



- (51) International Patent Classification:
G06F 9/38 (2006.01)
- (21) International Application Number:
PCT/US2014/013455
- (22) International Filing Date:
28 January 2014 (28.01.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/753,098 29 January 2013 (29.01.2013) US
- (71) Applicant: **ADVANCED MICRO DEVICES, INC.**
[US/US]; One AMD Place, P.O. Box 3453, Sunnyvale,
California 94088 (US).
- (72) Inventor: **YAZDANI, Reza**; 911 Matts Ct., Los Altos,
California 94024 (US).
- (74) Agent: **RANKIN, Rory D.**; Meyertons, Hood, Kivlin,
Kowert & Goetzel, P.C., P.O. Box 398, Austin, Texas
78767-0398 (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, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.

- (84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments (Rule 48.2(h))

(54) Title: **HARDWARE AND SOFTWARE SOLUTIONS TO DIVERGENT BRANCHES IN A PARALLEL PIPELINE**





FIG. 8

(57) **Abstract:** A system and method for efficiently processing instructions in hardware parallel execution lanes within a processor. In response to a given divergent point within an identified loop, a compiler arranges instructions within the identified loop into very large instruction words (VLIW's). At least one VLIW includes instructions intermingled from different basic blocks between the given divergence point and a corresponding convergence point. The compiler generates code wherein when executed assigns at runtime instructions within a given VLIW to multiple parallel execution lanes within a target processor. The target processor includes a single instruction multiple data (SIMD) micro-architecture. The assignment for a given lane is based on branch direction found at runtime for the given lane at the given divergent point. The target processor includes a vector register for storing indications indicating which given instruction within a fetched VLIW for an associated lane to execute.



TITLE: HARDWARE AND SOFTWARE SOLUTIONS TO DIVERGENT BRANCHES IN A PARALLEL PIPELINE

BACKGROUND

5

Technical Field

[0001] This disclosure relates to computing systems, and more particularly, to efficiently processing instructions in hardware parallel execution lanes within a processor.

Background

10 [0002] The parallelization of tasks is used to increase the throughput of computer systems. To this end, compilers may extract parallelized tasks from program code to execute in parallel on the system hardware. To increase parallel execution on the hardware, a processor may include multiple parallel execution lanes, such as in a single instruction multiple word (SIMD) micro-architecture. This type
15 of micro-architecture may provide higher instruction throughput for particular software applications than a single-lane micro-architecture or a general-purpose micro-architecture. Some examples of tasks that benefit from a SIMD micro-architecture include video graphics rendering, cryptography, and garbage collection.

[0003] In many cases, particular software applications have data parallelism in which the
20 execution of each work item, or parallel function call, is data dependent within itself. For example, a first work item may be data independent from a second work item, and each of the first and the second work items are simultaneously scheduled on separate parallel execution lanes within a SIMD micro-architecture. However, an amount of instructions executed within each of the first and the second work items may be data-dependent. A conditional test implemented as a branch instruction
25 may pass for the first work item, but fail for the second work item dependent on the data for each work item.

[0004] The efficiency of parallel execution may be reduced as the second work item halts execution and waits as the first work item continues with its ongoing execution. The inefficiency grows when only a few work items continue execution due to passed tests whereas most of the work
30 items are idle due to failed tests.

SUMMARY OF EMBODIMENTS

[0005] Systems and methods for efficiently processing instructions in hardware parallel execution lanes within a processor are contemplated. In various embodiments, a back-end compiler examines program instructions of a software application to arrange instructions and generate code to efficiently process the instructions on a target processor. The target processor may include multiple parallel execution lanes within a single instruction multiple data (SIMD) micro-architecture. The compiler may identify loops and corresponding basic blocks. A divergent point within the loop may include a branch instruction. For example, an if-elseif-else construct, an if-else construct, a case construct and so forth may be used in the program instructions within the identified loop. Multiple trace paths may be traversed during execution of the translated and compiled program instructions between the divergent point and a corresponding convergent point.

[0006] During compilation, in response to a given divergent point within an identified loop, the compiler may arrange instructions within the identified loop into one or more very large instruction words (VLIW's). At least one VLIW may include instructions intermingled from different basic blocks between the given divergence point and a corresponding convergence point. For example, if there is a Basic Block A with four instructions and a Basic Block B with six instructions between the given divergent point and the corresponding convergent point, the compiler may arrange the instructions into six VLIW's. The first four VLIW's may include one instruction from each of the Basic Block A and the Basic Block B. The first VLIW may include the first instruction from each of the Basic Block A and the Basic Block B. The second VLIW may include the second instruction from each of the Basic Block A and the Basic Block B, and so forth. The last two VLIW's may include an instruction from the Basic Block B grouped with a nop. The compiler may keep track of a program counter (PC) value that points to each of the VLIW's.

[0007] The compiler may generate code to insert with the translated and compiled program instructions. The inserted code when executed may assign at runtime instructions within a given VLIW to the multiple parallel execution lanes within the target processor. The assignment for a given lane may be based on branch direction found at runtime for the given lane at the given divergent point. Continuing with the above example, if the VLIW is the second VLIW created from the Basic Block A and the Basic Block B and the branch instruction for the given lane is taken, then

the given lane may be assigned the second instruction in Basic Block A within the second VLIW. If the branch instruction is not taken, then the given lane may be assigned the second instruction in Basic Block B within the second VLIW. In various embodiments, the VLIW's have a variable length. The inserted code when executed may update the size of a VLIW corresponding to a next
5 PC.

[0008] In some embodiments, a processor includes multiple parallel execution lanes within a single instruction multiple data (SIMD) micro-architecture. The processor may include a size register for storing a size of a variable-length VLIW. Control logic within the processor may fetch and decode in a respective cycle a number of instructions within a given VLIW equal to the stored
10 size. The multiple execution lanes may simultaneously execute the number of instructions within the given VLIW. The processor may additionally include a vector register with bit ranges associated with corresponding lanes of the multiple execution lanes. The bit ranges may store offsets. A given offset may identify a given instruction within a fetched VLIW for an associated lane to execute.

[0009] These and other embodiments will be further appreciated upon reference to the following
15 description and drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a generalized block diagram of one embodiment of single instruction multiple data (SIMD) pipeline execution flow.

20 [0011] FIG. 2 is a generalized block diagram of one embodiment of a control flow graph.

[0012] FIG. 3 is a generalized block diagram of one embodiment of execution sequences for a control flow graph.

[0013] FIG. 4 is a generalized block diagram of one embodiment of a logical layout of a SIMD micro-architecture for a processor.

25 [0014] FIG. 5 is a generalized block diagram of one embodiment of an object code arrangement.

[0015] FIG. 6 is a generalized flow diagram of one embodiment of a method for optimizing parallel execution of multiple work-items in a processor with compiler techniques.

[0016] FIG. 7 is a generalized flow diagram of one embodiment of a method for optimizing parallel execution of multiple work-items in a processor using hardware techniques.

30 [0017] FIG. 8 is a generalized block diagram of another embodiment of an object code arrangement.

[0018] While the embodiments are susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are herein described in detail. It should be understood, however, that drawings and detailed description thereto are not intended to limit the embodiments to the particular form disclosed, but on the contrary, the
5 embodiments are to cover all modifications, equivalents and alternatives falling within the spirit and scope of the embodiments as defined by the appended claims.

DETAILED DESCRIPTION OF EMBODIMENT(S)

[0019] In the following description, numerous specific details are set forth to provide a thorough
10 understanding of the embodiments. However, one having ordinary skill in the art should recognize that the embodiments might be practiced without these specific details. In some instances, well-known circuits, structures, and techniques have not been shown in detail to avoid obscuring the embodiments.

[0020] Referring to FIG. 1, a generalized block diagram illustrating one embodiment of a single
15 instruction multiple data (SIMD) pipeline execution flow 100 is shown. Instructions 102-108 may be fetched and sent to a SIMD pipeline with associated data. Multiple computation units within parallel vertical execution lanes are shown. Some of the computation units are active computation units 110. Other computation units are inactive computation units 112 due to being disabled during a given pipe stage. Control logic and storage elements, such as pipeline registers, are not shown for
20 ease of illustration.

[0021] The hardware computation units include hardware that performs the execution of a given instruction of a given work-item with associated data. This hardware may include an arithmetic logic unit that is configured to perform addition, multiplication, zero detect, a bit-wise shift, division, video graphics and multimedia instructions or other operations known to those skilled in
25 the art of processor design. Examples of processors with parallel execution lanes within a SIMD pipeline include graphics processing units (GPUs), digital signal processing (DSP), and so forth. In one embodiment, the SIMD pipeline may be located on a video card. In another embodiment, the SIMD pipeline may be integrated on the motherboard.

[0022] The SIMD pipeline may improve computing performance for a wide variety of data-
30 parallel applications used in gaming, entertainment, science and medical fields. Such applications

generally entail executing the same programs on large numbers of objects. Since each object is processed independently of other objects, but the same sequence of operations is used, a SIMD micro-architecture provides considerable performance enhancement. GPUs have also been considered for non-graphical calculations.

5 [0023] A software application may include a collection of function calls, or compute kernels, and internal functions. A software programmer may define the function calls, whereas the internal functions may be defined in a given library. For example, a software application may perform data processing on a two-dimensional (2D) array of data, such as an image file. The software application may perform an algorithm developed by a software programmer on a pixel-by-pixel basis of a 2D
10 image or an element-by-element basis of a two-dimensional matrix. A given function call may be invoked over an index space. An index space may also be referred to as a dimensional space. For a data-parallel software application, an N-Dimensional computation domain may define a one-, two-, or three-dimensional space, or index space. An example is the pixels within the 2D image.

