



(51) International Patent Classification:
G06F 9/30 (2006.01)

(21) International Application Number:
PCT/US2017/040537

(22) International Filing Date:
01 July 2017 (01.07.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/473,732 20 March 2017 (20.03.2017) US

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

(72) Inventors: **VALENTINE, Robert**; Rechov Hadganiot 33-5, 36054 Kiryat Tivon (IL). **BAUM, Dan**; Nizanim 28 A St, 34354 Haifa (IL). **SPERBER, Zeev**; 32nd Igal Alon St., 3092832 Zichron Yackov (IL). **CORBAL, Jesus**; Jordi Girona 1-3 Intel Labs p3, Nexus II, 08034 Barcelona (ES).

OULD-AHMED-VALL, Elmoustapha; 5000 W Chandler Blvd, M/S: CH7-401, Chandler, Arizona 85226 (US). **TOLL, Bret L.**; 2868 NE Lorie Dr., Hillsboro, Oregon 97124 (US). **CHARNEY, Mark J.**; 610 Waltham St, Lexington, Massachusetts 02421 (US). **HEINECKE, Alexander**; 60 Descanso Drive, Unit 3106, San Jose, California 95134 (US).

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

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

(54) Title: SYSTEMS, METHODS, AND APPARATUSES FOR TILE DIAGONAL

TILEDIAGONAL DESTINATION MATRIX (TILE) OPERAND IDENTIFIER, SOURCE IDENTIFIER
2402




FIG. 24

(57) Abstract: Embodiments detailed herein relate to matrix operations. In particular, tile diagonal support is described. For example, a processor is detailed having decode circuitry to decode an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and execution circuitry to execute the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand.

SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *of inventorship (Rule 4.17(iv))*

Published:

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

SYSTEMS, METHODS, AND APPARATUSES FOR TILE DIAGONAL

FIELD OF INVENTION

[0001] The field of invention relates generally to computer processor architecture, and, more specifically, to matrix manipulation.

BACKGROUND

[0002] Matrices are increasingly important in many computing tasks such as machine learning and other bulk data processing.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] 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:

[0004] **Figure 1** illustrates an embodiment of configured tiles;

[0005] **Figure 2** illustrates several examples of matrix storage;

[0006] **Figure 3** illustrates an embodiment of a system utilizing a matrix (tile) operations accelerator;

[0007] **Figures 4 and 5** show different embodiments of how memory is shared using a matrix operations accelerator;

[0008] **Figure 6** illustrates an embodiment of matrix multiply accumulate operation using tiles ("TMMA");

[0009] **Figure 7** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

[0010] **Figure 8** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

[0011] **Figure 9** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

[0012] **Figure 10** illustrates an embodiment of a subset of the execution of an iteration of chained fused multiply accumulate instruction;

[0013] **Figure 11** illustrates power-of-two sized SIMD implementations wherein the accumulators use input sizes that are larger than the inputs to the multipliers according to an embodiment;

- [0014] **Figure 12** illustrates an embodiment of a system utilizing matrix operations circuitry;
- [0015] **Figure 13** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles;
- [0016] **Figure 14** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles;
- [0017] **Figure 15** illustrates an example of a matrix expressed in row major format and column major format;
- [0018] **Figure 16** illustrates an example of usage of matrices (tiles);
- [0019] **Figure 17** illustrates an embodiment a method of usage of matrices (tiles);
- [0020] **Figure 18** illustrates an exemplary execution of a TILECONFIG instruction;
- [0021] **Figures 19(A)-(D)** illustrate examples of register(s);
- [0022] **Figure 20** illustrates an embodiment of a description of the matrices (tiles) to be supported;
- [0023] **Figure 21** illustrates an embodiment of method performed by a processor to process a TILECONFIG instruction;
- [0024] **Figure 22** illustrates a more detailed description of an execution of a TILECONFIG instruction using memory addressing;
- [0025] **Figure 23** illustrates exemplary pseudocode for an execution of a TILECONFIG instruction;
- [0026] **Figure 24** illustrates an exemplary execution of a TILEDIAGONAL instruction;
- [0027] **Figure 25** illustrates an embodiment of method performed by a processor to process a TILEDIAGONAL instruction;
- [0028] **Figure 26** illustrates a more detailed description of an execution of a TILEDIAGONAL instruction;
- [0029] **Figure 27** is exemplary pseudocode describing an embodiment of a method performed by a processor to process a TILEDIAGONALD instruction;
- [0030] **Figures 28(A)-(C)** illustrate an exemplary instruction format;
- [0031] **Figure 29** is a block diagram of a register architecture according to one embodiment of the invention;
- [0032] **Figures 30A-B** illustrate the in-order pipeline and in-order core;
- [0033] **Figures 31A-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;

[0034] **Figure 32** is a block diagram of a processor 3200 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention;

[0035] **Figures 33-36** are block diagrams of exemplary computer architectures; and

[0036] **Figure 37** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention.

DETAILED DESCRIPTION

[0037] In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description.

[0038] References in the specification to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0039] In many mainstream processors, handling matrices is a difficult and/or instruction intensive task. For example, rows of a matrix could be put into a plurality of packed data (e.g., SIMD or vector) registers and then operated on individually. For example, an add two 8x2 matrices may require a load or gather into four packed data registers depending upon data sizes. Then a first add of packed data registers corresponding to a first row from each matrix is performed and a second add of packed data registers corresponding to a second row from each matrix is performed. Then the resulting packed data registers are scattered back to memory. While for small matrices this scenario may be acceptable, it is often not acceptable with larger matrices.

[0040] I. High-Level Discussion

[0041] Described herein are mechanisms to support matrix operations in computer hardware such as central processing units (CPUs), graphic processing units (GPUs), and

accelerators. The matrix operations utilize 2-dimensional (2-D) data structures representing one or more packed regions of memory such as registers. Throughout this description, these 2-D data structures are referred to as tiles. Note that a matrix may be smaller than a tile (use less than all of a tile), or utilize a plurality of tiles (the matrix is larger than the size of any one tile). Throughout the description, matrix (tile) language is used to indicate operations performed using tiles that impact a matrix; whether or not that matrix is larger than any one tile is not typically relevant.

[0042] Each tile may be acted upon by different operations such as those that are detailed herein and include, but are not limited to: matrix (tile) multiplication, tile add, tile subtract, tile diagonal, tile zero, tile transpose, tile dot product, tile broadcast, tile row broadcast, tile column broadcast, tile multiplication, tile multiplication and accumulation, tile move, etc. Additionally, support for operators such as the use of a scale and/or bias may be used with these operations or in support of non-numeric applications in the future, for instance, OpenCL “local memory,” data compression/decompression, etc.

[0043] Portions of storage (such as memory (non-volatile and volatile), registers, cache, etc.) are arranged into tiles of different horizontal and vertical dimensions. For example, a tile may have horizontal dimension of 4 (e.g., four rows of a matrix) and a vertical dimension of 8 (e.g., 8 columns of the matrix). Typically, the horizontal dimension is related to element sizes (e.g., 2-, 4-, 8-, 16-, 32-, 64-, 128-bit, etc.). Multiple datatypes (single precision floating point, double precision floating point, integer, etc.) may be supported.

[0044] A. Exemplary usage of configured tiles

[0045] **Figure 1** illustrates an embodiment of configured tiles. As shown, there are four tiles 111, 113, 115, and 117 that are loaded from application memory 101. In this example, tiles T0 111 and T1 113 have M rows and N columns with 4 element bytes (e.g., single precision data). Tiles T2 115 and T3 117 have M rows and N/2 columns with 8 element bytes (e.g., double precision data). As the double precision operands are twice the width of single precision, this configuration is consistent with a palette, used to provide tile options, supplying at least 4 names with total storage of $16 \times N \times M$ bytes. Depending upon the instruction encoding scheme used, the number of tiles available varies.

[0046] In some embodiments, tile parameters are definable. For example, a “palette” is used to provide tile options. Exemplary options include, but are not limited to: the number of tile names, the number of bytes in a row of storage, the number of rows and columns in a tile, etc. For example, a maximum “height” (number of rows) of a tile may be defined as:

[0047] Tile Max Rows = Architectured Storage / (The Number of Palette Names * The Number of Bytes per row)

[0048] As such, an application can be written such that a fixed usage of names will be able to take advantage of different storage sizes across implementations.

[0049] Configuration of tiles is done using a tile configuration (“TILECONFIG”) instruction, where a particular tile usage is defined in a selected palette. This declaration includes the number of tile names to be used, the requested number of rows and columns per name (tile), and, in some embodiments, the requested datatype of each tile. In some embodiments, consistency checks are performed during the execution of a TILECONFIG instruction to determine that it matches the restrictions of the palette entry.

[0050] B. Exemplary tile storage types

[0051] **Figure 2** illustrates several examples of matrix storage. In (A), a tile is stored in memory. As shown, each “row” consists of four packed data elements. To get to the next “row,” a stride value is used. Note that rows may be consecutively stored in memory. Strided memory accesses allows for access of one row to then next when the tile storage does not map the underlying memory array row width.

[0052] Tile loads from memory and stores to memory are typically strided accesses from the application memory to packed rows of data. Exemplary TILELOAD and TILESTORE instructions, or other instruction references to application memory as a TILE operand in load-loop instructions, are, in some embodiments, restartable to handle (up to) 2*rows of page faults, unmasked floating point exceptions, and/or interrupts per instruction.

[0053] In (B), a matrix is stored in a tile comprised of a plurality of registers such as packed data registers (single instruction, multiple data (SIMD) or vector registers). In this example, the tile is overlaid on three physical registers. Typically, consecutive registers are used, however, this need not be the case.

[0054] In (C), a matrix is stored in a tile in non-register storage accessible to a fused multiple accumulate (FMA) circuit used in tile operations. This storage may be inside of a FMA, or adjacent to it. Additionally, in some embodiments, discussed below, the storage may be for a data element and not an entire row or tile.

[0055] The supported parameters for the TMMA architecture are reported via CPUID. In some embodiments, the list of information includes a maximum height and a maximum SIMD dimension. Configuring the TMMA architecture requires specifying the dimensions for each tile, the element size for each tile and the palette identifier. This configuration is done by executing the TILECONFIG instruction.

[0056] Successful execution of a TILECONFIG instruction enables subsequent TILE operators. A TILERELASEALL instruction clears the tile configuration and disables the TILE operations (until the next TILECONFIG instructions executes). In some embodiments, XSAVE, XSTORE, etc. are used in context switching using tiles. In some embodiments, 2 XCR0 bits are used in XSAVE, one for TILECONFIF metadata and one bit corresponding to actual tile payload data.

[0057] TILECONFIG not only configures the tile usage, but also sets a state variable indicating that the program is in a region of code with tiles configured. An implementation may enumerate restrictions on other instructions that can be used with a tile region such as no usage of an existing register set, etc.

[0058] Exiting a tile region is typically done with the TILERELASEALL instruction. It takes no parameters and swiftly invalidates all tiles (indicating that the data no longer needs any saving or restoring) and clears the internal state corresponding to being in a tile region.

[0059] In some embodiments, tile operations will zero any rows and any columns beyond the dimensions specified by the tile configuration. For example, tile operations will zero the data beyond the configured number of columns (factoring in the size of the elements) as each row is written. For example, with 64 byte rows and a tile configured with 10 rows and 12 columns, an operation writing FP32 elements would write each of the first 10 rows with 12×4 bytes with output/result data and zero the remaining 4×4 bytes in each row. Tile operations also fully zero any rows after the first 10 configured rows. When using 1K tile with 64 byte rows, there would be 16 rows, so in this example, the last 6 rows would also be zeroed.

[0060] In some embodiments, a context restore (e.g., XRSTOR), when loading data, enforces that the data beyond the configured rows for a tile will be maintained as zero. If there is no valid configuration, all rows are zeroed. XRSTOR of tile data can load garbage in the columns beyond those configured. It should not be possible for XRSTOR to clear beyond the number of columns configured because there is not an element width associated with the tile configuration.

[0061] Context save (e.g., XSAVE) exposes the entire TILE storage area when writing it to memory. If XRSTOR loaded garbage data in to the rightmost part of a tile, that data will be saved by XSAVE. XSAVE will write zeros for rows beyond the number specified for each tile.

[0062] In some embodiments, tile instructions are restartable. The operations that access memory allow restart after page faults. The computational instructions that deal

with floating point operations also allow for unmasked floating point exceptions, with the masking of the exceptions controlled by a control and/or status register.

[0063] To support restarting instructions after these events, the instructions store information in the start registers detailed below.

[0064] II. Matrix (Tile) Operation Systems

[0065] A. Exemplary Hardware Support

[0066] **Figure 3** illustrates an embodiment of a system utilizing a matrix (tile) operations accelerator. In this illustration, a host processor/processing system 301 communicates commands 311 (e.g., matrix manipulation operations such as arithmetic or matrix manipulation operations, or load and store operations) to a matrix operations accelerator 307. However, this is shown this way for discussion purposes only. As detailed later, this accelerator 307 may be a part of a processing core. Typically, commands 311 that are tile manipulation operator instructions will refer to tiles as register-register (“reg-reg”) or register-memory (“reg-mem”) format. Other commands such as TILESTORE, TILELOAD, TILECONFIG, etc., do not perform data operations on a tile. Commands may be decoded instructions (e.g., micro-ops) or macro-instructions for the accelerator 307 to handle.

[0067] In this example, a coherent memory interface 303 is coupled to the host processor/processing system 301 and matrix operations accelerator 405 such that they can share memory. **Figures 4 and 5** show different embodiments of how memory is shared using a matrix operations accelerator. As shown in **Figure 4**, the host processor 401 and matrix operations accelerator circuitry 405 share the same memory 403. **Figure 5** illustrates an embodiment where the host processor 501 and matrix operations accelerator 505 do not share memory, but can access each other’s memory. For example, processor 501 can access tile memory 507 and utilize its host memory 503 as normal. Similarly, the matrix operations accelerator 505 can access host memory 503, but more typically uses its own memory 507. Note these memories may be of different types.

[0068] The matrix operations accelerator 307 includes a plurality of FMAs 309 coupled to data buffers 305 (in some implementations, one or more of these buffers 305 are stored in the FMAs of the grid as shown). The data buffers 305 buffer tiles loaded from memory and/or tiles to be stored to memory (e.g., using a tileload or tilestore instruction). Data buffers may be, for example, a plurality of registers. Typically, these FMAs are arranged as a grid of chained FMAs 309 which are able to read and write tiles. In this example, the matrix operations accelerator 307 is to perform a matrix multiply operation using tiles T0, T1, and T2. At least one of tiles is housed in the FMA grid 309. In some embodiments, all tiles in an

operation are stored in the FMA grid 309. In other embodiments, only a subset are stored in the FMA grid 309. As shown, T1 is housed and T0 and T2 are not. Note that A, B, and C refer to the matrices of these tiles which may or may not take up the entire space of the tile.

[0069] **Figure 6** illustrates an embodiment of matrix multiply accumulate operation using tiles (“TMMA”).

[0070] The number of rows in the matrix (TILE A 601) matches the number of serial (chained) FMAs comprising the computation’s latency. An implementation is free to recirculate on a grid of smaller height, but the computation remains the same.

[0071] The source/destination vector comes from a tile of N rows (TILE C 605) and the grid of FMAs 611 performs N vector-matrix operations resulting in a complete instruction performing a matrix multiplication of tiles. Tile B 603 is the other vector source and supplies “broadcast” terms to the FMAs in each stage.

[0072] In operation, in some embodiments, the elements of matrix B (stored in a tile B 603) are spread across the rectangular grid of FMAs. Matrix B (stored in tile A 601) has its elements of a row transposed to match up with the columnar dimension of the rectangular grid of FMAs. At each FMA in the grid, an element of A and B are multiplied and added to the incoming summand (from above in the Figure) and the outgoing sum is passed to the next row of FMAs (or the final output).

[0073] The latency of a single step is proportional to K (row height of matrix B) and dependent TMMAs typically have enough source-destination rows (either in a single tile or across tile) to hide that latency. An implementation may also split the SIMD (packed data element) dimension M (row height of matrix A) across time steps, but this simply changes the constant that K is multiplied by. When a program specifies a smaller K than the maximum enumerated by the TMACC, an implementation is free to implement this with “masking” or “early outs.”

[0074] The latency of an entire TMMA is proportional to $N \cdot K$. The repeat rate is proportional to N. The number of MACs per TMMA instruction is $N \cdot K \cdot M$.

[0075] **Figure 7** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on signed sources wherein the accumulator is 2x the input data size.

[0076] A first signed source (source 1 701) and a second signed source (source 2 703) each have four packed data elements. Each of these packed data elements stores signed data

such as floating point data. A third signed source (source 3 709) has two packed data elements, each of which stores signed data. The sizes of the first and second signed sources 701 and 703 are half that of the third signed source (initial value or previous result) 709. For example, the first and second signed sources 701 and 703 could have 32-bit packed data elements (e.g., single precision floating point) while the third signed source 709 could have 64-bit packed data elements (e.g., double precision floating point).

[0077] In this illustration, only the two most significant packed data element positions of the first and second signed sources 701 and 703 and the most significant packed data element position of the third signed source 709 are shown. Of course, the other packed data element positions would also be processed.

[0078] As illustrated, packed data elements are processed in pairs. For example, the data of the most significant packed data element positions of the first and second signed sources 701 and 703 are multiplied using a multiplier circuit 705, and the data from second most significant packed data element positions of the first and second signed sources 701 and 703 are multiplied using a multiplier circuit 707. In some embodiments, these multiplier circuits 705 and 707 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source 709. The results of each of the multiplications are added using addition circuitry 711.

[0079] The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of the signed source 3 709 (using a different adder 713 or the same adder 711).

[0080] Finally, the result of the second addition is either stored into the signed destination 715 in a packed data element position that corresponds to the packed data element position used from the signed third source 709, or passed on to the next iteration, if there is one. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0081] **Figure 8** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on signed sources wherein the accumulator is 2x the input data size.

[0082] A first signed source (source 1 801) and a second signed source (source 2 803) each have four packed data elements. Each of these packed data elements stores signed data such as integer data. A third signed source (source 3 809) has two packed data elements, each of which stores signed data. The sizes of the first and second signed sources 801 and 803 are half that of the third signed source 809. For example, the first and second signed sources 801 and 803 could have 32-bit packed data elements (e.g., single precision floating point) the third signed source 809 could have 64-bit packed data elements (e.g., double precision floating point).

[0083] In this illustration, only the two most significant packed data element positions of the first and second signed sources 801 and 803 and the most significant packed data element position of the third signed source 809 are shown. Of course, the other packed data element positions would also be processed.

