

(19) **DANMARK**



Patent- og
Varemærkestyrelsen

(10) **DK/EP 3779681 T3**

(12) Oversættelse af
europæisk patentskrift

-
- (51) Int.Cl.: **G 06 F 9/30 (2018.01)**
- (45) Oversættelsen bekendtgjort den: **2024-07-08**
- (80) Dato for Den Europæiske Patentmyndigheds bekendtgørelse om meddelelse af patentet: **2024-04-10**
- (86) Europæisk ansøgning nr.: **20199012.4**
- (86) Europæisk indleveringsdag: **2019-02-13**
- (87) Den europæiske ansøgnings publiceringsdag: **2021-02-17**
- (30) Prioritet: **2018-03-28 US 201815938924**
- (62) Stamansøgningsnr: **19157044.9**
- (84) Designerede stater: **AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**
- (73) Patenthaver: **INTEL Corporation, 2200 Mission College Blvd., Santa Clara, CA 95054, USA**
- (72) Opfinder: **NARAYANAMOORTHY, Srinivasan, 1450 NE Alex Way, 293, Hillsboro, OR Oregon 97124, USA**
SATISH, Nadathur Rajagopalan, 700 Agnew Road, 135, Santa Clara, CA California 95054, USA
SUPRUN, Alexey, 17667 SW Cody Lane, Beaverton, OR Oregon 97007, USA
JANIK, Kenneth J., 530 NE Geraldine Drive, Hillsboro, OR Oregon 97124, USA
- (74) Fuldmægtig i Danmark: **RWS Group, Europa House, Chiltern Park, Chiltern Hill, Chalfont St Peter, Bucks SL9 9FG, Storbritannien**
- (54) Benævnelse: **ACCELERATOR TIL SPARSOM-TÆT MATRIXMULTIPLIKATION**
- (56) Fremdragne publikationer:
US-A- 5 673 407
US-A1- 2001 011 291
US-A1- 2004 133 617
US-A1- 2007 198 621
US-A1- 2015 277 904
US-A1- 2016 378 465
US-A1- 2018 046 900
MORAD AMIR ET AL: "Efficient Dense and Sparse Matrix Multiplication on GP-SIMD", 2014 24TH INTERNATIONAL WORKSHOP ON POWER AND TIMING MODELING, OPTIMIZATION AND SIMULATION (PATMOS), IEEE, 29 September 2014 (2014-09-29), pages 1-8, XP032682133, DOI: 10.1109/PATMOS.2014.6951895
YAVITS L ET AL: "Sparse Matrix Multiplication On An Associative Processor", IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, IEEE SERVICE CENTER, LOS ALAMITOS, CA, US, vol. 26, no. 11, 1 November 2015 (2015-11-01), pages 3175-3183, XP011586282, ISSN: 1045-9219, DOI: 10.1109/TPDS.2014.2370055 [retrieved on 2015-10-07]
KOANANTAKOOL PENPORN ET AL: "Communication-Avoiding Parallel Sparse-Dense Matrix-Matrix Multiplication", 2014 IEEE 28TH INTERNATIONAL PARALLEL AND DISTRIBUTED PROCESSING SYMPOSIUM,

Fortsættes ...

IEEE, 23 May 2016 (2016-05-23), pages 842-853, XP032926642, ISSN: 1530-2075, DOI: 10.1109/IPDPS.2016.117 [retrieved on 2016-07-18]

ACER SEHER ET AL: "Improving performance of sparse matrix dense matrix multiplication on large-scale parallel systems", PARALLEL COMPUTING, vol. 59, 6 October 2016 (2016-10-06), pages 71-96, XP029801056, ISSN: 0167-8191, DOI: 10.1016/J.PARCO.2016.10.001

KIRAN MATAM ET AL: "Sparse matrix-matrix multiplication on modern architectures", HIGH PERFORMANCE COMPUTING (HIPC), 2012 19TH INTERNATIONAL CONFERENCE ON, IEEE, 18 December 2012 (2012-12-18), pages 1-10, XP032383550, DOI: 10.1109/HIPC.2012.6507483 ISBN: 978-1-4673-2372-7

WU TAO ET AL: "An efficient sparse-dense matrix multiplication on a multicore system", 2017 IEEE 17TH INTERNATIONAL CONFERENCE ON COMMUNICATION TECHNOLOGY (ICCT), IEEE, 27 October 2017 (2017-10-27), pages 1880-1883, XP033340220, DOI: 10.1109/ICCT.2017.8359956 [retrieved on 2018-05-15]

DESCRIPTION

Description

FIELD OF INVENTION

[0001] The present disclosure pertains to the field of processing logic, microprocessors, and associated instruction set architectures, and more specifically, to an accelerator for sparse-dense matrix instructions.

BACKGROUND

[0002] Deep Learning is a class of machine learning algorithms. Deep learning architectures, such as deep neural networks, have been applied to fields including computer vision, speech recognition, natural language processing, audio recognition, social network filtering, machine translation, bioinformatics and drug design.

[0003] Inference and training, two tools used for deep learning, are tending towards low precision arithmetic. Maximizing throughput of deep learning algorithms and computations may assist in meeting the needs of deep learning processors, for example, those performing deep learning in a data center.

[0004] Sparse-dense matrix multiplication (SDMM) operations are useful in a deep learning context. But traditional CPU and GPU instruction set architectures require symmetric inputs having the same density, which limits the ability to gain a performance advantage by taking advantage of the sparsity of a sparse input matrix.

[0005] US 2018/046900 A1 relates to primitive operations of a sparse convolutional neural network accelerator. For example, a method, computer program product, and system perform computations using a processor is provided. A first instruction including a first index vector operand and a second index vector operand is received and the first index vector operand is decoded to produce first coordinate sets for a first array, each first coordinate set including at least a first coordinate and a second coordinate of a position of a non-zero element in the first array. The second index vector operand is decoded to produce second coordinate sets for a second array, each second coordinate set including at least a third coordinate and a fourth coordinate of a position of a non-zero element in the second array. The first coordinate sets are summed with the second coordinate sets to produce output coordinate sets and the output coordinate sets are converted into a set of linear indices.

[0006] Morad Amir et al.: "Efficient Dense and Sparse Matrix Multiplication on GP-SIMD" relates to dense and sparse matrix multiplication on GP-SIMD, a hybrid general purpose SIMD computer architecture that eliminates synchronization by in-memory computing, combining data storage and massively parallel processing.

[0007] Yavits L. et al.: "Sparse Matrix Multiplication on an Associative Processor" concerns sparse matrix multiplication algorithms, combining associative processor (AP) and baseline CPU processing to various levels.

[0008] Koanantakool Penporn et al.: "Communication-Avoiding Parallel Sparse-Dense Matrix-Matrix Multiplication" analyzes communication lower bounds and compares the communication costs of various classic parallel algorithms in the context of sparse-dense matrix-matrix multiplication. A communication-avoiding algorithms based on a 1D decomposition, called 1.5D, is discussed as outperforming 2D and 3D variants for sparse-dense multiplication.

[0009] Wu Tao et al.: "An efficient sparse-dense matrix multiplication on a multicore system" describes a sparse matrix storage format, block-based compressed storage format and coordinate format, and a thread-scalable computing kernel for sparse dense matrix multiplication.

SUMMARY

[0010] The present invention is defined in the independent claims. The dependent claims define embodiments thereof. Any "embodiment" or "example" which is disclosed in the following description but is not covered by the claims should be considered as presented for illustrative purpose only.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] 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:

Figure 1 is a block diagram illustrating processing components for executing a Sparse-dense matrix multiplication (SDMM) instruction, such as a virtual neural network instruction (SDMM), according to some embodiments;

Figure 2A is a block diagram illustrating a data flow for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figure 2B is a block diagram illustrating a data flow for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figure 3 is a block diagram illustrating execution circuitry to process a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figure 4 is a block diagram illustrating execution circuitry to process a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figure 5 is pseudocode illustrating execution circuitry to process a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figure 6 is a process flow diagram illustrating execution of a sparse-dense matrix multiplication (SDMM) instruction by a processor, according to some embodiments;

Figure 7 is a block diagram illustrating a format of a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments;

Figures 8A-8B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to some embodiments;

Figure 8A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to some embodiments;

Figure 8B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to some embodiments;

Figure 9A is a block diagram illustrating an exemplary specific vector friendly instruction format according to some embodiments;

Figure 9B is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the full opcode field according to one embodiment of the invention;

Figure 9C is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the register index field according to one embodiment of the invention;

Figure 9D is a block diagram illustrating the fields of the specific vector friendly instruction format that makes up an augmentation operation field according to one embodiment of the invention;

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

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 some embodiments;

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 some embodiments;

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;

Figure 12A is a block diagram of a single processor core, along with its connection to the on-die interconnect network and with its local subset of the Level 2 (L2) cache, according to some embodiments;

Figure 12B is an expanded view of part of the processor core in **Figure 12A** according to some embodiments;

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 some embodiments;

Figures 14-17 are block diagrams of exemplary computer architectures;

Figure 14 shown a block diagram of a system in accordance with one embodiment of the present invention;

Figure 15 is a block diagram of a first more specific exemplary system in accordance with an embodiment of the present invention;

Figure 16 is a block diagram of a second more specific exemplary system in accordance with an embodiment of the present invention;

Figure 17 is a block diagram of a System-on-a-Chip (SoC) in accordance with an embodiment of the present invention; and

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 some embodiments.