[0024] A function call may be matched with one or more records of data to produce one or more
15 work-items of computation. Therefore, two or more work-items may utilize the same instructions of the single function call, but operate on different records of data. A function call may include control flow transfer instructions that create forks, whereas a fork in a computer program typically creates a software thread, by common definition. A given instance of the function call at a given point in the index space may be referred to as a “work-item”. A work-item may also be referred to as a work
20 unit. Continuing with the above example, a work-item may operate with the one or more instructions in the function call on a record of data corresponding to a given pixel (a given index) of the 2D image. Typically, work-items have an associated unique identifier (ID).

[0025] The index space may define a total number of work-items that execute in parallel if there is sufficient hardware support. For example, the index space may define a number of 280 work-items,
25 but a GPU may support the simultaneous execution of 64 work-items at any given time. The total number of work-items may define a global work size. The work-items may be further grouped into work groups. Each work group may have a unique identifier (ID). The work-items within a given work group may be able to communicate with each other and synchronize execution and coordinate memory accesses. A number of work-items may be clustered into a wave front for simultaneous

execution on a GPU in a SIMD manner. Regarding the example above for 280 total work-items, a wave front may include 64 work-items.

[0026] Instructions 102-108 may be fetched and enter a SIMD pipeline with associated data. Instruction 104 may be a control flow transfer instruction, such as a conditional branch. The instruction 106 may be a first instruction in a path executed when the condition is true. The instruction 108 may be a first instruction in a path executed when the condition is false. For example, the branch instruction 104 may be associated with an IF statement in a high-level language program. The instruction 106 may be associated with a THEN statement in the high-level language program. The instruction 108 may be associated with an ELSE statement in the high-level language program.

[0027] Each of the computation units within a given row may be a same computation unit. Each of these computation units may operate on a same instruction, but different data associated with a different work-item. As shown, some of the work-items pass the test provided by the conditional branch instruction 104 and other work-items fail the test. Control logic within the SIMD pipeline may execute each of the available paths and selectively disable the execution units, such as the computation units, corresponding to work-items that did not choose the current path. For example, during execution of an If-Then-Else construct statement, within each column of a SIMD architecture are execution units configured to execute the “Then” (Path A) and the “Else” (Path B) paths.

[0028] The efficiency of parallel execution may be reduced as the first and the second work-items pause execution and wait as the third work-item continues with its ongoing execution. Therefore, not all of the computation units are active computation units 110 in a given row after execution of the branch instruction 104. As shown, one or more computation units are inactive computation units 112 that have been disabled for execution. If a large number of computation units are inactive during a given pipe stage, the efficiency and throughput of the SIMD core is reduced. In one embodiment, an “Else” path is a return for the function call. Execution of the function call ends and the corresponding work-item becomes idle. However, neighboring work-items in the SIMD core may continue executing.

[0029] Turning now to FIG. 2, a generalized block diagram illustrating one embodiment of a control flow graph 200 is shown. Generally speaking, control flow graphs may be used by compiler optimizers and static analysis tools. The control flow graph 200 may represent all paths that might

be traversed through the program or a portion of the program during its execution. In a control flow graph, each node in the graph represents a basic block. Most representations include an entry block, through which control enters into the control flow graph, and an exit block, through which control leaves the control flow graph.

5 [0030] During compilation, a software application may provide the control flow graph 200 with eight basic blocks (BB's) numbered from basic block 0 (BB 0) to basic block 7 (BB 7). Although eight basic blocks are shown, in other examples, another number of basic blocks may be used. For the control flow graph 200, basic block 1 is the entry block and basic block 6 is the exit block. Each one of the basic blocks 0-7 is a straight-line sequence of instructions with one entry point and one
10 exit point. The control flow graph 200 may represent a loop. Inside the loop, the control flow graph 200 may represent an IF-THEN-ELSE construct with basic blocks 1-4 and an IF construct with basic blocks 4-6.

[0031] Referring now to FIG. 3, a generalized block diagram illustrating one embodiment of execution sequences 300 for a control flow graph is shown. The execution sequence 310 represents
15 a typical execution sequence for the previously shown control flow graph 200 when assigned to a SIMD pipeline. The execution time for a single loop iteration is the sum of the execution times of each of the basic blocks within the loop, such as BB 1 to BB 6. However, for a given work-item and corresponding hardware execution lane, only one of BB 2 and BB 3 is executed. Similarly, BB 5 may be skipped for the given work-item. Although particular basic blocks may not be executed for
20 the given work-item, the associated execution times still contribute to the total sum of execution time for the loop iteration.

[0032] The execution sequence 320 represents an alternative execution sequence for the previously shown control flow graph 200 when assigned to a modified SIMD pipeline. The execution time for a single loop iteration is the sum of the execution times of the basic blocks
25 actually executed in the single loop iteration, rather than the sum of the execution times of each of the basic blocks within the loop. The execution sequence 320 transforms the execution of the control flow graph 200. The compiler may perform this transformation during compilation of the source code to object code.

[0033] In some embodiments, the compiler generates code at the end of each basic block wherein
30 when executed the code identifies a next basic block to execute. The generate code may be inserted

at the end of a basic block. Alternatively, a branch instruction may be inserted at the end of a basic block to transfer control flow to the additional code before transferring control to the next basic block. The additional code is represented by BB C, which refers to a central basic block. When executed, each branching basic block, such as BB 1, BB 4 and BB 6, identifies a next basic block to have control transferred to it. The identification is based on the branch resolution, which is further based on a particular record of data and the branch instruction. When executed, the BBC receives the address of the target basic block from each work-item and sets up thread registers to execute the respective target basic block. Across the different work-items being executed in the SIMD pipeline, a given work-item has a single target for a control flow transfer instruction, such as branch, jump, and case statements.

[0034] In the execution sequence 320, at the completion of BB 4, a first work-item may branch to BB 5 and pass the corresponding address to BB C. At the completion of BB 4, a second work-item may branch to BB 6 and pass the corresponding address to BB C. The compiler may create a dynamic very large instruction word (DVLIW) that includes an instruction from each of the target basic blocks to be simultaneously executed. When executed, the generated code by the compiler may update a size of a next DVLIW to execute corresponding to a next program counter (PC) value. Additionally, when executed, the generated code may update a mapping between a given work-item being executed in a given parallel execution lane and a pointer to an instruction within the next DVLIW to execute. The pointer may be an offset identifying an associated one of the instructions within a next DVLIW, which is to be fetched.

[0035] Referring to FIG. 4, a generalized block diagram illustrating one embodiment of a logical layout 400 of a SIMD micro-architecture for a processor is shown. The processor has a dynamic random access memory (DRAM) 450 for storage of data and instructions. In some embodiments, a given level of a cache memory subsystem is used in addition to the DRAM. As shown, the processor may have a relatively small cache memory subsystem 430 grouped with control logic 420 for each row of computation units. The data flow within the processor may be pipelined, although storage elements, such as pipeline registers, are not shown in order to simplify the illustration. In a given pipeline stage, a computation unit may be unused if an associated instruction in this stage is not executed based on a previous failed test, such as a not-taken branch.

[0036] The SIMD pipeline includes work-items 460 comprising lanes A-F. Each one of the lanes A-F may correspond to a respective one of the vertical parallel hardware lanes comprising computation units. The pipeline may additionally include a vector register 462. The vector register 462 may include an entry, a field, or a bit range for each one of the parallel execution lanes. Each entry may include a total number of bits comprising a first number of bits to identify a given trace executing on a respective work-item and a second number of bits to support special codes. The special codes may identify wait or sleep states, a loop exit state, a barrier identifier for stalling execution other than for ending a loop, an event identifier, and so forth. When a special code is not stored in a given entry, the value stored may identify a respective one of the instructions within the DVLIW for an associated lane to execute.

[0037] The program counter (PC) register 466 may store a pointer value or address pointing to a next DVLIW to fetch from memory, such as an i-cache. The processor may further include a size register 468 that stores the size, or the length, of the DVLIW. In some embodiments, the size may be an integer representing a number of instructions within the variable length DVLIW.

[0038] Each one of the instructions Instr A to Instr G within the DVLIW 464 represents an executing trace in the control flow graph. The compiler may lay out DVLIW's in memory, such as an i-cache. In one example, lane B within the work-items 460 may correspond to the second vertical execution lane from the left in the SIMD pipeline. The offset B stored in the vector register 462 may be associated with lane B and point to the first instruction within the DVLIW 464, which is Instr A. Therefore, lane B may receive Instr A to process. Similarly, lane A within the work-items 460 may correspond to the left-most vertical execution lane in the SIMD pipeline. The offset A stored in the vector register 462 may be associated with lane A and point to the last instruction within the DVLIW 464, which is Instr G. Therefore, lane A may receive Instr G to process.

[0039] Although not shown, the instruction cache (i-cache) may include one of multiple implementations to support the DVLIW's. The i-cache may include multiple smaller caches for fetching one or more instructions for a given single PC corresponding to a DVLIW. The same PC may index a valid instruction in one or more of the smaller caches depending on the size of the DVLIW. The I-cache may receive the size stored in the size register 468 in addition to the pointer or address value stored in the PC register 466. Alternatively, the I-cache may have multiple data ports

for accessing one or more instructions in a same valid cache line or cache set. Again, the number of valid instructions to fetch may equal the received size from the size register 468.