[0084] As illustrated, packed data elements are processed in pairs. For example, the data of the most significant packed data element positions of the first and second signed sources 801 and 803 are multiplied using a multiplier circuit 805, and the data from second most significant packed data element positions of the first and second signed sources 801 and 803 are multiplied using a multiplier circuit 807. In some embodiments, these multiplier circuits 805 and 807 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source (initial value or previous iteration result) 809. The results of each of the multiplications are added to the signed third source 809 using addition/saturation circuitry 811.

[0085] Addition/saturation (accumulator) circuitry 811 preserves a sign of an operand when the addition results in a value that is too big. In particular, saturation evaluation occurs on the infinite precision result between the multi-way-add and the write to the destination or next iteration. When the accumulator 811 is floating point and the input terms are integer, the sum of products and the floating point accumulator input value are turned into infinite precision values (fixed point numbers of hundreds of bits), the addition of the multiplication results and the third input is performed, and a single rounding to the actual accumulator type is performed.

[0086] Unsigned saturation means the output values are limited to a maximum unsigned number for that element width (all 1s). Signed saturation means a value is limited to the be in

the range between a minimum negative number and a max positive number for that element width (for bytes for example, the range is from -128 ($= -2^7$) to 127 ($=2^7-1$)).

[0087] The result of the addition and saturation check is stored into the signed result 815 in a packed data element position that corresponds to the packed data element position used from the signed third source 809, or passed on to the next iteration if there is one. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0088] **Figure 9** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on a signed source and an unsigned source wherein the accumulator is 4x the input data size.

[0089] A first signed source (source 1 901) and a second unsigned source (source 2 903) each have four packed data elements. Each of these packed data elements has data such as floating point or integer data. A third signed source (initial value or result 915) has a packed data element of which stores signed data. The sizes of the first and second sources 901 and 903 are a quarter of the third signed source 915. For example, the first and second sources 901 and 903 could have 16-bit packed data elements (e.g., word) and the third signed source 915 could have 64-bit packed data elements (e.g., double precision floating point or 64-bit integer).

[0090] In this illustration, the four most significant packed data element positions of the first and second sources 901 and 903 and the most significant packed data element position of the third signed source 915 are shown. Of course, other packed data element positions would also be processed if there are any.

[0091] As illustrated, packed data elements are processed in quadruplets. For example, the data of the most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 907, data from second most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 907, data from third most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 909, and data from the least significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 911. In some embodiments, the signed packed data elements of the first source 901 are sign extended and the unsigned packed data elements of the second source 903 are zero extended prior to the multiplications.

[0092] In some embodiments, these multiplier circuits 905-911 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source 915. The results of each of the multiplications are added using addition circuitry 911.

[0093] The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of the signed source 3 915 (using a different adder 913 or the same adder 911).

[0094] Finally, the result 919 of the second addition is either stored into the signed destination in a packed data element position that corresponds to the packed data element position used from the signed third source 915, or passed to the next iteration. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0095] **Figure 10** illustrates an embodiment of a subset of the execution of an iteration of chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on a signed source and an unsigned source wherein the accumulator is 4x the input data size.

[0096] A first signed source (source 1 1001) and a second unsigned source (source 2 1003) each have four packed data elements. Each of these packed data elements stores data such as floating point or integer data. A third signed source (initial or previous result 1015) has a packed data element of which stores signed data. The sizes of the first and second sources 1001 and 1003 are a quarter of the third signed source 1015. For example, the first and second sources 1001 and 1003 could have 16-bit packed data elements (e.g., word) and the third signed source 1015 could have 64-bit packed data elements (e.g., double precision floating point or 64-bit integer).

[0097] In this illustration, the four most significant packed data element positions of the first and second sources 1001 and 1003 and the most significant packed data element position of the third signed source 1015 are shown. Of course, other packed data element positions would also be processed if there are any.

[0098] As illustrated, packed data elements are processed in quadruplets. For example, the data of the most significant packed data element positions of the first and second sources 1001 and 1003 are multiplied using a multiplier circuit 1007, data from second most significant packed data element positions of the first and second sources 1001 and 1003 are

multiplied using a multiplier circuit 1007, data from third most significant packed data element positions of the first and second sources 1001 and 1003 are multiplied using a multiplier circuit 1009, and data from the least significant packed data element positions of the first and second sources 1001 and 1003 are multiplied using a multiplier circuit 1011. In some embodiments, the signed packed data elements of the first source 1001 are sign extended and the unsigned packed data elements of the second source 1003 are zero extended prior to the multiplications.

[0099] In some embodiments, these multiplier circuits 1005-1011 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source 1015. The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of the signed source 3 1015 using addition/saturation circuitry 1013.

[0100] Addition/saturation (accumulator) circuitry 1013 preserves a sign of an operand when the addition results in a value that is too big or too small for signed saturation. In particular, saturation evaluation occurs on the infinite precision result between the multi-way-add and the write to the destination. When the accumulator 1013 is floating point and the input terms are integer, the sum of products and the floating point accumulator input value are turned into infinite precision values (fixed point numbers of hundreds of bits), the addition of the multiplication results and the third input is performed, and a single rounding to the actual accumulator type is performed.

[0101] The result 1019 of the addition and saturation check is stored into the signed destination in a packed data element position that corresponds to the packed data element position used from the signed third source 1015, or passed to the next iteration. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0102] **Figure 11** illustrates power-of-two sized SIMD implementations wherein the accumulators use input sizes that are larger than the inputs to the multipliers according to an embodiment. Note the source (to the multipliers) and accumulator values may be signed or unsigned values. For an accumulator having 2X input sizes (in other words, the accumulator input value is twice the size of the packed data element sizes of the sources), table 1101 illustrates different configurations. For byte sized sources, the accumulator uses word or half-precision floating-point (HPFP) values that are 16-bit in size. For word sized sources,

the accumulator uses 32-bit integer or single-precision floating-point (SPFP) values that are 32-bit in size. For SPFP or 32-bit integer sized sources, the accumulator uses 64-bit integer or double-precision floating-point (DPFP) values that are 64-bit in size.

[0103] For an accumulator having 4X input sizes (in other words, the accumulator input value is four times the size of the packed data element sizes of the sources), table 1103 illustrates different configurations. For byte sized sources, the accumulator uses 32-bit integer or single-precision floating-point (SPFP) values that are 32-bit in size. For word sized sources, the accumulator uses 64-bit integer or double-precision floating-point (DPFP) values that are 64-bit in size in some embodiments.

[0104] For an accumulator having 8X input sizes (in other words, the accumulator input value is eight times the size of the packed data element sizes of the sources), table 1105 illustrates a configuration. For byte sized sources, the accumulator uses 64-bit integer.

[0105] As hinted at earlier, matrix operations circuitry may be included in a core, or as an external accelerator. **Figure 12** illustrates an embodiment of a system utilizing matrix operations circuitry. In this illustration, a plurality of entities are coupled with a ring interconnect 1245.

[0106] A plurality of cores 1201, 1203, 1205, and 1207 provide non-tile based instruction support. In some embodiments, matrix operations circuitry is provided in a core 1203, and in other embodiments matrix operations circuitry 1211 and 1213 is accessible on the ring interconnect 1245.

[0107] Additionally, one or more memory controllers 1223-1225 are provided to communicate with memory 1233 and 1231 on behalf of the cores and/or matrix operations circuitry.

[0108] **Figure 13** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles. Branch prediction and decode circuitry 1303 performs branch predicting of instructions, decoding of instructions, and/or both from instructions stored in instruction storage 1301. For example, instructions detailed herein may be stored in instruction storage. In some implementations, separate circuitry is used for branch prediction and in some embodiments, at least some instructions are decoded into one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals using microcode 1305. The branch prediction and decode circuitry 1303 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.

[0109] The branch prediction and decode circuitry 1303 is coupled to a rename/allocator circuitry 1307 which is coupled, in some embodiments, to scheduler circuitry 1309. In some embodiments, these circuits provide register renaming, register allocation, and/or scheduling functionality by performing 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 instruction for execution on execution circuitry out of an instruction pool (e.g., using a reservation station in some embodiments).

[0110] The scheduler circuitry 1309 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) scheduler circuitry 1309 is coupled to, or includes, physical register file(s) 1315. Each of the physical register file(s) 1315 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), tiles, etc. In one embodiment, the physical register file(s) 1315 comprises vector registers circuitry, write mask registers circuitry, and scalar registers circuitry. These register circuits may provide architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) 1315 is overlapped by a retirement circuit 1317 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 circuit 1317 and the physical register file(s) 1315 are coupled to the execution circuit(s) 1311.

[0111] 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 may also include separate instruction and data cache units and a shared L2 cache unit, 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.

[0112] The execution circuitry 1311 a set of one or more execution circuits 1321, 1323, and 1327 and a set of one or more memory access circuits 1325. The execution circuits 1321,

1323, and 1327 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 scalar circuitry 1321 performs scalar operations, the vector/SIMD circuitry 1323 performs vector/SIMD operations, and matrix operations circuitry 1327 performs matrix (tile) operations detailed herein.

[0113] The set of memory access units 1364 is coupled to the memory unit 1370, which includes a data TLB unit 1372 coupled to a data cache unit 1374 coupled to a level 2 (L2) cache unit 1376. In one exemplary embodiment, the memory access units 1364 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 1372 in the memory unit 1370. The instruction cache unit 1334 is further coupled to a level 2 (L2) cache unit 1376 in the memory unit 1370. The L2 cache unit 1376 is coupled to one or more other levels of cache and eventually to a main memory.

[0114] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement a pipeline as follows: 1) an instruction fetch circuit performs fetch and length decoding stages; 2) the branch and decode circuitry 1303 performs a decode stage; 3) the rename/allocator circuitry 1307 performs an allocation stage and renaming stage; 4) the scheduler circuitry 1309 performs a schedule stage; 5) physical register file(s) (coupled to, or included in, the scheduler circuitry 1307 and rename/allocate circuitry 1307 and a memory unit perform a register read/memory read stage; the execution circuitry 1311 performs an execute stage; 6) a memory unit and the physical register file(s) unit(s) perform a write back/memory write stage; 7) various units may be involved in the exception handling stage; and 8) a retirement unit and the physical register file(s) unit(s) perform a commit stage.

[0115] The core 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 1390 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.

[0116] 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).

[0117] **Figure 14** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles. Branch prediction and decode circuitry 1403 performs branch predicting of instructions, decoding of instructions, and/or both from instructions stored in instruction storage 1401. For example, instructions detailed herein may be stored in instruction storage. In some implementations, separate circuitry is used for branch prediction and in some embodiments, at least some instructions are decoded into one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals using microcode 1405. The branch prediction and decode circuitry 1403 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.

[0118] The branch prediction and decode circuitry 1403 is coupled to a rename/allocator circuitry 1407 which is coupled, in some embodiments, to scheduler circuitry 1409. In some embodiments, these circuits provide register renaming, register allocation, and/or scheduling functionality by performing 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 instruction for execution on execution circuitry out of an instruction pool (e.g., using a reservation station in some embodiments).

[0119] The scheduler circuitry 1409 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) scheduler circuitry 1409 is coupled to, or includes, physical register file(s) 1415. Each of the physical register file(s) 1415 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), tiles, etc. In one embodiment, the physical register file(s) 1415 comprises vector registers circuitry, write mask registers circuitry, and scalar registers circuitry. These register circuits may provide

architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) 1415 is overlapped by a retirement circuit 1417 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 circuit 1417 and the physical register file(s) 1415 are coupled to the execution circuit(s) 1411.

[0120] 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 may also include separate instruction and data cache units and a shared L2 cache unit, 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.

[0121] The execution circuitry 1411 a set of one or more execution circuits 1427 and a set of one or more memory access circuits 1425. The execution circuits 1427 perform matrix (tile) operations detailed herein.

[0122] The set of memory access units 1464 is coupled to the memory unit 1470, which includes a data TLB unit 1472 coupled to a data cache unit 1474 coupled to a level 2 (L2) cache unit 1476. In one exemplary embodiment, the memory access units 1464 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 1472 in the memory unit 1470. The instruction cache unit 1434 is further coupled to a level 2 (L2) cache unit 1476 in the memory unit 1470. The L2 cache unit 1476 is coupled to one or more other levels of cache and eventually to a main memory.

[0123] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement a pipeline as follows: 1) an instruction fetch circuit performs fetch and length decoding stages; 2) the branch and decode circuitry 1403 performs a decode stage; 3) the rename/allocator circuitry 1407 performs an allocation stage and renaming stage; 4) the scheduler circuitry 1409 performs a schedule stage; 5) physical register file(s) (coupled to, or included in, the scheduler circuitry 1407 and rename/allocate circuitry 1407 and a memory unit perform a register read/memory read stage; the execution circuitry 1411 performs an execute stage; 6) a memory unit and the physical register file(s) unit(s) perform a write back/memory write stage; 7) various units may be involved in the

exception handling stage; and 8) a retirement unit and the physical register file(s) unit(s) perform a commit stage.

[0124] The core 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 1490 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.

[0125] 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).

[0126] B. Layout

[0127] Throughout this description, data is expressed using row major data layout. Column major users should translate the terms according to their orientation. **Figure 15** illustrates an example of a matrix expressed in row major format and column major format. As shown, matrix A is a 2x3 matrix. When this matrix is stored in row major format, the data elements of a row are consecutive. When this matrix is stored in column major format, the data elements of a column are consecutive. It is a well-known property of matrices that $A^T * B^T = (BA)^T$, where superscript T means transpose. Reading column major data as row major data results in the matrix looking like the transpose matrix.

[0128] In some embodiments, row-major semantics are utilized in hardware, and column major data is to swap the operand order with the result being transposes of matrix, but for subsequent column-major reads from memory it is the correct, non-transposed matrix.

[0129] For example, if there are two column-major matrices to multiply:

$$\begin{array}{rcl}
 \begin{matrix} a & b \\ c & d \\ e & f \end{matrix} & \begin{matrix} g & i & k \\ h & j & l \end{matrix} & = & \begin{matrix} ag+bh & ai+bj & ak+bl \\ cg+dh & ci+dj & ck+dl \\ eg+fh & ei+fj & ek+fl \end{matrix} \\
 (3 \times 2) & (2 \times 3) & & (3 \times 3)
 \end{array}$$

[0130] The input matrices would be stored in linear memory (column-major) as:

a c e b d f

and

g h i j k l.

[0131] Reading those matrices as row-major with dimensions 2x3 and 3x2, they would appear as:

a c e and g h
b d f i j
k l

[0132] Swapping the order and matrix multiplying:

g h a c e ag+bh cg+dh eg+fh
i j * b d f = ai+bj ci+dj ei+fj
k l ak+bl ck+dl ek+fl

the transpose matrix is out and can then be stored in in row-major order:

ag+bh cg+dh eg+fh ai+bj ci+dj ei+fj ak+bl ck+dl ek+fl

and used in subsequent column major computations, it is the correct un-transposed matrix:

ag+bh ai+bj ak+bl
cg+dh ci+dj ck+dl
eg+fh ei+fj ek+fl

[0133] III. Exemplary Usage

[0134] **Figure 16** illustrates an example of usage of matrices (tiles). In this example, matrix C 1601 includes two tiles, matrix A 1603 includes one tile, and matrix B 1605 includes two tiles. This figure shows an example of the inner loop of an algorithm to compute a matrix multiplication. In this example, two result tiles, tmm0 and tmm1, from matrix C 1601 are used to accumulate the intermediate results. One tile from the A matrix 1603 (tmm2) is re-used twice as it multiplied by two tiles from the B matrix 1605. Pointers to load a new A tile and two new B tiles from the directions indicated by the arrows. An outer loop, not shown, adjusts the pointers for the C tiles.

[0135] The exemplary code as shown includes the usage of a tile configuration instruction and is executed to configure tile usage, load tiles, a loop to process the tiles, store tiles to memory, and release tile usage.

[0136] **Figure 17** illustrates an embodiment of usage of matrices (tiles). At 1701, tile usage is configured. For example, a TILECONFIG instruction is executed to configure tile usage including setting a numbers of rows and columns per tile. Typically, at least one matrix (tile) is loaded from memory at 1703.

[0137] IV. Exemplary Instructions

[0138] A. Tile Configuration

[0139] As discussed above, tile usage typically needs to be configured prior to use. For example, full usage of all rows and columns may not be needed. Not only does not configuring these rows and columns save power in some embodiments, but the configuration may be used to determine if an operation will generate an error. For example, a matrix multiplication of the form $(N \times M) \times (L \times N)$ will typically not work if M and L are not the same.

[0140] Detailed herein are embodiments of a matrix (tile) configuration (“TILECONFIG”) instruction and its execution. Prior to using matrices using tiles, in some embodiments, tile support is to be configured. For example, how many rows and columns per tile, tiles that are to be used, etc. are configured. A TILECONFIG instruction is an improvement to a computer itself as it provides for support to configure the computer to use a matrix accelerator (either as a part of a processor core, or as an external device). In particular, an execution of the TILECONFIG instruction causes a configuration to be retrieved from memory and applied to matrix (tile) settings within a matrix accelerator.

[0141] i. Exemplary Execution

[0142] **Figure 18** illustrates an exemplary execution of a TILECONFIG instruction. The TILECONFIG instruction format includes fields for an opcode and a memory address.

[0143] As illustrated, the TILECONFIG instruction uses the address as a pointer to a memory 1801 location containing the description of the matrices (tiles) to be supported 1803.

[0144] Execution circuitry 1811 of a processor/core 1805 performs the TILECONFIG by retrieving the description 1803 from memory 1801 via a memory controller 1815, configuring tiles for a palette (setting the number of rows and columns) in a tile configuration 1817, and marking that matrix support is in use. In particular, instruction execution resources 1811 are configured to use tiles as specified by setting tile configurations 1817. The instruction execution resources may also include a machine specific register or configuration register to indicate tile usage.

[0145] Tile configurations 1817 are set to indicate parameters per tile as indicated by the tile description 1803 via the execution of the TILECONFIG instruction. The set parameters are the number of rows and columns per tile. Additional values such as in-use and start values are also set. The tile configurations 1817 utilize one or more registers 1819 to store tile usage and configuration information.