DETAILED DESCRIPTION

[0012] In the following description, numerous specific details are set forth. However, it is understood that some embodiments 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.

[0013] References in the specification to "one embodiment," "an embodiment," "an example embodiment," etc., indicate that the embodiment described may include a feature, structure, or characteristic, but every embodiment may not necessarily include the feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a feature, structure, or characteristic is described about an embodiment, it is

submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic about other embodiments if explicitly described.

[0014] Disclosed embodiments maximize execution throughput of a sparse-dense matrix multiplication (SDMM) instruction having variable precision inputs, such as a virtual neural network (VNN) matrix multiplication instruction having one sparse matrix input and one dense matrix input. Using the circuitry disclosed herein, the disclosed SDMM instructions are expected to yield a performance gain over matrix multiplication instructions having symmetric operands.

[0015] Disclosed embodiments, by not using a conventional, symmetric matrix multiplication circuit, are expected to improve a computer's SDMM instruction throughput by a factor proportional to the sparsity of the sparse input matrix. As used herein, sparsity relates to the proportion of matrix elements having zero or null values. For example, an SDMM instruction operating on a sparse matrix having 0.125 sparsity with only 1/8 of the elements having non-zero values is expected to have an 8-fold increase in throughput. For another example, an SDMM instruction operating on a sparse matrix having 0.05 sparsity with only 1/20 of the elements having non-zero values is expected to have a 20-fold increase in throughput.

[0016] In some embodiments, SDMM instructions are executed by execution circuitry having SIMD processing lanes using a grid of fused multiply-add (FMA) circuits. The SIMD lane width can differ in different embodiments. For example, SIMD lane width can include any of 16 elements, 32 elements, 64 elements, and 128 elements, and the elements can be any of 8 bits, 16 bits, 32 bits, 64 bits, and 128 bits, without limitation. SIMD lanes are used to execute an instruction in parallel on multiple data elements.

[0017] In some embodiments, multi-bank memories are used for intermediate data and result storage by the multiple SIMD lanes. For example, a SIMD execution circuitry having 8, 64-bit SIMD lanes may use an 8-bank memory.

[0018] **Figure 1** is a block diagram illustrating processing components for executing a sparse-dense matrix multiplication (SDMM) instruction, such as a virtual neural network instruction (SDMVNNI), having asymmetric inputs, according to some embodiments. A virtual neural network instruction is applied in a deep learning context and is one type of instruction that can benefit from a sparse-dense matrix multiplication. There are other types of applications that may benefit from a SDMM instruction, such as Galois Field New Instructions (GFNI). Disclosed SDMM instructions are thus not meant to be limited to VNNI instructions. As illustrated, storage 101 stores an SDMM instruction 103 to be executed.

[0019] The SDMM instruction 103 is fetched from storage 101 by fetch circuit 105. The fetched SDMM instruction 107 is decoded by decode circuit 109. For example, decode circuit 109 receives the fetched SDMM instruction 107 from fetch circuit 105. The SDMM instruction format, as described further below and with respect to **Figures 7-9**, has fields to specify an opcode, a dense output matrix, a dense source matrix, and a sparse source matrix having a

sparsity of non-zero elements being less than one. Decode circuit 109 decodes the fetched SDMM instruction 107 into one or more operations. In some embodiments, this decoding includes generating a plurality of micro-operations to be performed by execution circuitry (such as execution circuitry 117). The decode circuit 109 also decodes instruction suffixes and prefixes (if used). Execution circuitry 117 is further described and illustrated below, including at least respect to **Figures 2-6** and **Figures 11-12**, below.

[0020] In some embodiments, register renaming, register allocation, and/or scheduling circuit 113 provides functionality for one or more of: 1) renaming logical operand values to physical operand values (e.g., a register alias table in some embodiments), 2) allocating status bits and flags to the decoded instruction, and 3) scheduling the decoded SDMM instruction 111 for execution on execution circuitry 117 out of an instruction pool (e.g., using a reservation station in some embodiments).

[0021] Registers (register file) and/or memory 115 store data as operands of decoded SDMM instruction 111 to be operated on by execution circuitry 117. Exemplary register types include writemask registers, packed data registers, general purpose registers, and floating point registers, as further described and illustrated below, at least with respect to **Figure 10**.

[0022] In some embodiments, write back circuit 119 commits the result of the execution of the decoded SDMM instruction 111.

[0023] **Figure 2A** is a block diagram illustrating a data flow for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments. As shown, SDMM instruction 202 has fields to specify an opcode 204 (SDMMVNNIW), a dense output matrix 206, a dense source matrix 210, and a sparse source matrix 208. As shown, the specified sparse source matrix 212 is logically a M=3 row by K=6-column matrix having a sparsity of non-zero elements being roughly equal to 16.67%. In other words, only 16.67% of the elements of the sparse source matrix have non-zero values. Processing the SDMM instruction 202 according to embodiments disclosed herein improves the throughput of the processor by up to six times, when compared to using a conventional symmetric matrix multiplication circuit. Disclosed embodiments avoid wasting processing cycles on zero-valued elements of the specified sparse source matrix 212. In some embodiments, the sparsity of the specified sparse source matrix 212 is limited to less than 20%. In the context of virtual neural networks, and as shown, the sparse and dense source matrices can represent an activation matrix and a weights vector.

[0024] In some embodiments, the specified sparse source matrix 212 is logically a M by K matrix, but only its non-zero elements are stored in memory in a compressed sparse row (CSR) or compressed sparse column (CSC) format, which in some embodiments is prepared in advance. CSC and CSR formats are further described below, at least with reference to **Figure 4**.

[0025] In some embodiments, SDMM instruction 202 further specifies an element size (here, opcode 204 includes a "W" suffix, specifying Word-sized elements). The format of SDMM

instruction is further illustrated and described below, at least with respect to **Figures 7-9**. In some embodiments, one or more of the identified matrices are stored in registers, such as in a register file of a processor, for example as illustrated and discussed below with reference to **Figure 10**. In some embodiments, one or more of the identified matrices are stored in a memory location.

[0026] As shown, specified sparse source matrix 212 is a matrix logically having M rows (equal to 3) and K columns (equal to 6), with non-zero elements at (0,0), (3,3), and (2,5). Hence, specified sparse source matrix 212 has a sparsity of around 16.67%, and processing the instruction according to disclosed embodiments provides an up to six-times improvement in processor throughput. Specified dense source matrix 214 has K rows (equal to 6), and N columns (equal to 6). The specified dense output matrix 216 is shown as having M rows and N columns. As described herein, capital letters, M, N, and K, are used to refer to the maximal dimensions of the matrices, whereas lower-case letters, m, n, and k, are used as indices to element positions within the matrices.

[0027] In operation, as shown by data flow indicators 218, 220 and 222, for each non-zero element at row m and column k of the specified sparse source matrix 212, execution circuitry generates a product of the non-zero element and all the corresponding dense elements at row k of the specified dense source matrix 214. Execution circuitry then accumulates each generated product with previous values of a corresponding output element at row m of the specified dense output matrix 216. In some embodiments, execution circuitry writes the accumulated sum to the corresponding elements of the specified dense output matrix 216. In some embodiments, execution circuitry writes the accumulated sum to a scratchpad memory (Not shown) before writing to the dense output matrix.