[0040] Referring now to FIG. 5, a generalized block diagram of one embodiment illustrating an object code arrangement 500 is shown. The code arrangement 500 illustrates one embodiment of object code layout that may be generated by the compiler for the control flow graph 200 and the accompanying execution sequence 320 shown in FIG. 2 and FIG. 3, respectively. Basic block code 504 represents code for respective basic blocks. For example, one copy of code for basic blocks 0-3 and 7 are shown laid out in layout 500. Two copies of code for basic blocks 4-5 are shown laid out in layout 500. Four copies of code for basic block 6 are shown in layout 500.

[0041] The code 502 may be generated and inserted by the compiler for entries of a region, such as a loop. When later executed, the code 502 may update the size of a next DVLIW and update mappings between the instructions within a fetched DVLIW and the parallel execution lanes in a target processor. For example, the size and vector registers within the target processor may have their stored contents updated by the executed code 502. As shown, the code 502 may be inserted at the beginning of basic blocks 0, 1 and 7.

[0042] The code 506 may be generated and inserted by the compiler for the transitions into divergent points, such as the end of basic blocks 1, 4 and 6. When later executed, the code 506 may determine the DVLIW size changes and corresponding mapping changes between the instructions within a fetched DVLIW and the parallel execution lanes in a target processor. Therefore, the size and mappings are updated at divergent and convergent points within the control flow graph. The compiler identifies the points where the size of the DVLIW's change and where the mappings change. Using the notation BBC (0, 1), wherein the first index indicates a trace identifier (ID) and the second index indicates a basic block (BB) ID, the identified points for inserting code 506 may exist at BBC (0, 1), BBC (0, 4), BBC (0,6), BBC (1, 4), BBC (1, 6), BBC (2, 6) and BBC (3, 6).

The trace ID in this example may be the same as a corresponding offset.

[0043] At the beginning of the object code arrangement 500, the code 502 performs initialization steps and may set the DVLIW size to 1. The offsets within each of the entries of the vector register may be set to point to a same instruction within the BB 0, such as an offset of 0. Therefore, the PC may be set to 0 or another appropriate starting address. As shown, there are four possible parallel traces, but the number of work-items may be independent. For example, the SIMD pipeline may

have 8, 16, 64 or another number of parallel execution lanes for processing assigned work-items. Each of the work-items within the SIMD pipeline has a stored offset of 0 in the vector register and will execute a same instruction. This same instruction for each work-item is an instruction from BB 0. The instructions in BB 0 are executed one-by-one by each work-item and the PC is incremented
5 after each instruction fetch.

[0044] After execution of BB 0 is complete, the loop entry block, BB 1, is next to be processed. The code 502 at the start of BB 1 maintains the DVLIW size as 1 and the offset for each work-item as 0. The offsets within each of the entries of the vector register may be set to point to a same instruction within the BB 1, such as an offset of 0. The PC may remain an incremented value at the
10 completion of BB 0. Each of the work-items within the SIMD pipeline has a stored offset of 0 in the vector register and will execute a same instruction. This same instruction for each work-item is an instruction from BB 1. The instructions in BB 1 are executed one-by-one by each work-item and the PC is incremented after each instruction fetch.

[0045] When executed, the code 506 at BBC (0, 1) which is at the end of BB 1, changes the stored
15 DVLIW size from 1 to 2. The instructions within BB 3 are now added to the fetched DVLIW's. Additionally, when executed, the code 506 at BBC (0, 1) sets the entries in the vector register for work-items branching to BB 3 to store the value 1. The entries in the vector register for work-items branching to BB 2 remain unchanged by continuing to store 0. Although the values of 0 and 1 are used in this manner, other numerical values may be used to set the corresponding indications and
20 mappings. At this point, the DVLIW has two instructions, which are intermingled from the two separate basic blocks BB 2 and BB 3. As the PC continues to be incremented, the fetched DVLIW's continue to include intermingled instructions from these two basic blocks until the processing of BB 2 completes. The compiler may have laid out the instructions as intermingled in this manner in memory in order to support the parallel execution of the instructions within the DVLIW's.

[0046] At the completion of BB 2, the DVLIW size remains 2. The stored offsets in the vector
25 register also maintain their values. However, now the offset 0 corresponds to instructions in BB 4, rather than the BB 2. At the completion of BBC (0, 4), the code 506 at the end of BB 4 in trace 0 updates the size register to store 3 and updates the entries for work-items branching to BB 6 to store 2. A DVLIW of length 3 is fetched from the i-cache after sending the single PC and the stored size
30 to the i-cache. The DVLIW's include intermingled instructions from BB 3 or BB 4, BB 5 and BB 6.

Work-items with an associated stored offset of 0 in the vector register obtain fetched instructions from BB (0, 5). Work-items with an associated stored offset of 1 in the vector register obtain results from either BB (1, 3) or BB (1, 4) depending on how far the PC has incremented. Work-items with an associated stored offset of 2 in the vector register obtain fetched instructions from BB (2, 6). The compiler had previously laid out the instructions in memory in this manner so that a single PC and the stored size indicate to the i-cache the type of DVLIW to fetch.

[0047] At the completion of BBC (1, 4), the code 506 at the end of BB 4 in trace 1 updates the size register to store 4 and updates the entries for work-items branching to BB (3, 6) to store 3. A DVLIW of length 4 is fetched from the i-cache after sending the single PC and the stored size to the i-cache. The DVLIW's include intermingled instructions from a first copy of BB 6, a single copy of BB 5, a second copy of BB 6 and a third copy of BB 6. Work-items with an associated stored offset of 0 in the vector register obtain fetched instructions from BB (0, 6). Work-items with an associated stored offset of 1 in the vector register obtain results from either BB (1, 5). Work-items with an associated stored offset of 2 in the vector register obtain fetched instructions from BB (2, 6). Work-items with an associated stored offset of 3 in the vector register obtain fetched instructions from BB (3, 6). The compiler had previously laid out the instructions in memory in this manner so that a single PC and the stored size indicate to the i-cache the type of DVLIW to fetch.

[0048] For each of BB (0, 6), BB (1, 6), BB (2, 6) and BB (3, 6), control flow at the end of BB 6 may either return to BB 1 for another iteration of the loop or terminate from the loop. An associated branch instruction and data within a corresponding record will determine at runtime the control flow direction. Some work-items may continue with another iteration and other work-items may exit the loop. A special code status may be stored in a corresponding entry in the vector register to indicate which path is taken. In response to determining a given work-item will continue another loop iteration and at least one other work-item of the multiple work-items is still processing code for an associated basic block, a sleep state encoding may be stored in the associated entry in the vector register for the given work-item.

[0049] In response to determining the given work-item will exit the loop, an exit state encoding may be stored in the associated entry in the vector register for the given work-item. Each of the sleep and exit state encodings is unique from the offsets used during loop iterations and unique from one another. In some embodiments, for the given work-item that is in the sleep or the exit state, the

code 506 halts execution of the given work-item and stores at least a next program counter (PC) and the work-item identifier (ID) in memory, such as a stack memory for fast retrieval.

[0050] The code 506 at the end of BB 6 may check the status of each of the work-items. In response to determining each work-item is in the sleep state or each work-item is in the exit state, the processor while executing code 506 may restart execution for each work-item by branching to a respective stored next PC. In response to determining each work-item is halted and at least one work-item is in a different state than another work-item, the processor while executing code 506 may restart execution for only work-items in a sleep state by branching to a respective stored next PC. If at least one work-item is still processing instructions within a basic block in the loop, then execution continues for the at least one work-item while the other work-items in a special state wait.

Work-items leaving the sleep state branch back to BB 1. Additionally, the code 502 at the beginning of BB 1 reinitializes the vector and size registers. Work-items leaving the exit state branch to BB 7. Additionally, the code 502 at the beginning of BB 7 accordingly reinitializes the vector and size registers.

[0051] In the above example, the loop has a single exit. In other cases with multiple exits, corresponding state information, such as at least a next PC and a work-item ID, may be stored in memory, such as a stack. Later, the state information may be retrieved, such as popping them from the stack, for restarts. Both work-items that are in a sleep state or an exit state may have state information stored in memory, such as a stack. As different work-items may exit the loop at different loop iterations, multiple entries with state information may be placed in memory, such as the stack. At restart time, when executed, the compiler-generated code may pop the state information and combine information for work-items that restart from a same next PC.

[0052] Turning now to FIG. 6, one embodiment of a method 600 for optimizing parallel execution of multiple work-items in a processor with compiler techniques is shown. For purposes of discussion, the steps in this embodiment and subsequent embodiments of methods described later are shown in sequential order. However, in other embodiments some steps may occur in a different order than shown, some steps may be performed concurrently, some steps may be combined with other steps, and some steps may be absent.

[0053] In block 602, a software program or subroutine may be located and analyzed. The program code may be written by a designer in a high-level language such as C or another language. This

software program may be written for compilation and execution of a parallel-data application, such as in the gaming, business, medical and other fields. Program code may refer to any portion of a software application, subroutine, dynamic linked library, or otherwise. A pathname may be entered at a command prompt by a user. Alternatively, a pathname may be read from a given directory location, or other, in order to begin compiling the source code. Instructions within the program code may be examined, translated, optimized, and further processed during compilation.

[0054] In some embodiments, the source code is statically compiled. In such embodiments, during a front-end compilation, the source code may be translated to an intermediate representation (IR). A back-end compilation step may translate the IR to machine code. The static back-end compilation may perform more transformations and optimizations. In other embodiments, the source code is compiled with a Just-In-Time (JIT) method. The JIT method may generate an appropriate binary code after obtaining the system configuration. With either method, the compiler may identify function calls, loops, traces within the loops, and basic blocks in the program code. One or more control flow graphs may be constructed during program analysis.