[0146] ii. Exemplary tile storage

[0147] **Figures 19(A)-(D)** illustrate examples of register(s) 1819. **Figure 19(A)** illustrates a plurality of registers 1819. As shown each tile (TMM0 1901 ... TMMN 1903) has a separate register with each register storing a row and column size for that particular tile. StartK and StartM are stored in separate registers 1911 and 1913. One or more status registers 1915 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0148] **Figure 19(B)** illustrates a plurality of registers 1819. As shown each tile has separate registers for its rows and columns. For example, TMM0 rows configuration 1921, TMM0 columns configuration 1923, StartK and StartM are stored in separate registers 1911 and 1913. One or more status registers 1915 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0149] **Figure 19(C)** illustrates a single register 1819. As shown, this register stores tile configurations (rows and columns per tile) 1931, StartK 1933, and StartM 1933 are stored in single register as packed data registers. One or more status registers 1915 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0150] **Figure 19(D)** illustrates a plurality of registers 1819. As shown, a single register stores tile configurations (rows and columns per tile) 1931. StartK and StartM are stored in separate registers 1911 and 1913. One or more status registers 1915 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0151] Other combinations are contemplated such as combining the start registers into a single register where they are shown separately, etc.

[0152] iii. Exemplary Stored Matrix (Tile) Description

[0153] **Figure 20** illustrates an embodiment of a description of the matrices (tiles) to be supported. In this example, each field is a byte. In byte[0], a palette ID 2001 is stored. The palette ID is used to index a palette table 1813 which stores, per palette ID, a number of bytes in a tile, and bytes per row of the tiles that are associated with this ID as defined by the configuration. Bytes 1-7 are reserved and are typically zero.

[0154] Bytes 8-9 store a value for a “startM” register 2003 and bytes 10-11 store a value for a “startK” register 2005. To support restarting instructions after these events, the instructions store information these registers. The startM indicates a row that should be used for restart. The startK indicates a position in the inner-product for relevant operations. The position in the row (the column) is not needed. Two-dimensional operations like the element-wise addition/subtraction/multiplication only use startM.

Three-dimensional operations use values from both startM and startK. Typically, operations that only require startM will zero startK when writing startM.

[0155] Any time an interrupted tile instruction is not restarted, in some embodiments, it is the responsibility of software to zero the startM and startK values. For example, unmasked floating point exception handlers might decide to finish the operation in software and change the program counter value to another instruction, usually the next instruction. In this case the software exception handler must zero the startM and startK values in the exception frame presented to it by the operating system before resuming the program. The operating system will subsequently reload those values.

[0156] Bytes 16-17 store the number of rows 2013 and columns 2015 for tile 0, bytes 18-19 store the number of rows and columns for tile 1, etc. In other words, each 2 byte group specifies a number of rows and columns for a tile. If a group of 2 bytes is not used to specify tile parameters, they should have the value zero. Specifying tile parameters for more tiles than the implementation limit or the palette limit results in a fault. Unconfigured tiles are set to the INIT state with 0 rows, 0 columns.

[0157] Finally, the configuration in memory typically ends with an ending delineation such as all zeros for several consecutive bytes.

[0158] iv. Exemplary Format(s)

[0159] An embodiment of a format for a TILECONFIG instruction is TILECONFIG Address. In some embodiments, TILECONFIG is the opcode mnemonic of the instruction. Address is a pointer to a matrix (tile) description in memory. In some embodiments, the address field is a R/M value (such as 2446).

[0160] In embodiments, encodings of the instruction include a scale-index-base (SIB) type memory addressing operand that indirectly identifies multiple indexed destination locations in memory (e.g., field 2450). In one embodiment, an SIB type memory operand may include an encoding identifying a base address register. The contents of the base address register may represent a base address in memory from which the addresses of the particular destination locations in memory are calculated. For example, the base address may be the address of the first location in a block of potential destination locations for an extended vector instruction. In one embodiment, an SIB type memory operand may include an encoding identifying an index register. Each element of the index register may specify an index or offset value usable to compute, from the base address, an address of a respective destination location within a block of potential destination locations. In one embodiment, an SIB type memory operand may include an encoding specifying a scaling factor to be applied

to each index value when computing a respective destination address. For example, if a scaling factor value of four is encoded in the SIB type memory operand, each index value obtained from an element of the index register may be multiplied by four and then added to the base address to compute a destination address.

[0161] In one embodiment, an SIB type memory operand of the form $vm32\{x,y,z\}$ may identify a vector array of memory operands specified using SIB type memory addressing. In this example, the array of memory addresses is specified using a common base register, a constant scaling factor, and a vector index register containing individual elements, each of which is a 32-bit index value. The vector index register may be a 128-bit register (e.g., XMM) register ($vm32x$), a 256-bit (e.g., YMM) register ($vm32y$), or a 512-bit (e.g., ZMM) register ($vm32z$). In another embodiment, an SIB type memory operand of the form $vm64\{x,y,z\}$ may identify a vector array of memory operands specified using SIB type memory addressing. In this example, the array of memory addresses is specified using a common base register, a constant scaling factor, and a vector index register containing individual elements, each of which is a 64-bit index value. The vector index register may be a 128-bit register (e.g., XMM) register ($vm64x$), a 256-bit (e.g., YMM) register ($vm64y$) or a 512-bit (e.g., ZMM) register ($vm64z$).

[0162] v. Exemplary Method(s) of Execution

[0163] **Figure 21** illustrates an embodiment of method performed by a processor to process a TILECONFIG instruction.

[0164] At 2101, an instruction is fetched. For example, a TILECONFIG instruction is fetched. An embodiment of the TILECONFIG instruction includes fields for an opcode and a memory address operand.

[0165] The fetched instruction is decoded at 2103. For example, the fetched TILECONFIG instruction is decoded by decode circuitry such as that detailed herein.

[0166] A description found at the memory address of the memory address operand is retrieved at 2105 and the decoded instruction is scheduled (as needed).

[0167] At 2107, the decoded instruction is executed by execution circuitry (hardware) such as that detailed herein. For the TILECONFIG instruction, the execution will cause execution circuitry to configure usage of tiles in a tile configuration (setting the number of rows and columns) and marking that matrix (tile) support is in use (active). For example, configuration one or more registers 1819. Tile support usage (e.g., “TILES_CONFIGURED”) is typically indicated by setting a bit in a status, control, or

machine specific register. In particular, instruction execution resources 1811 are configured to use tiles as specified by the retrieved configuration.

[0168] In some embodiments, the instruction is committed or retired at 2109.

[0169] **Figure 22** illustrates a more detailed description of an execution of a TILECONFIG instruction using memory addressing. Typically, this is performed by execution circuitry such as that detailed above after the description has been retrieved from memory. While not illustrated, in some embodiments, a check is first performed to determine if tiles are supported. Support is usually found by a CUID check.

[0170] At 2201, a determination of if the palette ID is supported is made. For example, does the CUID state that this ID is supported? If not, then a general protection fault occurs at 2203.

[0171] At 2205, a first tile specific grouping is read. For example, the number of rows and columns for tile 0 (T0) is read.

[0172] A determination of if the read grouping is valid is made at 2207. For example, if one the number of rows or columns (not both) is set 0, then the grouping is not valid and the configuration halts and tiles are not considered to be in use at 2203. Invalid groups occur, for example, when one of rows or columns (not both) are zero. Additionally, when a value for the number of rows is greater than the maximum of rows supported (this is found by dividing the tile byte size of the palette ID with the number of bytes per row for the palette ID as found in the palette table) as fault occurs. Another potential fault is when there are more names than supported.

[0173] If the read grouping is valid, then the tile associated with the read grouping is configured to use the number of rows and columns specified by the grouping in a tile configuration at 2211. The size of the elements in the tile are set by the palette table entry for the palette ID.

[0174] A determination of if all tiles of the retrieved configuration have been configured is made at 2213. For example, have all of the possible tile names been processed? In some embodiments, when the rows and columns for a particular tile are both 0, then all tiles have been processed.

[0175] When all tiles have not been configured, at 2215, the tile number is incremented such that the next tile in the configuration will be evaluated.

[0176] At 2217, the incremented tile's grouping is read. For example, the number of rows and columns for tile 1 (T1) is read. A determination of if the read grouping is valid is made at 2207, etc.

[0177] When all tiles have been configured, then the instruction completes at 2209. The tiles will be marked as being in use for matrix operations, for example, by setting an in-use indicator in a register.

[0178] vi. Exemplary Pseudocode

[0179] **Figure 23** illustrates exemplary pseudocode for an execution of a TILECONFIG instruction.

[0180] B. Tile Diagonal

[0181] Detailed herein are embodiments of a matrix (tile) diagonal (“TILEDIAGONAL”) instruction and its execution. A TILEDIAGONAL instruction is an improvement to a computer itself as it provides for support to populate the main diagonal of a matrix (tile) with a single instruction. In particular, the execution of the TILEDIAGONAL instruction causes execution circuitry to store the identified source operand to every element along the main diagonal of the destination matrix (tile) and zeros all other elements in configured rows. The size of the data values to be stored varies depending on the instruction and tile support. Exemplary sizes include, but are not limited to, 16-bit, 32-bit, 64-bit, 128-bit, 256-bit, etc. In some embodiments, elements of the destination matrix (tile) that are not on the diagonal are zeroed. This instruction may be used, for example, to generate a diagonal matrix, scalar matrix, or identity matrix.

[0182] i. *Exemplary Execution*

[0183] **Figure 24** illustrates an exemplary execution of a TILEDIAGONAL instruction. The TILEDIAGONAL instruction 2402 format includes fields for an opcode, a source operand identifier, and a destination matrix (tile) operand identifier (shown as “DESTINATION MATRIX (TILE)”).

[0184] The source operand 2404, as shown, identifies a register such as, for example, a general purpose register of a processor’s register file. The destination matrix (tile) operand fields represent a destination matrix (tile) 2410. As detailed earlier, a matrix (tile) may be stored in a collection of registers, locations in memory (e.g., as strided rows), or in other storage accessible to execution circuitry.

[0185] As shown, execution circuitry 2406 uses a grid of fused multiply-add units (GRID OF FMAS 2408) to execute a decoded TILEDIAGONAL instruction to store an identified source operand 2404 to every element along the main diagonal of destination matrix (tile) 2410.

[0186] Also shown are remaining (unconfigured) columns and rows being set to zero, which is done in some embodiments. In some embodiments, a matrix (tile) is configured to

use only a subset of the rows and columns possible. For example, a matrix (tile) may have up to 16 rows and columns to use, but only uses 4 of each. The configuration of each matrix (tile) is typically done by the execution of a configuration instruction prior to matrix (tile) usage. In this example, there are N columns and M rows possible.

[0187] *ii. Exemplary Format(s)*

[0188] An embodiment of a format for a TILEDIAGONAL instruction is OPCODE DESTINATION MATRIX (TILE) IDENTIFIER, SOURCE OPERAND IDENTIFIER. In some embodiments, TILEDIAGONAL [A]{B/W/D/Q} is the opcode mnemonic of the instruction where B/W/D/Q is an optional field to represent data element sizes (byte, word, double word, quadword) of the source scalar value and the destination matrix (tile) elements, and where A is an optional prefix that indicates that an antidiagonal is to be generated, rather than a main diagonal. DESTINATION MATRIX (TILE) IDENTIFIER is a field for the destination matrix (tile) operand. SOURCE OPERAND IDENTIFIER is a field for the source operand identifier. In some embodiments, the SOURCE OPERAND IDENTIFIER field is a R/M value (such as 2846), the destination matrix (tile) field is REG 2844, and the data element size is found in 2865.

[0189] In some embodiments, encodings of the instruction include a scale-index-base (SIB) type memory addressing operand that indirectly identifies multiple indexed destination locations in memory. In one embodiment, an SIB type memory operand may include an encoding identifying a base address register. The contents of the base address register may represent a base address in memory from which the addresses of the particular destination locations in memory are calculated. For example, the base address may be the address of the first location in a block of potential destination locations for an extended vector instruction. In one embodiment, an SIB type memory operand may include an encoding identifying an index register. Each element of the index register may specify an index or offset value usable to compute, from the base address, an address of a respective destination location within a block of potential destination locations. In one embodiment, an SIB type memory operand may include an encoding specifying a scaling factor to be applied to each index value when computing a respective destination address. For example, if a scaling factor value of four is encoded in the SIB type memory operand, each index value obtained from an element of the index register may be multiplied by four and then added to the base address to compute a destination address.

[0190] In one embodiment, an SIB type memory operand of the form vm32{x, y, z} may identify a vector array of memory operands specified using SIB type memory addressing. In

this example, the array of memory addresses is specified using a common base register, a constant scaling factor, and a vector index register containing individual elements, each of which is a 32-bit index value. The vector index register may be a 128-bit register (e.g., XMM) register (vm32x), a 256-bit (e.g., YMM) register (vm32y), or a 512-bit (e.g., ZMM) register (vm32z). In another embodiment, an SIB type memory operand of the form vm64{x, y, z} may identify a vector array of memory operands specified using SIB type memory addressing. In this example, the array of memory addresses is specified using a common base register, a constant scaling factor, and a vector index register containing individual elements, each of which is a 64-bit index value. The vector index register may be a 128-bit register (e.g., XMM) register (vm64x), a 256-bit (e.g., YMM) register (vm64y) or a 512-bit (e.g., ZMM) register (vm64z).

[0191] *iii. Exemplary Method(s) of Execution*

[0192] **Figure 25** illustrates an embodiment of method performed by a processor to process a TILEDIAGONAL instruction.

[0193] At 2501, an instruction is fetched. For example, a TILEDIAGONAL instruction is fetched. The TILEDIAGONAL instruction includes fields for an opcode, a source operand identifier, and a destination matrix (tile) operand identifier. In some embodiments, the instruction is fetched from an instruction cache. The opcode of the TILEDIAGONAL instruction indicates populating a main diagonal of an identified destination matrix (tile) operand is to occur, and a size of the data to be stored (written).

[0194] The fetched instruction is decoded at 2503. For example, the fetched TILEDIAGONAL instruction is decoded by decode circuitry such as that detailed herein.

[0195] Data values associated with the identified source operand of the decoded instruction are retrieved at 2505 and the decoded instruction is scheduled (as needed). For example, when the identified source operand is a memory location, the data from the indicated memory location is retrieved.

[0196] At 2507, the decoded instruction is executed by execution circuitry (hardware) such as that detailed herein. For the TILEDIAGONAL instruction, the execution will cause execution circuitry to store (write) the identified source operand to every element along the main diagonal of the destination matrix (tile). In some embodiments, unconfigured elements of rows of the destination matrix (tile) are zeroed as are elements not on the diagonal.

[0197] In some embodiments, the instruction is committed or retired at 2509.

[0198] **Figure 26** illustrates a more detailed description of an execution of a TILEDIAGONAL instruction. Typically, this is performed by execution circuitry such as that detailed above.

[0199] At 2602, a determination is made as to whether the destination matrix (tile) elements have the same size as the identified source operand. If not, then a fault is raised at 2604. 2602 is optional, as signified by its dashed borders in **Figure 26** (dashed borders are used herein to identify optional items.)

[0200] If a positive determination is made at 2602, the execution circuitry at 2606 loops with loop index x equal to zero (0) to the Minimum of (dest.rows and dest.columns). For example, if the destination matrix (tile) has four rows and five columns, x will go from zero to four. On each loop iteration, at 2608, the execution circuitry stores (writes) the identified source operand to the destination matrix (tile) at element $[x][x]$. At 2610 the execution circuit increments x , and determines whether at least one more row and at least one more column of the destination matrix (tile) remain. If so, the execution circuit returns to the start of the loop at 2606. If it is determined at 2610 that there is not at least one remaining row and at least one remaining column, the process ends. In this way, the execution circuit, over the course of the loop, stores (writes) the identified source operand to every element along the main diagonal of the destination matrix (tile).

[0201] *iv. Exemplary Pseudocode*

[0202] **Figure 27** is exemplary pseudocode describing an embodiment of a method performed by a processor to process a TILEDIAGONALD instruction. As shown in pseudocode 2702, the TILEDIAGONALD instruction includes an opcode, a source operand identifier SRC, and a destination operand TDEST to identify a destination matrix (tile). As shown, the pseudocode 2702 first causes the execution circuitry to generate a fault if any of three error checks fails. Then the pseudocode causes the processor to loop for a number of LOOP_ITERATIONS equaling the MINIMUM of the number of rows and the number of columns of the destination matrix (tile). At each iteration, the processor sets the double word at destination element $[x][x]$ to the value of the source operand. The TILEDIAGONALD opcode includes a “D” suffix, indicating that the elements of the destination matrix (tile) are each the size of a doubleword. Pseudocode 2704 operates similarly to pseudocode 2702, but has a “W” suffix, indicating that its destination matrix (tile) elements are each two bytes in size.

[0203] *v. Examples*

[0204] Example 1 A processor comprising: decode circuitry to decode an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and execution circuitry to execute the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand.

[0205] Example 2 The processor of example 1, wherein the opcode defines a size of each data element of the destination matrix operand.

[0206] Example 3 The processor of example 2, wherein the size of each data element of the destination matrix operand is a doubleword.

[0207] Example 4 The processor of example 2, wherein the size of each data element of the destination matrix operand is a word.

[0208] Example 5 The processor of any of examples 1-4, wherein the execution circuitry is further to zero any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.

[0209] Example 6 The processor of any of examples 1-5, wherein the destination matrix operand is a plurality of registers to represent a matrix.

[0210] Example 7 The processor of any of examples 1-5, wherein the execution circuitry is to fault upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.

[0211] Example 8 A method comprising: decoding an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and executing the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand.

[0212] Example 9 The method of example 8, wherein the opcode defines a size of each data element of the destination matrix operand.

[0213] Example 10 The method of example 9, wherein the size of each data element of the destination matrix operands is a doubleword.

[0214] Example 11 The method of example 9, wherein the size of each data element of the destination matrix operands is a word.

[0215] Example 12 The method of any of examples 8-11, further comprising zeroing any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.

- [0216] Example 13 The method of any of examples 8-12, wherein the identified destination matrix operand is a plurality of registers to represent a matrix.
- [0217] Example 14 The method of any of examples 8-13, further comprising faulting upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.
- [0218] Example 15 A non-transitory machine-readable medium storing an instruction which causes a processor to perform a method, the method comprising: decoding an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and executing the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand.
- [0219] Example 16 The non-transitory machine-readable medium of example 15, wherein the opcode defines a size of each data element of the destination matrix operand.
- [0220] Example 17 The non-transitory machine-readable medium of example 16, wherein the size of each data element of the destination matrix operand is a doubleword.
- [0221] Example 18 The non-transitory machine-readable medium of example 16, wherein the size of each data element of the destination matrix operand is a word.
- [0222] Example 19 The non-transitory machine-readable medium of any of examples 15-18, wherein the method further comprises zeroing any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.
- [0223] Example 20 The non-transitory machine-readable medium of any of examples 15-19, wherein the identified destination matrix operand is a plurality of registers to represent a matrix.
- [0224] Example 21 The non-transitory machine-readable medium of any of examples 15-20, wherein the method further comprises faulting upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.
- [0225] Example 22 A system comprising: a processor; and an accelerator coupled to the processor, the accelerator including: decode circuitry to decode an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and

execution circuitry to execute the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand.