[0028] **Figure 2B** is a block diagram illustrating a data flow for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments. As shown, SDMM instruction 252 has fields to specify an opcode 254 (SDMMVNNI), a dense output matrix 256, a dense source matrix 260, and a sparse source matrix 258. The specified sparse source matrix 262 is logically a M=4 by K=4 matrix, with non-zero elements in the first and third columns. Hence, the specified sparse source matrix 262 has a sparsity of 0.5. In other words, 50% of the elements of the specified sparse source matrix 262 have non-zero values. SDMM instruction 252, processed according to embodiments disclosed herein, improves the throughput of the processor by two times, when compared to using a conventional symmetric matrix multiplication circuit. Disclosed embodiments avoid wasting processing cycles on zero-valued source elements. In some embodiments, the sparsity of specified sparse source matrix 262 is limited to less than 10%. In some embodiments, one or more of the identified matrices are stored in registers, such as in a register file of a processor, for example as illustrated and discussed below with reference to **Figure 10**. In some embodiments, one or more of the specified sparse source, dense source, and dense output matrices are stored in a memory location.

[0029] As shown, specified sparse source matrix 262 is a matrix logically having M rows (equal

to 4) and K columns (equal to 4), with eight non-zero elements at column 0 and column 2. specified dense source matrix 270 has K rows (equal to 4), and N columns (equal to 4). As described herein, capital letters, M, N, and K, are used to refer to the maximal dimensions of the matrices, whereas lower-case letters, m, n, and k, are used to refer to indices of the element positions within the matrices.

[0030] In operation, according to some embodiments, for each non-zero element at row m and column k of specified sparse source matrix 262, execution circuitry generates a product of the non-zero element and each corresponding element at row k and column {0, n-1} of specified dense source matrix 264. In this embodiment, as shown at step 1 266, each of the non-zero elements of the first column of specified sparse source matrix 262 is multiplied by every element at corresponding first row of specified dense source matrix 264. As shown at step 2 268, each of the non-zero elements of the third column of the specified sparse source matrix 262 is multiplied by every element at the corresponding third row of the specified dense source matrix 264. Then, execution circuitry accumulates the products generated in step 1 266 and step 2 268 with previous values of corresponding output element at row m and column n of specified dense output matrix 270. For simplicity, here, the previous values of the output matrix are not shown, but are assumed to be zero. In some embodiments, execution circuitry writes the accumulated sums to the corresponding elements of the dense output matrix. In some embodiments, execution circuitry writes the accumulated sum to a scratchpad memory before writing to the dense output matrix.

[0031] **Figure 3** is a block diagram illustrating an execution circuit for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments. As shown, SDMM instruction 302 has fields to specify an opcode 304 (SDMMVNNI), a dense output matrix 306, a sparse source matrix 308, and a dense source matrix 310. As illustrated, specified sparse source matrix 312 has a sparsity of non-zero elements being, for illustration, around 0.07. In other words, around 7% of the elements of the specified sparse source matrix have non-zero values. Processing the SDMM instruction according to embodiments disclosed herein improves the throughput of the processor by roughly 14 times, when compared to using a conventional symmetric matrix multiplication circuit. Disclosed embodiments avoid wasting processing cycles on zero-valued source elements. In some embodiments, the sparsity of the sparse source matrix is limited to less than 10%. In some embodiments, one or more of the identified matrices are stored in registers, such as in a register file of a processor, for example as illustrated and discussed below with reference to **Figure 10**. In some embodiments, one or more of the identified matrices are stored in a memory location.

[0032] As shown, specified sparse source matrix 312 is a matrix having M rows (equal to 8) and K columns (equal to 2), with one non-zero elements at row 6 and column 1. Specified dense source matrix 314 has K rows (equal to 2), and N columns (equal to 6). The specified dense output matrix 316 is shown as having M rows (equal to 8) and N columns (equal to 6).

[0033] In operation, according to some embodiments, for each non-zero element at row m and column k of specified sparse source matrix 312, execution circuitry generates a product of the

non-zero element and each corresponding dense element at row k and column n of specified dense source matrix 314. As shown the non-zero element at element (6,1) of specified sparse source matrix 312 is multiplied by every element at corresponding row 1 of specified dense source matrix 314 using multipliers 318. Execution circuitry then generates an accumulated sum of the products generated by multipliers 318 and previous values of the corresponding elements of specified dense output matrix 316 using adders/accumulators 320. Here, the previous values of the output matrix are not shown, but, for simplicity, are assumed to be zero. In some embodiments, execution circuitry writes the accumulated sums to the corresponding elements of specified dense output matrix 316. In some embodiments, execution circuitry writes the accumulated sums to a scratchpad memory 322 before writing to specified dense output matrix 316.

[0034] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is further illustrated and discussed at least with respect to **Figures 4-6**, and **Figures 11-12**.

[0035] **Figure 4** is a block diagram illustrating an execution circuit 411 for processing a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments. As shown, SDMM instruction 402 has fields to specify an opcode 404 (SDMMVNNI), a dense output matrix 406, a sparse source matrix 408, and a dense source matrix 410. In the illustrated embodiment, the specified dense output, dense source, and sparse source matrices have dimensions $M=N=256$ and $K=512$ of 2-byte-precision entries and the specified sparse source matrix has a sparsity of 0.125.

COMPRESSED SPARSE Row / COMPRESSED SPARSE COLUMN FORMAT

[0036] In some embodiments, the specified sparse source matrix is stored in sparse format in memory 412, such that the execution circuit 411 reads the sparsely-formatted matrix, but only buffers the non-zero values in sparse buffers 414.

[0037] To avoid unnecessary memory accesses and to conserve memory space, however, some embodiments store only the non-zero elements of the sparse source matrix in memory 412 in compressed sparse row (CSR) or compressed sparse column (CSC) format. With CSR and CSC formats, only the non-zero elements of the sparse source matrix, organized in row-major format or column-major format, respectively, are stored. In some embodiments, the CSR or CSC-formatted sparse source matrices are prepared in memory 412 in advance of the operation by specialized hardware or software.

SPARSE BUFFERS 414

[0038] The specified sparse source matrix in the illustrated embodiment is logically a 256×512

matrix having 32 non-zero elements per column, i.e., at a sparsity of 0.125. The illustrated embodiment uses CSC format to store the sparse source matrix in the sparse buffers 414 and each column of the sparse source matrix is partitioned into 8 banks based on the row index M . In operation, 32 non-zero elements per column of the sparse source matrix are stored in 8 banks of the sparse buffers 414, equally distributed among them in the ideal case. In some embodiments, each of the sparse buffer entries uses up to five bytes, which includes two bytes to store a data value, and up to 3 bytes to specify a matrix position of the element within the specified sparse source matrix. In some embodiments, the specified sparse source matrix logically has $M=16$ rows and $K=16$ columns, and the 1-byte matrix position specifies an offset of the element within the 256 elements. In some embodiments, the specified sparse source matrix has $M=16$ rows and $K=16$ columns, and the 1-byte matrix position includes a nibble to specify the column and a nibble to specify the row in which the element is located.

MULTIPLIER ARRAY 416

[0039] In the illustrated embodiment, the multiplier array is of size 8×32 where the 8 rows of the multiplier array are connected to the 8 banks of the sparse buffers 414 providing the 8 multiplier values per cycle needed for multiplication. 32 elements of row k of the specified dense source matrix which form the multiplicand are broadcasted and multiplied across 8 elements of the sparse buffers 414, thus performing 256 multiplications per cycle. The 32-element partial product generated per bank is then accumulated in accumulator array 418. In some embodiments, the 32 elements of the specified dense source matrix are buffered in registers (not shown) before the multiplications. In some embodiments, as shown, the 32 elements are fed into multiplier array 416 as they are loaded from memory 412. In some embodiments, multiplier array 416 comprises a grid of fused-multiply-add (FMA) hardware units.

ACCUMULATOR ARRAY 418

[0040] As shown, accumulator array 418 includes 256 accumulators divided into 8 banks, each connected to the corresponding bank of sparse buffers 414 and multiplier array 416. In operations, accumulators (32 per bank) in accumulator array 418 accumulate the products generated by multiplier array 416 with previous values of corresponding elements of the dense output matrix specified by the dense output matrix 406 field of SDMM instruction 402.

ROUNDING AND SATURATION CIRCUIT 420

[0041] In some embodiments, the products generated by multiplier array 416 and accumulated by accumulator array 418 are high-precision intermediate results represented by at least twice as many bits as used by the data elements of the specified matrices. In some embodiments,

rounding and saturation circuit 420 saturates the intermediate results to a predefined maximum and rounds them to fit within the number of bits of elements of the dense output matrix specified by the dense output matrix 406 field of the SDMM instruction 402, which here is 16 bits.

[0042] In the case of floating point arithmetic, rounding and saturation circuit 420 may round the intermediate results according to the IEEE 754 floating point standard, established in 1985 and updated in 2008 by the Institute of Electrical and Electronics Engineers. The IEEE 754 floating point standard defines rounding rules to be applied, including round to nearest with ties to even, round to nearest with ties away from zero, toward zero, toward positive infinity, and toward negative infinity. In some embodiments, rounding and saturation circuit 420 includes a software-accessible rounding control register (not shown) to specify the rounding rule to apply.

