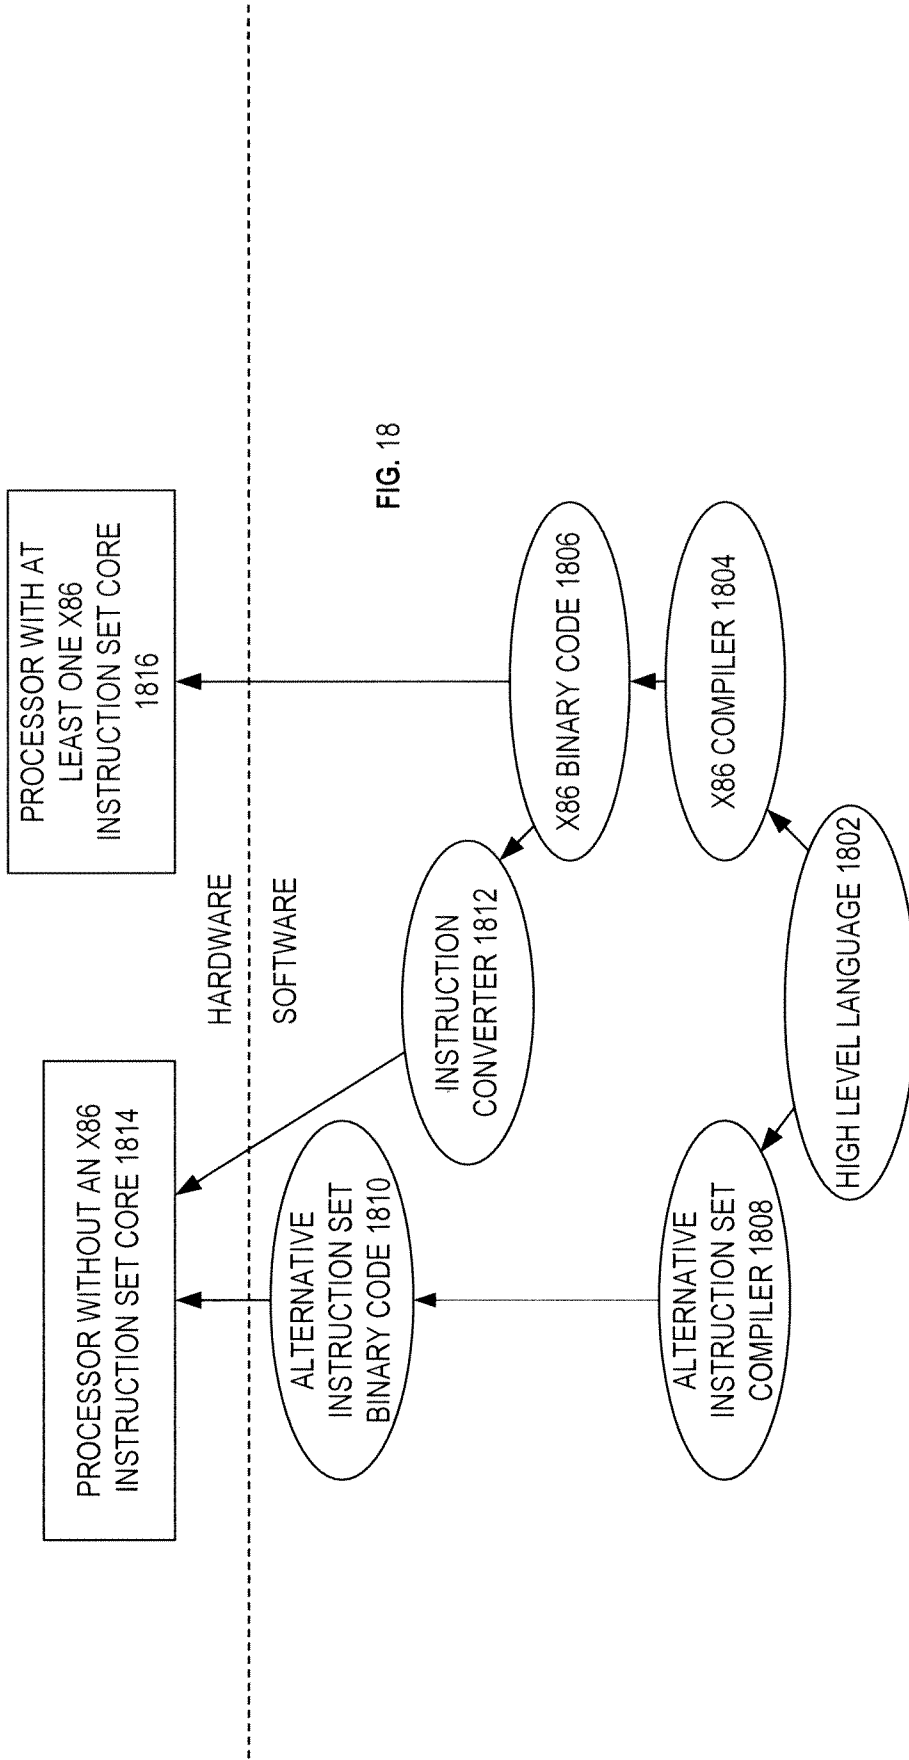


FIG. 18

**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/06(2006.01)i, G06F 9/30(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 9/06; G06F 15/00; G06F 9/302; G06F 9/22; G06F 7/00; H04N 7/12; G06F 7/38

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) &amp; Keywords: sum of absolute difference, SAD, packed, double block, second source;

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 7054895 B2 (IGOR M KOBA et al.) 30 May 2006 See the abstract, claims 1-45 and figures 1-6.	1-20
A	US 7516307 B2 (ABDALLAH MOHAMMAD A. et al.) 07 April 2009 See the abstract, claims 1-12 and figures 1-13.	1-20
A	US 2007-0263730 A1 (MAO ZENG et al.) 15 November 2007 See the abstract, claims 1-28 and figures 1-10.	1-20
A	US 05880979A A (MENNEMEIER LARRY M. et al.) 09 March 1999 See the abstract, claims 1-15 and figures 1-7b.	1-20

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

\* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

24 AUGUST 2012 (24.08.2012)

Date of mailing of the international search report

**27 AUGUST 2012 (27.08.2012)**

Name and mailing address of the ISA/KR

Korean Intellectual Property Office  
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan  
City, 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

BOK, Jin Yo

Telephone No. 82-42-481-5113



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US201 1/067071**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7054895 B2	30.05.2006	US 2003-005267 A1	02.01.2003
US 7516307 B2	07.04.2009	US 2002-006233 1 A1 US 6377970 B1	23.05.2002 23.04.2002
US 2007-0263730 A1	15.11.2007	CN 10 1438598 A EP 2025 175 A2 JP 2009-536773 A KR 10-0996337 B1 WO 2007- 134011 A2 WO 2007- 134011 A3	20.05.2009 18.02.2009 15.10.2009 23.11.2010 22.11.2007 22.11.2007
US 05880979A A	09.03.1999	WO 97-23822 A1	03.07.1997