[0226] Example 23 The system of example 22, wherein the opcode defines a size of each data element of the destination matrix operand.

[0227] Example 24 The system of any of examples 22-23, wherein the execution circuitry is further to zero any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.

[0228] Example 25 The system of any of examples 22-24, wherein the destination matrix operand is a plurality of registers to represent a matrix.

[0229] V. Detailed Exemplary Systems, Processors, and Emulation

[0230] Detailed herein are examples of hardware, software, etc. to execute the above described instructions. For example, what is described below details aspects of instruction execution including various pipeline stages such as fetch, decode, schedule, execute, retire, etc.

[0231] An instruction set includes one or more instruction formats. A given instruction format defines various fields (number of bits, location of bits) to specify, among other things, the operation to be performed (opcode) and the operand(s) on which that operation is to be performed. Some instruction formats are further broken down 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 a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands.

[0232] A. Exemplary Instruction Formats

[0233] 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.

[0234] VEX Instruction Format

[0235] VEX encoding allows instructions to have more than two operands, and allows SIMD vector registers to be longer than 128 bits. The use of a VEX prefix provides for three-operand (or more) syntax. For example, previous two-operand instructions performed operations such as $A = A + B$, which overwrites a source operand. The use of a VEX prefix enables operands to perform nondestructive operations such as $A = B + C$.

[0236] **Figure 28A** illustrates an exemplary instruction format including a VEX prefix 2802, real opcode field 2830, Mod R/M byte 2840, SIB byte 2850, displacement field 2862, and IMM8 2872. **Figure 28B** illustrates which fields from **Figure 28A** make up a full opcode field 2874 and a base operation field 2841. **Figure 28C** illustrates which fields from **Figure 28A** make up a register index field 2844.

[0237] VEX Prefix (Bytes 0-2) 2802 is encoded in a three-byte form. The first byte is the Format Field 2890 (VEX Byte 0, bits [7:0]), which contains an explicit C4 byte value (the unique value used for distinguishing the C4 instruction format). The second-third bytes (VEX Bytes 1-2) include a number of bit fields providing specific capability. Specifically, REX field 2805 (VEX Byte 1, bits [7-5]) consists of a VEX.R bit field (VEX Byte 1, bit [7] – R), VEX.X bit field (VEX byte 1, bit [6] – X), and VEX.B bit field (VEX byte 1, bit[5] – B). Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding VEX.R, VEX.X, and VEX.B. Opcode map field 2815 (VEX byte 1, bits [4:0] – mmmmm) includes content to encode an implied leading opcode byte. W Field 2864 (VEX byte 2, bit [7] – W) - is represented by the notation VEX.W, and provides different functions depending on the instruction. The role of VEX.vvvv 2820 (VEX Byte 2, bits [6:3]-vvvv) may include the following: 1) VEX.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) VEX.vvvv encodes the destination register operand, specified in 1s complement form for certain vector shifts; or 3) VEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. If VEX.L 2868 Size field (VEX byte 2, bit [2]-L) = 0, it indicates 128 bit vector; if VEX.L = 1, it indicates 256 bit vector. Prefix encoding field 2825 (VEX byte 2, bits [1:0]-pp) provides additional bits for the base operation field 2841.

[0238] Real Opcode Field 2830 (Byte 3) is also known as the opcode byte. Part of the opcode is specified in this field.

[0239] MOD R/M Field 2840 (Byte 4) includes MOD field 2842 (bits [7-6]), Reg field 2844 (bits [5-3]), and R/M field 2846 (bits [2-0]). The role of Reg field 2844 may include the following: encoding either the destination register operand or a source register operand (the rrr of Rrrr), or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 2846 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.

[0240] Scale, Index, Base (SIB) - The content of Scale field 2850 (Byte 5) includes SS2852 (bits [7-6]), which is used for memory address generation. The contents of SIB.xxx 2854 (bits [5-3]) and SIB.bbb 2856 (bits [2-0]) have been previously referred to with regard to the register indexes Xxxx and Bbbb.

[0241] The Displacement Field 2862 and the immediate field (IMM8) 2872 contain data.

[0242] B. Exemplary Register Architecture

[0243] **Figure 29** is a block diagram of a register architecture 2900 according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers 2910 that are 512 bits wide; these registers are referenced as zmm0 through zmm31. The lower order 256 bits of the lower 32 zmm registers are overlaid on registers ymm0-15. The lower order 128 bits of the lower 32 zmm registers (the lower order 128 bits of the ymm registers) are overlaid on registers xmm0-15.

[0244] General-purpose registers 2925 - 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.

[0245] Scalar floating point stack register file (x87 stack) 2945, on which is aliased the MMX packed integer flat register file 2950 - 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.

[0246] In some embodiments, tiles 2920 are supported using an overlay over physical registers. For example, a tile may utilize 16 1,024-bit registers, 32 512-bit registers, etc. depending on the implementation.

[0247] Alternative embodiments of the invention may use wider or narrower registers. Additionally, alternative embodiments of the invention may use more, less, or different register files and registers.

[0248] Exemplary Core Architectures, Processors, and Computer Architectures

[0249] 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. Detailed herein are circuits (units) that comprise exemplary cores, processors, etc.

[0250] C. Exemplary Core Architectures

[0251] In-order and out-of-order core block diagram

[0252] **Figure 30A** is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention. **Figure 30B** is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention. The solid lined boxes in **Figures 30A-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.

[0253] In **Figure 30A**, a processor pipeline 3000 includes a fetch stage 3002, a length decode stage 3004, a decode stage 3006, an allocation stage 3008, a renaming stage 3010, a scheduling (also known as a dispatch or issue) stage 3012, a register read/memory read stage 3014, an execute stage 3016, a write back/memory write stage 3018, an exception handling stage 3022, and a commit stage 3024.

[0254] **Figure 30B** shows processor core 3090 including a front end unit 3030 coupled to an execution engine unit 3050, and both are coupled to a memory unit 3070. The core 3090 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 3090 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.

[0255] The front end unit 3030 includes a branch prediction unit 3032 coupled to an instruction cache unit 3034, which is coupled to an instruction translation lookaside buffer (TLB) 3036, which is coupled to an instruction fetch unit 3038, which is coupled to a decode unit 3040. The decode unit 3040 (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 3040 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 3090 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 3040 or otherwise within the front end unit 3030). The decode unit 3040 is coupled to a rename/allocator unit 3052 in the execution engine unit 3050.

[0256] The execution engine unit 3050 includes the rename/allocator unit 3052 coupled to a retirement unit 3054 and a set of one or more scheduler unit(s) 3056. The scheduler unit(s) 3056 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 3056 is coupled to the physical register file(s) unit(s) 3058. Each of the physical register file(s) units 3058 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 3058

comprises a vector 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) 3058 is overlapped by the retirement unit 3054 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 3054 and the physical register file(s) unit(s) 3058 are coupled to the execution cluster(s) 3060. The execution cluster(s) 3060 includes a set of one or more execution units 3062 and a set of one or more memory access units 3064. The execution units 3062 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) 3056, physical register file(s) unit(s) 3058, and execution cluster(s) 3060 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) 3064). 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.

[0257] The set of memory access units 3064 is coupled to the memory unit 3070, which includes a data TLB unit 3072 coupled to a data cache unit 3074 coupled to a level 2 (L2) cache unit 3076. In one exemplary embodiment, the memory access units 3064 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 3072 in the memory unit 3070. The instruction cache unit 3034 is further coupled to a level 2 (L2) cache unit 3076 in the memory unit 3070. The L2 cache unit 3076 is coupled to one or more other levels of cache and eventually to a main memory.

[0258] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 3000 as follows: 1) the instruction fetch 3038 performs the fetch and length decoding stages 3002 and 3004; 2) the decode unit 3040

performs the decode stage 3006; 3) the rename/allocator unit 3052 performs the allocation stage 3008 and renaming stage 3010; 4) the scheduler unit(s) 3056 performs the schedule stage 3012; 5) the physical register file(s) unit(s) 3058 and the memory unit 3070 perform the register read/memory read stage 3014; the execution cluster 3060 perform the execute stage 3016; 6) the memory unit 3070 and the physical register file(s) unit(s) 3058 perform the write back/memory write stage 3018; 7) various units may be involved in the exception handling stage 3022; and 8) the retirement unit 3054 and the physical register file(s) unit(s) 3058 perform the commit stage 3024.

[0259] The core 3090 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 3090 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.

[0260] 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).

[0261] 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 3034/3074 and a shared L2 cache unit 3076, 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.

[0262] Specific Exemplary In-Order Core Architecture

[0263] **Figures 31A-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.

[0264] **Figure 31A** is a block diagram of a single processor core, along with its connection to the on-die interconnect network 3102 and with its local subset of the Level 2 (L2) cache 3104, according to embodiments of the invention. In one embodiment, an instruction decoder 3100 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 3106 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 3108 and a vector unit 3110 use separate register sets (respectively, scalar registers 3112 and vector registers 3114) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 3106, alternative embodiments of the invention may use a different approach (e.g., use a single register set or include a communication path that allow data to be transferred between the two register files without being written and read back).

[0265] The local subset of the L2 cache 3104 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 3104. Data read by a processor core is stored in its L2 cache subset 3104 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 3104 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 1024-bits wide per direction in some embodiments.

[0266] **Figure 31B** is an expanded view of part of the processor core in **Figure 31A** according to embodiments of the invention. **Figure 31B** includes an L1 data cache 3106A part of the L1 cache 3104, as well as more detail regarding the vector unit 3110 and the vector registers 3114. Specifically, the vector unit 3110 is a 32-wide vector processing unit (VPU) (see the 16-wide ALU 3128), 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 3120, numeric conversion with numeric convert units 3122A-B, and replication with replication unit 3124 on the memory input.

[0267] Processor with integrated memory controller and graphics

[0268] **Figure 32** is a block diagram of a processor 3200 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention. The solid lined boxes in **Figure 32** illustrate a processor

3200 with a single core 3202A, a system agent 3210, a set of one or more bus controller units 3216, while the optional addition of the dashed lined boxes illustrates an alternative processor 3200 with multiple cores 3202A-N, a set of one or more integrated memory controller unit(s) 3214 in the system agent unit 3210, and special purpose logic 3208.

[0269] Thus, different implementations of the processor 3200 may include: 1) a CPU with the special purpose logic 3208 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 3202A-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 3202A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 3202A-N being a large number of general purpose in-order cores. Thus, the processor 3200 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 3200 may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0270] The memory hierarchy includes one or more levels of cache within the cores 3204A-N, a set of one or more shared cache units 3206, and external memory (not shown) coupled to the set of integrated memory controller units 3214. The set of shared cache units 3206 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 3212 interconnects the integrated graphics logic 3208, the set of shared cache units 3206, and the system agent unit 3210/integrated memory controller unit(s) 3214, 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 3206 and cores 3202-A-N.

[0271] In some embodiments, one or more of the cores 3202A-N are capable of multi-threading. The system agent 3210 includes those components coordinating and operating cores 3202A-N. The system agent unit 3210 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 3202A-N and the integrated graphics logic 3208. The display unit is for driving one or more externally connected displays.

[0272] The cores 3202A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 3202A-N may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

[0273] D. Exemplary Computer Architectures

[0274] **Figures 33-36** 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, handheld 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.

[0275] Referring now to **Figure 33**, shown is a block diagram of a system 3300 in accordance with one embodiment of the present invention. The system 3300 may include one or more processors 3310, 3315, which are coupled to a controller hub 3320. In one embodiment, the controller hub 3320 includes a graphics memory controller hub (GMCH) 3390 and an Input/Output Hub (IOH) 3350 (which may be on separate chips); the GMCH 3390 includes memory and graphics controllers to which are coupled memory 3340 and a coprocessor 3345; the IOH 3350 is couples input/output (I/O) devices 3360 to the GMCH 3390. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory 3340 and the coprocessor 3345 are coupled directly to the processor 3310, and the controller hub 3320 in a single chip with the IOH 3350.

[0276] The optional nature of additional processors 3315 is denoted in **Figure 33** with broken lines. Each processor 3310, 3315 may include one or more of the processing cores described herein and may be some version of the processor 3200.

[0277] The memory 3340 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 3320 communicates with the processor(s) 3310, 3315 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface, or similar connection 3395.

[0278] In one embodiment, the coprocessor 3345 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 3320 may include an integrated graphics accelerator.

[0279] There can be a variety of differences between the physical resources 3310, 33155 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

[0280] In one embodiment, the processor 3310 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 3310 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 3345. Accordingly, the processor 3310 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 3345. Coprocessor(s) 3345 accept and execute the received coprocessor instructions.

[0281] Referring now to **Figure 34**, shown is a block diagram of a first more specific exemplary system 3400 in accordance with an embodiment of the present invention. As shown in **Figure 34**, multiprocessor system 3400 is a point-to-point interconnect system, and includes a first processor 3470 and a second processor 3480 coupled via a point-to-point interconnect 3450. Each of processors 3470 and 3480 may be some version of the processor 3200. In one embodiment of the invention, processors 3470 and 3480 are respectively processors 3310 and 3315, while coprocessor 3438 is coprocessor 3345. In another embodiment, processors 3470 and 3480 are respectively processor 3310 coprocessor 3345.

[0282] Processors 3470 and 3480 are shown including integrated memory controller (IMC) units 3472 and 3482, respectively. Processor 3470 also includes as part of its bus controller units point-to-point (P-P) interfaces 3476 and 3478; similarly, second processor 3480 includes P-P interfaces 3486 and 3488. Processors 3470, 3480 may exchange information via a point-to-point (P-P) interface 3450 using P-P interface circuits 3478, 3488. As shown in **Figure 34**, IMCs 3472 and 3482 couple the processors to respective memories, namely a memory 3432 and a memory 3434, which may be portions of main memory locally attached to the respective processors.

[0283] Processors 3470, 3480 may each exchange information with a chipset 3490 via individual P-P interfaces 3452, 3454 using point to point interface circuits 3476, 3494, 3486, 3498. Chipset 3490 may optionally exchange information with the coprocessor 3438 via a high-performance interface 3492. In one embodiment, the coprocessor 3438 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.

[0284] 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.

[0285] Chipset 3490 may be coupled to a first bus 3416 via an interface 3496. In one embodiment, first bus 3416 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another I/O interconnect bus, although the scope of the present invention is not so limited.

[0286] As shown in **Figure 34**, various I/O devices 3414 may be coupled to first bus 3416, along with a bus bridge 3418 which couples first bus 3416 to a second bus 3420. In one embodiment, one or more additional processor(s) 3415, such as coprocessors, high-throughput MIC processors, GPGPU's, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor, are coupled to first bus 3416. In one embodiment, second bus 3420 may be a low pin count (LPC) bus. Various devices may be coupled to a second bus 3420 including, for example, a keyboard and/or mouse 3422, communication devices 3427 and a storage unit 3428 such as a disk drive or other mass storage device which may include instructions/code and data 3430, in one embodiment. Further, an audio I/O 3424 may be coupled to the second bus 3416. Note that other architectures are possible. For example, instead of the point-to-point architecture of **Figure 34**, a system may implement a multi-drop bus or other such architecture.

[0287] Referring now to **Figure 35**, shown is a block diagram of a second more specific exemplary system 3500 in accordance with an embodiment of the present invention. Like elements in **Figures 34** and **35** bear like reference numerals, and certain aspects of **Figure 34** have been omitted from **Figure 35** in order to avoid obscuring other aspects of **Figure 35**.

[0288] **Figure 35** illustrates that the processors 3470, 3480 may include integrated memory and I/O control logic ("CL") 3572 and 3582, respectively. Thus, the CL 3572, 3582 include integrated memory controller units and include I/O control logic. **Figure 35** illustrates that not only are the memories 3432, 3434 coupled to the CL 3572, 3582, but also that I/O devices 3514 are also coupled to the control logic 3472, 3482. Legacy I/O devices 3515 are coupled to the chipset 3490.

[0289] Referring now to **Figure 36**, shown is a block diagram of a SoC 3600 in accordance with an embodiment of the present invention. Similar elements in **Figure 32** bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In **Figure 36**, an interconnect unit(s) 3602 is coupled to: an application processor 3610 which includes a set of one or more cores 3202A-N, cache units 3204A-N, and shared cache unit(s) 3206; a system agent unit 3210; a bus controller unit(s) 3216; an integrated memory controller unit(s) 3214; a set or one or more coprocessors 3620 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 3630; a direct memory access (DMA) unit 3632; and a display unit 3640 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 3620 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.

[0290] Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the invention may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

[0291] Program code, such as code 3430 illustrated in **Figure 34**, 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.

[0292] The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

[0293] One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as “IP cores”

may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

[0294] 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.

[0295] Accordingly, embodiments of the invention also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

[0296] E. Emulation (including binary translation, code morphing, etc.)

[0297] 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.

[0298] **Figure 37** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. **Figure 37** shows a program in a high level language 3702 may be compiled using an first compiler 3704 to generate a first binary code (e.g., x86) 3706 that may be natively executed by a processor with at least one first instruction set core 3716.

In some embodiments, the processor with at least one first instruction set core 3716 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 first compiler 3704 represents a compiler that is operable to generate binary code of the first instruction set 3706 (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one first instruction set core 3716. Similarly, **Figure 37** shows the program in the high level language 3702 may be compiled using an alternative instruction set compiler 3708 to generate alternative instruction set binary code 3710 that may be natively executed by a processor without at least one first instruction set core 3714 (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 3712 is used to convert the first binary code 3706 into code that may be natively executed by the processor without a first instruction set core 3714. This converted code is not likely to be the same as the alternative instruction set binary code 3710 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 3712 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 a first instruction set processor or core to execute the first binary code 3706.

CLAIMS

What is claimed is:

1. A processor comprising:
 decode circuitry to decode an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and
 execution circuitry to execute the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand and zero other data elements.
2. The processor of claim 1, wherein the opcode defines a size of each data element of the destination matrix operand.
3. The processor of claim 2, wherein the size of each data element of the destination matrix operand is a doubleword.
4. The processor of claim 2, wherein the size of each data element of the destination matrix operand is a word.
5. The processor of any of claims 1-4, wherein the execution circuitry is further to zero any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.
6. The processor of any of claims 1-5, wherein the destination matrix operand is a plurality of registers to represent a matrix.
7. The processor of any of claims 1-5, wherein the execution circuitry is to fault upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.