[0043] In some embodiments, each of the FMA hardware units in multiplier array 416 performs the rounding by itself. In some embodiments, each of the FMA hardware units in multiplier array 416 checks for saturation and performs the saturating itself.

SCRATCHPAD 422

[0044] As shown, execution circuit 411 includes scratchpad 422 to store intermediate execution results. In some embodiments, as here, scratchpad 422 is a 32 kB memory, partitioned into eight (8) banks, each connected one to one with the corresponding row of the multiplier array and bank of the accumulator. In some embodiments, all of the banks of scratchpad 422 communicate with corresponding banks of accumulator array 418 in parallel, yielding a high-bandwidth connection. Using the CSC format for multiplication, whereby only non-zero elements of the specified sparse source matrix are stored in sparse buffers 414 and supplied to multiplier array 416 obviates the need to use an expensive gather scatter circuit to gather and supply non-zero data to multiplier array 416 and accumulator array 418.

[0045] In some embodiments, as here, execution circuit 411 utilizes one or more single-instruction, multiple data (SIMD) processing lanes, for example, 8 lanes, to perform a same operation on multiple data elements at the same time. In some embodiments, a SIMD processing lane has a lane width of 32 elements, and 8 SIMD processing lanes are used to perform 256 operations on 256 elements of data.

[0046] In some embodiments, two or more SIMD processing lanes operate concurrently, and in parallel. The number of lanes in a SIMD processor, as well as the number of elements assigned to each lane, can vary, without limitation. According to some embodiments, a SIMD processing lane is implemented as having a lane width being one 8 elements, 16 elements, or 32 element s, with element widths of 8-bits, 16 bits, 32 bits, or 64-bits, without limitation.

[0047] In some embodiments, execution circuit 411 performs the SDMM instruction 402 by performing multiply-accumulate operations using fused multiply-add (FMA) hardware units to

generate the products of each non-zero element of the specified sparse source matrix and each of the elements in a corresponding row of the specified dense source matrix, and to accumulate products with the previous values of corresponding elements of the dense output matrix specified by the dense output matrix 406 field of SDMM instruction 402.

[0048] As used herein, the term "corresponding" has a different interpretation based on its context. In the context of generating the products, the corresponding elements of the dense source matrix specified by the dense source matrix 410 field that correspond to each non-zero element (m, k) of the sparse source matrix specified by the sparse source matrix 408 field of SDMM instruction 402 are the corresponding elements, (k, n) , in a corresponding row, k of the dense source matrix specified by the dense source matrix 410 field of the SDMM instruction 402. In the context of accumulating the products with previous contents of the dense output matrix specified by the dense output matrix 406 field of SDMM instruction 402. The corresponding elements of the specified dense output matrix are those at locations (m, n) .

[0049] Accordingly, execution circuit 411, by executing an SDMM instruction 402 specifying a sparse source matrix 408 having a sparsity of 0.125, improves the throughput of the processor in which it is incorporated by 8 times, when compared to using a symmetric matrix multiplication circuit.

[0050] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is further illustrated and discussed, at least with respect to **Figures 3-6**, and

Figures 11-12.

[0051] **Figure 5** is pseudocode illustrating operation of execution circuitry to process instructions calling for matrix multiplication, here, a virtual neural network instruction (VNNI). For illustrative purposes, pseudocode 502 and 504 illustrate implementation of a VNNI matrix multiplication instruction on symmetric source operands. As described herein, capital letters, M , N , and K , are used to refer to the maximal dimensions of the matrices, whereas lower-case letters, m , n , and k , are used to refer to indices of element positions within the matrices. Pseudocode 502 and pseudocode 504 both illustrate executing a VNNI instruction on a M -row by K -column ($M \times K$) source matrix A and a K -row by N -column ($K \times N$) source matrix B to generate results of a M -row by N -column ($M \times N$) output matrix C . Pseudocodes 506 and 508, on the other hand, illustrate executing the SDMM instruction according to disclosed embodiments, wherein only the non-zero elements of sparse source matrix A are processed, thereby increasing processor throughput in proportion to the sparsity of source matrix A .

[0052] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is illustrated and discussed at least with respect to **Figures 3-6**, and **Figures 11-12.**

[0053] **Figure 6** is a process flow diagram illustrating execution of a sparse-dense matrix

multiplication (SDMM) instruction by a processor, according to some embodiments. At 602, the processor fetches, from code storage using fetch circuitry, the sparse-dense matrix multiplication instruction having fields to specify an opcode, a dense output matrix, a dense source matrix, and a sparse source matrix having a sparsity of non-zero elements, the sparsity being less than one. The SDMM instruction fetched at 602 may be referred to as an asymmetric SDMM instruction, insofar as one source is a sparse matrix and the other source is a dense matrix. At 604, the processor decodes, by decode circuitry, the fetched SDMM instruction. At 606, the processor optionally schedules execution of the decoded SDMM instruction by a SIMD execution circuit. Operation 606 is optional, as indicated by its dashed border, insofar as scheduling execution of the decoded instruction may occur at a different time, or not at all. At 608, the processor executes, by execution circuitry, the decoded SDMM instruction to, for each non-zero element at row M and column K of the sparse source matrix, generate a product of the non-zero element and each corresponding dense element at row K and column N of the dense source matrix; and accumulates each generated product with a previous value of a corresponding output element at row M and column N of the dense output matrix. At 610, the processor optionally commits or retires the executed SDMM instruction. Operation 610 is optional, as indicated by its dashed border, insofar as it may occur at a different time, or not at all.

[0054] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is illustrated and discussed at least with respect to **Figures 3-6**, and **Figures 11-12**.

[0055] **Figure 7** is an exemplary format of a sparse-dense matrix multiplication (SDMM) instruction, according to some embodiments. As shown, SDMM instruction 700 includes opcode 702 (SDMMVNNI*), and fields to specify a dense output matrix 704, dense source matrix 706, and sparse source matrix 708. SDMM instruction 700 further includes optional fields to specify element size 710, and M, N, and K dimensions, 712, 714, and 716. In some embodiments, one or more of optional dimensions 712, 714, and 716 is specified by a software programmable model specific register (MSR), which may be a predetermined MSR. As shown in **Figure 7**, the optional element size 710 may be specified by part of the opcode, and the M, N, and K dimensions, 712, 714, and 716 may be specified by a MSR, as part of the opcode, or a combination thereof. Opcode 702 is shown as including an asterisk to indicate that it may optionally include additional prefixes or suffixes to specify instruction behaviors. For example, opcode 702 may include a suffix, such as "B," "W," "D," or "Q," to specify an element size of eight, sixteen, thirty-two, and sixty-four bits, respectively. If the SDMM instruction 700 does not specify any of the optional parameters, predetermined default values are used. The format of the SDMM instruction is further illustrated and described below with respect to **Figure 8A**,

Figure 8B, and **Figures 9A-D**.

PREPPING THE COMPRESSED SPARSE SOURCE MATRIX AHEAD-OF-TIME

[0056] Referring again to **Figure 2B**, the specified compressed sparse source matrix (CSR or CSC) in some embodiments is prepared in advance. CSC and CSR formats are further illustrated and described, at least with reference to **Figure 4**.

[0057] In some embodiments, SDMM instruction 700 is used to cause the processor to prepare a sparse source matrix in compressed format (either CSR, or CSC, as further illustrated and described with respect to **Figure 4**). In some embodiments, opcode 702 may include a prefix or a suffix, such as "PREP," to instruct the processor to prepare a compressed sparse source matrix. In such embodiments, the dense output matrix 704 field specifies a memory address at which to store the compressed sparse source matrix, and sparse source matrix 708 specifies a memory location at which a sparse data set is stored, the sparse data set comprising a large block of memory having a sparsity of valid elements, the validity being determined by the data values themselves (e.g., invalid being null, zero, or below a threshold value), or by a control field appended to and indicating validity of each element. The size of the sparse data set is specified in such embodiments by the dense source matrix 706 field.

[0058] In one example, SDMM instruction 700 includes an opcode 702 having a "PREP" suffix, specifies a dense output matrix 704 where to store the condensed sparse source matrix, uses the sparse source matrix 708 field to specify where a sparse data set comprising tens, hundreds, thousands, millions, or billions of data elements, with invalid elements having null values, and uses the dense source matrix 706 field to specify a size of the sparse data set. In response, the processor loads the data elements from the specified sparse source matrix 708 location, uses the dense source matrix 706 field to determine the size of the sparse data set, determines whether each data element is valid, and writes the valid elements, in compressed format (CSR or CSC) to the specified dense output matrix 704 location, which in some embodiments is a cache-line-aligned location. The processor thus packs the valid elements of the sparse data set in compressed format (CSR or CSC) to the specified output matrix, which can serve as a sparse source matrix for a subsequent SDMM instruction.