[0055] In various embodiments, the program code is compiled on a processor, such as a general-purpose processor. The program code may be compiled for a target processor, which includes a parallel micro-architecture, such as a SIMD micro-architecture. One or more associated record(s) of data may be assigned to a function call to produce one or more work-items.

[0056] Prior to detecting any divergent points in the program code, the compiler may lay out the analyzed and translated instructions in memory as they appear in the program code. Essentially, the compiler may be creating VLIW's with a size of one. If the compiler detects a divergent point within an identified loop (conditional block 604), then in block 606, the compiler may create very large instruction words (VLIW's). The compiler may create VLIW's by laying out in memory instructions intermingled from multiple basic blocks between the divergent point and a corresponding convergent point. One or more associated record(s) of data may be arranged and assigned to accompany the intermingled instructions within the created VLIW's to produce one or more associated work-items.

[0057] In block 608, generated code may be inserted wherein when executed the generated code maps an offset pointing to an instruction within a VLIW to a given lane of multiple parallel execution lanes within a target processor. Alternatively, the mapping may be between the offset and

a work-item ID. In block 610, generated code may be inserted wherein when executed updates a size of a next VLIW to fetch. In block 612, generated code may be inserted at the end of the loop wherein when executed stores state information for a lane of execution going to a sleep or an exit state. The generated code may be inserted in particular points in the program code as described in the earlier example. Basic blocks associated with divergent and convergent points may have more inserted code for maintaining updates to the mappings and DVLIW size.

[0058] Referring now to FIG. 7, one embodiment of a method 700 for optimizing parallel execution of multiple work-items in a processor using hardware techniques is shown. For purposes of discussion, the steps in this embodiment and subsequent embodiments of methods described later are shown in sequential order. However, in other embodiments some steps may occur in a different order than shown, some steps may be performed concurrently, some steps may be combined with other steps, and some steps may be absent.

[0059] In block 702, associated record(s) of data are assigned to compiled code to produce multiple work-items. In block 704, the work-items are scheduled to a target processor with a single instruction multiple data (SIMD) micro-architecture. In block 706, using an updated VLIW size and a single program counter (PC), a VLIW is fetched from memory, such as an i-cache, with a length equal to the updated VLIW size. The instructions in the VLIW may be from separate basic blocks between a divergent point and a convergent point within a loop.

[0060] In block 708, a vector register is read for mapping information between instructions within the fetched VLIW and parallel execution lanes in the processor executing the work-items. The mapping information may determine for a given work-item and corresponding execution lane, which instruction in the VLIW to process. In block 710, the instructions in the VLIW are simultaneously executed using the parallel execution lanes. If the end of a loop is reached for a given work-item (conditional block 712) and if no work-items are detected to be active (conditional block 714), then in block 716, respective state information is read for each of the work-items assigned to the parallel execution lanes. The state information may include at least a next PC and a work-item ID. The state information may be used for continuing execution within the parallel execution lanes. If the end of a loop is reached for a given work-item (conditional block 712) and if any work-item is detected to be active (conditional block 714), then in block 718, the state information for the given work-item is

stored for later use. The given work-item may have execution halted and may be placed in a sleep or an exit state.

[0061] Referring now to FIG. 8, a generalized block diagram of another embodiment illustrating an object code arrangement 800 is shown. The code arrangement 800 illustrates one embodiment of object code layout that may be generated by the compiler for the control flow graph 200 and the accompanying execution sequence 320 shown in FIG. 2 and FIG. 3, respectively. The code 502-506 may perform the same functionality as previously described.

[0062] The sizes of the basic blocks are known at compile time. The compiler may arrange and schedule the instructions of the program code to minimize initialization code and reduce the DVLIW size. In the earlier example using the control flow graph 200 and the object code arrangement 500, four copies of BB 6 are used. The compiler may reduce the number of traces, and accompanying offsets, by introducing gaps. The gaps may utilize nop operations.

[0063] The steps described earlier regarding code arrangement 500 may be used here. At the completion of BB 2, the DVLIW size remains 2. However, one of the two instructions within the fetched DVLIW's is a nop operation providing a gap in the arrangement 800. The offset 0 corresponds to the nop operations and the offset 1 corresponds to the instructions in BB 3. At the completion of BB 3, the DVLIW is reduced from 2 to 1. Now the offset 0 corresponds to the instructions in BB 4 and the offset 1 corresponds to the nop operations.

[0064] At the completion of BBC (0, 4), the code 506 at the end of BB 4 in trace 0 updates the size register from 1 to 2. The code 506 also updates the entries for work-items branching to BB 6 to store 1. A DVLIW of length 2 is fetched from the i-cache after sending the single PC and the stored size to the i-cache. The DVLIW's include intermingled instructions from BB 5 and BB 6. Work-items with an associated stored offset of 0 in the vector register obtain fetched instructions from BB (0, 5). Work-items with an associated stored offset of 1 in the vector register obtain fetched instructions from BB (1, 6).

[0065] At the completion of BB (1, 6), the code 506 updates the DVLIW size from 2 to 1. The work-items corresponding to BB (1,6) halt execution, store state information, and wait until restarting execution at a next PC. The next PC may point to BB 1 or BB 7. Other steps may be performed as described earlier. The code arrangement 800 illustrates in response to determining a first trace path between the detected given divergent point and the corresponding convergent point is

smaller than a second trace path between the given divergent point and the corresponding convergent point, the compiler may group nops with instructions corresponding to the second trace path in VLIW's created between the completion of the first trace path and the corresponding convergent point.

5 [0066] For each of the embodiments described earlier, further parallelization may occur with code being parallelized outside the loop with the code inside the loop. For example, the program code for BB 7 may be parallelized with BB 1 to terminate the program for work-items that complete the loop. Additionally, in response to detecting an offset corresponds to multiple resource-independent instructions within the VLIW, rather than a single instruction, an associated work-item and
10 execution lane may simultaneously execute the multiple resource-independent instructions within the execution lane.

[0067] Further, the compiler may decrease the DVLIW size using register allocation. Program code may include the following statement: $X = (A+B) + (B+M)$. Here, two operations use the same opcode. A first trace includes an ADD operation such as $T1 = A+B$. A second trace includes
15 an ADD operation such as $T2 = C+D$. There are work-items that use T1 from a basic block X (BB X) with an offset such as 0. There are other work-items that use T2 from BB Y with an offset such as 1. If a first operands pair "C" and "A", a second operands pair "B" and "D", and a results pair "T1" and "T2" are assigned to the same registers in BB X and BB Y, respectively, then the expression $r3 = r1 + r2$ may be used as a DVLIW with a size of 1. The corresponding offset may be
20 set to 0 to save on decoding time or free a slot.

[0068] It is noted that the above-described embodiments may comprise software. In such an embodiment, the program instructions that implement the methods and/or mechanisms may be conveyed or stored on a computer readable medium. Numerous types of media which are configured to store program instructions are available and include hard disks, floppy disks, CD-ROM, DVD,
25 flash memory, Programmable ROMs (PROM), random access memory (RAM), and various other forms of volatile or non-volatile storage. Generally speaking, a computer accessible storage medium may include any storage media accessible by a computer during use to provide instructions and/or data to the computer. For example, a computer accessible storage medium may include storage media such as magnetic or optical media, e.g., disk (fixed or removable), tape, CD-ROM, or DVD-
30 ROM, CD-R, CD-RW, DVD-R, DVD-RW, or Blu-Ray. Storage media may further include volatile

or non-volatile memory media such as RAM (e.g. synchronous dynamic RAM (SDRAM), double data rate (DDR, DDR2, DDR3, etc.) SDRAM, low-power DDR (LPDDR2, etc.) SDRAM, Rambus DRAM (RDRAM), static RAM (SRAM), etc.), ROM, Flash memory, non-volatile memory (e.g. Flash memory) accessible via a peripheral interface such as the Universal Serial Bus (USB) interface, etc. Storage media may include microelectromechanical systems (MEMS), as well as storage media accessible via a communication medium such as a network and/or a wireless link.

[0069] Additionally, program instructions may comprise behavioral-level description or register-transfer level (RTL) descriptions of the hardware functionality in a high level programming language such as C, or a design language (HDL) such as Verilog, VHDL, or database format such as GDS II stream format (GDSII). In some cases the description may be read by a synthesis tool, which may synthesize the description to produce a netlist comprising a list of gates from a synthesis library. The netlist comprises a set of gates, which also represent the functionality of the hardware comprising the system. The netlist may then be placed and routed to produce a data set describing geometric shapes to be applied to masks. The masks may then be used in various semiconductor fabrication steps to produce a semiconductor circuit or circuits corresponding to the system. Alternatively, the instructions on the computer accessible storage medium may be the netlist (with or without the synthesis library) or the data set, as desired. Additionally, the instructions may be utilized for purposes of emulation by a hardware based type emulator from such vendors as Cadence®, EVE®, and Mentor Graphics®.

[0070] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

WHAT IS CLAIMED IS

1. A non-transitory computer readable storage medium storing at least one program configured for execution by at least one processor of a computer system, wherein the at least one program comprising instructions executable to:

5 identify loops and corresponding basic blocks within a plurality of program instructions; and in response to a given divergence point within an identified loop, arrange a plurality of instructions within the identified loop into a plurality of very large instruction words (VLIW's), wherein at least one VLIW comprises instructions intermingled from different basic blocks between the given divergence point and a corresponding
10 convergence point.