8. A method comprising:
 - decoding an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and
 - executing the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand and zero other data elements.
9. The method of claim 8, wherein the opcode defines a size of each data element of the destination matrix operand.
10. The method of claim 9, wherein the size of each data element of the destination matrix operands is a doubleword.
11. The method of claim 9, wherein the size of each data element of the destination matrix operands is a word.
12. The method of any of claims 8-11, further comprising zeroing any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.
13. The method of any of claims 8-12, wherein the identified destination matrix operand is a plurality of registers to represent a matrix.
14. The method of any of claims 8-13, further comprising faulting upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.
15. A non-transitory machine-readable medium storing an instruction which causes a processor to perform a method, the method comprising:
 - decoding an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and

executing the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand and zero other data elements.

16. The non-transitory machine-readable medium of claim 15, wherein the opcode defines a size of each data element of the destination matrix operand.

17. The non-transitory machine-readable medium of claim 16, wherein the size of each data element of the destination matrix operand is a doubleword.

18. The non-transitory machine-readable medium of claim 16, wherein the size of each data element of the destination matrix operand is a word.

19. The non-transitory machine-readable medium of any of claims 15-18, wherein the method further comprises zeroing any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.

20. The non-transitory machine-readable medium of any of claims 15-19, wherein the identified destination matrix operand is a plurality of registers to represent a matrix.

21. The non-transitory machine-readable medium of any of claims 15-20, wherein the method further comprises faulting upon a determination of one of: the identified source operand having a different number of bytes than each element of the identified destination matrix operand, each element of the destination matrix operand having a different size than a size identifier included in the opcode, and the identified destination matrix operand having zero configured elements.

22. A system comprising:

a processor; and

an accelerator coupled to the processor, the accelerator including:

decode circuitry to decode an instruction having fields for an opcode, a source operand identifier, and a destination matrix operand identifier; and

execution circuitry to execute the decoded instruction to write the identified source operand to each element along a main diagonal of the identified destination matrix operand and zero other data elements.

23. The system of claim 22, wherein the opcode defines a size of each data element of the destination matrix operand.
24. The system of any of claims 22-23, wherein the execution circuitry is further to zero any remaining columns of the identified destination matrix operand and unconfigured rows of the identified destination matrix operand.
25. The system of any of claims 22-24, wherein the destination matrix operand is a plurality of registers to represent a matrix.