[0059] In another example, SDMM instruction 700 includes an opcode 702 having a "PREP" suffix, specifies a dense output matrix 704 where to store the condensed sparse source matrix, uses the sparse source matrix 708 field to specify where a sparse data set comprising tens, hundreds, thousands, millions, or billions of data elements, with each element including a control fields including at least a valid bit, and uses the dense source matrix 706 field to specify a size (i.e., number of elements) of the sparse data set. In response, the processor loads the data elements from the specified sparse source matrix 708 location, uses the dense source matrix 706 field to determine the size of the sparse data set, determines whether each data element is valid, and writes the valid elements, in compressed format (CSR or CSC) to the specified dense output matrix 704 location, which in some embodiments is a cache-line-aligned location. The processor thus packs the valid elements of the sparse data set in compressed format (CSR or CSC) to the specified output matrix, which can serve as a sparse source matrix for a subsequent SDMM instruction.

[0060] In some embodiments, a processor responding to an instance of SDMM instruction 700 having an opcode with a "PREP" suffix operates in the background by opportunistically scheduling execution of its requisite loads and stores only when a processor activity level is below a threshold value. In some embodiments, a processor responding to an instance of SDMM instruction 700 having an opcode with a "PREP" suffix generates an interrupt upon completion of compressing the specified sparse data set.

INSTRUCTION SETS

[0061] An instruction set may include one or more instruction formats. A given instruction format may define various fields (e.g., number of bits, location of bits) to specify, among other things, the operation to be performed (e.g., opcode) and the operand(s) on which that operation is to be performed and/or other data field(s) (e.g., mask). Some instruction formats are further broken down though the definition of instruction templates (or subformats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in each one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands. A set of SIMD extensions referred to as the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developer's Manual, September 2014; and see Intel® Advanced Vector Extensions Programming Reference, October 2014).

EXEMPLARY INSTRUCTION FORMATS

[0062] Embodiments of the instruction(s) described herein may be embodied in different formats. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

Generic Vector Friendly Instruction Format

[0063] 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.

[0064] Figures 8A-8B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to some embodiments. **Figure 8A** is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to some embodiments; while **Figure 8B** is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to some embodiments. 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.

[0065] While some embodiments will be described in which the vector friendly instruction format supports the following: a 64 byte vector operand length (or size) with 32 bit (4 byte) or 64 bit (8 byte) data element widths (or sizes) (and thus, a 64 byte vector consists of either 16 doubleword-size elements or alternatively, 8 quadword-size elements); a 64 byte vector operand length (or size) with 16 bit (2 byte) or 8 bit (1 byte) data element widths (or sizes); a 32 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); and a 16 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); alternative embodiments may support more, less and/or different vector operand sizes (e.g., 256 byte vector operands) with more, less, or different data element widths (e.g., 128 bit (16 byte) data element widths).

[0066] The class A instruction templates in **Figure 8A** include: 1) within the no memory access 805 instruction templates there is shown a no memory access, full round control type operation 810 instruction template and a no memory access, data transform type operation 815 instruction template; and 2) within the memory access 820 instruction templates there is shown a memory access, temporal 825 instruction template and a memory access, non-temporal 830 instruction template. The class B instruction templates in **Figure 8B** include: 1) within the no memory access 805 instruction templates there is shown a no memory access, write mask control, partial round control type operation 812 instruction template and a no memory access, write mask control, vsize type operation 817 instruction template; and 2) within the memory access 820 instruction templates there is shown a memory access, write mask control 827 instruction template.

[0067] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in **Figures 8A-8B**.

[0068] Format field 840 - a specific value (an instruction format identifier value) in this field

uniquely identifies the vector friendly instruction format, and thus occurrences of instructions in the vector friendly instruction format in instruction streams. As such, this field is optional in the sense that it is not needed for an instruction set that has only the generic vector friendly instruction format.

[0069] Base operation field 842 - its content distinguishes different base operations.

[0070] 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 enough 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 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).

[0071] 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.

[0072] 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.

[0073] Scale field 860 - its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base}$).

[0074] Displacement Field 862A- its content is used as part of memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$).

[0075] Displacement Factor Field 862B (note that the juxtaposition of displacement field 862A directly over displacement factor field 862B indicates one or the other is used) - its content is used as part of address generation; it specifies a displacement factor that is to be scaled by the size of a memory access (N) - where N is the number of bytes in the memory access (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{scaled displacement}$). Redundant low-

order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field 874 (described later herein) and the data manipulation field 854C. The displacement field 862A and the displacement factor field 862B are optional in the sense that they are not used for the no memory access 805 instruction templates and/or different embodiments may implement only one or none of the two.

[0076] Data element width field 864 - its content distinguishes which one of many data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

[0077] Write mask field 870 - its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other one embodiment, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in one embodiment, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field 870 allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While some embodiments are described in which the write mask field's 870 content selects one of many write mask registers that contains the write mask to be used (and thus the write mask field's 870 content indirectly identifies that masking to be performed), alternative embodiments instead or additional allow the mask write field's 870 content to directly specify the masking to be performed.

[0078] 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.

[0079] Class field 868 - its content distinguishes between different classes of instructions. With reference to Figures 8A-B, the contents of this field select between class A and class B instructions. In Figures 8A-B, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 868A and class B 868B for the class field 868 respectively in

Figures 8A-B).

INSTRUCTION TEMPLATES OF CLASS A

[0080] In the case of the non-memory access 805 instruction templates of class A, the alpha field 852 is interpreted as an RS field 852A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 852A.1 and data transform 852A.2 are respectively specified for the no memory access, round type operation 810 and the no memory access, data transform type operation 815 instruction templates), while the beta field 854 distinguishes which of the operations of the specified type is to be performed. In the no memory access 805 instruction templates, the scale field 860, the displacement field 862A, and the displacement scale field 862B are not present.

NO-MEMORY ACCESS INSTRUCTION TEMPLATES - FULL ROUND CONTROL TYPE OPERATION

[0081] 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 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).

[0082] SAE field 856 - its content distinguishes if 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.

[0083] Round operation control field 858 - its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 858 allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes a control register for specifying rounding modes, the round operation control field's 850 content overrides that register value.

No MEMORY ACCESS INSTRUCTION TEMPLATES - DATA TRANSFORM TYPE OPERATION

[0084] 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 many data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

[0085] In the case of a memory access 820 instruction template of class A, the alpha field 852 is interpreted as an eviction hint field 852B, whose content distinguishes which one of the eviction hints is to be used (in **Figure 8A**, temporal 852B.1 and non-temporal 852B.2 are respectively specified for the memory access, temporal 825 instruction template and the memory access, non-temporal 830 instruction template), while the beta field 854 is interpreted as a data manipulation field 854C, whose content distinguishes which one of a number of data manipulation operations (also known as primitives) is to be performed (e.g., no manipulation; broadcast; up conversion of a source; and down conversion of a destination). The memory access 820 instruction templates include the scale field 860, and optionally the displacement field 862A or the displacement scale field 862B.

[0086] Vector memory instructions perform vector loads from and vector stores to memory, with conversion support. As with regular vector instructions, vector memory instructions transfer data from/to memory in a data element-wise fashion, with the elements that are transferred is dictated by the contents of the vector mask that is selected as the write mask.

MEMORY ACCESS INSTRUCTION TEMPLATES-TEMPORAL

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

MEMORY ACCESS INSTRUCTION TEMPLATES - NON-TEMPORAL

[0088] Non-temporal data is data unlikely to be reused soon enough to benefit from caching in the 1st-level cache and should be given priority for eviction. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

INSTRUCTION TEMPLATES OF CLASS B

[0089] 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.

[0090] 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.

[0091] 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).

[0092] 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.

[0093] 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 many data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

[0094] 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 if 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.

[0095] About 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 these fields, the full opcode field 874 includes less than these fields in embodiments that do not support all of them. The full opcode field 874 provides the operation code (opcode).

[0096] 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.

[0097] 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.

[0098] The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments, 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. Programs written in a high level language would be put (e.g., just in time compiled or statically compiled) into an variety of different executable forms, including: 1) a form having only instructions of the class(es) supported by the target processor for execution; or 2) a form having alternative routines written using different combinations of the instructions of all classes and having control flow code that selects the routines to execute based on the instructions supported by the processor which is currently executing the code.