2. The non-transitory computer readable storage medium as recited in claim 1, further comprising instructions executable to in response to the given divergence point within the identified loop, assign at runtime instructions within a given VLIW to a plurality of parallel execution lanes within a target
15 processor comprising a single instruction multiple data (SIMD) micro-architecture.

3. The non-transitory computer readable storage medium as recited in claim 2, wherein assigning for a given lane is based on branch direction found at runtime for the given lane at the given divergent point.
20

4. The non-transitory computer readable storage medium as recited in claim 2, further comprising instructions executable to update a stored size of a VLIW corresponding to a next program counter (PC).

25 5. The non-transitory computer readable storage medium as recited in claim 4, further comprising instructions executable to in response to a first trace path being smaller than a second trace path between the detected given divergent point and the corresponding convergent point, group nops with instructions corresponding to the second trace path in VLIW's.

6. The non-transitory computer readable storage medium as recited in claim 4, wherein to assign at runtime instructions within the given VLIW to the plurality of parallel execution lanes, further comprising instructions executable to write offsets in particular bit-ranges within a vector register that are associated with corresponding lanes of the plurality of execution lanes, wherein the offsets
5 identify a given instruction within the fetched given VLIW for an associated lane to execute.
7. The non-transitory computer readable storage medium as recited in claim 6, further comprising instructions executable to in response to a given trace having reached the end of the identified loop:
write a sleep state in a corresponding bit-range within the vector register responsive to
10 determining at runtime the given trace is scheduled to branch back to the beginning of the identified loop; and
write an exit state in the corresponding bit-range within the vector register responsive to determining at runtime the given trace is scheduled to branch out of the identified loop.
15
8. The non-transitory computer readable storage medium as recited in claim 7, further comprising instructions executable to in response to the given trace being in the sleep state or the exit state:
halt execution of the given trace; and
store at least a next program counter (PC) and the work-item identifier (ID) for the given
20 trace.
9. The non-transitory computer readable storage medium as recited in claim 8, further comprising instructions executable to in response to each work-item being in the sleep state or each work-item being in the exit state, restart execution for each work-item at a respective stored next PC.
25
10. The non-transitory computer readable storage medium as recited in claim 8, further comprising instructions executable to in response to each work-item being halted and at least one work-item is being a different state than another work-item, restart execution for only work-items in a sleep state at a respective stored next PC.
30

11. A processor comprising:

a plurality of parallel execution lanes within a single instruction multiple data (SIMD) micro-architecture;

a size register configured to store a size of a very large instruction word (VLIW); and

5 control logic configured to fetch and decode in a respective cycle a number of instructions within a given VLIW equal to the stored size;

wherein the plurality of execution lanes are configured to simultaneously execute the number of instructions within the given VLIW.

10 12. The processor as recited in claim 11, wherein the processor further comprises a vector register configured to store offsets in particular bit-ranges that are associated with corresponding lanes of the plurality of execution lanes, wherein the offsets identify a given instruction within the fetched given VLIW for an associated lane to execute.

15 13. The processor as recited in claim 12, wherein a number of valid offsets stored in the vector register is equal to the size stored in the size register.

14. The processor as recited in claim 12, wherein in response to detecting an offset corresponds to a plurality of resource-independent instructions within the VLIW, the plurality of execution lanes are
20 further configured to simultaneously execute the plurality of instructions within an associated lane.

15. The processor as recited in claim 12, wherein in response to a given trace having reached the end of the identified loop and a corresponding bit range within the vector register indicating the given trace is in a sleep state or an exit state, the control logic is further configured to:

25 halt execution of the given trace; and

store at least a next program counter (PC) and the lane identifier (ID) for the given trace.

16. The processor as recited in claim 15, wherein in response to each lane of the plurality of execution lanes being in the sleep state or in the exit state, the control logic is further configured to
30 restart execution for each lane by branching to a respective stored next PC.

17. The processor as recited in claim 15, wherein in response to each lane of the plurality of execution lanes being in the sleep state or in the exit state, the control logic is further configured to restart execution for only lanes in a sleep state by branching to a respective stored next PC.

5

18. A method comprising:

identifying loops and corresponding basic blocks within a plurality of program instructions;
in response to a given divergence point within an identified loop, arranging a plurality of instructions within the identified loop into a plurality of very large instruction words (VLIW's), wherein at least one VLIW comprises instructions intermingled from different basic blocks between the given divergence point and a corresponding convergence point.

10

19. The method as recited in claim 17, wherein in response to the given divergence point within the identified loop, the method further comprises assigning at runtime instructions within a given VLIW to a plurality of parallel execution lanes within a target processor comprising a single instruction multiple data (SIMD) micro-architecture based on branch direction found at runtime for the given lane at the given divergent point.

15

20. The method as recited in claim 19, wherein to assign at runtime instructions within the given VLIW to the plurality of parallel execution lanes, the method further comprises storing indications that are associated with corresponding lanes of the plurality of execution lanes, wherein the indications identify a given instruction within the given VLIW for an associated lane to execute.