FIG. 1




FIG. 2

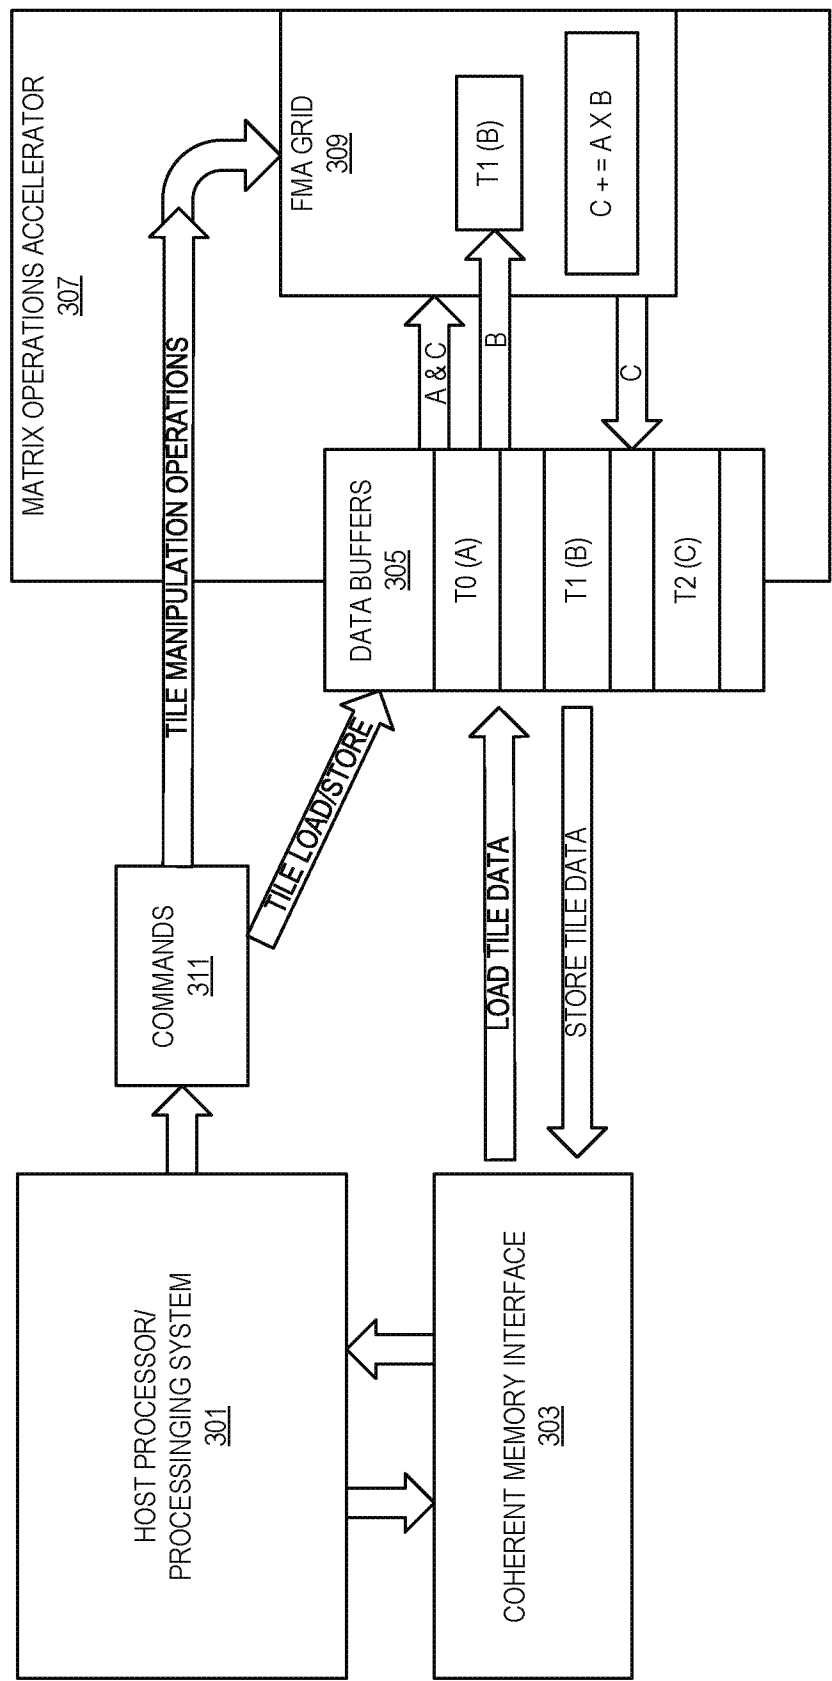


FIG. 3




FIG. 4

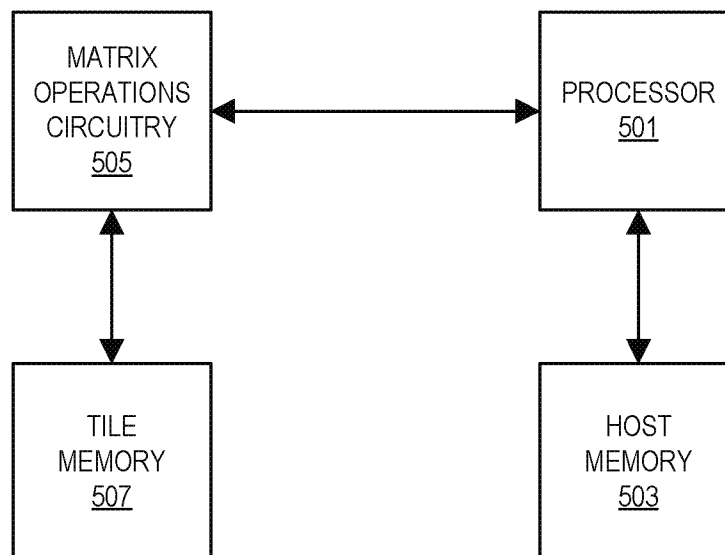


FIG. 5

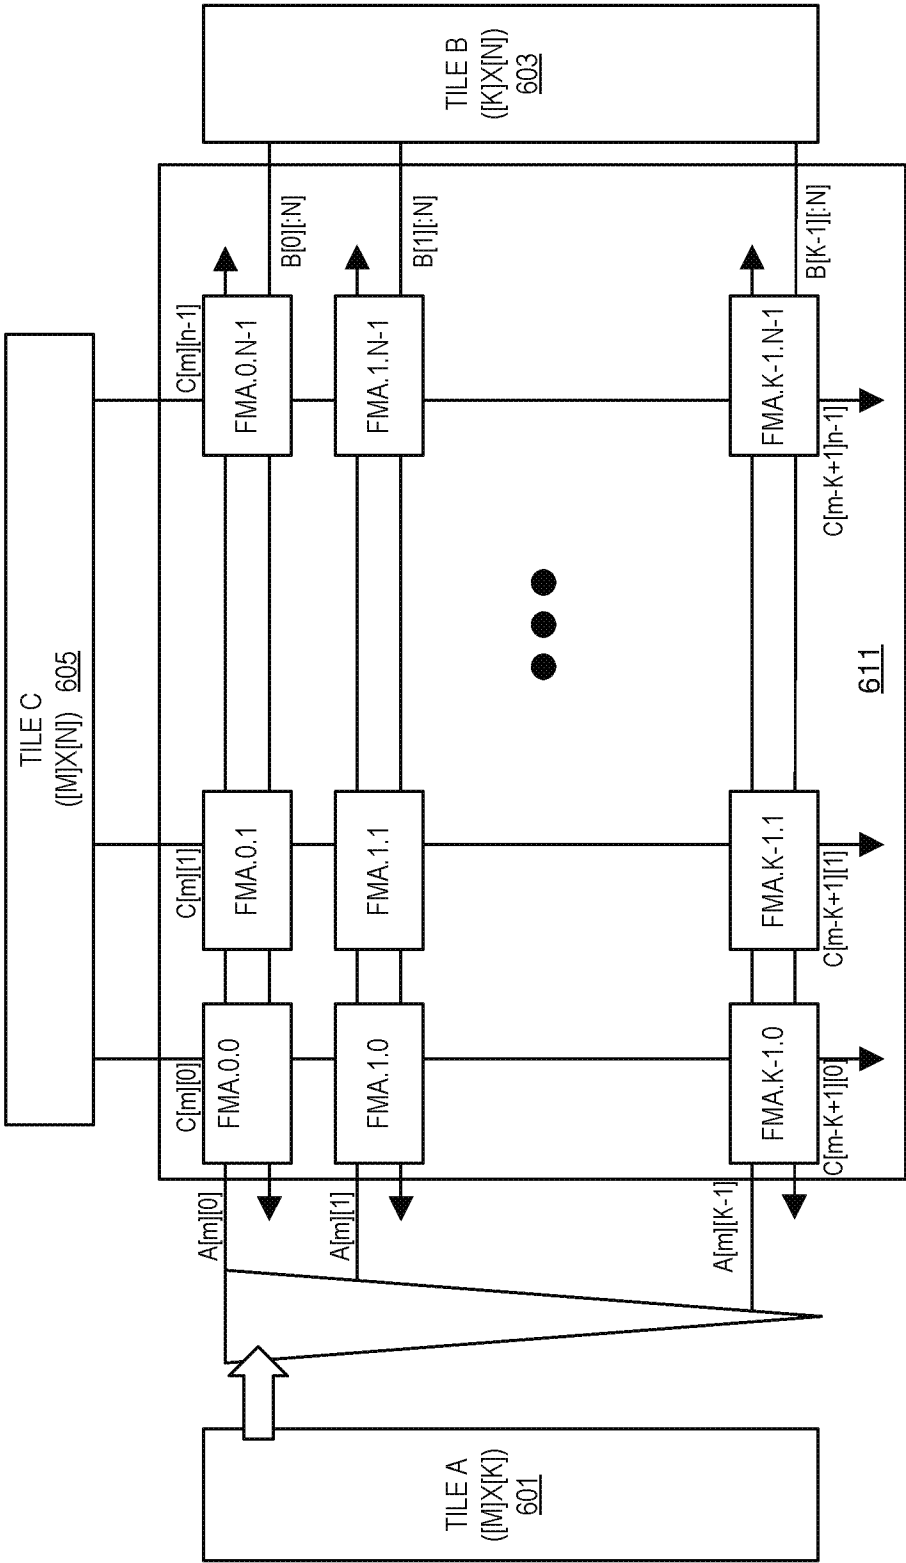


FIG. 6

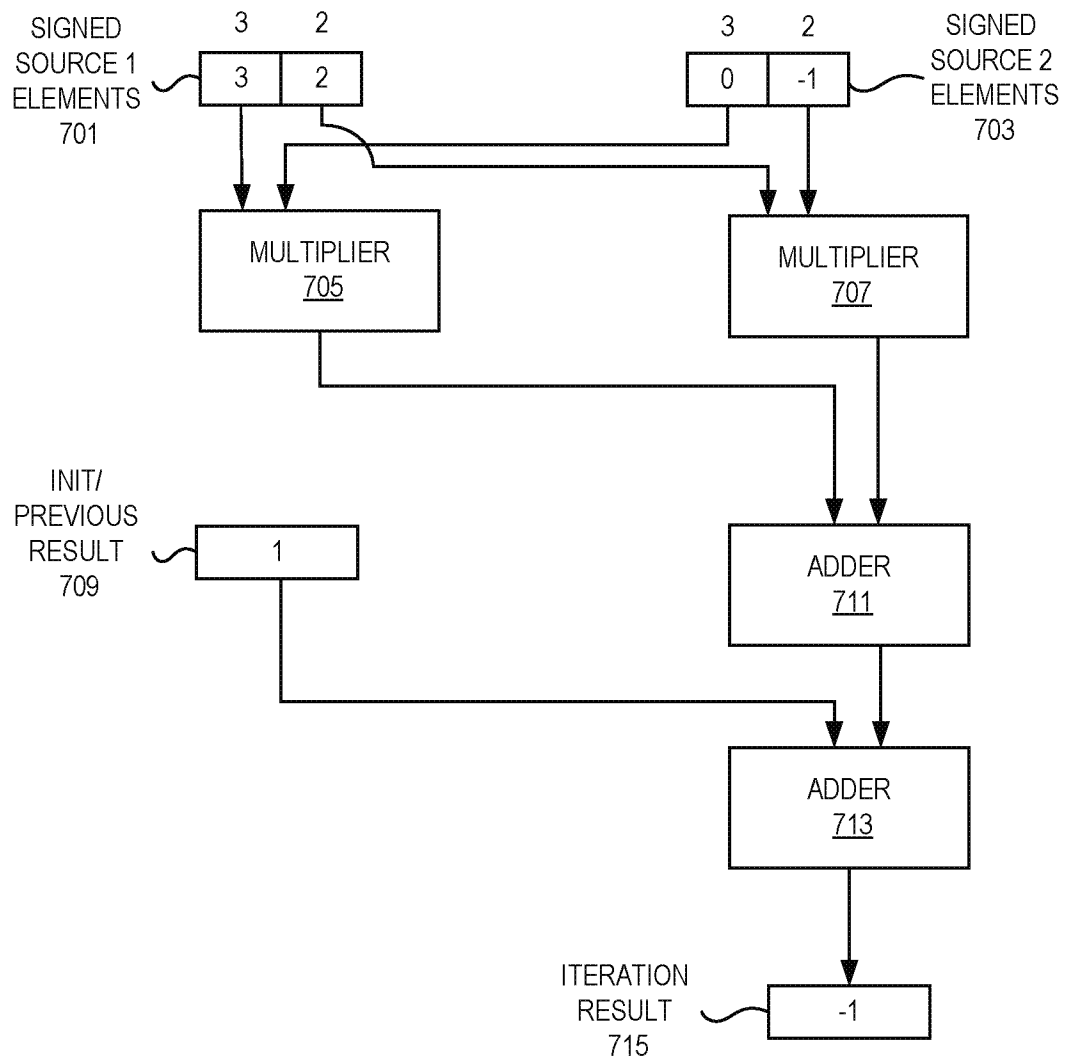


FIG. 7

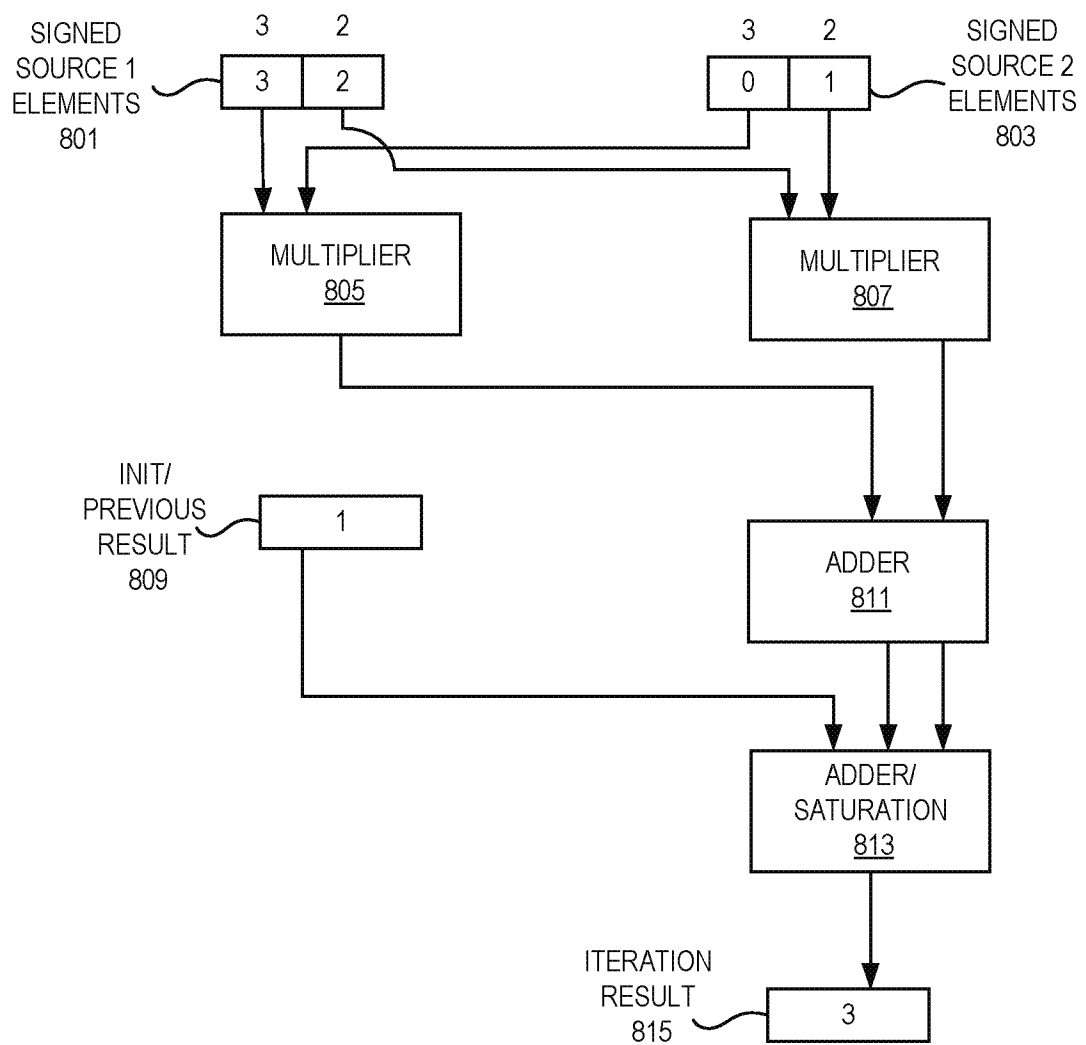


FIG. 8

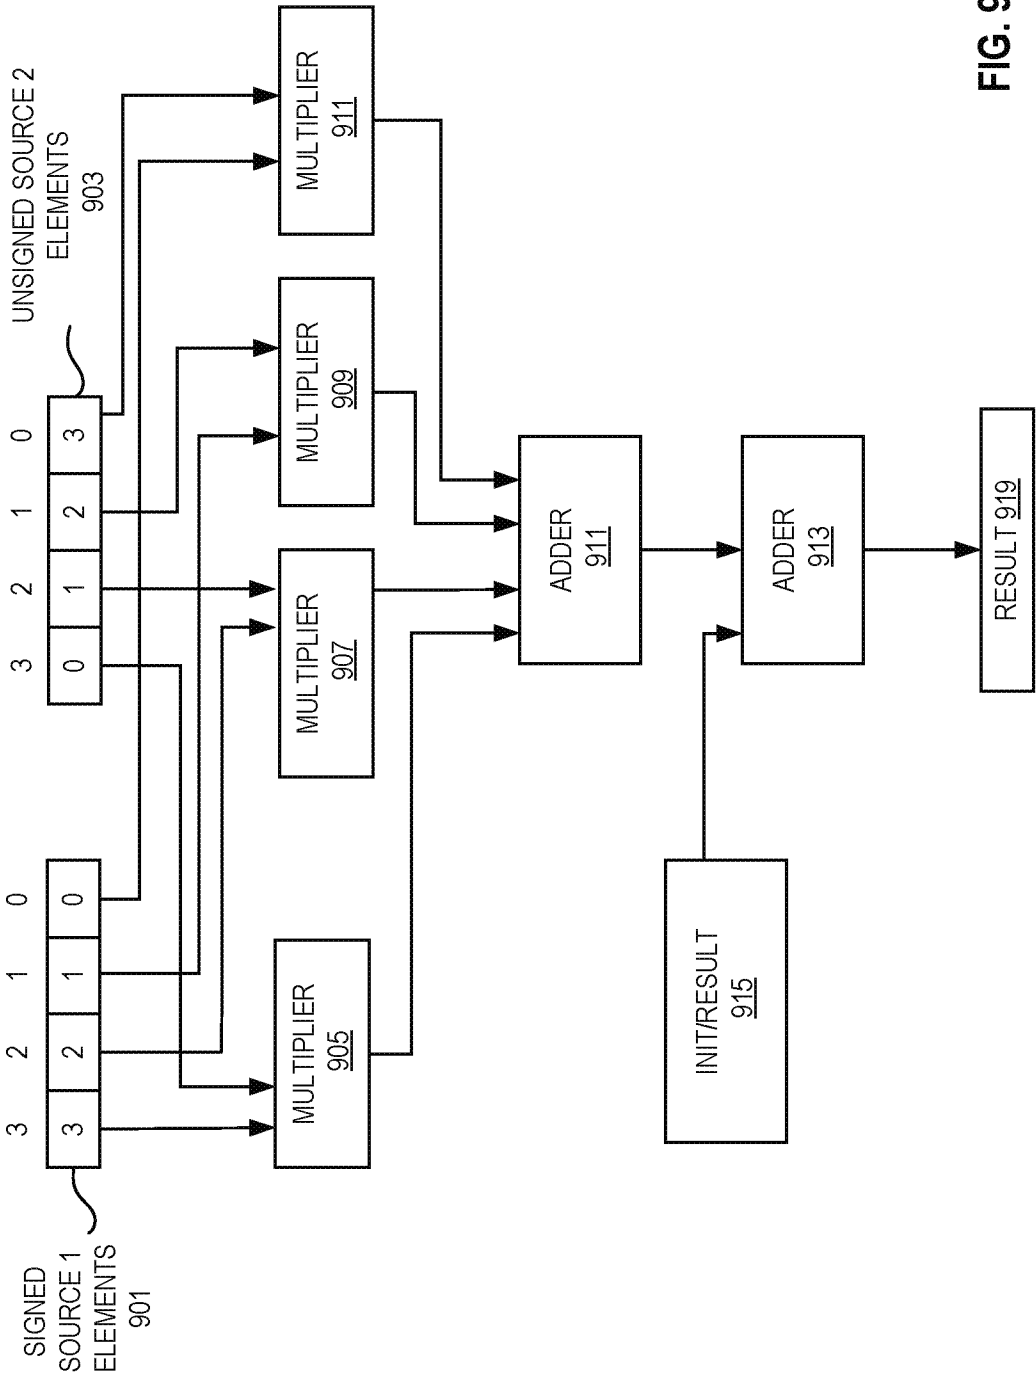


FIG. 9




FIG. 10

10/36

ACCUMULATOR 2X INPUT SIZES 1101

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	WORD/HPFP	16
WORD	16	INT32/SPFP	32
SPFP/INT32	32	INT64/DPFP	64

ACCUMULATOR 4X INPUT SIZES 1103

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	INT32/SPFP	32
WORD	16	INT64/DPFP	64