EXEMPLARY SPECIFIC VECTOR FRIENDLY INSTRUCTION FORMAT

[0099] **Figure 9A** is a block diagram illustrating an exemplary specific vector friendly instruction format according to some embodiments. **Figure 9A** 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 9A** map are illustrated.

[0100] It should be understood that, although some embodiments 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).

[0101] The generic vector friendly instruction format 800 includes the following fields listed below in the order illustrated in **Figure 9A**.

[0102] EVEX Prefix (Bytes 0-3) 902 - is encoded in a four-byte form.

[0103] 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).

[0104] The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

[0105] 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.

[0106] 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 some embodiments 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.

[0107] Opcode map field 915 (EVEX byte 1, bits [3:0] - mmmm) - its content encodes an implied leading opcode byte (0F, 0F 38, or 0F 3).

[0108] 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).

[0109] EVEX.vvvv 920 (EVEX Byte 2, bits [6:3]-vvvv)- the role of EVEX.vvvv may include the following: 1) EVEX.vvvv encodes the first source register operand, specified in inverted (1s complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in 1s complement form for certain vector shifts; or 3) EVEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. Thus, EVEX.vvvv field 920 encodes the 4 low-order bits of the first source register specifier stored in inverted (1s complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

[0110] 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.

[0111] 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.

[0112] Alpha field 852 (EVEX byte 3, bit [7] - EH; also known as EVEX.EH, EVEX.rs, EVEX.RL, EVEX.write mask control, and EVEX.N; also illustrated with α) - as previously described, this field is context specific.

[0113] Beta field 854 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s₂₋₀, EVEX.r₂₋₀, EVEX.rr1, EVEX.LL0, EVEX.LLB; also illustrated with $\beta\beta\beta$) as previously described, this field is context specific.

[0114] REX' field 810 - this is the remainder of the REX' field and is the EVEX.V' bit field (EVEX Byte 3, bit [3] - V') that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V', EVEX.vvvv.

[0115] Write mask field 870 (EVEX byte 3, bits [2:0]-kkk) - its content specifies the index of a register in the write mask registers as previously described. In one embodiment of the invention, the specific value EVEX.kkk=000 has a special behavior implying no write mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

[0116] Real Opcode Field 930 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

[0117] MOD R/M Field 940 (Byte 5) includes MOD field 942, Reg field 944, and R/M field 946. As previously described, the MOD field's 942 content distinguishes between memory access and non-memory access operations. The role of Reg field 944 can be summarized to two situations: encoding either the destination register operand or a source register operand, or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 946 may include the following: encoding the instruction operand that references a memory address, or encoding either the destination register operand or a source register operand.

[0118] 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.

[0119] 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.

[0120] Displacement factor field 862B (Byte 7) - when MOD field 942 contains 01, byte 7 is the displacement factor field 862B. The location of this field is that same as that of the legacy x86 instruction set 8-bit displacement (disp8), which works at byte granularity. Since disp8 is sign extended, it can only address between -128 and 127 bytes offsets; in terms of 64 byte cache lines, disp8 uses 8 bits that can be set to only four really useful values -128, -64, 0, and 64; since a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to disp8 and disp32, the displacement factor field 862B is a reinterpretation of disp8; when using displacement factor field 862B, the actual displacement is determined by the content of the displacement factor field multiplied by the size of the memory operand access (N). This type of displacement is referred to as disp8*N. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement is based on the assumption that the effective displacement is multiple of the granularity of the memory access, and hence, the redundant low-order bits of the address offset do not need to be encoded. In other words, the displacement factor field 862B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 862B is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to disp8*N. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset). Immediate field 872 operates as previously described.

FULL OPCODE FIELD

[0121] **Figure 9B** is a block diagram illustrating the fields of the specific vector friendly instruction format 900 that make up the full opcode field 874 according to one embodiment of the invention. Specifically, the full opcode field 874 includes the format field 840, the base operation field 842, and the data element width (W) field 864. The base operation field 842 includes the prefix encoding field 925, the opcode map field 915, and the real opcode field 930.

REGISTER INDEX FIELD

[0122] **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 VVV field 920, xxx field 954, and the bbb field 956.

AUGMENTATION OPERATION FIELD

[0123] **Figure 9D** is a block diagram illustrating the fields of the specific vector friendly instruction format that makes up the augmentation operation field 850 according to one embodiment of the invention. When the class (U) field 868 contains 0, it signifies EVEX.U0 (class A 868A); when it contains 1, it signifies EVEX.U1 (class B 868B). When U=0 and the MOD field 942 contains 11 (signifying a no memory access operation), the alpha field 852 (EVEX byte 3, bit [7] - EH) is interpreted as the rs field 852A. When the rs field 852A contains a 1 (round 852A.1), the beta field 854 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the round control field 854A. The round control field 854A includes a one bit SAE field 856 and a two bit round operation field 858. When the rs field 852A contains a 0 (data transform 852A.2), the beta field 854 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data transform field 854B. When U=0 and the MOD field 942 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 852 (EVEX byte 3, bit [7] - EH) is interpreted as the eviction hint (EH) field 852B and the beta field 854 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data manipulation field 854C.

[0124] 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]- S₀) 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]- S₂₋₁) is interpreted as the round operation field 859A, while when the RL field 857A contains a 0 (VSIZE 857.A2) the rest of the beta field 854 (EVEX byte 3, bit [6-5]- S₂₋₁) is interpreted as the vector length field 859B (EVEX byte 3, bit [6-5]- L₁₋₀). When U=1 and the

MOD field 942 contains 00, 01, or 10 (signifying a memory access operation), the beta field 854 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the vector length field 859B (EVEX byte 3, bit [6-5]- L₁₋₀) and the broadcast field 857B (EVEX byte 3, bit [4]- B).

Exemplary Register Architecture

[0125] **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

[0126] In other words, the vector length field 859B selects between a maximum length and one or more other shorter lengths, where each such shorter length is half the length of the preceding length; and instructions templates without the vector length field 859B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 900 operate on packed or scalar single/double-precision floating point data and packed or scalar integer data. Scalar operations are operations performed on the lowest order data element position in an zmm/ymm/xmm register; the higher order data element positions are either left the same as they were prior to the instruction or zeroed depending on the embodiment.

[0127] Write mask registers 1015 - in the embodiment illustrated, there are 8 write mask registers (k0 through k7), each 64 bits in size. In an alternate embodiment, the write mask registers 1015 are 16 bits in size. As previously described, in one embodiment of the invention, the vector mask register k0 cannot be used as a write mask; when the encoding that would normally indicate k0 is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

[0128] General-purpose registers 1025 - in the embodiment illustrated, there are sixteen 64-bit general-purpose registers that are used along with the existing x86 addressing modes to address memory operands. These registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

[0129] Scalar floating point stack register file (x87 stack) 1045, on which is aliased the MMX packed integer flat register file 1050 - in the embodiment illustrated, the x87 stack is an eight-element stack used to perform scalar floating-point operations on 32/64/80-bit floating point data using the x87 instruction set extension; while the MMX registers are used to perform operations on 64-bit packed integer data, as well as to hold operands for some operations performed between the MMX and XMM registers.

[0130] Alternative embodiments may use wider or narrower registers. Additionally, alternative embodiments may use more, less, or different register files and registers.

EXEMPLARY CORE ARCHITECTURES, PROCESSORS, AND COMPUTER ARCHITECTURES

[0131] Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

EXEMPLARY CORE ARCHITECTURES

IN-ORDER AND OUT-OF-ORDER CORE BLOCK DIAGRAM

[0132] **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 some embodiments. **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 some embodiments. The solid lined boxes in Figures 11A-B illustrate the in-order pipeline and in-order core, while the optional addition of the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline and core. Given that the in-order aspect is a subset of the out-of-order aspect, the out-of-order aspect will be described.

[0133] In **Figure 11A**, a processor pipeline 1100 includes a fetch stage 1102, a length decode stage 1104, a decode stage 1106, an allocation stage 1108, a renaming stage 1110, a scheduling (also known as a dispatch or issue) stage 1112, a register read/memory read stage 1114, an execute stage 1116, a write back/memory write stage 1118, an exception handling stage 1122, and a commit stage 1124.

[0134] **Figure 11B** shows processor core 1190 including a front end unit 1130 coupled to an execution engine unit 1150, and both are coupled to a memory unit 1170. The core 1190 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 1190 may be a special-purpose core, such as, for example, a network or communication core, compression engine, coprocessor core, general purpose computing graphics processing unit (GPGPU) core, graphics core, or the like.