20

25

1/8

SIMD Pipeline Execution Flow 100

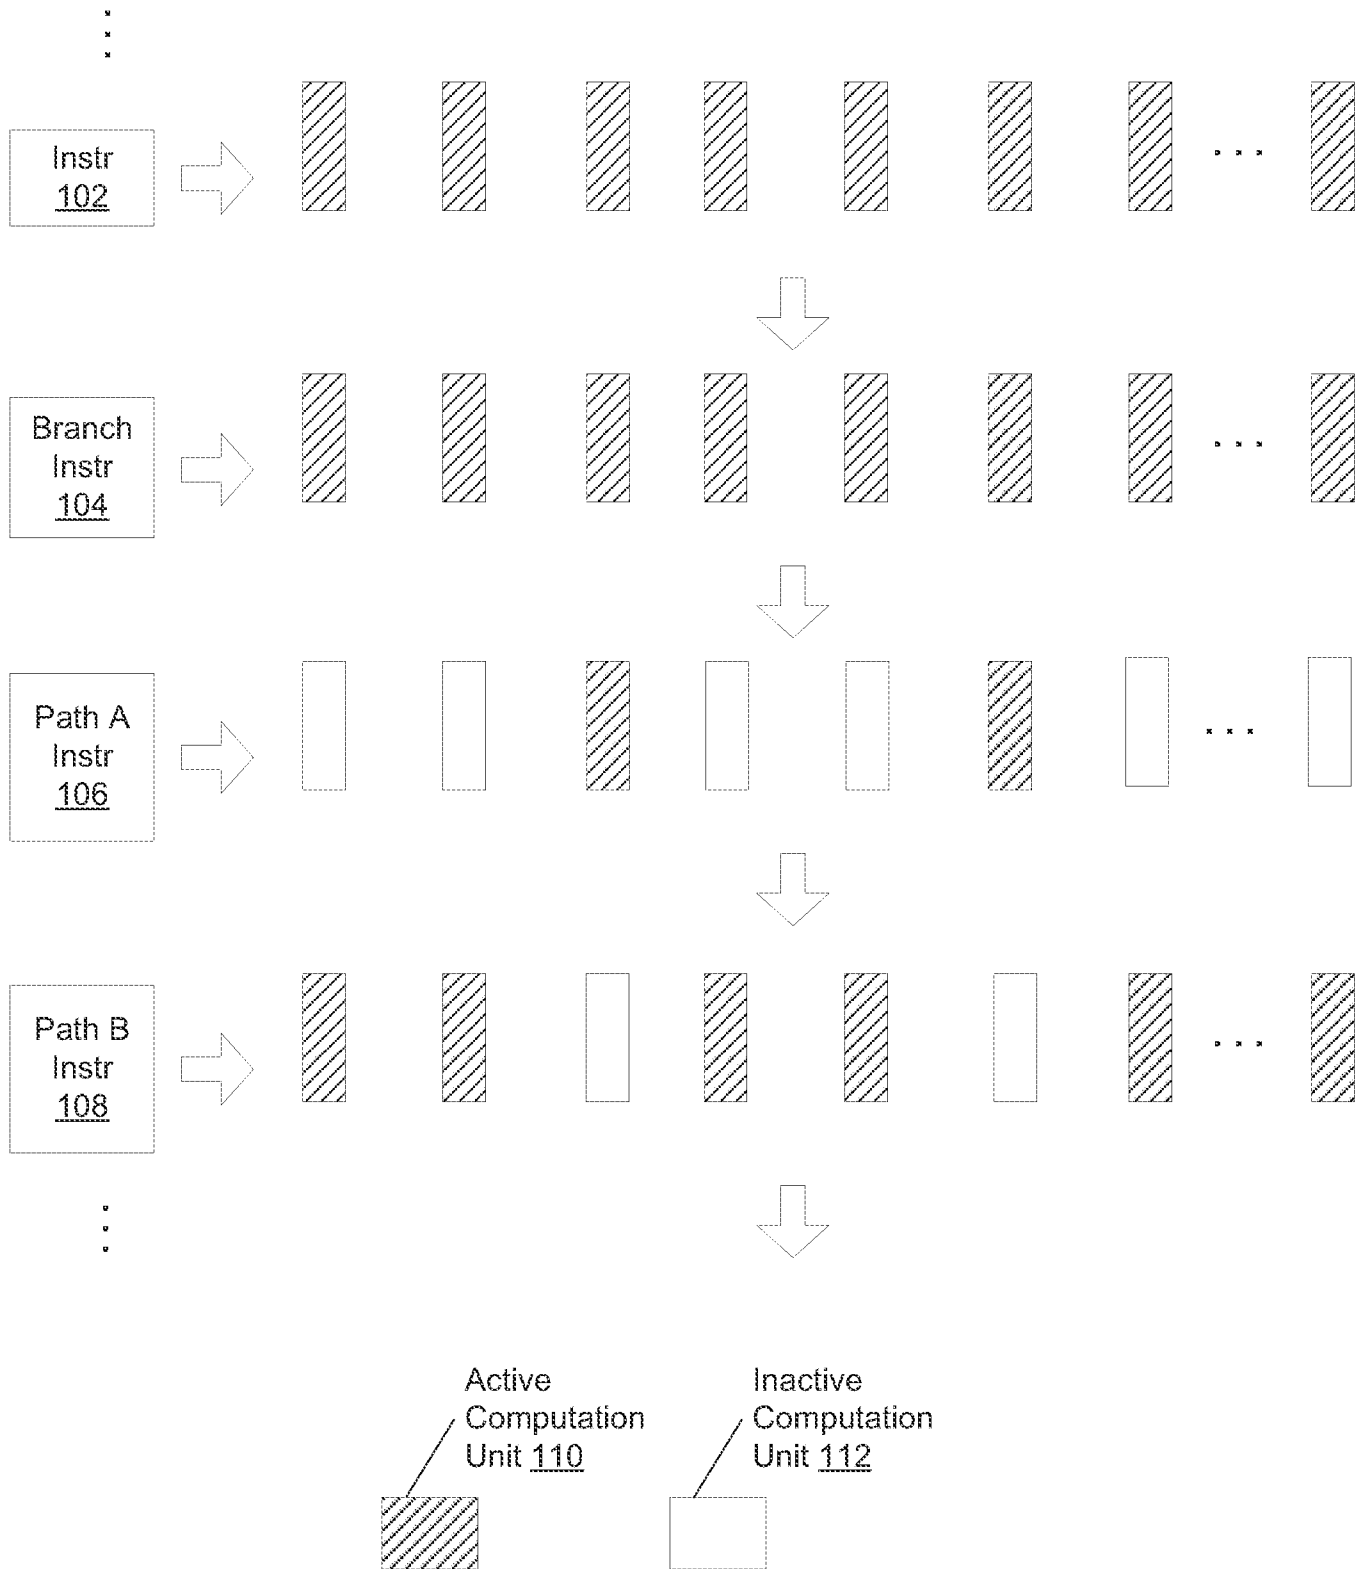


FIG. 1

2/8

Control Flow Graph 200

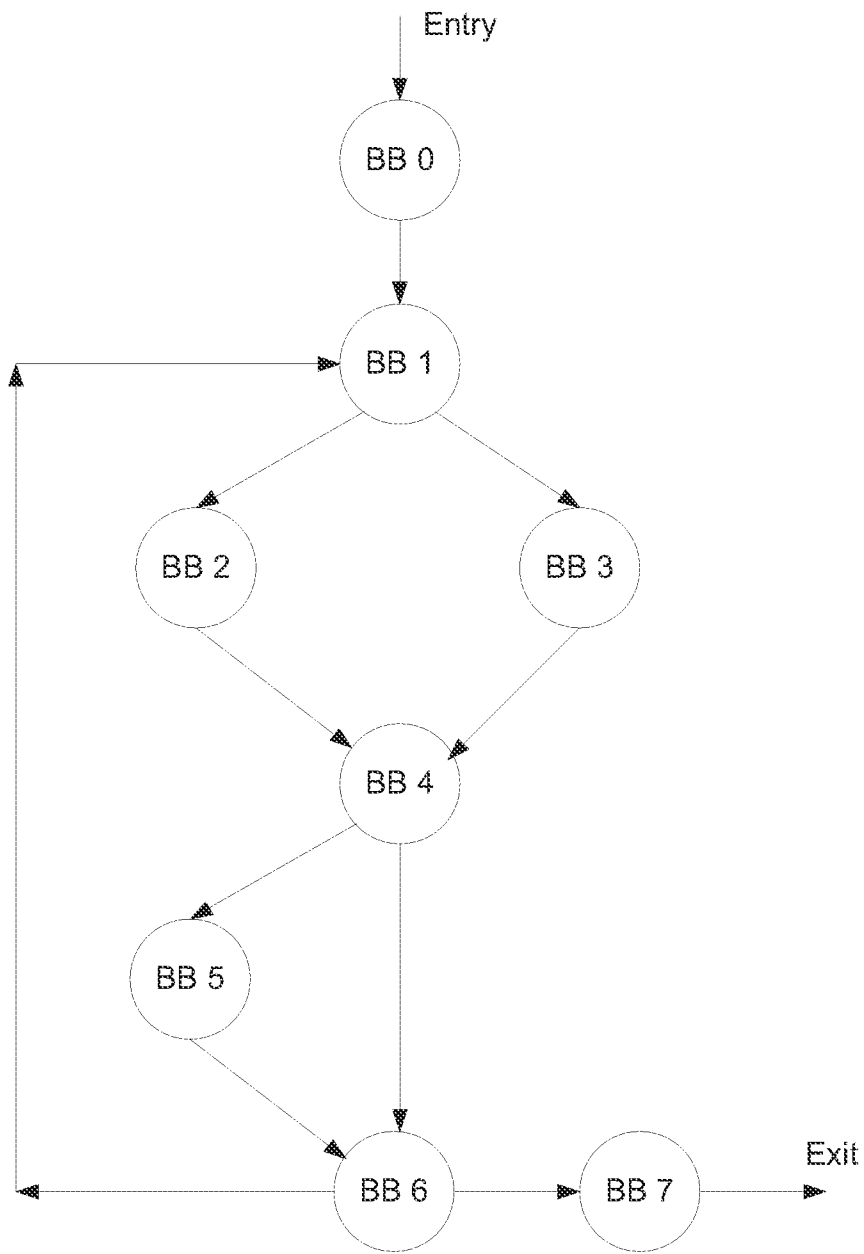


FIG. 2





FIG. 3

4/8

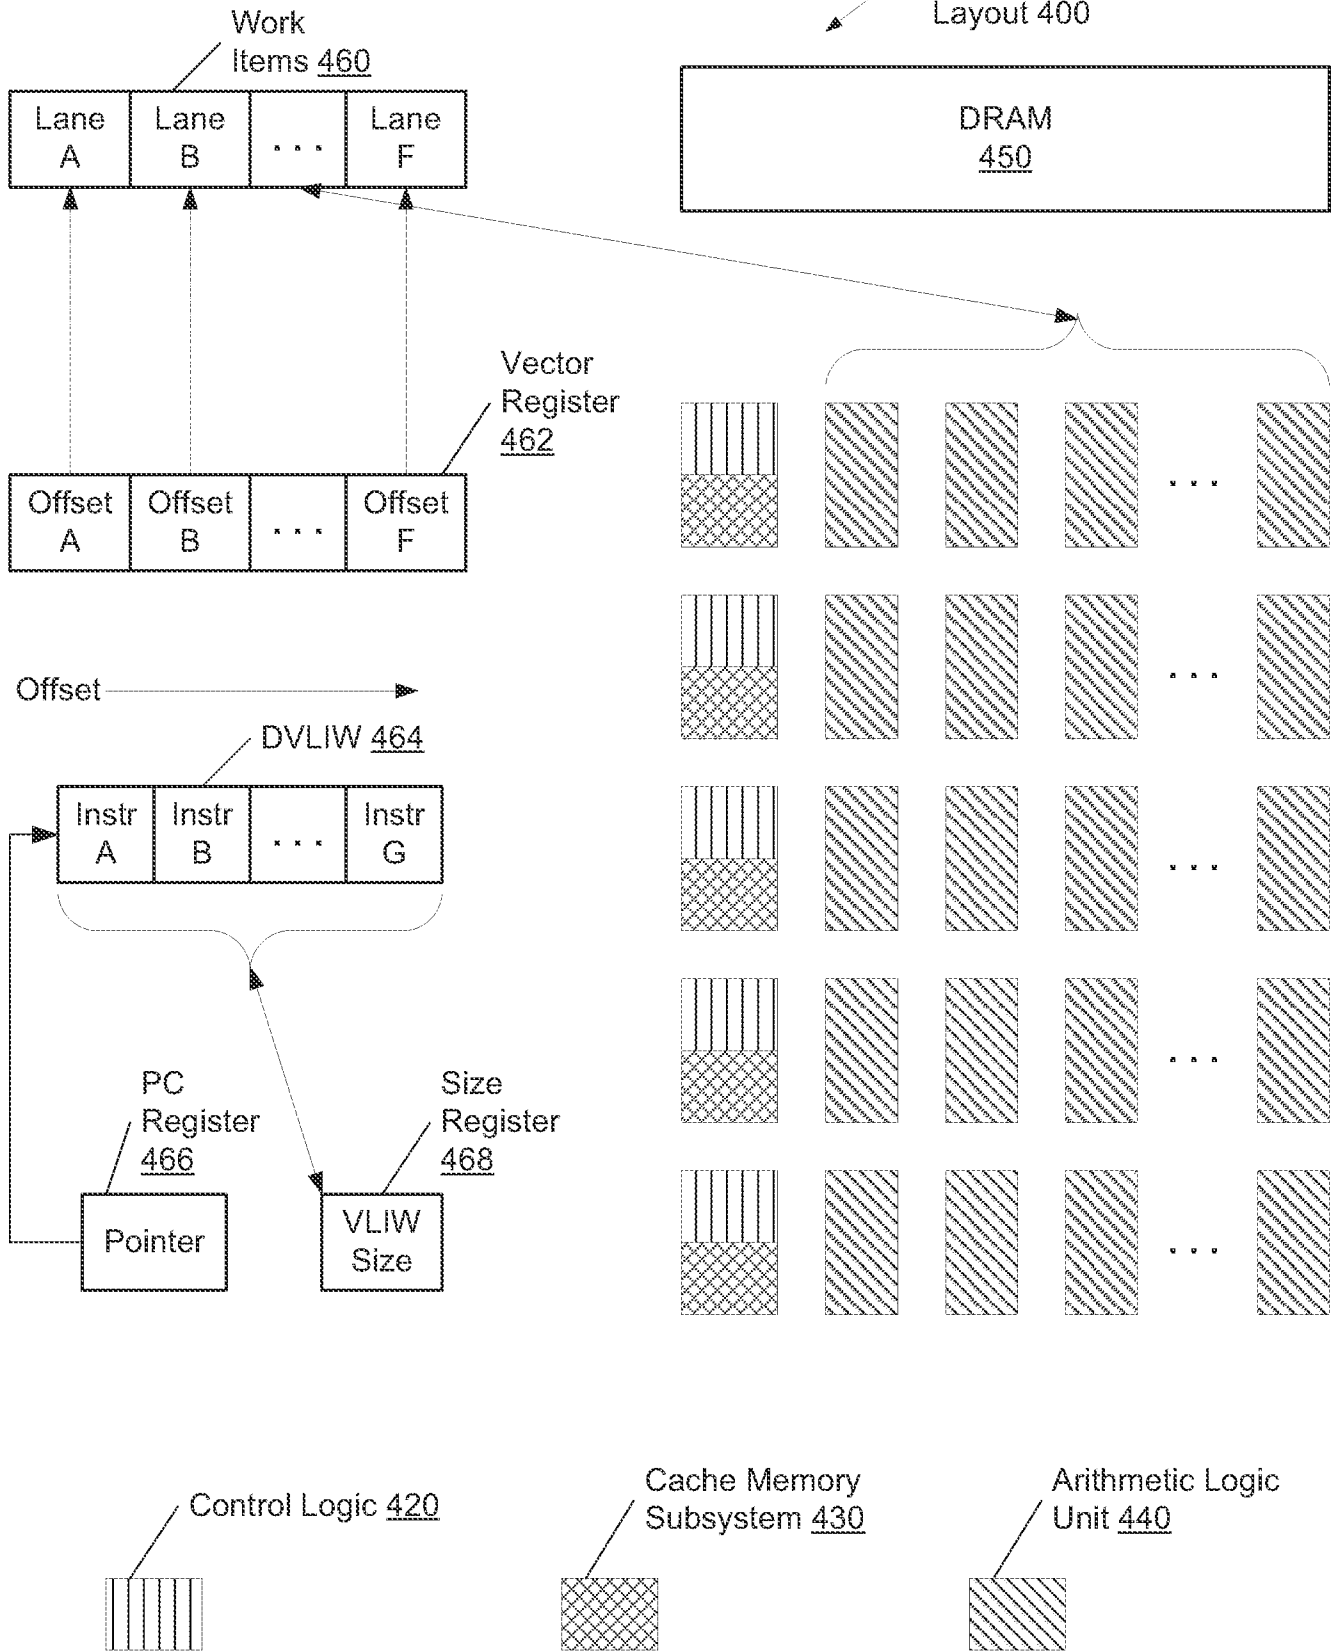


FIG. 4

5/8

Object Code Arrangement
500




FIG. 5

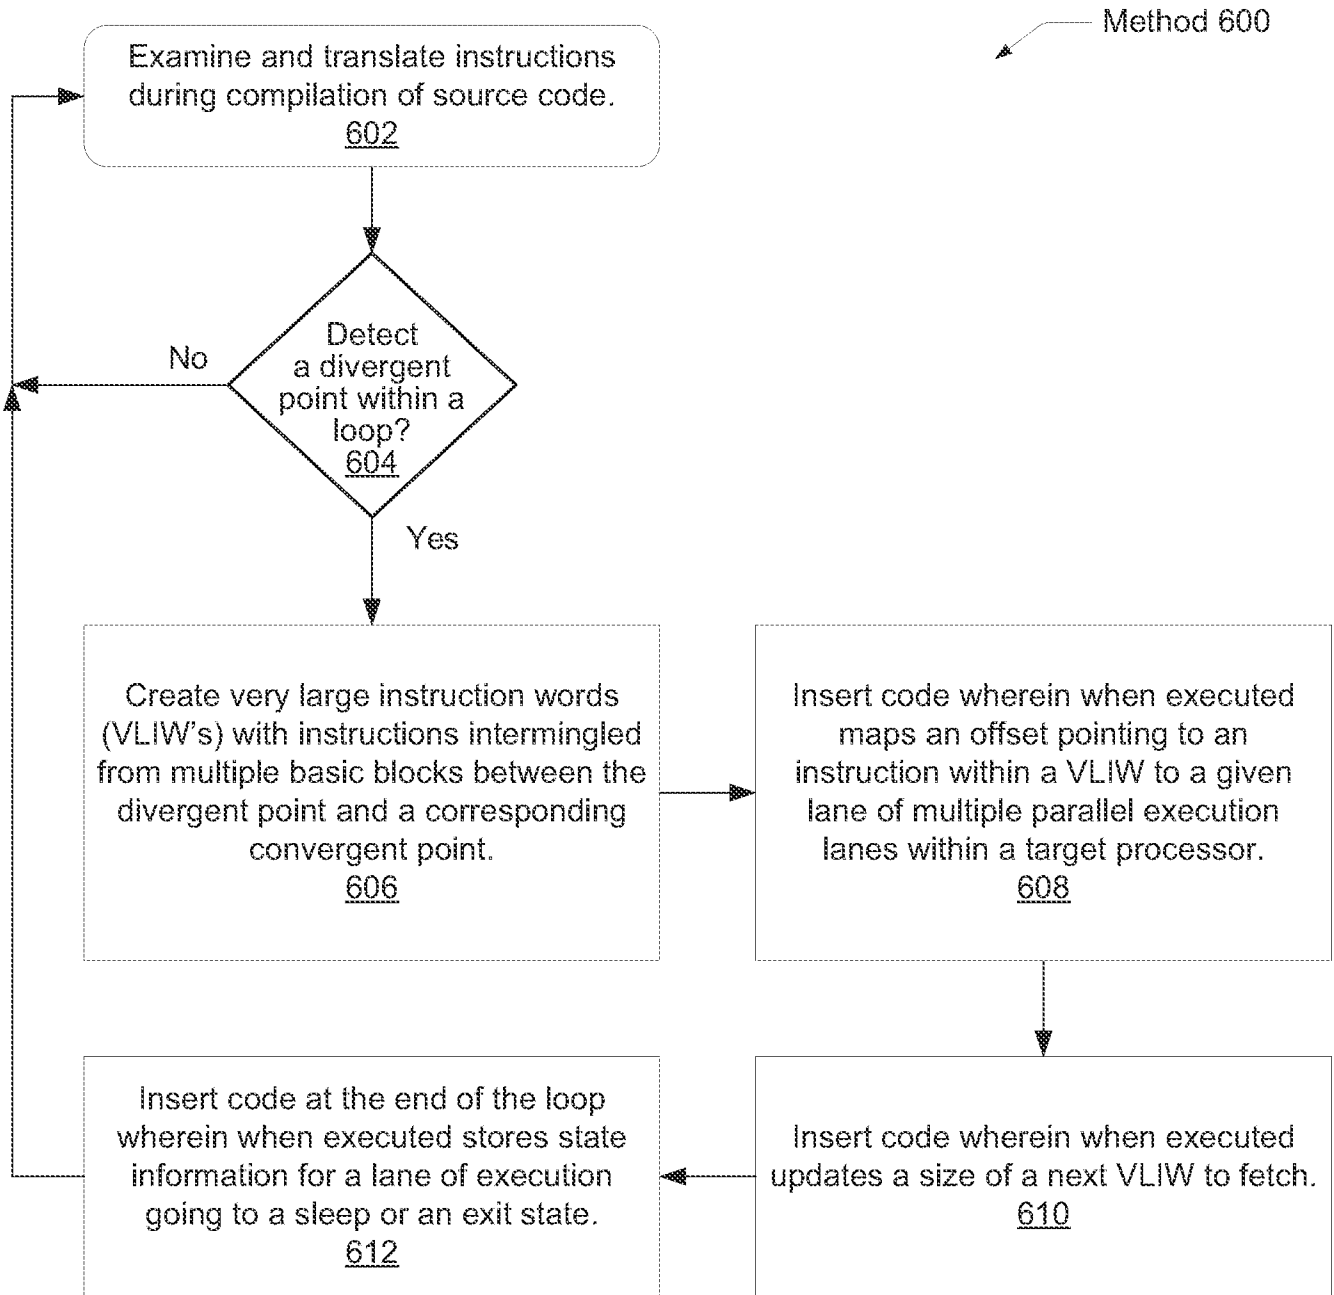


FIG. 6

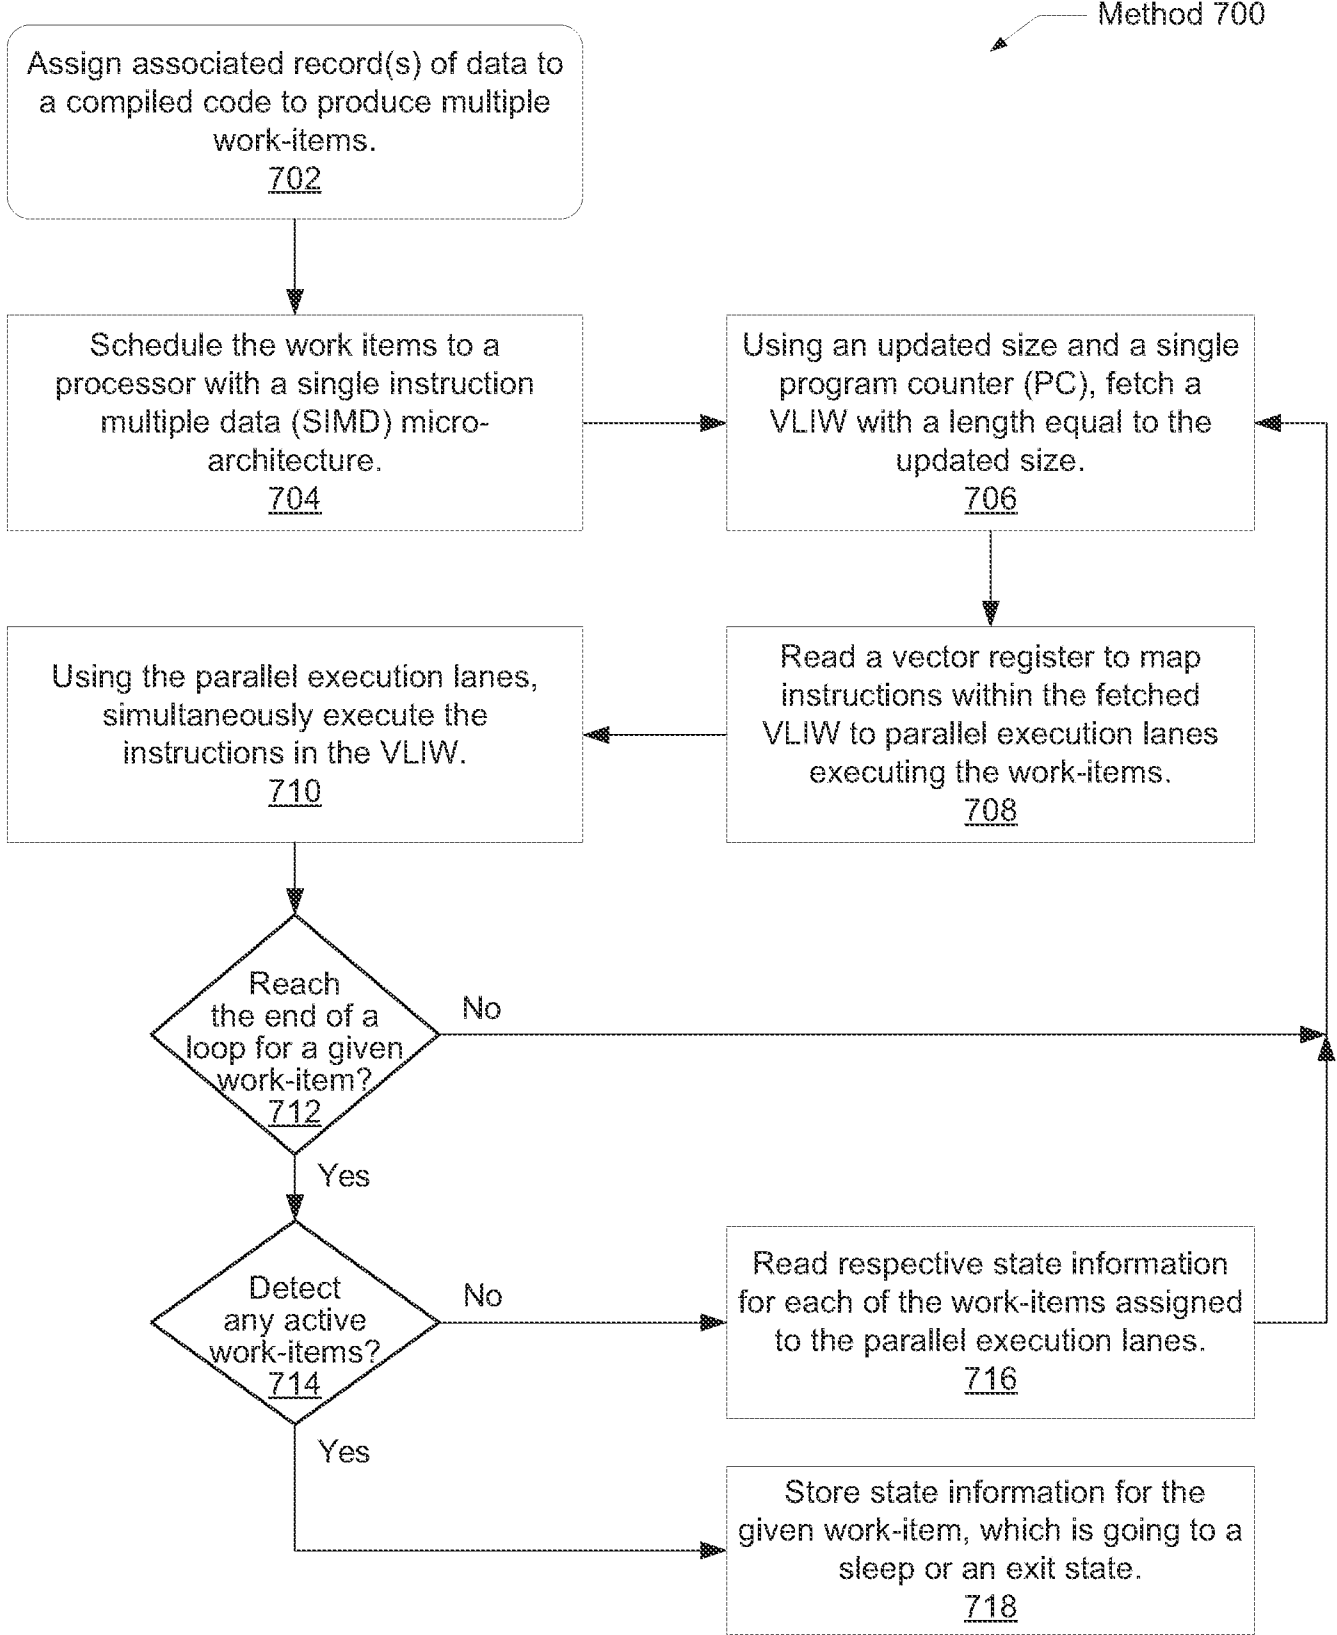



FIG. 7

8/8

Object Code Arrangement
800

Program
Counters



Offset 0
1




FIG. 8

A. CLASSIFICATION OF SUBJECT MATTER**G06F 9/38(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/38; G06F 15/00; G06F 9/00; G06F 9/45

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: parallel pipeline, branch instruction, VLIW

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2007-0011434 A1 (DAVID A. LUICK) 11 January 2007 See abstract, paragraphs [0004]-[0020], and figures 2-6.	1-20
A	US 2007-0174599 A1 (ERIC M. DOWLING) 26 July 2007 See abstract, paragraphs [0014]-[0027], and figures 2-9.	1-20