ACCUMULATOR 8X INPUT SIZES 1105

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	INT64/DPFP	64

FIG. 11

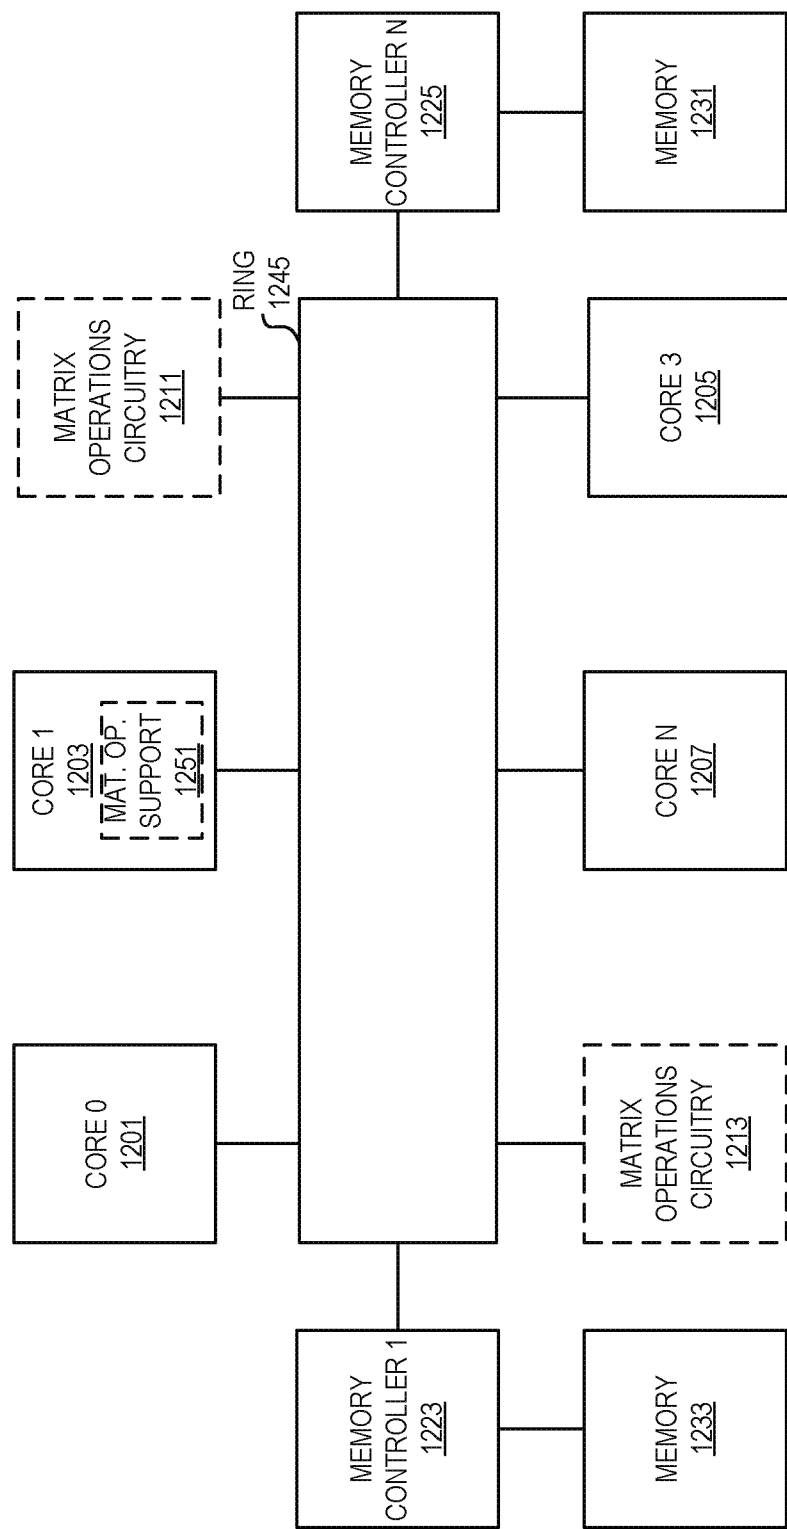


FIG. 12

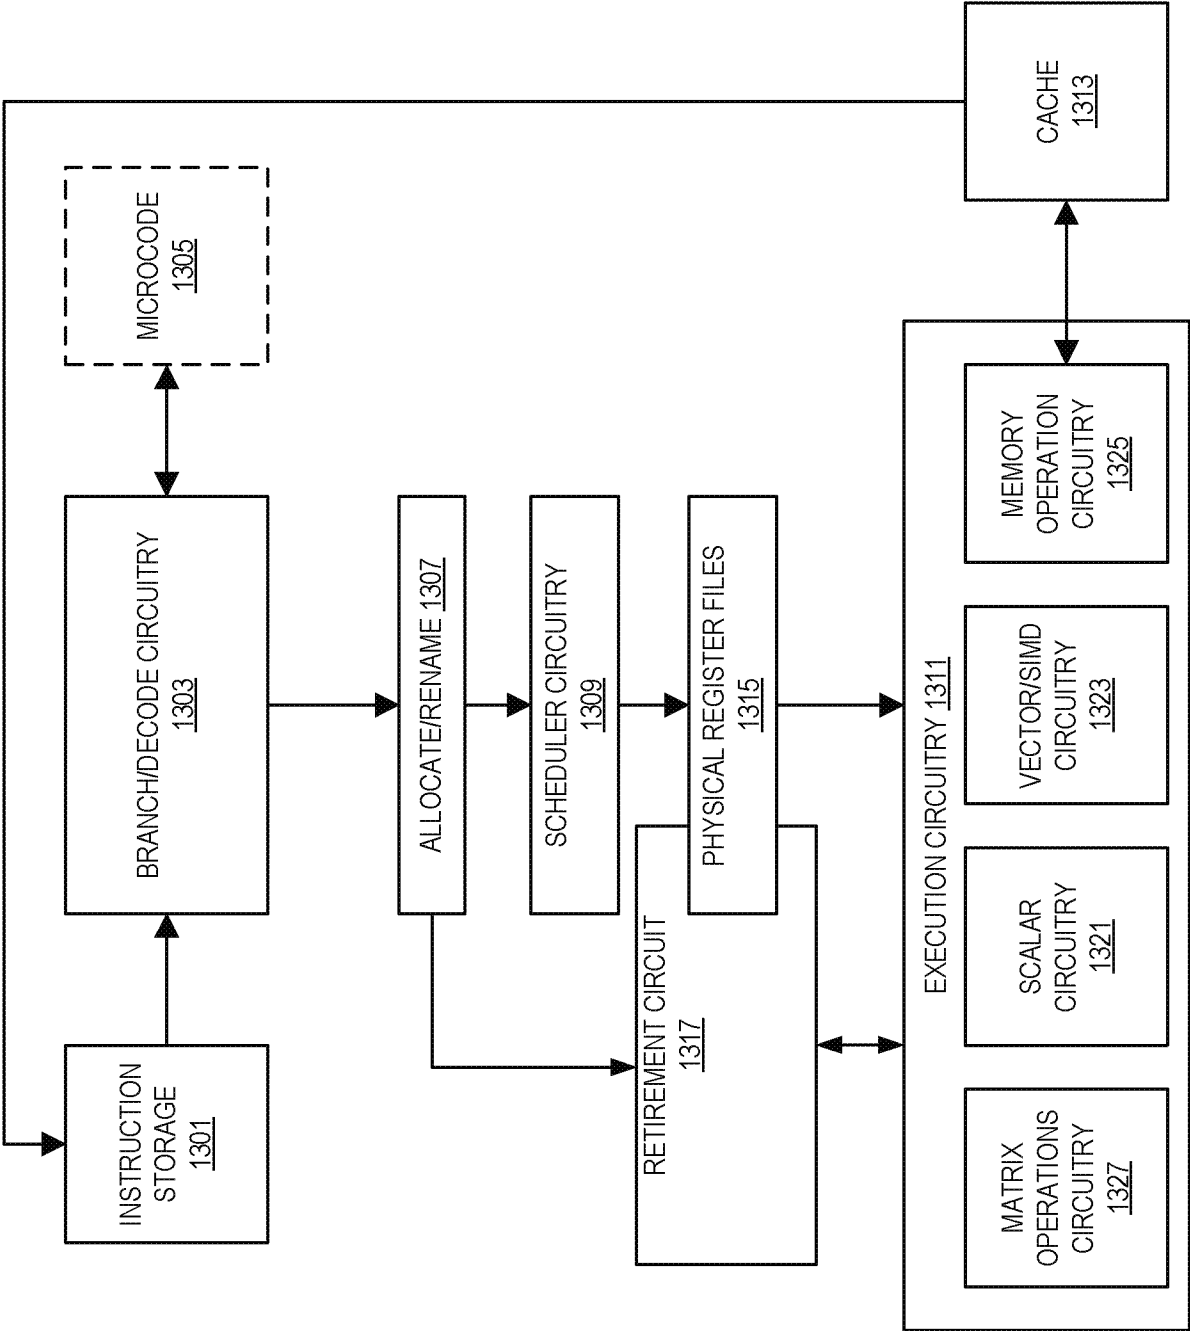


FIG. 13

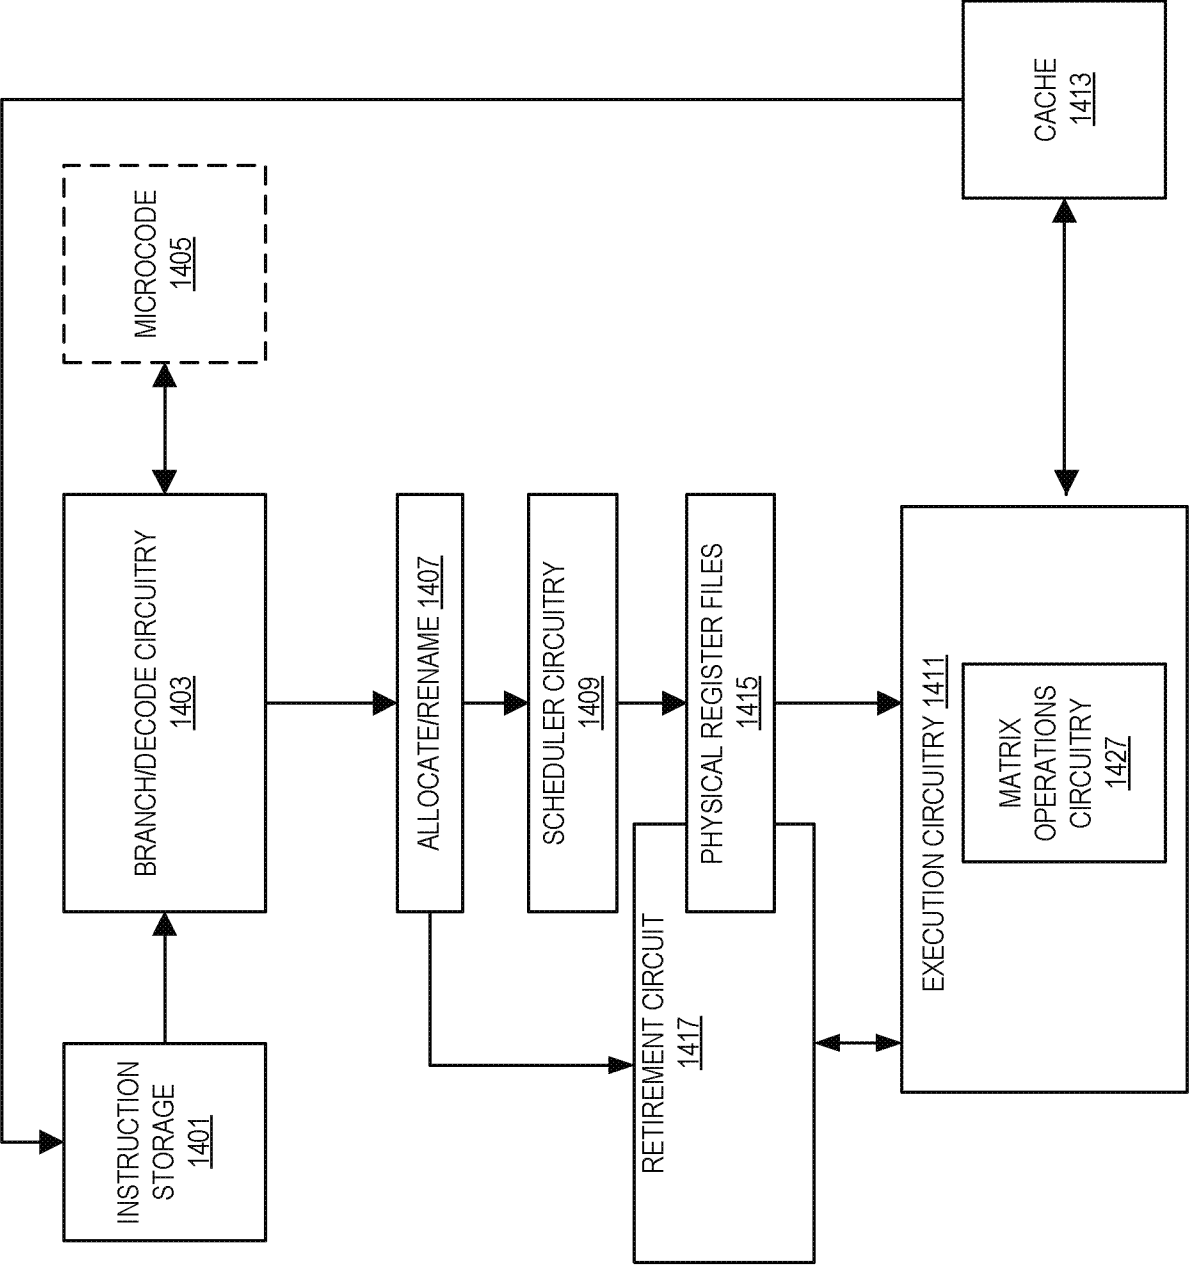


FIG. 14

$$A = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \end{bmatrix}$$


ADDR	VALUE
0	A_{11}
1	A_{12}
2	A_{13}
3	A_{21}
4	A_{22}
5	A_{23}

ROW MAJOR

ADDR	VALUE
0	A_{11}
1	A_{21}
2	A_{12}
3	A_{22}
4	A_{13}
5	A_{23}

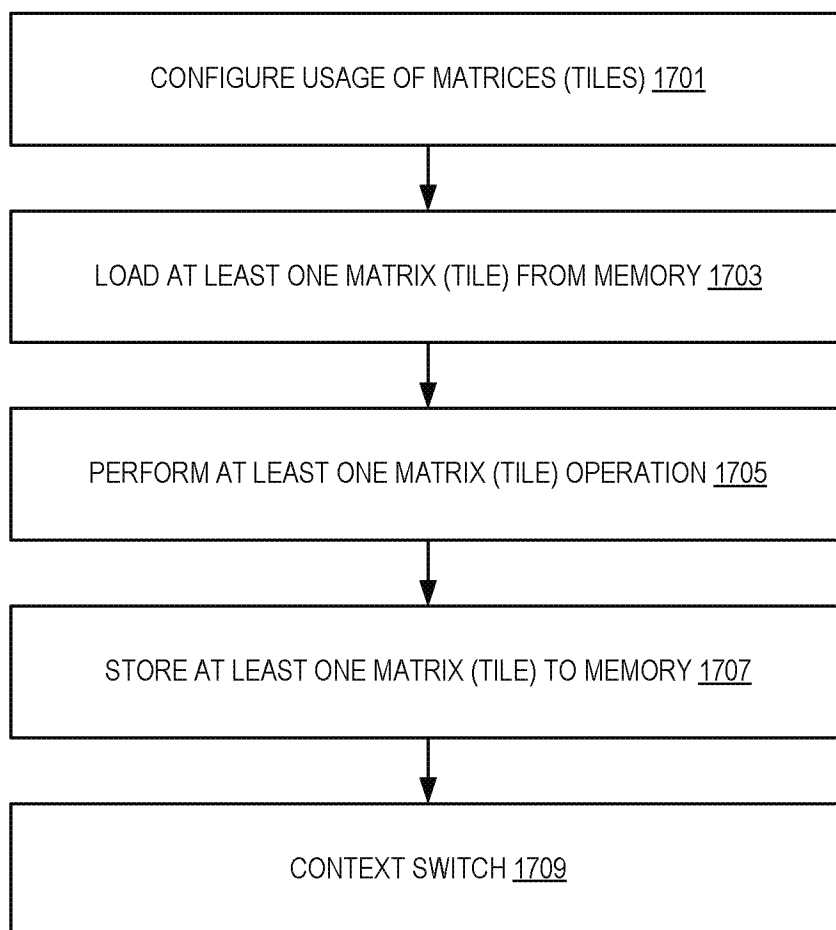
COLUMN MAJOR

FIG. 15



```
TILECONFIG [RAX]
// ASSUME SOME OUTER LOOPS DRIVING THE CACHE TILING (NOT SHOWN)
{
    TILELOAD TMM0, RSI+RDI // SRCDST, RSI POINTS TO C, RDI HAS
    TILELOAD TMM1, RSI+RDI+N // SECOND TILE OF C, UNROLLING IN SIMD DIMENSION N
    MOV KK, 0
    LOOP:
        TILELOAD TMM2, R8+R9 // SRC2 IS STRIDED LOAD OF A, REUSED FOR 2 TMM A INSTR.
        TILELOAD TMM3, R10+R11 // SRC1 IS STRIDED LOAD OF B
        TMMAPS TMM0, TMM2, TMM3 // UPDATE LEFT TILE OF C
        TILELOAD TMM3, R10+R11+N // SRC1 LOADED WITH B FROM NEXT RIGHTMOST TILE
        TMMAPS TMM1, TMM2, TMM3 // UPDATE RIGHT TILE OF C
        ADD R8, K // UPDATE POINTERS BY CONSTANTS KNOWN OUTSIDE OF LOOP
        ADD R10, K*R11
        ADD KK, K
        CMP KK, LIMIT
        JNE LOOP
        TILESTORE RSI+RDI, TMM0 // UPDATE THE C MATRIX IN MEMORY
        TILESTORE RSI+RDI+M, TMM1
    } // END OF OUTER LOOP
    TILERELASE // RETURN TILES TO INIT STATE
```

FIG. 16

**FIG. 17**

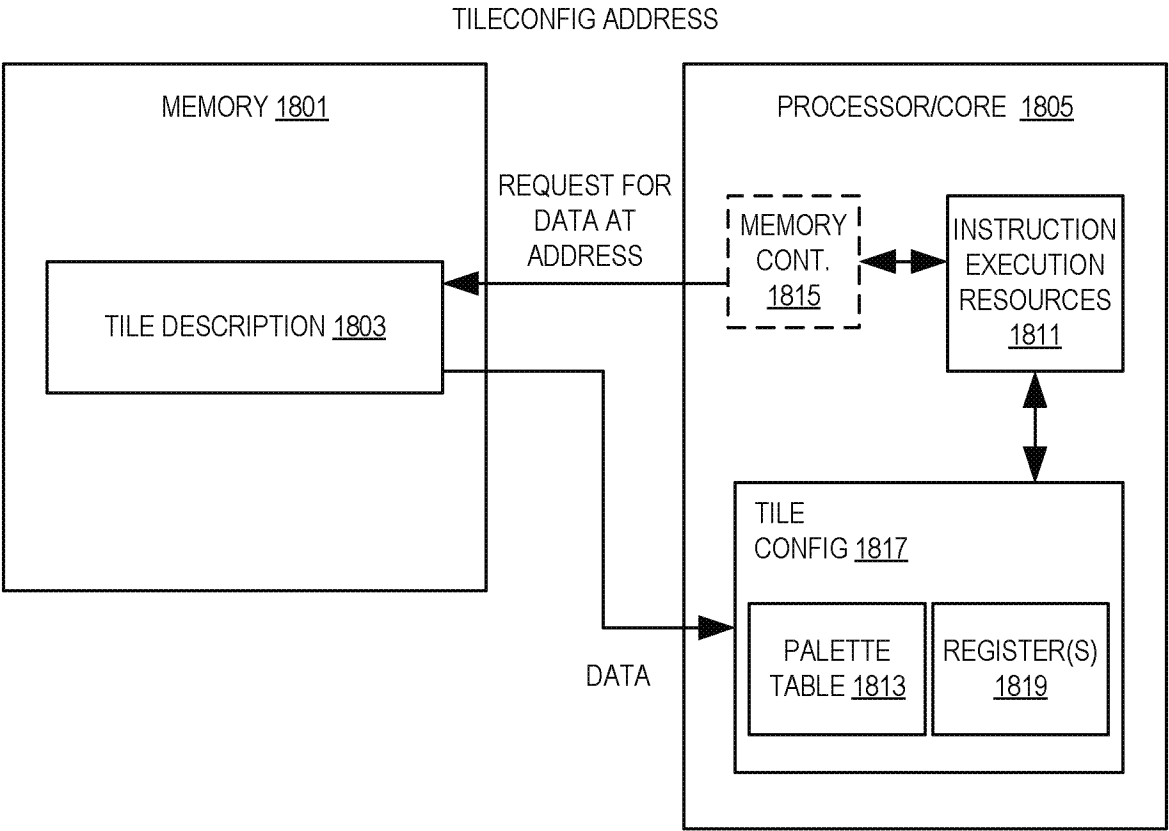


FIG. 18




FIG. 19(A)




FIG. 19(B)

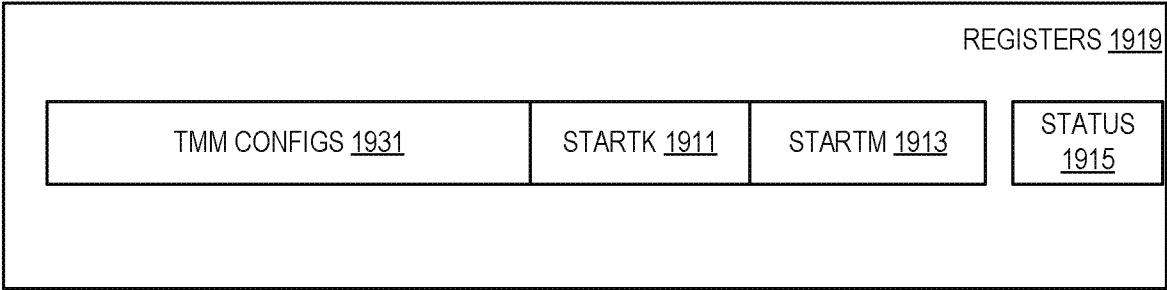


FIG. 19(C)


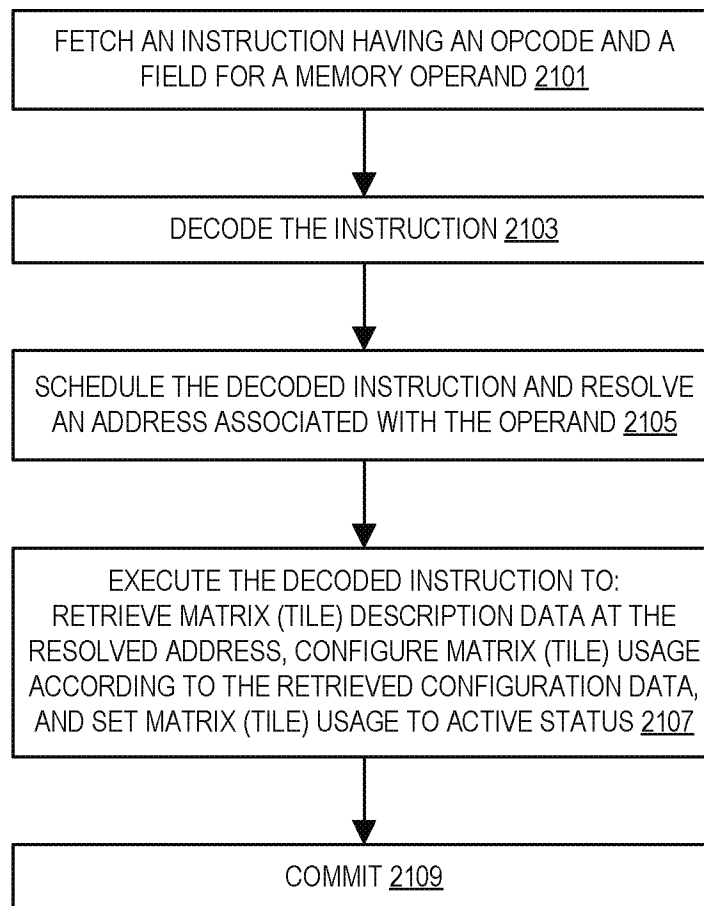


FIG. 19(D)

PALETTE ID <u>2001</u>	0
0	0
0	0
0	0
STARTM <u>2003</u>	
STARTK <u>2005</u>	
0	0
TMM0 ROWS <u>2013</u>	TMM0 COLUMNS <u>2015</u>
TMM1 ROWS	TMM1 COLUMNS
• • •	
TMM15 ROWS	TMM15 COLUMNS
0	

FIG. 20

**FIG. 21**

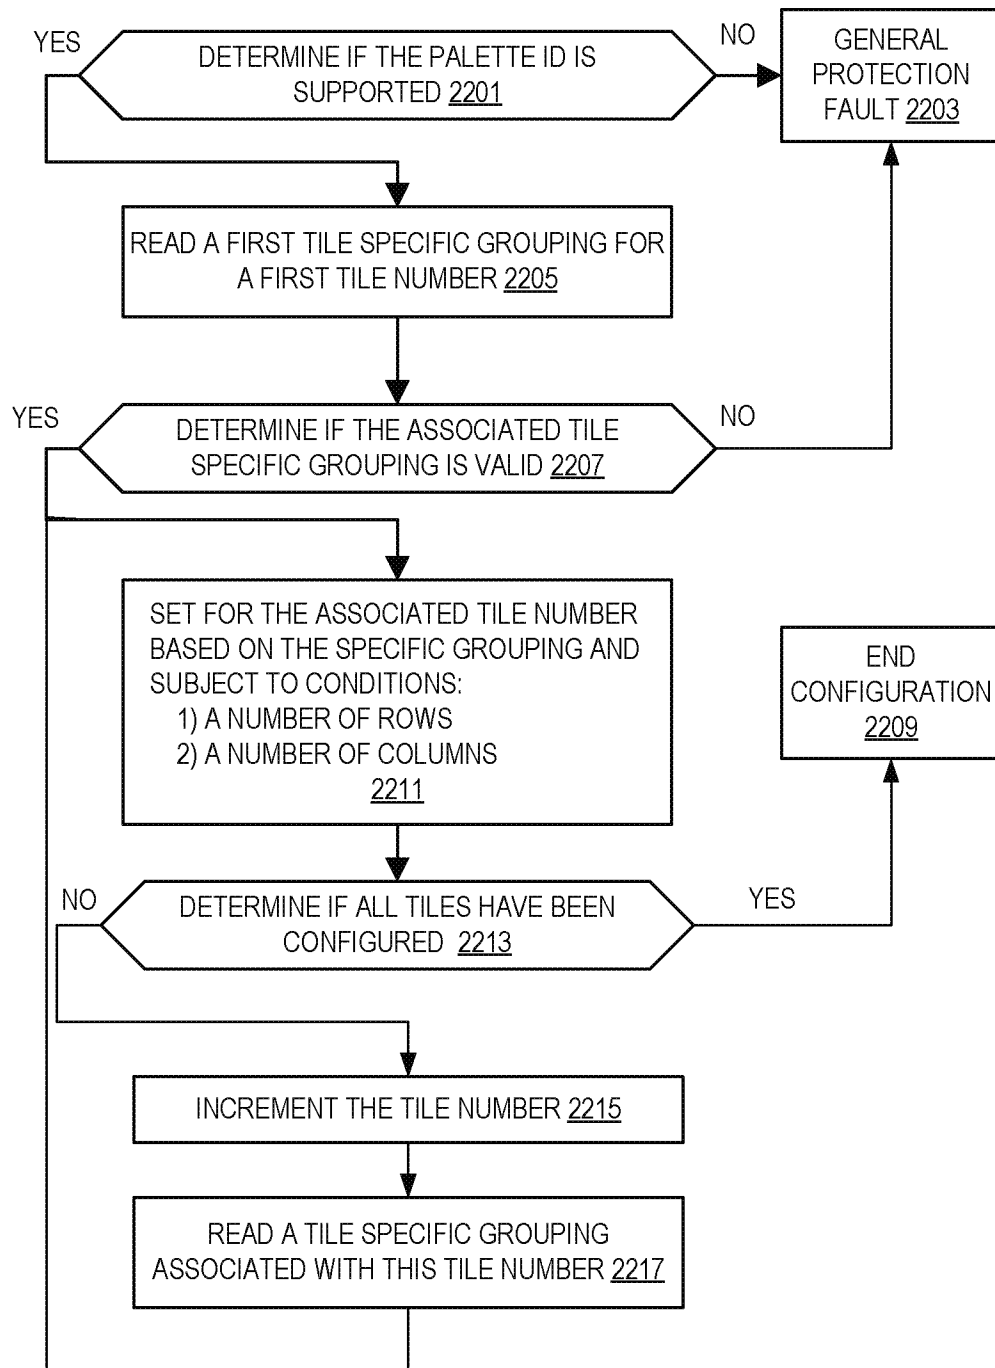


FIG. 22

```

// FORMAT OF MEMORY PAYLOAD. EACH FIELD IS A BYTE
// 0: PALETTE_ID
// 1-7: RESERVED (MUST BE ZERO)
// 8-9: STARTM(16B)
// 10-11: STARTK(16B)
// 12-15: RESERVED, (MUST BE ZERO)
// 16-17: TILE0ROWSTILE0COLS
// 18-19: TILE1ROWSTILE1COLS
// 20-21: TILE2ROWSTILE2COLS
// ...
// 46-47: TILE15ROWS TILE15COLS
// 48-63: 16B RESERVED, (MUST BE ZERO)
TILECONFIG MEM
PALETTE_ID := MEM.BYTE[0]
#GP IF PALETTE_ID IS AN UNSUPPORTED PALETTE // FROM CPUID
#GP IF MEM.BYTE[1..7] ISNONZERO
STARTM := MEM.WORD[4] // BYTES 8..9
STARTK := MEM.WORD[5] // BYTES 10..11
#GP IF MEM.BYTE[12..15] ISNONZERO

MAX_NAMES := IMPL.MAX_TILE_BYTES / PALETTE_TABLE[I].TILE_BYTES
MAX_NUM_ROWS := PALETTE_TABLE[PALETTE_ID].TILE_BYTES /
PALETTE_TABLE[PALETTE_ID].BYTES_PER_ROW
P := 16
FOR N IN 0 ...MAX_NAMES-1:
  T[N].ROWS := MEM.BYTE[P++]
  T[N].COLS := MEM.BYTE[P++]
  #GP IF T[N].ROWS == 0
  #GP IF T[N].COLS == 0
  #GP IF T[N].ROWS > MAX_NUM_ROWS
// ALL INSTRUCTIONS CHECK COLUMNS LIMITS BASED ON THEIR ELEMENT
// WIDTHS SO WE DO NOT ENFORCE COLUMN WITH RESTRICTION HERE.

WHILE P < 64: // CONFIRM REST OF BYTES IN 64B CHUNK ARE ZERO
  #GP IF MEM.BYTE[P] != 0
  P := P + 1
FOR N IN 0 ...MAX_NAMES-1:
  TILEZERO TILE[N]
  TILES_CONFIGURED := 1