[0135] The front end unit 1130 includes a branch prediction unit 1132 coupled to an instruction cache unit 1134, which is coupled to an instruction translation lookaside buffer (TLB) 1136, which is coupled to an instruction fetch unit 1138, which is coupled to a decode unit 1140. The decode unit 1140 (or decoder) may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decode unit 1140 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. In one embodiment, the core 1190 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 1140 or otherwise within the front end unit 1130). The decode unit 1140 is coupled to a rename/allocator unit 1152 in the execution engine unit 1150.

[0136] The execution engine unit 1150 includes the rename/allocator unit 1152 coupled to a retirement unit 1154 and a set of one or more scheduler unit(s) 1156. The scheduler unit(s) 1156 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 1156 is coupled to the physical register file(s) unit(s) 1158. Each of the physical register file(s) units 1158 represents one or more physical

register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. In one embodiment, the physical register file(s) unit 1158 comprises a vector registers unit, a write mask registers unit, and a scalar registers unit. These register units may provide architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) unit(s) 1158 is overlapped by the retirement unit 1154 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement unit 1154 and the physical register file(s) unit(s) 1158 are coupled to the execution cluster(s) 1160. The execution cluster(s) 1160 includes a set of one or more execution units 1162 and a set of one or more memory access units 1164. The execution units 1162 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. The scheduler unit(s) 1156, physical register file(s) unit(s) 1158, and execution cluster(s) 1160 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster - and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) 1164). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

[0137] The set of memory access units 1164 is coupled to the memory unit 1170, which includes a data TLB unit 1172 coupled to a data cache unit 1174 coupled to a level 2 (L2) cache unit 1176. In one exemplary embodiment, the memory access units 1164 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 1172 in the memory unit 1170. The instruction cache unit 1134 is further coupled to a level 2 (L2) cache unit 1176 in the memory unit 1170. The L2 cache unit 1176 is coupled to one or more other levels of cache and eventually to a main memory.

[0138] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 1100 as follows: 1) the instruction fetch 1138 performs the fetch and length decoding stages 1102 and 1104; 2) the decode unit 1140 performs the decode stage 1106; 3) the rename/allocator unit 1152 performs the allocation stage 1108 and renaming stage 1110; 4) the scheduler unit(s) 1156 performs the schedule stage 1112; 5) the physical register file(s) unit(s) 1158 and the memory unit 1170 perform the register read/memory read stage 1114; the execution cluster 1160 perform the execute stage 1116; 6) the memory unit 1170 and the physical register file(s) unit(s) 1158 perform the write

back/memory write stage 1118; 7) various units may be involved in the exception handling stage 1122; and 8) the retirement unit 1154 and the physical register file(s) unit(s) 1158 perform the commit stage 1124.

[0139] The core 1190 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 1190 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[0140] 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).

[0141] While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor also includes separate instruction and data cache units 1134/1174 and a shared L2 cache unit 1176, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

SPECIFIC EXEMPLARY IN-ORDER CORE ARCHITECTURE

[0142] Figures 12A-B illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip. The logic blocks communicate through a high-bandwidth interconnect network (e.g., a ring network) with some fixed function logic, memory I/O interfaces, and other necessary I/O logic, depending on the application.

[0143] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is illustrated and discussed at least with respect to **Figures 3-6**, and **Figures 11-12**.

[0144] **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 some embodiments. In one embodiment, an instruction decoder 1200 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 1206 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 1208 and a vector unit 1210 use separate register sets (respectively, scalar registers 1212 and vector registers 1214) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 1206, alternative embodiments 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).

[0145] The local subset of the L2 cache 1204 is part of a global L2 cache that is divided into separate local subsets, one per processor core. Each processor core has a direct access path to its own local subset of the L2 cache 1204. Data read by a processor core is stored in its L2 cache subset 1204 and can be accessed quickly, in parallel with other processor cores accessing their own local L2 cache subsets. Data written by a processor core is stored in its own L2 cache subset 1204 and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data. The ring network is bi-directional to allow agents such as processor cores, L2 caches and other logic blocks to communicate with each other within the chip. Each ring data-path is 1012-bits wide per direction.

[0146] **Figure 12B** is an expanded view of part of the processor core in **Figure 12A** according to some embodiments. **Figure 12B** includes an L1 data cache 1206A part of the L1 cache 1204, as well as more detail regarding the vector unit 1210 and the vector registers 1214. Specifically, the vector unit 1210 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 1228), which executes one or more of integer, single-precision float, and double-precision float instructions. The VPU supports swizzling the register inputs with swizzle unit 1220, numeric conversion with numeric convert units 1222A-B, and replication with replication unit 1224 on the memory input. Write mask registers 1226 allow predicating resulting vector writes.

[0147] Execution circuitry to execute the SDMM instruction according to disclosed embodiments is illustrated and discussed at least with respect to **Figures 3-6**, and **Figures 11-12**.

[0148] **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 some embodiments. 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.

[0149] 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 many process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0150] 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 (integrated graphics logic 1308 is an example of and is also referred to herein as special purpose logic), 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.

[0151] In some embodiments, one or more of the cores 1302A-N are capable of multithreading. 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.

[0152] 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 can execute only a subset of that instruction set or a different instruction set.

EXEMPLARY COMPUTER ARCHITECTURES

[0153] **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.

[0154] 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 couples input/output (I/O) devices 1460 to the GMCH 1490. Alternatively, one or both memory and graphics controllers are integrated within the processor (as described herein), the memory 1440 and the coprocessor 1445 are coupled directly to the processor 1410, and the controller hub 1420 in a single chip with the IOH 1450.

[0155] 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.

[0156] 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.

[0157] 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.

[0158] There can be a variety of differences between the physical resources 1410, 1415 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

[0159] In one embodiment, the processor 1410 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 1410 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 1445. Accordingly, the processor 1410 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 1445. Coprocessor(s) 1445 accept and execute the received coprocessor instructions.

[0160] Referring now to **Figure 15**, shown is a block diagram of a first more specific exemplary system 1500 in accordance with an embodiment of the present invention. As shown in **Figure**

15, multiprocessor system 1500 is a point-to-point interconnect system, and includes a first processor 1570 and a second processor 1580 coupled via a point-to-point interconnect 1550. Each of processors 1570 and 1580 may be some version of the processor 1300. In one embodiment of the invention, processors 1570 and 1580 are respectively processors 1410 and 1415, while coprocessor 1538 is coprocessor 1445. In another embodiment, processors 1570 and 1580 are respectively processor 1410 coprocessor 1445.

[0161] 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.

[0162] Processors 1570, 1580 may each exchange information with a chipset 1590 via individual P-P interfaces 1552, 1554 using point to point interface circuits 1576, 1594, 1586, 1598. Chipset 1590 may optionally exchange information with the coprocessor 1538 via a high-performance interface 1592. In one embodiment, the coprocessor 1538 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

[0163] 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.

[0164] 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.

[0165] 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 programable 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.

[0166] 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** to avoid obscuring other aspects of **Figure 16**.

[0167] **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.

[0168] 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 1302A-N, which include cache units 1304A-N, and shared cache unit(s) 1306; a system agent unit 1310; a bus controller unit(s) 1316; an integrated memory controller unit(s) 1314; a set or one or more coprocessors 1720 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; a 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.

[0169] Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Some embodiments 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.

[0170] 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.

[0171] 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 programming language. In any case, the language may be a compiled or interpreted language.

[0172] 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 make the logic or processor.

[0173] 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.

[0174] Accordingly, some embodiments also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

EMULATION (INCLUDING BINARY TRANSLATION, CODE MORPHING, ETC.)

[0175] In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

[0176] **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 some embodiments. 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 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 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 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 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.