A	US 2010-0050164 A1 (VAN DE WAERDT JAN-WILLEM et al.) 25 February 2010 See abstract, paragraphs [0037]-[0041], and figures 3-4.	1-20
A	US 6892293 B2 (HOWARD G. SACHS et al.) 10 May 2005 See abstract, column 3, line 21-column 5, line 4, and figures 2-8.	1-20

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

* Special categories of cited documents:

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

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

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

29 May 2014 (29.05.2014)

Date of mailing of the international search report

30 May 2014 (30.05.2014)

Name and mailing address of the ISA/KR

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

Facsimile No. +82-42-472-7140

Authorized officer

JI, Jeong Hoon

Telephone No. +82-42-481-5688



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/013455

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0011434 A1	11/01/2007	CN 1598762 A	23/03/2005
		CN 1598762 C0	06/06/2007
		US 2005-0066148 A1	24/03/2005
		US 2009-044049 A1	12/02/2009
		US 7124318 B2	17/10/2006
		US 7454654 B2	18/11/2008
US 2007-0174599 A1	26/07/2007	US 06157988 A	05/12/2000
		US 2004-0068643 A1	08/04/2004
		US 6647467 B1	11/11/2003
		US 7197603 B2	27/03/2007
		US 7398358 B2	08/07/2008
US 2010-0050164 A1	25/02/2010	EP 2106584 A1	07/10/2009
		WO 2008-072178 A1	19/06/2008
US 6892293 B2	10/05/2005	DE 69424370 D1	15/06/2000
		DE 69424370 T2	15/02/2001
		DE 69430018 D1	04/04/2002
		DE 69430018 T2	21/11/2002
		DE 69433124 D1	09/10/2003
		DE 69433124 T2	27/05/2004
		EP 0652509 A1	10/05/1995
		EP 0652509 B1	10/05/2000
		EP 0974894 A2	26/01/2000
		EP 0974894 A3	02/08/2000
		EP 0974894 B1	27/02/2002
		EP 1186995 A1	13/03/2002
		EP 1186995 B1	03/09/2003
		EP 1369775 A1	10/12/2003
		HK 1044056 A1	16/04/2004
		US 05794003 A	11/08/1998
		US 2003-0079112 A1	24/04/2003
US 2003-0191923 A1	09/10/2003		
US 6360313 B1	19/03/2002		
US 7039791 B2	02/05/2006		