```

FIG. 23

TILEDIAGONAL DESTINATION MATRIX (TILE) OPERAND IDENTIFIER, SOURCE IDENTIFIER
2402


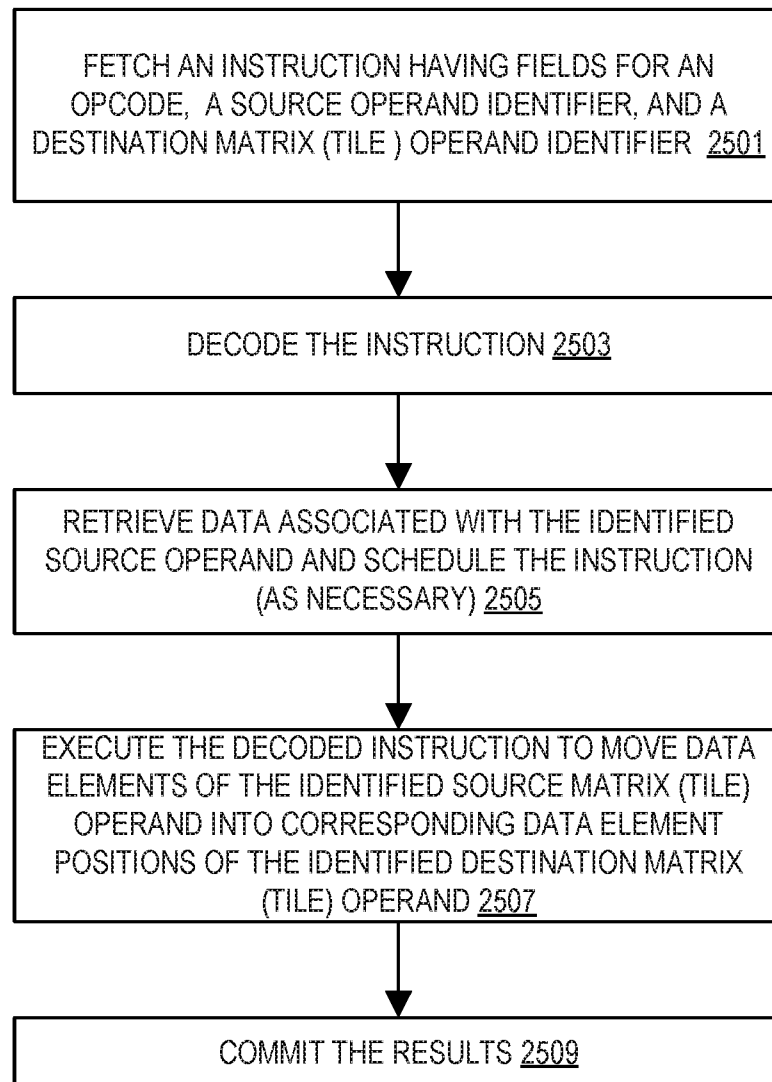
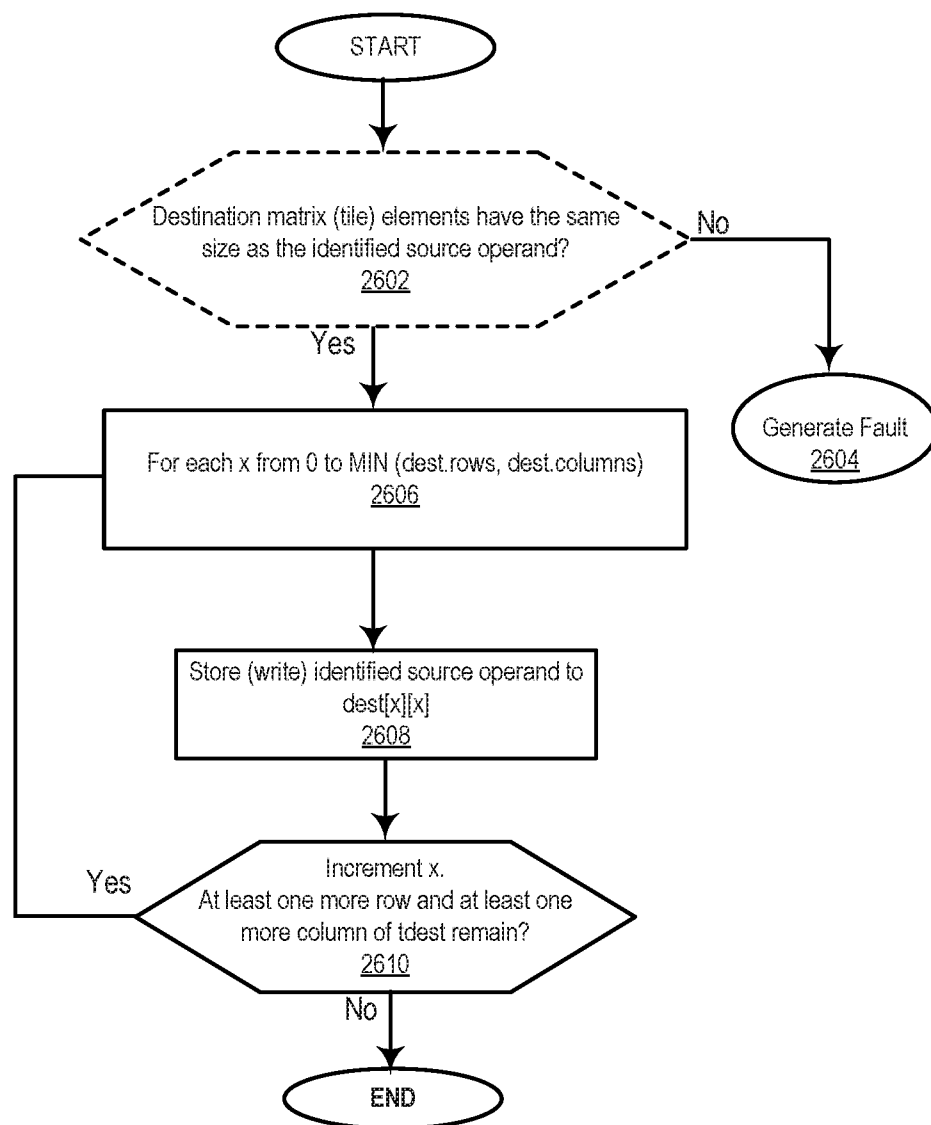


FIG. 24

24/36

**FIG. 25**

25/36

**FIG. 26**

2702

TILEDIAGONALD TDEST, SRC

#GP IF TILES_CONFIGURED == 0

#GP IF TDEST.NUM_BYTES_PER_ELEMENT != 4

#GP IF TDEST.NUM_BYTES_PER_ELEMENT != SRC.NUM_BYTES

LOOP_ITERATIONS = MIN(TDEST.NUM_ROWS, TDEST.NUM_COLUMNS)

FOR X IN 0...LOOP_ITERATIONS-1:

TDEST[X]X] ← SRC.

ZERO_UNUSED_ELEMENTS (TDEST)

2704

TILEDIAGONALW TDEST, SRC

#GP IF TILES_CONFIGURED == 0

#GP IF TDEST.NUM_BYTES_PER_ELEMENT != 2

#GP IF TDEST.NUM_BYTES_PER_ELEMENT != SRC.NUM_BYTES

LOOP_ITERATIONS = MIN(TDEST.NUM_ROWS, TDEST.NUM_COLUMNS)

FOR X IN 0...LOOP_ITERATIONS-1:

TDEST[X]X] ← SRC.

ZERO_UNUSED_ELEMENTS (TDEST)

FIG. 27






FIG. 28C





FIG. 29



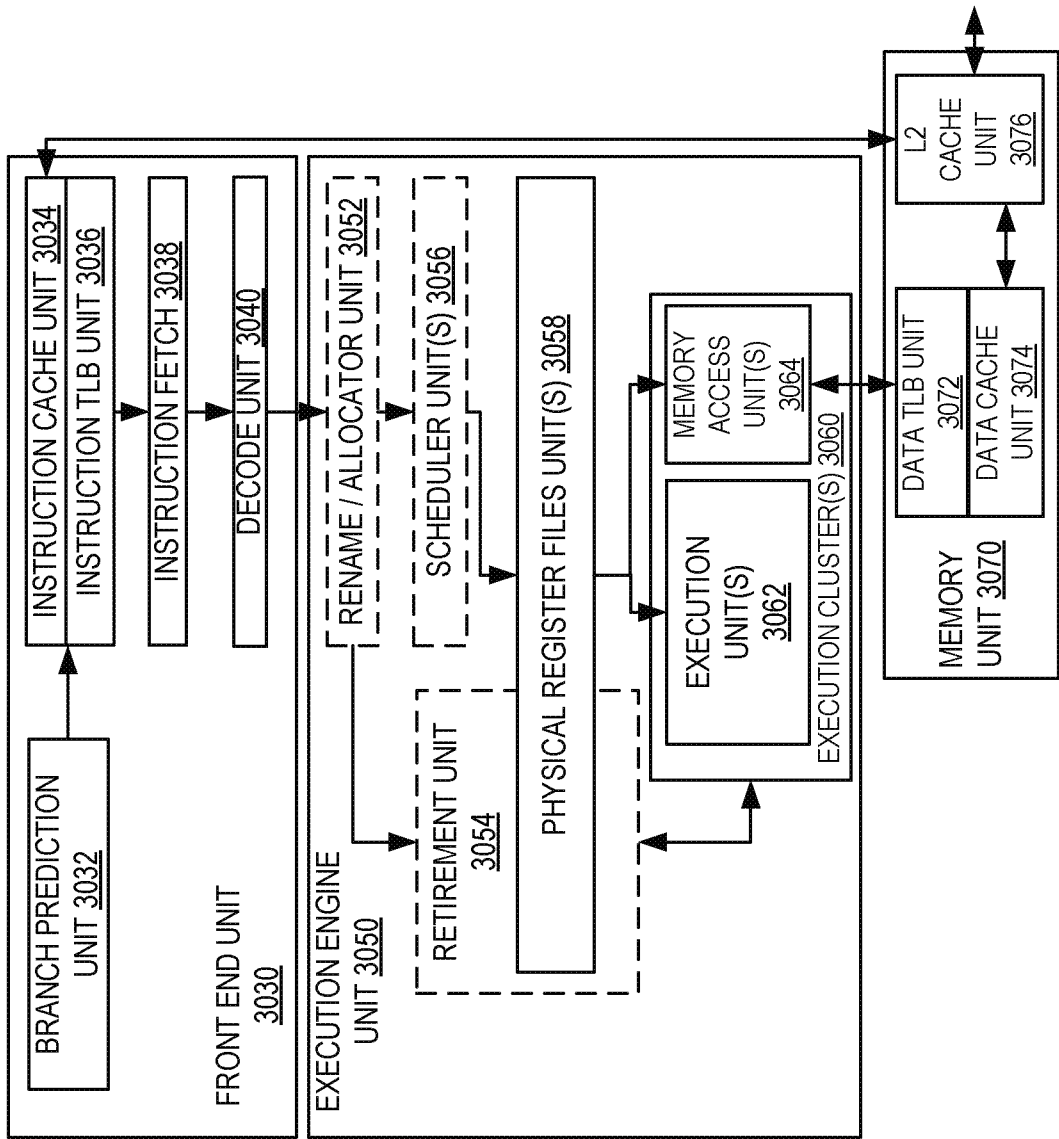



FIG. 31A





FIG. 31B






FIG. 32




FIG. 33




FIG. 34




FIG. 35




FIG. 36

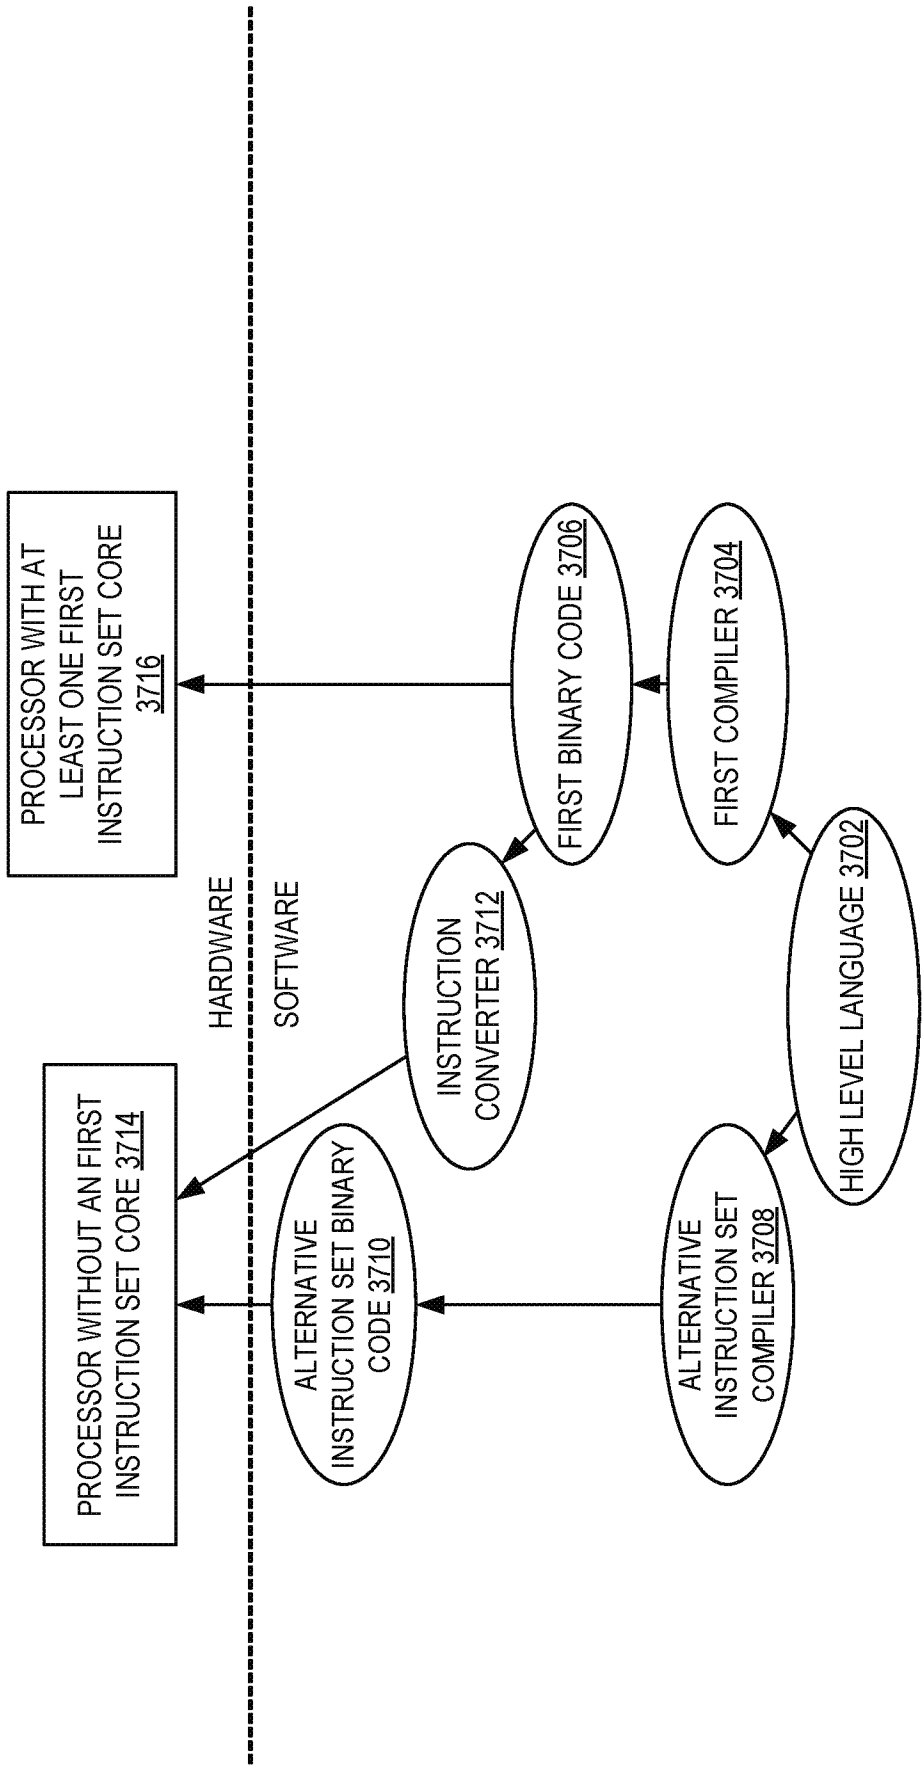


FIG. 37

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/040537**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/30(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 9/30; G01R 31/28; G11C 7/00; G06F 15/80; G06F 15/82; G06F 17/14; G06F 7/38Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: matrix, write, diagonal, source, destination, operand**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008-0208942 A1 (NARA WON et al.) 28 August 2008 See paragraphs [0017], [0040]-[0042]; claim 1; and figure 6.	1-25
Y	US 6212112 B1 (DAVID NAURA et al.) 03 April 2001 See claims 1, 8.	1-25
A	US 2006-0190517 A1 (MIGUEL A. GUERRERO) 24 August 2006 See paragraphs [0054]-[0056]; and figure 7.	1-25
A	US 6877020 B1 (JOSEPH P. BRATT et al.) 05 April 2005 See paragraphs [0374]-[0378]; and figure 75.	1-25
A	KR 10-2011-0079495 A (INTERNATIONAL BUSINESS MACHINES CORPORATION) 07 July 2011 See paragraphs [0022]-[0054]; and figure 2.	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 December 2017 (20.12.2017)

Date of mailing of the international search report

20 December 2017 (20.12.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

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



Facsimile No. +82-42-481-8578

Authorized officer

LEE, Chang Ho

Telephone No. +82-42-481-8288



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/040537

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008-0208942 A1	28/08/2008	US 7797362 B2 WO 2008-103885 A2 WO 2008-103885 A3	14/09/2010 28/08/2008 16/10/2008
US 6212112 B1	03/04/2001	FR 2787233 A1 FR 2787233 B1	16/06/2000 16/02/2001
US 2006-0190517 A1	24/08/2006	None	
US 6877020 B1	05/04/2005	None	
KR 10-2011-0079495 A	07/07/2011	CN 102053948 A CN 102053948 B JP 2011-100452 A JP 5689282 B2 US 2011-0107060 A1 US 8539201 B2	11/05/2011 11/09/2013 19/05/2011 25/03/2015 05/05/2011 17/09/2013