REFERENCES CITED IN THE DESCRIPTION

Cited references

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US2018046900A1 [0005]

Non-patent literature cited in the description

- **MORAD AMIR et al.**Efficient Dense and Sparse Matrix Multiplication on GP-SIMD, [\[0006\]](#)
- **YAVITS L. et al.**Sparse Matrix Multiplication on an Associative Processor, [\[0007\]](#)
- **KOANANTAKOOL PENPORN et al.**Communication-Avoiding Parallel Sparse-Dense Matrix-Matrix Multiplication, [\[0008\]](#)
- **WU TAO et al.**An efficient sparse-dense matrix multiplication on a multicore system, [\[0009\]](#)
- Architectures Software Developer's Manual20140900 [\[0061\]](#)
- Advanced Vector Extensions Programming Reference20141000 [\[0061\]](#)

Patentkrav

1. Processor (1300), der omfatter:
en cache (1306) til at lagre data
- 5 mindst én kerne (1302), der er koblet til cachen (1306), idet den mindst ene kerne (1302) omfatter:
udførelseskredsløb (411, 1150) til at udføre multiplikations-
akkumuleringsoperationer med en første kildematrix og en anden
kildematrix for at generere en resultatmatrix, der reagerer på
- 10 en matrixmultiplikationsinstruktion, idet matrixmultiplikationsinstruktionen omfatter en flerhed af
felter, herunder et første felt til at angive en opcode, et
andet felt til at identificere resultatmatricen, et tredje felt
til at identificere den første kildematrix og et fjerde felt til
- 15 at identificere den anden kildematrix, hvor den første kildematrix er en sparsom matrix med ikke-nul-dataelementer, der befinder sig ved visse positioner, og hvor den anden kildematrix omfatter en tæt matrix, hvor én eller flere af de identificerede
matricer er lagret i en registerfil i processoren (1300),
- 20 hvor kun den første kildematrix er lagret i et komprimeret format, der identificerer positionerne af ikke-nul-dataelementerne i den første kildematrix, idet udførelseskredsløbet (411, 1150) ydermere omfatter:
en flerhed af multiplikations-akkumuleringskredsløb (416, 418)
- 25 til at udføre en flerhed af multiplikation-additionsoperationer for at multiplicere ikke-nul-dataelementerne i den første kildematrix med tilsvarende dataelementer i den anden kildematrix, der identificeres på basis af positionerne i komprimeret format, for at generere en flerhed af produkter,
- 30 addere flerheden af produkter til akkumulerede værdier for at generere dataelementer i resultatmatricen og udføre afrunding på hver af de akkumulerede værdier i henhold til en afrundingsmåde, der specificeres af matrixmultiplikationsinstruktionen,
- 35 hvor flerheden af multiplikations-akkumuleringskredsløb (416, 418) i udførelseskredsløbet (411, 1150) for hvert ikke-nul-dataelement ved en række og en kolonne i den første kildematrix skal generere et produkt af ikke-nul-dataelementet i den første

kildematrix og alle tilsvarende dataelementer ved en række i den anden kildematrix, idet rækken i den anden kildematrix svarer i antal til kolonnen i den første kildematrix, og derefter akkumulere hvert genereret produkt med tidligere akkumulerede værdier af et tilsvarende dataelement ved en række i resultatmatricen, idet rækken i resultatmatricen svarer i antal til rækken i den første matrix.

2. Processor (1300) ifølge krav 1, hvor det komprimerede sparsomme matrixformat omfatter et komprimeret sparsomt rækkeformat, CSR-format, eller et komprimeret sparsomt kolonneformat, CSC-format.

3. Processor (1300) ifølge krav 1 eller 2, der ydermere omfatter et instruktionshentningskredsløb (1138) til at hente instruktionen og en dekoder (1140) til at afkode matrixmultiplikationsinstruktionen.

4. Processor (1300) ifølge et af kravene 1 til 3, hvor den første kildematrix og den anden kildematrix omfatter en maskinlæringsaktiveringsmatrix og/eller en vægtvektor.

5. Fremgangsmåde, der omfatter:
lagring af data i en processors cache
udførelse ved hjælp af mindst én kerne i processoren, der er koblet til cachen, af multiplikations-akkumuleringsoperationer med en første kildematrix og en anden kildematrix for at generere en resultatmatrix, der reagerer på en matrixmultiplikationsinstruktion, idet
matrixmultiplikationsinstruktionen omfatter en flerhed af felter, herunder et første felt til at angive en opcode, et andet felt til at identificere resultatmatricen, et tredje felt til at identificere den første kildematrix og et fjerde felt til at identificere den anden kildematrix, hvor den første kildematrix er en sparsom matrix med ikke-nul-dataelementer, der befinder sig ved visse positioner, og hvor den anden kildematrix omfatter en tæt matrix, hvor én eller flere af de identificerede matricer er lagret i en registerfil i processoren,

hvor kun den første kildematrix er lagret i et komprimeret format, der identificerer positionerne af ikke-nul-dataelementerne i den første kildematrix, hvor udførelsen af en multiplikations-akkumuleringsoperation ydermere omfatter:

- 5 udførelse af en flerhed af multiplikation-additionsoperationer for at multiplicere ikke-nul-dataelementerne i den første kildematrix med tilsvarende dataelementer i den anden kildematrix, der identificeres på basis af positionerne i komprimeret format, for at generere en flerhed af produkter,
- 10 addering af flerheden af produkter til akkumulerede værdier for at generere dataelementer i resultatmatricen, og udførelse af afrunding på hver af de akkumulerede værdier i henhold til en afrundingsmåde, der specificeres af matrixmultiplikationsinstruktionen,
- 15 hvor der for hvert ikke-nul-dataelement ved en række og en kolonne i den første kildematrix genereres et produkt af ikke-nul-dataelementet i den første kildematrix og alle tilsvarende dataelementer ved en række i den anden kildematrix, idet rækken i den anden kildematrix svarer i antal til kolonnen i den første
- 20 kildematrix, og hvert genereret produkt derefter akkumuleres med tidligere akkumulerede værdier af et tilsvarende dataelement ved en række i resultatmatricen, idet rækken i resultatmatricen svarer i antal til rækken i den første matrix.

- 25 6. Fremgangsmåde ifølge krav 5, hvorved det komprimerede sparsomme matrixformat omfatter et komprimeret sparsomt rækkeformat, CSR-format, eller et komprimeret sparsomt kolonneformat, CSC-format.

- 30 7. Fremgangsmåde ifølge krav 5 eller 6, der ydermere omfatter hentning af instruktionen og afkodning af matrixmultiplikationsinstruktionen.

8. Fremgangsmåde ifølge et af kravene 5 til 7, hvorved den
- 35 første kildematrix og den anden kildematrix omfatter en maskinlæringsaktiveringsmatrix og/eller en vægtvektor.

9. Maskinlæsbart medie med programkode, der er lagret derpå, som, når det udføres af en processor (1300) ifølge et af kravene 1 til 4, får processoren til at udføre fremgangsmåden ifølge et af kravene 5 til 8.

DRAWINGS

Drawing




FIG. 1

Sparse-Dense Matrix Multiplication Instruction <u>202</u> SDMM			
Opcode^ SDMMVNNIW <u>204</u>	Dense Output Matrix <u>206</u>	Sparse Source Matrix <u>208</u>	Dense Source Matrix <u>210</u>




FIG. 2A

Sparse-Dense Matrix Multiplication Instruction <u>252</u> SDMM			
Opcode^ SDMMVNNI <u>254</u>	Dense Output Matrix <u>256</u>	Sparse Source Matrix <u>258</u>	Dense Source Matrix <u>260</u>




FIG. 2B

Sparse-Dense Matrix Multiplication Instruction <u>302</u> SDMM			
Opcode^ SDMMVNNI <u>304</u>	Dense Output Matrix <u>306</u>	Sparse Source Matrix <u>308</u>	Dense Source Matrix <u>310</u>






FIG. 3

**FIG. 4**

```

502
VNNI(C, A, B) {           // C(M x N), A(M x K), B(K x N)
  for (k in 0 to K) {      // for each flow instance
    for (m in 0 to M) {    // for each of M rows of A
      for (n in 0 to N) {  // for each of N columns of B
        C(m, n) += A(m,k) * B(k,n);
      } } } }

504
VNNI(C, A, B) {           // C(M x N), A(M x K), B(K x N)
  for ((m, k) in (0,0)-(M,K)) { // for each element of A at (m, k)
    for (n in 0 to N) {      // for each column of B in corresponding row
      C(m, n) += A(m,k) * B(k,n);
    } } }

506
SDMMVNNIW(C, A, B) {      // C(M x N), A(M x K), B(K x N)
  for (each nonzero A(m, k)) { // Use CSR or CSC to find
    for (n in 0 to N) {      // for each column of B in corresponding row
      C(m, n)[31:0] += A(m,k)[15:0] * B(k,n)[15:0];
    } } }

508
SDMMVNNI(C, A, B) {      // C(M x N), A(M x K), B(K x N)
  for (each nonzero A(m, k)) { // Use CSR or CSC to find
    for (n in 0 to N) {      // for each column of B in corresponding row
      C(m, n)[15:0] += A(m,k)[15:0] * B(k,n)[15:0];
    } } }

```

FIG. 5




FIG. 6




FIG. 7

FIG. 8B




FIG. 9D





FIG. 10






FIG. 11B

FIG. 12A





FIG. 12B



PROCESSOR 1300





FIG. 13

**FIG. 14**




FIG. 15




FIG. 16




FIG. 17

